Learning—Adjustment with persistent noise

14.126 Game Theory
Mihai Manea
Muhamet Yildiz

Main idea

- There will always be small but positive probability of mutation.
- Then, some of the strict Nash equilibria will **not** be "stochastically stable."

General Procedure

Stochastic Adjustment

1. Consider a game.
2. Specify a state space Θ, e.g., the # of players playing a strategy.

1.

	A	B
A	2,2	0,0
B	0,0	1,1

2. Θ = {AA,AB,BA,BB}

Stochastic Adjustment, continued

3. Specify an adjustment dynamics, e.g., best-response dynamics, with a transition matrix P, where

   $P_{\theta,\xi}$ = Pr(θ at t+1|ξ at t)

   ϕ = a probability distribution, a column vector.

3.

P =

	AA	AB	BA	BB
AA	1	0	0	0
AB	0	0	1	0
BA	0	1	0	0
BB	0	0	0	1

Stochastic Adjustment, continued

4. Introduce a small noise: Consider P^ε, continuous in ε and

   $P^\varepsilon \to P$ as $\varepsilon \to 0$.

   Make sure that there exist a unique ϕ_ε^* s.t.

   $\phi_\varepsilon^* = P^\varepsilon \phi_\varepsilon^*$.

4.

P^ε =

	AA	AB	BA	BB
AA	$(1-\varepsilon)^2$	$(1-\varepsilon)\varepsilon$	$(1-\varepsilon)\varepsilon$	ε^2
AB	$(1-\varepsilon)\varepsilon$	ε^2	$(1-\varepsilon)^2$	$(1-\varepsilon)\varepsilon$
BA	$(1-\varepsilon)\varepsilon$	$(1-\varepsilon)^2$	ε^2	$(1-\varepsilon)\varepsilon$
BB	ε^2	$(1-\varepsilon)\varepsilon$	$(1-\varepsilon)\varepsilon$	$(1-\varepsilon)^2$

$\phi_\varepsilon^* = (1/4,1/4,1/4,1/4)^T$.

Stochastic Adjustment, continued

5. Verify that $\lim_{\varepsilon\to 0}\phi_\varepsilon^* = \phi^*$ exists; compute ϕ^*. (By continuity $\phi^* = P\phi^*$.)
6. Check that ϕ^* is a point mass, i.e.,

$$\phi^*(\theta^*) = 1$$

for some θ^*.

The strategy profile at θ^* is called *stochastically stable equilibrium*.

Kandoori, Mailath & Rob

Coordination game

	A	B
A	(2,2)	(0,0)
B	(0,0)	(1,1)

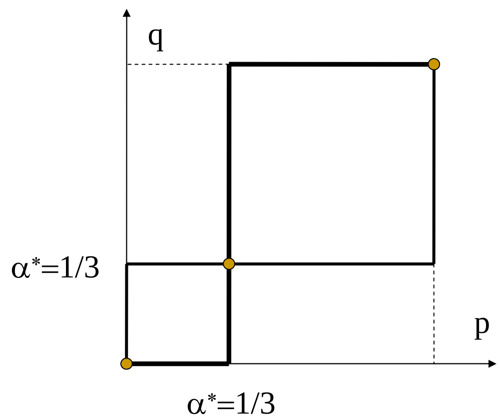

Adjustment Process

- N = population size.
- θ_t = # of players who play A at t.
- $u_A(\theta_t) = \theta_t/N\ u(A,A) + (N-\theta_t)/N\ u(A,B)$
- $\theta_{t+1} = P(\theta_t)$, where

$$P(\theta_t) > \theta_t \Leftrightarrow u_A(\theta_t) > u_B(\theta_t)\ \&$$
$$P(\theta_t) = \theta_t \Leftrightarrow u_A(\theta_t) = u_B(\theta_t).$$

- Example:

$$P(\theta_t) = BR(\theta_t) = \begin{cases} N \text{ if } u_A(\theta_t) > u_B(\theta_t) \\ \theta_t \text{ if } u_A(\theta_t) = u_B(\theta_t) \\ 0 \text{ if } u_A(\theta_t) < u_B(\theta_t) \end{cases}$$

Noise

- Independently, each agent with probability 2ε mutates, and plays either of the strategies with equal probabilities.

$$P^{\varepsilon} = \begin{bmatrix} (1-\varepsilon)^N & (1-\varepsilon)^N & \dots & (1-\varepsilon)^N & \varepsilon^N & \dots & \varepsilon^N \\ N(1-\varepsilon)^{N-1}\varepsilon & N(1-\varepsilon)^{N-1}\varepsilon & \dots & N(1-\varepsilon)^{N-1}\varepsilon & N(1-\varepsilon)\varepsilon^{N-1} & \dots & N(1-\varepsilon)\varepsilon^{N-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ N(1-\varepsilon)\varepsilon^{N-1} & N(1-\varepsilon)\varepsilon^{N-1} & \dots & N(1-\varepsilon)\varepsilon^{N-1} & N(1-\varepsilon)^{N-1}\varepsilon & \dots & N(1-\varepsilon)^{N-1}\varepsilon \\ \varepsilon^N & \varepsilon^N & \dots & \varepsilon^N & (1-\varepsilon)^N & \dots & (1-\varepsilon)^N \end{bmatrix}$$

- $\varphi^*(\varepsilon)$ = invariant distribution for P^{ε}.

P^{ε}

$$\left[\begin{array}{ccc|ccc} (1-\varepsilon)^N & \cdots & (1-\varepsilon)^N & \varepsilon^N & \cdots & \varepsilon^N \\ N(1-\varepsilon)^{N-1}\varepsilon & \cdots & N(1-\varepsilon)^{N-1}\varepsilon & N(1-\varepsilon)\varepsilon^{N-1} & \cdots & N(1-\varepsilon)\varepsilon^{N-1} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \binom{N}{N^*-1}(1-\varepsilon)^{N-N^*+1}\varepsilon^{N^*-1} & \cdots & \binom{N}{N^*-1}(1-\varepsilon)^{N-N^*+1}\varepsilon^{N^*-1} & \binom{N}{N^*-1}(1-\varepsilon)^{N^*-1}\varepsilon^{N-N^*+1} & \cdots & \binom{N}{N^*-1}(1-\varepsilon)^{N^*-1}\varepsilon^{N-N^*+1} \\ \hline \binom{N}{N^*}(1-\varepsilon)^{N-N^*}\varepsilon^{N^*} & \cdots & \binom{N}{N^*}(1-\varepsilon)^{N-N^*}\varepsilon^{N^*} & \binom{N}{N^*}(1-\varepsilon)^{N^*}\varepsilon^{N-N^*} & \cdots & \binom{N}{N^*}(1-\varepsilon)^{N^*}\varepsilon^{N-N^*} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \varepsilon^N & \cdots & \varepsilon^N & (1-\varepsilon)^N & \cdots & (1-\varepsilon)^N \end{array}\right]$$

$N^* = \lceil a^*N \rceil < N/2$.

Invariant Distribution

- $N^* = \lceil \alpha^* N \rceil < N/2.$
- $D_A = \{\theta | \theta \geq N^*\}$; $D_B = \{\theta | \theta < N^*\}$;
- $q_{AB} = Pr(\theta_{t+1} \in D_A | \theta_t \in D_B)$;
- $q_{BA} = Pr(\theta_{t+1} \in D_B | \theta_t \in D_A)$;
- $$p_A(\varepsilon) = \sum_{\theta \in D_A} \varphi_\varepsilon^*(\theta)$$
- $p_B(\varepsilon) = 1 - p_A(\varepsilon)$.

Invariant distribution, continued

- $$\begin{bmatrix} p_A(\varepsilon) \\ p_B(\varepsilon) \end{bmatrix} = \begin{bmatrix} 1 - q_{BA} & q_{AB} \\ q_{BA} & 1 - q_{AB} \end{bmatrix} \begin{bmatrix} p_A(\varepsilon) \\ p_B(\varepsilon) \end{bmatrix}$$
- $$\begin{bmatrix} -q_{BA} & q_{AB} \\ q_{BA} & -q_{AB} \end{bmatrix} \begin{bmatrix} p_A(\varepsilon) \\ p_B(\varepsilon) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$
- $$\frac{p_A(\varepsilon)}{p_B(\varepsilon)} = \frac{q_{AB}}{q_{BA}}$$

Invariant distribution, continued

$$\frac{p_A(\varepsilon)}{p_B(\varepsilon)}=\frac{q_{AB}}{q_{BA}}$$

$$=\frac{\binom{N}{N^*}\varepsilon^{N^*}(1-\varepsilon)^{N-N^*}+\binom{N}{N^*+1}\varepsilon^{N^*+1}(1-\varepsilon)^{N-N^*-1}+\cdots}{\binom{N}{N^*-1}\varepsilon^{N-N^*+1}(1-\varepsilon)^{N^*-1}+\binom{N}{N^*-2}\varepsilon^{N-N^*+2}(1-\varepsilon)^{N^*-2}+\cdots}$$

$$=\frac{\binom{N}{N^*}\varepsilon^{N^*}(1-\varepsilon)^{N-N^*}+o\left(\varepsilon^{N^*}\right)}{\binom{N}{N^*-1}\varepsilon^{N-N^*+1}(1-\varepsilon)^{N^*-1}+o\left(\varepsilon^{N-N^*+1}\right)}\cong\frac{\binom{N}{N^*}}{\binom{N}{N^*-1}}\frac{1}{\varepsilon^{N-2N^*+1}}\to\infty.$$

Proposition

If N is large enough so that N* < N/2, then limit φ^* of invariant distributions puts a point mass on $\theta = N$, corresponding to all players playing A.

Replicator dynamics
&
Evolutionary stability

14.126 Game Theory
Muhamet Yildiz

Road Map

1. Evolutionarily stable strategies
2. Replicator dynamics

Notation

- $G = (S,A)$ a symmetric, 2-player game where
  - S is the strategy space;
  - $A_{i,j} = u_1(s_i,s_j) = u_2(s_j,s_i)$.
- x,y are mixed strategies; $u(x,y) = x^{T}Ay$; $u(s,y)$.
- $ax+(1-a)y$.
- $u(ax+(1-a)y,z) = au(x,z) +(1-a)u(y,z)$
- $u(x,ay+(1-a)z) = au(x,y) +(1-a)u(x,z)$

ESS

Definition: A (mixed) strategy x is said to be *evolutionarily stable* iff, given any $y \neq x$, there exists $\varepsilon_y > 0$ s.t.

$$u(x,(1-\varepsilon)x+\varepsilon y) > u(y,(1-\varepsilon)x+\varepsilon y)$$

for each ε in $(0,\varepsilon_y]$.

- Each player is endowed with a (mixed) strategy.
- Assumes that population is a state
- Asks whether a strategy (state) is robust to evolutionary pressures.
- Disregards effects on future actions.

Alternative Definition

Fact: x is evolutionarily stable iff, $\forall y \neq x$,

1. $u(x,x) \geq u(y,x)$, and
2. $u(x,x) = u(y,x) \Rightarrow u(x,y) > u(y,y)$.

<u>Proof:</u> Define

$$F(\varepsilon,y) = u(x,(1-\varepsilon)x+\varepsilon y) - u(y,(1-\varepsilon)x+\varepsilon y)$$
$$= u(x-y, (1-\varepsilon)x+\varepsilon y)$$
$$= (1-\varepsilon)\, u(x-y,x) + \varepsilon\, u(x-y,y).$$

Hawk-Dove game




	Hawk	Dove
Hawk	(1-c,1-c)	(2,0)
Dove	(0,2)	(1,1)

1. $c < 1$
2. $c > 1$

ESS-NE

- If *x* is an ESS, then (*x*,*x*) is a Nash equilibrium.
- In fact, (*x*,*x*) is a proper equilibrium.
- If (*x*,*x*) is a strict Nash equilibrium, then *x* is ESS.

Rock-Scissors-Paper

	R	S	P
R	0,0	1,-1	-1,1
S	-1,1	0,0	1,-1
P	1,-1	-1,1	0,0

- Unique Nash Equilibrium (s^*,s^*) where

  $s^* = (1/3,1/3,1/3)$
- s^* is not ESS.

ESS in role-playing games

- Given (S^1, S^2, u_1, u_2), consider the symmetric game $(\underline{S}, u)$ where
  - $\underline{S} = S^1 \times S^2$;
  - $u(\underline{x}, \underline{y}) = [u_1(x_1, y_2) + u_2(x_2, y_1)]/2 \ \forall \underline{x} = (x_1, x_2), \underline{y} = (y_1, y_2) \in \underline{S}$.

Theorem: $\underline{x}$ is an ESS of $(\underline{S}, u)$ iff $\underline{x}$ is a strict Nash equilibrium of (S^1, S^2, u_1, u_2).

Replicator Dynamics

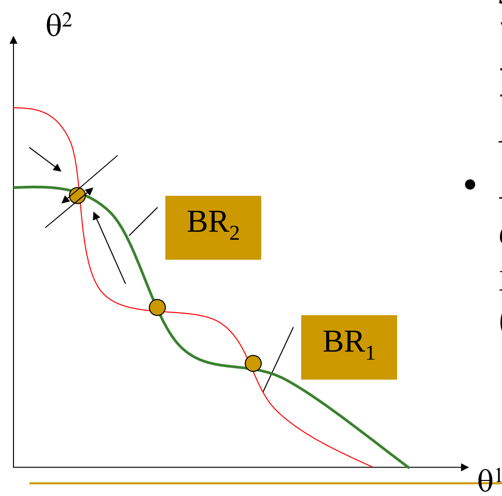

- Steady state: $F(\theta) = \theta$
- A steady state θ of F is *stable* iff, $\forall$neighborhood U of θ, $\exists$nbrd U_1 s.t., if $\theta_0 \in U_1$, $F_t(\theta_t) \in U$ $\forall t>0$.
- A steady state θ of F is *asymptotically stable* iff, it is stable and $\exists U$: $\theta_0 \in U \Rightarrow \lim_t F_t(\theta_t) = \theta$.

Replicator dynamics

- $p_i(t)$ = #people who plays s_i at t;
- $p(t)$ = total population at t.
- $x_i(t) = p_i(t)/p(t)$; $x(t) = (x_1(t),\dots, x_k(t))$.
- $u(x,x) = \Sigma_i\, x_i u(s_i,x)$.
- Birthrate for s_i at t is $\beta + u(s_i,x(t))$; death rate=δ.
- $\dot{p}_i = [\beta + u(s_i, x) - \delta]p_i$
- $\dot{p} = [\beta + u(x, x) - \delta]p$
- $\dot{x}_i = [u(s_i, x) - u(x, x)]x_i$
- $\dot{x}_i = u(s_i - x, x)x_i$

Example

- Consider (S,A) where $A = \begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix}$
- $u(s_1,x) = a_1 x_1$;
- $u(x,x) = (x_1,x_2)A(x_1,x_2)^T = a_1 x_1^2 + a_2 x_2^2$
- $u(s_1-x,x) = (a_1 x_1 - a_2 x_2) x_2$
- $$\dot{x}_1 = (a_1 x_1 - a_2 x_2) x_1 x_2$$

Examples

- Replicator dynamics in prisoners' dilemma
- Replicator dynamics in chicken
- Replicator dynamics in the battle of the sexes.

Observations

- $\frac{d}{dt}\left[\frac{x_i}{x_j}\right] = \frac{\dot{x}_i}{x_j} - \frac{x_i}{x_j}\frac{\dot{x}_j}{x_j} = [u(s_i,x) - u(x,x)]\frac{x_i}{x_j} - \frac{x_i}{x_j}[u(s_j,x) - u(x,x)]\frac{x_j}{x_j}$

$$= [u(s_i,x) - u(s_j,x)]\frac{x_i}{x_j}$$

- If u becomes $\underline{u} = au+b$, then Replicator dynamics becomes

$$\dot{x}_i = \underline{u}(s_i - x, x)x_i = au(s_i - x, x)x_i$$

Rationalizability

- $\xi(.,x_0)$ is the solution to replicator dynamics starting at x_0.

Theorem: If a pure strategy i is strictly dominated (by y), then $\lim_t \xi_i(t,x_0) = 0$ for any interior x_0.

Proof: Define $v_i(x) = \log(x_i) - \Sigma_j y_j \log(x_j)$. Then,

$$\frac{dv_i(x(t))}{dt} = \frac{\dot{x}_i}{x_i} - \sum_j y_j \frac{\dot{x}_j}{x_j} = u(s_i - x, x) - \sum_j y_j u(s_j - x, x) = u(s_i - y, x).$$

Hence, $v_i(x(t)) \to -\infty$., i.e., $x_i(t) \to 0$.

Theorem: If i is not rationalizable, then $\lim_t \xi_i(t,x_0) = 0$ for any interior x_0.

Theorems

Theorem: Every ESS x is an asymptotically stable steady state of replicator dynamics.

(If the individuals can inherit the mixed strategies, the converse is also true.)

Theorem: If x is an asymptotically stable steady state of replicator dynamics, then (x,x) is a perfect Nash equilibrium.