Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 13.
Adaptive Play

Adaptive Play

We have so far applied the stochastic stability framework to single-population interactions, i.e. symmetric games.

We now examine a learning protocol based on Young (1993 Ecta, 1998) known as *adaptive play* which applies to:

- Possibly asymmetric games,
- Choices based on fragmented information about the *history* of play,
- Any population size, not just large N.

Population Game

- Two disjoint populations, *row* and *column*.
- Every period two players are selected, one from each population.
- They play a two-person game with strategy set X for the row player and Y for the column player.
- They sample from a bounded memory of the history of play between the two populations and (myopically) best respond to this information.

Fragmented Information

The history of play is a vector

$$h^t = \left((x^{t-m}, y^{t-m}), \ldots (x^{t-1}, y^{t-1})\right)$$

where x^{t-1} $\left(y^{t-1}\right)$ is the most recent choice by a row (column) player and m is the memory length.

- Each revising player at time t draws a random sample of size s (without replacement) from h^t, i.e. from the last m plays of the other population.

- The revising row player computes the frequency $p(y)$ of each action y in her sample. Similarly, for the column player.

Best Response Dynamic

Myopic best response:

- With high probability $1 - \varepsilon$, a revising agent best responds to its sample.
- With low probability ε, it chooses uniformly at random from its strategy set.

This produces a finite Markov chain on the state space $Z = (X \times Y)^m$, which is the set of truncated histories of length m.

The process $P^{m,s,\varepsilon}$ is called **adaptive play** with memory m, sample size s and error rate ε.

Convention

A *social convention* is a self-enforcing pattern of behavior:

- Everyone prefers to conform to the pattern when everyone else conforms.
- Once the convention is established, everyone believes everyone else will conform.

See Thomas Schelling The Strategy of Conflict (1960) and David Lewis Convention: A Philosophical Study (1967).

In the language of adaptive play, a convention is a history of the form

$$h^* = \big((x^*, y^*), \ldots (x^*, y^*)\big)$$

where (x^*, y^*) is a strict Nash equilibrium of the game.

As far as anyone can recall, row has always played x^* and column has always played y^*.

Absorbing States

Proposition 12.1 A state is an absorbing state of the unperturbed process $P^{m,s,0}$ if and only if it is a convention.

Proof. Suppose $h^T = h^* = \big((x^*, y^*), \ldots (x^*, y^*)\big)$.

All possible samples for a row player consist of s plays of y^*.

Since (x^*, y^*) is a strict Nash equilibrium of the game, the unique BR is x^*.

Similarly, for column.

Thus the convention is perpetuated. Iterating: $h^t = h^T = h^*$ for all $t > T$, so h^* is absorbing.

Now suppose $h^T = \big((x, y), \ldots (x, y)\big)$ is <u>not</u> a convention. Then there is a positive probability that either row or column player draws a sample to which x or y respectively is not a BR. Hence h^T is not an absorbing state.

Global Convergence

Proposition 12.2 Consider a coordination game. For $s \leq \frac{1}{2}m$, the unperturbed process converges almost surely to a convention from any initial state.

Proof.

- We claim there exists a probability at least $p > 0$ of transiting from any state to a convention in at most T periods.
- The probability of not transiting to a convention in λT periods is then at most $(1-p)^\lambda$ which goes to 0 as $\lambda \to \infty$.
- This would establish the proposition.

Global Convergence

Proof.

- To establish the claim, consider an arbitrary initial state:

$$\begin{array}{l} x^1 \; x^2 \; x^{m-s+1} \; \dots x^m \\ y^1 \; y^2 \; \underbrace{y^{m-s+1} \; \dots y^m}_{\hat{y}} \end{array}$$

- There is a positive probability that revising row players draw the sample $\hat{y} = (\hat{y}_1, \hat{y}_2, \dots \hat{y}_m)$ for the next s periods (because $s \leq \frac{1}{2}m$) and plays the same BR to it, say a, in each of these periods.

Global Convergence

Proof.

- This produces a history of the form:

$$\begin{array}{ccc} \dots & \dots & \underbrace{a\,a\dots a}_{\hat{a}} \\ \dots & \underbrace{\hat{y}_1\,\hat{y}_2\,\dots\hat{y}_m}_{\hat{y}} & \dots \end{array}$$

- Now row can sample $\hat{y}$ once more and play a, while column can sample $\hat{a}$ and play the best response a (as this is a coordination game).

Global Convergence

Proof.

- This produces a history of the form:

$$\begin{array}{ccc} \ldots & a & \underbrace{a\,a\ldots a}_{\hat{a}} \\ \ldots & \ldots & a \end{array}$$

- This history contains at least as many plays of a by column as the history one period before.
- Therefore, row can still draw a sample to which the BR is a.
- Also, column can draw sample $\hat{a}$ and best respond with a.
- Iterating this argument leads to convergence to a convention with positive probability in finite time, as claimed.

Stochastic Stability in 2×2 Games

Propositions 12.1 and 12.2 together imply that the conventions are the only recurrence classes of the unperturbed process.

Which ones are stochastically stable?

Theorem 12.3. Consider adaptive play in a 2×2 coordination game with sufficiently incomplete information ($s/m \leq 1/2$).

For s and m sufficiently large, the unique stochastically stable state is the risk-dominant convention.

Risk Dominance

- A (possibly) asymmetric 2×2 coordination game:

	A	B
A	$\underline{a_1}$, $\underline{a_2}$	$\underline{c_1}$, b_2
B	b_1, $\underline{c_2}$	$\underline{d_1}$, $\underline{d_2}$

  with $a_i > b_i$ and $d_i > c_i$ for $i = 1, 2$.

  Definition. (A, A) *risk dominates* (B, B) if:

$$(a_1 - b_1)(a_2 - b_2) \geq (d_1 - c_1)(d_2 - c_2),$$

  i.e. the product of the deviation losses is at least as high for (A, A) (Harsanyi and Selten 1988).

- If the inequality is strict then (A, A) strictly risk dominates (B, B).

Risk Dominance & Stochastic Stability

Proof. Mixed NE determines tipping point for each player.

Let p (q) be the proportion of row (column) players choosing A.

The mixed NE is given by:

$$p^* = \frac{d_1 - c_1}{(d_1 - c_1) + (a_1 - b_1)} \;,\; q^* = \frac{d_2 - c_2}{(d_2 - c_2) + (a_2 - b_2)}.$$

Define recurrence classes E_1 as 'All A' ($p = q = 1$) and E_2 as 'All B' ($p = q = 0$).

$$r(E_2, E_1) = \lceil s \times \mathsf{min}(p^*, q^*) \rceil \;,\; r(E_1, E_2) = \lceil s \times \mathsf{min}(1 - p^*, 1 - q^*) \rceil.$$

For s sufficiently large we can ignore integer issues.

Risk Dominance & Stochastic Stability

Proof. Hence E_1 is stochastically stable if and only if

$$\min(p^*, q^*) \leq \min(1 - p^*, 1 - q^*). \tag{1}$$

Case 1. $p^* \leq q^*$: Then (1) becomes $p^* \leq 1 - q^*$ or

$$\begin{aligned}
\frac{d_1 - c_1}{(d_1 - c_1) + (a_1 - b_1)} &\leq \frac{a_2 - b_2}{(d_2 - c_2) + (a_2 - b_2)} \\
(d_1 - c_1)[(d_2 - c_2) + (a_2 - b_2)] &\leq (a_2 - b_2)[(d_1 - c_1) + (a_1 - b_1)] \\
(d_1 - c_1)(d_2 - c_2) &\leq (a_2 - b_2)(a_1 - b_1),
\end{aligned}$$

i.e. if E_1 is risk dominant.

Case 2. $p^* > q^*$: Then (1) becomes $q^* \leq 1 - p^*$ or $p^* \leq 1 - q^*$, as before. So the same argument applies.

Stochastic Stability in Larger Games

	E_1	E_2	E_3
E_1	5, 60	0, 0	0, 0
E_2	0, 0	7, 40	0, 0
E_3	0, 0	0, 0	100, 1

Risk dominance and stochastic stability can differ:

- min resistance E_1-tree: $\frac{1}{61} + \frac{7}{107}$.
- min resistance E_2-tree: $\frac{5}{105} + \frac{1}{41}$.
- min resistance E_3-tree: $\frac{5}{105} + \frac{7}{107}$.

E_1 is risk dominant, but E_2 is stochastically stable (because of indirect paths in the spanning tree).

Stochastic Stability in Larger Games

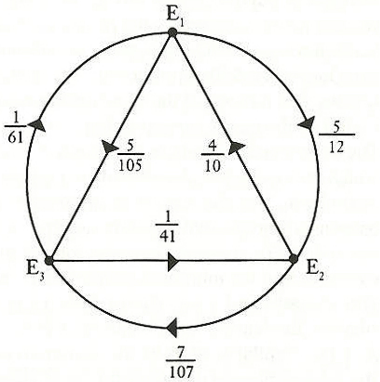

Risk dominance and stochastic stability can differ:

- min resistance E_1-tree: $\frac{1}{61} + \frac{7}{107}$.
- min resistance E_2-tree: $\frac{5}{105} + \frac{1}{41}$.
- min resistance E_3-tree: $\frac{5}{105} + \frac{7}{107}$.

Stochastic Stability in Larger Games

	E_1	E_2	E_3
E_1	2, 6	0, 1	1, 0
E_2	1, 0	3, 4	0, 1
E_3	0, 1	1, 0	4, 3

There is no guarantee a risk dominant equilibrium exists:

- E_3 risk dominates E_2: $(4-0)(3-0) > (3-1)(4-1)$.
- E_2 risk dominates E_1: $(3-1)(4-1) > (2-1)(6-1)$.
- E_1 risk dominates E_3: $(2-0)(6-0) > (4-1)(3-1)$.

Stochastic Stability in Larger Games

An additional issue:

- So far, the calculation of resistances has been simple.
  - To calculate the transition $E_2 \to E_1$, we ask how many deviations to E_1 are needed for it to become a BR to switch from E_2 to E_1.

- When off-diagonal payoffs are nonzero, we not only need to account for indirect paths through the spanning tree (when calculating the stochastic potential of each recurrence class), but also the fact that the least cost path between recurrence classes may be indirect:
  - To calculate the transition $E_2 \to E_1$, we also need ask what is the minimum number deviations of all kinds (including to E_3) needed for it to become a BR to switch to E_1.

Weak Acyclicity & Convergence

Definition. A game is *weakly acylic* if from every node there exists a directed path in the best reply graph to a sink (i.e. a node from which there is no exiting edge).

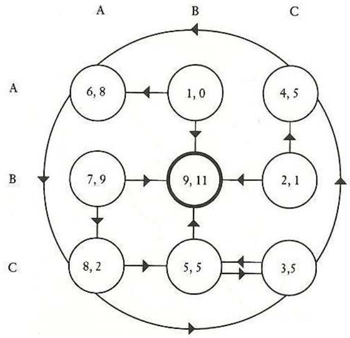

Theorem 12.4. Consider an n-person game that is weakly acyclic. If s/m is sufficiently small, then the unperturbed process $P^{m,s,0}$ converges with probability one to a convention from any initial state and the stochastically stable states are the conventions with minimum stochastic potential.

Cyclic Games

- What if the recurrence classes of the process include not only absorbing states but limit cycles?
- Consider the following (two-population) game:

	A	B	C	D
A	3, 3	4, 1	1, 4	-1, -1
B	1, 4	3, 3	4, 1	-1, -1
C	4, 1	1, 4	3, 3	-1, -1
D	-1, -1	-1, -1	-1, -1	0, 0

- Under the best response dynamic, the game produces a single absorbing state (D, D) and a best-response cycle involving A, B and C.

Stochastically Stable Limit Cycles

- The resistances of the transitions between (D, D) and the cycle are calculated as the path of least resistance from (D, D) to some state in the cycle and *vice versa*.
- The weakest point in the cycle is where A, B and C occur with the same frequency.
- For N large enough $\lceil \frac{1}{6}N \rceil < \lceil \frac{11}{14}N \rceil$ and hence the limit cycle is stochastically stable.
- Intuitively, as $\varepsilon \to 0$ it is far easier to enter the cycle from the sink (D, D) than it is to exit from the cycle to the sink.