NORTHWESTERN UNIVERSITY

Creativity, Unawareness, and Cooperation in Game Theory

A DISSERTATION

SUBMITED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLEMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Arnau Bages Amat

EVANSTON, ILLINOIS

December 2008

ABSTRACT

Creativity, Unawareness, and Cooperation in Game Theory

Arnau Bages Amat

Game Theory is the branch of applied mathematics that studies the strategic interaction among intelligent agents. So far, standard Game Theory literature has interpreted “intelligent agents” exclusively as “rational agents”. This work points out that this interpretation is an important limitation since intelligence consists of more abilities, some of which may play an important role in strategic interactions. In particular, it introduces the concept of *creativity* (an agent’s capacity of finding viable modifications he was previously unaware of to a given game) and emphasizes its relevance in Game Theoretic settings.

In order to properly model the introduction of new ideas that some players may be unaware of we can not use a standard formalization where the structure of the game is assumed to be common knowledge among the players. Therefore, the first chapter of this dissertation is devoted to develop a formalization for extensive-form games where players might be unaware of relevant facts of the game. This formalization generalizes the standard one allowing players to have different subjective perceptions of the game at different points.

Besides of providing a proper definition of creativity in strategic settings, the second chapter also presents a *bluffing creative method* that exploits an *asymmetric awareness effect* to easily modify the outcome of a given game in the desired direction by using creativity. It is shown that this method can be employed to sustain cooperation in the Finitely Repeated Prisoners' Dilemma.

The third chapter does not deal with creativity or unawareness but departs from the implicit assumption used by the repeated games literature that the probability of repeating the stage game is completely exogenous. It is shown that when this assumption is not satisfied payoff vectors below the minmax might be possible to sustain whereas others above it might not. Moreover, examples and applications are provided where erroneously making this assumption can lead to completely wrong predictions. Furthermore, a new force able to generate cooperative behavior is identified. This force is the fear or desire that the game stops being repeated and, since it does not rely on punishments, it does not require monitoring to operate.

Acknowledgements

I would like to thank my advisors, Jeffrey Ely, Alessandro Pavan, Marciano Siniscalchi, and Michael Whinston as well as Wojciech Olszewski and Tomasz Strzalecki for their valuable time and comments. Fundación Ramón Areces provided me financial support during the first three years of my graduate program and Fundación Caja Madrid during the last two.

Dedicated to my parents and to Marina.

Table of Contents

List of Figures

CHAPTER 1

Extensive Games with Unawareness

1.1 Introduction

The aim of this chapter is to present a game theoretic formalization which allows for the possibility that some players might be unaware of relevant aspects of a given interaction. Existing standard models implicitly assume that the structure of the game is common knowledge among the players and hence, these are assumed to be automatically aware of every relevant aspect of the interaction. Still, there are many instances where we may think that this assumption is unsatisfactory. One of the major ones that ultimately motivated this work is to be able to study the strategic implications of creativity, as will be shown in chapter 2.

Moreover, asymmetric awareness may play a crucial role in strategic interactions. For example, an unaware player can be surprised by a new action, the incursion of a new player or something unexpectedly affecting the final payoff. In standard games, the biggest surprise a rational player might encounter is an opponent playing irrationally but, outside of that, there is no room for surprises such as the ones described above affecting the bare structure of the game.

The strategic value of this asymmetric awareness might be considerable. An aware player can decide to maintain another one unaware in order to keep him undertaking a particular set of actions that he would not take if he was aware of them. He could also decide to eventually surprise the unaware player (for instance, undertaking an action the rival was unaware of) taking advantage that the rival could not foresee it coming. He might even trick another player to change his actions by making a false claim (as will be shown in the second chapter).

Another example of an interaction that exploits the asymmetric awareness is a swindle. Generally, the way a swindle starts is with somebody tricking someone else to take some sort of action he was previously unaware of with the promise that, by doing it, he will obtain important future gains. After performing this action, the naïve individual enters into a new unfamiliar world he has not been capable to fully understand, especially because he is not given the time for analyzing it. The swindler counts on that so the naïve person will be unaware of how the swindle works and he just needs to finish the game as planned from the very beginning. If everything went alright, only then the conned person will understand the game he was really playing. I will later present an example to illustrate a formalization that can be interpreted as a swindle.

In the case of a swindle, it is obvious that unawareness plays a major role in the evolution of the game. The asymmetry between the awareness levels of the players is what makes it feasible. Without unawareness, only a very stupid ("not rational") person would fall for a swindle. However, even a hundred per cent rational individual can be swindled if he is not careful enough and only because he is not aware of everything in this world.

This is not the only case where lack of unawareness might lead to the same consequences as bounded rationality. In fact, there is an increasing literature in bounded rationality and some of it deals with different sorts of "myopic agents" that are unable to foresee some relevant features of the game. This myopia could be caused not by a problem with the rational capacity of the agent but by the fact that he is unaware of something significant. If this is the true cause of the myopia, then the right way to model these situations is using a game with awareness like the one presented in this chapter. Applying this model and an appropriate solution concept, you can derive the maximizing behavior of the unaware players instead of having to assume it by using a rule-of-thumb, as it is usually done in the current literature.

1.2 Constructing a Model for a Game with Awareness

Since the formalization of a game that I intend to build is a generalization of the standard extensive-form one, it is convenient to start by recalling the formal definition of this standard extensive-form representation. The following definition is based on the one provided by Osborne and Rubinstein(1994) but making the set of actions explicit as in Halpern and Rego (2006).

Definition: A (finite) *extensive-form game* is a tuple ($N, M, H, P, f_c, \{\Psi_i : i \in N\}$, $\{u_i : i \in N\}$), where

- N is the finite set of players.
- M is the finite set of actions (or movements) available to players (and nature) in the game.

- H is a finite set of finite sequences of actions that is closed under prefixes, meaning that if $h \in H$ and *h'* is a prefix of *h*, then $h' \in H$. In words, every element of *H* is a history of actions. The nodes in a game tree can be identified with the histories in *H* where each node is characterized by the sequence of actions required to reach it. We will denote by *Z* the set of runs (i.e. histories that are not a prefix of any other history). Let $M_h = \{m \in M : h \cdot \langle m \rangle \in H\}$ (where · denotes concatenation of sequences); M_h is the set of actions available after history *h*.
- $P : (H - Z) \rightarrow N \bigcup \{c\}$ is a function associating a player or nature (denoted by *c*) to each non-terminal history. $P(h) = i$ means that player *i* takes an action after history *h* and $P(h) = c$ that nature makes a movement. Let $H_i = \{h : P(h) = i\}$ denote the set of all histories after which player *i* takes an action.
- f_c is a function that assigns a probability measure f_c (· | h) on M_h to every history for which $P(h) = c$. Intuitively, f_c (· | h) gives the probability of each of nature's movements at history *h*.
- Ψ_i is a partition of H_i with the property that if *h* and *h'* are in the same cell of the partition then $M_h = M_{h'}$. This implies that the same set of actions is available for every history belonging to the same cell of the partition. In words, if two histories are in the same cell of Ψ_i, then they are identical

from i's point of view and i must consider both possible whenever one of them is reached. A cell $\psi \in \Psi_i$ will be called an i-information set.

- $u_i : Z \rightarrow \mathrm{R}$ is a function assigning player i a payoff to each run of the game.

An assumption that I will also make throughout the chapter is that players will have *perfect recall*; i.e. they do not forget the actions they have previously taken or the information they have hold. This assumption can be formally stated in the following way:

A1. If h and h' are in the same i-information set and h_1 is a prefix of h such that $P(h_1) = i$, then there exists a prefix h_1' of h' such that h_1 and h_1' are in the same information set; additionally, if $h_1 \cdot \langle m \rangle$ is a prefix of h (meaning that m was the action taken when h_1 was reached in h) then $h_1' \cdot \langle m \rangle$ is a prefix of h'.

To allow for the possibility that players might at some point be unaware of some aspects of the game, the generalized formalization of the model must admit the option that a player does not know the whole structure of the game. Therefore, this formalization must allow for subjective visions or perceptions of the game. These subjective visions that a player holds at a particular moment might evolve once the player becomes aware of a new fact that proves that the previous perception was wrong or incomplete. Therefore, to capture the possibility that players might be unaware, we need a formalization that also indicates the perception of the game hold by the player who has to move at that particular time[1].

We want to find a formalization to properly model a strategic interaction where players might be unaware of some relevant facts. This formalization will be called a *game with awareness*[2] and I will assume that each player, at each point of the interaction views it as a game (possibly with awareness). By definition of unawareness, a player does not know he is unaware of a particular fact. Thus, since a game is supposed to capture all the relevant information the player holds, he must consider his vision of the game as the true one even when it is not. Therefore, to formalize a game with awareness we will require a set of games with awareness capturing the different subjective visions of the game hold at particular times by different players. In addition to that, we will also require a set of functions linking the players who move with their subjective visions of the game (which might be another game).

An implicit assumption of the model will be that players are always *certain* about the awareness level of their opponents at every point in the game. Here, by *certain* I do not mean that they know the true awareness level of the opponents, but that, from their subjective point of view, they are confident about it. For instance, if some player is aware of something that some other player is not aware, this second player can not know the awareness level of the first one because then she would not be unaware of it. Moreover, it will be assumed that this lack of uncertainty is common knowledge among players. Even

[1] Of course, we will also want to impose some restrictions relating the subjective vision with the real game being played.
[2] Sometimes I will also indistinctly refer to it as a "game with unawareness".

with this assumption, the model is flexible enough to formalize situations where a player is not sure about the awareness level of another player. In order to do that, we only need to use a trick similar to the one introduced by Harsanyi (1967-68) of reinterpreting a game with incomplete information as a game with imperfect information. In this approach, one imagines that the players' awareness levels are determined by nature with an ex ante probability distribution assumed to be common knowledge among the players within that awareness level.

I will start by naming *awareness level* this notion of "subjective vision" or "perception" of the interaction. Formally, an awareness level is the collection of all the relevant facts regarding to a particular interaction that some player is aware of at a certain time. Using the formalization presented in this work, it should be possible to represent an awareness level as a game with awareness. In addition to that, since the awareness level of a player might change along the course of it, sometimes it will be more convenient to talk about *identities* (a decision maker at a decision point in a dynamic game) than about players. Hence, a player can have different identities with potential different awareness levels, but each identity will have a single awareness level associated to it.

In any game with awareness there may be some identities that do not consider that some identities might be unaware of any aspects of the game. For this reason, it will be possible and convenient to represent the awareness level of these identities using a standard game without awareness. Then, there might also be some identities that think that other identities are unaware of some aspects of the game. Again, there might also be some identities that think that other identities think that other identities are unaware of some aspects of the game and so on. It will be practical to classify awareness levels according to their relative awareness degree measured by the number of these iterations.

Define a *First Order Awareness Level* as an awareness level which does not consider that any identity can be unaware of any aspect of the game. Of course, this does not preclude the existence of unaware players in the game since this same identity could actually be unaware of something without even suspecting it. As mentioned above, the game from this subjective point of view can be represented as a standard game without unawareness. $\Gamma^{r^1} = (N^{r^1}, M^{r^1}, H^{r^1}, P^{r^1}, f_c^{r^1}, \{\Psi_i^{r^1} : i \in N^{r^1}\}, \{u_i^{r^1} : i \in N^{r^1}\})$ will denote an extensive form representation of the game according to the first order awareness level r^1. Let G^{1m} be the collection of the extensive form representations of all the first order awareness levels that are possible or considered possible by an identity in a given game.

Similarly, we will define a *Second Order Awareness Level* as an awareness level that knows that there exists some identity that is unaware of some feature of the game whose subjective perception can be represented as a game belonging to G^{1m}. Moreover, we will require that a second order awareness level cannot be aware that some identity is at the same time unaware of something and aware that some other identity is unaware of something else[3]. Generally, for $k > 1$, define a *k-order awareness level* as an awareness level that considers that i) there is some identity that is unaware of some fact that has a

[3] We require this condition in order to avoid that a higher order of awareness could also be considered a second order of awareness.

(k-1)-*order awareness level* ii) there is no identity unaware of some fact with a *k-order awareness level*[4].

We will call Γ^{r^k} an extensive-form representation of a game according to a k-order awareness level r^k and G^{km} will be the collection of the extensive form representations of all k-order awareness levels that are possible or considered possible by an identity in a given game. Throughout this analysis we will assume that the set of awareness levels that are possible or considered possible by some identity is finite.

When $k > 1$, Γ^{r^k} can no longer be represented as a standard extensive game without awareness. To represent this game in extensive-form we will first require some extra notation. For $q < k$, we will call $G^{qr^k} \subseteq G^{qm}$ the set of the extensive form representations of all the q-order awareness levels that a k-order awareness level r^k considers possible. It is important to note that each awareness level should have one representation and only one. Additionally, we also require a function F^{r^k} that maps a history h in Γ^{r^k} such that $P(h) = j$ to a pair (Γ^h, ψ), where $\Gamma^h \in G^{1r^k} \bigcup ... \bigcup G^{(k-1)r^k} \bigcup \Gamma^{r^k}$ and ψ is a j-information set in game Γ^h. The intuition is that when the awareness level of player i is such that it can be represented through Γ^{r^k}, function F^{r^k} tells you at each history h the game that player i thinks the identity j who moves at h considers to be the true game being played and ψ consists of the set of histories in Γ^h that j considers possible at that particular time. When writing down the proper definition of a game with

[4] Notice that in every order of awareness there might be several awareness levels, each of them not knowing the existence of the others. So, an awareness level can think that an identity with an awareness level of his same order is in fact of lower order.

awareness, we will also impose that whenever player i considers that identity j is fully aware he will assign game Γ^{r^k} and the information set where h belongs to this identity.

Therefore, to define a game with unawareness we first need to define a set of games representing the subjective visions of identities that are unaware of different features of the game being played. Since these identities might at the same time consider that some other identities might be unaware of some fact of which they are aware, this iterative process continues until we reach some identities with a first order awareness level. This allows us to sequentially define a game from lower to higher awareness levels.

A game with awareness can either be the representation of the subjective view of the game of an identity or the game from the perspective of an outside modeler that knows absolutely every aspect of the game, including the awareness levels of the identities. When Γ^{r^k} represents the subjective vision of a player at some point of the game, the set $G^{r^k} \equiv G^{1r^k} \cup ... \cup G^{(k-1)r^k}$ is the set of all representations of the game corresponding to the lower awareness levels that this player considers that some identities might have. On the other hand, the whole structure of the game actually played from the point of view of an omniscient modeler can also be represented by this framework. Since this is the real game and all the remaining subjective visions must be somehow related to it, it will be useful to name it as Γ^m. However, it is worth remarking that each identity considers his particular vision of the game as being the true one until he becomes aware of some new fact. This explains why a subjective vision of the game must be described as if it was a true game (or a modeler's game).

Given a modeler game Γ^m, all the subjective visions of the game that an identity might hold must satisfy some consistency conditions relating it to "reality". In particular, there are two sets of requirements that the games in G^m must satisfy. The first ones refer to the desired property that players remember what they did, what they knew and what they were aware of. These conditions generalize perfect recall to games with unawareness. The second set of conditions concern the fact that, at a given point, a player can only undertake actions that actually exist and he is aware of.

Definition: A game with unawareness Γ^{r^k} has *generalized perfect recall* if the following two conditions hold for any game $\Gamma^q \in G^{r^k} \bigcup \Gamma^{r^k}$:

B1. If *h* and *h'* are in the same *i*-information set in Ψ^q and h_1 is a prefix of *h* in Γ^q such that $P^q(h_1) = i$, then there is a prefix h_1' of *h'* such that h_1 and h_1' are in the same information set in Ψ^q or $F^q(h_1) = F^q(h_1')$; additionally, if $h_1 \cdot \langle m \rangle$ is a prefix of *h* (so that *m* was the action performed when h_1 was reached in *h*) then $h_1' \cdot \langle m \rangle$ is a prefix of *h'*.

B2. If in game Γ^q history h_1 is a prefix of *h* such that $P^q(h) = P^q(h_1) = i$ and $F^q(h) = (\Gamma^z, \psi)$, then there exist histories h_1' and $h' \in \psi$ in Γ^z such that h_1' is a prefix of h', $P^z(h') = P^z(h_1') = i$ and $M^z_{h_1'} \subseteq M^q_{h_1}$; moreover, $F^q(h_1) = F^z(h_1')$.

Condition B1 is the generalization of the perfect recall condition for games without unawareness to extend it to all possible identities in all possible games. The difference is that now it can also be the case that the awareness level of a player changes from h_1 to h and we must take this into account. To interpret condition B2 think on the modeler's game (i.e. $q = m$). In this case, B2 says that the modeler assumes that a player does not forget his previous actions or the awareness level he held at that time (which in turn implies that he does not forget what he was aware before). Nevertheless, since each unaware identity also considers his vision of the game as being the modeler's one, he will also impose this condition and, so, it must hold for all awareness levels.

There are a few additional restrictions that we must impose to a game with awareness Γ^{r^k} that must hold for any game $\Gamma^q \in G^{r^k} \bigcup \Gamma^{r^k}$ in order to be properly defined and capture certain features of unawareness:

C1. Let $F^q(h) = (\Gamma^h, \psi^h)$, for any $h' \in \psi^h$, if $h' \cdot \langle m \rangle \in \Gamma^h$, then $h \cdot \langle m \rangle \in \Gamma^q$.

C2. Let $F^q(h) = (\Gamma^h, \psi^h)$, for any $h' \in \psi^h$, if $h' \cdot \langle m \rangle \notin \Gamma^h$, then $h \cdot \langle m \rangle \notin \Gamma^q$.

C3. If h and h' belong to the same information set in Ψ^q, then $F^q(h) = F^q(h')$.

C1 is a fundamental condition. Again, it must be interpreted from the point of view of a modeler ($q = m$) and it reflects the fact that a player can only take actions that do actually exist (i.e. they also belong to the modeler's game)[5]. This condition establishes the link between the real game being played and the subjective vision of a player in a

[5] Notice that this assumption does not prevent to model a situation where a player *thinks* he can do action z but, upon trying it, realizes that he actually cannot: you can still model it using two different nodes, the action "try z" and an additional awareness level.

particular moment. Besides, since each identity views his game as the true one and knows that other identities also view their game as the true one, we want to impose it for every particular awareness level. Yet, we must remark that the awareness level q can be unaware of something resulting in the player moving at h (whose awareness level was erroneously assigned by q) doing some movement not specified in Γ^q.

Condition C2 means that if a higher awareness level knows that an identity is unaware of a particular action at the time of taking a decision, he also knows that this identity cannot take this particular action. This is an implication of the fact that a player cannot take an action if he is not aware of it. Finally, C3 captures the fact that if according to an awareness level two histories belong to the same information set, it must be the case that in both histories this awareness level assigns the same awareness level and information set to the corresponding identities.

Now we are at last ready to define a game with unawareness:

Definition: A *game with awareness* Γ^{r^k} representing an interaction from the point of view of an awareness level r^k of order k, is a tuple $(N^{r^k}, M^{r^k}, H^{r^k}, P^{r^k}, f_c^{r^k},$ $\{\Psi_i^{r^k} : i \in N^{r^k}\}, \{u_i^{r^k} : i \in N^{r^k}\}, G^{r^k}, F^{r^k})$ where,

- $N^{r^k}, M^{r^k}, H^{r^k}, P^{r^k}, f_c^{r^k}, \Psi_i^{r^k}$ and $u_i^{r^k}$ are defined as in a standard game without unawareness.

- G^{r^k} is a set of games with awareness of orders lower than *k* containing a unique game for each possible awareness level.
- F^{r^k} is a function that maps a history *h* in Γ^{r^k} such that $P^{r^k}(h) = i$ to a pair (Γ^h, ψ), where $\Gamma^h \in G^{r^k}$ and ψ is an *i*-information set in game Γ^h or $\Gamma^h = \Gamma^{r^k}$ and ψ is the *i*-information set in game Γ^{r^k} where *h* belongs.
- Conditions B1, B2, C1, C2 and C3 hold for any game $\Gamma^q \in G^{r^k} \bigcup \Gamma^{r^k}$.

Example

I will illustrate the representation of a game with awareness with a simple example that will help to clarify the concepts introduced in the previous section.

Consider the standard game without awareness represented in Figure 1.1. The only Subgame Perfect Nash Equilibrium (SPNE) of this game involves player *B* choosing action *c* whenever history $\langle b \rangle$ is reached and player *A*, forecasting this behavior, choosing action *a*. Now, assume that player *A* is not aware of the existence of action *c*. If this is the case, he will be unable to predict that player *B* will play *c* and, instead, he will assume that he will play *d*. Based on this wrong forecast, now player *A* might choose to play action *b* expecting a payoff of 1.

To model this interaction where one of the players is unaware of the existence of a particular action we can use the formalization proposed in the previous section. There are two possible awareness levels: A^1, the first order awareness level of the unaware player *A* and B^2, the second order awareness level of the fully aware player *B*. Besides, since this player is fully aware of everything, in this case this would also be the awareness level of the omniscient modeler.

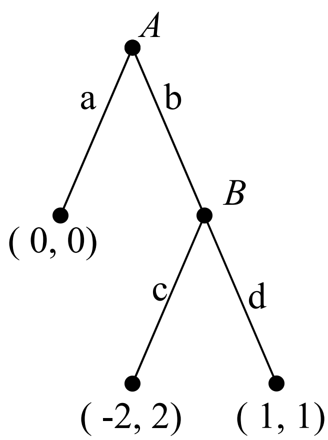

Figure 1.1: Game Γ^{B^2}

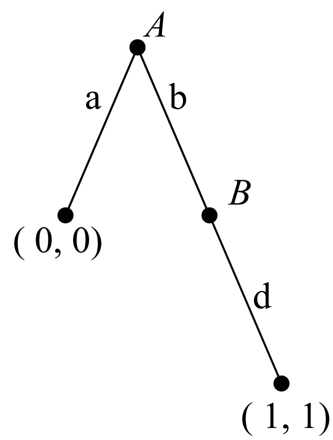

Figure 1.2: Game Γ^{A^1}

The extensive-form representation of the game according to A^1 is depicted in Figure 1.2. Note that, since A^1 is a first order awareness level, Γ^{A^1} already captures all the relevant facts of the interaction. Nonetheless, to represent the game from the point of view of B^2, the figure depicted in Γ^{B^2} is not sufficient. Γ^{B^2} incorporates all the information regarding $N^{B^2}, M^{B^2}, H^{B^2}, P^{B^2}, f_c^{B^2}, \Psi_i^{B^2}$ and $u_i^{B^2}$ but, in order to capture the fact that player A is unaware of action c, we need the set $G^{B^2} = \{\Gamma^{A^1}\}$ and the function F^{B^2} defined as follows[6]:

$$F^{B^2}(\langle\,\rangle) = (\Gamma^{A^1}, \langle\,\rangle)$$

$$F^{B^2}(\langle b\rangle) = (\Gamma^{B^2}, \langle b\rangle)$$

[6] Sometimes, I will use loose notation by referring by Γ^{r^k} to both, the full description of the game with awareness including G^{r^k} and F^{r^k} and the part of it that does not include them and can be depicted as a standard game.

This example can be thought as a simple swindle: player B, the swindler, has set this con to player A, who is told that he can get a payoff of 1 by doing a certain mutually beneficial action. Nevertheless, player A is unaware of the fact that, by doing this action, the swindler will be able to take action c and steal from him.

Now, suppose that in reality player B is unaware that player A is in fact aware of everything and setting a trap for him. To model this new situation the game from the point of view of player B obviously does not change and, consequently, can still be represented by Γ^{B^2}.Nonetheless, we now need to add a third order of awareness in the game to capture the awareness level A^3 of player A and the modeler. The representation according to A^3 is the one depicted in Figure 1.3 plus the set $G^{A^3} = \{\Gamma^{A^1}, \Gamma^{B^2}\}$ and the function F^{A^3} defined as:

$$F^{A^3}(\langle\,\rangle) = (\Gamma^{A^3}, \langle\,\rangle)$$

$$F^{A^3}(\langle b\rangle) = (\Gamma^{B^2}, \langle b\rangle)$$

$$F^{A^3}(\langle b, c\rangle) = (\Gamma^{A^3}, \langle b, c\rangle)$$

It is important to realize that, in order to model this new situation, we still need the representation Γ^{A^1} even if no player actually holds this awareness level. This is because player B still believes that player A holds it.♦

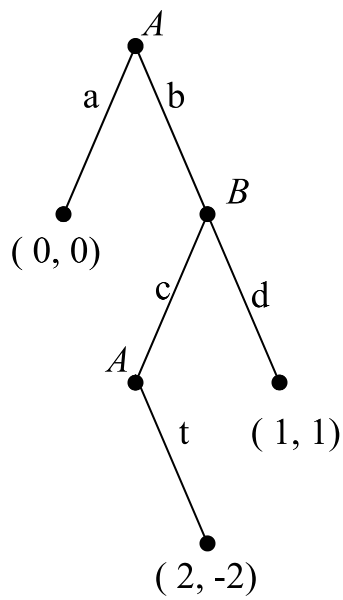

Figure 1.3: Game Γ^{A^3}

1.3 Partial Strategies

In a similar way that we need a set of games to define a game with unawareness, we also require a set of "partial strategies" to define a full strategy in a game with unawareness. However, we do not require considering a partial strategy for each game in G^m. The reason for that is because the strategy of an unaware player will be determined in the game he thinks he is playing. On the other hand, sometimes it will be convenient to take into account more than one partial strategy for a single player with a single awareness level. This is so because an identity of a player who has just become aware of new facts has absolutely no control over the strategy of another identity of the same player who would have became aware of the same facts after a different history[7]. Therefore, instead of considering players, our analysis can be simplified by studying the behavior of *relevant* identities.

Definition: A *relevant identity* for player i in game with awareness Γ^k is an identity of player i in Γ^k such that:

i) $P^k(h) = i$ and $F^k(h) = (\Gamma^k, \psi)$, where $h \in \psi$.

ii) If h_1 is a prefix of h such that $P^k(h_1) = i$, it must be the case that $F^k(h_1) = (\Gamma^q, \psi')$, where $\Gamma^q \in G^k$ and ψ' is an i-information set in Γ^q.

In words, a relevant identity is the first identity of a player in a given game that knows that he is actually playing that game and not some other with a lower awareness level. A relevant identity is an identity that chooses the strategy that the player will follow thereafter until the end of the game or the next awareness change. Under generalized perfect recall, there is no reason why they should otherwise reconsider their strategy later in the game. Since some of the relevant identities will find themselves in a situation that cannot be considered the beginning of a subgame, the following definition will be handy:

[7] This situation is similar to a Bayesian game where it is the case that it is equivalent to consider a set of independent types each maximizing on his own or a single individual maximizing for all possible types before knowing the actual realization of it.

Definition: A *partial subgame* is a subset of the game that begins at a single information set and contains all successor decision nodes, and only these nodes.

The difference with a subgame is that a partial subgame does not need to begin in a single node and it might break some information sets. At this point, it is interesting to observe that, under generalized perfect recall, a partial subgame will not break any information set of the player who makes the first move in the partial subgame.

We are concerned about the partial subgames whose first identity is a relevant one. Formally, in game Γ^m, for each player *i*, call T_i the set of partial subgames belonging to some game $\Gamma^k \in G^m \bigcup \Gamma^m$ starting at an information set ψ such that (i) $P^k(h) = i$ and $F^k(h) = (\Gamma^k, \psi)$ for every $h \in \psi$ and (ii) if h_1 is a prefix of *h* such that $P^k(h_1) = i$ it must be the case that $F^k(h_1) = (\Gamma^q, \psi')$, where $\Gamma^q \in G^k$ and ψ' is an *i*-information set in Γ^q. Call j_{i,t_i} a relevant identity, i.e. the first identity of player *i* in an element $t_i \in T_i$.

Definition: Let $\Psi(j_{i,t_i})$ denote the collection of relevant identity j_{i,t_i}'s information sets and $M_{\psi(j_{i,t_i})}$ the set of actions possible at information set ψ in partial subgame t_i. A *partial strategy* $\sigma_{j_{i,t_i}}$ for identity j_{i,t_i} is a function mapping information sets in $\Psi(j_{i,t_i})$ to a probability distribution over $M_{\psi(j_{i,t_i})}$.

A *complete strategy* σ_i for player *i* is the collection of the partial strategies of all his relevant identities. I.e.: $\sigma_i = \{\sigma_{j_{i,t_i}} : t_i \in T_i\}$

1.4 Solution Concept

Finding the right solution concept for a game with unawareness is not an easy task. Standard solution concepts used in games without unawareness need, at least, to be generalized in order to make sense. Moreover, which generalization would be the appropriate one is also context dependent, like it is for their standard versions.

A solution concept for the games we are studying must capture the fact that an identity might be unaware of some aspects of the game and, so, her best response cannot consider the strategy of a player with a higher awareness level. This, in turn, implies that a player might be "surprised" by an opponent with higher awareness playing optimally but making some apparently stupid decisions from the point of view of an unaware player. When this happens, the unaware player may find himself playing outside what he thought the equilibrium path was. Therefore, we might want to apply a solution concept that imposes some conditions outside the equilibrium path. This, for instance, implies that a generalization of Nash equilibrium might be inappropriate when we are dealing with games with unawareness and that, on the other hand, a concept generalizing subgame perfect Nash equilibrium (SPNE) would be more desirable.

Here, I will develop a general procedure that, when a solution concept is chosen, allows picking up the partial strategies consistent with the rationale behind the solution concept subject to the awareness limitations faced by the players[8]. Once we have this set of partial strategies, we can use them from the perspective of an outside omniscient modeler to predict the "reasonable" outcomes of the entire game.

Loosely speaking, a solution concept is a criterion we use to decide which strategies are *likely* to be played by *rational* players[9]. In other words, it basically allows us to rule out those strategies that, for some reason, we do not think the players will select. When considering games with unawareness, we have to take into account that players make their decisions subject to their current awareness level. Moreover, they also have to consider the awareness level of their opponents in order to predict their actions and optimally respond to those.

Although this complicates even further once we take into account higher order considerations mixed with different awareness levels, the truth is that the structure we have applied to construct the game allows us to deal with these considerations in a treatable way. The notion of orders of awareness lets us to fully hierarchize the awareness levels of the players. Since identities with a first order awareness level see the game as a standard game without unawareness, the usual solution concept that we want to generalize would immediately tell you which strategies this identity is *likely* to use. Once we have the set of all *consistent* strategies for all the identities with a first order awareness level, we can proceed to find the *likely* strategies for the ones of second order.

---

[8] In fact, nothing prevents us to use different solution concepts for different awareness levels. It might be the case that we find that rationalizability best captures the expected behavior of some unaware identities whereas for the aware ones a generalized SPNE might be more appropriate. If we have a reasonable justification for the game at hand, this sort of mixed solution concepts is perfectly acceptable in this model.

[9] The words "likely" and "rational" are emphasized because its meaning is highly dependent on the chosen solution concept.

To do that, we must notice that these identities will take as given the set of strategies used by identities with a lower order of awareness that they are aware of. Then, we must proceed by dismissing the strategies of the players with the higher order of awareness as the rationality of the solution concept would suggest. This procedure can be iteratively applied to higher orders of awareness until we reach the highest one, which is the modeler's. Then, the set of *likely* strategies will be fully identified for all identities.

It will be particularly interesting to apply the above mentioned procedure to generalize an extensive-form solution concept widely used such as Subgame Perfect Nash Equilibrium. This is what I will proceed to do in the following subsection.

1.4.1 Generalized Subgame Perfect Nash Equilibrium

Generalizing the concept of subgame perfect Nash equilibrium (SPNE) for games with possible unaware players is not conceptually complex. Nonetheless, its proper formalization requires an intricate algorithm and cumbersome notation. Hence, I will start by explaining the intuition that hopefully will help understanding the formalization.

Since first order awareness levels consider that they play a standard game without unawareness, we can apply the usual definitions of Nash equilibrium and SPNE to identify the set of Generalized Subgame Perfect Nash Equilibrium (GSPNE) equilibrium strategies for these identities with a first order awareness level. When these GSPNE are unique in each awareness level, an identity with a second order awareness level (who expects an unaware identity to play according to what this identity considers an equilibrium strategy) will know exactly how unaware identities behave. Considering this as given, he can compute his best response to these strategies and to the strategies of the other fully aware identities. This allows us to easily define a generalized Nash Equilibrium and GSPNE for this second order awareness level.

When there are multiple GSPNE in a first order awareness level, an identity of higher order can expect the unaware identities to play any of those equilibria. The question is which one will they pick. As we are talking about equilibria, we will just assume that, in the same fashion that players know their opponents' strategies and coordinate in a single equilibrium, they also agree on the equilibrium played by awareness levels of lower order and coordinate accordingly. If this assumption is too strong will depend on the particular application at hand. When there are not many orders of awareness and there are few cases of multiplicity of equilibrium, it seems a reasonable solution concept.

To deal with the multiplicity of GSPNE, we need to use *profiles of equilibria* for a specific order of awareness k. Each one of these profiles contains a particular GSPNE for every awareness level of order k or lower with the peculiarity that all of them have been derived using the strategy profiles specified in the ones with lower orders. Therefore, a profile of equilibria of order k contains a set of equilibria for awareness levels of this order and lower, each of them consistent with the same set of strategies for identities with a lower order of awareness.

Since we have already been able to compute the first order GSPNE, we are able to organize them in profiles of equilibria. Now, we only need a procedure that, given a profile of equilibria of order k–1, provides us the set of GSPNE for all the awareness levels of order *k*. With those, we can generate all the *k*-order equilibrium profiles and iterate the process until we reach the order of awareness of the omniscient modeler. The set of GSPNE (together with the equilibrium profiles that generated them) of the modeler's game can be considered the solution of the entire game and, depending on the context and the number of GSPNE, it will offer us a more or less accurate prediction of the outcome of the interaction.

Given an awareness level r^k whose vision of the game can be represented as game Γ^{r^k}, call $J^{qr^k}\ (q \leq k)$ the set of all relevant identities with an awareness level of order *q* that this awareness level knows and let $J^{r^k} = \bigcup_{q=1}^{k} J^{qr^k}$. Define a *generalized strategy profile* of Γ^{r^k} to be a set of partial strategies $\vec{\sigma} = \{\sigma_{j_{i,t_i}} : j_{i,t_i} \in J^{r^k}\}$.

Call $EU_{j,\Gamma^k}(\vec{\sigma})$ [10] the expected payoff of identity j_{i,t_i} in game Γ^k given that he plays his partial strategy $\vec{\sigma}_j$ and the remaining relevant identities use their partial strategies $\vec{\sigma}_{-j}$, where $\vec{\sigma}_{-j}$ denotes the set of all partial strategies in $\vec{\sigma}$ but $\vec{\sigma}_j$.

To apply the above described iterative process, we must start by the identities with first order awareness levels. For these awareness levels, since the game can be expressed as a regular game without awareness, the standard definition of Nash equilibrium and Subgame Perfect Nash equilibrium apply without modifications. Also,

[10] To simplify notation, I will omit the subscripts of j_{i,t_i} whenever no confusion can arise.

observe that, in a first order awareness level r^1, the set of players N^{r^1} coincides with the set of relevant identities J^{r^1}.

Definition: A generalized strategy profile $\vec{\sigma}^*$ is a *generalized Nash equilibrium* of game $\Gamma^{r^1} = (N^{r^1}, M^{r^1}, H^{r^1}, P^{r^1}, f_c^{r^1}, \{\Psi^{r^1}{}_i : i \in N^{r^1}\}, \{u^{r^1}{}_i : i \in N^{r^1}\})$ with a first order awareness level r^1 if for every relevant identity $j \in J^{r^1}$ and local strategy σ_j for j in Γ^{r^1},

$$EU_{j,\Gamma^{r^1}}(\vec{\sigma}^*) \geq EU_{j,\Gamma^{r^1}}(\sigma_{j,}, \vec{\sigma}_{-j}^*)$$

Because we are dealing with extensive-form games and the reasons previously explained, it is more convenient to work with a solution concept imposing constrains outside the "equilibrium path". So, we can generalize the notion of subgame perfection:

Definition: A generalized strategy profile $\vec{\sigma}^*$ is a *generalized subgame perfect Nash equilibrium (GSPNE)* of game Γ^{r^1} with a first order awareness level r^1 if it induces a generalized Nash equilibrium in every subgame of Γ^{r^1}.

Next, we need to find the set of GSPNE for games with higher orders of awareness. To do that, we start by finding all the GSPNE for each of the games with a first order awareness level using the above definitions and describe a process that will allow us to build the equilibrium profiles from lower to higher awareness levels.

The notation might become slightly cumbersome: we have to index each GSPNE by the order, the awareness level within the order and the number of the equilibrium within the awareness level. In addition to that, since an equilibrium will be a profile of strategies for the full aware players plus a set of profiles of equilibria (i.e. strategies) of lower order awareness levels, it will be convenient to introduce this information when indexing an equilibrium. A profile of order k GSPNE will be denoted by EP_v^k where v is a sequence of k elements whose n^{th} element is the n^{th} order equilibrium profile index. EP_v^k contains one GSPNE for each awareness level of order k and lower[11]. Finally, we will denote by $E_{v,r,n}$ the GSPNE n of the awareness level r of order k based on the equilibrium profile $EP_v^{k\text{-}1}$.

Given an order of awareness $k > 1$, take a profile of lower order equilibria $EP_v^{k\text{-}1}$ and consider an awareness level r^k of order k. The identities who share this awareness level will take the strategies of the identities with a lower order awareness level as given and then look for their own strategies which are a best response to those of the unaware identities and the remaining fully aware identities.

Definition: Given $EP_v^{k\text{-}1}$, a generalized strategy profile $\vec{\sigma}^*$ is a *generalized Nash equilibrium* of game $\Gamma^{r^k} = (N^{r^k}, M^{r^k}, H^{r^k}, P^{r^k}, f_c^{r^k}, \{\Psi_i^{r^k} : i \in N^{r^k}\}, \{u_i^{r^k} : i \in N^{r^k}\}, G^{r^k}, F^{r^k})$ with a k-order awareness level r^k if

[11] Notice that, since there might exist a continuum of equilibria, there might be infinitely many equilibrium profiles. To deal with this, it will be convenient to consider sets of similar equilibria in building the equilibrium profiles so you can keep its number finite.

- for every fully aware relevant identity $j \in J^{kr^k}$ and local strategy σ_j for j in Γ^{r^k},

$$EU_{j,\Gamma^{r^k}}(\vec{\sigma}^*) \geq EU_{j,\Gamma^{r^k}}(\sigma_{j,}, \vec{\sigma}_{-j}^*)$$

- all other relevant identities $j' \in J^{qr^k}$ for $q < k$ play according to $EP_v^{k\text{-}1}$.[12]

Definition: A generalized strategy profile $\vec{\sigma}^*$ together with the lower order equilibrium profile $EP_v^{k\text{-}1}$ is a *generalized subgame perfect Nash equilibrium (GSPNE)* of game Γ^{r^k} with a k-order awareness level r^k if they induce a generalized Nash equilibrium in every subgame of Γ^{r^k}.

Once we found all the GSPNE based on $EP_v^{k\text{-}1}$ for this awareness level r^k, we can assign an index n to each of them and so they will be identified by $E_{v,r,n}$. Then we can repeat this exercise for all the remaining awareness levels of order k (keeping the profile $EP_v^{k\text{-}1}$fixed). As soon as we have identified all the GSPNE for all the awareness levels of this order, we can construct the k-order equilibrium profiles based on $EP_v^{k\text{-}1}$. We will call them $EP_{v^\wedge}^{k}$ where the first k – 1 elements of $v^\wedge$ will be identical to those of v and the k^{th} one indexes each one of the possible k-order equilibria profiles.

Now, we can repeat all the above described procedure for another profile $EP_w^{k\text{-}1}$, find for each awareness level r^k all the GSPNE consistent with this profile of lower

[12] Notice that $EP_v^{k\text{-}1}$ contains strategies of identities not belonging to r^k. Therefore, the identities with this awareness level are unaware of those other identities and their equilibrium strategies. However, notice that these will not be needed in order to compute $EU_{j,\Gamma^{r^k}}(\sigma)$.

awareness equilibria and construct $EP_{w^\wedge}{}^k$. Doing this for all the equilibrium profiles of order $k - 1$ allows us to obtain all the possible GSPNE for all the k-order awareness levels and compute all the k-order equilibrium profiles $EP_v{}^k$.

Starting by a first order awareness level, we can iterate this process for higher orders until we solve the GSPNE for the highest order (i.e. the modeler's game). Then, we can say that we have found all the "equilibria" of this game with unaware players. Although the tedious notation in this section makes this procedure seem a bit complicated, in practice, when the number of possible awareness levels is reduced, it becomes straightforward to apply, as the following example illustrates.

Example (Cont'd)

Again, consider the game Γ^{A^3}. In this game there are two relevant identities for player B: $\langle b \rangle$ in game Γ^{B^2} (that we will call B^2) and in game Γ^{A^1} (that we will call B^1). Player A has also two relevant identities: the initial identities in both games Γ^{A^1} and Γ^{A^3} that we will call A^1 and A^3 respectively. To find all the GSPNE of the game we have to start by the identities with a first order awareness level, in this case identities A^1 and B^1. These identities believe that the game they are playing is Γ^{A^1}, which is a standard game without unawareness whose unique SPNE has identity B^1 playing d and identity A^1 playing b. Since there is only one first order awareness level (A^1) which has a unique SPNE, we will have a single first order equilibrium profile $EP_1{}^1$ with a single element on it: $E_{A^1,1}$.

Next, we can use this first order equilibrium profile to compute the second order equilibrium profiles. The only second order awareness level in the game is B^2, which can be represented by Γ^{B^2}. Since according to EP_1^1 we know that player *A* will choose action *b* in game Γ^{B^2}, the only GSPNE of the game has relevant identity B^2 playing *c*. Again, we call this unique GSPNE $E_{1,B^2,1}$ and it is the unique element of the unique second order equilibrium profile $EP_{1,1}^2$.

This second order equilibrium profile tells us that player *B* will choose action *c* and, hence, the unique GSPNE of Γ^{A^3} has player *A* playing actions *b* and *t*. Now, we have fully characterized the equilibrium of the game and the outcome that we can forecast is that player *A* will receive a payoff of 2 units whereas player *B* will receive -2 units. In a swindle context it is equivalent to saying that the swindler will expect the other player to be naïve and be tricked by his swindle but he ends up falling for the trap of the opponent.♦

1.4.2 Existence of GSPNE

This subsection establishes one of the most desired properties of a solution concept: its existence under fairly general conditions. Given the hierarchical structure of a game with unawareness and the GSPNE, its existence can be proven by generalizing the standard proofs for games without awareness. Therefore, sufficient conditions for existence are analogous to those of standard games: finiteness of the strategy sets and the number of players. The proof is constructed by induction proving the existence of equilibrium for the first order awareness levels and showing that, by plugging the restricted equilibrium strategies of unaware players to higher order awareness levels, equilibria for these levels can also be found.

Proposition 1: Every game $\Gamma^{r^k} = (N^{r^k}, M^{r^k}, H^{r^k}, P^{r^k}, f_c^{r^k}, \{\Psi_i^{r^k} : i \in N^{r^k}\}, \{u_i^{r^k} : i \in N^{r^k}\}, G^{r^k}, F^{r^k})$ with a *k*-order awareness level r^k in which for every game $\Gamma^q \in G^{r^k} \bigcup \Gamma^{r^k}$ the sets N^q, M^q and G^q are finite has a GSPNE.

In order to prove this proposition we will need a couple of lemmas. The first one appears in Mas-Colell, Whinston and Green (1995) (henceforth MWG) and its proof uses the Kakutani fixed point theorem:

Lemma 1 (Proposition 8.D.3 in MWG): A Nash equilibrium exists in game $\Gamma = [N', \{S_i\}, \{u_i(\cdot)\}]$ if for all $i = 1, ..., N'$

(i) S_i is a nonempty, convex and compact subset of some Euclidean space $\Re^M$

(ii) $u_i(s_1, ..., s_i)$ is continuous in $(s_1, ..., s_i)$ and quasiconcave in s_i.

Lemma 2: Consider a game with unawareness Γ^{r^k} ($k > 1$) in which for every game $\Gamma^q \in G^{r^k} \bigcup \Gamma^{r^k}$ the sets N^q, M^q and G^q are finite. Each lower order equilibrium profile $EP_v^{k\text{-}1}$ generates at least one generalized Nash Equilibrium.

Proof of Lemma 2:

Given the equilibrium strategies for the unaware players in $EP_v^{k\text{-}1}$, the game with unawareness Γ^{r^k} can be viewed as a standard game without unawareness Γ'. In this standard game, the players are all the fully aware relevant identities $j \in J^{kr^k}$ plus all the unaware identities that are restricted to play according to the partial strategy specified in $EP_v^{k\text{-}1}$.

Since $EP_v^{k\text{-}1}$ specifies a single partial strategy for each unaware identity, the restricted strategy sets for this identities will be a nonempty, convex and compact subset of some Euclidean space $\Re^M$. Moreover, the set of mix strategies of the fully aware relevant identities will also have these properties. In addition to that, the payoff functions $u_i(\sigma_1,...,\sigma_i) = \sum_{s\in S}[\prod_{k=1}^{N^q}\sigma_k(s_k)]u_i(s)$ are continuous in $(\sigma_1,...,\sigma_i)$ and quasiconcave in σ_i for all $i = 1,...,N'$. Henceforth, all the conditions for Lemma 1 are satisfied and it follows the result of Lemma 2. ♦

Proof of Proposition 1:

Start by considering a first order awareness level r^1 that can be represented through Γ^{r^1}. As argued before, this is a standard game without unawareness. Therefore, since the number of players and the number of actions is finite so will be the strategy sets and the classic existence result (for instance, Proposition 8.D.2 in MWG) implies that Γ^{r^1}

has at least one (generalized) Nash equilibrium. Furthermore, since by the same reasons a Nash equilibrium must also exist for every subgame in Γ^{r^1}, it follows that a (generalized) Subgame Perfect Nash equilibrium must also exist for Γ^{r^1}.

Since this is true for all first order awareness levels in a game and there are finitely many of them, we are able to construct all the different first order equilibrium profiles EP_v^1. Then, Lemma 2 implies that for every EP_v^1 a generalized Nash equilibrium must also exist for each second order awareness level r^2. Again, since this result must also be true for every subgame of every Γ^{r^2}, it must be the case that a GSPNE exist for every game with a second order awareness level.

As long as the number of possible awareness levels is finite (which is granted by the assumption that the set G is finite for every awareness level) an inductive argument implies that a GSPNE must exist for any awareness level of any order. ♦

1.5 Application: Bounded Rationality vs. Unawareness

Historically, Industrial Organization has been one of the fields where Game Theoretic models have been more successfully applied. Even thought it is generally the case that there are no awareness asymmetries among players in an Industrial Organization setting, we can find some instances where these asymmetries exist and players try to take advantage of them. The most common of these asymmetries arise when some consumers are unaware of some relevant fact concerning the market interaction and firms, knowing it, seek to exploit this unawareness to increase its profits.

It turns out that some of these situations have already been studied within the literature of bounded rationality. However, this literature does not formally consider that these players are unaware of some facts but, instead, that they fail to play optimally because they are not fully rational. In concrete, they regard this "bounded rationality" as the inability to foresee some aspects of the game which makes them behave (suboptimally) without taking them into account. For this reason, they call these players "myopic". The myopia of these agents can go from the failure to consider hidden add-on prices (Gabaix and Laibson (2005)) to the fact that firms might use current behavior to discriminate in the future (see Fudenberg and Villas-Boas (2005) for a complete survey) or even sell this information to third firms (Taylor (2004)).

As you can see, the difference between their idea of "myopia" and unawareness is very subtle, if it exists at all. In fact, to describe myopic players, the authors of these papers even use the world "unaware"[13], although in an informal way and not as it is formally defined in Decision Theory. The question is whether the intuition of unawareness that these authors had in mind indeed captures the same idea that is formally stated in Decision Theory and used in this work. From my point of view, at least in some of these models where myopia plays a role, the agents are actually fully rational (able to optimize) and if they fail to find the best strategy is only because they are unaware of some features of the interaction. If my perspective is correct, then the right model to

[13]Gabaix and Laibson (2005) even employ it as a synonym to explain the meaning of "myopic" in their abstract.

formalize this awareness asymmetry would be one like the proposed in this work and we would not be talking about "bounded rationality" anymore, but about "unaware" agents.

To correctly identify whether the problem relies on the rational ability of the agent or in his awareness level, we have to look for the reason behind his misperception. For instance, take an example of an *add-on* from Gabaix and Laibson (2005): if, at the moment to purchase a printer, an agent does not consider the price of the cartridges that he will have to purchase in the future because he thinks they are not relevant, then the problem is in his rationality. If, on the other hand, he does not think about them just because at that moment it does not cross his mind that they will be important in the future, then it means that he is unaware of this cost at the time of purchase. A substantial difference between both scenarios is provided by cheap talk: If somebody tells the agent that he will later have to purchase cartridges, a bounded rational agent should not, in principle, change his behavior whereas a fully rational agent who was previously unaware of this fact will now consider its cost in his decision. Gabaix and Laibson explicitly assume in their model that firms can unshroud information about these add-ons to the myopic consumers and that, if they decide to do so, those ones start taking them into account in their decisions, exactly like a rational unaware player would do[14].

---

[14] In some other models it is not so clear whether the players are unaware or bounded rational. For instance, Della Vigna and Malmendier (2004, 2006) study the case where some players are incapable to correctly predict a change in their preferences. In particular, they analyze the situation where consumers have time-inconsistent preferences and are partially naïve about it. An example would be a consumer that, at the moment of joining a health club, thinks that he will use it more frequently than he will actually do. This overestimation can be caused because the consumer fails to realize at that time that he will later feel "lazy" about going to the gym or, on the other hand, he might be aware of it but just underestimate his laziness. The former would be an instance of unawareness whereas the suboptimal behavior in the second case

If we consider that myopia is not a problem related with rationality but with awareness, the formalization proposed here opens a new door to properly study these situations. As well described by Ellison (2006), the bounded rationality literature has usually modeled these cases using a "rule-of-thumb" approach:

> "The approach of the rule of thumb literature is to simply posit rules of thumb that consumers are assumed to follow. One can think of this as similar to the game-theoretic approach, but skipping the part of the argument in which one posits utility functions and derives the behaviors as optimizing choices."

However, the model here presented permits to properly formalize the fact that some consumers might be unaware of some aspects of the game and derive the behavior of these consumers as an optimal one (given their awareness level) without having to assume it out of nowhere. In truth, the rule of thumb assumed in these papers is not actually taken from "out of nowhere" but from the intuition of its authors. As a matter of fact, it turns out that the behavior that they assume for the myopic players coincides with the optimal behavior that we could derive from a rational unaware player using a game with unawareness and the appropriate solution concept. This is so because, without formalizing it, the authors where implicitly solving a game with unawareness to find the

---

would be more likely caused by wrong priors. In their notation, the first situation would correspond to $\hat{\beta}=1$ and the second to $\hat{\beta}\in(\beta,1)$.

optimal behavior of the unaware players and then just plugging it into the game without unawareness of the fully aware players. Therefore, the formalization presented here allows modeling these situations and obtaining a formal justification of the behavior of myopic agents.

1.6 Related Literature

Right now, there is quite a large literature on unawareness. Dekel, Lipman and Rustichini (1998) showed that standard state spaces are not adequate for modeling unawareness. Following that paper, a growing literature on unawareness both in economics and computer science has shown that more elaborated models are required (Fagin and Halpern (1988), Modica and Rustichini (1994, 1999), Halpern (2001)). Generalizations of Aumann (1976) epistemic state space capturing non-trivial unawareness among multi-agents were later independently introduced by Li (2006) and Heifetz, Meier and Schipper (2006a).

Since then, there have been several attempts to model a game with possible unaware players. Copic and Galeotti (2007) construct a model of awareness in normal-form games with incomplete information. In addition to using a normal-form game, their work departs from the present one by the fact that players and actions are supposed to be common knowledge among the players. Feinberg (2004, 2005) also studies normal-form games with incomplete awareness and provides an extension of Nash equilibrium concept for his games. Heifetz, Meier and Schipper (2006b) study a generalized state-space model with interactive unawareness and probabilistic beliefs. Ozbay (2006) and Filiz (2006) also study games with unawareness where one of the players might communicate with the other one to change her awareness. Filiz provides one of the first applications of games with unawareness as she incorporates it into Contract Theory.

From all attempts to model a game with possible unaware players, Halpern and Rego's (2006) is the closest to the one presented here. They also consider extensive-form games and propose a generalization of the Nash equilibrium concept and prove its existence. The key idea is the same in both formalizations: to describe the game from the point of view of every agent at every possible history. In spite of this, the way the formalization is created in each model is conceptually opposed. Their formalization of a game with unawareness is based on a *true* standard extensive game without unawareness Γ. Given this standard game Γ, they then define the different awareness levels as being sets of runs in Γ plus possible extra movements from nature to determine them.

On the other hand, the formalization proposed here does not start from any *true* game without awareness but, it directly captures the different perceptions (or awareness levels) that players might hold by constructing games with unawareness. Then, it introduces function $F(.)$ to interlink these different perceptions and a couple of consistency conditions to ensure that players only take actions that exist and they are aware of. Thus, this formalization does not constrain the awareness level of a player to be a set of runs of a true game and it allows it to take almost any desired form. This extra degree of flexibility permits that unawareness is not restricted to being unaware of particular actions. Instead, unaware players might have a completely wrong perception of the game. For instance, they may not be aware that there is a problem that precludes certain actions to be played or some fact they did not consider that ultimately affects the final payoffs. What is more, I do not require that players are fully aware of all the histories that lead to a particular situation. Hence, I can intuitively model situations where players are unaware of some actions previously taken by the opponents.

In addition to a function similar to F (.), Halpern and Rego's formalization uses an *awareness function* to describe the awareness of players at each nonterminal history. They use this awareness function to create what they call *augmented games* that describe the subset of nodes of the *true* game that a player knows at each particular time. The formalization presented here does not require this extra awareness function and all the further conditions that it entails. This allows having a simpler representation of a game with unawareness that makes it more similar to the standard representation of a game.

Additionally, the introduction of *orders of awareness* turns out to be extremely convenient, not only to formalize an interaction, but especially to construct solution concepts consistent with the different awareness levels of the agents. This causes the construction of solution concepts to be entirely different in both formalizations. Whereas my formalization uses this hierarchy of awareness levels to sequentially find the strategies consistent with the rationale of the equilibrium concept from lower to higher awareness levels, their formalization treats each of their *augmented games* independently from one another.

Besides, I believe that it is impossible to know all the possible elements that might play a role in a given interaction. Every time a player becomes aware of a new fact the perception of the game may change. This is particularly important when players might be creative and find new ideas they were previously unaware of that modify the preestablished structure of the game. In these situations, a player reaches a higher order of awareness and the whole game changes. The model I propose allows introducing this modification by simply adding a new game with a higher order of awareness to the previous game. All lower order games may remain unchanged since the other identities might still have the same level of awareness. On the other hand, Halpern and Rego's formalization is based on the existence of a *true* standard game over which the game with awareness is constructed. Therefore, the discovery of a new idea means that the *true* game is no longer true and it needs to be replaced by a new one which, in turn, implies that the whole game needs to be modified. This explains why I do not feel comfortable with the use of this *true* game.

1.7 Conclusions

The model presented in this chapter provides a useful tool to formally study strategic interactions where some of the involved agents may be unaware of certain relevant aspects of it. Furthermore, thanks to the use of hierarchies of awareness, the model can be considered as a natural generalization of the standard one used in Game Theory: it keeps its whole structure and it just adds the minimum required elements to appropriately capture the subjective awareness levels hold by players at each time. I believe that, from the overall existing literature, this makes this formalization the closest one to the standard one.

Moreover, this hierarchy of awareness levels also suggests the natural way to generalize the different solution concepts regularly used in game theory to capture the implications of unawareness. Even though these generalized solution concepts require a cumbersome notation, the intuition behind them is fairly simple and they should only be slightly more complicated to compute than their standard versions.

Here, only one of these generalizations, the one for subgame perfect Nash equilibrium, is formally presented. Similar generalizations can be obtained for other solution concepts following similar steps. One that could be particularly interesting is rationalizability. Its appeal would be that, since it is not an equilibrium concept, it does not require that players know which equilibrium is being played by the opponents. This assumption, which may be already strong for standard games, it becomes a bit more awkward when awareness is involved: how could somebody who has just become aware of a relevant fact possibly know in which equilibrium of the new game have the other players coordinated to? On the downside, since we are dealing with extensive-form games, the right concept to generalize would be *extensive-form rationalizability* (Pearce, (1984)), which, since it is already significaly more complex to compute than its normal-form version, it would make it even more complicated and less popular for games with unawareness.

In this chapter, a few possible applications of games with awareness are already briefly discussed. However, one of the major doors that this model opens is the possibility to study creativity and the introduction of new ideas in a game. Up to now, game theory has considered rationality as the unique relevant facet of intelligence that may play a role in strategic interactions. Probably, one of the reasons for this is that standard games preclude the possibility of new alternative ideas by assuming that the entire structure of the game is common knowledge from the very beginning. Nonetheless, once we allow for the possibility that a player is unaware of some facts, the effects of the introduction of a new creative idea can be studied. Since creativity is extremely interesting per se, I dedicate the next chapter to it.

CHAPTER 2

Introducing Creativity in Game Theory

2.1 Introduction

Game Theory is the branch of applied mathematics that studies the strategic interaction among intelligent agents. So far, standard Game Theory literature has interpreted "intelligent agents" as "rational agents". By "rational" we intend to capture the fact that, given a game, agents try to obtain their best possible individual outcome conditional on what they expect their opponents to play. However, "rationality", as normally understood in Game Theory, fails to capture all the different dimensions of intelligence. For instance, a fully rational player is not allowed to search alternative ways to improve his expected outcome *outside* of the game he is facing as he knows it. Nevertheless, in some circumstances, we may expect that a smart agent who is unhappy with the interaction's anticipated outcome might study possible alternative ways to improve upon it. Therefore, when modeling this situation as a standard game with rational players who cannot take any action not specified in the game, we are implicitly putting a bound to their intelligence.

One of the main purposes of this chapter is to raise consciousness about this limitation of the present literature and propose a way to overcome it. In particular, as opposed to *rationality* (behaving optimally inside a given game), I will introduce the concept of *creativity* (capacity of finding viable modifications to a game). By introducing this new ability, which may be crucial in many strategic interactions, we will be able to better analyze, understand and predict the strategic behavior of intelligent agents that are not only rational but also creative[15].

When we try to model a specific interaction by writing down a game, this one needs to capture all the relevant aspects that are significant in this interaction. However, in real life situations there are potentially infinite factors that could play a role and it is obvious that players cannot possibly have all of them in mind. Some of these aspects can be considered irrelevant and players simply abstract from them. Nonetheless, there is another set of non irrelevant factors which have not even crossed the mind of the players that could somehow affect the interaction at hand. In Decision Theory, it is said that in this case the agent is *unaware* of these particular aspects.

Technically speaking, an agent is *unaware* of something if he does not know it, he does not know that he does not know it, and so on *ad infinitum*. Therefore, being unaware is different than just "not knowing" something. Even if a player does not know something (for instance, the exact size of a rival army), he can take it into account as long as he knows that he does not know it. As Harsanyi (1967-68) showed, these situations with incomplete information can be modeled using a game with imperfect information. On the other hand, when players are unaware of some aspect of the game, (for instance, that the rival army possesses a particular new technology) they obviously can not take it into account and so, they play as it did not exist. For that reason, a change in the awareness of

[15] Notice that there other traits of intelligence outside of rationality and creativity that could also have a relevance in strategic interactions. One of them is empathy, which may have a role in explaining the behavior observed in, for instance, experiments about the celebrated Ultimatum game. In this case, empathy helps an agent to correctly predict the opponent's true payoffs or utility, which may incorporate a notion of "fairness", and behave accordingly.

a player modifies the whole perception he previously hold of the entire situation (or game) and consequently, it might have a dramatic impact in any given interaction.

Nevertheless, regular games, as used in Game Theory, do not allow players to be unaware of any feature of the game. Hence, we cannot rely on the standard tools used in Game Theory to model a situation in which one of the players is aware of something the other player is not. But this is precisely the case when a creative player has been able to find a new way he was previously unaware of to improve his expected payoff of the game. In addition to that, in lots of occasions, the value of this innovation relies precisely on the fact that it is not expected by the opponents.

Therefore, when studying creativity (and the strategic value of "surprises"), applying a formalization of a game that allows the possibility that players might have different (possibly changing) awareness levels is crucial. In concrete, I will use the formalization of games with possible unaware players that generalizes the standard extensive-form game used in Game Theory developed in the first chapter. This formalization of a game incorporating unawareness is a useful tool that provides creative players a framework to compute the optimal way to introduce their innovations. For instance, an investor who has just found a new investment strategy might want to carefully choose the best time to use it. When taking this decision he must consider that, the moment he first uses it, other investors might become aware of it, which, most certainly, will have an effect on its future profitability.

Creativity can take infinitely many forms. Therefore, it is beyond the aim of this work to study all the ways creativity can be introduced in a game. Nonetheless, there are several generic procedures that indicate a particular form to apply creativity in a game. I will call them *creative methods*. Particularly, I will concentrate in the study of a one of these creative methods that exploits an asymmetric awareness effect to enhance its effectiveness.

This creative method, which indicates a promising direction where to look for new ideas making creativity easier to apply, might lead players to change their actions in the way intended by the creative one. Nonetheless, this does not imply that the opponents are necessarily hurt by the utilization of a creative method since it can also be mutually beneficial. An example of this latter case would be that it may allow sustaining cooperation in the Finitely Repeated Prisoners' Dilemma (henceforth FRPD).

The existing models used in the literature to obtain cooperation in the FRPD suppose that players are *stupid* in some way or another. In several models it is assumed that some player might have a non-arbitrary form of bounded rationality (that, for instance, forces him to play Tit-for-Tat[16]) or a grain of unawareness (he might not be aware of the possibility of defecting[17]). Conversely, I present a method where cooperation can be achieved by agents whose intelligence is not limited by the formalization of the game and are able to make use of their creative thinking. This is precisely its beauty: it puts forward an effective scheme that fully rational and creative agents can employ to obtain cooperation in the FRPD without having to assume that the opponent is stupid in some predetermined non-arbitrary way.

[16] See Kreps, Milgrom, Roberts and Wilson (1982).
[17] See Feinberg (2004).

2.2 Creativity

Creativity is the ability to generate new useful ideas or associations between existing ideas. In a game theoretic setting, I will define *creativity* as an agent's ability to explore and discover new viable modifications or alternatives he was previously unaware of to a given interaction or game with awareness. Once an agent finds a new viable modification, his perception of the game changes and he sees it as a different game (with awareness). These *modifications* can affect any of the different aspects of the game: action sets, players, information sets, payoffs, etc. For example, an agent may discover a new action that allows him to obtain a higher payoff or a new communication device that solves coordination problems.

Creativity may play a major role in many strategic interactions. Whenever any of the involved parties is not plenty satisfied with the payoff he anticipates, he may try to find suitable alternatives that will allow him to obtain a better outcome. The nonconformist human nature makes this search not an uncommon enterprise and this is precisely how progress usually takes place, not exclusively in strategic interactions. Creativity is, undoubtedly, a fundamental part of it.

In many cases, the creative player will return empty-handed of this search or *creative thinking*, but, in many others, the result will be a brand new vision of the game or understanding of the world. Moreover, when one of the players achieves this new level of awareness, it does not necessarily need to be shared with the remaining participants in the interaction. Then, an asymmetric awareness between the players with probably crucial strategic implications is created.

Using a formalization of a game allowing possibly unaware players, we are able to properly model these situations. Hence, this formalization is a powerful tool at the hands of a creative player that, once he has developed a new idea, he is able to write down a model to study the repercussions of the different possible implementation methods and moments and then decide to adopt the most convenient ones (if any).

So far, the only subfield of Game Theory that has utilizing creativity, probably without considering it so, is Mechanism Design. The job of mechanism designers is precisely to combine their creativity and their knowledge of Game Theory to create a game that should generate a specific desired outcome in a particular situation. So far, the literature has been using all the standard elements of a game to build the mechanisms to achieve its results. However, as far as I know, it has never studied or exploited the effect of asymmetric awareness on its mechanisms. Although this asymmetry can prove to be extremely powerful in delivering the desired outcomes, there may be a couple of reason why awareness considerations have not been taken into account before. The first one is that it does not yet exist a commonly accepted formalization to represent a game with possibly unaware players. The second reason is that a mechanism that relies on asymmetric awareness can only work as long as this asymmetry persists. Once the players become aware of the full mechanism, which in most cases happens after its first implementation, it can no longer take advantage of it.

This can be illustrated through a very old and well known example of mechanism design: the Judgment of Solomon[18]. As the story goes, King Solomon was approached by two women, both claiming to be the mother of a single baby. Each woman offered exactly the same story and, so, this could not be used to distinguish the true mother. After some deliberation, Solomon came up with the following idea: he claimed that there was only one fair solution consisting in splitting the baby in two, each woman receiving half of it. Upon hearing this verdict, one of the women begged the king not to kill the baby but to give it to the other woman instead whereas the other one claimed to be happy with the resolution. After hearing the mothers' reactions, King Solomon correctly identified the first one as the true mother of the baby and resolved to give it to her.

Solomon made great use of his creativity to design this mechanism that exploits the unawareness of the players. At no point was Solomon's idea to actually split the baby. He only made the claim as a stratagem to uncover the truth by tricking the parties into revealing themselves. The key of the stratagem is that the women were initially unaware that their reactions to the verdict would be used to determine the true mother. For this reason, if on the following day two other women who knew the story came with exactly the same poser, Solomon's stratagem would no longer work. Since the mechanism relied on the initial unawareness of the players, once these know how it works, they would both claim that they prefer that the baby is not killed but given to the other one instead.

Unfortunately, history is not always fair and, failing to realize the fact that his mechanism was only intended to work once, King Solomon's name is nowadays more

[18] The story comes from the Old Testament of the Bible in the book of Kings (chapter 3 verses 16-28).

associated to a decision consisting in destroying the matter subject of a dispute than to a creative use of unawareness to construct a successful mechanism.

2.3 Creative Methods

I define a *creative method* as a generic procedure indicating a particular form to apply creativity in a game in order to modify it with the objective to achieve a more favorable outcome.

A very simple and widely used creative method is to look for feasible punishments (rewards) you were previously unaware of and use them to threaten (give an incentive to) other players to induce them to play actions more beneficial for your interests. In order for this creative method to work, the creative player needs to find a large enough punishment or reward that at the same time also needs to be credible. For instance, whereas threatening with putting a bomb might be considered a large punishment according to most standards, in most situations it would be unlikely to work due to its low credibility.

Other creative methods may be focused on the payoffs like, for instance, attempting to find a way to interlink or distort the payoffs to line up interests (maybe through a contract). Others may be centered on the players, such as introducing in the game a new player with a particular function (a referee, for instance). In general, they can be constructed around any element of a game[19].

---

[19] A creative method that might produce devastating outcomes to the opponents is the generation and subsequent spread of rumors. A rumor can take virtually any form, which makes them easy to conceive and

When a player uses creativity to modify a game to one where he expects to obtain a higher payoff, he can do it in two different ways. The first one is with a "real" modification of the game that will indeed be realized at some point during the game. The second way is with a claim affecting a part of the game that will not belong to the new equilibrium path. In the latter case, the modification does not need to be a "real" one and the creative player might only be bluffing about it. These modifications have the advantage that, since they will not actually take place, in many cases they are much easier to find (i.e. they do not require a large dose of creativity). On the downside, for this very same reason, they will also have a *plausibility* problem that makes them more unlikely to succeed on its goal.

In the previous paragraph, by "plausibility" I refer to the likelihood that player assigns to a threat being true and it is related to the feasibility of the punishment specified in the threat. Plausibility should not be confused with *credibility*, which concerns the likelihood that the threat will be actually carried out. For instance, whereas it could seem plausible that somebody would be able to kill himself unless he is given $10, the threat has a low credibility.

However, when assessing this plausibility problem, we also must take into account that sometimes there exists an *awareness asymmetry effect* that helps to overcome these plausibility issues. This effect will be properly explained and studied in the next subsections. There, we will introduce a "bluffing" creative method more intricate

difficult to predict, and they are typically malicious. Although they have always existed, the fabulous spreading power that new technologies such as Internet provide makes rumors exceptionally dangerous nowadays.

than the ones discussed above that will exploit this awareness asymmetry effect to be more plausible and less demanding in terms of creativity at the same time.

2.3.1 A Bluffing Creative Method

Suppose you want somebody to change the action you expect him to play. If he is rational, a good way to achieve this is by offering him the right incentives, either through rewards or through punishments. But, what if there are no available rewards or punishments in the game? Since a game is already supposed to capture all the relevant elements of the interaction the players are aware of, you will be forced to look for this punishments or rewards *outside* of the game. This means that you will have to search possibilities you were initially unaware of. Needless to say, the fact that the players are not aware of these possibilities precisely implies that they are not easy to find since, if they had been obvious, the game would have already incorporated them.

But then again, what if you can neither find feasible punishments/rewards you were previously unaware of? Then, you may look for unfeasible ones. As I will show, that a punishment/reward is unavailable does not mean that it is useless. The fact that the other might be uncertain about its feasibility, or even the mere possibility that it is not common knowledge among players that the punishment is not feasible, might be enough to make them change their actions. Therefore, to achieve our goal we do not need a feasible punishment that we were unaware of, which might be very difficult to find in most contexts, but just some punishment that opponents might consider possible, even if they also considers it to be highly unlikely a priori[20].

Here, I present a bluffing creative method that relies on punishments[21] that players were previously unaware of. The key point is that these punishments do not need to be feasible. This makes this creative method less demanding in terms of creativity because, in general, there are far more unfeasible punishments than feasible ones. Besides, the reason why you might be unaware of them is precisely because they are unfeasible.

Once you find this (possibly unfeasible) punishment using your creativity, the next thing that you need to do is to coerce the opponent to play what you want him to play under the threat of the punishment. Then, whether the trick will be successful or not will depend on the opponent's beliefs. At first, he might find that it is highly unlikely that your menace is truly feasible. However, on the other hand, he might also find that, if it

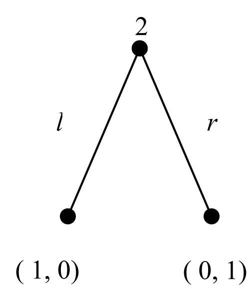

Figure 2.1: Game Γ^U

[20] The meaning of "a priori" in this sentence will later become clear.
[21] A similar bluffing creative method could be based on rewards instead of punishments. There are a few differences, but most of what follows would also be true for this alternative creative method.

was not feasible, it would be extremely unlikely that you were aware of it. Therefore, by Bayesian update, he might find that the probability of the menace being true is high enough to change his actions towards your preferred ones.

I will illustrate this with a simple clarifying example: Consider a game like the one depicted in Figure 2.1. Player 2 prefers to take action *l* and player 1 (who, in principle, can not take any action in the game) would like to persuade him to change to action *r* instead. Now, if player 1 tells player 2 that his family belongs to the mafia and choosing *l* is going to have bad consequences for him, will player 2 reconsider his decision? In one hand, he might consider that *a priori* it is extremely unlikely that player's 1 claim is true. For instance, he might assess that the possibility that a family is related to the mafia is only one in a million. So, unless the threat is extremely high, the expected value of taking action *l* will still exceed the expected value of playing *r*. This would imply that a rational player 2 would not switch to action *r*. Nevertheless, in the previous analysis we made an implicit assumption that is unlikely to hold: the fact that player 2 estimates that the probability that a family is related to the mafia is one in a million does not imply that the probability of this event conditional on the fact that player 1 claimed it (and thus, was aware of it) is the same one. For instance, while it seems reasonable to assume that if the family belongs to the mafia player 1 would be aware of the fact that, by proclaiming it, player 2 could actually reconsider his decision, it is not that clear that the same would be true if player's 1 family did not belong to the mafia. Accordingly, suppose that player 2 assesses that there is only ten possibilities out of a million that a player 1 whose family does not belong to the mafia would be aware that by claiming it he could make player 2 pick action *r*. Then, if player 1 makes this claim, using Bayes' rule, player 2 will estimate that it has almost a 10% chance of being true (far higher that the previous one in a million). Now, a moderate punishment will be sufficient to make him reconsider his decision.

The key that makes this creative method function is the fact that players for whom the punishment is unfeasible are initially unaware of the existence of the punishment and only a small set of creative ones are able to come up with it. If the whole structure of the game was common knowledge among all types of players, everyone would make the claim, the Bayesian update would not alter the prior probability and the opponent would not modify his course of action. Therefore, the existence of asymmetric awareness creates a multiplicative effect on the plausibility of the threat. I will call this effect *asymmetric awareness effect*.

In standard Game Theory, when all players are implicitly assumed to be fully aware of everything, there is not much that a player can obtain by lying. This is because everybody can lie and thus, you will not be able to fool anybody. Nevertheless, under asymmetric awareness, not everybody can lie because not everybody is aware of the lie. This fundamental difference explains why to model this creative method and the asymmetric awareness effect we cannot draw on the standard representation of a game. Instead, we need to employ a formalization of a game allowing for unaware players. Hence, I will use the enhanced extensive-form representation of a game that deals with unawareness developed in the first chapter.

2.3.2 A Model

Consider two players facing a game they initially deem can be represented as the one in Figure 2.1. Assume this is common knowledge among them but, that at the same time they both know that there may also exist some other elements about which they are unaware of that could play a relevant role in the game.

Since player 1 supposes that his opponent is rational and will thus pick action *l*, he has a particular interest in exploring alternative ways that can *change* the game so that player 2 would decide on action *r* instead. A way to achieve this is by using the creative method above described. To do so, he needs to find a possible punishment that player 2 would not know if it was feasible and suits the particular context of the game at hand.

To give a different example that does not rely on the mafia, let's suppose that player 1 is a lawyer and tells player 2 (who is not a lawyer) that there exist some laws that would cause a significant loss to player 2 if this one would choose *r* and player 1 would sue him. In particular, player 1 claims that the *real* game is not like the one in Figure 2.1 but the one depicted in Figure 2.2, where the parameter $a > 0$ accounts for the

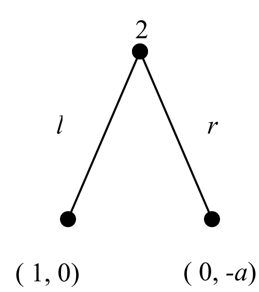

Figure 2.2: Game Γ^C

expected losses that player 2 would suffer after going to the trial[22]. Assume also that checking if these laws do in fact exist is not possible or too costly for player 2.

If player 2 is smart enough, once player 1 has made his claim, he will consider the possibility that this one would be taking advantage of his unfamiliarity with the laws to trick him to play *r* with a lie. For that reason, after player 1 claimed that there exist those laws, the subjective perception of the game being played for a clever player 2 will not be the one in Figure 2.2 but a more complex one with unaware identities.

To properly model this situation, we will use the extensive-form representation of a game with possible unaware players presented in Bages-Amat (2008). When using this formalization, we need to describe an extensive game with unawareness for each possible awareness level and, in addition to that, create a set of functions that, at each awareness level, assign an awareness level and a set of histories to each identity.

As in this game there are two different possible awareness levels (aware and unaware of the potential existence of the laws) we will need two games with awareness to fully describe it. The first game will correspond to the unaware players and it will be the one described in Figure 2.1. Since this is a standard game without unawareness, we do not need any other element to fully describe the game from this awareness level viewpoint. This representation will be called Γ^U.

[22] In this model I will suppose that player 1 automatically sues player 2 after this one plays *r* when the punishment is feasible. Introducing the possibility of choosing not to sue would only make the analysis more cumbersome without adding any insight.

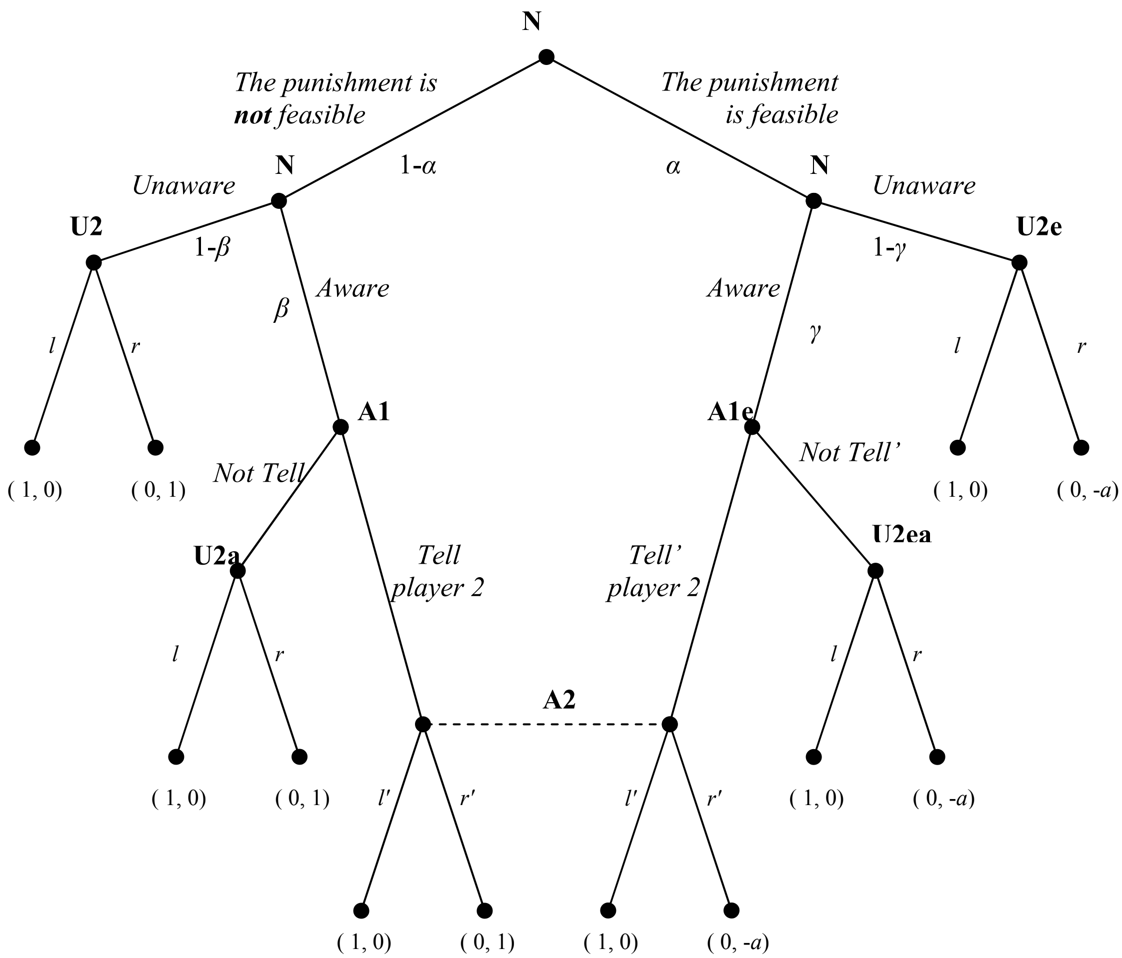

Figure 2.3: Game Γ^A

The representation of the game from the perspective of the identities that are fully aware of the structure of the game will be slightly more complicated and will require Figure 2.3. There, to capture the uncertainty that player 2 has over the true existence of the punishment (i.e., the laws), *nature* makes a first movement in the game and we will denote by α the *a priori* probability of the existence of those laws. Conditional on this event, player 2 will assess the probability that player 1 is aware of the possible (but not necessary the actual) existence of the laws. We will denote by β the probability that player 1 would be aware of the possibility of the existence of the laws when they do not really exist[23]. Similarly, we will define γ as the corresponding conditional probability when the laws do in fact exist. In many contexts, like the mafia one, it will be reasonable to assume that $\gamma = 1$. In general, to have the multiplicative asymmetric awareness effect, it will be enough to have $\beta < \gamma$.

Figure 2.3 captures all the information of Γ^A (the representation of the game from the point of view of a fully aware identity) relative to N^A, M^A, H^A, P^A, $f_c^A, \{\Psi_i^A : i \in N^A\}$ and $\{u_i^A : i \in N^A\}$. However, since this awareness level knows that some identities view the game as Γ^U, this information needs to be incorporated in the description of the game. Γ^U will be the only element in G^A. F^A will be a function mapping the identities U2u, U2a, U2e and U2ea to the empty history in Γ^U and all the remaining identities to their corresponding information set in game Γ^A. This means that identities U2u, U2a, U2e and U2ea view the game as in Γ^U whereas the remaining identities view the game as in Figure 2.3. This completes the full representation of the game capturing the awareness levels of the players at each relevant moment.

In this game there are four relevant identities (i.e. the first identity of a player in a game or the first after a change in unawareness): A1 and A1e for player 1 and U2[24] and

[23] In fact, β is a more complex figure. As I will later explain, there might be other possible claims that a creative player 1 can make to try to convince the opponent to change his action. Yet, he will only use the one that he estimates that will generate better results. However, for now I will abstract of this in order to facilitate the understanding of the intuition behind the creative method.

[24] U2 will denote the identity corresponding to the empty history in Γ^U.

A2 for player 2. As, by generalized perfect recall, a relevant identity is never a successor of another relevant identity of the same player with the same (or higher) awareness level, we can treat them as if they were independent players. Henceforth, we need to define partial strategies for each of them: Identity U2's pure partial strategies are {*l, r*}, A2's are {*l', r'*}, A1's are {*tell*, *not tell*} and A1e's are {*tell', not tell'*}.

2.3.3 Solution of the Game

To start by the simplest case, assume that it is common knowledge that the following inequality holds:

$$P(laws \quad exist \,|\, 1 \quad is \quad aware) \equiv \mu = \frac{\alpha\gamma}{(1-\alpha)\beta + \alpha\gamma} > \frac{1}{1+a}$$

This inequality implies that the probability that player 2 assigns to the real existence of the laws conditional on 1 being aware of its possible existence (μ) is high enough so that if player 2 expects that player 1 will make the claim whenever he is aware of the possible existence of the laws, he will optimally pick *l'*.

If the values of all the probabilities are common knowledge and the inequality holds, then this condition will be satisfied. However, even if the value of β is not common knowledge, if, for instance, it is common knowledge that $\gamma > \beta$ and $\alpha > (1+a)^{-1}$, the condition will still be satisfied.

Since some of the players of this game might be unaware of some features of it, in order to predict the expected behavior of the players it is desirable to use a solution concept that takes awareness considerations into account. A solution concept that satisfies this condition is *generalized subgame perfect Nash equilibrium* (GSPNE) as defined in the first chapter. Following the steps detailed in there to identify all the GSPNE of a given game with awareness, the first thing we need to do is to find all the standard SPNE of all the games corresponding to a first order awareness level (i.e. in this case only Γ^U). The single SPNE of Γ^U consists on identity U2 choosing *r*.

Taking this strategy for the unaware identities as given, we can look for the GSPNE in Γ^A, the game of the fully aware players. Plugging in the fact that the unaware identities U2u, U2a, U2e and U2ea will play *r*, we can view Γ^A as a standard 2x2x2 game. This game has three Nash equilibria in pure strategies that are also GSPNE of Γ^A. In one of them, player 1 always tells player 2 that there exist those laws and A2 always plays *l'*. All the other equilibria (i.e. the other two in pure strategies and all the ones using mixed strategies) involve some identities playing a weakly dominated strategy with positive probability. These equilibria are the ones where A2 always chooses action *r'* and, for some reason, it is more likely that A1 decides to make the claim than A1e. Clearly, in any of these equilibria, identity A1e has to play his weakly dominated strategy "not telling" with strictly positive probability. Therefore, these equilibria are not particularly robust and could be "eliminated" by using some refinement of the solution concept.

Recall that to compute all the GSPNE of the game we assumed that it was common knowledge that $\mu > (1+a)^{-1}$. Still, in most situations it might be difficult to justify that there is some degree of common knowledge regarding the values of *α, β* and γ. Once player 2 is told about the punishment he was previously unaware of, he realizes that the game is like the one in Figure 2.3. Then, he must personally assess the values of these probabilities and player 1 can only try to guess player 2's assessments.

However, to obtain the main results that we wanted to show, it is not necessary to have any degree of common knowledge about these probabilities. Since the action "telling" is (weakly) dominant for identities A1 and A1e, player 1 does not need to make any assessment about the other player's assessments: regardless of the values of the probabilities, "telling" will always be optimal. These probabilities only affect identity A2's final decision, who will use them together with the probability that each identity of player 1 decides to "tell" to compute the expected value of each action. If, for instance, he thinks that each identity of player 1 will take his weakly dominant action "tell" and he estimates that $\mu = \frac{\alpha\gamma}{(1-\alpha)\beta + \alpha\gamma} > \frac{1}{1+a}$, then his optimal decision will be to play *l'*.

Even if we do not eliminate the equilibria that uses weakly dominated strategies, the point that I wanted to show has been proven: there is an equilibrium of the game where a smart creative player facing the game in Figure 2.1 is able to get another player to switch to his favorite action (even when this player is hurt by the action). This shows that this bluffing creative method might work, even without any feasible punishment available. Finally, notice that without asymmetric awareness, the multiplicative effect would not take place and, if the probability of existence of the punishment is small[25], player 2 would never switch his action to *l*.

[25] By small I mean that $\alpha < 1/(1+a)$

Moreover, even thought a general model can not be constructed because creativity is context specific and can take endless forms, the example above is sufficiently general to show how the bluffing creative method would work under a lot of different kinds of both, games and claims. In particular, Γ^U does not need to take the specific form of Figure 2.1 but it could also be another game where player 1 wants the opponent to switch to another action. Also, if the claim asserts that if the opponent takes the undesired action there will be a subsequent new game, all that matters is the equilibrium payoff of this game. This will be considered as the punishment (*a* in the example). In all these cases, the asymmetric awareness effect will behave in exactly the same way.

2.3.4 Further Comments

There are a couple of points that should be noticed though. The first one is that, given this result, once a player knows how this creative method works, we might expect that he will be willing to use it whenever he can take advantage of it. Therefore, a smart opponent in the position of player 2 might be suspicious when hearing this sort of claim if he thinks player 1 knew about this creative method. When assessing the plausibility of a particular claim, player 2 must estimate the probability that player 1 would make it when it is not true. We can break this probability down in three different necessary conditions:

- that player 1 was aware of the creative method (given the claim was not true)
- that player 1 comes up with this specific claim (given the claim was not true)
- that player 1 does not come up with a better claim (given the claim was not true)

In order to facilitate the explanation and the understanding of the mechanism behind the creative method, what we have been calling "β" was only the second of these elements. Yet, it might be the case that player 1 would not use this creative method simply because he did not know that it existed. This is precisely the first of the probabilities that player 2 must assess. If player 2 knows that player 1 read this research (or is aware of it) then he will assign probability one to this event. However, as the fact that, as far as I know, this creative method has never been studied before would suggest, its existence might be fairly unknown. If this is the case, player 2 will, in general, estimate this probability as being zero or close to zero. If so, the multiplicative effect of the asymmetric awareness is maximal and player 2 will be more inclined to believe that the claim is in fact true.

On the other hand, once a player is aware of how this creative method operates, in order to use it he needs to look for potentially convincing claims that he thinks that might work. However, he may find more than one possible claim and he will only use the one (or the combination) that he assesses to be the most likely to succeed. Therefore, what we have been calling "β" should actually be the product of the probability of these three conditional events instead of only the second of them.

The second point I wanted to mention is that I have been deliberately ambiguous on the role played by *Nature*, which has been divided in two separate steps. The first one "decides" the true state of the word whereas the second sets the awareness level of the players. The first step is equivalent to the usual role that nature plays in games of asymmetric information of choosing the *payoff relevant type* among two different types of player 1. However, the second step is conceptually different from the first one (and the ones you can typically find in the literature without unawareness). An important difference is that, whereas in standard contexts players are not allowed to change their type during the game, the awareness level of a player can vary in a game for different reasons. These include communication with other players, introspection and plain *luck*. As a result, what determines the awareness of a player does not depend only on *luck* (or *Nature*): an intelligent player may be able to change it simply by thinking carefully enough. Moreover, the more *intelligent* a player is, the more likely it is that he finds elements he was previously unaware of. Considering this, step two can also be decomposed in two parts: in the first, *Nature* sets the "*intelligence*" of the player (among different types or intelligence levels) and in the second decides whether each different "*intelligent-type*" becomes aware with a different probability. Assuming that this intelligence will have no other effect in the game, we can simplify it by collapsing these two movements in a single one, as it has been done in the game of Figure 2.3.

2.4 Cooperation in the Finitely Repeated Prisoners' Dilemma

Among the multiple environments where this creative method can be applied, one that is particularly interesting due to its historic relevance in Game Theory is the case of the Finitely Repeated Prisoners' Dilemma.

A straightforward application of the above discussed creative method implies that a creative player facing a one-shoot Prisoners' Dilemma can persuade his opponent to cooperate by making an appropriate and convincing claim. However, when the Prisoners'

Dilemma is repeated finitely many times, the application of the creative method generates some other interesting results. First, not only a creative player might convince an opponent to cooperate in some of the periods but, depending on the claim he makes, it might also be in his own interest to cooperate for a certain number of periods. As a consequence, the application of the creative method is not necessarily harmful for the naive non-creative player and it might generate mutually beneficial outcomes. A second important result is that cooperation might arise even when the opponent is convinced that the claim is false, as long as this is not common knowledge.

To illustrate these two points, consider a creative player making a claim of the form: "after *x* repetitions there will be a punishment/reward conditional on both having cooperated in all previous periods". If it is not common knowledge that the other player did not believe this statement, players might be willing to cooperate for a certain time[26]. This is because, even if the second player did not believe the claim, he will still be willing to cooperate to make the first one believe that he might have. Then, the first player will also want to cooperate to make the second believe that he thinks that the other believed the claim and so on; much like in a reputation model. Eventually, one of the players may defect putting an end to the cooperative phase. Nevertheless, since there has been some periods of mutually beneficial cooperation, each player might end up obtaining an overall payoff above the one of the unique equilibrium of the game when players are fully rational but not creative.

[26] A lot of different types of claim may as well work. For instance, that in the stage game there is a third action similar to defect but with worse consequences. Even if the opponent would supposedly learn the action after the hypothetical case of seeing it, cooperation would still be possible thanks to the menace of a second Nash Equilibrium with strictly worse payoffs.

Most of the existent literature on the FRPD tries to rationalize cooperation by assuming that players' rationality is somehow bounded in a specific way. The most important work in this direction is due to Kreps, Milgrom, Roberts and Wilson (1982) and Fudenberg and Maskin (1986). They show how incomplete information about players' options, motivations or behavior can explain cooperation. As Neyman (1999) points out, their approaches can be considered as a perturbation of the strategy sets of the players in which at least one of the players is restricted to use mixed strategies that with positive probability chose a particular non-arbitrary strategy[27]. Radner (1986) explores several departures from the strict Nash equilibria "rationality". In one of these departures, players are uncertain about the degree of cooperativeness of the opponent. In another one, players are satisfied to get close in utility to the best response to the rival's strategy. In this case, as the number of repetitions increases, the corresponding sets of equilibria include those with longer cooperation.

On the other hand, Feinberg (2004) presents a purely syntactic framework with unawareness that he uses to show that a grain of unawareness can lead to cooperation in the FRPD. However, in his model cooperation is only attained when there is a very peculiar form of unawareness; namely, that one of the players might be unaware of the possibility of defecting.

Finally, Neyman (1999) uses a completely different approach that shows that cooperation can be sustained in the FRPD when we drop the assumption that the number of repetitions is common knowledge among players. He shows that to achieve this result

[27] Another paper using a similar approach is due to Vega-Redondo (1994).

only exponentially small departures from common knowledge are required. Nevertheless, by doing so, the FRPD becomes virtually an infinitely repeated Prisoners' Dilemma which makes the result not as surprising as one could initially expect.

There are two main differences between the existing literature and the present approach. The first one is that in most of the literature, to attain cooperation it is necessary to assume that players might have a specific form of bounded rationality. On the other hand, I present a model where agents not only are not supposed to be bounded rational, but they possess an intelligence that lets them be creative in addition to rational. The second main difference is that whereas the aim of the existing literature is to explain the cooperative behavior observed in experiments and real life situations, this work proposes a mechanism that creative players may try to implement to attain this cooperative outcome.

2.4.1 The Game

The previous ideas can be properly formalized using a game with awareness as defined in Bages-Amat (2008). In order to construct a game with awareness, we need a set of games with different awareness levels corresponding with the diverse views that players might hold at a particular moment. The natural way to start is by describing the game from the point of view of a player that thinks that the game at hand is a standard Finitely Repeated Prisoners' Dilemma. In particular, I will use the following version of the one-shoot Prisoners' Dilemma where risk neutral players have to choose between the actions *defect* or *cooperate*. Figure 2.4 shows the payoffs associated with each combination of actions.

	Defect	*Cooperate*
Defect	0 , 0	*b* , -*a*
Cooperate	-*a* , *b*	1 , 1

Figure 2.4: Prisoners' Dilemma

Numbers *a* and *b* satisfy the inequalities

$$a > 0, \qquad b > 1, \qquad b - a < 2 \tag{1}$$

These inequalities imply that, in the one-shoot game, each player prefers to *defect* regardless of the action taken by the opponent (i.e. *defect* is a strictly dominant strategy). As a result, the unique Nash equilibrium of this one-shoot game is both players defecting and obtaining a payoff of (0, 0). Moreover, a standard backward induction argument shows that when this game is repeated finitely many times the unique subgame perfect Nash equilibrium is the repetition of this unique Nash equilibrium of the one-shoot game in every single period.

This is the stage game that both players are presented and assumed to play for T periods, where T is a finite number. Clearly, the outcome of the unique subgame perfect Nash equilibrium is very unsatisfactory when compared to the payoff the players would obtain by cooperating. Thus, an intelligent player might attempt to use his creativity to transform the game into one where he could achieve higher payoffs. As we showed in the previous section, an easy way to accomplish this is by using a bluffing creative method similar to the one introduced there. There are many different types of claims that might potentially persuade players to cooperate for some periods. Here, I will analyze the following one:

Claim A: "If we both cooperate in every period $t \in \{1,...,T\}$, each of us will receive an extra payment of $x > b - 1$ in $T + 1$"

Of course, the creative player needs to find a good justification of the reason why this happens; but this is precisely the creative part, which is context dependent. This justification must be as persuasive as possible because it directly affects the probability the opponent gives to the statement being true. As we will see, the higher this probability is, the higher the willingness to cooperate will also be.

In order to simplify the analysis, even thought x may be considered a variable that player 1 might choose, here we will only consider the case where x is a parameter fixed by the used justification of the claim that cannot be modified by the players.

Once the creative player (that we will also call player 1) makes this claim, the view of the game of the opponent (player 2) switches to a second order awareness level in a way very similar to the one discussed in the previous section. The game from the point of view of this second order awareness level can be represented as a game with awareness analogous to the one in Figure 2.3 but for the FRPD. As in the game presented there, this one also has equilibria where player 1 uses the weakly dominated strategy of not making the claim. Since these equilibria are already discussed in the previous section and do not add anything important or interesting, I will concentrate in the equilibria where player 1 always uses the (weakly) dominant actions of making the claim.

Following the notation of that section, I will denote by μ the probability that player 2 gives to the statement being true once the claim is made. Using Bayesian update we can obtain:

$$P(claim \quad is \quad true \mid 1 \quad is \quad aware) \equiv \mu = \frac{\alpha\gamma}{(1-\alpha)\beta + \alpha\gamma}.$$

The first result that we can derive follows trivially from the fact that, if the opponent does not cooperate, the best possible response is to defect too.

Proposition 2: Given Claim A, for any finite *T* and any parameter values *a, b,* and *x* satisfying (1), there always exists a GSPNE where both players defect in every period.

Notice that the existence of this equilibrium, whose proof is straightforward and, for this reason, will be omitted, is independent of the value of μ. This implies that, regardless of the fact that μ is common knowledge or not, the equilibrium still exists.

This proposition only establishes the fact that, since the claim states that the extraordinary payoff *x* will only be realized if both players cooperate, if a player expects that the other will defect, the best he can do is to defect too. Notice though that this is not a general result and it will not hold under a large variety of alternative claims.

Nonetheless, we are interested in equilibria that allow a certain degree of cooperation and where players expect that the opponent might cooperate with positive probability. I will start by studying the simplest situation when the game is only played once and the value of μ is strictly positive and common knowledge. As we can deduce from the results of the previous section, under these conditions the creative player may induce the opponent to cooperate.

Proposition 3: Let $T = 1$ and the value of $\mu > 0$ be common knowledge. Given any parameter values a, and b satisfying (1), for a large enough x there exists a GSPNE where player 2 cooperates.

The proof is omitted because it is a particular case of Proposition 4 below. However, this result should not come as a surprise because it is just a direct application of the bluffing creative method delivering outcomes identical to the ones of the previous section. In this instance, it makes no sense for a creative player (i.e. a player who makes the claim when it is not true) to cooperate because in the last period there is no reason for doing so. Nevertheless, when the Prisoners' Dilemma is played more than once, there is an incentive for the creative player to cooperate in the first periods to support his claim and increase his payoffs by sustaining longer cooperation from player 2.

Proposition 4: Let the number of repetitions T be finite and the value of $\mu > 0$ be common knowledge. Given any parameter values a, and b satisfying (1), for a large enough x there exists a GSPNE where player 2 cooperates in every period and player 1 cooperates in each period $t < T$.

Proof of Proposition 4:

First, notice that for unaware identities, by the classic result using backward induction we know that the only strategy that can belong to a GSPNE is to defect in every period. Consider now the following partial strategies for relevant identities that are aware of *Claim A*:

- player 1 when the claim is not true: tell the claim and cooperate in every period $t < T$ if and only if both players cooperated in all previous periods and defect otherwise. Defect in the last period.
- player 1 when the claim is true: tell the claim and cooperate in the first period and in the following ones if and only if both players cooperated in all previous periods and defect otherwise.
- player 2: cooperate in the first period and in the following ones if and only if both players cooperated in all previous periods and defect otherwise.

Now, to verify that these partial strategies generate a GSPNE of the game with awareness, we need to check that it does not exist a profitable deviation for any identity.

When somebody defected, we expect the opponent to defect in all subsequent periods. Hence, these strategies are optimal outside the equilibrium path.

Clearly, when the claim is not true, player 1 can not improve his expected payoff because by defecting earlier, player 2 would defect in all remaining periods. When the claim is in fact true, player 1 can not benefit by deviating in any period as long as $x > b - 1$. Also, since player 2 always cooperates if player 1 has been cooperating, player's 1 expected payoff will be strictly positive regardless of whether the claim is true or not. Hence, making the claim at the beginning of the game is the optimal action.

Finally, the most profitable deviation for player 2 will be to defect in period $t = T$ (if $\mu > \frac{b-1}{b}$ or $T = 1$) or in period $t = T - 1$ (if $\mu < \frac{b-1}{b}$ and $T > 1$).

If player 2 decides to cooperate in every period, he expects to obtain a payoff of $(T-2)+1+\mu(1+x)+(1-\mu)(-a)$. Alternatively, if he defects in $t = T$ he expects to obtain $(T-1)+\mu(b)$ and if $T > 1$ and he defects in $t = T - 1$, he expects to obtain a payoff of $(T-2)+b$. Therefore, the expected payoff of cooperating always will be the maximal one if

$$x \geq x^* = \max\{\frac{b-1-\mu+(1-\mu)a}{\mu}, \frac{(b-1)\mu+(1-\mu)a}{\mu}\}$$

Notice that $x^* > b-1$ and henceforth, for any $x \geq x^*$ cooperating in every period will be the optimal behavior of player 2, which concludes the proof. ♦

2.4.2 Two-Sided Reputation

When the value of μ is common knowledge, as we have been considering so far, the reputation effect only arises from the side of player 1 who wants to be seen as a truth teller. However, when we drop the assumption that that μ is common knowledge, a reputation effect can also develop for player 2. In other words: when player 1 does not know the value of μ, even if player 2 believes that the claim is false with probability one, he might still be willing to cooperate for a certain number of periods in order to gain the reputation of being gullible. I will illustrate this with the following version of the FRPD:

Consider again the Prisoners' Dilemma as represented in Figure 2.4 where the parameter values satisfy condition (1). Now, suppose player 1 makes *claim A* but this time assume that the value of μ is not common knowledge but its cumulative distribution function is. In particular, in order to obtain close form solutions, let's assume that μ is distributed uniformly between 0 and 1 and that $b = a + 1$ (i.e. the cooperation cost is independent of the rival's action)[28].

It will be useful to start by considering the case with only two periods to later use the results we will obtain to find equilibria for any larger but finite amount of periods. When $T = 2$, for some parameter values satisfying (1) and for some $\mu_1, \mu_2 \in (0,1)$, there exist a GSPNE where:

- If the claim is true, player 1 cooperates always as long as every player cooperated in the previous period.
- If the claim is not true, player 1 uses the mix strategy consisting in cooperating with probability $p \in [0,1]$ in $t = 1$ and defecting in the last period.
- If $\mu < \mu_1$, player 2 will defect in both periods.
- If $\mu_1 \leq \mu < \mu_2$, player 2 will cooperate in period 1 and defect in period 2.

[28] Making these particular assumptions obviously reduces the generality of the example but it simplifies significantly the algebra allowing to obtaining nice close form specifications of the equilibria. Similar results also hold when these assumptions are not satisfied.

- If $\mu \geq \mu_2$, player 2 will cooperate in period 1 and in period 2 as long as player 1 also cooperated in the first period.

To show that these strategies are in fact part of an equilibrium, let's start by showing that the strategies of player 2 are an optimal response to those of player 1. Since cooperating in the second period can only be optimal if both players cooperated in the first one, there are only three possible pure strategies that can be optimal:

1. Never cooperate and obtaining an expected payoff of:

$$\pi^0(\mu, p) = \mu b + (1-\mu)(pb)$$

2. Cooperate only in the first period:

$$\pi^1(\mu, p) = \mu(1+b) + (1-\mu)[p + (1-p)(-a)]$$

3. Cooperate in both periods (in t = 2 only if both players cooperated in t = 1):

$$\pi^2(\mu, p) = \mu(2+x) + (1-\mu)[p(1-a) + (1-p)(-a)]$$

Notice that, for $p \in [0,1]$, all these expected payoff functions are lineal and strictly increasing in μ and the more cooperative a strategy is, the more increasing its expected payoff function is in μ. That is, for a fixed $p \in [0,1]$:

$$0 < \frac{\partial \pi^0(\mu, p)}{\partial \mu} < \frac{\partial \pi^1(\mu, p)}{\partial \mu} < \frac{\partial \pi^2(\mu, p)}{\partial \mu}$$

We will call μ_1 the value of μ for which $\pi^0(\mu, p) = \pi^1(\mu, p)$ and μ_2 the value for which $\pi^1(\mu, p) = \pi^2(\mu, p)$. For a given p, it is easy to check that these values are:

$$\mu_1 = \frac{a}{1+a}; \qquad \mu_2 = \frac{ap}{1+x-b+ap}$$

Suppose now that, for a given *p*, it is the case that $\mu_1 < \mu_2$[29]. In this situation, a player 1 that knows that the claim was not true will play the above specified mixed strategy only if he is indifferent between cooperating and defecting in the first period. By doing some algebra, you can check that this will be the case only when $\mu_2 = b^{-1}$. This implies that our supposition $\mu_1 < \mu_2$ will be satisfied if and only if $a < 1$. Using the value of μ_2 we can characterize the value of *p* as: $p = \frac{x-a}{a^2}$.

In order for these strategies to truly configure an equilibrium we need to check a couple of extra conditions. The first one is that μ_1, μ_2, and *p* belong to the interval [0, 1] and are proper probabilities. For the parameter values satisfying the above conditions, this will always be true for μ_1 and μ_2 and it will be also satisfied for *p* if and only if $x < ab$. The other condition we need to check is that, when the claim is in fact true, player 1 will actually be willing to cooperate in the first period and in the second (if player 2 did not defect earlier). It can be checked that this will be the case as long as $x \geq 1$.

The expected payoff of a creative player 1 who made the claim up will be 1, which is higher than the expected one if the claim is not made, which is zero. Also, the expected payoff of a truth teller player 1 will be also strictly positive and, as a result, the

[29] Slightly different results are obtained for the case of $\mu_1 \geq \mu_2$ because, depending on μ, player 2's strategy of cooperating only in the first period is dominated by one of the other two pure strategies. It turns out that this case takes only place when $a \geq 1$, which is the exact complementary condition of when $\mu_1 < \mu_2$.

action "making the claim" will be (weakly) dominant. Therefore, the above strategies configure indeed a GSPNE.

To show that the set of parameter values satisfying all the previous conditions is nonempty, take, for instance, $a = 2/3$, $b = 1 + 2/3$ and $x = 1$. With these parameter values we would obtain $\mu_1 = 2/5$, $\mu_2 = 3/5$ and $p = 3/4$ generating a well defined GSPNE.

Using the previously described GSPNE for the case when $T = 2$, we can show that, when the Prisoners' Dilemma is repeated more than twice, under some parameter values there exist equilibria where every player cooperates at least until $t = T - 2$. In particular, even if player 2 believes that the claim is false with probability one, as long as this is not common knowledge, he might cooperate for $T - 2$ periods.

To prove it, we can start by showing that for $T = 3$ there exists a GSPNE where, if both players cooperated in the first period, then they play according to the equilibrium described above for the two remaining periods and they defect otherwise. Since the strategies used in the last two periods are already part of a GSPNE, we only need to check that cooperation in the first period is optimal for both players. An inductive argument will then allow us to generalize this result for any larger finite T.

It can be checked that player 1 who knows that the claim is false will cooperate in the first period as long as $a \leq 1$. Moreover, if he cooperates when the claim is false, he will have even more incentives to do it when it is true; as a result, $a \leq 1$ is also a sufficient condition in this case.

On the other hand, player 2 will have the highest incentive to deviate when $\mu = 0$. Since, in this instance, it can be checked that cooperating in the first period will still be part of player 2's best response as long as $x \geq \frac{a^3}{b} + a$, under this condition player 2 will always cooperate in the first period.

Observe that the parameter values that satisfied the above example also satisfy these two conditions and, as a consequence, there also exists a GSPNE of the game where both players cooperate in the first of the three periods.

Hence, there is a GSPNE of the game when $T = 3$ where every player cooperates in the first period. An inductive argument can now show that then there is also a GSPNE of the FRPD where, by making claim A, every player will cooperate until $t = T - 2$. In this equilibrium, the expected payoffs of a creative player who made up the claim will be of $T - 1$, which is close to the payoff that he would obtain if both players cooperated in every period (i.e. T). As the number of periods increases, not only the expected payoff of a creative player 1 gets closer to the full cooperation payoff, but also player's 2 expected payoff does because the number of periods with two-sided cooperation also increases.

Example in Industrial Organization

Consider a cartel more than fifteen years ago that has been illegally colluding over a significant amount of time in a small European country. They suddenly learn that in 3 years the borders will be open to more powerful competition from the rest of Europe. The members of the cartel immediately know that after the integration their market shares will fall and become insignificant. We can model this situation as a FRPD where cooperation is impossible per se. Imagine also that the government will not allow a merger until after the market is open. In order to restore collusion in the cartel for a few more periods one of the firms might try to use the creative method above described and make the following claim:

"I have become aware of the following scoop: as soon as we enter the European Market, the antitrust authorities will pass a law according to which if one of the members of a cartel reports the other members this one will be rewarded and the remaining severely punished".

All of a sudden, as we have seen in the above model, the threat of this new large punishment might allow the cartel to cooperate for a few additional time.

2.5 Applications of Creativity

As intelligent human beings, we are capable to use creativity to improve our everyday life by generating new ideas to confront the different problems and challenges that we face. When we do it, we are making the most of our creativity. Examples are endless; frequently, its implementation is simple: you have a problem and, after thinking about it, you might find a creative solution. Probably, this solution is to take a certain action which will benefit you. If we are in a strategic setting, the consequences of this new idea must be studied from a strategic point of view. The creative agent not only has to take into account the direct effects of the new idea and the possible reactions of the remaining agents, but she must also realize that other agents that were initially unaware of the new idea, might become aware of it in the course of the interaction. Therefore, all the strategic consequences of the introduction of a new idea can become vastly complex.

Nevertheless, most of the times, the decision of whether to implement it and, if so, when will be a trivial one because its expected consequences will be unambiguously beneficial (or detrimental) for the creative player. For instance, this is usually the case of a mutually beneficial idea or a firm deciding to implement a breakthrough innovation. Although those are typically successful applications of creativity, the fact that the decision regarding its implementation is straightforward makes them uninteresting from a strategic viewpoint.

Because of its intrinsic nature, creativity is difficult to analyze from a general perspective. Once a creative idea has been generated, the world is no longer the same for anyone who has been exposed to such idea. Since it is impossible to return to the previous state (assuming, of course, that the idea is not forgotten) the effect of its introduction can only take place once within the set of involved agents. This implies that most of the interesting applications (especially from a strategic perspective) of a particular creative idea will be very specific, not only in terms of the particular creative idea at hand, but also from the fact that it can only be new once.

2.5.1 Applications in Industrial Organization

Even though a creative idea can only take place once, it is also true that every idea was once a creative idea. For instance, the first time a new pricing strategy, selling mechanism, product feature, advertising channel, rewards program, cost reduction strategy, etc. is introduced, creativity is involved. This explains why, although there are plenty of examples where creativity has been used in Industrial Organization settings, they are usually extremely specific[30].

However, when a firm decides whether to use a new idea or not, there are a number of factors it must consider and some of them are fairly general. In particular, it must be careful about the perverse effects that this new idea might have in the long run once other companies adapt it: the benefits of being the first firm to utilize a new idea might be offset by the posterior negative impact that its use by other firms can produce. This might be either because the industry reaches a new equilibrium where every firm is worse or because some asymmetries make some players better off and some worse off.

The first case can take place, for instance, when firms are able to obtain more information about the consumers. This may lead to more intense competition between firms like through the application of third-degree or behavioral-based price discrimination methods[31]. A rational firm able to predict this outcome might therefore decide not to use a new idea generating this perverse effect. On the other hand, since customers could benefit from this more intense competition, the welfare implications of the fact that this new creative idea is not put into practice could be negative.

When there are asymmetries among firms, the spreading of a new idea might benefit some and harm others. For instance, in the Airline industry the introduction of the frequent flyers programs increased the flyers loyalty which was beneficial for the larger

[30] This is especially true when the players involved are firms that learn (and copy successful innovations) from other industries.
[31] See Fudenberg and Villas-Boas (2006) for a complete survey on Behavioral-Based Price Discrimination.

carriers with broader networks. If prior to its implementation, a small carrier developed the idea aimed to increase the loyalty of its customers it must have also considered that similar programs would also benefit other carriers. Then, it must have foreseen that larger ones would benefit more which, in turn, could end up placing the initial introducer in a worse position. Therefore, if this is the case, the creative firm has two options: do not apply the new idea and hope that other firms do not come up with a similar one or try to change the situation to make the implementation of the idea favorable. In the frequent flier programs this could be achieved, for example, by establishing alliances with other carriers to benefit from interlinked programs.

On the other hand, since the introduction of a successful creative idea can report great benefits to a creative agent, this generates incentives to invest time and other resources in finding new ideas. For instance, in the industry, besides investments in R+D, you might also find firms specialized in providing these services to other firms. Part of the job of consulting firms is to make their client aware of problems they face and ways to overcome them to increase their profitability. Sometimes, the consultant will come up with brand new ideas but, in many occasions, it will use his experience and knowledge about similar situations to “adapt” old ones. Therefore, these firms play two different roles: creation of new ideas and diffusion of successful ones.

2.5.2 Applications outside Industrial Organization

Industrial Organization is probably the field where the application of Game Theory has been more successful. However, the fact that the firms usually know their environment very well and the set of actions they have available for interaction is highly limited (for instance, by laws), does not leave a lot of room for unawareness and, therefore, for being creative. That is why, as opposed to the standard Game Theory, it is easier to find applications in other areas where players do not know their environment as well and have far more possible actions available for interaction at their disposal.

A lot of military settings meet these characteristics. In a war anything goes and the information is far from being perfect. Most of the times, the sides ignore and are unaware of a lot of relevant facts. For instance, each side might be unaware of certain features of the technology of the rivals, their tactics, alliances, information obtained by spies… All this implies that most military interactions can only be modeled using a game with possible unaware players. Moreover, there is a lot of room to use creativity to develop new technologies, weapons, tactics, intelligence and counterintelligence methods…

2.5.3 Swindles

One of the most elaborate uses of creativity also related with Mechanism Design consists in exploiting asymmetric awareness by putting forward a plausible scheme to trick some other agent to voluntarily execute a set of actions that he would not perform if he was fully aware of all of its consequences. This plot is usually referred to as a "swindle", although it does not always need to involve a financial motivation and/or unknown victims[32].

To devise a swindle, the swindler must construct a fake game around the set of actions that he wants the victim to take. By "fake" I mean that, for instance, some of the actions do not need to be feasible but it must look plausible from the point of view (i.e. the awareness level) of the victim. Moreover, the actions he is intended to play must also be optimal from his viewpoint. In other words, they must belong to the equilibrium path of the fake game.

This explanation of a swindle does not rely on bounded rationality. I.e.: the reason why it works is not because the victims are not fully rational, but because they are simply not aware of the whole structure of the game they are playing. There are a number of reasons why a swindle might not work. The first one is because there might be multiplicity of equilibria in one of the games involved. However, most of the reasons have to do with the heterogeneity of the possible victims: different victims might have different awareness levels or different types (in the standard sense of the word: different payoffs, risk aversion, action sets…)

2.6 Concluding Remarks

In order to use creativity, you first need to be able to *think outside the box*, where "the box" is the game tree you think you are playing. In other words, you must drop the

[32] The first chapter of this dissertation uses the example of a swindle to illustrate the representation and subsequent analysis of a game with awareness. The interested reader can check that example to see a proper formalization of a plain swindle.

implicit assumption made so far in Game Theory that the game already captures all the relevant facts of a given interaction. At that time you will realize that in almost any real life interaction there are many potential aspects you are initially unaware of that can have an important impact to it. Only then you will be able to take advantage of your creativity to find the alternatives that best suit your interests.

Game Theory should be something more than a mere tool to analyze strategic interactions in order to make relatively accurate predictions of their outcome. Game theorists must be able to combine their knowledge of Game Theory with their creativity and seek to influence a given interaction by finding ways to improve upon an unsatisfactory expected outcome. This is partially the task of Mechanism Design and, by acknowledging the role of creativity (i.e. by accepting that games can be modified), we must be capable to bring this subfield a step closer to its goal. However, creativity should not only be used to build a full game with several desired features out of nothing, but it should also be used to modify and improve existing interactions.

It is common in the field to regard Game Theory exclusively from the perspective of and outside modeler. Nevertheless, Game Theory can also be a powerful tool at the hands of a player who, in his effort to obtain the best possible outcome, will not hesitate to employ his full intelligence and not only his rationality. This is why evaluating this work as any other work in Game Theory solely in terms of its predictive power over a particular type of interaction is not appropriate[33]. Instead, this work aims to expand the use of Game Theory as an instrument that agents can exploit to impact strategic interactions and reach better results.

In this line, the job of economists is not only to make predictions. Our job also consists in finding solutions to the problems we face and ways to improve the expected outcomes. The formalization of a game with possible unaware players proposed in the first chapter allows evaluating the strategic implications of these solutions and finding the optimal implementation method and time. Furthermore, this formalization may also assist a creative agent to identify the elements of the game that need to be modified or added in order to achieve the desired result.

Interactions, and thus games, are neither static nor fixed. Recognizing this and the decisive role that creativity may play in strategic interactions also allows a better understanding and study of those by increasing the explanatory power of game theoretic models. Admitting that creative thinking may allow reaching better outcomes permits properly analyzing strategic interactions and their evolution from a new perspective. I believe that creativity has a huge potential and acknowledging it brings Game Theory to a new level, not only making it able to study a larger number of situations, but also by expanding the role of game theorist who may improve the outcome of a game by exploiting their creativity.

[33] Recall that it has never been the goal of this research to predict a creative idea. Nonetheless, a game with awareness is a fundamental device to predict the effect of the different possible methods to implement a creative idea (once the creative player is already aware of this idea) and, accordingly, choose the best one.

CHAPTER 3

Non-Exogenous Probability of Repeating a Game

3.1 Introduction

Everyday, each of us interacts with others in many different situations. Some of these interactions may only take place once, but a good number of them are not new for us and they will probably be repeated in the future with exactly the same people and circumstances. When studying these repeated interactions, it is convenient to consider them as a global unit because it is likely that the outcome of past interactions will affect future ones.

Repeated Games is an important branch of Game Theory that has been devoted to the study of this particular kind of interactions. The literature on this field can be divided in two main areas. One of them only considers games for which there is a fixed horizon and players know the exact number of times the interaction will be repeated. These games are known as "finitely repeated games". On the other hand, the second area studies situations for which the number of times that the interaction will be repeated is uncertain. In this case, we say that there is an infinite-horizon and these games are called "infinitely repeated games".

There is a vast amount of work dedicated to the study of infinitely repeated games. An important subset of it aims to characterize the set of payoff vectors that can be sustained as equilibrium of an infinitely repeated game when players are sufficiently patient. The main results of this work are captured by the celebrated "folk theorems".

These theorems prove that any feasible individually rational payoff vector can be sustained as equilibrium payoff vector of the infinitely repeated game when the discount factor is close enough to one[34].

As its name indicates, the intuition behind the folk theorems has been known for a long time. Nevertheless, it remained to Fudenberg and Maskin (1986) to prove the "perfect folk theorem" with perfect monitoring. Later on, other papers considered cases without perfect monitoring. So, Fudenberg, Levine and Maskin (1994) proved the folk theorem with imperfect public information, meaning that players do not perfectly observe the action played by the opponent but they do observe some common variable related with their actions. More recently, Ely and Valimaki (2002) and Piccione (2002) have independently shown that cooperation in the Prisoners' Dilemma can be attained even without public monitoring. Nonetheless, these papers investigated only cases with almost-perfect monitoring and most of their arguments rely heavily on this fact. Finally, Matsushima (2004) proves the folk theorem for the Prisoners' Dilemma with conditionally independence of the signals.

At this point, someone could wonder how well the notable results of the folk theorems regarding infinitely repeated games can describe real life situations that are unlikely to be repeated that many times. The response theory gave to this concern is that, despite its name, infinitely repeated games are able to model situations where the stage game does not need to be repeated infinitely many times. To see how, consider the discount factor δ that players use to discount future payoffs. This discount factor may

[34] A dimensionality condition is also required. See Fudenberg and Maskin (1986).

represent pure time preference. In this case, if r is the rate of time preference and Δ is the period length, then δ is equal to $e^{-r\Delta}$. However, it can also take into account the possibility that the game terminates at the end of each period. For instance, suppose that there is a probability μ that the game continues to the next period. Under these circumstances, one unit of utility next period will be collected only if the game lasts one additional period. Hence, it is worth nothing with probability $(1-\mu)$ and $\delta = e^{-r\Delta}$ with probability μ. This is equivalent to an expected discount factor $\delta' = \mu\delta$. Redefining $\mu' = 1$ and $r' = r - \ln(\mu)/\Delta$, you can see that this situation is the same as the one represented by the regular discount factor. Thus, this shows that infinitely repeated games can represent games that terminate in finite time with probability 1[35].

It is our job to ask ourselves whether this answer is satisfying enough. We must question whether this result implies that infinitely repeated games can be a good approximation of real life interactions and, consequently, provide us with valid predictions or not. To address this subject, it is useful to have a close look to the assumptions of the model. An implicit and necessary assumption of the previous reasoning is that the probability that the game is repeated, μ, is independent of the players' actions. This implies that players have absolutely no influence on whether the game takes place again in the future in the exact same way or not. Among other implications, this rules out that players can decide if they want to keep playing the game

[35] As stated in Fudenberg and Tirole (1991): "the key is that the conditional probability of continuing one more period should be bounded away from 0".

or not, but it also excludes the possibility that the payoffs attained by players have the slightest effect on the likelihood that the interaction is repeated.

Since, by looking at the world surrounding us, it becomes apparent that the instances when this assumption will be actually satisfied are the exception, it is crucial to study in which degree it affects the results. In order to do that, we must see how much these results change once we drop this strong assumption. However, as far as I know, the repeated games literature has made no attempt to analyze this issue or to study what happens if players can influence the probability that a game is repeated. One of the aims of this work is to address this important question and analyze the causes of the changes that may occur and how they will impact the outcome of a game.

3.2 Non-Exogenous Probability of Repeating a Game

There are many possible reasons why a set of agents may stop repeating a particular interaction. Sometimes, as in the cases considered in the standard repeated games literature, these reasons are completely out of control of the involved agents. In these situations, we say that the game ends by entirely exogenous reasons. This could be the case, for instance, if one of the players unexpectedly died: every day he might expect that he will die with a certain small probability but he ignores the exact date of his decease. In this case, the end of the game has nothing to do with the actions played by the players (assuming, of course, that those actions do not affect their life expectancy).

However, in most situations we may expect that players' actions can, somehow, have an influence on the chances of repeating the exact same interaction in the future. If this is the case, players should consider not only the payoff related consequences of their actions, but also their effect on future gains as a result of variations on the probabilities of keep repeating the same game. In assessing all these consequences, it is crucial to know the causes and the implications that stop repeating the particular interaction has for each one of the players. The causes may be diverse and the interaction can either conclude forever or evolve to a new different one[36]. Whichever is the case, it may bring the players to a worse situation (for instance, if their partner fires them) or it could as well be beneficial (their unique competitor fails in bankruptcy and they inherit a monopoly). Depending on these implications, players will have incentives towards increasing or reducing the provability that the game stops being repeated by each particular cause.

The effect that different actions have on the chances that a game is repeated can take many different forms, but it can be classified depending on whether it is caused directly by the actions or indirectly through the payoffs. We say that the effect is direct when the actions per se influence the likelihood that the game is repeated. For instance, this may happen when an agent decides to walk away and stop repeating a given interaction, but also when his specific actions affect this likelihood. The actions that the members of an illegal cartel undertake to preserve its secrecy and prevent regulatory intervention constitute an example of this later situation. On the other hand, we say that the effect is indirect when the probability of repeating a game depends on the payoffs attained by the players, but not on the particular actions taken by them. In this case, if

[36] For instance, this could be due to one of the players (or an external one) finding a useful creative idea affecting the interaction. See Bages-Amat (2008).

different combinations of actions result in the same final payoffs, the probability of repeating the game should be the same across all these different combinations. There are also some situations that are a combination of both effects and it could also be the case that a player considers that his payoffs are not large enough and decides to stop playing the repeated game to do something else expecting to obtain higher payoffs.

Regardless of whether the effect is direct or indirect, its implications can play a major role in any strategic interaction. Failing to realize how different actions affect the likelihood of repeating a game may lead an unwary modeler to reach entirely wrong predictions about the expected outcome of an interaction. For instance, consider a voluntary interaction where a player has an alternative to it, as most interactions are. This alternative could be to play the same game with a different player, to play a different game with the same or another player, or not to play a game at all. If the other agent tries to somehow improve his current payoff by taking advantage of this player, he might take this alternative and, as a result, the expected future payoffs of the agent who tried to exploit the opponent will be hurt. This is an excellent reason to treat well recurring customers, even if it has a cost, and avoid the temptation of increasing current period profits by taking advantage of them. If a modeler did not consider that an unsatisfied costumer may switch to another seller, he may erroneously conclude that the seller will exploit the information asymmetries (for instance, regarding the quality of the product) that are present at the moment of the transaction. Many moral hazard and adverse selection problems can be solved without the use of contracts when the interactions are repeated and players have alternatives[37].

When players can take actions that directly affect the probability of keep repeating a particular stage game, studying the interaction as a repeated game may not always be the best option. In this case, regarding it as a dynamic game right away may be the best mode to treat it. Nevertheless, when this effect is only through the payoffs, it might still be convenient to study it from the perspective of a repeated game in a setting similar but more general than the usually employed in the standard repeated games literature. This is what I will proceed to carry out in the following sections.

3.3 Payoff-Dependent Probability of Repeating a Game

In many situations, payoffs attained by players in the stage game affect the probability of repeating it as it is. For instance, if a player has an alternative to playing the game and estimates that its expected payoff is not large enough, he might drop it and take the alternative. It could also be the case that firms in a certain market are obtaining a level of profits such that outside firms opt to step in. Or maybe regulatory institutions decide to intervene and change a game because they assess that the profits are too high, too low, or too unfairly split.

---

[37] Notice that here I am not invoking the folk theorems. When the minmax payoff of a player is attained when he takes advantage of an other one, the folk theorems only guarantee that there are equilibria where this player receives at least as much as when he takes advantage of the other one. But this does not solve the adverse selection or moral hazard problems. On the contrary, as I will later show, the existence of an alternative can lead a player to not take advantage of the opponent even if this means attaining payoffs below his minmax.

When current payoffs have an impact on the probability that a game is repeated, rational players that are aware of this fact should take it into account when choosing their actions. Moreover, since not only their own payoffs but also the opponent's ones may have this impact, an *altruism* effect (positive or negative) is generated. In particular, the opponent's payoffs may indirectly enter in the expected discounted payoff of the players because the expected number of times that the stage game is repeated is a function of these payoffs. This implies that, in the payoff space of the stage game, the indifference curves of a player may no longer be straight lines orthogonal to his own payoff axis. Under these circumstances, players have preferences over the opponent's payoffs and, as a consequence, they may be willing to forego some of their own current payoff if this is compensated by an increase/decrease on the current payoff of the opponents.

As far as I know, repeated games literature has always been focused on games where the probability of repeating a game was completely exogenous. When this is the case, a player only takes into account his own payoffs and is completely indifferent among two payoff vectors of the stage game that give him the same profit regardless of the one of the opponent. Therefore, this literature has never considered the possibility that this altruism effect existed. Nonetheless, under the presence of this effect, some of its classic results, including the folk theorems, no longer neither hold nor can be applied. Maybe the most important difference that arises is that, in this setup, payoffs below the minmax value of the stage game might be possible to sustain in equilibrium at the same time that others above it might not. At this point, it might be helpful to prove this claim with illustrative examples.

Example 1

The game in Figure 3.1 represents a simple interaction where two players simultaneously decide if the other player earns 1 or 2 units. Since a player's action has no effect on his payoff, any combination of actions generates a Nash equilibrium of this stage game. Furthermore, according to the standard repeated games literature, this implies that, when the probability that the game is repeated is independent of the actions of the players, any strategy should also be part of an equilibrium of the repeated game. As a result, any payoff vector in the feasible payoff set (i.e. where each player gets a payoff belonging to [1, 2]) can be sustained as an equilibrium payoff vector of the repeated game.

	N	*F*
N	2 , 2	1 , 2
F	2 , 1	1 , 1

Figure 3.1: Game 1

However, things start to change when players know that the probability of repeating the game depends on the played actions. Suppose that the value that players assign to not playing the game in smaller than the equivalent of 1 unit per period, so that players always strictly prefer to play the game. If the discount factor is $\delta > 0$ [38] and the probability that the game is repeated is strictly increasing in the payoffs of the players, then not every strategy will still be part of an equilibrium of the repeated game. In concrete, a payoff vector close enough to (1, 1), the minmax payoff vector of the stage game, will not be possible to be sustained as an equilibrium payoff anymore.

To show this claim, it will be convenient to first consider the minmax payoff vector of the stage game. Suppose that, by using some strategies, it was possible to sustain (1, 1) as an equilibrium payoff vector. This would imply that, in equilibrium, each period all players must play *F* with probability 1. Yet, if one of the players decides to deviate and play *N* every period, he would gain from this deviation. This is so because the opponent would not be able to punish him since he is already obtaining his minmax value. Moreover, this deviation improves the payoff that the other player obtains and, therefore, at the same time increases the probability of repeating the game. Since the value they give to each period they are not playing the game is smaller than 1, they will strictly benefit by deviating.

In addition to that, the increase in the probability of keep playing the game that players obtain by playing *N* always has a strictly positive value. Thus, a stronger result is also true because a similar reasoning applies to any payoff vector close enough to (1, 1). So, these players would be willing to forego a smaller payoff above the minmax in order to increase the probability to stay in the game.

---

[38] From here on, when referring to the discount factor δ, it accounts for pure time preference and only for the probability of finish the game that is completely exogenous. Additionally, the assumption that $\delta > 0$ is also made in the whole chapter, even if it is not specifically mentioned.

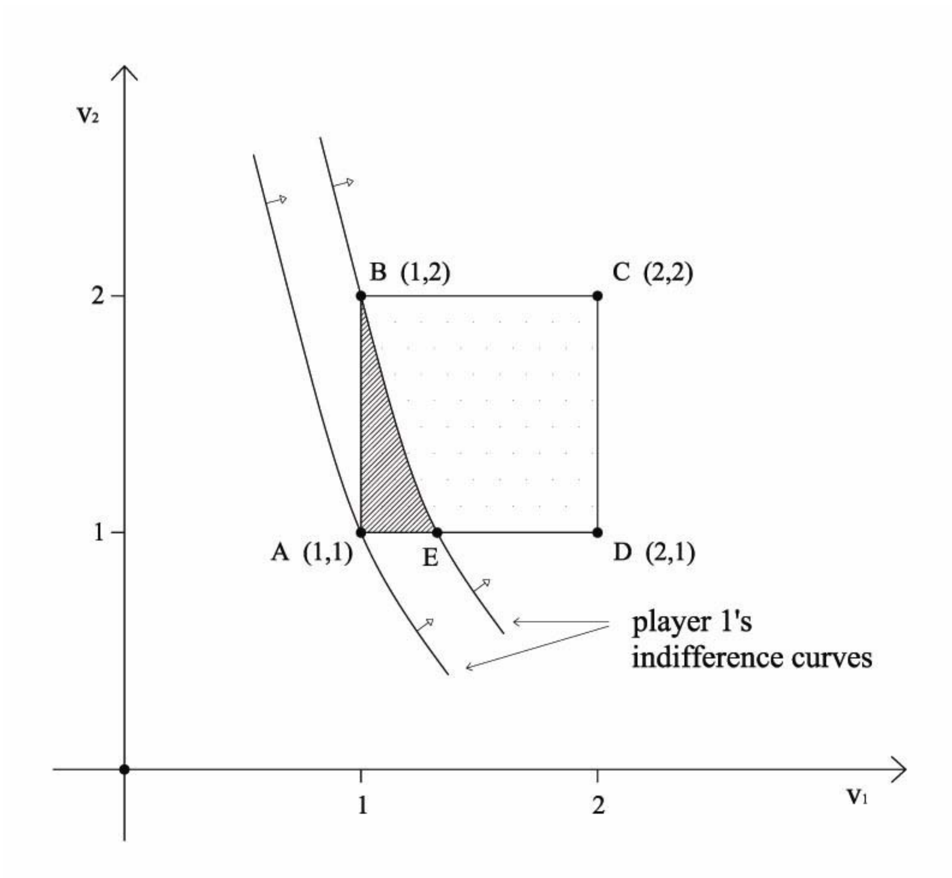

Figure 3.2: Set of feasible payoff vectors

If we assume that, in addition to strictly increasing, the probability that the game is repeated is continuous in the payoffs of the players, we can illustrate this example using Figure 3.2. In this figure, the horizontal axis represents player 1's stage game payoffs and the vertical axis player 2's ones. Square *ABCD* represents the set of feasible payoffs and point *A* is the minmax payoff vector. However, the altruism effect causes that the indifference curves of player 1 are no longer vertical lines. In particular, player 1 will now prefer payoff vector *B* over any payoff vector at the left of the indifference curve that goes through *B*. Henceforth, it will be impossible to implement any payoff vector in this area as an equilibrium payoff vector because player 1 would always deviate and play *N* with probability 1. ♦

This example, simple as it is, already shows that two classic results of the repeated games literature need not to be true anymore when the probability of repeating the game is not completely independent from the actions taken by the players.

Proposition 5: The repetition of a stage game Nash equilibrium no longer needs to be a Nash equilibrium of the repeated game when the probability of repeating the game is not independent from the players' payoffs.

Proposition 6: For any discount factor $\delta > 0$, even if the dimension of the set of feasible payoffs equals the number of players, it might not be possible to sustain all the payoff vectors above the minmax payoff vector of the stage game as a Nash equilibrium of the repeated game when the probability of repeating the game is not independent from the players' payoffs.

Proposition 6 means that "the perfect folk theorem" proved by Fudenberg and Maskin (1986) needs to be modified in order to hold when the probability to end a game depends on the actions of the players. It is worth noticing that most of the assumptions of the game in example 1 are not necessary to obtain this result. As a matter of fact, only in games for which the correspondent to the dashed area *ABE* in Figure 3.2 for the game in Example 1 is reduced to the minmax payoff vector of the stage game the perfect folk theorem will still be valid[39].

Additionally, at the same time that some payoff vectors above the minmax of the stage game may not be possible to sustain in equilibrium, this altruism effect may also cause that some payoff vectors below it could now be sustained. This is so because a player might prefer a slightly lower payoff if this increases the probability that the game is repeated. Therefore, a payoff below the minmax can be strictly preferred by some player to any other above it that, on the other hand, induces a lower or higher payoff to some other players. This implies the following result:

Proposition 7: When the probability of repeating a game is not independent of the players' payoffs, there may be equilibria of the repeated game where some players obtain a payoff strictly lower than their minmax value of the stage game.

Example 2

	C
A	100 , 100
B	110 , 10

Figure 3.3: Lord and serf

[39] A relevant example where this happens is the Prisoners' Dilemma where the probability of repeating the game is increasing in the payoffs and players prefer to keep playing the game.

The degenerated game in Figure 3.3 will be useful to illustrate and prove proposition 7[40]. Clearly, the minmax payoff vector of this game is (110, 10): player 1 can warranty a payoff of 110 units for himself just by playing *B*. When this game is repeated and the payoff received by the players doesn't affect the probability of repeating it, playing *B* is still the only possible equilibrium action for player 1. Since action *B* strictly dominates action *A*, in principle he has no reason for ever switching to play *A*. Nonetheless, if the probability of stop playing the game does depend on the payoff received by player 1, this might no longer be the case. For instance, suppose that player 1 is a lord and player 2 is one of his serfs. If the lord decides to feed the serf a bit better (i.e. to play action *A*), which costs him 10 units, it increases this ones' survival chances. So, it will be more likely that he will keep enjoying the fruits of his labors in the future. If the change in the survival probability is large enough, playing *A* will be the only optimal action.

Thus, notice that, just by taking into account the fact that the payoffs affect the probability of repeating the game, the unique equilibrium of the game switched from always playing *B* to always playing *A*. This is an example of how the predictions can diametrically change if the modeler fails to realize that the probability of repeating a game is not independent from the payoffs.♦

[40] By "degenerated" I mean that player 2 has only one possible action and thus, cannot take any decision at all. Adding a second action for player 2 that affected only slightly the payoffs would not change the conclusions and the analysis would not be as clear.

The previous analysis suggests that the minmax payoff of the stage game may no longer be appropriate to determine the set of payoff vectors that can be sustained in an equilibrium of the repeated game. When the probability of repeating a game depends on the payoffs there are two significant differences that cause this. The first one is that the fact that the payoffs of the opponents enter in the expected discounted utility creates the altruism effect that transforms the indifference curves of the players in “real curves”. This implies that the minmax utility can no longer be expressed as a single value. Instead, when the probability of repeating a game is a continuous function of the payoffs, the minmax utility should be rather expressed as an indifference curve.

The second important difference is that in this setup the future expected payoffs of all players affect the shape of the indifference curves. Therefore, the one period minmax profile of actions against a player depends on the strategies that players will use in the subsequent periods. For instance, a consequence of this is that the minmax profile of actions may be a function of the number of periods that the players will use these minmax strategies.

Ideally, we would like to construct a generalized folk theorem to fully characterize the set of payoff vectors that can be sustained in equilibrium of a game where the probability of repeating it is not exogenous. However, the folk theorems are limit results in the sense that they are only valid when the discount factor is very close to 1. We must recall that this discount factor has two components: it accounts for pure time preference but also for the probability that the game is repeated. Since here we are concerned precisely about the consequences that this probability depends on the actions of the players, it makes no sense conceptually to study this limit case. If we forced these probabilities to converge to 1 (even if it was at different rates), all the interesting effects that we try to study created by the fact that they depend on the payoffs would vanish until, in the limit, the standard folk theorems would be valid again.

3.4 A New Cooperative Force

Game Theory literature has identified so far two main forces that can drive players to cooperate in repeated games[41]. The first of these forces takes place when, by some reason, the payoff functions of the players are interconnected. When this happens, players share the payoffs and, as a consequence, care about the outcome reached by the other player. For instance, if the players are two competing firms, this could be achieved through an explicit contract or an exchange of shares. Nonetheless, this sort of behaviors is usually regulated and antitrust authorities may not allow it.

The second major force to sustain cooperation in a repeated game is through the use of threats of punishments and/or promises of rewards. The use of these punishments and rewards has been vastly studied in the repeated games literature and it is the mechanism exploited by the folk theorems to sustain a particular payoff vector as the equilibrium payoff vector. Continuing with the example of the firms, a duopoly may be able to sustain the monopoly price in equilibrium by using the threat of a price war. As the folk theorems show, by using punishments and rewards every feasible and

[41] Here, unless otherwise noticed, by "cooperate" I mean that a player voluntarily takes a costly action to benefit another player.

individually rational payoff may be sustained in equilibrium for a large enough discount factor, proving that this method can be extremely effective. Nevertheless, a shortcoming of this mechanism is that it requires some sort of monitoring to be implemented; without it, it would be impossible to know when to punish or reward a particular player.

Nevertheless, there is a third force which so far has captured little or no attention at all from the Game Theory literature and that, in certain situations, may allow to sustain cooperation among the involved agents. This force is the fear or the desire that the game stops being repeated. When the probability of keep repeating the game is completely exogenous, this force cannot influence the actions of the players and, hence, it is unable to sustain cooperation in a game by itself. In contrast, if this probability depends on the actions of the players, these ones may deliberately modify their behavior with the purpose of changing this probability in the desired direction. In some situations, this may imply that players will be willing to cooperate (and pay some cost in the current period) in order to increase the total expected discounted utility by affecting the expected number of repetitions of the game.

If players have the option of voluntarily stop repeating an interaction, they may try to use this possibility as a punishment to threat the opponents and force them to take certain actions. Nonetheless, this threat will only be credible when the consequences of voluntarily stop repeating the game are at least as good as repeating the game. Thus, in many occasions this will be an empty threat, like when an employee demands a huge wage increase under the menace of leaving the job without actually having an alternative option. However, even if players do not use empty threats, it might still be possible to sustain cooperation in situations like this one. The reason for this is that the latent menace of a potential arrival of an alternative option is sufficient to create an incentive to increase the payoff of the other players: if these ones enjoy a higher payoff, they will be more reticent to accept an alternative outside option when this comes.

Notice that, even though it causes damage to the other player, leaving the repeated interaction is not used as a punishment in this case. The reason why a player stops repeating the game is not to punish the opponent, but because he expects to benefit from the change. In addition to that, and despite its name, the altruism effect does not mean that players are truly altruists. They are still purely rational one hundred per cent egoist and the only motivation to undertake these costly actions to benefit the remaining players is because this increases his own total expected discounted payoff.

The second application contained in the following section presents a model that fits these features. In it, players form partnerships which they may voluntarily drop after receiving an alternative option that they estimate will represent an improvement to them. The model shows how the altruism effect may cause partners to cooperate even if they are not punished when they do not.

Here, I would like to emphasize the relevance of this type of situations. Most real life interactions are voluntary and, precisely for this reason, agents decide to participate in them because they obtain a benefit out of it. This remarks the significance of the altruism effect that we have been discussing because it emerges in a large number of interactions in which it can play a critical role.

Nonetheless, the possibility of voluntary exit is not a requirement for this force to operate. If the probability of repeating a stage game is increasing (decreasing) in the payoffs of the other players and those prefer to keep (stop) repeating the game, the altruism effect introduced in the previous sections will make players willing to increase the payoffs of the opponents, even if this has a cost in terms of their own current payoff. Since the reason why players cooperate is because they actually prefer to do it, punishments or rewards are not required to enforce it. This implies that monitoring is not required to attain this cooperation. Therefore, these incentives to cooperate may allow sustaining cooperation among players in setups where punishments fail to achieve it due to the lack of monitoring.

3.5 Applications

This section presents several applications of repeated interactions where the probability that the interaction is repeated is not exogenous. In these applications, the predictions that a modeler would make may be entirely erroneous if he fails to realize or to consider that the probability that the interaction is repeated depends on the actions of the players.

The first two applications show that players might be willing to cooperate even without being punished when they do not. On the contrary, the third application is an example of a situation where the fact that the probability of repeating the interaction depends on the payoffs causes the players to become more confrontational, with one of them trying to expel the other one. The forth application illustrates an apparent paradox caused by enemies that, despite this fact, do not want the rival to disappear.

3.5.1 Cooperation without Monitoring

Consider an interaction where the probability of keep repeating the stage game depends on a subset of the actions of the players. Assume that the actions that maximize the probability of keep repeating the game are costly and non-observable. Also suppose that these actions do not affect the stage payoffs of the remaining players. In this setup, even if there exist other observable actions, the only signal that a player receives regarding the effort level of the opponents is that they are still playing the game. The longer they have been playing, the more likely the other players have been making an effort. However, this signal is not sufficient to implement a punishment scheme able to make players undertake the highest effort to keep repeating the game. This is so because the only possible "bad" signal is that the game is over and, consequently, after receiving it, it is no longer possible to punish any player.

Now, suppose that there is also another subset of actions that affect the stage game payoffs which are somehow possible to monitor (lets call them the observable actions). In this game, even if it is not possible to punish the players for the non-observable actions, it might be the case that they decide to take the costly non-observable actions in order to increase the probability of repeating the game. This will be the case when its cost is smaller than its expected benefit. The latter depends on the increase of the probability of keep playing the game, the expected value of being out of the game and the expected value of keep playing the game. Hence, a player may be able to increase the incentives to make the opponent make the non-observable effort by increasing this one's expected payoffs. Thus, if a player is interested in keep playing the game, an *altruism effect* is created that makes this player willing to forego some of his own current payoff to increase the opponents' expected payoff. The reason for this is that, if the other players expect to receive more out of the stage game, these players will also have higher incentives to increase the non-observable effort which, in turn, increases the probability of keep repeating the game and, as a result, his future expected gains.

Example 3

	Defect	*Cooperate*
Defect	0 , 0	*b* , -*a*
Cooperate	-*a* , *b*	1 , 1

Figure 3.4 Prisoners' Dilemma

Consider a situation involving two individuals playing the version of the Prisoners' Dilemma presented in Figure 3.4 where numbers *a* and *b* satisfy the inequalities:

$$a > 0, \qquad b > 1, \qquad b - a < 2 \qquad (1)$$

This stage game is played repeatedly with players discounting future payoffs with a discount rate $\delta > 0$. The game ends each period with a certain probability and, every period, each player can incur into an effort that has a cost $c > 0$ to increase this probability. If both players take this effort, the probability of keep repeating the game is μ_1; if only one of them takes it, this probability is μ_2, and μ_3 if non takes it, where $\mu_1 > \mu_2 > \mu_3$.

This effort is non-observable for the players. In other words, players do not know whether the opponent is taking it or not. So, at the end of each period, players only know their own payoff (and thus, if the opponent cooperated or not in the Prisoners' Dilemma part of the game) and whether the game will be played or not for at least one more period. The expected discounted value of a player if the game is no longer repeated is *U*.

The solution concept that I will use is Nash Equilibrium. However, since the whole game is also the only subgame of the game, the set of Nash Equilibria will coincide with set of Subgame Perfect Nash Equilibria.

Depending on the parameter values, this game has two interesting equilibria. The first one involves both players defecting always and never incurring into the cost to increase the probability of repeating the game. A necessary and sufficient condition for this equilibrium to exist is that $U \geq \dfrac{c(\delta\mu_3 - 1)}{(1-\delta)(\mu_2 - \mu_3)}$, which is a negative number.

Nevertheless, for some parameter values, it also exist a Nash Equilibrium of the game where both players cooperate in each period and, at the same time, take the non-observable effort to increase the probability of repeating the game. To see that this can be sustained as an equilibrium, consider the following strategy: start cooperating and making the effort and keep doing it as long as both players cooperated in all previous periods. Using this strategy when the other player is also employing it will be optimal if the following two conditions are satisfied in the equilibrium path:

i) That defecting is never optimal. This will be the case as long as the expected discounted value of cooperating is larger than the maximal one a player can obtain by defecting. I.e.:

$$\frac{1-c+(1-\mu_1)U}{1-\delta\mu_1} \geq \underset{i\in\{1,2\},j\in\{2,3\}}{Max} \; b+\delta\mu_i \frac{(1-\mu_j)U}{1-\delta\mu_j} \tag{2}$$

ii) That making the effort always dominates not making it. In order for this condition to be satisfied, the discounted expected utility must be higher when making the effort. Since it is not possible to punish a player who is not taking this effort, the cost of the effort must be compensated by the increase of the probability repeating the game. I.e.:

$$\frac{1-c+(1-\mu_1)U}{1-\delta\mu_1} \geq \frac{1+(1-\mu_2)U}{1-\delta\mu_2} \tag{3}$$

For instance, when $U = 0$ and $c < 1$, these two conditions are equivalent to $\delta \geq \frac{b+c-1}{b\mu_1}$ and $\delta \geq \frac{c}{\mu_1 + (1-c)\mu_2}$ respectively, which are satisfied for a large range of parameter values. On the other hand, in this case, when $\delta < \frac{b-1}{b\mu_3}$, cooperation is impossible to sustain without the non-observable effort because players would always prefer to defect. Consequently, when $x \equiv \max\left\{\frac{b+c-1}{b\mu_1}, \frac{c}{\mu_1+(1-c)\mu_2}\right\} < \frac{b-1}{b\mu_3}$, for the discount factors belonging to the interval $\left[x, \frac{b-1}{b\mu_3}\right]$, cooperation in the Prisoners' Dilemma can only be attained if players make the effort to increase the probability to keep repeating the game. What is more, at the same time, players will only make this effort if the cooperative outcome is attained in the Prisoners' Dilemma.♦

As this example shows, players might be willing to take a costly action even if they cannot be punished by the other players when they fail to take it. Moreover, when the probability of repeating the game is not independent of the actions of the players, the altruism effect might increase the incentives to cooperate[42] making cooperation possible in setups where, without the non-observable effort, it would be impossible.

Although the assumptions of this model might seem difficult to be met in real life situations, it fits very well in a particular sort of interactions: the illicit ones. A reason why repeated illicit relations exist is because they are beneficial to its members. However, by its illicit nature, these members are also required to make an effort to try to avoid that the relationship is discovered and, as a consequence, probably dissolved with possible severe punishments. Sometimes, due to its secrecy, this effort can be difficult to observe. Hence, the model above presented may be a good description of these situations. Examples would include an illegal cartel or even a love affair.

[42] These incentives can be measured, for instance, in terms of the minimum discount factor needed to sustain collusion.

3.5.2 Cooperation with Adverse Selection and Moral Hazard

In this section I will present a model where players can freely form partnerships that benefit both players and may last several periods. In this model, when two players meet and contemplate forming a partnership, they ignore how productive the other player is going to be. This means that there is an adverse selection problem. In particular, I will assume that there are only two different types of players: highly productive ones that can choose whether to make an effort which provides a benefit to the other player, and low productive ones who cannot do any effort at all (or this one is useless). I will also assume that there is no difference among a highly productive player not making the effort and a low productive one. Hence, in addition to an adverse selection problem, there is a moral hazard one because a highly productive player may decide not to make any effort.

This model will also endogenize the decision of leaving the partnership (and hence, to stop repeating the stage game). With certain probability, every period each player may receive an outside proposal to form an alternative partnership from a new player of unknown type. This is a take it or leave it offer and, if he decides to accept it, the previous partnership terminates leaving the mismatched player in a strictly worse situation (for instance, waiting for a new partnership offer while receiving lower payoffs).

As we will see, the threat of a potential alternative partnership proposal to your current partner generates the altruism effect above discussed making players to be willing to undertake some effort to increase the payoff of the opponent. Therefore, the moral hazard and adverse selection problems present in the model might be overcome without the need of contracts or punishments.

Model

In this game there are two possible types of players (highly and low productive) that may form partnerships to play the following stage game: if a player is highly productive he can decide to make an effort that costs him $c \in (0, \frac{1}{2})$. If a player is not highly productive, he cannot make any effort. Every player extracts a utility of 1 unit just for belonging to the partnership and, if the partner makes an effort, the player receives an additional unit. At the beginning of a partnership, players ignore which type of player their partner is. A priori, they only know that with probability α this will be highly productive. Players also discount future payoffs with a discount factor $\delta \in (0,1)$.

The notion of equilibrium that I will use is Subgame Perfect Nash Equilibrium (SPNE). As opposed to Nash Equilibrium, this concept has the advantage that players cannot use the punishment of leaving a partnership as a threat to obtain what they want because, if they leave the partnership when they do not have an alternative offer, they will end up in a worse situation and hence, threatening on leaving the partnership cannot be part of a SPNE.

When players do not consider the possibility that their current partner leaves them to form another partnership or when they fail to realize that this decision could be influenced by the payoffs he currently receives, in equilibrium they will never make a continuous effort unless the partner makes some effort too. To see why, notice that the only threat that under this situation can make a player incur into a repeated effort is that the other stops making this effort. However, this threat can only work when a player is actually making some effort. Hence, a highly productive player will never make a continuous effort when his partner is not highly productive too.

These results can change dramatically when players take into account that after every period each player may receive an outside proposition to create an alternative exclusive partnership with a new unknown player. In concrete, each player may receive this proposal with probability $(1 - \mu)$, independently of whether the partner also received a proposal or not. Players ignore if the new player is highly productive or not. The proposal is a take it or leave it offer and accepting it implies the end of the previous partnership. The mismatched player is left in a worse position, no longer enjoying the benefits of being in a partnership. Specifically, in order to simplify the computations, I will assume that he receives a payoff of 0 units per period until he finds a new partner and that the probability of finding this new partner is also $(1 - \mu)$ per period.

Under these circumstances, for some parameter values there is a SPNE of the game where a highly productive type always makes the effort regardless of the actions of the partner. Moreover, if the decision rule to accept the offers used by players is to leave a partnership unless the partner always made the effort, the only optimal strategy for a highly productive type consists in making always the effort. As a consequence, we can show that, conditional on players using this decision rule for accepting offers, for some parameter values there is a unique SPNE of the game where players make the highest possible effort.

In particular, the equilibrium strategy of the highly productive players will be to make the effort as long as he has made it in all previous periods, regardless of the actions of the partner. The only instance where a highly productive player will not make the effort is outside the equilibrium path after, by mistake, failing to make the effort in a previous period.

The decision rule that both players will use is to always accept an incoming offer unless the partner made the effort in all previous periods. Given the strategy used by highly productive players, this is an optimal decision rule because if the partner did not make the effort in all previous periods, according to the equilibrium strategies, he will not make an effort in the future. Thus, switching to a new partner of unknown type will generate at least the same expected payoffs. On the other hand, if the current partner was making the effort in every single period, we cannot improve by switching to a new unknown partner. What is more, even outside the equilibrium path this is an optimal decision rule and, consequently it can be part of a SPNE.

Now, we can show that, given this decision rule, for some parameter values the strategy of the highly productive players is the only optimal one. To do so, it is convenient to start by comparing the expected gains that, in equilibrium, this player would obtain by making the effort in the worst possible scenario (i.e. when the partner is not making any effort) with the ones he would obtain by taking the best possible deviation.

Call V^{+} the maximum expected discounted payoff that a player can obtain if he has made the effort in all previous periods and does not expect the partner to make any effort, V^{-} when he has not made the effort in all the previous periods, and V when he starts a new partnership with an unknown partner. Also, let $V^{+}(e)$ be the expected discounted payoff of a player that has made the effort in all previous periods, does not expect the partner to make any effort and decides to make the effort and $V^{+}(0)$ when he decides not to make the effort. Finally, let U denote the expected discounted payoff of a player who has no partnership. We have:

$$V^{+}(e) = 1 - c + \delta(1-\mu)V + \delta\mu[(1-\mu)V^{+} + \mu V^{+}] \quad (4)$$

$$V^{+}(0) = 1 + \delta(1-\mu)V + \delta\mu[(1-\mu)U + \mu V^{-}] \quad (5)$$

A sufficient condition for making the effort to be optimal regardless of whether the partner makes it or not is that it is optimal to make it when the partner does not do it; i.e. $V^{+}(e) \geq V^{+}(0)$. Since it must be the case that $V^{+} \geq V^{-}$, replacing V^{-} for V^{+} in (5) and subtracting the resulting inequality from (4), we can obtain the following sufficient condition for making the effort to be optimal:

$$\delta\mu(1-\mu)(V^{+} - U) \geq c \quad (6)$$

Suppose that, after not making the effort for a single period and if he did not receive an offer in that same period, our partner would still be willing to give us a second chance and, if we started making the effort again, not dismiss us in case of receiving an offer later on. Then, what this condition says is that if the cost of making the effort for one period (c) is smaller than the expected benefit of making it (i.e. eliminating the risk that in the next period somebody comes and takes our partner away), then making the effort should be optimal.

If making the effort is optimal when you have always been making it, then it must be the case that $V^{+}(e) = V^{+}$. Substituting $V^{+}(e)$ in equation (4) and solving it for V^{+} we can obtain:

$$V^{+} = \frac{1 - c + \delta(1-\mu)V}{1-\delta\mu} \qquad (7)$$

Since, while without a partner, the probability of finding a new partner is $(1 - \mu)$, the expected payoff at the beginning of a period of someone without a partner can be written as:

$$U = \frac{\delta(1-\mu)V}{1-\delta\mu} \qquad (8)$$

Using equations (7) and (8) we can rewrite condition (6) as:

$$\delta\mu(1-\mu)\frac{1-c}{1-\delta\mu} \geq c \qquad (9)$$

which can be rearranged as:

$$\delta \geq \frac{c}{\mu(1-\mu+\mu c)} \qquad (10)$$

Therefore, inequality (10) gives us a sufficient condition regarding the parameters of the model that, when satisfied, implies that making the effort in every period is optimal. For instance, when the cost of the effort is ¼ and so is the probability of having an offer from a new partner, then for any discount factor above 0,762 it is optimal to always make the effort.

Moreover, when the parameter values of the model satisfy condition (10) and players use the above mentioned decision rule to accept incoming proposals of partnership, making always the effort is the only optimal strategy for a highly productive type. To see why, notice that since condition (6) is satisfied, making the effort is optimal regardless of the action of the other player. Consequently, there is no punishment that this one can impose that would make a player stop making the effort[43].

Even when condition (6) is not met and making the effort when the opponent is not doing it is not optimal, it might still be possible that two highly productive players forming a partnership keep making effort. In this case, the typical Nash Reversion technique that (for a large enough discount factor) makes possible to sustain cooperation in the infinitely repeated Prisoners' Dilemma can be used to warranty that players keep making effort through punishments. However, even in this situation, realizing that the length of the partnership in not independent of the payoffs can make a difference: the altruism effect caused by the threat of a newcomer can be combined with Nash Reversion punishments to reduce the minimal discount factor required to sustain a mutual effort.♦

This model captures interesting features of oligopolistic industries where partnerships among subsets of the players can benefit its members. This is the case of the Airline industry and most industries where firms can take advantage of forming larger interconnected networks, like the ones related with logistics. In these industries it is frequent to see firms creating new partnerships and breaking old ones.

[43] Notice, however, that for this parameter values this is the unique SPNE of the game conditional on that decision rule. There are also other decision rules that can be optimal for which there are other optimal strategies.

Another situation that can be partially described by a model like this one is couple relationships. There are several reasons why people treat their partner the best they can. One of them, probably the most important one, is because they care for one another and their utility functions are interconnected. Another one is to obtain rewards and/or avoid punishments from the partner. Both of these reasons have been vastly studied in Economics and Game Theory. Nevertheless, this model exemplifies the third force to cooperate introduced in section 4. This motivation to give the best to your partner comes from the fear that she/he will leave you if somebody else can convince her/him that he/she can provide her/him a higher utility. The best way to minimize this risk is by providing the highest possible utility to your partner from the very beginning.

3.5.3 Market Entry

Consider a monopoly that receives the entry of a new competitor. The incumbent may decide to try to cooperate and share the monopoly profits, to compete and obtain lower profits (Cournot outcome) or to engage in a price war generating significant losses to both firms (for instance, due to fixed costs). From the folk theorems we know that, when the probability of keep repeating the game is exogenous and players are sufficiently patient, any feasible payoff vector above the minmax can be sustained as the equilibrium payoff vector. Therefore, if in this case the minmax payoffs are the ones corresponding to the price war, any of the above mentioned possibilities (and anyone in between) except the price war can take place in equilibrium when the probability of keep repeating the game is exogenous.

Now, suppose that the financial structure of the entrant is far weaker than the one of the incumbent. In other words, that after a few periods without making profits or after a certain amount of losses the entrant will be forced to leave the market. Under these circumstances, the best option (the only one that maximizes expected discounted profits) for the incumbent might be to engage in a price war to force the entrant to leave the market and recover the monopoly.

Example 4

Consider a market with an inelastic demand of 10 units per period. In this market there are 2 firms that have a production cost of $1 per unit produced and a fix cost of staying in the market of $1 per period. At the beginning of each period, these firms set a price which can only be either $1 or $2 per unit. The firm that sets a lower price covers all the demand at that price and the other one does not produce any unit. If both firms set the same price, they equally split the demand.

Hence, if both firms set a price of $1 per unit, they both end up selling 5 units covering the variable costs and losing $1. Similarly, if both firms decide to charge $2 per unit, they also sell 5 units but this time they make $4 in profits each. If one firm sets a price of $1 and the other $2 per unit, the first firm sells all 10 units and the other none but both lose $1 in the process. Since if one firm sets a price of $1 the other one will certainly lose $1, $\underline{v} = (-1,-1)$ is the Minmax payoff vector of the stage game.

When the probability of repeating the game is exogenous, the folk theorem in Fudenberg and Maskin (1986) tells us that for any feasible payoff vector where each player obtains strictly more than -1, there is a discount factor $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a Subgame Perfect Nash Equilibrium of this game with these payoffs.

However, this result can drastically change when the probability to finish the game is not independent from the actions of the players. In particular, suppose that player 2 can sustain a maximum of three straight periods of losses. After this, he will be forced to exit the market and player 1 will inherit the monopoly of the market. Once player 1 enjoys the monopoly, he expects to obtain a profit of \$9 per period by charging \$2 per unit. It can be easily checked that, under these circumstances, for all $\delta \in (\sqrt[3]{1/2}, 1)$, in any Subgame Perfect Nash Equilibrium of this game player 1 will charge a price of \$1 for the first three periods making both firms lose \$1 per period to force player 2 to exit the market.♦

In this example of a market entry, we did not allow for the possibility of voluntary exit, as the exit of the entrant was forced by the circumstances of the game. If the entrant had the possibility of voluntarily exit the market after checking that the incumbent had no intention to accommodate his entry, he would have definitely leaved the market. Moreover, after seeing the incumbent's behavior, other possible entrants will think it twice before entering this market in the future.

Once again, as this example illustrates, taking into account that the actions of the players may affect the probability of repeating in the game may totally change the predictions. Hence, every time we model a situation as a repeated game we must be careful to consider whether the probability of keep repeating it is truly exogenous and out of the reach of the involved agents.

In particular, the regulation implications of a model like this one may be critical: for instance, a regulatory agency failing to realize that the behavior of the incumbent is aimed to eliminate current and future competitors, may view the price war engaged by this one as a signal of healthy competence that benefits the consumers. However, the consequences of these two different scenarios are diametrically opposed, and so are the regulatory needs. In addition to that, this model can be easily extended to show that this price war may have negative welfare repercussions.

This belligerent behavior has been observed, for instance, in the Airline industry. Some network carriers decided to fight the entry of low cost carriers forcing them to exit from their most profitable routes. Some other ones ignored this threat and low cost carriers nowadays enjoy a significant market share of routes originating in the Hubs of these unwary carriers.

3.5.4 Cold War

In a Cold War setting players are enemies in the common use of the world: they have incompatible interests, defend opposing ideals or hold conflicting views of how the future should look like. Each player is a threat for the survival of the rival and they both draw support from individuals who want the enemy to be annihilated. Nonetheless, if they succeed in eliminating the rival, they may loss the support of those who just backed them because of the fear to the opponent and, under the new state of events, prefer to switch to less radical positions. The combination of both effects might generate a scenario where players decide to play "soft" actions against the enemies to minimize the probability to destroy them even if the payoffs of playing "tough" would be higher.

Example 5

In this cold war setting, two enemies that discount future payoffs with a discount factor δ are repeatedly playing the stage game represented in Figure 3.5 where they can decide whether to play a tough or a soft action against the enemy.

	Soft	*Tough*
Soft	3 , 3	2 , 5
Tough	5 , 2	4 , 4

Figure 3.5: Cold war

In a single period, a player strictly prefers to play tough on the enemy and obtain a higher payoff in the current period. In other words, considering only the stage game payoffs, playing tough is a strictly dominant action. However, the actions of the players also have an impact on the probability of repeating the game: if a player plays tough it is more likely that the enemy will be destroyed and, as a consequence, they would stop repeating the game. Whether this is something desirable or not for the players will depend on their expected payoffs once the game is over. In turn, those may depend on the particular form the game ended. It could be argued that, in some situations, the player that destroys the opponent will have a higher payoff once the war is over. Nevertheless, here I am interested in the case where both players are strictly worse when the game is no longer played. A justification for this is the one given above: once the rival is eliminated, a player might lose an important part of the internal support and be relegated to a worse position.

In order to simplify the algebra and obtain closed form solutions, I will assume that, if at least one of the players plays tough, the probability to keep playing the game one more period is $\mu \in (0,1)$, whereas if both played soft, they will play for at least another period with probability 1[44]. Also, I will assume that the expected utility of the players once the game is over is 0.

If players do not consider that their actions have an impact on the probability of keep playing the game, the unique Subgame Perfect Nash Equilibrium requires that both players use their strictly dominant action of playing tough. Nonetheless, when they realize that playing soft may increase future expected payoffs by increasing the expected number of repetitions of the stage game, this may become an optimal action. To see it, I will prove that there exists an equilibrium of the game where both players play soft as long as everyone played soft before. In this equilibrium, this strategy will be optimal as long as its expected discounted payoff exceeds the one of the best possible deviation.

[44] Setting this probability to be 1 is without loss of generality since, as we discussed in the introduction, if there is an exogenous constant probability to stop repeating the game, this can be incorporated in the discount factor.

Since this deviation consists in playing always though, the condition that needs to be satisfied can be stated as follows:

$$\frac{3}{1-\delta} \geq 5 + \delta\mu \frac{4}{1-\delta\mu} \qquad (11)$$

The left hand side of expression (11) represents the expected discounted payoffs of the equilibrium strategy when the rival is also employing it and the right hand side is the analogous when a player decides to deviate and play tough thereafter.

Rearranging terms, inequality (11) will be satisfied as long as:

$$\mu \leq \frac{5\delta - 2}{\delta(2-\delta)} \qquad (12)$$

A necessary condition for the right hand side of (12) to be positive (so that μ can be a proper probability and the equilibrium can exist) is that $\delta \geq 0{,}4$. Thus, when the discount factor is large enough and the probability of repeating the game when at least one player plays tough is below the bound in (12), there is a Subgame Perfect Nash Equilibrium where both players take the strictly dominated action of playing soft. ♦

The previous example is only a simple model to illustrate the seemingly paradox that enemies might prefer to play soft against each other (and obtaining lower current payoffs in the process) to avoid the annihilation of the enemy. The existence of these *perverse* incentives imply that, under certain conditions, less radical positions may, in fact, have a higher motivation to eradicate the enemy.

The Cold War between the USA and the Soviet Union is a prominent but not the unique example of a situation fitting the characteristics above discussed. The existence of terrorist groups can increase the support to other terrorists groups and even democratic parties that represent the ideological ideas more opposed to the ones of the terrorists. At the same time, the larger the support these parties obtain, the stronger the terrorist group becomes. However, if either of them disappeared, the less radical supporters of both groups would switch to less polarized positions.

3.6 Conclusions

Repeated Games literature has achieved a number of important results regarding interactions that are repeated over time. Among them, the folk theorems are specially celebrated because they are able to fully characterize the set of payoff vectors that can be sustained in equilibrium when the game is infinitely repeated and the players are sufficiently patient. Nevertheless, an important implicit assumption of the models they use is that the probability that the game is repeated is completely exogenous and independent of the actions of the players. Henceforth, they rule out the possibility that players could influence this probability by any mean. Since this strong assumption is unlikely to hold in most applications of repeated games, we must check how it affects the results and what the consequences of approximating an interaction disregarding this effect are.

In this research project we showed that some of the classic results of the repeated games literature are no longer true when the probability of repeating a game depends on the payoffs. In particular, the repetition of a Nash Equilibrium of the stage game may no longer be an equilibrium of the repeated game and some payoff vectors above the minmax might not be possible to sustain at the same time others below it might be.

Nevertheless, the most disturbing result is that, in some cases, the predictions that a modeler might reach when the effect is disregarded may be diametrically opposed to the ones he would obtain if the effect was taken into account. The conclusion is that, before using a repeated game to model a particular interaction, we must carefully check whether the probability that the game is repeated is absolutely exogenous. If the answer is no, then we must assess how, in this particular case, it might affect the behavior of the players.

Another important finding is that, in situations where the probability of repeating a game depends on the payoffs, an altruism effect emerges that may behave as a force to cooperate. Since this cooperative force does not rely on punishments, it works even without the presence of any kind of monitoring. As a consequence, we are able to show that it is able to sustain cooperation in environments where there is no monitoring and, consequently, a punishment scheme would be inoperative.

This research has pointed out a concern that should be considered whenever we try to model a situation as a repeated game. However, there is still a significant amount of research to be done in order to improve the understanding of how a non-exogenous probability of repeating a game affects the results of the repeated games literature in general setups.

If the probability of repeating a game is a function of the actions of the players and, somehow, we are able to modify this function, among other applications, we could be able to construct mechanisms to make cooperation easier to sustain. In these situations, it could be in the own interest of the players (or a regulator) to alter the function that determines the probability of repeating a game. For instance, the members of a monetary union can establish some economic standards that, if they are not met by one of the members, this one would be expelled. This could enhance cooperation among its members in order to avoid the costs of losing one of them. It would be interesting to have a look at all the new possibilities that playing with this function opens to creative agents and in the field of Mechanism Design.

This work is just a first step in a territory that remains still mostly unexplored. There are lots of questions that have yet to find an answer. Nevertheless, it has already shown that, in some situations, a more general setup than the one usually considered by the Repeated Games literature might be required. Moreover, it has also pointed out a new force to achieve cooperation, not using punishments, but through fear to the end of the game. Another interesting result is that this new way to cooperate can be added to punishments in order to enhance cooperation in some games and it can even be applied to achieve it when punishments fail to do it.

References

[1] Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics* 4, 1236-1239.

[2] Copic, J. and A. Galeotti (2007). Awareness Equilibrium. Mimeo.

[3] Della Vigna, S. and U. Malmendier (2004). Contract design and self-control: theory and evidence. *Quarterly Journal of Economics*, 119, 353-402.

[4] Della Vigna, S. and U. Malmendier (2006). Paying not to go to the gym. *American Economic Review*, 96, 694-719.

[5] Dekel, E., Lipman, B. and A. Rustichini (1998). Standard state-space models preclude unawareness. *Econometrica* 66, 159-173.

[6] Ellison, G. (2006). *Bounded rationality in Industrial Organization*, Richard Blundell, Whitney Newey, and Torsten Persson (eds.) Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press.

[7] Ely, J. and J. Valimaki (2002). A robust folk theorem for the Prisoners' Dilemma. *Journal of Economic Theory* 102, 84-105.

[8] Fagin, R. and J. Halpern (1988). Belief, awareness, and limited reasoning. *Artificial Intelligence* 34, 39-76.

[9] Feinberg, Y. (2004). Subjective reasoning - games with unawareness. Mimeo: Stanford Graduate School of Business.

[10] Feinberg, Y. (2005). Games with incomplete awareness. Mimeo: Stanford Graduate School of Business.

[11] Filiz, E. (2006). Incorporating awareness into Contract Theory". Mimeo: Columbia University.

[12] Fudenberg, D., D. Levine, and E. Maskin (1994). The Folk Theorem in repeated games with imperfect public information. *Econometrica* 62, 997-1040.

[13] Fudenberg, D. and E. Maskin (1986). The Folk Theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533-556.

[14] Fudenberg, D. and J. Tirole (1991). *Game Theory*, MIT Press, Cambridge.

[15] Fudenberg, D. and M. Villas-Boas (2007). *Behavior-based price discrimination and customer recognition*. Economics and Information Systems ed. T. Hendershott, Elsevier, Oxford.

[16] Gabaix, X. and D. Laibson (2005). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, forthcoming.

[17] Halpern, J. (2001). Alternative semantics for unawareness. *Games and Economic Behavior* 37, 321-339.

[18] Halpern, J. and L. Rego (2006). Extensive games with possibly unaware players. *Conference on Autonomous Agents and Multiagent Systems*, forthcoming.

[19] Harsanyi, J. (1967-68). Game with incomplete information played by Bayesian players. *Management Science* 14: 159-82, 320-34, 486-502.

[20] Heifetz, A., Meier, M. and B. C. Schipper (2006a). Interactive unawareness. *Journal of Economic Theory* 130, 78-94.

[21] Heifetz, A., Meier, M. and B. C. Schipper (2006b). Unawareness, beliefs and games. Unpublished manuscript.

[22] Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the Finitely Repeated Prisoners' Dilemma. *Journal of Economic Theory* 27(2), 245-252.

[23] Li, J. (2006). Information structures with unawareness. Mimeo: Duke University.

[24] Mas-Colell, A., Whinston M. D. and J. R. Green (1995). *Microeconomic Theory*, Oxford University Press.

[25] Matsushima, H. (1990). On the theory of repeated games with private information: Part I. *Economics Letters* 35, 253-256.

[26] Matsushima, H. (2004). Repeated games with private monitoring: two players. *Econometrica* 72, 823-852.

[27] Modica, S. and A. Rustichini (1999). Unawareness and partitional information structures. *Games and Economic Behavior* 27, 265-298.

[28] Modica, S. and A. Rustichini (1994). Awareness and partitional information structures. *Theory and Decision* 37, 107-124.

[29] Neyman, A. (1999). Cooperation in repeated games when the number of stages is not commonly known. *Econometrica* 67, 45-64.

[30] Osborne, M. J. and A. Rubinstein (1994). *A course in Game Theory*, MIT Press, Cambridge, Mass.

[31] Ozbay, E. Y. (2006). Unawareness and strategic announcements in games with uncertainty. Mimeo: New York University.

[32] Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029-1050.

[33] Piccione, M. (2002). The repeated Prisoners' Dilemma with imperfect private monitoring. *Journal of Economic Theory* 102, 70-83.

[34] Radner, R. (1986). Can bounded rationality resolve the Prisoners' Dilemma? *Contributions to Mathematical Economics*, edited by A. Mas-Colell and W. Hildenbrand, Amsterdam: North-Holland, 387-399.

[35] Taylor, C.R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, 34, 223-246.

[36] Vega-Redondo, F. (1994). Bayesian boundedly rational agents play the Finitely Repeated Prisioners' Dilemma. *Theory and Decision*, 36:2, 187-206.