2016 - Problem set - Evolutionary Game Theory

Learning in Games

Let $\mathcal{S}^i$ the set of pure strategies of player i in a game of interest. We are interested in analyzing **learning procedures/ behavior rules**, that is:

Definition 1 $\rho^i : H \to \Delta(\mathcal{S}^i)$

where H is the set of all the possiblle histories generated by the game.

Fictitious Play

We need some definitions: Player i has an initial weight function $\kappa_0^i : \mathcal{S}^{-i} \to \mathbb{R}_+$. This weight is updated by adding 1 to the weight of each opponent strategy each time it is played, so that:

$$\kappa_t^i(s^{-i}) = \kappa_{t-1}^i(s^{-i}) + \begin{cases} 1 & \text{if} \quad s_{t-1}^{-i} = s^{-i} \\ 0 & \text{if} \quad s_{t-1}^{-i} \neq s^{-i} \end{cases}$$

The probability that player i assigns to player $-i$ playing s^{-i} at date t is given by:

$$\gamma_t^i(s^{-i}) = \frac{\kappa_t^i(s^{-i})}{\sum_{\tilde{s}^{-i} \in \mathcal{S}^{-i}} \kappa_t^i(\tilde{s}^{-i})}$$

We define **Fictitious Play** in the following way:

Definition 2 *Fictitious Play is defined as any rule* $\rho_t^i(\gamma_t^i)$ *that assigns* $\rho_t^i(\gamma_t^i) \in BR^i(\gamma_t^i)$

(remember the Best Response definition:)

Definition 3 $u^i(.,s^{-i})$ *is strictly concave. The best response of player i to a profile, denoted* $BR^i(s^{-i})$*, is*

$$BR^i(s^{-i}) = arg\ max_{\tilde{s}^i} u^i(\tilde{s}^i, s^{-i})$$

Consider the game Matching Pennies:

		Player 2: H	Player 2: T
Player 1	H	$(1,-1)$	$(-1,1)$
	T	$(-1,1)$	$(1,-1)$

Question 1

Assume that the "initial weight", that is, the initial priors/ samples that the 2 players hold are the following: they have observed the following outcomes:

- (H,T): 2 realizations
- (T,H): 1.5 realizations

Deduce the initial samples of opponent's play for each player. Describe the observed outcomes for the following 7 rounds, assuming that both players follow Ficititious Play.

Convergence of Fictitious Play? .

<u>**Question 2**</u>

Find the unique (mixed-strategy) Nash Equilibrium of this game.

Definition 4 *The Marginal Empirical Distribution of j's play is*

$$d_t^j(s^j) = \frac{\kappa_t(s^j) - \kappa_0(s^j)}{t}$$

What are the Marginal Empirical Distribution of each player is converging to? (as $t \to \infty$). Compare those with the mixed-strategy Nash Equilibrium.

<u>**Question 3**</u>

Does it mean that Fictitious Play Marginal Empirical Distribution always converge to a Nash Equilibrium, as in the previous example? Let us study the Shapley example:

	L	M	R
T	$0,0$	$1,0$	$0,1$
M	$0,1$	$0,0$	$1,0$
D	$1,0$	$0,1$	$0,0$

We assume that this game has a unique mixed-strategy Nash Equilibrium, for each player to play $(1/3, 1/3, 1/3)$.

Assume that the initial weights are such that players choose (T, M) in the first round. What are the following 7 rounds, if both players use fictitious play? Does it help you say whether in the long run, the marginal empirical distributions converge to the unique Nash Equilibrium of the game?

Utility-Based criteria to assess learning procedures We have seen in class that Fictitious Play does not necessarily satisfies Universal Consistency (definition in the lecture slides). We want to see if we can find another learning procedure that satisfies it.

Smooth Fictitious Play

We study Smooth Fictitious Play in Matching Pennies, for player 1. We remind his payoff table (we do not care about player 2 now):

$$\begin{array}{cc} & \begin{array}{cc} H & T \end{array} \\ \begin{array}{c} H \\ T \end{array} & \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \end{array}$$

We aim at studying Smooth Fictitious Play using the entropy smoothing function seen in class. It boils down to the following best response function for player 1:

$\forall s^1 \in \mathcal{S}^1$, given the strategy profile of the other player(s) is σ^{-1}:

$$BR^1(\sigma^{-1})[s^1] = \frac{exp\big[(1/\lambda)u^1(s^1, \sigma^{-1})\big]}{\sum_{r^1 \in \mathcal{S}^1} exp\big[(1/\lambda)u^1(r^1, \sigma^{-1})\big]}$$

This is the exponential smooth Best Response. In Smooth Fictitious play, σ^{-i} are the frequency of play up to time t: γ_t^i.

We remind here the definition of ϵ- universal consistency:

Definition 5 *A rule ρ^i is ϵ-universally consistent if for any ρ^{-i}*

$$\limsup_{T \to \infty} \max_{\sigma^i \in \Delta(\mathcal{S}^i)} u_i(\sigma^i, \gamma_t^i) - \frac{1}{T} \sum_t u^i(\rho_t^i(h_{t-1})) \leq \epsilon \textit{ almost surely}$$

We wish to check if this property holds for the learning rule specified above, against a **specific rule** of the opponent: assume player 2 plays H when t is odd and T when t is even.

Question 4

Assume player 1 initial weights are $(1, \frac{1}{2})$. Compute the smooth Fictitious Play Best Response of Player 1 (that is, which weights he puts on playing H and T) in odd and in even periods.

Does the $\epsilon-$consistency requirement holds at least against this one specific rule of the opponent for some $\epsilon > 0$?

Limit of Smooth Fictitious Play and Calibration We saw in class that Smooth Ficititious Play does not solve the problem of Convergence, as the simulations of the Shapley example have shown. We want to find a calibrated version of the smooth fictitious play studied in the previous question.

Question 5

Take the learning procedure studied before (exponential smooth fictitious play in the matching penny game) and write the equations characterizing a calibrated version of this rule.

Comment Calibration helps for convergence in this specific example, see simulations in class, but not in general, see Jordan example.

Evolutionary Theory

Consider the following game:

$$\begin{array}{c c} & \begin{array}{ccc} 1 & 2 & 3 \end{array} \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \begin{pmatrix} 6,6 & 0,5 & 0,0 \\ 5,0 & 7,7 & 5,5 \\ 0,0 & 5,5 & 8,8 \end{pmatrix} \end{array}$$

Question 6

What are the 3 strict pure strategy Nash Equilibria of this game?

Dynamics Assume that you have a population of 101 agents (you have only one type of agents, the game is symmetric). The state $s \in \mathbb{R}^3$ is a $3-$ dimensional vector that describes the proportion of agents that played each of the 3 possible pure strategies in the previous period (thus the entries of the vector sum up to 1).

Each period, 2 agents are chosen at random to play, and they do in the following way (with $\epsilon > 0$ given):

- With probability $(1 - \epsilon)$: players choose at random (equal weights) between the possible Best Response to the state s.
- With probability ϵ they choose at random between any of the 3 pure strategies.

Then the state updates.

Question 7

We do not want to write the transition matrix of the Markov Chain because the state space is too big, but we assume that the model has 3 ergodic classes:

- $A \equiv$ everybody playing 1
- $B \equiv$ everybody playing 2

- $C \equiv$ everybody playing 3

(note that each of those ergodic classes is a singleton, a single state, but it does not have to be the case in general). Create a table with the least resistance between each ergodic class (you have to consider ALL possible paths from one to another to find the path with least resistance).

Question 8

Using the table you created in the previous question, analyze all the possible trees and determine which are the least resistance trees, and thus, conclude about which classes are stochastically stable states.