DECLARATORIA SULLA TESI DI DOTTORATO
Da inserire come prima pagina della tesi

Il/la sottoscritto/a

COGNOME	Gamba
NOME	Astrid
Matr.	1094978

Titolo della tesi:

How Cognitive Limitations and Behavioral Traits Affect Stable Economic Outcomes

Dottorato di ricerca in	Economics
Ciclo	XXI
Tutor del dottorando	Prof. Pierpaolo Battigalli
Anno di discussione	2010

DICHIARA

sotto la sua responsabilità di essere a conoscenza:

1) che, ai sensi del D.P.R. 28.12.2000, N. 445, le dichiarazioni mendaci, la falsità negli atti e l'uso di atti falsi sono puniti ai sensi del codice penale e delle Leggi speciali in materia, e che nel caso ricorressero dette ipotesi, decade fin dall'inizio e senza necessità di nessuna formalità dai benefici previsti dalla presente declaratoria e da quella sull'embargo;

2) che l'Università ha l'obbligo, ai sensi dell'art. 6, comma 11, del Decreto Ministeriale 30 aprile 1999 prot. n. 224/1999, di curare il deposito di copia della tesi finale presso le Biblioteche Nazionali Centrali di Roma e Firenze, dove sarà consentita la consultabilità, fatto salvo l'eventuale embargo legato alla necessità di tutelare i diritti di enti esterni terzi e di sfruttamento industriale/commerciale dei contenuti della tesi;

3) che il Servizio Biblioteca Bocconi archivierà la tesi nel proprio Archivio istituzionale ad Accesso Aperto e che consentirà unicamente la consultabilità on-line del testo completo (fatto salvo l'eventuale embargo);

4) che per l'archiviazione presso la Biblioteca Bocconi, l'Università richiede che la tesi sia consegnata dal dottorando alla Società NORMADEC (operante in nome e per conto dell'Università) tramite procedura on-line con contenuto non modificabile e che la Società Normadec indicherà in ogni piè di pagina le seguenti informazioni:

- tesi di dottorato: *How Cognitive Limitations and Behavioral Traits Affect Stable Economic Outcomes*;
- di *Gamba Astrid*;
- discussa presso l'Università commerciale Luigi Bocconi – Milano nell'anno *2010;*

- La tesi è tutelata dalla normativa sul diritto d'autore (legge 22 aprile 1941, n.633 e successive integrazioni e modifiche). Sono comunque fatti salvi i diritti dell'Università Commerciale Luigi Bocconi di riproduzione per scopi di ricerca e didattici, con citazione della fonte;

5) che la copia della tesi depositata presso la NORMADEC tramite procedura on-line è del tutto identica a quelle consegnate/inviate ai Commissari e a qualsiasi altra copia depositata negli Uffici dell'Ateneo in forma cartacea o digitale e che di conseguenza va esclusa qualsiasi responsabilità dell'Ateneo stesso per quanto riguarda eventuali errori, imprecisioni o omissioni nei contenuti della tesi;

6) che il contenuto e l'organizzazione della tesi è opera originale realizzata dal sottoscritto e non compromette in alcun modo i diritti di terzi (legge 22 aprile 1941, n.633 e successive integrazioni e modifiche), ivi compresi quelli relativi alla sicurezza dei dati personali; che pertanto l'Università è in ogni caso esente da responsabilità di qualsivoglia natura, civile, amministrativa o penale e sarà dal sottoscritto tenuta indenne da qualsiasi richiesta o rivendicazione da parte di terzi;

7a) che la tesi di dottorato non è il risultato di attività rientranti nella normativa sulla proprietà industriale, non è stata prodotta nell'ambito di progetti finanziati da soggetti pubblici o privati con vincoli alla divulgazione dei risultati; non è oggetto di eventuali registrazioni di tipo brevettale o di tutela, e quindi non è soggetta a embargo;

Data: 03Febbraio 2010

F.to (indicare nome e cognome): Astrid Gamba

How Cognitive Limitations and Behavioral Traits Affect Stable Economic Outcomes

Astrid Gamba

Bocconi University

February 3, 2010

To my father

Acknowledgements

Most of all, I would like to thank my advisor Pierpaolo Battigalli. I also thank Matthias Messner, Topi Miettinen, Philippe Jehiel and Nenad Kos for useful suggestions.

Contents

Chapter 1

Introduction

This thesis is made of three chapters which have two main threads. The first is the research subject: all these works have in common the study of *non-standard types* in a game theoretic setting. The second concerns the methodology: the common feature of the three chapters is the underlying *adaptive approach.* Both these two aspects constitute a substantial departure from the neoclassical approach to the study of economic behavior.

Non standard types There is a large amount of experimental results showing that agents do not always behave in the way standard theories would predict. It seems that the *homo economicus,* fully rational and self-interested, is not representative of the majority of the economic agents. The big challenge for the economic theory is to construct frameworks and provide models to describe and possibly explain aspects of the economic behavior that neoclassical theories cannot capture. In this thesis, we take on this challenge.

The main characteristic of this work is the assumption that the population of all economic agents includes *non-standard types.* For 'non-standard' we mean non self-interested and/or non fully rational in the neoclassical sense. Indeed, on one side, we assume that economic behavior is not driven only by preferences on consumption bundles (or lotteries over them). On the other side, we assume that agents do not always behave in a fully rational manner according to their preferences. The first feature characterizes the second and third chapters, where we analyse the *evolution of other-regarding preferences.* The second feature mainly regards the first chapter where we introduce a model where part of the players are assumed to be *boundedly rational* and part of them fully rational. However, the two aspects overlap with each other throughout the thesis. In all the three chapters we deal with a population which is *heterogenous* along some dimensions. Either we have heterogeneity in the preference types or in the degree of players' sophistication. In all the three chapters, we do have standard types (in the neoclassical sense) interacting with non-standard ones in the same strategic contexts.

In the first chapter (*Binary cursed equilibrium*) each player together with his payoff type is endowed with a *cursedness type* which represent the degree of his strategic sophistication.[1] Players might be either cursed or non-cursed (sophisticated in the standard sense). A cursed player, does not play in the way standard theory would predict because he holds incorrect conjectures on the opponents' play and best reply to them. In particular, a cursed player underestimate the informational content of other players' play.

In the second chapter (*A note on the evolution of other-regarding preferences*) we study the evolution of *altruistic preferences* in the context of a centipede game. In that model, player might be either self-interested or altruistic. An altruistic

[1] This work is strictly related to Eyster and Rabin (2005) which introduces the concept of "cursed equilibrium".

type does not maximize only his (expected) material payoff but he cares also for the material payoffs of his co-players. To consider only one preference type other than selfish is obviously restrictive. However, the strong assumption that there are only altruists and selfish players in the population is enough for the specific purposes of the note. Of course, if we want to construct a more general model on the evolution of preferences, we must take into consideration further types as alternative to the self-interested one.

That is the reason why in the third chapter ("*The evolution of preferences in a game of life*") we consider all possible types of other-regarding preferences which differ from the selfish type. We express a very generic functional form for the players' utility which consists in the (expected) material payoff plus a disposition term. This term might represent, for example, players' inclination toward the material payoffs of the co-players, care for egualitarian outcomes, reciprocity, envy and whichever subjective consideration or emotion may drive agents' behavior. For a selfish player this disposition term is simply equal to zero.

The adaptive approach There is a strong connection between these research subjects and the methodology we use. Throughout the thesis we adopt an *adaptive approach* to study economic behavior. On one hand, in Chapter 1, we provide a justification in terms of (implicit) *learning* for the existence of non-sophisticated types and we describe the strategic interaction among differently sophisticated players with a solution concept that has a plausible learning interpretation (namely *binary cursed equilibrium*). On the other hand, in Chapter 2 and 3, we intend to provide *evolutionary foundations* for the existence of heterogeneous preference types adopting an *indirect evolutionary approach* (introduced by Guth and Yaari

(1992) and Guth (1995)). Learning and evolutionary considerations are proper tools to justify from a theoretical point of view the widely observed fact that agents do not reason in the same way and/or their behavior is not driven by the same preferences. Learning processes develop in the short run: given their preferences players' *behavior adjusts* to conjectures on opponents' play which are formed on the basis of past observations. Differently, evolutive processes develop in the long run: *preferences adjust* over a long-run evolutionary dynamics on the basis of their relative fitnesses.

In the first chapter (*Binary cursed equilibrium*) we will use only the first notion of adaptation. Players are assumed to differ exclusively in the way they reason about their co-players' play. To analyse strategic interactions where some players have cognitive limitations (incorrect conjectures) we cannot refer to the notion of Nash equilibrium (or refinements) which is based on strong rationality assumptions. This is the reason why we do choose a weaker solution concept. Indeed, binary cursed equilibrium is a particular *Self-confirming equilibrium* (SCE)[2]. SCE describes stable outcomes 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 on whether they are correct or not (*beliefs'confirmation*). We will point out that binary cursed equilibrium can be seen as a stable outcome resulting from repeated and anonymous interactions. Players differ *a priori* in their access to relevant (ex post) information about past plays so that less informed players end up in equilibrium having incorrect conjectures about their co-players

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

behavior. The underlying hypothesis is that players learn how to play by their experience. What players observe after each play along their adjustment process is not always enough to lead them to behave optimally as the standard theory would predict. We do observe sub-optimal behavior because players have limited access to (ex post) information and this prevents them from holding *correct conjectures* (in equilibrium) on the opponents' play.

In the last two chapters we adopt the *indirect evolutionary approach* to provide theoretical justifications for the existence of non-standard preferences. The indirect evolutionary approach studies whether certain *preferences* are evolutionarily successful. Suppose there is an heterogenous population composed of different preference types. Individuals endowed with their specific preferences are repeatedly and randomly drawn to play a basic game with material payoffs. In each round, players behave rationally, maximizing their expected utility associated to their own preferences. The evolutive success of a certain preference type is evaluated on the basis of the material payoffs (*objective fitness*) induced by the equilibrium profile of strategies adopted. Agents whose preferences lead to higher material payoffs (higher fitness) tend to reproduce faster than those with lower material payoffs (lower fitness).

The literature on the evolution of preferences adopts Nash equilibrium (or variants of it) as a rule to describe behavior in any relevant state of the long-run evolutionary dynamics. Differently, in the second chapter (*Note on the evolution of other-regarding preferences*) we adopt SCE to pin down the evolutive fitness of a preference type. The underlying assumption is that evolution develop through a *two-speeds adaptive process.* There is a short-run process whereby (given a distribution of preference types) players' behavior adjust untill they reach the play of a

SCE. These stable states, corresponding to the play of a SCE, in turn constitute the rounds of an overall *long term evolutionary process* along which the population composition adjusts according to a fitness criterion. We will show that the departure from the assumption that a Bayesian Nash equilibrium is played in each relevant state of the long-term evolutionary dynamics provides new insights on the evolution of altruistic preferences.

In the last chapter ("*The evolution of preferences in a game of life*"), we introduce a model which intend to describe the evolution of other-regarding preferences assuming that economic agents choose the games to play (within a large but finite set of games) together with their strategies within each game. We still adopt the indirect evolutionary approach. However, we depart from the adoption of SCE to characterize stable outcomes of the short-run adaptive processes (within each game). We assume that the game-specific plays in each relevant state of the long-run evolutionary dynamics corresponds to a Bayesian Nash equilibrium. It is worth noticing that this occurs only because this assumption makes the model, which is already complex, more treatable.

Chapter 2

Binary Cursed Equilibrium

2.1 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 neglet 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 true 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. 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 reinterpret and to provide a learning foundation for PCE. He shows that any PCE may be seen as an *analogy-based expectation equilibrium.*

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

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*).[3] In a PCE and in a FCE players typically hold *incorrect joint beliefs* on opponents' types and actions but their marginal beliefs are correct. 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

[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

[3]This is what Battigalli (1987) calls "conjectural equilibrium". Fudenberg and Levine (1993) coined the term "self-confirming equilibrium" with reference to the specific case where the terminal node of the extensive game tree is observed *ex post.* We use SCE in a more general way, allowing for coarser ex post information structure.

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 and the correct conjecture).

In this paper we show that this pessimism with respect to the possibility of justifying partial sophistication is not warranted. In particular, we will provide a natural interpretation of partial cursedness 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 under some restrictions on the game being played 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 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 cursedness 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 (ex post) information about the outcomes of the game which he is playing. This implies that he is not able to link his opponents' actions to their payoff types and he forms the simplest conjecture consistent with his (ex post) observations. 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 model (i.e. a share which is constant across all payoff types of all players), such that the BC model 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.

2.2 Cursed Equilibrium

Throughout the paper we will consider an environment with only two players. Denote the set of their conceivable payoff types, independently distributed, by $\Theta_i = \left\{\theta_i^A, \theta_i^B\right\}$, $i = 1, 2$ 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 heterogeneous population of agents who can play in role j. Player j is randomly drawn from this population. 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 i the set of possible actions is $A_i = \{a_i', a_i''\}$.

A behavior strategy for player i is a mapping from the set of his payoff types to the set of his mixed actions. Denote by σ_i the "true" strategy of player i, $\sigma_i : \Theta_i \to \triangle A_i$. Let $\Sigma_i \equiv (\triangle A_i)^{\Theta_i}$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 $\overline{\sigma}_i(a_i) := \sum_{\theta_i \in \Theta_i} p^i(\theta_i)\sigma_i(a_i|\theta_i)$, $i = 1, 2$, $a_i \in A_i$. That is, $\overline{\sigma}_i(a_i)$ is the fraction of agents in population i playing 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. According to PCE the beliefs of a player j are given by the convex combination $\widehat{\sigma}_i = \chi\overline{\sigma}_i + (1-\chi)\sigma_i$, that is 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_j} the degree of cursedness of type θ_j of player j. 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 responds to the correct conjecture as in a Bayesian Nash equilibrium. For intermediate values of χ a *partially cursed equilibrium (PCE)* is played.

2.2.1 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[4] (SCE). If

[4]See Fudenberg and Levine (1993) and Dekel *et al.* (2004). 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.

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 characterizes the stationary states of 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 the link between partially cursed equilibrium and SCE is perhaps less clear. 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 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.

2.3 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 behavioral strategy of player j

is a mapping from the set of types T_j into the set of his mixed actions. As above, the set of every player's conceivable actions does not change across types. Hence, allowing for randomized strategies, we define $\rho_j : T_j \to \triangle A_j$, $j = 1, 2$. 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 of action a_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 (θ_i, NC) *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 (θ_i, C) *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 $\overline{\rho}_i \equiv (\overline{\rho}_i(\theta_i))_{\theta_i \in \Theta_i}$, $i = 1, 2$.

It is worth stressing that we do *not* make any implicit assumption on players' awareness of other players' cursedness. In other words we do not assume that the conjecture on action a_j being played by player j with payoff type θ_j which an uncursed player i has in mind is structured as the convex combination between $\rho_j(a_j|\theta_j, C)$ and $\rho_j(a_j|\theta_j, NC)$. In equilibrium, we only require that an uncursed player holds a correct conjecture on the payoff-type dependent strategy of his co-player *regardless of his co-player cursedness type*, while a cursed player holds a correct conjecture on the probability with which each action of the co-player is played regardless of his co-player payoff type *and* cursedness type.

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, for some simple environments, Eyster and Rabin's PCE concept can be seen as a reduced form of BCE, this will allow us to conclude that for those environments also PCE can be seen as a steady state of a learning process.

Our learning framework has two possible interpretations: either players differ in their ex post information structure or they differ in their strategic sophistication. We explain how these differences naturally deliver situations where (in the long run) part of the agents will be able to infer all details about their opponents strategies while other agents 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. Hence, payoff types are independent and their statistical distributions are commonly known[5]. Each player i knows his own payoff type, the set of payoff types of the opponent Θ_j, his

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

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 each agent inherits 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 actions 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 non doctrinaire beliefs about the statistical distribution of actions and types in the population of agents playing the opponents' role.[6] Through Bayesian updating they learn the true correlation between actions and types of the opponents. On the other hand, cursed players

[6]See Fudenberg and Levine (1993a)

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.[7] 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.[8] 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.

[7] 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.

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

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

2.4.1 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, player 1 cannot be cursed: $T_1 = \Theta_1 \times \{NC\}$, where $\Theta_1 = \left\{\theta_1^A, \theta_1^B\right\}$. On the other hand, the type set of player 2 coincides with his cursedness type set $T_2 = \left\{\overline{\theta}_2\right\} \times B_2$. where $B_2 = \{C, NC\}$. Let β be the probability that player 2 is cursed and $(1 - \beta)$ the probability that he is non-cursed. A behavioral strategy for player 1 is a function $\rho_1 : \Theta_1 \rightarrow \triangle A_1$; a behavioral strategy for player 2 is a function $\rho_2 : B_2 \rightarrow \triangle A_2$. We denote by $r_2 : B_2 \rightarrow A_2$ 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-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 strictly 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.

Let $\chi = \overline{\chi}$, $\overline{\chi} \in (0,1)$. 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 $a_2 \in A_2$, $\rho_2(a_2) = \beta\rho_2(a_2|C) + (1-\beta)\rho_2(a_2|NC) = \sigma_2(a_2)$.

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.*

2.4.2 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 whether 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, $\Theta_i = \{\theta_i^A, \theta_i^B\}$, $i = 1, 2$. Suppose that payoff types are equally likely. 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 (r) is $\overline{\sigma}_2(l) = 1/2$ ($\overline{\sigma}_2(r) = 1/2$), according to the average behavior $\overline{\sigma}_2$ of player 2 that player 1 has in mind. 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 $\beta_{\theta_1^A} = 0$, while to replicate the $\overline{\chi}$-cursed equilibrium strategy of type θ_1^B, we need to take $\beta_{\theta_1^B} = 1$. 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.$\Diamond$

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. Suppose that payoff types are equally likely. 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* $\boldsymbol{\beta}$ *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.*

2.5 Conclusions

In this paper we have discussed a possible theoretical interpretation of Partially Cursed Equilibrium. 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 learning foundation. In particular, we have discussed two 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.

2.6 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, Pierpaolo**: "Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali", *unpublished dissertation*, 1987, Università Commerciale "L.Bocconi", Milano.

[2] **Battigalli, Pierpaolo**; **Gilli, Mario and Molinari, M.Cristina**: "Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey", *Ricerche Economiche* 1992, *3-4*, 335-378.

[3] **Battigalli, Pierpaolo and Guaitoli, Danilo**: "Conjectural Equilibrium and Rationalizability in a game with incomplete information", in *Decision, Games and Market*s, Kluwer Academic Publihis, Dodrecht, 1997.

[4] **Dekel, Eddie; Fudenberg, Drew and Levine, David K.** "Learning to Play Bayesian Games", Games and Economic Behavior 2004, *46*, 282-303.

[5] **Eyster, Eric and Rabin, Matthew**: "Cursed Equilibrium", *Econometrica*, 2005, *73*, 1623-1672.

[6] **Fudenberg, Drew**: "Advancing beyond *Advances in Behavioral Economics*", *Journal of Economic Literature*, 2006, *44*, 694-711.

[7] **Fudenberg, Drew and Levine, David K.**: "Steady State Learning and Nash Equilibrium", *Econometrica*, 1993a, *61*, 547-573.

[8] **Fudenberg, Drew and Levine, David K.**: "Self Confirming Equilibrium", *Econometrica*, 1993b, *61*, 523-545.

[9] **Fudenberg, Drew and Kreps, David**: "Learning in extensive-form games. I. Self-confirming equilibria", *Games and Economic behavior*, 1995, *8*, 20-55.

[10] **Jehiel, Philppe**: "Analogy-based expectation equilibrium", *Journal of Economic Theory, 2005, 123*, 81-104.

[11] **Jehiel, Philpe and Koessler, Frederic**: "Revisiting games with incomplete information with analogy-based expectation", *Games and Economic behavior*, 2008, *62*, 533-557.

[12] **Miettinen, Topi**: "Learning foundation for the Cursed Equilibrium", *Economic Letters*, 2007, forthcoming.

Chapter 3

A Note on the Evolution of Other-regarding Preferences

3.1 Introduction

Most of the literature which studies the evolutionary foundations of other-regarding preferences adopts the *indirect evolutionary approach* (Guth and Yaari (1992) and Guth (1995)). While the standard approach in evolutionary game theory consists in investigating whether a *strategy* is robust to evolutionary selection pressures, the indirect approach studies whether certain *preferences* are evolutionarily successful. Suppose there is an heterogenous population composed of different preference types (eg. altruistic and selfish). Individuals endowed with their specific preferences are repeatedly and randomly drawn to play a *basic game with material payoffs* (eg. monetary payoffs). In each round, players behave rationally, maximizing their expected utility associated to their own preferences (which does not necessarily match the underlying expected material payoff). The evolutive success of a certain preference type is evaluated on the basis of the material payoffs (*objective fitness*) induced by the profile of strategies adopted. Agents whose preferences lead to higher material payoffs (higher fitness) tend to reproduce faster than those with lower material payoffs (lower fitness).

It should be noted that the use of this treatment poses a potentially difficult problem. In fact, while in the traditional approach individuals are identified with strategies (each agent is programmed to play a specific strategy), in the indirect evolutionary approach individuals are identified with preference types. Consequently, we need to choose a *rule which maps profile of preferences into behavior* to evaluate the evolutive fitness of a certain preference type given the types of others. Potentially, we can think of the evolution of preferences as a two-speed dynamic process. There is a *short term adaptation process* whereby, given the preferences composition of the population and the information structure, players adjust their behavior untill they reach some plausible stationary states (equilibrium play). These states in turn constitute the rounds of an overall *long term evolutionary process* along which the population composition adjusts according to a fitness criterion. Thus, the choice of the static solution concept which captures players' behavior in any relevant state of the evolutionary dynamics (when the distribution of preferences is given) becomes crucial.

The literature on the evolution of preferences adopts Nash equilibrium (or variants of it) as a rule to describe behavior in any state of the evolutionary dynamics. A common feature of these studies is the result that evolution can favor non-materialistic preferences *if* players observe their opponents' preference types at least to some degree. For example, Guth and Yaari (1992) argue that *observability* is the driving force for the evolution of interdependent preferences (altruistic, reciprocal,...) by means of a commitment effect. Bester and Guth (1998) show that when the context exhibits strategic complementarities and the players can observe their opponent's preference type natural selections favors altruism. Conversely, if opponents' preferences are not observable evolutionary forces favor preferences which coincide with the material payoff (Ok and Vega-Redondo (2001) and Ely and Ylankaya (2001)). In particular, Dekel, Ely and Ylankaya (2007) assume that in each state of the long term dynamics agents play a Bayesian Nash equilibrium given their preferences and the information about others' preferences. They show that, if players know the distribution of preference types in the population but do not observe the opponent's preferences any non-Nash equilibrium outcome of the underlying game with fitness payoffs can be destabilized by an entering population with materialistic preferences.

In the next section, we will discuss an application of the indirect evolutionary approach to the Centipede Game. We will show that by adopting the solution concept of self-confirming equilibrium to capture the play in the relevant state of the dynamics there is room for altruistic preferences to evolve even if preferences are unobservable. Before going into the details of the model it is worth discussing two preliminary observations.

The first point we want to make is that it is hard to justify the use of Nash equilibrium to capture the limit play of the short run adjustment process. Indeed, if we follow the *eductive interpretation* of Binmore (1987) and we assume that in each state of the long term dynamics agents play the basic game *only once*, we need to explain how they can logically derive and play a Bayesian Nash equilibrium (BNE). To justify the play of a BNE, we must necessarily rely on assumptions like rationality and common certainty of rationality, given the common knowledge of the (Bayesian) game. First of all, these assumptions naturally deliver rationalizable profile of strategies for Bayesian games and they do not imply in general that a BNE is played. Secondly, even if they only deliver the unique BNE, they remain strong epistemological assumptions. Alternatively, we can assume that the equilibrium play in each state of the long term dynamics is the result of a (short run) adaptive process (*adaptive justification*)[1]. In this case, we would need to be more explicit about the short-term learning dynamics to show that the play surely converges to a certain BNE. In particular, we should specify what players can observe regarding the outcomes of previous interactions to explain how they could end up in equilibrium holding a *common and correct belief about the play* (whatever their preferences). For example, in static games, under the assumption of private values, if a player observes the actions played in every round by the opponent he can learn the correct probabilities of each of the opponent's actions[2]. Differently, if the underlying game is dynamic to observe only the actions of the opponent may not be enough. We should be aware that there can exist stationary states that do not correspond to BNE. Indeed, given a certain ex post information structure the players' conjectures about the opponents' behavior can be confirmed even if they are not correct (and common to all players). With these considerations in mind, we may want to adopt a *weaker* solution concept than BNE.

Secondly, we want to stress that observability emerges in the above cited studies as a necessary condition to sustain the evolutive success of non-materialistic preferences only because it is assumed that in every state of the long-term dynamics

[1] This is an implicit assumption in Guth and Yaari (1992), Bester and Guth (1998) and Dekel, Ely and Ylankaya (2007).

[2] If we had interdependent values (but it is not the typical case in the cited literature) we would need to assume, for example, that players have access to some public joint statistics on actions *and* preference types of the opponent.

a BNE is played. Indeed, if preferences are not observable, the correctness of conjectures about opponents' behavior implies that selfish players (expected material payoff maximizers) will always obtain the highest possible payoff. Consequently, they will unavoidably perform either the same or better in terms of fitness than the non-material utility maximizers. If instead preferences are observable, the altruists do not succumb when facing a selfish, by discriminating their behaviour dependingly on the type of the player they face.

We suggest that self-confirming equilibrium (SCE) is an appropriate solution concept to describe the play in the relevant states of the evolutionary dynamics[3]. Since self-confirming equilibrium has a natural learning foundation, it is suitable to represent the stationary states of the short run adaptive processes. Essentially, the SCE describes 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 revise those conjectures, independently on whether they are correct or not (*conjectures' confirmation property*). The feature of a SCE which matters here is that it allows situations where players hold *heterogenous beliefs about the play* as long as these beliefs are not contradicted by the evidence. From an adaptive perspective we can justify such an equilibrium situation by arguing that, by behaving differently, individuals with different preferences may accumulate different experience through the learning process (if they rely only on their

[3]We will consider the version of SCE proposed by Dekel, Fudenberg and Levine (2004), applied to extensive form games, similar in the spirit to the notion of conjectural equilibrium proposed by Battigalli (1987) and Battigalli and Guaitoli (1997). The original version of SCE introduced in Fudenberg and Levine (1993) is not appropriate in this context: there the ex post information structure is such that players can observe both the moves of the opponents and of nature. But if this is the case then ex post observability of preference types is implied.

own observations). We show in the next section that by weakening in this sense the assumption of correctness of conjectures we can reach further insights into the evolution of non-materialistic preferences.

3.2 The indirect evolutionary model

3.2.1 The scenario

Consider a population composed of individuals with *heterogeneous preference types* and *heterogeneous beliefs* whose members interact with each other in pairs. Each player i can be either an altruistic or a selfish type, that is for every i, $\Theta_i = \{\alpha, \sigma\}$, where α means altruistic and σ selfish. An altruistic type aims at maximizing the joint (material) payoff, whereas a selfish type aims at maximizing his own material payoff. We define the material consequences of players' actions with the functions $\mathbf{m}_i : Z \to \mathbb{R}$, $i \in N$, where Z is the set ot terminal histories. While the utility of a selfish player i coincide with his material payoff, $U_{i,\sigma}(z) = \mathbf{m}_i(z)$, the preferences of an altruistic player i are represented by utility functions of the form: $U_{i,\alpha}(z) = \mathbf{m}_1(z) + \mathbf{m}_2(z)$.

We denote q the measure of altruistic types in the population. We do not make any explicit assumption on the players' knowledge of q: players might not have any clue of the distribution of preference types in the population. Actually, they might not even be aware of the existence of preference types different from their own.

Denote $\mu_{i,\theta_i}(\cdot) \in \Delta(S_j)$ the belief of player i with preference type θ_i on the strategy of player j.

We assume that players endowed with their preference types and their beliefs are *repeatedly* and *randomly* drawn to play a two-player game with material payoffs (monetary payoffs). We assume that each individual can end up with a probability of 1/2 either in the role of player 1 or player 2. Suppose that players do not observe at all the preference type of the individual they are matched to play with. Moreover, suppose that, after having played, they can observe only the actions actually taken by their co-player. We assume that in each round, given a preference distribution q and the information structure the agents play a SCE of the underlying game. That is, players best respond to their (heterogeneous) conjectures on the opponent's behavior by maximizing their *perceived* expected utility and the information revealed ex post, after the equilibrium play, confirm their conjectures. Each SCE equilibrium play determines the *objective fitness* of each preference type involved in the game by means of the material payoff obtained by the player endowed with that preference type. We evaluate the evolution of altruistic preferences by considering a *replicator dynamics*. Denote $\overline{m}_\theta(q)$ the average payoff of preference types θ in a typical state of the evolutionary dynamics where the fraction of altruists in the population is q. Call $\overline{m}(q)$ the current average payoff in the population. The growth rate for the population fraction q with altruistic preferences equals the difference between the current average payoff of individuals with these preferences ($\overline{m}_\alpha(q)$) and the current average payoff in the population ($\overline{m}$), that is: $\frac{dq}{dt} = [\overline{m}_\alpha(q) - \overline{m}(q)]\, q$.

With this scenario in mind, we will next discuss the example of the (three-stages) Centipede Game. We will consider an evolutionary dynamics where in each state, given a preference distribution q, players play a SCE of the Centipede Game, given the ex post information structure specified above and their beliefs on the opponent's play in that state. We will show that starting within a large subset of the simplex of all possible beliefs such an evolutionary dynamics can favor altruistic preferences.

3.2.2 The evolutionary Centipede Game

Consider this three periods version of the Centipede Game:

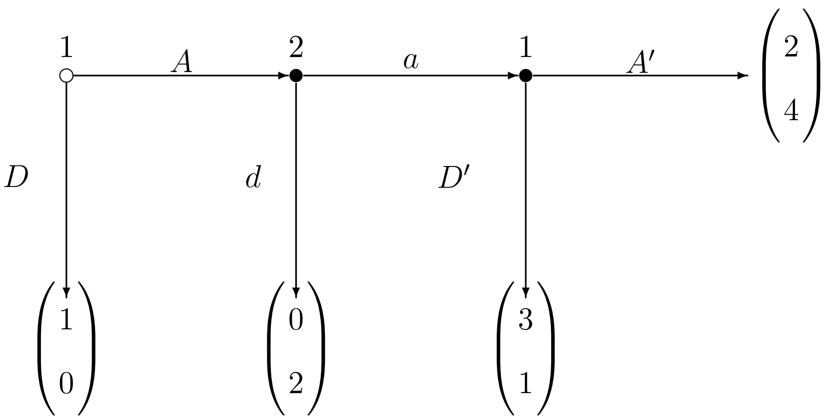

Figure 1. Centipede Game with Monetary Payoffs.

Notice that the numbers attached to the terminal nodes represent players' material payoffs (money) and do not necessarily coincide with their utilities.

Suppose that a selfish individual drawn to play in the role of player 2 has a system of beliefs on the strategies of the co-player 1 satisfying the following conditions:

i) his beliefs on the strategies of player 1 have full support, that is supp $\mu_{2,\sigma}(\cdot) = S_1$;

ii) he attaches a small probability to action A' being played after A, that is

$\mu_{2,\sigma}(AA') < 1/3$.

Suppose that a selfish individual drawn to play in the role of player 1 has a system of beliefs concerning the strategies that might be played by individuals drawn to play in 2's position satisfying the followings:

iii) his beliefs on the strategies of player 2 have full support that is $\operatorname{supp} \mu_{1,\sigma}(\cdot) = S_2$;

(iv) he attaches a small probability to action a being played by 2, that is $\mu_{1,\sigma}(a) < 1/3$.

Such a system of beliefs constitute a "large" relatively open subset of the simplex of all possible beliefs. The unique restriction is that all selfish individuals believe that the strategies AA' and a are unlikely to be played. According to these beliefs a selfish individual who has to play in 2's role will choose d, while a selfish individual who has to play in 1's position will choose DD'.

We do not have to make any restriction on the system of beliefs of an altruistic type. Indeed, whatever his beliefs, given that he wants to maximize the expected joint payoffs, he will play "across" whenever is his turn to move.

To sum up, with the systems of beliefs described above, selfish players choose to stop whenever possible, while altruistic players choose always to continue when they have the possibility to do so. In every state of the evolutionary dynamics (i.e. for every q) this profile of strategies constitutes a SCE of the underlying game. In fact, each preference type of each player is maximizing his (perceived) expected utility and given the ex post information structure none of them revise his beliefs on the equilibrium strategies of the opponent. The probability that the strategy AA' is chosen by player 1 coincides with the probability that the type drawn to play in 1's role is altruistic, that is q, while the probability that DD' is chosen coincides with the probability that the type drawn to play in 1's role is selfish, i.e. $(1-q)$. Similarly, the probability that a is chosen coincide with the probability that the type drawn to play in 2's role is altruistic (q), while the probability of d is exactly $(1-q)$. From an adaptive perspective, by sticking with going down whenever they can, selfish players do not have the opportunity to learn these exact probabilities. Expecting D' after a with high probability, they keep playing d after A (when they are in 2's role), so that they are prevented from learning that actually D' is never chosen after Aa. Similarly, expecting d after A with high probability, they keep playing "down" immediately (when they are in 1's role) so that they are prevented from learning that the probability of a is q. Their wrong conjectures will always be confirm. Differently, altruistic types may end up having in equilibrium a complete picture of the behavior of the opponent but still the correctness of their beliefs is not an issue. Given their preferences they choose to go "across" whenever they have the possibility to do so, whatever their beliefs. Whether they learn the exact frequencies of each action or not, they would behave altruistically anyway. What is crucial is that selfish individuals *do not learn* these exact frequencies and so, dependingly on the value of q, they might perform worse than altruistic individuals.

We evaluate the evolution of altruistic preferences by considering a replicator dynamics, where the replicators are here the preference types.

Recall that the equation of the replicator dynamics is the following: $\frac{dq}{dt} = [\overline{m}_{\alpha}(q) - \overline{m}(q)]\, q$.

We need to compute $\overline{m}_\alpha(q)$, the average SCE-payoff of an altruistic type, averaged across the roles he can take and across the preference types he can face. The average payoff in the sub-population of altruists is then: $\overline{m}_\alpha(q) = (\frac{1}{2}2 + \frac{1}{2}4)q = 3q$ (with probability $(1-q)$ he faces a selfish individual and he gets 0 whatever his role). Similarly, the average payoff of a selfish type is: $\overline{m}_\sigma(q) = (\frac{1}{2}1 + \frac{1}{2}2)q + (\frac{1}{2}1 + \frac{1}{2}0)(1-q)$. So, we can compute the average payoff in the population, that is: $\overline{m}(q) = 2q^2 + \frac{1}{2}q + \frac{1}{2}$. Hence, the dynamics is represented by the following equation: $\frac{dq}{dt} = [\overline{m}_\alpha(q) - \overline{m}(q)]\, q = \left[\frac{5}{2}q^2 - 2q^3 - \frac{1}{2}q\right]$, shown in Figure 2.

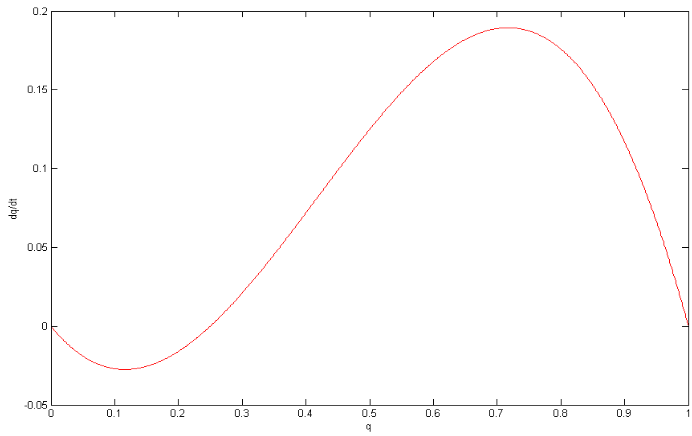

Figure 2. Replicator Dynamics For The Population Share q of Altruistic Types.

If we have an heterogeneous population (where the measure of altruists is high enough (>1/4)) individuals with altruistic preferences perform better than individual with selfish preferences and they can invade the whole population. More interestingly, Figure 2 and the reasoning above tell us that *a population of all altruists is not vulnerable with respect to the injections of a small share of selfish individuals.*

3.3 Conclusion

The literature on the evolution of preferences displayed a common result: conditional on playing some (Bayes) Nash equilibrium in every state of the evolutionary dynamics (when the preferences are fixed), only selfish individuals survive *if* the preference types are not observable. We have shown that by adopting the weaker solution concept of SCE, which allows for heterogenous beliefs across preference types, there is room for altruistic preferences to evolve even if preferences are not observable. Hence, adopting self-confirming equilibrium as a rule to pin down behavior in the relevant states of the evolutionary dynamics is promising and on top of all considerations it has a natural learning foundation.

It is worth noticing that we selected a *particular* self-confirming equilibrium and that, of course, there may be many for the same ex post information structure. However, we allowed for a very large set of initial beliefs imposing some restrictions only on the system of beliefs of selfish types. Nonetheless, the restrictions we have imposed are quite plausible: selfish players believe that the opponent will behave in the same way as they would behave if they were in his shoes. The implicit assumption is that they use introspection to form their beliefs.

Moreover, we have assumed that players in every state of the long term process play a self-confirming equilibrium *as if* they learned to play it. That is, we did not model explicitly the short-term learning dynamics that would lead to play that specific self-confirming equilibrium. By virtue of the fact that a self-confirming equilibrium can be typically learnt in an adaptive way, it would be interesting to describe the short-term behavioral adaptation. This analysis would more extensively support the choice of self-confirming equilibrium as a rule of mapping preferences into behavior in the relevant state of the long-term evolutive process.

References

[1] **Bacharach, Michael; Guerra, Gerardo and Zizzo, Daniel J.** “Is Trust Self-Fulfilling? An Experimental Study”, Working Paper, 2001.

[2] **Battigalli, Pierpaolo**. "Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali", unpublished dissertation, Milano: Bocconi University, 1987.

[3] **Battigalli, Pierpaolo and Guaitoli, Danilo.** "Conjectural Equilibria and Rationalizability in a Game with Incomplete Information", in *Decisions, Games and Markets*, P. Battigalli, A. Montesano and F. Panunzi, Eds.,Kluwer Academic Publishers, Dordrecht, 1997, 97-124.

[4] **Bester, Helmut and Guth, Werner.** "Is Altruism Evolutionary Stable?", *Journal of Economic Behavior and Organization*, 1998, *34*, 193-209.

[5] **Binmore, Kenneth**. "Modeling Rational Players: Part 1", *Economics and Philosophy*, 1987, *3*, 179-214.

[6] **Dekel, Eddie; Fudenberg, Drew and Levine, David K.** "Learning to Play Bayesian Games", Games and Economic Behavior 2004, *46*, 282-303.

[7] **Dekel, Eddie; Ely, Jeffrey C. and Ylankaya, Okan**. "Evolution of Preferences," *Review of Economic Studies*, 2007, *74*(3), 685-704.

[8] **Frank, Robert H.** "If Homo Economicus Could Choose His Own Utility Function, Would He Want one With a Conscience?", *The American Economic Review*, 1987, *77*(4), 593-604.

[9] **Fudenberg, Drew and Levine, David K.** "Self-Confirming Equilibrium", *Econometrica*, 1993, *61*(3), 523-545.

[10] **Gul, Faruk and Pesendorfer, Wolfgang.** "The Canonical Type Space for Interdependent Preferences", Working Paper, 2005.

[11] **Guth, Werner and Kliemt, Hartmut.** "The Indirect Evolutionary Approach: Bridging the Gap between Rationality and Adaptation", *Rationality and Society*, 1998, *10*(3), 377-399.

[12] **Guth, Werner and Yaari, Menahem**. "Explaining Reciprocal Behavior in a Simple Strategic Game". in *Explaining Process and Change-Approaches to Evolutionary Economics, University of Michigan Press*,1992, 23-24.

[13] **Guth, Werner and Yaari, Menahem.** "An Evolutionary Approach to Explaining Cooperative Behavior by Reciprocal Incentives", *International Journal of Game Theory*, 1995, *24*, 323-344.

[14] **Heifetz, Aviad; Shannon, Chris and Spiegel, Yossi.** "The Dynamic Evolution of Preferences", *Economic Theory*, 2007, *32*(2), 251-286.

[15] **Ely, Jeffrey C. and Ylankaya, Okan.**"Nash Equilibrium and the Evolution of Preferences", *Journal of Economic Theory*, 2001, *97*(2), 255-272.

[16] **Herold, Florian and Kuzmics, Christoph.** "Evolution of Preferences Under Perfect Observability: Almost Anything is Stable", Working Paper, 2008.

[17] **Herold, Florian and Kuzmics, Christoph.** "Evolutionary Stability of Discrimination Under Observability", *Games and Economic Behavior*. Forthcoming.

[18] **Hofbauer, Josef and Sigmund, Karl.** "The Theory of Evolution and Dynamical Systems", *Cambridge University Press*, 1998.

[19] **Ok, Efe and.Vega Redondo, Fernando**. "On the Evolution of Individualistic Preferences: An Incomplete Information Scenario", *Journal Of Economic Theory*, 2001, *97*, 231-254.

[20] **Rabin, Matthew**. "Incorporating Fairness into Game Theory and Economics," *American Economic Review*, 1993, *83*(5), 1281-1302.

[21] **Robson, Arthur J. and Vega-Redondo, Fernando.** "Efficient Equilibrium Selection in Evolutionary Games with Random Matching", *Journal of Economic Theory*, 1996, *70*(1), 65-92.

[22] **Samuelson, Larry.** "Introduction to the Evolution of Preferences", *Journal of Economic Theory*, 2001, *97*, 225-230.

[23] **Sethi, Rajiv and Somanathan**, **E.** "Preference Evolution and Reciprocity, *Journal of Economic Theory*, 2001, *97*, 273-297.

[24] **Weibull, Jorgen W.** "Evolutionary Game Theory", *Cambridge, MA,The MIT Press*, 1995.

Chapter 4

The Evolution of Preferences in the Game of Life

4.1 Introduction

In this chapter we present the project "The evolution of Preferences in the Game of Life". We intend to study the evolution of other-regarding preferences using an indirect evolutionary approach and assuming that individuals can play many different games. We assume that there is ah heterogenous population made of individuals with various preference types who have to choose which games to play and which strategy to take within each game. For a given statistical distribution of preference types in the population, choices of strategies within games and choices of games must be in equilibrium at the same time. Such equilibrium configurations constitute the relevant state of the long-run evolutionary dynamics along which preferences adjust. The evolutionary success of a preference type is evaluated on the basis of its objective fitness, that is the average material payoff of all the individuals with that preference type.

To the best of my knowledge, the literature on the evolution of preferences has been exclusively focused on the study of the evolutionary performance of preference types within the context of a *single* game. The most general model that comes to my mind is the one introduced by Dekel, Ely and Ylankaya (2007). The authors discussed the evolution of "subjective" (non-material) preferences considering the set of all possible modulus affine transformation utility functions. Still they assumed that the evolution of such preferences unfolds within a single game belonging to the class of static two-by-two simmetric games. Hence, the major departure of our work from the previous studies is the assumption that individuals play a *game of life,* meaning that they potentially play at the same time many different games belonging to an arbitrary and large set of strategic interactions. We do not impose any restriction on the classes of games that might be played: they could be either static or dynamic and the number of players (roles) might vary across games. We do believe that such an extension in the study of the evolution of preferences can provide more insights on the relative evolutionary performances of different preference types and so a theoretical justification to the existence of heterogeneous agents. Some preference types might benefit from playing certain games while they underperform with respect to others in some other games. Trivially, "tougher" players choose "tougher games". By assuming that evolution takes place within the play of a single game we cannot develop this intuition. By considering the game of life, instead, we can fully take into account eventual cross benefits that some preference type might exploit by playing various games.

As we have seen in the previous chapter, most of the literature on the evolution of preferences concludes that some degree of observability of the preference type of the co-players is needed in order to justify the existence of an heterogenous population. Suppose that there are only altruist and selfish agents in the population who meet each other in the play of a single (two by two) game. If we assume that before playing players do not observe each other's preference type and they play a (Bayesian) Nash equilibrium, it is hard to conclude that an altruist can perform better (in average) than a selfish agent. We have shown in the previous chapter that if we weaken the solution concept observability is not needed to justify the survival of altruists. Here, we want to go further. Similarly to Dekel, Ely and Ylankaya (2007) and Heifetz, Shannon and Spiegel (2006) we assume that preferences can be represented by a generic utility function. In particular, we assume that the utility each player gets from a strategy profile is given by his material payoff and a disposition term which accounts for any kind of subjective preferences (e.g. altruism, reciprocity, envy etc). We will stick with the assumption that players do not receive any signal of the co-players' preference types when matched to play. Our intuition is that individuals with similar preference types tend to concentrate on the same subsets of games. This would imply that, on one hand, non-selfish individuals can gain advantages from playing with types similar to their own and, on the other hand, selfish individuals cannot exploit completely their strategic advantages. In a nutshell, the distance between the evolutive performance of a selfish agent and the performance of a non-material payoff maximizer would be tempered.

It is worth noticing that to study the evolution of preferences in a very generic game of life pose many difficulties. Individuals belong to an heterogeneous population with all possibile preference types. Each of them faces many different games. Given a distribution of preference types in the overall population and given their conjectures on the equilibrium plays within each game, individuals choose the games to play via expected utility calculations. Their choices of games influence the statistical distribution of preference types within each game and consequently determine a distribution over equilibrium strategy profiles. This complex structure leaves few space for adopting solution concept weaker than (Bayesian) Nash equilibrium to capture the play within each game. For simplicity, we will assume that the equilibrium play within games is described by a Bayesian Nash equilibrium. For what concerns the equilibrium play across games we assume that a Quantal Response Equilibrium (QRE) is played (McKelvey and Palfrey (1995)). The Quantal Response Equilibrium assumes that better strategies are played more often than worse strategies but best strategies are *not always* played. The intuition is that the individual's expected utility from a strategy profile is subject to a random error. This implies that players do not compute the exact best responses to their conjectures. We will apply this concept to the choices of games: the probability of choosing a given game is a smoothly increasing function ot the average payoff for that game.

The Bayesian Nash equilibrium plays within games and the quantal response equilibrium in the (mixed) choices of games together form the overall equilibrium configuration in each relevant state of the long-run evolutionary dynamics. In this states we can pin down the average (material) payoff of each individual and hence the average (material) payoff of each preference type in the population. The relative objective fitness of each preference type is then the drive for the evolution of preferences *via* a replicator dynamics.

In the next section we describe the formal model we have in mind.

4.2 The model

Consider a population composed of individuals with *heterogeneous preference types.* This population is large but finite. Call Θ the set of all possible preference types in the population and $\theta \in \Theta$ the generic type. We only assume that $\Theta \subseteq \mathbb{R}$ and it contains a neighborhood of 0. We call μ the probability distribution on Θ, representing the distribution of preferences in the overall population (in a given state of the long-run evolutionary dynamics). We denote by $H = \{1, ..., h\}$ the finite set of individuals and by $i \in H$ the generic individual. There is a finite set of games individuals might play, $G = \left\{g^1, ..., g^L\right\}$ with tipical element g^l. The available games are the same for all individuals. We denote the finite set of players (roles) of the generic game g^l by $N^l = \{1, ..., n_l\}$ and by $j \in N^l$ the generic player of g^l.

Together with a preference type, every individual $i \in H$ is endowed with a profile of game-specific probability distributions over the set of possible roles, namely $\rho_i = (\rho_i^l)_{l=1,...L}$. We denote by $(\rho_i^l(j))_{j \in N^l}$ the array of probability measures over roles in game g^l for individual i, where $\rho_i^l(j)$ is the probability that individual i takes role j in game g^l according to the distribution ρ_i^l. Obviously, for every individual i and every role $j \in N^l$, $\rho_i^l(j) \geq 0$ and $\sum_{j \in N^l} \rho_i^l(j) = 1$. We assume that roles and preference types are independent. The idea is that each individual takes in g^l a certain role according to some personal characteristics. These personal characteristics, as for example the sex, the social status, the occupation etc., determine the propensity of individual i to fit some roles instead of others. Trivially, if an agent is a male he will hold (with probability one) the position of the man in a 'battle of the sexes' game. Similarly, if we consider a firm active in a 'entry game', it might be either an incumbent or an entrant depending upon its position in the market. It may also be the case that in a certain 'entry game' (e.g. a certain market) a firm can be qualified as an incumbent while in another 'entry game' (e.g. another market) it can be qualified as an entrant. Moreover, we need to take into account the possibility that an individual does not always play the same game in the same role but he might play in different roles with different frequencies.

We focus first on *players'* interaction at the level of the generic game g^l and we then discuss *individuals'* choice of games. Indeed, given the (expected) equilibrium play in each game and the associated average utility, each individual chooses which games to play. An agent plays *more frequently* games from which he expects higher (subjective) utility.

Before going into details, it is worth clarifying some crucial features of the model. In particular, we need to explain where the statistical distribution of preference types of the generic player j in game g^l comes from. In a standard analysis of the evolution of preferences, the preference type of player j would be randomly drawn from μ, the distribution of preference types in the *overall* population. Differently, in our context, the preference type of a player in game g^l is *randomly drawn* from a distribution which has as observation space *a subset* of the overall population. This subset is *endogenously* determined. Indeed,a generic player j of a game g^l is an individual i with a certain preference type θ_i who has *chosen* to play g^l (at least with positive probability) according to expected utility considerations and takes the role of player j according to a certain (individual specific) probability distribution. Hence, the fraction of individuals in the role of player j in game g^l with preference type θ_j is influenced by the fraction of preference type θ_j in the overall population, by the frequency with which each individual with that preference type plays the game g^l and by the probability distributions over roles in game g^l of all individuals with preference type θ_j active in g^l. We denote by $q_j^l \in \Delta(\Theta_j)$ the statistical distribution of θ_j within the sub-population j in game g^l. We call q the profile of such statistical distributions, $q^l = (q_j^l)_{j \in N^l}$.

4.2.1 Equilibrium play within games

Preferences over strategy profiles

Consider the generic game g^l. Every player j in game g^l is endowed with a preference type $\theta_j \in \Theta$. Call S_j the set of possible strategies of player j with generic element s_j. [1]We call σ_j a mixed strategy of player j, defined as $\sigma_j : \Theta_j \to \Delta(S_j)$, $\sigma_j(s_j|\theta_j)$ being the probability that agent j chooses strategy (or action) $s_j \in S_j$ when his preference type is θ_j. Denote $\Sigma_j = (\Delta S_j)^{\Theta_j}$ the set of feasible strategies for player j. We call σ the generic strategy profile and Σ the set of such profiles, $\Sigma = \prod_{j \in N^l} \Sigma_j$.

We assume that the preferences of a player over strategy profiles are a distortion of the true (material) payoff. The wedge between a player's objectives and actual payoffs is the *disposition* of the player toward a strategy profile and it depends on his preference type. This disposition term is private information to player j. We assume that the utility of player j with preference type θ_j from strategy profile σ is:

$$u_j(\sigma, \theta_j) = m_j(\sigma) + \beta_j(\sigma, \theta_j)$$

where $m_j(\sigma)$ is the material payoff from strategy profile σ while $\beta_j(\sigma, \theta_j)$ is the disposition of player j from strategy profile σ when his preference type is θ_j,

[1] For static games s_j is a feasible action. For what concerns dynamic games, we treat them in their *normal form* representation. This does not pose any problem as long as we are using BNE as a solution concept.

$\beta_j : \Sigma \times \Theta_j \to \mathbb{R}$. We assume that when $\theta_j = 0$, player j's utility form strategy profile σ coincides with j's material payoff $m_j(\sigma)$, that is for every j, $\beta_j(\sigma, 0) = 0$. That is to say that when $\theta_j = 0$, j is a selfish player.

Bayesian Nash Equilibrium play

Given the profile of statistical distributions q associated to the basic game g^l, we do have a Bayesian game $\Gamma^l(q)$. We assume that aggregate play of the sub-population involved in game g^l corresponds to a Bayesian Nash equilibrium (BNE) of $\Gamma^l(q)$. Each individual upon being selected to play holds a correct conjecture about the distribution of his opponents' play and chooses a mixed strategy that is a best reply to this conjecture given his own preference type. Call $B^l(q)$ the set of all the BNE of the (bayesian) game $\Gamma^l(q)$.

4.2.2 Equilibrium play across games

In the previous section we described the equilibrium play *within* each single game. We now discuss the *overall* equilibrium play in each relevant state of the long-run evolutionary dynamics, which includes the individuals' choices of games. The adoption of Nash equilibrium (or refinements of it) would not fit our purposes. Indeed, it would not be plausible to assume that each individual chooses a *unique* game to play and precisely the game which gives him in expected terms the highest utility among all. It is more realistic to assume that an individual gets into different strategic contexts in his everyday life and chooses to handle *some* of them. Among the strategic interactions he decides to stick with, there might be some which he decides to face more frequently than others. It is quite natural to assume that the higher the utility he expects to get from a game the higher the frequency with which he plays that game with respect to others. The choice of one game over another then becomes dependent of the magnitude of the difference in terms of expected utility. Still inferior games might be played with positive probability. This is a feature that cannot be captured by strong rationality assumptions as the ones involved by the Nash equilibrium notion.

On the contrary, the probabilistic choice approach of the *quantal response equilibrium* introduced by McKelvey and Palfrey (1995) can provide us with the proper framework to describe these aspects. Since the quantal response equilibrium describes the probabilistic choice of strategies in a given game, we need to modify the framework so that it can fit the description of the choice of games given the BNE plays in every single game. In our context, the decision makers are the *individuals* themselves with their preference types and their probability distributions over roles within games.

The quantal response equilibrium in the choices of games

We assume that all individuals have access to the set of all available games $G = \{g^1, ..., g^L\}$. We denote g_{il} the choice of individual i to play game g^l only. We call such a choice *pure life strategy.* We denote λ_i a *mixed life strategy* of individual i, which specifies the probability measures attached by individual i to each of the available games, that is $\lambda_i = (\lambda_i(g^l))_{l=1,...L}$, $\lambda_i \in \Lambda_i$. Denote by λ_{il} the probability that individual i plays game g_l, i.e. $\lambda_{il} = \lambda_i(g^l)$. Call $\Lambda = \prod\limits_{i \in H} \Lambda_i$ the set of mixed life strategy profiles. We denote by $g = (g_1, ..., g_h)$ the profile of pure life strategies of all individuals in the population and by $\lambda = (\lambda_1, ..., \lambda_h)$ the profile of their mixed life strategies.

Given a mixed life strategy profile λ individual i's average utility from λ is

$\pi_i(\lambda) = \sum_g p(g)\pi_i(g)$ where $p(g) = \prod_{i \in H} \lambda_i(g_i)$ is the probability distribution over pure life strategy profiles induced by λ. Denote by $\pi_{il}(\lambda)$ the average utility of individual i of adopting the pure life strategy g_{il} when the other individuals in the population use λ_{-i}, that is $\pi_{il}(\lambda) = \pi_i(g_{il}, \lambda_{-i})$. We assume that for each pure life strategy g_{il} there is an additional privately observed *payoff disturbance* ε_{il} so that the disturbed average utility of individual i from g_{il} becomes:

$$\widehat{\pi}_{il}(\lambda) = \pi_{il}(\lambda) + \mu_i \varepsilon_{il}$$

where μ_i , a strictly positive real number, is the error rate.

We denote by ε_i individual i's profile of random errors, $\varepsilon_i = (\varepsilon_{i1}, ..., \varepsilon_{iL})$ and we assume that this vector is distributed according to a joint distribution with density function $f_i(\varepsilon_i)$. The marginal density of f_i exists for each ε_{il} and errors are unbiased, i.e. $E(\varepsilon_i) = 0$. If these two properties are satisfied for all i, we call $f = (f_1, ..., f_h)$ *admissible.*

We assume that each individual chooses life strategy g_{il} when $\widehat{\pi}_{il}(\lambda) \geq \widehat{\pi}_{ik}(\lambda)$ for all $k = 1, ..., L$. Given this decision rule, the probability that individual i chooses game g_{il} is:

$$P_{il}(\pi_i) = \Pr\left[\tfrac{\pi_{il} - \pi_{ik}}{\mu_i} > \varepsilon_{il} - \varepsilon_{ik}, \text{for all } k \neq l\right]$$

Hence, $\pi = (\pi_1, ..., \pi_h)$ and $f = (f_1, ..., f_h)$ induce a probability distribution over the actual choices of each individual. Let us assume that for every individual the random errors are independent and that all f_i's have an extreme value distribution. These assumptions imply that the probability that individual i selects game g^l given π_i is:

$$P_{il}(\pi_i) = \frac{\exp(\pi_{il}/\mu_i)}{\sum_{k=1}^{L} \exp(\pi_{ik}/\mu_i)}$$

P_{il} which maps π_i into $\Lambda_i$is called *logistic quantal response function.*

Definition 1 *Let f be admissible. A Quantal Response Equilibrium (QRE) is any vector $\lambda^* \in \Lambda$ such that for all $i \in H$ and $l = 1, ..., L$,*

$$\lambda^*_{il} = P_{il}(\pi_i(\lambda^*)).$$

Matching technology

A QRE in the choices of games induces a probability distribution over profiles of pure life strategies (games). Since every individual $i \in N$ is endowed with a preference type, from the probability distribution over profiles of games we can derive for each game g^l a probability distribution of preference types in the sub-population active in g^l. We assume that players (roles) are independently drawn from this distribution. Recall that, together with a preference type, every individual is also endowed with a distribution over the roles of each game. So it might be the case that in game g^l an individual i is drawn to play in a role j in which he does not want or simply he cannot play, i.e. $\rho^l_i(j) = 0$. We call a matching of this sort *incompatible*. Incompatible matchings are not enforced, meaning that players simply do not play the game untill a compatible matching occurs. The probability that a player j has a certain preference type θ_j is the sum of the probabilities of those compatible matches where the individuals drawn to play in the role of that player are endowed with preference type θ_j. So, given μ and a game $g^l \in G$, given the profile of the mixed life strategies of all individuals and given the individual-specific probability distributions over roles within games, we can derive the induced probability measures of preference types in each sub-population involved in a role j of g^l, i.e. $(q^l_j(\theta_j))_{\theta_j \in \Theta}$.

4.2.3 The overall equilibrium configuration

So an equilibrium configuration of the game of life is characterized by a distribution of preference types in the population μ, a QRE λ^* which represents a distribution of preference types over games and together with μ induces a distribution of preference types within games and a set of BNE $B^l(q^*)$ within each game g^l. Hence, we can name such equilibrium configurations by the triple $(\mu, \lambda^*, (B^l(q^*))_{l=1,...L})$.

Definition 2 *In each relevant state of the long-run evolutionary dynamics a triple* $(\mu, \lambda^*, (B^l(q^*))_{l=1,...,L})$ *is an equilibrium configuration if*
i) λ^* *is a quantal response equilibrium and*
ii) $q^* = q(\lambda^*, \mu)$ *induce* $B^l(q^*)$ *such that no individual* i *has incentive to change* λ_i^*.

Fitness: evolution of preferences and stable configurations

We follow the indirect evolutionary approach which adopt fitness as a criterion to evaluate the evolutionary success of a certain preference type. The fitness of a preference type in a given state of the long-run evolutionary dynamics is the average of all *material* payoffs of individuals with that preference type. In our context the average material payoff of a preference type is influenced by:

i) the distribution of preferences in the overall population μ;

ii) the profile of distributions over roles of each individual in each game ρ;

iii) the equilibrium probability measures of pure life strategy profiles λ^* and

iv) the equilibrium play within each game.

We assume that preferences evolve according to a replicator dynamics. As in the standard literature, we assume that an overall configuration is *stable* if all types receive the same average fitness and if the statistical distribution of preferences resists entry by mutants.

4.3 Expected results

We do believe that this project is promising along two main directions. First of all, the innovative approach of the game of life instead of the play of a single game provides a more complete description of the *relative advantages* that different preference types can gain when facing different strategic situations. Some preference types may be evoutionary successful in some games and succumb in others. We expect that individuals will choose games from which they expect an higher average payoffs. Given the equilibirum configuration we have set, these expected payoffs coincide with the true average payoffs. Hence, individuals will choose more frequently games that give them an higher objective fitness.

Secondly, it is plausible to expect that similar preference types will tend to concentrate on the *same subsets* of games. This implies, for example, that altruists will meet more frequently other altruists than selfish players, so that they can reach efficient outcomes and exploit strategic advantages that allow them to survive. It is worth noticing that observability of others' preference types (before playing) will not be necessary for preference types other than selfish to survive in the long-run. It goes without saying that the necessary condition of observability is the major limit in the studies on the evolution of preferences. It is indeed plausible to assume that players receive some kind of signals in certain strategic context, for example when they are facing repeatedly the same subset of individuals (trivially when they know eachother). Though, in many strategic situations players' identities and characteristics are anonymous. Our model will promisingly overcome this difficulty.

We expect to get a stable configuration with an heterogenous population. Plausibly, there will be a large fraction of selfish individuals. At the same time though, there is room for other preference types to survive, because they frequently meet individuals with characteristics similar to their own.

References

[1] **Dekel, Eddie; Ely, Jeffrey C. and Ylankaya, Okan**. "Evolution of Preferences" *Review of Economic Studies*, 2007, *74*(3), 685-704.

[2] **Guth, Werner and Kliemt, Hartmut.** "The Indirect Evolutionary Approach: Bridging the Gap between Rationality and Adaptation", *Rationality and Society*, 1998, *10*(3), 377-399.

[3] **Guth, Werner and Yaari, Menahem**. "Explaining Reciprocal Behavior in a Simple Strategic Game". in *Explaining Process and Change-Approaches to Evolutionary Economics, University of Michigan Press*,1992, 23-24.

[4] **Guth, Werner and Yaari, Menahem.** "An Evolutionary Approach to Explaining Cooperative Behavior by Reciprocal Incentives", *International Journal of Game Theory*, 1995, *24*, 323-344.

[5] **Heifetz, Aviad; Shannon, Chris and Spiegel, Yossi.** "The Dynamic Evolution of Preferences", *Economic Theory*, 2007, *32*(2), 251-286.

[6] **McKelvey, Richard D. and Palfrey, Thomas.** "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior*, 1995, *10*(1), 6-38.