Lecture 8

Wednesday, 14 September 2022 1:51 PM

Dynamics in Games:

A particular dynamics is a behavioural model for players, i.e., how they should play over time.
I.e., For $t \in [0, \infty)$, dynamics specifies the strategy $x(t)$, $y(t)$ played by the two players.

Ideally would like:

(i) the dynamics should "converge" quickly to a stable state

(ii) the players should not require excessive information (e.g., other player's payoffs) to follow the dynamics

(iii) player actions should be "natural": not play obviously bad strategies, not solve computationally hard problems

(iv) ideally players should play pure strategies in each round, i.e., $x(t) = e_i$ for $i \in [n]$

Dynamics also referred to as learning in games, since players "learn" stable strategies.

E.g., a bad dynamics would be that one player knows all payoffs, computes a MNE σ, & tells all other players to play according to σ...

2-Player Game: Fictitious Play.

Let x, y be the two players, S_x, S_y be the pure strategies. In FP, at each time step, each player plays a pure strategy $x(t)$, $y(t)$. Think of each as a vector in $\Delta_{|S(x)|}$

Define $\bar{x}(t) = \frac{1}{t+1}\sum_{\tau=0}^{t} x(\tau)$

$\bar{y}(t) = \frac{1}{t+1}\sum_{\tau=0}^{t} y(\tau)$

This is a mixed strategy for both players, or the empirical distribution of pure strategies played until time t.

Fictitious Play:

- in round zero, each player picks an arbitrary pure strategy & plays it. $x(1) \in S_x$, $y(1) \in S_y$.
  (technically, $x(1) = e_i$ for some $i \in |S_x|$,
  $y(1) = e_i$ for some $i \in |S_y|$)
- in round t, each player picks a pure best-response to the empirical distribution of pure strategies played until round $t-1$

$$x(t) \in \arg\max_{s_x \in S_x} u_x(s_x, \bar{y}(t-1))$$

$$y(t) \in \arg\max_{s_y \in S_y} u_y(s_y, \bar{x}(t-1))$$

How does this do on our list of desiderata?

? not sure if it converges, or to what

\+ players only need to know their own payoffs & what others have played

\+ easy to compute, no coordination required

\+ pure strategies

An Example:

	L	R
U	3\3	0\0
D	4\0	1\1

(D, R) obtained by IEDS, hence only NE.

Say $x(0) = U$, $y(0) = L$

	$x(t)$	$y(t)$	$\bar{x}(t)$	$\bar{y}(t)$
$t=0$	U	L	$(1,0)$	$(1,0)$
$t=1$	D	L	$(\frac{1}{2},\frac{1}{2})$	$(1,0)$
$t=2$	D	L	$(\frac{1}{3},\frac{2}{3})$	$(1,0)$
$t=3$	D	L	$(\frac{1}{4},\frac{3}{4})$	$(1,0)$
$t=4$	D	L	$(\frac{1}{5},\frac{4}{5})$	$(1,0)$
$t=5$	D	R	$(\frac{1}{6},\frac{5}{6})$	$(\frac{5}{6},\frac{1}{6})$
$t=6$	D	R	$(\frac{1}{7},\frac{6}{7})$	$(\frac{5}{7},\frac{2}{7})$

...

In this case, FP seems to converge to a pure NE.

Example 2:

	L	R
U	0\0	1\1
D	1\1	0\0

Note that game has 2 pure NE & 1 mixed NE

Say $x(0) = U$, $x(0) = L$

	$x(t)$	$y(t)$	$\bar{x}(t)$	$\bar{y}(t)$
$t=0$	U	L	$(1,0)$	$(1,0)$
$t=1$	D	R	$(\frac{1}{2},\frac{1}{2})$	$(\frac{1}{2},\frac{1}{2})$
$t=2$	D	R	$(\frac{1}{3},\frac{2}{3})$	$(\frac{1}{3},\frac{2}{3})$
$t=3$	U	L	$(\frac{1}{2},\frac{1}{2})$	$(\frac{1}{2},\frac{1}{2})$
$t=4$	U	L	$(\frac{3}{5},\frac{2}{5})$	$(\frac{3}{5},\frac{2}{5})$
$t=5$	D	R	$(\frac{1}{2},\frac{1}{2})$	$(\frac{1}{2},\frac{1}{2})$

...

So, doesn't converge to a pure NE in this case.

Defn A sequence of pure strategies $(x(t), y(t))$ converges to (x^*, y^*) if $\exists T : \forall t \geq T$, $x(t) = x^*$, $y(t) = x^*$

(thus for each coordinate i, $\lim_{t\to\infty} \bar{x}_i(t) \to x_i^*$

$\lim_{t\to\infty} \bar{y}_j(t) \to y_j^*$)

OR, if $(x(t), y(t))$ converges to (x^*, y^*), then $(\bar{x}(t), \bar{y}(t))$ converges to (x^*, y^*) in the typical sense
$\forall \varepsilon \; \exists T : \|y^* - \bar{y}(t)\|_\infty \leq \varepsilon \quad \forall t \geq T$

Thus, if $(x(t), y(t))$ converges to (x^*, y^*), then
$\forall \varepsilon \; \exists T \; \forall t \geq T$, $\|Ry^* - R\bar{y}(t)\|_\infty \leq \varepsilon$
$\forall \varepsilon \; \exists T \; \forall t \geq T$, $\|Cx^* - C\bar{x}(t)\|_\infty \leq \varepsilon$

Lemma: (i) If the sequence of pure strategies $(x(t), y(t))$ in FP converges to (x^*, y^*), then (x^*, y^*) is a PNE

(ii) If (x^*, y^*) is a strict PNE, and $\exists t : (x(t), y(t)) = (x^*, y^*)$, then $(x(\tau), y(\tau)) = (x^*, y^*) \; \forall \tau \geq t$

Proof: (i) Assume for contradiction $(x(t), y(t)) \to (x^*, y^*)$ but (x^*, y^*) is not NE. $\exists i, j : x_j^* > 0$, $(Ry^*)_i - (Ry^*)_j = \varepsilon > 0$

- we know that $\exists T_0$ s.t. $\forall t \geq T_0$, $\|Ry^* - R\bar{y}(t)\|_\infty < \varepsilon/4$
  $\Rightarrow |(Ry^*)_i - (R\bar{y}(t))_i| < \varepsilon/4$
  and $|(Ry^*)_j - (R\bar{y}(t))_j| < \varepsilon/4$

  it follows that $(R\bar{y}(t))_i - (R\bar{y}(t))_j \geq \varepsilon/2 \quad \forall t \geq T_0$
  but this is a contradiction, since if $(R\bar{y}(t))_i > (R\bar{y}(t))_j$, $x_j(t+1) = 0 \quad \forall t \geq T_0$. But $x(t) \to x^*$, and $x_j^* > 0$ by assumption.

(ii) $x(t) = x^*$, $y(t) = y^*$

Thus $x^* R\bar{y}(t-1) \geq \hat{x} R\bar{y}(t-1) \quad \forall \hat{x} \in S_x$

$x^* R y^* > \hat{x} R y^* \quad \forall \hat{x} \in S_x \setminus x^*$

$\Rightarrow x^* R\bar{y}(t) > \hat{x} R\bar{y}(t) \quad \forall \hat{x} \in S_x \setminus x^*$

Hence, $x(t+1) = x^*$, $y(t+1) = y^*$

Lemma: If the sequence $(\bar{x}(t), \bar{y}(t))$ converges to (x^*, y^*), then (x^*, y^*) is a MNE.

Proof: Say (x^*, y^*) is not MNE.
Then $\exists i, j \in [m] : x_j^* > 0$ but $(Ry^*)_i > (Ry^*)_j$
Say $\varepsilon = (Ry^*)_i - (Ry^*)_j > 0$
$\exists T : \forall t \geq T$, $\|Ry^* - R\bar{y}(t)\| \leq \varepsilon/4$

$\Rightarrow |(Ry^*)_i - (R\bar{y}(t))_i| \leq \varepsilon/4$

$\& \; |(Ry^*)_j - (R\bar{y}(t))_j| \leq \varepsilon/4$

Thus, $(R\bar{y}(t))_i - (R\bar{y}(t))_j \geq \varepsilon/2$

But then $\forall t \geq T$, $x_j(t) = 0$

This contradicts that $\bar{x}(t) \to x^*$, since $x_j^* > 0$.

Theorem: FP converges under the following conditions:

(i) G is a 2-player 0-sum game

(ii) G is solvable under strict IEDS

(iii) G is an exact potential game.

Coarse - Correlated Equilibrium

So far, we've thought of a MNE $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ as a tuple of mixed strategies.
This defines a distribution over pure strategy profiles:

$$\sigma(s = (s_1, \dots, s_n)) = \prod_{i=1}^{n} \sigma_i(s_i)$$

Thus, σ is a product distribution over pure strategy profiles. Since it is a product distribution, the players pick their strategies from their individual distributions, w/o coordinating with other players.

Suppose however we consider distributions over S that are not product distributions.

So now, $\sigma \in \Delta_{|S|}$ is a distribution over all pure strategies.

$\forall i, t_i \in S_i, \quad \sigma(t_i) = \sum_{s : s_i = t_i} \sigma(s)$

$\forall i, t_{-i} \in S_{-i}, \quad \sigma(t_{-i}) = \sum_{s : s_{-i} = t_{-i}} \sigma(s) = \sum$

Example: Traffic Game

	S	G
S	0\0	0\10
G	10\0	-100\-100

3 equilibria: (S,G), (G,S) and MNE $\left(\frac{10}{11}, \frac{1}{11}\right)$

Realistically, one would expect one player to stop half the time, & the other player to stop the other half.

So, a good distribution would be: $\sigma(S, G) = 1/2$, $\sigma(G, S) = 1/2$.

This is obviously not possible as a MNE.

Defn: $\sigma \in \Delta_{|S|}$ is a Coarse Correlated Equilibria (CCE), aka a Hannan Equilibria, if $\forall i, \forall \hat{s}_i \in S_i$,

$$u_i(\sigma) \geq \sum_{\substack{s \in S \\ s = (s_i, s_{-i})}} u_i(\hat{s}_i, s_{-i}) \cdot \sigma(s) = \sum_{s_{-i} \in S_{-i}} \sigma(s_{-i})\, u_i(\hat{s}_i, s_{-i})$$

One interpretation is that, given σ, my utility for playing according to σ is at least as much as if I played a pure strategy, and others according to the distribution.

However, in a product distribution, each player independently chooses a pure strategy, which determines the distribution. If distribution is not a product distribution, how do we achieve this? This seems to assume some correlation among the players.

Recall the traffic game. In real-life, we do have a correlating device: a traffic signal.

Thus, a CCE σ corresponds to the following
- a distribution σ over S is fixed and known to all.
- a correlating device picks a pure strategy s from this distribution
- without knowing s, each agent i chooses either to:
  (i) play a pure strategy s_i', or
  (ii) learn s_i & play it.

If σ is a CCE, each player chooses the second option.

Note: every MNE is a CCE.

Finally, a CCE can be easily computed in polytime:

Let σ_s be a variable for all $s \in S$.

Then $\forall i, \forall s_i'$, $\sum_{s \in S} u_i(s)\,\sigma_s \geq \sum_{s \in S} u_i(s_i', s_{-i})\,\sigma_s$

$\forall s \in S \quad \sigma_s \geq 0$

$\sum_s \sigma_s = 1$