Cursed Equilibrium: a Theoretical Reinterpretation and Its Learning Foundation*

Astrid Gamba†

Preliminary version
May 26, 2009

Abstract

Eyster and Rabin (2005) introduce the notion of *Cursed Equilibrium,* which assumes that players in a strategic setting underrate the informational content of opponents' play, that is they underestimate the degree to which opponents' actions are correlated to their private information. This concept - in particular *partial* cursedness - has been criticized in the literature on the basis of the fact that it seems difficult to justify it in terms of learning.
In this paper we show that *partial* cursedness may be seen as a reduced form for the interaction of players who are either *fully* "cursed" or *fully* "non-cursed". Given that full cursedness has a quite natural learning interpretation, our result indirectly provides a learning foundation also for partial cursedness.

*Most of all, I would like to thank my advisor Pierpaolo Battigalli. Moreover, I also thank Matthias Messner for useful suggestions. It goes without saying that all mistakes are mine.

†Bocconi University, Department of Economics. Contact: astrid.gamba@unibocconi.it

1 Introduction

In this paper we provide a theoretical reinterpretation and a learning foundation for the concept of χ-*cursed equilibrium* introduced in Eyster and Rabin (2005). Several authors have questioned the existence of a plausible learning foundation for this equilibrium concept (Eyster and Rabin (2005), Fudenberg (2006), Crawford and Irriberri (2006), Jehiel and Koessler (2008)).
The χ-cursed equilibrium is based on the assumption that players in a strategic setting under-evaluate the informational content of other players' play. In particular, they underestimate the degree to which opponents' actions are correlated with their private information. In a χ-cursed equilibrium each payoff type of each player best responds to a convex combination of two conjectures: the *correct conjecture* that each opponent plays his type-dependent strategy and the "*cursed*" *conjecture* according to which *each* type of each of the other players plays the *same* mixed action, which corresponds to his average distribution of actions. Note that each player correctly predicts the equilibrium distribution of opponents' actions, but he does not identify the connection between types and actions. The parameter χ represents the extent to which each player believes that any other player is playing his average strategy rather than his type contingent strategy. For $\chi = 1$, players are said to be *fully cursed*, and the resulting equilibrium is called *fully cursed equilibrium* (FCE). For $\chi = 0$, the equilibrium coincides with the Bayesian Nash equilibrium (BNE). For intermediate values of χ players are playing a *partially cursed equilibrium* (PCE). The authors show that PCE outperforms BNE in organizing experimental evidence.
Despite the fact that PCE is powerful in rationalizing experimental data it is difficult to justify it both on an intuitive and on a theoretical level. In particular, it is not clear how players could at the same time be aware of the type dependence of their opponents' strategies and still form partially cursed beliefs about these strategies. It has already been pointed out in the literature that the difficulty to combine these two contradictory aspects also constitutes a serious obstacle to finding a satisfactory learning foundation for PCE. Fudenberg (2006), Jehiel and Koessler (2008) and the authors themselves questioned the possibility of finding a reasonable learning foundation for the concept of PCE. In particular, Fudenberg (2006) argues that the amount of cursedness should decline as players become more experienced.

Thus, it is not very plausible to see an intermediate degree of cursedness as an equilibrium phenomenon. Eyster and Rabin (2005) in the discussion of their concept point out that, while the fully cursed equilibrium might be justifiable in terms of learning, "*whether one could find a learning story combined with assumptions about a priori partial strategic sophistication, that would provide foundations for (our) exact specification of partially cursed equilibrium, seems more doubtful*".[1]

To the best of our knowledge there exists only one other paper, Miettinen (2007), which tries to provide a learning foundation for PCE. He shows that any PCE may be seen as an *analogy-based expectation equilibrium.* However, this way of rationalizing partial sophistication depends on players actually believing that their opponents' play depends in a more complex way on their types than it actually does.[2] Hence, in that paper partial cursedness seems to correspond to a form of 'hyper-sophistication' rather than a form of bounded rationality and thus goes against the original spirit of PCE.

As we discuss in this paper, the cursed equilibrium concept is a refinement of *self-confirming equilibrium* (SCE), something that has gone unnoticed in previous studies. SCE characterizes stationary states of plausible learning processes. In a SCE players are best responding to their conjectures (*rationality condition*) and the information revealed ex post, after the equilibrium play, would not induce them to change those conjectures, independently of whether they are correct or not (*beliefs'confirmation*). In a PCE and in a FCE players typically hold *incorrect joint beliefs* on opponents' types and actions. So, both PCE and FCE can be seen as SCE under specific assumptions on the ex post information structure which naturally deliver a situation where players best respond to wrong conjectures on the opponents' true type dependent strategies. In the light of this, the question becomes: how can we justify such a refinement of SCE? There is a reasonable answer for FCE: players have naive conjectures which do not take into account that the opponents' behavior depends on the opponents' private information. On the contrary, PCE does not have such a plausible justification. No primitive theoretical assumption is provided in support of the specific form of the beliefs structure of partially cursed players (i.e. the convex combination of the naive

[1] See Eyster and Rabin (2005), pp 1633-1634

[2] Essentially, the equivalence between PCE and analogy based expectation equilibrium relies on an enlargement of the type space in a payoff irrelevant way. Partial sophistication means that players incorrectly believe that each of their opponent's equilibrium play varies over the subsets of types over that player's payoff is constant

and the correct conjecture).
In this paper we show that this pessimism with respect to the possibility of justifying PCE is not warranted. In particular, we will provide a natural interpretation of partial sophistication in terms of full cursedness. More precisely, we show that partial cursedness may be seen as a 'reduced form' for a game in which players are either *cursed* or *non-cursed.* Put differently, we will define a game, in which each player is characterized by a two-dimensional type: his payoff type and his (payoff irrelevant) cursedness type. We will refer to this set up as *binary cursedness* (BC). In line with Eyster and Rabin (2005) the cursedness type determines only how players form their beliefs about their opponents' strategies. Unlike in PCE though we allow the cursedness parameter to take only the two extreme values. We show that for any PCE with parameter χ there exists a distribution over the type space of our game such that for each payoff type of each player the average equilibrium behavior of its cursed and its non-cursed version corresponds to this payoff type's partially cursed equilibrium behavior.
An immediate implication of this result is that providing a justification in terms of learning for partial sophistication reduces to finding a learning foundation for full cursedness. As has already been conjectured by several authors, the latter problem has a very natural solution. Here, we propose two different specifications of a learning foundation to PCE. In particular, we show that in a learning context fully cursed behavior may simply reflect the fact that a player has limited information about the outcomes of the game which he is playing, so that he is not able to link his opponents' actions to their payoff types. So, a PCE can be seen as a SCE of a game where *better informed* and *less informed* players are strategically interacting. Alternatively, fully cursed behavior may be delivered by the lack of ability to process the available information. Cursed players hold complete beliefs over actions and types which are structured as the product of the marginal over actions and the marginal over types. According to this second interpretation, a PCE can be seen as a SCE of a game where all players have the same (ex post) information structure but some of them are *boundedly rational.*

The paper is organized as follows. In Section 2 we describe first the framework and equilibrium concept of Eyster and Rabin (2005) and then we discuss the relation between cursed equilibrium and self-confirming equilibrium. In Section 3 we introduce our binary cursedness set up and define the corresponding equilibrium, which we call *binary cursed equilibrium* (BCE). Afterwards, we discuss the learning foundation for our environment. In Section 4 we prove two equivalence results between PCE and BCE. First, in Section 4.1 we establish a particularly strong equivalence result for a simple environment (two-players with one-sided asymmetric information): for each χ and each PCE for this parameter, there exists a fixed fraction of fully cursed types in the BC framework (i.e. a share which is constant across all payoff types of all players), such that the BC framework has an equilibrium in which the average behavior of each payoff type of each player coincides with the behavior of this player in the χ-partially cursed equilibrium we are considering (and vice versa). Section 4.2 considers a more general environment with two sided asymmetric information. We establish that in such environments equivalence holds only if both in the Eyster and Rabin framework as well as in our set up the cursedness parameters are allowed to vary across players and payoff types. Finally, Section 5 concludes.

2 Cursed Equilibrium

Throughout the paper we will consider an environment with only two players. Denote the set of their conceivable payoff types, independently distributed, $\Theta_i = \left\{\theta_i^A, \theta_i^B\right\}$, $i = 1, 2$ and Θ the set of payoff types' profiles. Let $p^i \in \Delta\Theta_j$ be the belief of player i on player j's payoff type, $i \neq j$. In this set up $p^i(\theta_j)$ should be interpreted as an objective measure. Imagine that for each player j there is a large population of agents who can play in his role, and each agent is randomly chosen together with a realization θ_j of the private information of player j. With independent payoff types and under the assumption that the statistical distribution of θ_j is known by agents in population i, $i \neq j$, $p^i(\theta_j)$ reflects the probability that the agent drawn from population j to play against i is of type θ_j.
For each player the set of possible actions is payoff type independent. Let us assume that both players have only two possible (pure) actions: $A_i = \{a_i', a_i''\}$.
A strategy for player i is a mapping from the set of his payoff types to the set of his possible actions. Denote σ_i the true strategy of player i, $\sigma_i : \Theta_i \rightarrow \triangle A_i$. Denote $\Sigma_i \equiv (\triangle A_i)^{\Theta_i}$and $\Sigma \equiv \Sigma_1 \times \Sigma_2$. Let $\sigma_i(a_i|\theta_i)$ be the probability that type θ_i of player i plays action a_i according to the type contingent strategy $\sigma_i(\theta_i)$. Denote $\overline{\sigma}_i$ the average strategy of player i, averaged over player i's payoff types, that is $\overline{\sigma}_i(a_i) := \sum_{\theta_i \in \Theta_i} p^i(\theta_i)\sigma_i(a_i|\theta_i)$, $i = 1, 2$, $a_i \in A_i$.

$\overline{\sigma}_i$ represents the cursed conjecture of player i's opponent according to which each type θ_i of player i randomizes according to the marginal probabilities on actions derived from the true strategy σ_i. That is, $\overline{\sigma}_i(a_i)$ is the frequency player i's opponent expects action a_i to be played with by player i when he is following the strategy σ_i. According to PCE each player believes that with probability χ *any type* of the opponent behaves according with the average frequency, $\overline{\sigma}_i$, and with probability $(1-\chi)$ he behaves according to his type-specific strategy, σ_i. Denote $\widehat{\sigma}_i$ the convex combination of these two conjectures parametrized by χ, which describes the extent to which a player is cursed, that is $\widehat{\sigma}_i = \chi\overline{\sigma}_i + (1-\chi)\sigma_i$. A χ-cursed player j incorrectly believes that type θ_i of player i plays action a_i with probability $\widehat{\sigma}_i(a_i|\theta_i) = \chi\overline{\sigma}_i(a_i) + (1-\chi)\sigma_i(a_i|\theta_i)$. As Eyster and Rabin themselves point out, it is possible to generalize such a set up allowing for χ to vary across players and payoff types. Denote χ_{θ_i} the degree of cursedness of type θ_i of player i. For the time being though, we consider the special case with a unique degree of cursedness, χ, for all players and all payoff types.
Player i's payoff function $u_i : A \times \Theta \to \mathbb{R}$ depends on players' action profile $a \in A \equiv A_1 \times A_2$ and their types.

Definition 1 *A behavioral strategy profile σ is a* partially cursed equilibrium *(or χ-*cursed equilibrium*) if for each i, $\theta_i \in \Theta_i$, each action a_i^* such that $\sigma_i(a_i^*|\theta_i) > 0$,*

$$a_i^* \in \arg\max_{a_i \in A_i} \sum_{\theta_j \in \Theta_j} \sum_{a_j \in A_j} p^i(\theta_j) \left[\chi\overline{\sigma}_j(a_j) + (1-\chi)\sigma_j(a_j|\theta_j)\right] u_i(\theta_i, \theta_j; a_i, a_j)$$

For $\chi = 1$, each type of each player best responds to the cursed conjecture and the equilibrium is called *fully cursed equilibrium (FCE)*. For $\chi = 0$ each player best respond to the correct conjecture as in a Bayesian Nash equilibrium. For intermediate values of χ a *partially cursed equilibrium (PCE)* is played.

2.1 Cursed Equilibrium and Self-Confirming Equilibrium

What Eyster and Rabin (2005) and the related literature cited above did not notice is that PCE is a refinement of the self-confirming equilibrium

concept[3] (SCE). If we want to provide a learning foundation to PCE we need to elaborate further on the relation between PCE and SCE. The self-confirming equilibrium concept is a static solution concept which rationalizes the stationary states of a plausible learning processes. Essentially, the SCE represents situations where players choose best replies to their conjectures on the opponents' play (*rationality condition*) and the information on the equilibrium play revealed *ex post* , after that the choices have been made, does not induce them to change those conjectures, independently of whether they are correct or not (*conjectures' confirmation property*). The key idea of a SCE is that individuals might have incorrect conjectures about others' behavior as long as these conjectures are not contradicted by the evidence. To verify if a situation is a SCE we need to make explicit assumptions on what players can observe ex post. We argue that cursed equilibrium can be seen as a particular SCE under two specific assumptions about the players' information structure:
(i) players know *ex ante* the objective probabilities of the states of nature θ;
(ii) players observe *ex post* the actions played by the opponents but *not* their types.
Assumption (i) justifies the fact that in equilibrium players hold correct marginal beliefs on opponents' payoff types. Assumption (ii) implies that if the learning process converges to a stationary state players learn the correct marginal probabilities of opponents'actions by observing the long-run frequencies. Given this information structure, players do not have the possibility to learn the connection between types and actions and, consequently, they typically hold *wrong joint conjectures* on opponents' types and actions in the stationary state.
It is immediate to see that fully cursed equilibrium is a SCE, while it is less clear the link between partially cursed equilibrium and SCE. In a PCE players are assumed to hold beliefs expressed as the convex combination ($\widehat{\sigma}_i$) of the naive conjecture ($\overline{\sigma}_i$) and the correct conjecture (σ_i) on each opponent i's behavior. At first sight, it may seem unplausible that players could to some extent be aware of the type dependence of their opponents' strategies and still form partially cursed beliefs about these strategies. The key point is that we do *not* have to assume that a player has in mind the two different

[3]See Fudenberg and Levine (1990). See also Battigalli (1987) and Battigalli and Guaitoli (1988) for the related conjectural equilibrium concept and the survey of Battigalli *et al.* (1992) for a discussion on the relevance of these two concepts for the analysis of adaptive processes in repeated interactions contexts.

objects, the naive and the correct conjectures, and combines them according to his degree of cursedness χ. Indeed, the convex combination is just a way to represent the player's incorrect beliefs about the opponents' equilibrium play. It is a device to represent his partial strategic sophistication. How we can justify this particular characterization of players' beliefs is a different issue which needs to be analyzed further. What we want to point out here is just that the fact that players essentially hold correct marginal beliefs and incorrect joint beliefs can be explained in the light of the specific assumptions on the information structure stated above. Cursed equilibrium can be seen as a refinement of SCE but still this specific refinement needs to be justified. This is not an issue with fully cursed equilibrium since it is based on a clear primitive theoretical assumption and it has a plausible learning story behind. Players hold naive conjectures on opponents' play which do not take into account that other players' behavior is linked with their private information. They update such conjectures from time to time, they learn the correct marginal distributions of opponents'actions but they end up in equilibrium having wrong conjetures on the true type contingent strategies of their opponents. Differently, partially cursed equilibrium does not have such an intuitive and plausible justification. To provide a theoretical justification and a learning foundation to PCE is the purpose of the next sections. We will provide a natural interpretation of partial sophistication in terms of full cursedness. More precisely, we show that any game with partially cursed players may be seen as a 'reduced form' for a game in which players are either *cursed* or *non cursed.* An immediate implication of this result is that providing a justification in terms of learning for partial sophistication reduces to finding a learning foundation for full cursedness.

3 Binary Cursedness

In this section we construct a different set up where players are drawn from heterogenous populations which differ not only in the private information but also in the way players form their conjectures on their opponent's play. Players have either *naive* conjectures on the equilibrium play of the opponent, which do not take into account that the opponent's strategy is (payoff) type-specific, or they fully identify the type contingency of the opponent's strategy. We do not allow for intermediate levels of cursedness: either a player fully believes that each payoff type of the opponent is playing *that payoff type* contingent strategy or he fully believes that *every payoff type* of the opponent is following through that player's average behavior. Essentially, we enlarge the type space in order to admit for each player a *cursedness type*, besides his payoff type. This dimension of the type does not influence the payoffs, but just the nature of the beliefs each player holds about the opponent's behavior.

Denote B_j the set of player j's beliefs parameters, $B_1 \equiv B_2 \equiv \{C, NC\}$, where C means cursed and NC non-cursed. Hence, each player essentially has a two-dimensional type $t_j \in T_j = \Theta_j \times B_j$. Each payoff type θ_j of player j is cursed with a certain probability β which represents exactly the frequency of cursed individuals in the whole population. In this new framework, a strategy of player j is a mapping from the set of types T_j into the set of his possible actions. As above, the set of every player's conceivable actions does not change across types. Hence, allowing for randomized strategies, we define $\rho_j : T_j \rightarrow \triangle A_j$, $j = 1, 2$. We denote the probability that player j of type t_j plays action a_j, under strategy $\rho_j(t_j)$, as $\rho_j(a_j|t_j)$. Denote by $\overline{\rho}_j$ the average strategy of player j, averaged across the two dimensions of type t_j of player j. The frequency action a_j is played with by player j is

$$\overline{\rho}_j(a_j) := \sum_{\theta_j \in \Theta_j} p^i(\theta_j) \left[\beta \rho_j(a_j|\theta_j, C) + (1-\beta)\rho_j(a_j|\theta_j, NC)\right].$$

The expected utility of a cursed player i when his payoff type is θ_i and he plays action a_i is equal to:

$$\sum_{\theta_j \in \Theta_j} \sum_{a_j \in A_j} p^i(\theta_j) \overline{\rho}_j(a_j) u_i(\theta_i, \theta_j; a_i, a_j).$$

The expected utility of a non-cursed player i when his payoff type is θ_i and he plays action a_i is equal to:

$$\sum_{\theta_j \in \Theta_j} \sum_{a_j \in A_j} p^i(\theta_j) \left[\beta \rho_j(a_j|\theta_j, C) + (1-\beta)\rho_j(a_j|\theta_j, NC)\right] u_i(\theta_i, \theta_j; a_i, a_j).$$

Definition 2 *A behavioral strategy profile* $\rho \in \prod_{i=1,2} (\Delta A_i)^{T_i}$ *is a* binary cursed equilibrium *(or* β-cursed equilibrium*) if for each* i*,*

i) for each $(\theta_i, NC) \in T_i$ *and each action* a_i^* *such that* $\rho_i(a_i^*|\theta_i, NC) > 0$*,*

$$a_i^* \in \arg\max_{a_i \in A_i} \left\{ \sum_{\theta_j \in \Theta_j} \sum_{a_j \in A_j} p^i(\theta_j) \left[\beta \rho_j(a_j|\theta_j, C) + (1-\beta)\rho_j(a_j|\theta_j, NC) \right] \times \right.$$

$$\left. u_i(\theta_i, \theta_j; a_i, a_j) \right\}$$

and

ii) for each $(\theta_i, C) \in T_i$ *and each action* a_i^* *such that* $\rho_i(a_i^*|\theta_i, C) > 0$*,*

$$a_i^* \in \arg\max_{a_i \in A_i} \sum_{\theta_j \in \Theta_j} \sum_{a_j \in A_j} p^i(\theta_j) \overline{\rho}_j(a_j) u_i(\theta_i, \theta_j; a_i, a_j).$$

Given a β-cursed equilibrium profile, we can express the *average equilibrium behavior* of each type θ_i of each player i, averaged across his cursedness types, as:

$$\overline{\rho}_{\theta_i} := \beta \left[\rho_i(\theta_i, C) \right] \oplus (1-\beta) \left[\rho_i(\theta_i, NC) \right]$$

and the average equilibrium profile as $\overline{\rho} \equiv (\overline{\rho}_1, \overline{\rho}_2)$, where $\overline{\rho}_i \equiv (\overline{\rho}_i(\theta_i))_{\theta_i \in \Theta_i}$, $i = 1, 2$.

3.1 Learning Interpretation for Binary Cursed Equilibrium

In this section we will argue that BCE admits a very natural interpretation in terms of learning. Since in the following sections we are able to show that Eyster and Rabin's PCE concept can be seen as a reduced form of BCE this will allow us to conclude that also PCE can be seen as a steady state of a

learning process.
Our learning framework has two possible specifications: either players differ in their ex post information structure or they differ in their a priori strategic sophistication. We explain how these differences naturally deliver situations where (in the long run) part of the players will be able to infer all details about their opponents strategies while other players will miss some relevant aspects. Notice that such an approach could not be applied directly to Eyster and Rabin (2005) PCE concept. In their framework it is assumed that *all* types of *all* players hold beliefs which are an average of correct and incorrect beliefs about the type contingent strategies of their opponents.
We focus on an adapted version of *fictitious play*. Imagine that players are *anonymously* interacting among each other. They are randomly drawn each period from a population composed of heterogeneous agents. So each individual called to play in the role of player $i = 1, 2$ is drawn from sub-population i together with a certain characteristic $\theta_i \in \Theta_i$, which describes his payoff type. Payoff types are independently distributed and their statistical distributions are commonly known[4]. Each player i knows his own payoff type, the set of payoff types of the opponent Θ_j, his own payoff function $u_i : \Theta \times A \to \mathbb{R}$, and that, after having played, he will receive further information on the behavior of the opponent. We assume that players inherit entirely the past observations of agents who played in their same role in past rounds, so that they can rely on the full lenght of history to form their beliefs.
We describe now the two different specifications of this learning framework.

i) *Different ex post information structures*
We assume that each sub-population is split up into two groups of agents. These groups differ in their *ex post information structures*, that is, in what agents can observe after each round of play. There are *less informed* agents who observe only the *opponent's action* and *better informed* agents who observe the *opponent's action and payoff type*. Agents from the first group, who are supposed to represent cursed players, cannot infer anything about the payoff type of their opponents and, therefore, they cannot identify the type dependence of the strategy they are following. To formalize this feature, it is natural to assume that they start with complete beliefs (on actions and types of the opponents) that are the product of the marginals over ac-

[4] Recall that we have given to beliefs an objective meaning, reflecting the true statistical distribution of the opponent's payoff type in the population.

tions and types. Since they cannot observe their opponent's payoff type they have no information which would induce them to change the product structure of their beliefs. Indeed, they start with arbitrary weights on the opponent's actions (recall that they know the probability of each payoff type of the opponent). On the next round they update these weights and revise (marginal) conjectures on opponents' actions on the basis of the new observation. Each round, the agent who is called to play inherits from the previous rounds the empirical frequencies of each action of the opponent collected in the past. Hence, after many rounds, he learns the frequencies with which the opponent is playing each of his conceivable actions. They do not know *who plays what*, but they know what the opponent played on average in the past and they best respond to this average behavior. So, they will typically have wrong conjectures about the type contingent strategy of the opponent. On the other hand, in the long-run, their marginal conjectures about actions must be correct.

Better informed agents, who represent non-cursed players, can observe the opponent's action and type, so that they can unequivocally identify action and payoff type of the opponent. They start with a distribution over the pairs of type and action of the other players which has to be consistent with the known distribution over types. After one round the information "type θ_j played action a_j" is inherited, and registered by the next agent playing in the role of player i. It is intuitive that in any steady state of this learning process players must know the type contingent strategy of the opponent.

ii) *Same ex post information structures but different structures of beliefs (boundedly rational agents)*

Assume that all players have access to information about actions and types of the individuals who played in the role of their opponent in the past. The candidates for non-cursed players are able to infer the correlation between actions and types of the opponents while the candidates for cursed players cannot. We can formalize this by assuming that sophisticated players start with complete beliefs on the action and type of the opponent, so that through Bayesian updating they learn the true correlation between actions and types of the opponents. On the other hand, cursed players hold complete beliefs over actions and types which are structured as the product of the marginal over actions and the marginal over types. That is, they do not have the information processing ability to identify the type contingency of the opponent's strategies even if they have sufficient data.[5] Given that they cannot link observed actions to the payoff type of their opponents, it is reasonable to assume that they only keep track of the frequencies of played actions without trying to connect them to their opponents types. The marginals over the actions of the other players end up being correct, while the joint beliefs are typically incorrect.

Note that any steady state of our learning dynamics in each of the two environments can be theoretically interpreted as a self-confirming equilibrium.[6] Indeed, this concept describes outcomes where players are best responding to their conjectures (*rationality condition*) and the information revealed ex post, after the equilibrium play, would not induce them to change those conjectures, independently of whether they are correct or not (*beliefs'confirmation*). In equilibrium, non-cursed players have correct beliefs on the type contingent strategies of the opponents. Cursed players have wrong conjectures on the joint distribution over actions and types but since they cannot revise these conjectures-because they do not have either the possibility (case (i))-or the ability to do so (case (ii))-additional information would not change their wrong beliefs on the type contingent strategies of the other players. Given these correct and, respectively, wrong conjectures on the opponents' behavior, each player is best responding to the other players' equilibrium strategies.

[5] Note that the difference between this framework and the previous one is that, in the environment described above, individuals would have the *ability to infer* the correlation, but since they do not observe the realized payoffs they have no *possibility to revise* their complete beliefs in a type contingents way. On the contrary, in the current environment, individuals are in a sense *boundedly rational*: they are simply not able to infer the correlation even if they have the possibility to do it. They update only the marginals over the actions, not being able to learn any correlation with the types of the opponent.

[6] We are referring to the version proposed by Battigalli (1987), Battigalli and Guaitoli (1997) and discussed in Battigalli *et al.*(1992)

4 On the relation between Partially Cursed Equilibrium and Binary Cursed Equilibrium

4.1 One-sided incomplete information

For the sake of simplicity, let us consider first the situation where player 1 has two payoff types, while player 2 only one. Player 1 is informed about the true state of nature, while player 2 is not and holds beliefs $p^2 \in \Delta\Theta_1$ which coincide with the commonly known statistical distribution of θ_1 in population 1. Call (σ_1, σ_2) the generic χ-cursed equilibrium for a given χ of such a bayesian game with incomplete information on the side of player 2.
Consider the BC extension of such a game. Obviously, since player 2's payoff type set Θ_2 is a singleton, so that player 1 knows the true state of nature, we can assume that player 1 is non-cursed with probability one, that is, his cursedness type set is a singleton, $T_1 = \left\{\theta_1^A, \theta_1^B\right\} \times \{NC\}$. Differently, the type set of player 2 coincides with his cursedness type set $T_2 = \left\{\overline{\theta}_2\right\} \times \{C, NC\} \equiv B_2$. Let β be the probability that player 2 is cursed and $(1-\beta)$ the probability that he is non-cursed. A strategy for player 1 is a mapping from his *payoff* type set into the set of probability measures on the set of his possible actions $\rho_1 : \Theta_1 \to \triangle A_1$, while the mapping for player 2 goes from his beliefs parameter set into the set of probability measures on the set of his possible actions, $\rho_2 : B_2 \to \triangle A_2$. Let us call $r_2 : B_2 \to A_2$ a pure strategy for player 2.

Obviously, in the situation where $\chi = 0$ ($\chi = 1$) the equivalence between the equilibrium behavior of the χ-cursed player 2 and the average equilibrium behavior of player 2 in the BC-game is immediately shown by taking $\beta = 0$ ($\beta = 1$).

We establish now the equivalence between the average equilibrium behavior in the BC-game, where the population of agents in the role of player 2 is partly cursed and partly non-cursed and the equilibrium behavior which would arise if *each individual* in the role of player 2 were χ-cursed.

Proposition 1 *Let* $(\rho_1, \rho_2) \in (\Delta A_1)^{\Theta_1} \times (\Delta A_2)^{B_2}$ *be a* β*-cursed equilibrium strategy profile of a BC-game, extension of a Bayesian game with incomplete information on the side of player 2. There exists* $\chi \in [0,1]$ *such that* $(\rho_1, \overline{\rho}_2)$ *is a* χ*-cursed equilibrium of the underlying Bayesian game.*

Proof

Suppose that player 2 is playing in equilibrium a pure strategy r_2. Given that (ρ_1, r_2) is an equilibrium strategy profile of the BC-game for a given β, if $r_2(C)$ and $r_2(NC)$ are the type contingent (pure) actions played in equilibrium by player 2, the following two conditions must hold:

$$\sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\overline{\rho}_1(a_1) u_2(\theta_1; a_1, r_2(C)) \geq$$

$$\sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\overline{\rho}_1(a_1) u_2(\theta_1; a_1, r_2(NC))$$

and

$$\sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1) u_2(\theta_1; a_1, r_2(NC)) \geq$$

$$\sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1) u_2(\theta_1; a_1, r_2(C))$$

The first inequality coincides with the equilibrium condition for player 2 in a χ-Cursed Equilibrium with $\chi = 1$, while the second inequality coincides with the equilibrium condition for player 2 in a χ-cursed equilibrium with $\chi = 0$. In other words, $(\rho_1, r_2(C))$ is a χ-cursed equilibrium strategy profile when $\chi = 1$, while $(\rho_1, r_2(NC))$ is a χ-cursed equilibrium strategy profile when $\chi = 0$.

Given the linearity in χ of the expected utility of player 2, there must exist a degree of cursedness $\widehat{\chi} \in [0,1]$ such that the $\widehat{\chi}$-cursed player 2 is indifferent between $r_2(C)$ and $r_2(NC)$, i.e.:

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\left[\widehat{\chi}\overline{\rho}_1(a_1) + (1-\widehat{\chi})\rho_1(a_1|\theta_1)\right] u_2(\theta_1; a_1, r_2(C)) =$$

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\left[\widehat{\chi}\overline{\rho}_1(a_1) + (1-\widehat{\chi})\rho_1(a_1|\theta_1)\right] u_2(\theta_1; a_1, r_2(NC))$$

Hence, any mixing between actions $r_2(C)$ and $r_2(NC)$ is a best response of a $\widehat{\chi}$-cursed player 2 to ρ_1. Thus,.we can take exactly $\overline{\rho}_2 \equiv \beta\cdot[r_2(C)] \oplus (1-\beta)\cdot[r_2(NC)]$, the average equilibrium behavior of player 2 in the BC-game, averaged across his cursedness types, as the best response of a $\overline{\chi}$-cursed player 2.to ρ_1. $(\rho_1, \overline{\rho}_2)$ is a χ-cursed equilibrium for $\chi = \overline{\chi}$.

Suppose now that player 2 is playing a randomized strategy ρ_2. Let (ρ_1, ρ_2) be the equilibrium strategy profile of a BC-game. Suppose that both (cursedness) types of player 2 are playing mixed actions. Call a_2 and a_2' the two feasible actions of player 2. It must be the case that both types of player 2 are indifferent between a_2 and a_2', so that the following holds:

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\overline{\rho}_1(a_1)u_2(\theta_1; a_1, a_2) = \sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\overline{\rho}_1(a_1)u_2(\theta_1; a_1, a_2')$$

and

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1)u_2(\theta_1; a_1, a_2) = \sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1)u_2(\theta_1; a_1, a_2')$$

Given ρ_1, the two conditions above coincide with the equilibrium condition for a χ-cursed player 2, respectively, for $\chi = 1$, the former, and for $\chi = 0$, the latter. Hence, by linearity of the expected payoffs in χ, it must hold that, for any $\chi \in [0, 1]$,

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\left[\chi\overline{\rho}_1(a_1) + (1-\chi)\rho_1(a_1|\theta_1)\right] u_2(\theta_1; a_1, a_2) =$$

$$\sum_{\theta_1\in\Theta_1}\sum_{a_1\in A_1} p^2(\theta_1)\left[\chi\overline{\rho}_1(a_1) + (1-\chi)\rho_1(a_1|\theta_1)\right] u_2(\theta_1; a_1, a_2')$$

Thus, we can take $\overline{\rho}_2 = \beta \cdot [\rho_2(C)] \oplus (1-\beta) \cdot [\rho_2(NC)]$ as the χ-best response of player 2 to ρ_1, so that $(\rho_1, \overline{\rho}_2)$ is a χ-cursed equilibrium of the underlying bayesian game.■

Proposition 2 *Let $(\sigma_1, \sigma_2) \in (\triangle A_1)^{\Theta_1} \times \Delta A_2$ be a χ-cursed equilibrium strategy profile of a Bayesian game with incomplete information on the side of player 2. There exists $\beta \in [0,1]$ and $\rho_2 \in (\Delta A_2)^{B_2}$ such that, for each $a_2 \in A_2$, $\overline{\rho}_2(a_2) = \sigma_2(a_2)$ and (σ_1, ρ_2) is a β-cursed equilibrium of the BC extension of the original Bayesian game.*

Proof

Given that (σ_1, σ_2) is an equilibrium strategy profile, for any a_2 in the support of σ_2 and $a_2' \in A_2$ the following inequality holds:

$$\sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1) \left[\chi \overline{\sigma}_1(a_1) + (1-\chi)\sigma_1(a_1|\theta_1)\right] \times$$

$$[u_2(\theta_1; a_1, a_2) - u_2(\theta_1; a_1, a_2')] \geq 0$$

Let us call d this difference which is a function of χ.

Fix $\overline{\chi} \in (0,1)$. Let us suppose that $d(\overline{\chi}) > 0$. Then it must be the case that $\sigma_2(a_2) = 1$. Given the linearity of the expected utility in χ either $d(\chi)$ is monotonically increasing or decreasing in χ. If $d(\chi)$ is increasing (decreasing) in χ, then $d(1) > 0$ ($d(0) > 0$). If $d(1) > 0$, the β-cursed equilibrium strategy of player 2, ρ_2, derived from σ_2, must be such that $\rho_2(a_2|C) = 1$. If we take $\beta = 1$, we get $\overline{\rho}_2(a_2) = 1 = \sigma_2(a_2)$. If $d(\chi)$ is decreasing in χ so that $d(0) > 0$, ρ_2 must be such that $\rho_2(a_2|NC) = 1$. If we take $\beta = 0$, we get $\overline{\rho}_2(a_2) = 1 = \sigma_2(a_2)$.

Suppose now that $d(\overline{\chi}) = 0$, the $\overline{\chi}$-cursed player 2 is indifferent among his two possible actions. As above, depending on the expected utility being increasing or decreasing in χ and, consequently, depending on $\rho_2(\cdot)$, take $\beta \in [0,1]$ such that, for each $\alpha_2 \in A_2$, $\rho_2(a_2) = \beta\rho_2(a_2|C) + (1-\beta)\rho_2(a_2|NC) = \sigma_2(a_2)$. Since we have taken an arbitrary $\overline{\chi}$, then for any $\chi \in (0,1)$ we must be able to find a β such that the β-average equilibrium behavior of player 2 in the BC-game exactly replicates the χ-cursed equilibrium behavior of player 2 in the original game.■

Remark: *Notice that the preceding propositions only tell us that for any given β-cursed equilibrium there exists a parameter χ for which there exists a χ-cursed equilibrium which coincides with the β-average equilibrium behavior of the BC-game we started with and vice versa. It is rather straightforward to show that this statement can be strengthened in the following sense: for any equilibrium strategy profile ρ of a BC-game with parameter β denote the corresponding χ-parameter and χ-cursed equilibrium strategy profile of the underlying bayesian game by $\chi(\beta,\rho)$ and $\sigma(\beta,\rho)$, respectively; similarly let $\beta(\chi,\sigma)$ and $\rho(\chi,\sigma)$ be the β-parameter and equilibrium strategy of the BC-game with parameter β associated with the pair (χ,σ), where σ is some χ-cursed equilibrium strategy profile of a bayesian game. Then the following holds: for any pair (β,ρ) we have $\beta(\chi(\beta,\rho),\sigma(\beta,\rho)) = \beta$ and, analogously, for any pair (χ,σ), $\chi(\beta(\chi,\sigma),\rho(\chi,\sigma)) = \chi$. While interesting in itself, this observation is not really central for our purposes. Thus we omitt its proof.*

4.2 Two-sided incomplete information

The equivalence between players' equilibrium behavior in the two frameworks is less immediate if it is the case that both players have more than one payoff type. Let us assume that also player 2 holds some payoff relevant private information, so that also the degree of cursedness of player 1 matters now. Recall that, under the assumption of one-sided incomplete information, given a β-cursed equilibrium, we only needed to find if there was a χ which solved the unique indifference condition of the unique payoff type of the uninformed player. Differently, with two sided incomplete information, we must find a χ for each payoff type of each player which leaves him indifferent, so that we face a *system* of indifference conditions.
In this section we show that it is not necessarily true that for any β constant across players and/or across types for each player we are able to find a constant χ such that the χ-cursed players' equilibrium behavior replicates the average equilibrium behavior of both players in the BC-game, and, viceversa, that for any χ, unique for all players and types, we are able to replicate their equilibrium behavior with their β-average behavior in the BC-game.

We show first that, for some constant χ, we can find a χ-cursed equilibrium outcome that cannot be reproduced as the average equilibrium outcome of the BC-game.

Counter-example 1

Let us consider a game with two players, each with two payoff types, $\Theta_i = \{\theta_i^A, \theta_i^B\}$, $i = 1, 2$. Both players, and both types of each player, are characterized by the same degree of cursedness, $\chi \in (0, 1)$. The payoff matrices are the following:

θ_1^A, θ_2^A	l	r
U	$1,2$	$0,1$
D	$0,1$	$2,0$

θ_1^A, θ_2^B	l	r
U	$0,1$	$1,2$
D	$2,0$	$0,1$

θ_1^B, θ_2^A	l	r
U	$0,1$	$3,0$
D	$1,2$	$0,1$

θ_1^B, θ_2^B	l	r
U	$3,0$	$0,1$
D	$0,1$	$1,2$

Given that l is dominant for θ_2^A and r is dominant for θ_2^B, in equilibrium, the weight that the χ-cursed belief of player 1 put on action l is $\overline{\sigma}_2(l) = 1/2$, according to the average behavior $\overline{\sigma}_2$ of player 2 that player 1 has in mind. Similarly, $\overline{\sigma}_2(r) = 1/2$. Hence, we can compute the difference among the expected utility that a χ-cursed player 1 with payoff type θ_1^A gets from action U and the expected utility he gets from action D as:

$$\sum_{\theta_2 \in \Theta_2} \sum_{a_2 \in A_2} p^1(\theta_2) \left[\chi \overline{\sigma}_2(a_2) + (1-\chi)\sigma_2(a_2|\theta_2)\right] \times$$

$$\left[u_1(\theta_1^A, \theta_2; U, a_2) - u_2(\theta_1^A, \theta_2; D, a_2)\right] =$$

$$1 - \tfrac{3}{2}\chi = d_A(\chi).$$

Similarly, we can compute the difference among the expected utility that a χ-cursed player 1 with payoff type θ_1^B gets from action U and the expected utility he gets from action D as:

$$\sum_{\theta_2 \in \Theta_2} \sum_{a_2 \in A_2} p^1(\theta_2) \left[\chi \overline{\sigma}_2(a_2) + (1-\chi)\sigma_2(a_2|\theta_2)\right] \times$$

$$\left[u_1(\theta_1^B, \theta_2; U, a_2) - u_2(\theta_1^B, \theta_2; D, a_2)\right] =$$

$$2\chi - 1 = d_B(\chi).$$

While $d_A(\cdot)$ is decreasing in χ, $d_B(\cdot)$ is increasing in χ. Moreover, $d_A(\chi) > 0$ iff $\chi < 2/3$, while $d_B(\chi) > 0$ iff $\chi > 1/2$. So, if we fix any $\overline{\chi} \in (1/2, 2/3)$, both the $\overline{\chi}$-cursed type θ_1^A and the $\overline{\chi}$-cursed type θ_1^B of player 1 would choose U in a $\overline{\chi}$-cursed equilibrium. But, being $d_A(\cdot)$ decreasing in χ and $d_B(\cdot)$ increasing in χ, in order to replicate in the BC-game the $\overline{\chi}$-cursed equilibrium behavior of type θ_1^A, we need to take $\beta_{\theta_i^A} = 0$, while to replicate the $\overline{\chi}$-cursed equilibrium strategy of type θ_1^B, we need to take $\beta_{\theta_i^B} = 1$. Hence, we have found a range of χ for which we cannot reproduce the equilibrium behavior of both payoff types of player 1 when they are partially cursed, taking the same proportion β of fully cursed players in the sub-population of agents with payoff type θ_1^A as in the sub-population of agents with payoff type θ_1^B.$\diamondsuit$

This example suggests that, possibly, a way to establish the equivalence in the two sided incomplete information case is to allow for more degrees of freedom of the parameter β. The next step is, indeed, to check if, letting the parameter β to vary across players and types, we are able to find a unique χ through wich we can replicate the average equilibrium behavior of every type of every player in the BC-game. The next example shows that we are not able to always find such a χ.

Counter-example 2

Let us consider a BC-game where both players have two payoff types as above and, besides, they can be either cursed or non-cursed. Let us recall that the cursedness type is payoff irrelevant.
We focus on player 1 and fix a vector β such that $\beta_{\theta_1^A} = 0$ and $\beta_{\theta_1^B} = 1$.

θ_1^A, θ_2^A	l	r
U	$1, 1$	$0, 0$
D	$0, 2$	$2, 1$

θ_1^A, θ_2^B	l	r
U	$0, 0$	$1, 1$
D	$2, 1$	$0, 2$

θ_1^B, θ_2^A	l	r
U	$0, 2$	$2, 1$
D	$1, 1$	$0, 0$

θ_1^B, θ_2^B	l	r
U	$2, 1$	$0, 2$
D	$0, 0$	$1, 1$

Player 2, when he has payoff type θ_2^A, plays l with probability one, while, when he has payoff type θ_2^B, plays r, independently on his cursedness type. So, given this type contingent strategy for player 2, the non-cursed type θ_1^A of player 1 would play U in equilibrium. Differently, the cursed type θ_1^A, who believes that player 2 is playing the average $\overline{\rho}_2 = (1/2, 1/2)$,would play D. On the contrary, the non-cursed type θ_1^B would play D, while the cursed type θ_1^B would play U, i.e.

$$r_1(\theta_1^A, NC) = r_1(\theta_1^B, C) = U \quad \text{and} \quad r_1(\theta_1^A, C) = r_1(\theta_1^B, NC) = D.$$

Hence, since we have fixed $\beta_{\theta_1^A} = 0$ and $\beta_{\theta_1^B} = 1$, the average equilibrium behavior of type θ_1^A is the same as the average equilibrium behavior of type θ_1^B and it consists in playing U with probability one.

Let us check if we can replicate the average equilibrium behavior of player 1 with a unique, i.e. type independent, χ.
Player 1 prefers U to D for different values of χ, depending on whether he has payoff type θ_1^A or θ_1^B. When his payoff type is θ_1^A, player 1 prefers U to D if and only if $\chi \leq 2/3 \equiv \underline{\chi}$, while when his payoff type is θ_1^B, he prefers U to D if and only if $\chi \geq 6/7 \equiv \overline{\chi}$. Therefore, we cannot replicate the average equilibrium behavior of the two payoff types of player 1 with a unique degree of cursedness for both.$\Diamond$

This second counterexample suggests that we should allow for *more degrees of freedom* also on the χ parameter's side. Indeed, it is plausible to assume that different players with different payoff types may be characterized by different degrees of cursedness. They might have, for instance, different abilities in interpreting opponents'actions and connecting them to opponents' private information. This different abilities may come from the amount of experience collected in the past (before this particular repeated interaction started) or from the particular position they have in the interaction. It is not natural to assume that agents playing in different roles and different payoff types makes mistakes to the same extent. Thus, on the side of a BC-game, we allow for a payoff-type dependent probability of being cursed and we call $\boldsymbol{\beta}$ the vector of such probabilities, $\boldsymbol{\beta} = ((\beta_{\theta_i})_{\theta_i \in \Theta_i})_{i=1,2}$. Similarly, we allow for a payoff-type dependent degree of cursedness and we call $\boldsymbol{\chi}$ the profile of such parameters, $\boldsymbol{\chi} = ((\chi_{\theta_i})_{\theta_i \in \Theta_i})_{i=1,2}$. The next two propositions generalize, respectively, proposition 1 and proposition 2 to the environment of two-sided incomplete information. Since the proofs to propositions 3 and 4 are very similar to those of propositions 1 and 2, we do not report them here. They are available under request though.

Proposition 3 *Let $(\rho_1, \rho_2) \in \prod_{i=1,2} (\Delta A_i)^{T_i}$ be a $\boldsymbol{\beta}$-cursed equilibrium of a BC-game, extension of a game with two-sided incomplete information.There exists $\boldsymbol{\chi}$ such that the strategy profile $(\sigma_1, \sigma_2) \in \prod_{i=1,2} (\Delta A_i)^{\Theta_i}$, with $\sigma_i(\theta_i) = \overline{\rho}_{\theta_i}$ for each $\theta_i \in \Theta_i$, of each player i, is a $\boldsymbol{\chi}$-cursed equilibrium of the underlying Bayesian game.*

Proposition 4 *Let $(\sigma_1, \sigma_2) \in \prod_{i=1,2} (\Delta A_i)^{\Theta_i}$ be a $\boldsymbol{\chi}$-cursed equilibrium of a bayesian game with two-sided incomplete information.There exists β such that $(\rho_1, \rho_2) \in \prod_{i=1,2} (\Delta A_i)^{T_i}$, which satisfies that $\overline{\rho}_{\theta_i} \equiv \sigma_i(\theta_i)$ for each θ_i of each player i, is a $\boldsymbol{\beta}$-cursed equilibrium of the BC extension of the original Bayesian game.*

5 Conclusions

In this paper we have discussed a possible interpretation of Partially Cursed Equilibrium both from a theoretical point of view and in terms of learning. We have shown that individual partial cursedness may be reinterpreted in terms of binary cursedness. The main advantage of the concept of binary cursedness is that on the individual level all players are either cursed or non-cursed and both these types of behavior admit a very natural learning foundation. In particular, we have discussed to possible specification of the learning framework. According to the first interpretation, non-cursed behavior can be seen as the steady state behavior of a player who has detailed information about past play. A cursed player instead may be interpreted as a player whose information does not allow him to infer the relation between his opponent's actions and payoff types. According to the second specification, the difference between cursed and non-cursed players concerns the ability to process the information they hold on past play.

6 Appendix

Proof of proposition 1: partially mixed strategies

Let us suppose now that only one type of player 2 is playing a mixed action. Then we have either

$$\text{(i)} \sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\overline{\rho}_1(a_1)u_2(\theta_1; a_1, a_2) = \sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\overline{\rho}_1(a_1)u_2(\theta_1; a_1, a_2')$$

or

$$\text{(ii)} \sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1)u_2(\theta_1; a_1, a_2) = \sum_{\theta_1 \in \Theta_1} \sum_{a_1 \in A_1} p^2(\theta_1)\rho_1(a_1|\theta_1)u_2(\theta_1; a_1, a_2')$$

where (i) is the equilibrium condition for a cursed player 2, while (ii) is the equilibrium condition for a non-cursed player 2. If (i) holds, for $\chi = 1$ $(\rho_1, \overline{\rho}_2)$ with $\overline{\rho}_2$ such that, for each $a_2 \in A_2$, $\overline{\rho}_2(a_2) = \beta\rho_2(a_2|C) + (1-\beta)\rho_2(a_2|NC)$ is a χ-cursed equilibrium. If (ii) holds, for $\chi = 0$ $(\rho_1, \overline{\rho}_2)$ is a χ-cursed equilibrium. Hence in both cases we are able to find a χ in the unit interval such that the χ-cursed player 2's equilibrium behaviour replicates the average equilibrium behaviour of player 2 in the BC-game.■

References

[1] BATTIGALLI, P. (1987): "Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali", *unpublihed dissertation*, Università Commerciale "L.Bocconi", Milano.

[2] BATTIGALLI, P., GILLI, M. and C.MOLINARI (1992): "Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey", *Ricerche Economiche*, 3-4, 335-378.

[3] BATTIGALLI, P. and P.GUAITOLI (1997) "Conjectural Equilibrium and Rationalizability in a game with incomplete information", in *Decision, Games and Market*s, Kluwer Academic Publihis, Dodrecht.

[4] DEKEL, E., FUDENBERG, D. and D.K. LEVINE (2004): "Learning to play Bayesian games", *Games and Economic behavior* 46, 282-303.

[5] EYSTER, E. and M.RABIN (2005): "Cursed Equilibrium", *Econometrica*, 73, 1623-1672.

[6] FUDENBERG, D. (2006): "Advancing beyond *Advances in Behavioral Economics*", *Journal of Economic Literature*, 44, 694-711.

[7] FUDENBERG, D. and D.K.LEVINE (1993a): "Self Confirming Equilibrium", *Econometrica*, 61, 523-545.

[8] FUDENBERG, D. and D.KREPS (1995): "Learning in extensive-form games. I. Self-confirming equilibria", *Games and Economic behavior*, 8, 20-55.

[9] JEHIEL, P. (2005): "Analogy-based expectation equilibrium", *Journal of Economic Theory,* 123, 81-104.

[10] JEHIEL, P. and F.KOESSLER (2008): "Revisiting games with incomplete information with analogy-based expectation", *Games and Economic behavior*, 62, 533-557.

[11] MIETTINEN, T. (2007): "Learning foundation for the Cursed Equilibrium", *working paper*, Max Planck Institute of Economics.