Notes on Self-Confirming Equilibrium
Course Econ 201B

Rodolfo G. Campos
UCLA

February 1, 2005

1 Intuition - Big Picture

Suppose that we are interested in an equilibrium concept which describes the equilibrium strategies of unsophisticated players who learn to play a game. These players are not very sophisticated in the sense that they don't use equilibrium concepts such as subgame perfection when conjecturing how other players will play. Instead, they form beliefs about how other players are going to play at their information sets. These beliefs may be completely unwarranted. The players' opponents do not necessarily play optimally in these beliefs. The players themselves are unsophisticated but not stupid, meaning that they do have to play optimally at their own information sets given their beliefs.

Now assume that players are able to learn from previous rounds of the game. Learning means that players observe how the game was played in the past and adjust their beliefs to match the observed play by their opponents. But this adjustment only takes place at those information sets that are reached in the course of the game. Beliefs at all other information sets never adjust to the thruth and may be anything. We may ask ourselves which equilibrium notion might be able to describe an equilibrium of such a learning process. The answer to this question is a Self-confirming equilibrium.

2 Notation

To define a Self-comfirming Equilibrium (SCE) we need to define beliefs about the opponents' play for each player. Hence, we need to introduce new notation. Further, it is necessary to distinguish between information sets that are reached in equilibrium play and those that are not, adding to the notational complexity.

To introduce the notation, it is useful to consider an example (Fudenberg and Kreps, 1988) of a 3 player game.

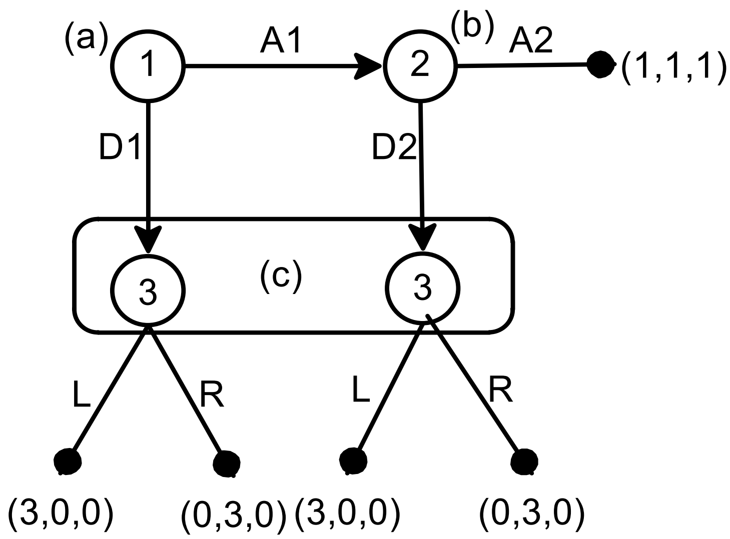

Example 1

To illustrate some of the definitions I will also use a modification of this game. A 2-player version of the game where player 1 plays at the lower information set is shown in the next figure.[1]

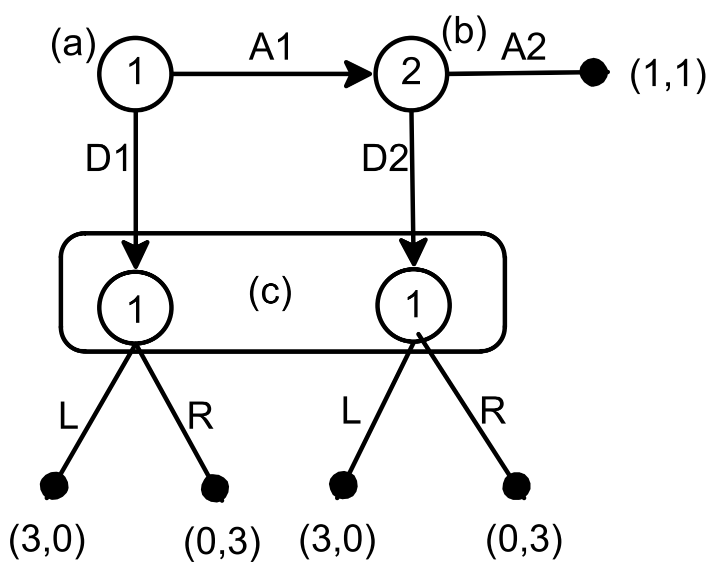

Example 2

[1] This game does not exhibit perfect recall, meaning that player 1 forgets what he played at the first information set. But this shouldn't trouble us too much since this version of the game will only be used as an example for notation.

2.1 Strategies

Pure strategies for player i are denoted $s_i \in S_i$. In the first example, $S_1 = \{A_1, D_1\}$, $S_2 = \{A_2, D_2\}$ and $S_3 = \{L, R\}$. Mixed strategies are denoted $\sigma_i \in \sum_i$. In the first example, $\sigma_1 = (\sigma_1(A_1), \sigma_1(D_1))$, $\sigma_2 = (\sigma_2(A_2), \sigma_2(D_2))$, $\sigma_3 = (\sigma_3(L), \sigma_3(R))$. In the second example $S_1 = \{A_1L, A_1R, D_1L, D_1R\}$ and $S_2 = \{A_2, D_2\}$. Mixed strategies are $\sigma_1 = (\sigma_1(A_1L), \sigma_1(A_1R), \sigma_1(D_1L), \sigma_1(D_1R))$ and $\sigma_2 = (\sigma_2(A_2), \sigma_2(D_2))$.

2.2 Information sets

Information sets for player i are denoted H_i. In the first example, $H_1 = \{a\}$, $H_2 = \{b\}$, $H_3 = \{c\}$. Each player has only one information set at which he is called to play. In the second example, $H_1 = \{a, c\}$, $H_2 = \{b\}$. To denote an element of H_i we use h_i. In the second example where player 1 plays at nodes a and c we could have $h_1 = a$ or $h_1 = c$ while for player 2 necessarily $h_2 = b$.

2.3 Information sets reached with positive proability

The notation for information sets reached with positive probability under strategy profile σ is $\bar{H}(\sigma)$. In Example 1 consider the strategy profile

$$\tilde{\sigma} = (\sigma_1(A_1), \sigma_1(D_1), \sigma_2(A_2), \sigma_2(D_2), \sigma_3(L), \sigma_3(R)) = (1, 0, 1, 0, p, 1-p) \tag{1}$$

with $p \in [0, 1]$. This means that player 1 is playing A_1, player 2 is playing A_2 and player 3 is doing anything. In this case, only information sets a and b are reached. Hence, $\bar{H}(\tilde{\sigma}) = \{a, b\}$. If $\hat{\sigma}$ is such that player 1 is playing D_1 with probability 1 (with the other players doing anything) then $\bar{H}(\hat{\sigma}) = \{a, c\}$. If all players mix then $\bar{H}(\sigma) = H_1 \cup H_2 \cup H_3 = \{a, b, c\}$.

2.4 Behavior strategies

A behavior strategy for player i is denoted $\pi_i \in \Pi_i$. In Example 1 behavior strategies are the same as mixed strategies because every player plays only once, thus $\pi_1 = (\sigma_1(A_1), \sigma_1(D_1))$, $\pi_2 = (\sigma_2(A_2), \sigma_2(D_2))$, $\pi_3 = (\sigma_3(L), \sigma_3(R))$. In Example 2 a behavior strategy for player 1 is $\pi_1 = (\pi_1(A_1), \pi_1(D_1), \pi_1(L), \pi_1(R))$ (compare this to player 1's mixed strategy) and for player 2 $\pi_2 = (\sigma_2(A_2), \sigma_2(D_2))$. Sometimes it will be useful to denote a behavior strategy at an information set $h_i \in H_i$ as $\pi_i(h_i)$. In the last example, $\pi_1(a) = (\pi_1(A_1), \pi_1(D_1))$, $\pi_1(c) = (\pi_1(L), \pi_1(R))$ and $\pi_1 = (\pi_1(a), \pi_1(c))$.

2.5 A map from mixed to behavior strategies

Denote the behavior strategy that player i uses at information set h_i under mixed strategy σ_i as $\hat{\pi}(h_i|\sigma_i)$. By Kuhn's Theorem, this mapping is well defined. Take Example 2. Consider

$$\sigma_1 = (\sigma_1(A_1L), \sigma_1(A_1R), \sigma_1(D_1L), \sigma_1(D_1R)) = \left(\frac{1}{6}, \frac{1}{2}, \frac{1}{3}, 0\right) \tag{2}$$

Then $\hat{\pi}(a|\sigma_1) = \left(\frac{2}{3}, \frac{1}{3}\right)$ and $\hat{\pi}(c|\sigma_1) = \left(\frac{1}{2}, \frac{1}{2}\right)$, meaning that under this mixed strategy, player 1 plays A_1 with probability $\frac{2}{3}$ at node a and L with probability $\frac{1}{2}$ at information set c. We can define $\hat{\pi}(\sigma_i)$ as the vector containing all $\hat{\pi}(h_i|\sigma_i)$ for every $h_i \in H_i$. In the example, $\hat{\pi}(\sigma_1) = \left(\frac{2}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}\right)$. We can further aggregate information by defining $\hat{\pi}(\sigma)$ as the vector containing $\hat{\pi}(\sigma_i)$ for all players. In our example, assume that for player 2 we have $\sigma_2 = (\sigma_2(A_2), \sigma_2(D_2)) = \left(\frac{1}{2}, \frac{1}{2}\right)$. Then $\hat{\pi}(\sigma_2) = \left(\frac{1}{2}, \frac{1}{2}\right)$ and

$$\hat{\pi}(\sigma) = (\hat{\pi}(\sigma_1), \hat{\pi}(\sigma_2)) = \left(\frac{2}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right) \tag{3}$$

2.6 Distribution over terminal nodes

A distribution over terminal nodes derived from a strategy profile is denoted as $\hat{\rho}(\pi)$. In a finite game this distribution is a vector with as many elements as terminal nodes exist, and where the sum of elements is one. It can be constructed from a mixed strategy by setting $\hat{\rho}(\sigma) = \hat{\rho}(\hat{\pi}(\sigma))$. Enumerate terminal notes as z_1, z_2, z_3, z_4, z_5 as in the figure. Then

$$\hat{\rho}(\pi) = (\rho(z_1|\pi), \rho(z_2|\pi), \rho(z_3|\pi), \rho(z_4|\pi), \rho(z_5|\pi)) \tag{4}$$

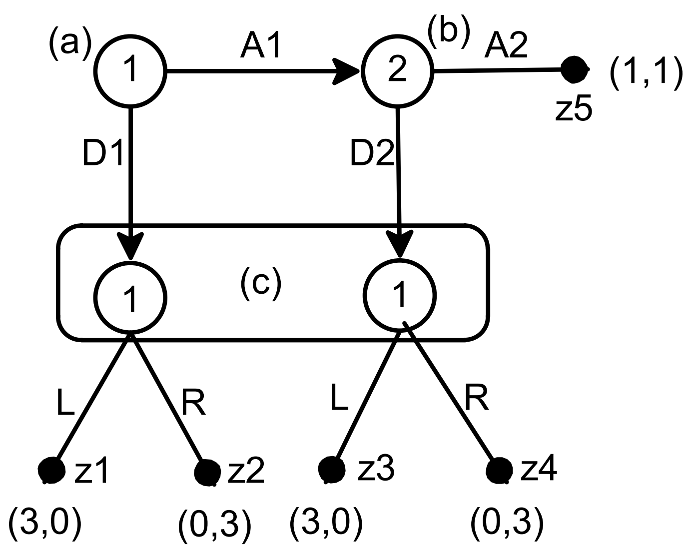

Example 2 with terminal nodes

To calculate $\hat{\rho}(\sigma)$ it is necessary to calculate the probabilities of reaching each terminal node under the mixed strategy profile σ. If we use $\hat{\pi}(\sigma) =$

$(\hat{\pi}(\sigma_1), \hat{\pi}(\sigma_2)) = \left(\frac{2}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right)$, as before, then $\rho(z_1) = \rho(z_2) = \frac{1}{3} \times \frac{1}{2} = \frac{1}{6}$. $\rho(z_3) = \rho(z_4) = \frac{2}{3} \times \frac{1}{2} \times \frac{1}{2} = \frac{1}{6}$. $\rho(z_5) = \frac{2}{3} \times \frac{1}{2} = \frac{1}{3}$. Hence,

$$\hat{\rho}(\sigma) = \hat{\rho}(\hat{\pi}(\sigma)) = \left(\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{3}\right) \tag{5}$$

2.7 Beliefs and Expected utility

To model the beliefs about opponents′ play, the notation μ_i is used. It is convenient to think of μ_i as a probability measure over the actions at each information set in Π_{-i}.[2] In Example 2, player 1′s beliefs are $\mu_1 = (\mu_1(A_2), \mu_1(D_2))$ and player 2′s beliefs are $\mu_2 = (\mu_2(A_1), \mu_2(D_1), \mu_2(L), \mu_2(R))$. Sometimes it will be useful to denote beliefs at some information set $h_{-i} \in H_{-i}$ as $\mu_i(h_{-i})$. For example, in Example 2, $\mu_2(a) = (\mu_2(A_1), \mu_2(D_1))$, $\mu_2(c) = (\mu_2(L), \mu_2(R))$ and $\mu_2 = (\mu_2(a), \mu_2(c))$. Notice that by defining beliefs in this way, they are formally equivalent to a behavior strategy of the opponents. μ_2 is formally equivalent to π_1 in the example.

Preferences can be expressed as expectations given beliefs

$$u_i(\pi_i|\mu_i) = \sum_z u_i(z)\,\hat{\rho}(z|\pi_i, \mu_i) \tag{6}$$

This formula takes the expectation of payoffs at terminal nodes, calculating the probabilities of reaching each node from the distribution that arises from considering a behavior strategy profile $\pi = (\pi_i, \mu_i)$ that uses the player′s beliefs about opponents′ play. In Example 2, if player 1 believes that player 2 is going to play D_2, i.e. $\mu_1 = (0, 1)$, and he plays $\pi_1 = (\pi_1(A_1), \pi_1(D_1), \pi_1(L), \pi_1(R)) = \left(1, 0, \frac{1}{2}, \frac{1}{2}\right)$, then $u_1(\pi_1|\mu_1) = 3 \times 1 \times 1 \times \frac{1}{2} + 0 \times 1 \times 1 \times \frac{1}{2} = \frac{3}{2}$. Expected utility can be extended to mixed strategies as well in the sense that

$$u_i(\sigma_i|\mu_i) = u_i(\hat{\pi}(\sigma_i)|\mu_i) \tag{7}$$

2.8 Observation Function

We need some notation for what each player is allowed to conjecture about what other players are playing. The following nasty-looking expression serves this purpose

$$\Pi_{-i}(\sigma_{-i}|J) \equiv \{\pi_{-i}|\pi_j(h_j) = \hat{\pi}(h_j|\sigma_j), \forall h_j \in H_{-i} \cap J\} \tag{8}$$

This expression defines the subset of behavior strategies consistent with the fact that player i′s opponents are playing σ_{-i} at the information sets in J. J is called the observation function. The intuition is that if players observe what happens at the information sets in J, then their conjectures about the strategies of their opponents should lie in the set $\Pi_{-i}(\sigma_{-i}|J)$.

[2] This definition of beliefs differs from the definition given in the lecture notes. In the lecture notes beliefs are a probability measure over Π_{-i}. The way I defined them, beliefs are elements of Π_{-i}.

In Example 2 consider $J = \{a\}$ and procced with the calculation of $\Pi_{-i}(\sigma_{-i}|J)$ for player 2. $\Pi_{-2}(\sigma_{-2}|\{a\}) = \Pi_{-2}(\sigma_1|\{a\})$

$$\Pi_{-2}(\sigma_1|\{a\}) \equiv \{\pi_1|\pi_1(h_1) = \hat{\pi}(h_1|\sigma_1), \forall h_1 \in \{a,c\} \cap \{a\}\} \tag{9}$$

The only $h_1 \in \{a,c\} \cap \{a\}$ is $h_1 = a$. Thus,

$$\Pi_{-2}(\sigma_1|\{a\}) = \{\pi_1|\pi_1(a) = \hat{\pi}(a|\sigma_1)\} \tag{10}$$

If player 1 is playing the mixed strategy $\sigma_1 = (\sigma_1(A_1L), \sigma_1(A_1R), \sigma_1(D_1L), \sigma_1(D_1R)) = \left(\frac{1}{6}, \frac{1}{2}, \frac{1}{3}, 0\right)$ then $\hat{\pi}(a|\sigma_1) = \left(\frac{2}{3}, \frac{1}{3}\right)$ (this was calculated before). Therefore,

$$\Pi_{-2}(\sigma_1|\{a\}) = \left\{\pi_1|\pi_1(a) = \left(\frac{2}{3}, \frac{1}{3}\right)\right\} \tag{11}$$

$$\Pi_{-2}(\sigma_1|\{a\}) = \left\{\left(\frac{2}{3}, \frac{1}{3}, x, 1-x\right), x \in [0,1]\right\} \tag{12}$$

So, what does this mean? It means that if player 1 is using the strategy $\sigma_1 = \left(\frac{1}{6}, \frac{1}{2}, \frac{1}{3}, 0\right)$ and player 2 only observes information set a, then the only restriction he can put on player 1's behavior strategy is that he is playing with probabilities $\left(\frac{2}{3}, \frac{1}{3}\right)$ at node a. He can't say anything about what behavior strategy is used by 1 at node c.

Alternatively, we might wonder what happens if we set $J = \{c\}$.[3] In that case we get

$$\Pi_{-2}(\sigma_1|\{c\}) = \left\{\left(x, 1-x, \frac{1}{2}, \frac{1}{2}\right), x \in [0,1]\right\} \tag{13}$$

Finally, if $J = \{a, c\}$ then

$$\Pi_{-2}(\sigma_1|\{a,c\}) = \left\{\left(\frac{2}{3}, \frac{1}{3}, \frac{1}{2}, \frac{1}{2}\right)\right\} \tag{14}$$

consists of a single element.

3 Definitions of equilibrium concepts

Now that we have chewed through the notation we can define equilibrium concepts with the new notation.

Definition 1 *(Nash Equilibrium) A Nash Equilibrium is a mixed profile σ such that for all i there exist beliefs μ_i such that*

1. $u_i(\sigma_i|\mu_i) \geq u_i(s_i'|\mu_i)$, for all $s_i' \in S_i$

2. $\mu_i \in \Pi_{-i}(\sigma_{-i}|H)$

In terms of the lecture notes, condition 2 hat to be written as $\mu_i(\Pi_{-i}(\sigma_{-i}|H)) = 1$.[4] The definition of a Nash Equilibrium implies that beliefs have to be correct

[3] We can do the calculation even though in our example it doesn't make much sense to think that player 2 observes c and does not observe a.

[4] Recall that my definition of beliefs treats them as elements of Π_{-i} instead of a probability distribution over Π_{-i}.

at all information sets.

Definition 2 *(Unitary SCE) A Unitary SCE is a mixed profile σ such that for all i there exist beliefs μ_i such that*

1. $u_i(\sigma_i|\mu_i) \geq u_i(s_i'|\mu_i)$, *for all* $s_i' \in S_i$

2′. $\mu_i \in \Pi_{-i}\left(\sigma_{-i}|\bar{H}(\sigma)\right)$

In the lecture notes 2′ is written $\mu_i\left(\Pi_{-i}\left(\sigma_{-i}|\bar{H}(\sigma)\right)\right) = 1$. The definition of a Unitary SCE implies that beliefs have to be correct at all information sets that are reached with positive probability under the equilibrium profile σ.

Definition 3 *(Heterogeneous SCE) A Heterogeneous SCE is a mixed profile σ such that for all i there exist beliefs μ_i for every $s_i \in$support(σ_i), such that*

1. $u_i(\sigma_i|\mu_i) \geq u_i(s_i'|\mu_i)$, *for all* $s_i' \in S_i$

2″. $\mu_i \in \Pi_{-i}\left(\sigma_{-i}|\bar{H}(s_i, \sigma_{-i})\right)$

In the lecture notes 2″ is written $\mu_i\left(\Pi_{-i}\left(\sigma_{-i}|\bar{H}(s_i, \sigma_{-i})\right)\right) = 1$. The definition of a Heterogeneous SCE implies that beliefs have to be correct at all information sets that are reached with positive probability if the profile of play is (s_i, σ_{-i}) where s_i is an element in the support of σ_i. The key idea of this equilibrium concept is that when we consider a mixed strategy for player i we treat each pure strategy in its support as if it was a different player i playing it. This means that the beliefs associated with pure strategy $s_i' \in$ support(σ_i) might differ from the beliefs associated with $s_i'' \in$ support(σ_i).

Notice that all three definitions are basically the same and only differ in the expression

$$\Pi_{-i}(\sigma_{-i}|J) \tag{15}$$

where the observation function takes values $J = H, \bar{H}(\sigma), \bar{H}(s_i, \sigma_{-i})$. Thus, the equilibrium concepts are, in a way, indexed by the observation function.

4 Examples

4.1 Nash Equilibrium vs. Unitary SCE

Consider the game in Example 1.

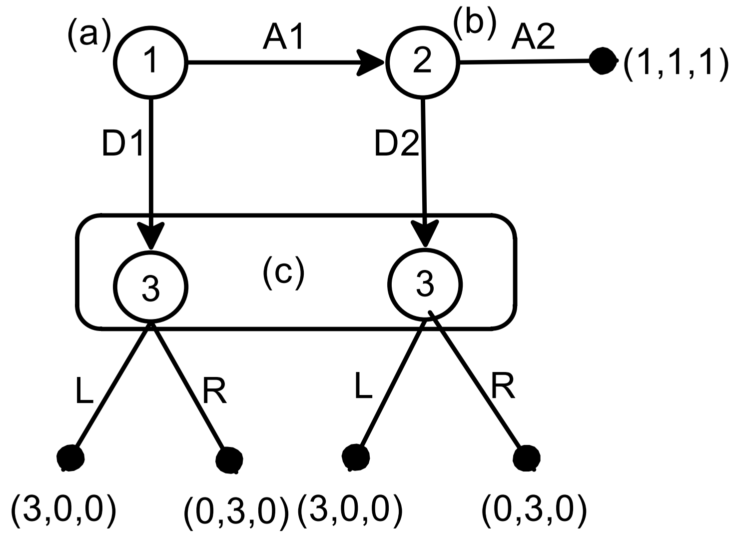

Example 1

In this game there is a Unitary SCE in which players 1 and 2 play across. Strategies are $\sigma_1 = (\sigma_1(A_1), \sigma_1(D_1)) = (1, 0)$, $\sigma_2 = (\sigma_2(A_2), \sigma_2(D_2)) = (1, 0)$, and $\sigma_3 = (\sigma_3(L), \sigma_3(R)) = (x, 1-x)$ with $x \in [0, 1]$. Beliefs which make these strategies optimal are $\mu_1 = (\mu_1(A_2), \mu_1(D_2), \mu_1(L), \mu_1(R)) = (1, 0, 0, 1)$,

$\mu_2 = (\mu_2(A_1), \mu_2(D_1), \mu_2(L), \mu_2(R)) = (1, 0, 1, 0)$ and

$\mu_3 = (\mu_3(A_1), \mu_3(D_1), \mu_3(A_2), \mu_3(D_2)) = (1, 0, 1, 0)$. These beliefs are correct on the equilibrium path $\bar{H}(\sigma) = \{a, b\}$. Beliefs of players 1 and 2 differ at information set c, which is not on the equilibrium path. Further, there is no Nash equilibrium with players 1 and 2 playing across. To check this, construct the normal form of the game and check the best responses.

Player 1 plays A_1

	L	R
A_2	1,**1**,**1**	**1**,1,**1**
D_2	**3**,0,**0**	**0**,**3**,0

Player 1 plays D_1

	L	R
A_2	**3**,**0**,**0**	0,**3**,**0**
D_2	**3**,**0**,**0**	**0**,**3**,**0**

The bold numbers show best the responses for all three players. As can be seen, there is no Nash Equilibrium where both players play across.

4.2 Unitary vs. Heterogeneous SCE

Consider the following game.

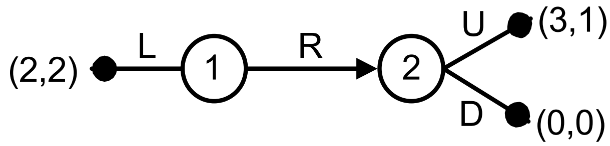

Selten Game

In this game there is a unitary SCE in which player 1 plays L, believing that player 2 plays D with at least some probability, and another one in which player 1 plays R and player 2 plays U. There is no Unitary SCE in which player 1 mixes. However, there exist Heterogeneous SCE of this type. Take $\sigma_1 = \left(\frac{1}{2}, \frac{1}{2}\right)$ and $\sigma_2 = (1, 0)$. There exist beliefs for each incarnation of player 1 such that playing a pure strategy in the support is optimal. For example, for the player who plays L beliefs are $\mu_1^L = (0, 1) \neq \sigma_2$, which is allowed since the information set at which player 2 plays is not in $\bar{H}(L, \sigma_2)$. Beliefs for the incarnation who plays right have to be correct, $\mu_1^R = (1, 0) = \sigma_2$.

References

[1] Fudenberg, D and Levine, D.K. (1993), Self-Confirming Equilibrium, Econometrica, 61: 523-546.

[2] Fudenberg, D and Levine, D.K. (1993), Steady State Learning and Nash Equilibrium, Econometrica, 61: 547-573.

[3] Fudenberg, D and Levine, D.K. (1998), The Theory of Learning in Games, MIT Press.

[4] Rahman, D. (2003), Notes on Sequential and Self Confirming Equilibrium, unpublished manuscript, UCLA.