Sequential Equilibria of Games with Infinite Sets of Types and Actions

by

Roger Myerson and Philip J. Reny*

Draft notes May 2011

Abstract: We formulate a definition of *basic sequential equilibrium* for multi-stage games with infinite type sets and infinite action sets, and we prove its general existence. We then explore several difficulties of this basic concept and propose concepts of *essential* and *extended sequential equilibrium* to resolve such difficulties while maintaining general existence.

* 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.
Most detailed formulation of "perfect bayesian eqm" in Fudenberg-Tirole 1991.
Harris-Stinchcombe-Zame 2000 explored definitions with nonstandard analysis.

It is well understood that sequential equilibria of an infinite game can be defined by taking limits of sequential equilibria of finite games that approximate it.
The problem is to define what finite games are good approximations.
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 must try to define a class of finite approximations that yield limit-equilibria which include reasonable equilibria and exclude unreasonable equilibria.
Here we present the best definitions that we have been able to find, but of course this depends on intuitive judgments about what is "reasonable".
Others should explore alternative definitions.

Dynamic multi-stage games $\Gamma=(\Theta,N,A,T,\tau,v,p)$

Θ = {*initial state space*}, $i \in N$ = {*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 period k}.

T_{ik} = {possible informational *types* for player i at period k}, disjoint sets.

Algebras (closed under finite $\cap$ and complements) of *measurable* subsets are specified for Θ and T_{ik}, including all one-point sets as measurable.

If Θ is a finite set, then all subsets of Θ and T_{ik} are measurable.

There is a finitely additive *probability measure* p on the measurable subsets of Θ.

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

The subscript, <k, denotes the projection onto periods before k. For example,

$A_{<k} = \times_{h<k} \times_{i\in N} A_{ih}$ = {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.

Any player i's information at any period k is specified by a *type function*

$\tau_{ik}:\Theta\times A_{<k} \to T_{ik}$ such that, $\forall a\in A$, $\tau_{ik}(\theta,a_{<k})$ is a measurable function of θ.

Assume *perfect recall*: $\forall ik\in L$, $\forall m<k$, $\exists\rho_{ikm}:T_{ik}\to T_{im}\times A_{im}$ such that

$\rho_{ikm}(\tau_{ik}(\theta,a_{<k})) = (\tau_{im}(\theta,a_{<m}),a_{im})$, $\forall\theta\in\Theta$, $\forall a\in A$, and

$\{t_{ik}| \rho_{ikm}(t_{ik})\in R_{im}\times\{a_{im}\}\}$ is measurable in T_{ik}, $\forall$ measurable $R_{im}\subseteq T_{im}$, $\forall a_{im}\in A_{im}$.

Each player i has a bounded *utility function* $v_i:\Theta\times A\to\mathbb{R}$ such that

$v_i(\theta,a)$ is a measurable function of $\theta\in\Theta$, $\forall a\in A$. Bound $\Omega \geq |v_i(\theta,a)|$, $\forall(i,\theta,a)$.

Problems of finite support mixtures in approximating games

Example. (Kuhn) Consider a zero-sum game in which player 1 chooses a number a_1 in [0,1], and player 2 chooses a continuous function, f, from [0,1] into itself whose Lebesgue integral must be 1/2. Player 1's payoff is $f(a_1)$.

Player 1 can guarantee 1/2 by choosing a_1 uniformly from [0,1] and player 2 can guarantee 1/2 by choosing $f(x) = 1/2$ for all x.

But in any finite approximation in which player 2 can choose a function that is zero at each of player 1's finitely many available actions, player 1's equilibrium payoff is 0.

One solution is to replace ordinary action sets with mixed action sets.

Algebras (closed under finite $\cap$ and complements) of *measurable* subsets are specified also for each A_{ik}, including all one-point sets as measurable.

Each $\tau_{ik}(\theta,a_{<k})$ and each $v_i(\theta,a)$ is assumed jointly measurable in (θ,a).

Let $\tilde{\Theta} = \Theta\times(\times_{ik}\,[0,1])$. In the state $\tilde{\theta} = (\theta,(\tilde{\theta}_{ik})_{ik\in L})$, nature draws θ from the given p, and, for each ik, draws $\tilde{\theta}_{ik}\in[0,1]$ independently from Lebesgue measure.

This defines a new distribution $\tilde{p}$ on $\tilde{\Theta}$. No player observes any of the new $\tilde{\theta}_{ik}$.

Let $\tilde{A}_{ik}$={measurable maps $\tilde{a}_{ik}$:[0,1] $\rightarrow A_{ik}$} be player ik's set of mixed actions.

Then $\tilde{v}_i(\tilde{\theta},\tilde{a}) = v_i(\theta,\tilde{a}(\tilde{\theta}))$ and $\tilde{\tau}_{ik}(\tilde{\theta},\tilde{a}_{<k}) = (\tau_{ik}(\theta,\tilde{a}(\tilde{\theta})_{<k}), \tilde{a}_{i,<k})$ are measurable in $\tilde{\theta}$ for each profile of mixed actions $\tilde{a}$.

All our analysis could be done for this model with $\tilde{A},\tilde{\Theta},\tilde{p}$ instead of A,Θ,p.

Problems of spurious knowledge in approximating games

Some approximations of the state space Θ can change the information structure and give players information that they do not have in the original game.

Example. Θ includes two independent uniform [0,1] random variables:
the first (1's cost to sell some object) is observed only by player 1,
the second (2's value of buying this object) is observed only by player 2.
For any integer $m>0$, we might finitely approximate this game by one in which the state space includes 10^{2m} equally likely pairs of numbers:
the first number, observed by 1, ranges over the 10^m m-digit decimals in [0,1];
the second, observed by 2, is a 2m-digit decimal in which the first m digits range over all m-digit decimals but the last m digits repeat the first number in the pair.
As $m\to\infty$, these pairs uniformly fill the unit square in the state space Θ,
but they represent games in which player 2 knows 1's type.

To avoid distorting the information structure, we will keep the original state space Θ and probability distribution p in all finite approximations.
A player should not know anything about the state and past history in a finite approximation that he does not know in the given infinite game.
To avoid giving spurious knowledge to any player here, information is made finite by finitely partitioning each dated-player's type space T_{ik} independently.

Problems of spurious signaling in approximating games

Example. The state θ is a 0-or-1 random variable, which is observed by player 1 as his type t_1. Then player 1 chooses a number a_1 in [0,1], which is subsequently observed by player 2.

Consider a finite approximation in which 1 observes t_1, and then can choose any m-digit decimal number in [0,1] whose m-th digit is t_1.

Then 1's choice would reveal his observation, even if he wanted to conceal it.

To avoid such spurious signaling, we should assume that the finite action choices available to player 1 do not depend on his observed type information.

That is, the finite subset of actions A_{ik} that are available to player i at period k in a finite approximation should not depend on what ik has observed in T_{ik}.

In some games, the ability to choose some special action in A_{ik} or the ability to distinguish some special subsets of T_{ik} may be particularly important.

So we consider convergence of a net of approximations (not just a sequence) that are indexed on all finite subsets of A_{ik} and all finite partitions of T_{ik}.

In such a net, any finite collection of actions and any finite partition of types can be assumed available when we are determining the properties of the limits.

Problems of spurious concealment in approximating games

Example. The state θ is a 0-or-1 random variable, which is observed by player 1 as his type t_1. Then player 1 chooses a number a_1 in [0,1], which is subsequently observed by players 2 and 3.

Consider a finite approximation in which 1 observes t_1, and then can choose any m-digit decimal number in [0,1].

Suppose that player 1 wants to share information with 2 but conceal it from 3.

Such concealed signaling would be possible in a finite approximation where 2 can observe 1's action exactly while 3 can observe only its first m-1 digits.

But in the real game, any message that 1 sends to 2 is also observed by 3, and it should not be possible for 1 and 2 to tunnel information past 3.

We can avoid such spuriously concealed communication by taking limits of finite approximations in which every player's ability to observe increases faster than any player's ability to choose different actions.

So we first take limits as partitions of type-space T_{ik} become infinitely fine, and we then take limits as the finite subsets of feasible actions expand to fill A_{ik}.

Observable events and action approximations

For any $a \in A$ and any measurable $R_{ik} \subseteq T_{ik}$, let $P(R_{ik}|a) = p(\{\theta | \tau_{ik}(\theta, a_{<k}) \in R_{ik}\})$.
(For k=1, we could write $P(R_{i1})$ for $P(R_{i1}|a)$, ignoring the trivial $a_{<1} = \emptyset$.)

The set of *observable events for i at k* that can have positive probability is
$\mathcal{Q}_{ik} = \{R_{ik} \subseteq T_{ik} | R_{ik}$ is measurable and $\exists a \in A$ such that $P(R_{ik}|a) > 0\}$.
Let $\mathcal{Q} = \cup_{ik \in L} \mathcal{Q}_{ik}$ (a disjoint union) denote the set of all events that can be observed with positive probability by some dated player.

We extend this notation to describe observable events that can have positive probability when players are restricted to actions in some subsets of the A_{ik}.
An *action approximation* is any $C = \times_{ik \in L} C_{ik}$ such that each C_{ik} is a nonempty finite subset of A_{ik}, and so $C \subseteq A$.

For any action approximation $C \subseteq A$, let
$\mathcal{Q}_{ik}(C) = \{R_{ik} \subseteq T_{ik} | R_{ik}$ is measurable and $\exists c \in C$ such that $P(R_{ik}|c) > 0\}$.
Let $\mathcal{Q}(C) = \cup_{ik \in L} \mathcal{Q}_{ik}(C)$ denote the set of all events that can be observed with positive probability by some dated player when all players use actions in C.

Action approximations are partially ordered by inclusion.
If $C \supseteq C^o$ then C is a better approximation than C^o to the true action sets A.
Fact. If $C \supseteq C^o$ then $\mathcal{Q}(C) \supseteq \mathcal{Q}(C^o)$.
$\mathcal{Q}$ is the union of all $\mathcal{Q}(C)$ over all finite approximations C.

Finite approximations of the game

An *information approximation* is any $S = \times_{ik\in L} S_{ik}$ such that each S_{ik} is a finite partition of measurable subsets of T_{ik}.
(So elements of each S_{ik} are disjoint measurable sets with union T_{ik}.)

Let $\mathcal{U}(S_{ik})$ denote the set of nonempty unions of sets in S_{ik}.
Let $\mathcal{U}(S) = \cup_{ik\in L}\, \mathcal{U}(S_{ik})$ (a disjoint union).
Information approximations are partially ordered by inclusion of their unions.
If $\mathcal{U}(S) \supseteq \mathcal{U}(S^o)$ then S is a better approximation than S^o to the true type sets T.

An action approximation C and an information approximation S together define a *finite approximation* (C,S) of the game.
Any s_{ik} in S_{ik} is a possible *type* of dated player ik in this finite approximation.

(C,S) satisfies *perfect recall* iff $\forall ik\in L$, $\forall s_{ik}\in S_{ik}$, $\forall m<k$, $\exists s_{im}\in S_{im}$, $\exists d_{im}\in C_{im}$ s.t.
$\{(\theta,a)|\ \tau_{ik}(\theta,a_{<k})\in s_{ik},\ a_{im}\in C_{im}\} \subseteq \{(\theta,a)|\ \tau_{im}(\theta,a_{<m})\in s_{im},\ a_{im}=d_{im}\}$.
Let $\mathcal{F}$ denote the set of finite approximations with perfect recall.
Fact. For any action approximation C and any information approximation S^o, there exists an information approximation S such that $\mathcal{U}(S) \supseteq \mathcal{U}(S^o)$ and (C,S) satisfies perfect recall, so that $(C,S)\in\mathcal{F}$.

Strategy profiles for finite approximations

Here let $(C,S)\in\mathcal{F}$ be a given finite approximation of the game.

A *strategy profile* for the finite approximation (C,S) is any $\sigma = (\sigma_{ik})_{ik\in L}$ such that each $\sigma_{ik}:S_{ik}\rightarrow\Delta(C_{ik})$.

So $\sigma_{ik}(c_{ik}|s_{ik})\geq 0,\ \forall c_{ik}\in C_{ik},\ \sum_{\gamma\in C\_ik} \sigma_{ik}(\gamma|s_{ik})=1,\ \forall s_{ik}\in S_{ik}$.

Let $[t_{ik}]$ denote the element of S_{ik} containing $t_{ik}\in T_{ik}$.

Given $(C,S)\in\mathcal{F}$, for any measurable $Z\subseteq\Theta$ and $c\in C$, let

$P(Z,c|\sigma) = \int_{\theta\in Z} (\prod_{ik\in L} \sigma_{ik}(c_{ik}|[\tau_{ik}(\theta,c_{<k})])) p(d\theta)$.

For any observable $R_{ik}\in\mathcal{U}(S_{ik})$, let

$P(R_{ik}|\sigma)= \sum_{c\in C} P(\{\theta\in\Theta|\ \tau_{ik}(\theta,c_{<k})\in R_{ik}\},c|\sigma)$.

A *totally mixed* strategy profile σ for (C,S) has $\sigma_{ik}(c_{ik}|s_{ik})>0\ \forall c_{ik}\in C_{ik},\ \forall s_{ik}\in S_{ik}$.

Any totally mixed σ yields $P(R_{ik}|\sigma) > 0,\ \forall R_{ik}\in \mathcal{U}(S_{ik})\cap\mathcal{Q}_{ik}(C),\ \forall ik\in L$.

With $P(R_{ik}|\sigma) > 0$, let $P(Z,c|\sigma,R_{ik}) = P(Z\cap\{\theta|\tau_{ik}(\theta,c_{<k})\in R_{ik}\},c|\sigma)/P(R_{ik}|\sigma)$.

We can extend this probability function to define

$P(Y|\sigma,R_{ik}) = \sum_{c\in C} P(Y(c),c|\sigma,R_{ik})$ for any $Y\subseteq\Theta\times A$ such that,

for each $c\in C$, $Y(c) = \{\theta|\ (\theta,c)\in Y\}$ is a measurable subset of Θ.

Approximate sequential equilibria

Here let $(C,S) \in \mathcal{F}$ be a given finite approximation of the game, and let σ be a totally mixed strategy profile for (C,S).

With the probability function P defined above, we can define

$V_i(\sigma|R_{ik}) = \sum_{c\in C} \int_{\theta\in\Theta} v_i(\theta,c)\, P(d\theta,c|\sigma,R_{ik}),\ \ \forall R_{ik} \in \mathcal{U}(S_{ik}) \cap \mathcal{Q}_{ik}(C)$.

For any $ik \in L$ and any $c_{ik} \in C_{ik}$, let (σ_{-ik},c_{ik}) denote the strategy profile that differs from σ only in that i chooses action c_{ik} at k with probability 1.

Changing the action of i at k does not change the probability of i's types at k, so

$P(R_{ik}|\sigma_{-ik},c_{ik}) = P(R_{ik}|\sigma) > 0$, $\forall R_{ik} \in \mathcal{U}(S_{ik}) \cap \mathcal{Q}_{ik}(C)$.

So we can similarly define $V_i(\sigma_{-ik},c_{ik}|R_{ik})$ to be i's *sequential value* of choosing action c_{ik} at period k, given the observation R_{ik}, if others apply the σ strategies.

For $\varepsilon>0$, σ is an *ε-approximate sequential equilibrium* for the finite approximation (C,S) iff σ is a totally mixed strategy for (C,S) and

$V_i(\sigma_{-ik},c_{ik}|s_{ik}) \le V_i(\sigma|s_{ik}) + \varepsilon,\ \ \forall ik \in L,\ \ \forall s_{ik} \in S_{ik} \cap \mathcal{Q}_{ik}(C),\ \ \forall c_{ik} \in C_{ik}$.

This inequality for types implies the same inequality for their unions:

$V_i(\sigma_{-ik},c_{ik}|R_{ik}) \le V_i(\sigma|R_{ik}) + \varepsilon,\ \ \forall ik \in L,\ \ \forall R_{ik} \in \mathcal{U}(S_{ik}) \cap \mathcal{Q}_{ik}(C),\ \ \forall c_{ik} \in C_{ik}$.

Fact. Any finite approximation has an ε-approximate sequential equilibrium.

Assessments

Let $\mathcal{Y}$ denote the set of all *outcome events* $Y \subseteq \Theta \times A$ such that $\{\theta | (\theta,a) \in Y\}$ is a measurable subset of Θ, $\forall a \in A$.

For any finite approximation (C,S), let $\mathcal{W}(C,S) = \cup_{ik \in L} C_{ik} \times (\mathcal{Q}_{ik}(C) \cap \mathcal{U}(S_{ik}))$.

The sequential values that are compatible with (C,S) are indexed on $\mathcal{W}(C,S)$.

Let $\mathcal{W} = \cup_{ik \in L} A_{ik} \times \mathcal{Q}_{ik}$ (the union of all such $\mathcal{W}(C,S)$ domains).

For the bounded utility functions, let Ω be such that $|v_i(\theta,a)| \le \Omega$ $\forall i$, $\forall \theta$, $\forall a$.

An *assessment* is a pair (μ,ω), where the vector μ specifies conditional probabilities $\mu(Y|Q) \in [0,1]$ $\forall Y \in \mathcal{Y}$, $\forall Q \in \mathcal{Q}$, and the vector ω specifies sequential values $\omega_i(a_{ik}|R_{ik}) \in [-\Omega,\Omega]$ $\forall ik \in L$, $\forall a_{ik} \in A_{ik}$, $\forall R_{ik} \in \mathcal{Q}_{ik}$.

So (μ,ω) is in $[0,1]^{\mathcal{Y} \times \mathcal{Q}} \times [-\Omega, \Omega]^{\mathcal{W}}$, which is compact in product topologies.

An *assessment-test* is a pair of sets $\Phi = (\Phi_1, \Phi_2)$ such that $\Phi_1 \subseteq \mathcal{Y} \times \mathcal{Q}$, $\Phi_2 \subseteq \mathcal{W}$, and both Φ_1 and Φ_2 are finite sets.

An assessment test Φ is *compatible* with a finite approximation (C,S) iff $\Phi_1 \subseteq \mathcal{Y} \times (\mathcal{Q}(C) \cap \mathcal{U}(S))$ and $\Phi_2 \subseteq \mathcal{W}(C,S)$.

Compatibility means that all elements of Φ are conditioned on events that can be observed with positive probability in (C,S), and Φ_2 considers only actions in C.

Fact. If Φ is compatible with (C^o,S^o) and $C^o \subseteq C$ and $\mathcal{U}(S^o) \subseteq \mathcal{U}(S)$ then Φ is compatible with (C,S).

Basic sequential equilibria

An assessment (μ,ω) is a *basic sequential equilibrium* iff: for every $\varepsilon > 0$, for every finite assessment-test $\Phi=(\Phi_1,\Phi_2)$, for every action approximation C^o, there exists an action approximation $C \supseteq C^o$ such that,
for every information approximation S^o,
there exists an information approximation S such that $\mathcal{U}(S) \supseteq \mathcal{U}(S^o)$, $(C,S)\in\mathcal{F}$,
the test Φ is compatible with (C,S),
and (C,S) has some ε-approximate sequential equilibrium σ such that
$|\mu(Y|Q)-P(Y|\sigma,Q)| \leq \varepsilon,\ \forall(Y,Q)\in\Phi_1$ and
$|\omega_i(a_{ik}|R_{ik})-V_i(\sigma_{-ik},a_{ik}|R_{ik})| \leq \varepsilon,\ \forall(a_{ik},R_{ik})\in\Phi_2$.

Theorem. The set of basic sequential equilibria is nonempty and is a closed subset of $[0,1]^{\mathcal{Y}\times\mathcal{Q}} \times [-\Omega, \Omega]^{\mathcal{W}}$ with the product topology.

The proof is based on Tychonoff's Theorem on compactness of products of compact sets. Kelley, 1955, p143.

Defining basic sequential equilibria as limits of nets

A set of assessments U is *open* (in the product topology) iff:

$\forall(\mu,\omega)\in U$, $\exists$an assessment-test $\Phi=(\Phi_1,\Phi_2)$ and $\exists\varepsilon>0$ such that

U contains all points (μ',ω') such that $|\mu'(Y|Q)-\mu(Y|Q)| \leq \varepsilon$, $\forall(Y,Q)\in\Phi_1$

and $|\omega'_i(a_{ik}|R_{ik})-\omega_i(a_{ik}|R_{ik})| \leq \varepsilon$, $\forall(a_{ik},R_{ik})\in\Phi_2$.

A set of assessments is *closed* iff its complement in $[0,1]^{\mathcal{Y}\times\mathcal{Q}} \times [-\Omega, \Omega]^{\mathcal{W}}$ is open.

Let $E(\varepsilon,C,S)$ denote the set of ε-approximate sequential equilibria for the finite approximation (C,S). Given any $\sigma\in E(\varepsilon,C,S)$:

let $m(Y|Q,\sigma,C,S) = P(Y|\sigma,Q)$ if $Q\in\mathcal{Q}(C)\cap\mathcal{U}(S)$, else $m(Y|Q,\sigma,C,S) = 0$;

let $w_i(a_{ik}|R_{ik},\sigma,C,S) = V_i(\sigma_{-ik},a_{ik}|R_{ik})$ if $(a_{ik},R_{ik})\in\mathcal{W}(C,S)$, else $w_i(a_{ik}|R_{ik},\sigma,C,S)=0$.

Action and information approximations are directed sets, partially ordered by inclusion ($C^o\subseteq C$ or $\mathcal{U}(S^o)\subseteq \mathcal{U}(S)$), where the union of any pair is beyond both.

(For positive numbers ε, "beyond" means closer to 0.)

For any net of sets indexed by a directed set, a *subnet limit* is any point x such that, for any open set containing x and any given index value, the net must include a set, with an index beyond the given index, that contains a point in the open set. Let "lim" here denote the set of such subnet limits.

Then our basic sequential equilibria are

$\lim_{\varepsilon>0} \lim_{C\subset A} (\lim_{S:(C,S)\in\mathcal{F}} \{(m(\sigma,C,S),w(\sigma,C,S))|\ \sigma\in E(\varepsilon,C,S)\})$.

Elementary properties of basic sequential equilibria

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

Let (μ,ω) be a basic sequential equilibrium.
Then μ has the following general properties of conditional probabilities, for any outcome events Y and Z and any observable events R_{ik} and R_{jm}:
$\mu(Y|R_{ik})\in[0,1]$, $\mu(\Theta\times A|R_{ik}) = 1$, $\mu(\emptyset|R_{ik}) = 0$ (*probabilities*);
if $Y\cap Z=\emptyset$ then $\mu(Y\cup Z|R_{ik}) = \mu(Y|R_{ik}) + \mu(Z|R_{ik})$ (*finite additivity*);
$\mu(Y|R_{ik}) = \mu(Y\cap I(R_{ik})|R_{ik})$ (*conditional support*);
$\mu(Y\cap I(R_{jm})|R_{ik}) = \mu(Y\cap I(R_{ik})|R_{jm})\ \mu(I(R_{jm})|R_{ik})$ (*Bayes consistency*).

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 $\Theta\times A$ for a basic sequential equilibrium can be defined by $\mu(Y) = \mu(Y|T_{ik})$, $\forall Y\subseteq\Theta\times A$, $\forall ik\in L$.
The unconditional marginal distribution of μ on Θ is the given prior p:
$\mu(Z\times A) = p(Z)$, for any measurable $Z\subseteq\Theta$.

Fact (sequential rationality). If (μ,ω) is a basic sequential equilibrium then,
$\forall ik\in L$, $\forall R_{ik}\in\mathcal{Q}_{ik}$, $\forall c_{ik}\in A_{ik}$, $\int v_i(\theta,a)\ \mu(d\theta,da|R_{ik}) \geq \omega_i(c_{ik}|R_{ik})$.
The left-hand side integral is *player i's equilibrium payoff conditional on* R_{ik}.

Games with perfect information

If $\Theta=\{\theta_0\}$, then a dynamic multi-stage game has *perfect information* iff:
$\forall k\in\{1,...,K\}$, $|A_{ik}|>1$ for at most one player $i\in N$, and
$\forall ik\in L$, τ_{ik}: $\{\theta_0\}\times A_{<k} \to T_{ik}$ is one to one.
(No two players make choices at the same date, and all players observe the past history of play.)

Theorem. Every dynamic multi-stage game of perfect information with $\Theta=\{\theta_0\}$ has a basic sequential equilibrium, (μ,ω), in which all prior probabilities are 0 or 1. That is, $\mu(\{\theta_0\}\times B)\in\{0,1\}$, $\forall B\subseteq A$.

(Proof Sketch: $\forall C$, $\exists S^o$ such that (C,S) has perfect information and a single state of Nature for every $\mathcal{U}(S)\supseteq\mathcal{U}(S^o)$. Hence, (C,S) has a pure strategy sequential equilibrium which, being pure, has the requisite property. The property therefore holds in a basic sequential equilibrium limit.)

Limitations of step-strategy approximations

Example: θ is uniform [0,1], player 1 observes θ, chooses $a_1 \in [0,1]$,
$v_1(\theta,a_1)=1$ if $a_1=\theta$, $v_1(\theta,a_1)=0$ if $a_1 \neq \theta$.
In any finite approximation, player 1's finite action set cannot allow any positive probability of c_1 matching θ exactly, and so 1's expected payoff must be 0.
Player 1 would like to use the strategy "choose $a_1=\theta$," but this strategy can be only approximated by step functions when 1 has finitely many feasible actions.
Step functions close to this strategy yield very different expected payoffs because the utility function is discontinuous.
If we gave player 1 an action that simply applied this strategy, he would use it!
(It may be hard to choose any real number exactly, but easy to say "I choose θ.")

Thus, adding a *strategic action* that implements a strategy which is feasible in the limit game can significantly change our sequential equilibria.
This problem follows from our principle of finitely approximating information and actions separately, which is needed to prevent spurious signaling.

Example (Akerlof): θ uniform [0,1], 1 observes $\tau_1=\theta$, 1 chooses $a_1 \in [0, 1.5]$,
2 chooses $a_2 \in \{0,1\}$, $v_1(\theta,a_1,1) = a_1-\theta$, $v_2(\theta,a_1,1) = 1.5\theta-a_1$, $v_i(\theta,a_1,0) = 0$.
Given any finite set of strategies for 1, we could always add some strategy b_1 such that $\theta < b_1(\theta) < 1.5\theta \;\; \forall\theta$, and $b_1(\theta)$ has probability 1 of being in a range that has probability 0 under the other strategies in the given set.

Finite additivity of solutions

Example: Consider a game where the winner is the player who chooses the smallest strictly positive number. In any finite approximation, let each player choose the smallest strictly positive number that is available to him.
In the limit, we get $\mu(0 \le a_1 \le x) = 1, \forall x>0$, but $\mu(a_1=0) = 0$.
The probability measure μ is not countably additive, only finitely additive, and this finite additivity lets us represent an infinitesimal positive action a_1.
Expected utilities are well-defined with finite additivity, as utility is bounded.

Suppose the players are $\{1,2\}$, and player 2 observes a_1 before choosing a_2.
For any possible value of a_1, 2 would always choose $a_2<a_1$ after observing this a_1 in any sufficiently large finite approximation. So in the limit, we get:
$\forall x>0$, $\mu(a_2<a_1|a_1=x) = 1$ and $\mu(0<a_1 \le x) = 1$.
But multiple equilibria can have any prior P(2 wins) = $\mu(a_2<a_1)$ in $\{0,1\}$.
By considering a subnet of finite approximations in which 1 always has smaller actions than 2, we can get an equilibrium with $\mu(a_2<a_1) = 0$ even though $\mu(a_2<a_1|a_1=x) = 1 \ \forall x>0$ and $\mu(a_1>0)=1$.
In this subnet, 1 always chooses a_1 where 2's strategy has not converged, and so the outcome cannot be derived by backward induction from 2's limiting strategy.

Strategic actions avoid such problems, as would considering subnets where later players' actions grow faster than earlier players' (fast inner limits for later $C_{\bullet k}$).

Strategic entanglement (Milgrom-Weber 1985)

Example: θ is uniform [0,1], observed by both players 1 and 2, who then simultaneously play a battle-of-sexes game where $A_1=A_2=\{1,2\}$,
$v_i(\theta,a_1,a_2) = 2$ if $a_1=a_2=i$, $v_i = 1$ if $a_1=a_2\neq i$, and $v_i = 0$ if $a_1\neq a_2$.
Consider finite-approximation equilibria such that, for a large integer m, they do $a_1=a_2=1$ if the m'th digit of θ is odd, $a_1=a_2=2$ if the m'th digit of θ is even.
As $m\to\infty$, these equilibria converge to a limit where the players randomize between actions (1,1) and (2,2), each with probability 1/2, independently of θ.
In the limit, the players' actions are not independent given θ in any positive interval. They are correlated by commonly observed infinitesimal details of θ.
Given θ in any positive interval, player 1's expected payoff is 1.5 in eqm,
but 1's sequential value of deviating to $a_1=1$ is $\omega_1(1) = 0.5(2)+0.5(0) = 1$,
and 1's sequential value of deviating to $a_1=2$ is $\omega_1(2) = 0.5(0)+0.5(1) = 0.5$.

Thus, the sequential-rationality inequalities can be strict for all actions.
To tighten the sequential-rationality lower bounds for conditional expected payoffs in equilibrium, we could consider also deviations of the form: "deviate to action c_{ik} when the equilibrium strategy would select an action in the set B_{ik}" for any $c_{ik}\in A_{ik}$ and any $B_{ik}\subseteq A_{ik}$.

In the next example, strategic entanglement is unavoidable, occurring in the only basic sequential equilibrium.

Strategic entanglement (cont'd) (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 $-|a_1|^2$.

Approximations in which 1's action set is {-1,…,-2/m,-1/m,1/m,2/m,…,1} have a unique subgame perfect (hence sequential) equilibrium in which player 1 chooses ±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.
Player 3's and player 4's strategies are entangled in the limit.

The limit of every approximation produces (the same) strategic entanglement.

Strategic entanglement, type 2 (Harris, Stinchcombe, Zame 2000)
Example: θ is uniform [0,1], payoff irrelevant, observed by both players 1 and 2, who then simultaneously play the following 2x2 game.

	L	R
L	1,1	3,2
R	2,3	0,0

Consider finite approximations in which, for large integers m, each player observes the first m-1 digits of the ternary expansion of θ.
Additionally, player 1 observes whether the m'th digit is 1 or not and player 2 observes whether the m'th digit is 2 or not.
Consider equilibria in which each player i chooses R if the m'th digit of the ternary expansion of θ is i and chooses L otherwise. As $m \to \infty$, these converge to a correlated equilibrium of the 2x2 game where each cell but (R,R) obtains with probability 1/3, independently of θ. Not in the convex hull of Nash equilibria!
The type of strategic entanglement ("type 2") generated here is impossible to generate in some approximations—it depends on the fine details—unlike the entanglement ("type 1") in the previous two examples.
Curiously, viewing this and the "Strategic Entanglement I" example as Bayesian games as in Milgrom-Weber (1985), the first type-1 entanglement can arise as a limit of Bayes-Nash equilibria, while the present type-2 entanglement cannot.

Spuriously exclusive coordination

Example: Three players choose simultaneously from [0,1].
Players 1 and 2 wish to match one another, receiving 1 if they do but 0 otherwise, and to mismatch player 3, each receiving an additional 1 if they individually mismatch 3 and an additional 0 otherwise.
Player 3 wishes to match player 1, receiving 1 if he matches and 0 otherwise.

Consider finite approximations in which players 1 and 2 share an action that is not available to player 3, e.g., players 1 and 2 always have available the action that is one-half the smallest positive action available to 3.
Then there are pure strategy equilibria in which 1 and 2 coordinate on that special common action and player 3 chooses any action, $a_3=1$ for example.

Hence, there is a basic sequential equilibrium (μ,ω) such that μ assigns probability one to any open set containing the action profile (0,0,1) and such that players i=1,2 each receive the equilibrium payoff $\int v_i(a)\mu(da)=2$, while player 3 receives the equilibrium payoff $\int v_3(a)\mu(da)=0$.

As with type 2 strategic entanglement, the exclusivity of the coordination occurring here depends on the fine details of the approximation.

Resolving type-2 strategic entanglement and spurious exclusion

To avoid the problems of type 2 strategic entanglement and spuriously exclusive coordination, one could insist that equilibria be robust to *all* approximations.
But this is not compatible with a general existence result (e.g., two players each wish to choose the smallest strictly positive number).
One resolution is to require a weaker form of robustness to the approximation.

Essential sequential equilibria

The set of *essential sequential equilibria* is the smallest closed set of assessments containing every set of assessments M that is minimal with respect to the following property: M is closed and there exists a selection $\sigma_{(\bullet)}$ such that
$\sigma_{(\varepsilon,C,S)} \in E(\varepsilon,C,S)$ for every (ε,C,S) and
$\lim_{\varepsilon>0} \lim_{C\subset A} (\lim_{S:(C,S)\in\mathcal{F}} \{(m(\sigma_{(\varepsilon,C,S)},C,S),w(\sigma_{(\varepsilon,C,S)},C,S))\}) \subseteq M$.

Theorem. The set of essential sequential equilibria is nonempty and is a closed subset of the set of basic sequential equilibria.

(The proof, using Zorn's lemma and Tychonoff's Theorem, establishes the existence of at least one such minimal set M that is nonempty.)

Another possible solution: Require that type sets and action sets which are "the same" must have the same finite approximations. (How to recognize sameness?)

Using strategic actions to solve problems of step-strategy approximations

One way to avoid the problems of step-strategies is to permit the use of strategic actions. This must be done with care to avoid the problem of spurious signaling.

Algebras (closed under finite $\cap$ and complements) of *measurable* subsets are specified also for each A_{ik}, including all one-point sets as measurable.

Each $\tau_{ik}(\theta,a_{<k})$ and each $v_i(\theta,a)$ now assumed jointly measurable in (θ,a).

If Θ is a finite set, then all subsets of A_{ik} are measurable.

A *strategic action* for $ik \in L$ is a measurable function, $a_{ik}{}^*{:}T_{ik} \rightarrow A_{ik}$.

Let $A_{ik}{}^*$ denote ik's set of strategic actions.

The set of *intrinsic actions* A_{ik} is a subset of $A_{ik}{}^*$ because for each $a_{ik} \in A_{ik}$, $A_{ik}{}^*$ contains the constant strategic action taking the value a_{ik} on T_{ik}.

For every $a^* \in A^*$ and every $\theta \in \Theta$, let $\alpha(\theta,a^*)$ denote the action profile $a \in A$ that is determined when the state is θ and each player plays according to the strategic action profile a^*. Note that for each $a^* \in A^*$, $\alpha(\theta,a^*)$ is measurable in θ.

To ensure perfect recall, define new type spaces and type maps:

$T_{ik}{}^* = T_{ik} \times (\times_{h<k} A_{ih}{}^*)$, with measurable subsets whose slices contained in T_{ik} defined by any $a_{<k}{}^*$ are measurable, and $\tau_{ik}{}^*(\theta,a_{<k}{}^*) = (\tau_{ik}(\theta,\alpha(\theta,a^*)_{<k}), \times_{h<k} a_{ih}{}^*)$.

Extend each $v_i(\theta,a)$ from $\Theta \times A$ to $\Theta \times A^*$ by defining $v_i(\theta,a^*) = v_i(\theta,\alpha(\theta,a^*))$.

Note that for every $ik \in L$ and every $a^* \in A^*$, $\tau_{ik}{}^*(\theta,a_{<k}{}^*)$ and $v_i(\theta,a^*)$ are measurable functions of θ.

This defines a dynamic multi-stage game, Γ*, with perfect recall and we may therefore employ all of the notation and definitions we previously developed. $\mathcal{Q}_{ik}$* is the set of observable events for player i at period k that can have positive probability in Γ*. $\mathcal{Q}$* is the (disjoint) union of these $\mathcal{Q}_{ik}$* over all ik.
Approximations of Γ* are defined, as before, by (C,S), where C is any finite set of (strategic) action profiles in A* and S is any information approximation of all the T_{ik}*.

Perfect recall in Γ*

In a finite approximation of Γ*, perfect recall means that a player ik remembers, looking back to a previous date m<k, the partition element containing his information-type at that date and the strategic action chosen there.
If the partition element contains more than one of his information types and his strategic action there is not constant across the projection of those types onto T_{im}, then he will not recall the action in A actually taken at the previous date.
Thus, there is perfect recall with respect to strategic actions but not with respect to intrinsic actions.
Of course, if the strategic action chosen is a constant function, i.e., equivalent to an intrinsic action, then the action too is recalled.

Misrepresentations and incentive compatibility

The Akerlof example shows that applying basic sequential equilibrium to Γ^* is inadequate as a solution to Γ because of spurious signaling.

To avoid spurious signaling we take a direct approach.

Given an approximation (C,S) of Γ^*, a *misrepresentation* for $ik \in L$ is a T_{ik}^*-measurable function $r_{ik}:T_{ik}^* \rightarrow T_{ik}^*$ such that for each $s_{ik} \in S_{ik}$, the range $r_{ik}(s_{ik})$ is contained in some element, $[r_{ik}(s_{ik})]$, of S_{ik}.

(Misrepresentations permit players' information types in Γ^* to mimic other of their information types.)

In the approximation (C,S), define $\sigma_{ik} \circ r_{ik}$ so that, for any information type s_{ik}, $\sigma_{ik} \circ r_{ik}$ assigns the probability $\sigma_{ik}(c_{ik}|[r_{ik}(s_{ik})])$ to the strategic action $c_{ik} \circ r_{ik} \in A_{ik}^*$.

That is $\sigma_{ik} \circ r_{ik}(c_{ik} \circ r_{ik}|s_{ik}) = \sigma_{ik}(c_{ik}|[r_{ik}(s_{ik})])$, $\forall c_{ik} \in A_{ik}^*$, $\forall s_{ik} \in S_{ik}$.

The strategic action $c_{ik} \circ r_{ik}$ specifies action $c_{ik}(r_{ik}(t_{ik}))$ for any type t_{ik} in s_{ik}.

Given an approximation (C,S) of Γ^* and $\varepsilon > 0$,

a totally mixed strategy σ is ε-*incentive compatible* iff

$V_i(\sigma_{-ik}, \sigma_{ik} \circ r_{ik}|s_{ik}) \leq V_i(\sigma|s_{ik}) + \varepsilon$, $\forall ik \in L$, $\forall s_{ik} \in S_{ik} \cap \mathcal{Q}_{ik}(C)$, $\forall$misrepresentatn r_{ik}.

Note. The containment condition on $r_{ik}(s_{ik})$ ensures that every ε-approximate sequential equilibrium of (C,S) is ε-incentive compatible whenever $C \subseteq A$.

Maximal incentive compatible sets of strategic actions

Any subset B of A* is *incentive compatible* iff $B \supseteq A$ and
for every $\varepsilon>0$ and every finite subset C^o of B,
there is a finite set C such that $C^o \subseteq C \subseteq B$, and such that
for every information approximation S^o,
there is an information approximation S such that $\mathcal{U}(S) \supseteq \mathcal{U}(S^o)$, $(C,S) \in \mathcal{F}$ and
(C,S) has an ε-incentive compatible ε-approximate sequential equilibrium.

The subset B of A* is *maximally incentive compatible* if it is incentive compatible and no subset of A* strictly containing B is also incentive compatible.

Note. Because B=A is incentive compatible and arbitrary unions of nested incentive compatible sets are incentive compatible, Zorn's lemma yields the existence of at least one maximally incentive compatible set.

Let $\mathcal{Y}^*$ be the set of all $Y \subseteq \Theta \times A^*$ such that $\forall a^* \in A^*$, $\{\theta | (\theta,a^*) \in Y\}$ is a measurable subset of Θ, and let $\mathcal{W}^* = \cup_{ik \in L} A_{ik}{}^* \times \mathcal{Q}_{ik}{}^*$.

Let $E^*(\varepsilon,C,S)$ denote the set of ε-incentive compatible ε-approximate sequential equilibria for the finite approximation (C,S) of Γ*. (For $\varepsilon>0$, if $B \subseteq A^*$ is incentive compatible, then $\{C \subseteq B | \{S | E^*(\varepsilon,C,S) \neq \emptyset, (C,S) \in \mathcal{F}\}$ is cofinal$\}$ is cofinal in B.)

To resolve the problem of step-strategies, we offer the following definition.

Extended sequential equilibria

The set of *extended sequential equilibria* of Γ is the smallest closed set of assessments in $[0,1]^{\mathcal{Y}^* \times \mathcal{Q}^*} \times [-\Omega, \Omega]^{\mathcal{W}^*}$ that contains, for every maximally incentive compatible subset B of A*, the set,
$\lim_{\varepsilon>0} \lim_{C\subset B} \lim_{S:(C,S)\in\mathcal{F}} \{(m(\sigma,C,S),w(\sigma,C,S))| \sigma\in E^*(\varepsilon,C,S)\}$.

(Note. E*(ε,C,S) may be empty for some (ε,C,S).)

Theorem. The set of extended sequential equilibria is nonempty and is a closed subset of $[0,1]^{\mathcal{Y}^* \times \mathcal{Q}^*} \times [-\Omega, \Omega]^{\mathcal{W}^*}$ with the product topology.

The proof is based on the previous observation that there is at least one nonempty maximally incentive compatible subset, B say, of A*.
Then, in particular, $\lim_{\varepsilon>0} \lim_{C\subset B} \lim_{S:(C,S)\in\mathcal{F}} \{(m(\sigma,C,S),w(\sigma,C,S))| \sigma\in E^*(\varepsilon,C,S)\}$ is nonempty, by Tychonoff's theorem.

(Note. If A is finite, there is an equivalence between the sets of extended and basic sequential equilibria of Γ because every a* is a step strategy.)

The following simultaneously resolves *all* of the problems identified above.

Essential extended sequential equilibria

The set of *essential extended sequential equilibria* of Γ is the smallest closed set of assessments in $[0,1]^{\mathcal{Y}^* \times \mathcal{Q}^*} \times [-\Omega, \Omega]^{\mathcal{W}^*}$ containing every set of assessments M that is minimal with respect to the following property: M is closed and there is a selection $\sigma_{(\bullet)}$ such that $\sigma_{(\varepsilon,C,S)} \in E^*(\varepsilon,C,S)$ whenever $E^*(\varepsilon,C,S) \neq \emptyset$, and such that $\lim_{\varepsilon>0} \lim_{C \subset B} \lim_{S:(C,S) \in \mathcal{F} \text{ and } E^*(\varepsilon,C,S) \neq \emptyset} \{(m(\sigma_{(\varepsilon,C,S)},C,S), w(\sigma_{(\varepsilon,C,S)},C,S))\} \subseteq M$ for every maximally incentive compatible $B \subseteq A^*$.

Theorem. The set of essential extended sequential equilibria of Γ is nonempty and is a closed subset of the set of extended sequential equilibria of Γ.

The proof, using Zorn's lemma and Tychonoff's Theorem, establishes the existence of at least one such minimal set M that is nonempty.

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.