Columbia University

*Department of Economics
Discussion Paper Series*

**Recursive Structure and Equilibria in Games
with Private Monitoring**

Massimiliano Amarante

Discussion Paper #:0102-48

***Department of Economics
Columbia University
New York, NY 10027***

April 2002

Recursive Structure and Equilibria in Games with Private Monitoring

Massimiliano Amarante*
Columbia University†

Revised: 2001

Abstract

In each stage of a repeated game with private monitoring, the players receive payoffs and privately observe signals which depend on the players' actions and the state of world. I show that, contrary to a widely held belief, such games admit a recursive structure. More precisely, I construct a representation of the original sequential problem as a sequence of static games with incomplete information. This establishes the ground for a characterization of strategies and, hence, of behavior in interactive-decision settings where private information is present. Finally, the representation is used to give a recursive characterization of the equilibrium payoff set, by means of a multi-player generalization of dynamic programming.

*I am very grateful to In-Koo Cho, Larry Epstein, Denis Gromb, Stephen Morris, Paolo Siconolfi, Lones Smith and Max Stinchcombe for several insights and suggestions. Also, I wish to thank the participants to the seminars at MEDS, NYU, Columbia University, Caltech, UCLA, University of Rochester, University of Texas-Austin, Northwestern Summer Microeconomics Conference 98, Summer in Tel Aviv 98, and NASM98.

†Columbia University, Dept of Economics, 1019 IAB, 420W 118th St., NY, NY 10027.
Email: ma734@columbia.edu.

1 Introduction

In each stage of a repeated game with imperfect private monitoring, the players receive payoffs and privately observe signals which depend on the players' actions and the state of the world. This class of games has recently received renewed attention from both an applied perspective and a more abstract one. For example, in the context of oligopolistic markets a question naturally arises about the possibility of collusion among competing firms which observe neither the prices nor the quantities of their opponents but only their own (Stigler [20]). When this is the only available information and consumers' demand is stochastic, each firm may attribute a fall in its own sales either to a low level of consumers' demand or to an unobserved price cut by an opponent.

In the perfect information case, where each firm's decision (a price in a Bertrand competition, a quantity in a Cournot) is commonly observed, collusion among firms can be (noncooperatively) enforced by each firm's threat to revert to a lower price or to a higher quantity if others do not act according to the collusive plan. However, when firms' decisions are not commonly observed, the use of strategies of this form may lead to undesirable punishments. In fact, the decision of a firm to revert, say, to a lower price can be based only on a low level of observed sales. Since this may occur even when all other firms adhered to the collusive plan, such a behavior would lead to a decrease in everyone's profit, and may even jeopardize any possibility of collusion at all.

In the case of homogeneous products – where the price of the good serves as a publicly observed signal about the competitors' quantities – the problem has been successfully investigated by Green and Porter [10] and Abreu, Pearce and Stacchetti [1], who have shown that collusive behavior may emerge as an equilibrium, and have characterized such a behavior in terms of the commonly observed price. Moreover, Fudenberg, Levine and Maskin established that a Perfect Folk-Theorem holds for this class of games (any feasible and individually rational payoff can be obtained as an outcome of a perfect equilibrium as the discount factor tends to one). Hence, for large discount factors, collusive behavior is neither less profitable nor less feasible than in the perfect information case.

However, the question of whether collusion can be sustained has gone unanswered when no such public signal is available. Still, the presence of private information seems to be a salient feature of many oligopolistic markets.

The same widespread presence of private information makes many other economic environments suitable of analysis within the framework of games with private monitoring. Examples range from various types of agency problems to bargaining models to models with information lags about the opponents' actions.

From a more theoretical point of view, the model of games with private monitoring appears as the natural laboratory for jointly analyzing two of the main problems economic theory has focused on during the last decades: the behavior of individuals who interact each other and recognize that others' decisions might be based on information unavailable to them, and the behavior of individuals

who are engaged in long-term relationships.

In spite of the economic interest in this class of problems, a general theory has not yet appeared. Several papers (Bhaskar and van Damme [5], Compte [6], Kandori and Matsushima [12], Matsushima [14], Sekiguchi [19]) have argued that the main difficulty for the construction of such a theory relies in that, in contrast to the public information case, games with private monitoring lack of a recursive structure. By this is meant the possibility of representing a problem involving maximization over sequences as a sequence of static maximization problems, the main advantage being in that it leads to a transparent characterization of the decision maker's behavior (e.g., Stokey and Lucas [21], in a decision theoretic framework; Fudenberg, Holmstrom and Milgrom [7], in a contracting framework). In interactive-decision settings, the achievement of a recursive structure – and, consequently, of appropriate generalizations of dynamic programming – has led to comprehensive theories both in the perfect (Fudenberg and Maskin [9]) and in the imperfect public information case (Abreu, Pearce and Stacchetti [2]; Fudenberg, Levine and Maskin [8]) as well as in the case of stochastic games (Mertens and Parthasarathy [16]).

It is natural, then, to ask if such a structure is irretrievably lost when private information is present.

The wide range of economic situations that naturally lend themselves to the framework of games with private monitoring has motivated the importance of understanding such a class of games. With this motivation in mind, the question about the recursive structure is primarily a question about how to look at such problems. For if a recursive structure were not achievable, that would mean that we have to look for new concepts, since anything that is known for other classes of repeated games comes from (or can be formulated in terms of) a recursive structure. In fact – under the presumption that a recursive structure does not exist – the literature devoted to the problem has mainly been confined either to two-period games or to quite specific questions.

In this paper, a recursive structure for games with imperfect private monitoring is provided. The representation that arises from this finding (the sequence of static games above) leads, perhaps not surprisingly, to a close connection to games of incomplete information,¹ thus establishing the ground for a characterization of strategies and, hence, of behavior in interactive-decision settings where private information is present. Finally, the representation is used to give a recursive characterization of the equilibrium payoff set, by means of a multi-player extension of dynamic programming.

Related literature. In recent years, several papers have been devoted to repeated games with imperfect private monitoring. Absent a theory, the main goal

¹This should not be surprising as Harsanyi ([11]) modeled games of incomplete information as games with imperfect private monitoring. In his model, players receive private signals about the move of the player “Nature”.

has been to devise a series of examples that could cast light on the features distinguishing this class of games from the more understood ones. Within specific contexts, questions have been asked about particular classes of strategies (Compte [6]), about the possibility of attaining efficiency (Ayoagi [4], Bhaskar and van Damme [5], Sekiguchi [19]), about the role of communication when private information is present (Kandori and Matsushima [12]) as well as about the classical theme of repeated games that of “repetition leads to cooperation” (Matsushima [14], Mailath and Morris [13]). All these papers, with the exception of [13], have argued that a recursive structure does not exist (see section 3.2), and consequently assert the impossibility of extending dynamic programming methods to the analysis of these games.

In a different vein, Mertens (Mertens [15]) suggested that a general model of repeated games might display a recursive structure, by proposing a class of “Entrance Laws”² (a class of static games, in the terminology used above) for the repeated game. He also gave a detailed analysis for the two-person zero-sum case.

Summary of the Argument. In the present paper, the problem is approached by introducing a family of sequences of approximating games, whose limits are the original game. The properties of such a family lead, in a constructive way, to a space of entrance laws for general n -person nonzero-sum repeated games with imperfect private monitoring. Also, it is sketched how such a construction can be extended to more general games. An implication of this procedure is a recursive characterization of the equilibrium payoff set.

The argument developed in the paper can be briefly summarized as follows.

1. **State variable:** In a sequential decision problem, an obvious state variable is the vector of signals the decision-maker has received up to the moment in which he makes his decision. Similarly, an obvious state variable for a repeated game is the I -tuple whose i -th component is player i ’s vector of signals, and I is the number of players. This is called a “state of the world”.
2. **Strategic-form continuation games:** A strategic-form continuation game (a game with continuation payoffs) is associated to each state of the world.
3. **Games with Incomplete Information:** Given a state of the world, each player has, generally speaking, only partial information about the true continuation game.³ This information is summarized by a conditional probability distribution over the set of possible continuation games. The set of all possible continuation games at a given state of the world, along with such conditionals defines a game of incomplete information.
4. **Common Knowledge:** It will be shown that (“at an equilibrium” of the repeated game) this game of incomplete information is consistent, and by a theorem of Mertens and Zamir ([17], Cor. 4.7 and Thm. 5.3) it is common knowledge among the players.

²This terminology is due to Mertens [15].

³This is so because the true continuation game depends also on his opponents’ signals, which are, generally speaking, unknown to him.

5. Method of Successive Approximations: This construction is then used to define sequences whose elements are finitely repeated games (called auxiliary games) so that each sequence converges to the original repeated game. The properties of these sequences are then investigated and used to construct an appropriate space of entrance laws for the repeated game.
6. The Equilibrium Payoff Set: From the procedure outlined so far it will follow that the equilibrium payoff set of the repeated game is the infinite intersection of a family of equilibrium payoff sets, each associated with an auxiliary game.

Section 2 describes the basic model of discounted repeated games with imperfect private monitoring. After a brief summary of the main results achieved for the imperfect public information case and a discussion about the claim of nonexistence of a recursive structure, section 3 introduces, rather informally, some of the concepts the present analysis is built on. Section 4 contains the main definitions and develops the approximating procedure. There, a representation of the space of entrance laws is given. Section 5 is devoted to the study of the equilibrium payoff set. Section 6 discusses some extensions, and Section 7 concludes.

2 The Model

In the stage game both the payoffs and the signals received by each player depend on the realization of a random variable in a set Ω , which is called the basic domain of uncertainty. An action profile in the stage game determines a lottery over Ω ; then, an element in Ω is selected by such a lottery, and both payoffs and signals are assigned.

More precisely, the stage game G has the following features:

- (i) A finite set of players I ;
- (ii) Finite action spaces A_i , for each player $i \in I$. An action profile is an element $a \in A = \times_i A_i$;
- (iii) The basic domain of uncertainty is $\Omega \subset R^d$ ($d < \infty$), and is assumed to be compact. (Ω, B_Ω) is a measurable space, with B_Ω referring to the Borel sets of Ω .
- (iv) A profile $a \in A$ induces a measure $\mu(a)$ on Ω . We assume that $\mu(a)$ is absolutely continuous with respect to some measure μ on R^d , for any $a \in A$.
- (v) When profile a is played and ω realized, player i gets a payoff $g_i(\omega, a)$. Such a payoff is nonnegative (WLOG) and bounded from above by some $c \in R$.

A mixed action m_i for player i is a measure over A_i . The set of all mixed actions for player i is denoted by $\Delta(A_i)$, and $\Delta(A)$ stands for $\times_i \Delta(A_i)$, whose generic element is denoted by m . The measures $\mu(a)$'s are extended to the Borel field of $\Delta(A)$ in the obvious way, and $g_i(\omega, m)$ defined correspondingly. When profile m is played, player i 's expected payoff is denoted by $E_\Omega [g_i(\omega, m)]$.

In the repeated game G^∞ , each player maximizes the discounted sum of his per-period payoffs; $\delta \in (0, 1)$ denotes their common discount factor. The signalling structure in G^∞ is specified as follows:

(vi) For each player i , a (deterministic, for simplicity) signalling function $\theta_i : \Omega \rightarrow \Theta_i$ is defined by $\theta_i(\omega) = \theta_i \in \Theta_i$. Such functions are onto (WLOG), and Θ_i is called the space of signals for player i .

To avoid technicalities, we assume for most of the paper that Θ_i is a finite set, $\forall i \in I$. The main advantage for doing so is that the assumption guarantees that the auxiliary games of Section 4 have a finite normal-form. Hence, we can invoke Nash theorem to ensure existence of equilibria (see, however, Section ?? about removing this assumption).

(vii) The functions $\theta_i(\cdot)$'s are assumed to be common knowledge among the players.

The realization ω which occurs at the end of stage $t - 1$ is denoted by ω_t . Thus, player i 's information at the stage t -th is given by $\theta_i^t = \{\theta_{i,1}, \dots, \theta_{i,t}\} \in \Theta_i^t = \times_{s=1}^t \Theta_i$, where $\theta_{i,t} = \theta_i(\omega_t)$. A strategy for player i in G^∞ is $\{\sigma_i^t\}_{t=0}^\infty$, $\sigma_i^t : \Theta_i^t \rightarrow \Delta(A_i)$.

An object that will play a role in this paper is the signalling function for the game. This is the function $\theta : \Omega \rightarrow \Theta = \times_{i \in I} \Theta_i$ defined by $\theta(\cdot) = (\theta_1(\cdot), \dots, \theta_I(\cdot))$. We make Θ into a measurable space so that $\theta(\cdot)$ is measurable. By $P(\cdot \mid m)$ we denote the pushforward $\theta_* \mu(m) = \mu(m) \circ \theta^{-1}$ of $\mu(m)$, and by $P^t(\cdot \mid \cdot)$ the product measure on $\Theta^t = \times_{i \in I} \Theta_i^t$. Finally, we denote by P the pushforward of μ , and by P^t the corresponding product measure on Θ^t .

Note that, in order to ensure a high degree of flexibility, we have not specified the nature of the spaces Ω and Θ that appear in the above description. Consequently, several classes of games can be casted within the present framework by means of a simple reformulation of the relevant variables. For example, in the present study we are mainly interested in the type of problems described in the introduction. In those settings, an important source of private information comes from a players knowledge of his own actions. Clearly, such a knowledge must be included in the space Θ_i of player i 's signals. Nevertheless, sometimes the problem already contains, among its data, spaces Ω and Θ , but player i 's privately observed actions are not included in the Θ_i 's. In such a case, a way to correctly reformulate the problem consists in defining new spaces $\tilde{\Omega} = \Omega \times A$ and $\tilde{\Theta} = \Theta \times A$, and in writing player i 's signalling function as $\tilde{\theta}_i = (\theta_i, a)$.

A general example of this model is a repeated game where, at the end of each stage, players observe neither ω , nor the actions played, nor their own payoffs, but only some imperfectly correlated signal in some space Θ_i . A case of particular interest is that of observable payoffs: θ_i is the payoff received by player i after a stage is played, and this is the only information available to him. An example of this specification is Stigler's model of oligopoly described in the introduction. Clearly, the model contains as special cases (a) the imperfect public information case, identified by $\theta_i(\omega) = \omega, \forall \omega \in \Omega, \forall i \in I$; (b) the perfect information case, where $\Omega = \Delta(A)$ and $\theta_i(m) = m, \forall i \in I$.

A simple example that captures many of the typical features of repeated games with imperfect private monitoring is as follows. Let G be a 2×2 normal form game whose payoffs depend on the realization of a random variable $\omega \in \Omega$, and assume that $\Omega = \{(L, L), (L, H), (H, L), (H, H)\}$. After G is played and ω is realized, player 1 is told the first component of ω , and player 2 the second. Therefore, $\Theta_i = \{L, H\}$, $i = 1, 2$, and after the signal is received, player i has a nontrivial partition on $\Omega = \Theta = \Theta_1 \times \Theta_2$.

3 Preliminaries

3.1 Repeated Games with Public Information

This section contains a quick, and by no means comprehensive, overview of the results achieved in the public information case. The main purpose of this part is instrumental: by discussing the public information case both some concepts that will be later developed are introduced and it will be explained why it has been suggested that a recursive structure might not be achievable in the private monitoring case.

Repeated games with imperfect public information are characterized by the specification $\theta_i(\omega) = \omega, \forall i \in I$ (that is the realization ω is publicly observed; section 2 above). This class of games has been extensively studied by Abreu, Pearce, and Stacchetti ([2]; hereafter APS) and Fudenberg, Levine and Maskin ([8]).

Since the purpose of this section is only expository, the reader may assume – in order to avoid some qualifications – that ω is the only information available to each player. In such a case, we do not need to distinguish between pure, mixed and behavior strategies. Alternatively, the reader can think of the games in the text as decision problems, and cast the explanation in the standard dynamic programming framework.

In [2], it was shown that any equilibrium payoff vector $v \in V$ of the repeated game, can be written as

$$v = (1 - \delta)E_{\Omega} [g(\omega, \sigma^0)] + \delta \int_{\Omega} q(\omega) d\mu(\sigma^0)$$

where, σ^0 is the time-0 component of an equilibrium profile σ that generates v , and $q = (q_1, \dots, q_I)$ is a measurable function $q : \Omega \rightarrow V$.

A flavor of logic underlying the proof may be obtained from the following considerations. In this framework, after G is played, ω is realized and observed by all players. We have, then, a subgame starting at such a realization of ω . A strategy for player i maps $\omega^t \mapsto \Delta(A_i)$. Under the assumption of full support of μ ([2]), every Nash equilibrium is subgame perfect in the subgames thus defined. As a consequence, if σ is a Nash equilibrium of G^{∞} , we immediately obtain – as of time 0 – the above decomposition as soon as we assign to $q(\omega)$ an equilibrium value for the subgame starting at ω .

Now consider the restriction of the given equilibrium profile σ to the subgame starting at ω . Since the initial state in G^∞ is free, we can start by playing this profile at time 0, and still get an equilibrium. That is, such a restriction is an equilibrium profile of G^∞ as well. Hence, q has range in V . In other words, $q(\omega)$ is an equilibrium payoff vector of G^∞ as well, for any $\omega \in \Omega$. Combining this fact with the previous observation, it follows that each $q(\omega)$ can be decomposed in the same fashion as above.

In essence, this is the recursive structure. The sequential game is represented by a static one. This, in turn, can be decomposed into a current game and a continuation game, and the continuation game is subject to the same decomposition. Moreover (see below), an equilibrium profile in the repeated game is recovered as a sequence of static equilibria, each associated to a static game.

This results suggests (as already noted by APS) a representation of the game with imperfect public information as a stochastic game. To see this, suppose, to fix ideas, that $\Omega = \{\omega^1, \omega^2, \omega^3, \omega^4\}$. Then, the above decomposition permits us to associate to each ω^k , $k = 1, 2, 3, 4$, a continuation game, $CG(\omega^k, f^k)$, where payoffs are defined by

$$q_i(\omega^k) = (1 - \delta)E_\Omega [g_i(\omega', \sigma)] + \delta \int_\Omega q_i^k(\omega') d\mu(\sigma) \quad (1)$$

and $range(q^k) \subseteq V$.

Then, the problem players face can be described according to the picture below

where the diagram means that players are told that they will play the game G at stage 0, and then they will enter the set $CG(\omega) = \{CG(\omega^k, q^k)\}$. After ω is realized, they are publicly informed of such a realization, and if, say, ω^1 is realized they will play the game $CG(\omega^1, q^1)$ in stage 1, etc.. Moreover, for each of the continuation games (auxiliary one-shot games, in the terminology of APS) we can perform the same decomposition, and so on.

This procedure leads us to recover the equilibrium behavioral strategies of the repeated game as sequences of Nash equilibria of the various continuation games we construct along the above decomposition.

3.2 The Problem of the Recursive Structure in Games with Private Monitoring

The claim of the nonexistence of a recursive structure in repeated games with imperfect private monitoring relies, essentially, on the inability of deterministically specifying the continuation play when private information is present. To see this, let us get back to the public information case and consider what happens in that setting. Recall that strategies are maps from the signals to the set of (mixed) actions, and fix an equilibrium profile σ of the repeated game. In that setting, after the first stage is played some ω is realized and publicly observed. It follows that beliefs of Bayesian players are concentrated on that ω and hence on the strategy profile that according to σ follows that ω . In other words, each player can infer from the realization ω the opponents' continuation strategies deterministically. When private information is present, this link is broken: player i 's continuation strategy depends on the signal he received, and this signal is unknown to player j . A consequence is that the *actual* continuation profile is not common knowledge among the players. This last fact has been interpreted as a suggestion that a recursive structure is not achievable in the presence of private information ([5], [12], [19]). The argument goes, essentially, as follows. The continuation play depends on the signals received by each player as well as on the conditional probabilities that each player has about the opponents' signals, and hence on their continuation strategies. Now suppose that we propose some profile σ , and we want to check whether or not it is an equilibrium for the repeated game. To do so, let us consider a deviation of player i from the proposed profile in the first stage. Then, while each player $j \neq i$ will compute his conditionals according to the given profile σ , player i will derive his conditionals by taking his own deviation into account. The implication hence derived is that, at such a point, not only the actual continuation profile is not common knowledge, but not even the family of conditionals is. This last circumstance, finally, seems to prevent any attempt to decide whether the continuation strategies conforming to the proposed profile σ may constitute an equilibrium at all, and hence the ability to check whether the hypothesized deviation is profitable or not. In other words, in the presence of private information what seems to be irremediably broken is the link between the incentive structure in the first period and that of subsequent periods.

A couple of considerations are now probably in order. First, it should be noted that common knowledge of an equilibrium profile is a necessary condition "locally", that is, it is required to hold at an equilibrium (Aumann and Brandenburger [3]). Second, like any game with private information, a careful analysis of the problem outlined above would require the consideration of a hierarchy of beliefs, which is not carried out above.

Before proceeding, it may help to restate the above argument by using the representation in terms of stochastic games given in the previous subsection. Suppose that each of two players can get one of two signals L and H . Then,

the picture given above becomes

The interpretation is, as before, that players play the component game G , and then they enter the set of games $G(\Theta) = \{CG(\theta) \mid \theta \in \Theta\}$, with the payoffs in $CG(\theta)$ being given by $(1 - \delta)E[g_i(\omega, \sigma_i^1, \sigma_j^1)] + \delta \int_{\Omega} f_i(\omega, \theta) d\mu(\sigma_i^1, \sigma_j^1)$. If,

say, $\omega = \theta = (L, H)$ is realized, then player 1 observes L and knows that the true game is either $CG((L, L))$ or $CG((L, H))$, while player 2 observes H and knows that it is either $CG((L, H))$ or $CG((H, H))$. We can see that the actual continuation game (the one associated with the actual realization of θ) is not common knowledge among the players.

The argument above can then be restated by saying that if we suppose that the decomposition into current and continuation game could actually be performed, it would lead to a continuation game that is not common knowledge among the players. This, in turn, would prevent us from checking whether the proposed profile was indeed an equilibrium.

Note, however, that the figure itself seems to suggest that we might think of the continuation game in a different way. In fact, given a family of conditional probability distributions $\{p(\theta_j \mid \theta_i)\}_{i \in I}$, the pair $(\{p(\theta_j \mid \theta_i)\}_{i \in I}, G(\Theta))$ defines a game of incomplete information where player 1 is told the row where the true continuation game lies, and player 2 the column

This is tantamount to saying that the continuation play is now determined stochastically. (Note, also, that this description contains the public information case as a special case).

In the next section, I will show that these considerations do hold water, and that this procedure does lead us to recover a recursive structure for games with private monitoring by considering the appropriate state space for the repeated game.

3.3 Games of Incomplete Information: a first check of consistency

A game of incomplete information is a triple $\left\{ \Gamma(\theta), (p_i)_{i \in I}, \Theta = \times_{i \in I} \Theta_i \right\}$, where Θ_i is the set of types for player i , $\{p_i\}$ is a family of probability distributions $p_i(\cdot \mid \theta_i)$ over Θ_{-i} , each representing player i 's beliefs over other players' types given that is type is θ_i , and $\forall \theta \in \Theta$, $\Gamma(\theta)$ is a normal form game. A strategy for player i is $\sigma_i : \Theta_i \rightarrow \Delta(A_i)$. In these games the notion of type accounts for all the information that a player has about the game, and that it is not common knowledge among the players.

A game of incomplete information is said to be consistent if there is a probability distribution $P(\cdot)$ over $\Theta = \times_{i \in I} \Theta_i$, such that the p_i 's are the conditionals computed from $P(\cdot)$ by Bayes' formula. If such a P exists, it can be given the meaning of a common prior ([17]). Then, from the viewpoint of the NE, the game is equivalent to an extensive form game where "Nature" moves first selecting a vector of types θ according to P , then each player is informed about his own type and the game is played (Harsanyi [11]). Mertens and Zamir ([17], Cor. 4.7) have shown that if such a consistent probability P exists, then it is common knowledge among the players. Hence ([17], Thm. 5.3), the above extensive form game is common knowledge as well.

For future reference, fix an equilibrium profile of the two-stage game in the example above. Such an equilibrium can be determined from the normal-form. Then, *at that equilibrium* that profile is common knowledge among the players ([3]), and its first component induces a probability distribution P^1 over the set of games in stage 1. After stage 0, each player receives his private signal and computes his conditional from P^1 . Hence, the game of incomplete information played in stage 1 is consistent by construction, and by the preceding it is common knowledge among the players.

4 The Recursive Structure

Before we address the question of existence of a recursive structure, it might help the reader if we summarize here the main notational conventions we will be using in the remaining of the paper. An equilibrium profile in G^∞ will be denoted by σ , and σ^k will refer to the k -th component of σ (σ^k is the strategy profile played in stage k). Occasionally, the same notation will be used to indicate equilibrium profiles in the auxiliary games, $AG(k)$, introduced below. No confusion will result. A continuation payoff functions at stage k is typically

denoted by f_k . N stands for the set of natural numbers, which indexes stages in G^∞ .

The goal of this section is to show that:

a) For any $k \in N$, any equilibrium payoff vector achievable starting from stage k in the repeated game G^∞ can be achieved as an equilibrium of a static game Γ^k . Γ^k is a game with continuation payoffs.

b) Any equilibrium strategy profile that attains such a payoff in Γ^k is the stage k -th component of the corresponding equilibrium profile in the repeated game.

c) For any $k \in N$, Γ^k and Γ^{k+1} are recursively linked, in the sense that the expected equilibrium continuation payoffs in Γ^k are the (expectation of the) equilibrium payoffs in Γ^{k+1} .

These three properties can be taken as a definition of recursive structure. They are the multi-player analog of the familiar dynamic programming representation of one-person decision problems. The family of all sequences $\{\Gamma^k\}_{k \in N}$ satisfying the above properties is called the space of entrance laws.

The idea underlying the determination of the family of sequences $\{\Gamma^k\}$'s can be intuitively explained as follows. Consider the requirements a), b) and c), and, for σ an equilibrium profile in the repeated game, denote by $v(\sigma) |_k$ the stage k -th expected equilibrium payoff corresponding to σ in G^∞ . In addition, denote by $v_k(\cdot, f_k)$ an equilibrium payoff of a (static) game Γ^k with continuation payoffs defined by some function f_k . Then, for any equilibrium profile σ in G^∞ , the requirements a) and b) are equivalent to a set of equations in the unknowns f_k 's like

$$\begin{aligned} v(\sigma) &= v_0(\sigma^0, f_0) \\ v(\sigma) |_1 &= v_1(\sigma^1, f_1) \\ &\dots \\ v(\sigma) |_k &= v_k(\sigma^k, f_k) \\ &\dots \end{aligned} \tag{2}$$

We are concerned with the solutions of the system (2). To this end, we will be using a method of successive approximations. If we focus on the first equation only, the method consists in determining its solution, f_0 , as the uniform limit of a sequence of functions $\{f_0^n\}_{n \in N}$. Roughly speaking, the method used here consists in constructing, for each equilibrium profile σ in G^∞ , a family of sequences $\{f_0^n\}, \dots, \{f_k^n\}, \dots$ (one for each equation) so that the recursive relation (requirement c)) is satisfied at each step of the construction, $f_k^n \rightarrow f_k$ (uniformly) as $n \rightarrow \infty$ for each $k \in N$, and the f_k 's solve (2). Clearly, we must show that it is so for each equilibrium σ in G^∞ .

Note, however, that this way of describing the problem is perhaps suggestive but rather vague. We have said nothing about the domains of the f_k 's, and hence about the form the games Γ^k 's. In fact, these are themselves unknowns to be determined.

Above all, the main difficulty to carry out such a program is that, in order to solve the above equations, we need to specify at which equilibrium profile we are solving. Since, clearly, to determine the equilibria is part of the problem, we need a procedure that generates the equilibrium profile in the equations at the same time it generates the above limits.

Such a procedure is based on the introduction of sequences of finitely repeated games, $\{AG(k)\}$ – whose typical element is called an auxiliary game of order k . Their definition requires the notion of “state of the repeated game” and that of “type” of a player. The latter arises immediately in the context we are interested in, and conforms to that introduced by Harsanyi for games with incomplete information.

Definition 1 A type θ_i^k for player i at $k \in N$ is the stream of signals he has received up to k : $\theta_i^k = \{\theta_{i,1}, \dots, \theta_{i,k}\} \in \Theta_i^k$.

Definition 2 A state of the world for the game G^∞ at stage $k \in N$ is an element $\theta^k \in \Theta^k = \times_{i \in I} \Theta_i^k$.

That is, a state of the world at stage k is an I -tuple of types.

The next definition is that of auxiliary game of order k .

First, given a $\mu \times P^k$ -measurable function $f_k : \Omega \times \Theta^k \rightarrow [0, c]^I$, associate to each state of the world $\theta^k \in \Theta^k = \times_{i \in I} \Theta_i^k$ a strategic-form continuation game, $CG_{f_k}(\theta^k)$, by

$$\theta^k \mapsto \left\{ I, A_i, u_i = (1 - \delta)E_\Omega [g_i(\omega, a)] + \delta \int_\Omega f_{k,i}(\omega, \theta^k) d\mu(a) \right\}$$

Then,

Definition 3 Given a repeated game G^∞ , an auxiliary game of order k , $AG_{f_{k-1}}(k)$, is a k -stage game with the following features:

i) In the first $k-1$ stages the payoffs are determined by the stage game G (payoffs in G are multiplied by $1 - \delta$);

ii) At the end of each stage l , $l < k-1$, player i receives a private signal $\theta_{i,l+1} = \theta_i(\omega_{l+1})$, $l = 0, \dots, k-2$;

iii) the last stage ($l = k-1$) is defined by a family of games $G_{f_{k-1}} = \{CG(\theta^{k-1})\}_{\theta^{k-1} \in \Theta^{k-1}}$, and if $\bar{\theta}^{k-1}$ is the true state of the world, payoffs in the last stage are determined according to $CG(\bar{\theta}^{k-1})$.

In each $AG_{f_{k-1}}(k)$, each player has a discount factor δ , and maximizes the discounted sum of his per-period payoffs. Note that when $\theta_{i,l}$ is only imperfectly related to ω_l , players have only partial information about the true game in the $(k-1)$ -th stage summarized by a conditional probability $P^k(\theta^{k-1} \mid \theta_i^{k-1} = \bar{\theta}_i^{k-1})$.

In the sequel, when referring to an $AG_{f_{k-1}}(k)$, we will occasionally drop either the subscript or the reference to k . No confusion should result. In the simple game with two players and two signals described above, an auxiliary game of order 2 is like the one depicted at the end of section 3.2. With reference to same example, an auxiliary game of order 3 is depicted in the following figure

The main motivation for the introduction of the sequences $\{AG(k)\}$'s relies on the following transparent property::

Fact Consider the set of all I -player δ -discounted (finitely or infinitely) repeated games with bounded payoffs whose stage game has the same action spaces A_i . It is a normed linear space with

$$\sup_{i \in I, \sigma \in \Sigma} |v_i(\sigma)|$$

where $v_i(\sigma)$ is player i 's payoff in the repeated game generated by a profile σ .

It is immediate that every sequence $\{AG(k)\}$ converges, as $k \rightarrow \infty$, uniformly to G^∞ in this norm.

The properties stated in the next Lemma and in the Corollary thereafter are the first step to show that requirements a), b) and c) are satisfied.

Lemma 4 *For any game $AG_{f_k}(k+1)$ and any equilibrium profile $s = (s^0, \dots, s^k)$ of $AG_{f_k}(k+1)$, there exists a game $AG_{f_{k-1}}(k)$ which generates the same equilibrium payoff vector as s in $AG_{f_k}(k+1)$. Moreover, such a payoff is generated*

in $AG_{f_{k-1}}(k)$ by the equilibrium profile which is the restriction of s to the first k stages.

Proof. Consider an $AG_{f_k}(k+1)$, and its associated normal-form. By Nash, equilibria exist. Let s be an equilibrium profile. To ease notation, set $dZ^t = dP^t(\theta^t \mid s^{t-1}, s^{t-2}, \dots, s^0)$. Player i 's payoff corresponding to s is given by

$$(1-\delta) \sum_{t=0}^{k-1} \delta^t \int_{\Theta_i^t} \int_{\Theta_{-i}^t} E_{\Omega} [g_i(\omega, s(\theta_i^t, \theta_{-i}^t))] dZ^t + \\ \delta^k \int_{\Theta_i^k} \int_{\Theta_{-i}^k} \left[(1-\delta) E_{\Omega} [g_i(\omega, (s(\theta^k)))] + \delta \int_{\Omega} f_{i,k}(\omega, \theta^k) d\mu(s(\theta^k)) \right] dZ^k$$

Let $F_k = \{f_k \mid f_k : \Omega \times \Theta^k \rightarrow [0, c]^I\}$. To prove the Lemma we need only to show that there exists a function in $\hat{f}_{k-1} \in F_{k-1}$ so that the above payoff can be expressed as a payoff in $AG_{\hat{f}_{k-1}}(k)$. In fact, since – from the normal-form – the equilibrium condition (player i 's incentive constraint) is satisfied, such a payoff would necessarily be an equilibrium payoff in $AG_{\hat{f}_{k-1}}(k)$ (this is immediately checked by applying the transformation below to both sides of the incentive constraint). Noting that the expression in square brackets defines a function

$$\tilde{f}_i(\theta^k) = \left[(1-\delta) E_{\Omega} [g_i(\omega, (s(\theta^k)))] + \delta \int_{\Omega} f_{i,k}(\omega, \theta^k) d\mu(s(\theta^k)) \right]$$

we can write the player i 's payoff corresponding to s in $AG_{f_k}(k+1)$ as

$$(1-\delta) \sum_{t=0}^{k-1} \delta^t \int_{\Theta_i^t} \int_{\Theta_{-i}^t} E_{\Omega} [g_i(\omega, s(\theta_i^t, \theta_{-i}^t))] dZ^t + \delta^k \int_{\Theta^k} \tilde{f}_i(\theta^k) dZ^k$$

First, note that clearly $\tilde{f}_i : \Theta^k \rightarrow R$ is P^k -summable and $\text{range}(\tilde{f}_i) \subseteq [0, c]$ ($\text{range}(g_i) \subseteq [0, c]$, along with $f_{i,k}(\cdot, \theta^k) : \Omega \rightarrow [0, c]$ and $\delta < 1$, imply $\text{range}(\tilde{f}_i) \subseteq [0, c]$). Now, given the $\tilde{f}_i : \Theta^k \rightarrow [0, c]$, we show that there exists a $\mu \times P^{k-1}$ -measurable function $\hat{f}_i : \Omega \times \Theta^{k-1} \rightarrow [0, c]$, such that

$$\int_{\Theta^{k-1}} \int_{\Omega} \hat{f}_i(\omega, \theta^{k-1}) d\mu(s^{k-1}(\theta^{k-1})) dZ^{k-1} = \int_{\Theta^k} \tilde{f}_i(\theta^k) dZ^k \quad (3)$$

Let $\pi_1 : \Omega \times \Theta^{k-1} \rightarrow \Omega$ be the projection on the first factor, and let $\pi_2 : \Omega \times \Theta^{k-1} \rightarrow \Theta^{k-1}$ be the projection on the second. Let $h : \Omega \times \Theta^{k-1} \rightarrow \Theta^k$ be defined by $h = (\theta \circ \pi_1) \otimes \pi_2$, that is $h(\omega, \theta^{k-1}) = (\theta \circ \pi_1(\omega, \theta^{k-1}), \pi_2(\omega, \theta^{k-1}))$.

It is immediate that h is measurable (being the tensor product of measurable mappings) and that so is $\hat{f}_i : \Omega \times \Theta^{k-1} \rightarrow [0, c]$ defined by $\hat{f}_i = \tilde{f}_i \circ h$. Clearly, such an $\hat{f}_i$ satisfies (3).

By using (3), we can now write the payoff as

$$(1 - \delta) \sum_{t=0}^{k-1} \delta^t \int \int_{\Theta_i^t \Theta_{-i}^t} E_{\Omega} [g_i(\omega, s(\theta_i^t, \theta_{-i}^t))] dZ^t + \delta^k \int_{\Theta^{k-1}} \int_{\Omega} \hat{f}_i(\omega, \theta^{k-1}) d\mu(s^{k-1}(\theta^{k-1})) dZ^{k-1}$$

or

$$(1 - \delta) \sum_{t=0}^{k-2} \delta^t \int \int_{\Theta_i^t \Theta_{-i}^t} E_{\Omega} [g_i(\omega, s(\theta_i^t, \theta_{-i}^t))] dZ^t + \delta^{k-1} \int \int_{\Theta_i^k \Theta_{-i}^k} \left[(1 - \delta) E_{\Omega} [g_i(\omega, (s(\theta^{k-1})))] + \delta \int_{\Omega} \hat{f}_i(\omega, \theta^{k-1}) d\mu(s(\theta^{k-1})) \right] dZ^{k-1}$$

■

For future reference, it is useful to make here the following observation. Set $F_{k-1}^0 = F_{k-1} = \{f_{k-1} \mid f_{k-1} : \Omega \times \Theta^{k-1} \rightarrow [0, c]^I \text{ and } f_{k-1} \text{ is } \mu \times P^{k-1} \text{ measurable}\}$, and notice that the proof of the Lemma defines a correspondence $I_k : F_k \rightarrow F_{k-1}$. Obviously, $I_k(F_k) = F_{k-1}^1 \subseteq F_{k-1}^0$. If we start from the set of possible auxiliary games of order $k+2$, we can then apply the procedure twice and obtain the set $F_{k-1}^2 = I_k(I_{k+1}(F_{k+1}))$. In a similar fashion, we can define the sets $F_{k-1}^n, \forall n \in N$. Then, it follows immediately (by induction) that $F_{k-1}^{n+1} \subseteq F_{k-1}^n, \forall n \in N$.

As an immediate consequence of the lemma, we obtain the following corollary.

Corollary 5 *$\forall k \in N$, and any NE profile s of $AG_{f_k}(k+1)$, there exists a sequence $\{(\Gamma^t, P^t)\}_{t=0}^{k-1}$ of games with incomplete information, such that the component $\{s_i^t(\theta_i^t)\}_{i \in I}$ of s is a NE of (Γ^t, P^t) . In addition, $\{(\Gamma^t, P^t)\}_{t=0}^{k-1}$ is defined in a recursive way and the equilibrium payoff generated by s is a NE payoff of $(\Gamma^0, P^0) = \Gamma^0$, achieved by the equilibrium (in Γ^0) action profile $\{s_i^0\}_{i \in I}$.*

Proof. Let $(s^0, s^1, \dots, s^{k-1})$ be a Nash equilibrium in $AG_{f_k}(k+1)$. Recall that the last stage in $AG_{f_k}(k+1)$ is defined by a family of strategic-form games $G_{f_k} = \{CG(\theta^k)\}_{\theta^k \in \Theta^k}$. Since s is a NE, it induces a common prior P^k over this family of games. When players enter the last stage – the one defined by the family of strategic-form games $\{CG(\theta^k)\}_{\theta^k \in \Theta^k}$, each has a conditional probability that depends on the signals he has received. At s , all this conditionals come from the common prior P^k . It follows that the game with incomplete

information $\left\{ \left\{ CG(\theta^k) \right\}_{\theta^k \in \Theta^k}, P^k \right\}$ is common knowledge among the players. Clearly, s^{k-1} is a *NE* of such a game (by contradiction, using the fact that s is a *NE* for $AG_{f_k}(k+1)$ as established in the normal-form). Now, consider the problem the players face in $AG_{f_k}(k+1)$ at the beginning of stage $k-1$. Each player has, at that point, a conditional probability about the true state of the world at $k-1$, and – for the same reason as before – all these conditionals come from a common prior P^{k-1} . It is immediate to see that the last two stages in $AG_{f_k}(k+1)$ along with P^{k-1} define a game with incomplete information, and that (s^{k-1}, s^k) is a *NE* of such a game. Moreover, by using the same argument as in the preceding Lemma, the equilibrium payoff corresponding to (s^{k-1}, s^k) can be recovered as an equilibrium payoff of a game with incomplete information $\left\{ G_{\hat{f}_{k-1}} = \left\{ CG(\theta^{k-1}) \right\}_{\theta^{k-1} \in \Theta^{k-1}}, P^{k-1} \right\}$, and such a payoff is obtained by the equilibrium profile s^{k-1} (in $\left\{ G_{\hat{f}_{k-1}} = \left\{ CG(\theta^{k-1}) \right\}_{\theta^{k-1} \in \Theta^{k-1}}, P^{k-1} \right\}$). Finally, s^{k-1} and P^{k-1} induce the common prior P^k on $G_{f_k} = \left\{ CG(\theta^k) \right\}_{\theta^k \in \Theta^k}$, and s^k is an equilibrium of $\left\{ \left\{ CG(\theta^k) \right\}_{\theta^k \in \Theta^k}, P^k \right\}$. By considering the problem players face at stages $k-2$, $k-3$, etc., and proceeding in the same fashion, the statement follows. ■

The property expressed by Lemma 4 allows us to determine, at the same time, both the equilibrium profiles in the repeated game and, for each equilibrium, the functions that solve the equations (2).

Recall that $F_k = \left\{ f_k \mid f_k : \Omega \times \Theta^k \rightarrow [0, c]^I \right\}$. According to definition 3, any $f_k \in F_k$ defines an auxiliary game of order k , $AG_{f_k}(k)$. Such a game has a finite normal-form; hence (Nash), we can determine the equilibria of such a game. Lemma 4 then says that for each equilibrium we can associate to f_k a function $f_{k-1} \in F_{k-1}$. By applying the same procedure for all the f_k 's in F_k , it remains defined, as noted above, a correspondence $I_{k-1} : F_k \rightarrow F_{k-1}$.

We can now construct the successive approximations. First, consider all the possible auxiliary games of order 1, and let $AG_{f_0}(1)$, $f_0 \in F_0$, be one of those. In these games, the choice of the continuation function is restricted only by the feasibility condition. Denote by $\{\sigma^0(f_0)\}$ the set of Nash equilibrium profiles associated to f_0 in $AG_{f_0}(1)$. Now, consider all the possible $AG_{f_1}(2)$, $f_1 \in F_1$. Each $AG_{f_1}(2)$ has a finite normal form. Its equilibria are pairs $(\sigma^0, \sigma^1)(f_1)$. By Lemma 4, to each f_1 and each $(\sigma^0, \sigma^1)(f_1)$ we can associate a function $f_0^1 \in I_1(f_1) \subseteq F_0$ so that $\sigma^0(f_0^1)$ is an equilibrium of $AG_{f_0^1}(1)$. By repeating the procedure for each $f_1 \in F_1$, we determine a subset of all possible auxiliary games of order 1. Such a subset has the property that, for each game, both the equilibrium profiles and the continuation functions are consistent with equilibrium behavior in the second stage, when the continuation function are unrestricted in the second stage. Similarly, starting from all possible games

of order 3, we first determine a subset of the games of order 2 so that, for each game in the subset, equilibrium profiles and continuation functions are consistent with equilibrium behavior in the third stage. Note that now equilibrium profiles are of the type $(\sigma^0, \sigma^1)(\cdot)$. Then, by applying again Lemma 4, we further restrict the set of auxiliary games of order 1. Continuing in this way, we generate profiles $(\sigma^0, \sigma^1, \dots, \sigma^m)$ so that if $(\sigma^0, \sigma^1, \dots, \sigma^m)$ is an equilibrium profile of $AG_{f_m}(m)$, then $(\sigma^0, \sigma^1, \dots, \sigma^{m-1})$ is an equilibrium profile of $AG_{f_{m-1}^1}(n-1)$, $f_{m-1}^1 \in I_m(f_m)$. At the same time, we have been generating sequences $\{f_0^n\}, \dots, \{f_k^{n-k}\}, \dots$ by means of the relation $f_{j-1}^1 \in I_j(f_j)$.

As $n \rightarrow \infty$, a profile $(\sigma^0, \sigma^1, \dots, \sigma^n)$ tends to a profile σ that is feasible in the repeated game as it respects, by construction, the signalling structure of the latter. In light of the convergence property of the $\{AG(k)\}$'s (the Fact above), to show that these limiting profiles are equilibria (and all the equilibria) of the repeated game, it suffices to show that for each sequence $\{f_k^{n-k}\}$ and for each k the limit exists. Even though a direct proof will be given only below (Theorem 6 – Lemma 7), assume for a moment that it is, indeed, so. That is, suppose that all the sequences $\{f_k^{n-k}\}$ converge, as $n \rightarrow \infty$, and that the profiles σ thus determined are the equilibrium profiles of the repeated game. Then, if $f_k^{n-k} \rightarrow f_k^*$ as $n \rightarrow \infty$, we can define an auxiliary game of order k , $AG_{f_k^*}(k)$, so that each equilibrium payoff vector in $AG_{f_k^*}(k)$ is an equilibrium payoff vector in G^∞ , and each equilibrium profile in $AG_{f_k^*}(k)$ coincides with the first k components of the corresponding equilibrium in G^∞ . Moreover, this is so for any $k \in N$.

At this point, to complete the proof that all the requirements a), b) and c) are satisfied, we need only to apply Corollary 5. In other words, we have

1. For any $k \in N$, any equilibrium payoff vector achievable starting from stage k in the repeated game G^∞ can be obtained as an equilibrium of a (consistent) static game with incomplete info (Γ^k, P^k) . For any $k \in N$, (Γ^k, P^k) is a game with continuation payoffs.

2. An equilibrium profile σ^k that obtains such a payoff in (Γ^k, P^k) is the stage k th-component of an equilibrium profile that attains that same equilibrium payoff in G^∞ (starting from stage k).

3. (The recursive relation) For each equilibrium σ in G^∞ , the equilibrium continuation payoffs in $(\Gamma^k, P^k)(\sigma)$ are the expected equilibrium payoff in $(\Gamma^{k+1}, P^{k+1})(\sigma)$.

The method of successive approximations described above can be summa-

rized by means of the following diagram

$$\begin{array}{ccccccc}
& & & & F_3 & \dots & \\
& & & & \downarrow I_3 & & \\
& & & F_2 & F_2^1 & \dots & \\
& & & \downarrow I_2 & \downarrow I_2 & & \\
& & F_1 & F_1^1 & F_1^2 & \dots & \\
& & \downarrow I_1 & \downarrow I_1 & \downarrow I_1 & & \\
F_0 & F_0^1 & F_0^2 & F_0^3 & \dots & &
\end{array}$$

where (as observed above) for any k and n in N , $F_k^{n+1} \subseteq F_k^n$.

The following theorem establishes the existence of all the limits F_k^∞ .

Theorem 6 *For any $k \in N$, and any $n \in N$, each F_k^n is nonempty and compact. Hence, for any $k \in N$, F_k^∞ is nonempty.*

Proof. First observe that F_k is weak*-compact for each $k \in N$. By the upper hemicontinuity of the equilibrium correspondence, the correspondence $I_k : F_{k+1} \rightrightarrows F_k$ constructed in Lemma 4 is nonempty valued and upperhemicontinuous, for each $k \in N$. Hence, the projective limit of the projective sequence $\{F_k, I_k\}$ is nonempty. ■

A brief comment is in order. In the proof above, we showed existence of the limits by showing that the projective limit of the projective sequence $\{F_k, I_k\}$ is nonempty. Recall that such a projective limit is the set of sequences (a subset of $\times_k F_k$) with the property that if f_k is the k th – element in a sequence, then it belongs to the image according to I_k of the $(k+1)$ th – element of that sequence. In other words, such a projective limit is the set $F^\infty = \left\{ \{f_k\} \in \times_k F_k \mid f_k \in I_k(f_{k+1}) \right\}$. The property of the projective limit expresses the inductive relation c) at the beginning of this section. Once an f_k is known, we can define inductively, for each $t > k$, the continuation payoff functions by means of the relation $f_k \in I_k(f_{k+1})$.

5 The equilibrium payoff set

An immediate consequence of the procedure developed in the previous section is the following characterization of the equilibrium payoff set as the limit of a sequence of nested compact sets.

Let $V_0 = [0, c]^I$ and $F_0 = \{f = (f_1, \dots, f_I) \mid f : \Omega \rightarrow [0, c]^I, f \text{ measurable}\}$. Also, recall that $F_k = \left\{ f \mid f : \Omega \times \Theta^k \rightarrow [0, c]^I, f \text{ measurable} \right\}$, and that a choice of $f_k \in F_k$ defines (definition 3) an auxiliary game of order k , $AG_{f_k}(k)$. Denote by $V_{k,f}$ the set of Nash equilibrium payoff vectors of such an auxiliary game, that is $V_{k,f} = \{NE \text{ payoffs for } AG_{f_k}(k)\}$. Finally, let $V_k = \bigcup_{f_k \in F_k} V_{k,f}$.

Theorem 7 Let V be the equilibrium payoff set of the repeated game G^∞ , and let $V_\infty = \bigcap_{n=0}^\infty V_n$. Then, (i) V_k is nonempty and compact (in R^I), $\forall k \in N$; (ii) $V_k \supseteq V_{k+1}$; (iii) $V = V_\infty = \bigcap_{k=0}^\infty V_k$.

Proof. (i) By Nash, $V_{k,f} \neq \emptyset$. To show that V_k is compact, first note that V_0 is compact, and so is each F_k . Let $\{v_k^n\}_{n=1}^\infty$ be a sequence of elements in V_k . To such a sequence, there corresponds a sequence $\{f_k^n\}_{n=1}^\infty$ of elements in F_k . Then, it suffices to show that for $f_k^n \rightarrow f_k$ as $n \rightarrow \infty$, $v_k^n \rightarrow v_k \in V_k$, which is just the upper hemicontinuity of the correspondence $AG_{f_k}(k) \rightrightarrows V_{k,f}$.

(ii) Let $v_2 \in V_2$. Since v_2 is in V_2 , there exists an $f_1 : \Omega \times \Theta \rightarrow V$ and a profile $\sigma = (\sigma^0, \sigma^1)$ such that σ is a Nash equilibrium of $AG_{f_1}(2)$. By Lemma 4 we can define a function $\hat{f} \in F_0$ (that is $\hat{f} : \Omega \rightarrow [0, c]^I$) such that

$$v_2 = (1 - \delta)E_\Omega [g_i(\omega, \sigma^0)] + \delta \int_\Omega \hat{f}(\omega) d\mu(\sigma^0) \quad (4)$$

and v_2 is a Nash equilibrium of $AG_{\hat{f}}(1)$. Hence, $v_2 \in V_1$. By induction (Lemma 4 and the observation thereafter), $V_{k+1} \subseteq V_k$.

(iii) If $v \in V$, then, since the functions in F_n are restricted only by the feasibility condition, there exists an $f_n \in F_n$ such that v is a equilibrium payoff of $AG_{f_n}(n)$. Hence, $v \in V_n$. Since it is so for any $n \in N$, $v \in V_\infty$.

Let $v^* \in V_\infty$. Then, there exists a sequence $\{f_n^*\}$, $f_n^* \in F_n$, such that v^* is an equilibrium of $AG_{f_n^*}(n)$, for any $n \in N$. Let σ^* be the strategy profile constructed by means of the sequence of auxiliary games, $\{AG_{f_n^*}(n)\}$, which supports v^* . Suppose, $v^* \notin V$. Then, for some i , there exists a strategy σ_i in G^∞ for player i such that $v_i(\sigma_i, \sigma_{-i}^*) = v_i^* + \varepsilon$, for some $\varepsilon > 0$. ($v_i(\sigma_i, \sigma_{-i}^*)$ is the payoff achieved by player i in G^∞ when he plays σ_i and his opponents play σ_{-i}^*). In particular, there exists a set $K \subseteq N$ such that $\sigma_i \neq \sigma_i^*$ on K . Let $\underline{k}$ be the least element in K . Pick $n > \underline{k}$. Since $v_i(\sigma_i, \sigma_{-i}^*) = v_i^* + \varepsilon$, there exists an $f_n \in F_n$ such that player i achieves $v_i(\sigma_i, \sigma_{-i}^*)$ in $AG_{f_n}(n)$ when the restriction to the first n stages of $(\sigma_i, \sigma_{-i}^*)$ is played in $AG_{f_n}(n)$. We claim that for any $n \in N$, $\varepsilon \leq \delta^n c$. In fact, if $\varepsilon > \delta^n c$, player i can play according to σ_i for the first $n-1$ stages in $AG_{f_n^*}(n)$. Since $AG_{f_n}(n)$ and $AG_{f_n^*}(n)$ differ only because of the last stage, and in both cases the opponents play according to σ_{-i}^* , by doing so player i can ensure himself the same payoff in the first $n-1$ stages in both games. Hence, the payoff $\tilde{v}_i$ that he gets in $AG_{f_n^*}(n)$ by playing σ_i against σ_{-i}^* is at least as big as $v_i(\sigma_i, \sigma_{-i}^*) - \delta^n c$. Hence, $\tilde{v}_i \geq v_i(\sigma_i, \sigma_{-i}^*) - \delta^n c = v_i^* + \varepsilon - \delta^n c > v_i^*$, which contradicts σ^* being an equilibrium of $AG_{f_n^*}(n)$. Since, we can do the same reasoning for any $n > \underline{k}$, $\varepsilon \leq \delta^n c$ for any $n > \underline{k}$, which contradicts the assumption $\varepsilon > 0$. ■

Three easy implications of this above proof are probably worth to be stressed.

First, while existence of equilibria is not an issue in the present setting (the infinite repetition of a Nash equilibrium of G is an equilibrium of G^∞), it is worth noticing that the proof is itself an existence proof. To this end,

it suffices to observe that the family $\{V_k\}$ has the finite intersection property. The observation might reveal useful when the model is enlarged to allow for a stochastic game aspect (i.e., the component game G is allowed to vary with the state).

Second, the proof offers a way to approximate the equilibrium payoffs of G^∞ , a fact that might come useful when dealing with efficiency issues. Let f_n be any given function such that $f_n : \Omega \times \Theta^n \rightarrow [0, c]^I$. According to definition 3, such a choice defines an auxiliary game of order n , $AG_{f_n}(n)$. Let $V_n(f_n)$ be the equilibrium payoffs set of this $AG_{f_n}(n)$, and let v_{n, f_n} be an element in this set. For any choice of f_n and any n , the auxiliary game admits a finite normal form. Hence, up to some n , $V_n(f_n)$ can be computed. Then, it follows from above that the approximation in computing an equilibrium payoff vector of the repeated game with an element of $V_n(f_n)$ is less than $\delta^n c$, where δ is the discount factor and c bounds the payoff function. Furthermore, this bound does not depend on the particular choice of f_n . In other words,

$$V_n \subset \{x \in R^I \mid d(x, V) < \delta^n c\}$$

where $d(x, V) = \inf_{v \in V} \|x - v\|$, $\|\cdot\|$ denotes the Euclidean norm, and, as before V is the equilibrium payoff set of G^∞ .

Finally, one can express F^∞ , the projective limit of the previous section, as the largest invariant set of an operator on $\times_k F_k$. This observation provides some ground, at least in principle, for a computational approach to the entrance laws and to the equilibrium payoff set (which is the image of F_0^∞ under the Nash equilibrium correspondence). Formally, we have

Proposition 8 *There exists an operator $T : \times_k F_k \longrightarrow \times_k F_k$ such that the largest invariant set of T is F^∞ .*

Proof. For each k , we have correspondences $I_k : F_{k+1} \rightrightarrows F_k$ ($I_0 : F_0 \longrightarrow [0, c]^I$). For each $x \in \times_k F_k$, let g_x be the graph of the correspondence $I_k(x(k))$, $k \in N$. Such a graph is an element of $[0, c]^I \times \left(\times_k F_k\right)$. Then, the correspondence that associates x with such a graph is a correspondence from $\times_k F_k$ to $[0, c]^I \times \left(\times_k F_k\right)$. It is the canonical extension of the family $\{I_k\}$. By composing such a correspondence with the projection $[0, c]^I \times \left(\times_k F_k\right) \longrightarrow \left(\times_k F_k\right)$, we obtain a correspondence $T : \times_k F_k \longrightarrow \times_k F_k$. It follows from theorem 6 that T has an invariant set, and that its largest invariant set is F^∞ . ■

6 Extensions

6.1 Compact signal spaces

If we allow for infinite signal spaces, Θ_i , neither existence of equilibria of the $AG_{f_n}(n)$'s (for each $f_n \in F_n$) nor the upper hemicontinuity of the equilibrium correspondence are guaranteed any longer. Hence, in order to preserve the validity of our reasonings, we have to strengthen our assumptions. One possibility is to replace assumption (vi) with the following set.

- (vi.1) Each Θ_i is a complete, separable metric space;
- (vi.2) The signalling function $\theta_i : \Omega \rightarrow \Theta_i$ is continuous;
- (vi.3) For each $k \in N$, P^k is absolutely continuous with respect to the product of its marginals.

In such a case, the same argument as in Theorems 1 and 2 in Milgrom and Weber ([18]), ensure existence of equilibria of the $AG_{f_n}(n)$'s (for each $f_n \in F_n$) as well as the upper hemicontinuity of the equilibrium correspondence.

6.2 More general games

Modulo technical assumptions, our procedure can be extended to account for a large variety of features like deterministic state-transitions, stochastic state-transitions, stochastic signalling functions, etc. In fact, even in these circumstances, we can immediately give the same definitions of a type and a state of the world we gave above as soon as we appropriately enlarge those spaces. We, then, obtain the same construction and invoke Mertens-Zamir's theorem to establish the existence of a consistent probability on the set of possible states of the world at n , for any $n \in N$. The Mertens-Zamir consistent probability along with a specification of continuation payoffs will define a continuation game of incomplete information. We can then introduce the concept of auxiliary game of order n , appropriately modified to allow for the physical environment to change (either stochastically, or deterministically) from one stage to another. The successive approximations will then lead, as in section 4, to the space of consistent entrance laws, and to the determination of the equilibrium payoffs set of the repeated game.

7 Conclusions

In the Introduction, we said that one of the main motivations for the present work was the claim, present in most of the literature, of nonexistence of a recursive structure in games with private monitoring. As pointed out in that literature, this would have implied the inability of using dynamic programming methods and, hence, of describing a player's problem in any arbitrary stage and contingency. Consequently, the construction of a comprehensive theory for this class of games appeared out of reach. It goes without saying that a statement with such consequences called for a direct answer.

In this paper, we showed the existence of a recursive structure. In our representation, the strategic problem each player faces in any stage and contingency is described by means of a (static) game of incomplete information. By the very definition of a recursive structure, one such a game, $\{G(\Theta^k), P^k\}$, represents all the information that is needed to analyze the repeated game starting from any given state and contingency. The problem being recursive, all the other games, $\{G(\Theta^{k+n}), P^{k+n}\}$, are recovered from $\{G(\Theta^k), P^k\}$ in an inductive way. In the paper, we also provided (by means of the successive approximations) a constructive approach to the determination of these games, and hence to the representation itself.

Our construction emerges in a very natural way. With our definitions, the space of types is just the space of signals for the various players, and a state of the world is just a collection of types. These definitions came naturally from the definition of a player's strategy (a mapping from his information to his mixed action), and, in fact, our state space is the minimal one that ensures that all the games in the class we considered have perfect recall.

One first implication of our finding is the characterization of the equilibrium payoff set given above. More broadly, we feel that the main contribution of this paper is to have clarified the core concepts at the basis of a theory, and in which context the various questions are to be addressed. For instance, our setting delivers an immediate answer to the question of the common knowledge of the continuation play, which was very obscure before.

References

- [1] D. Abreu, D. Pearce and E. Stacchetti, Optimal Cartel Equilibria with Imperfect Private Monitoring, *J. of Econ. Theory* 39 (1986), 251-269.
- [2] D. Abreu, D. Pearce and E. Stacchetti, Toward a Theory of Discounted Repeated Games with Imperfect Monitoring, *Econometrica* 58 (1990), No. 5, 1041-1063.
- [3] R. J. Aumann and A. Brandenburger, Epistemic Conditions for Nash Equilibrium, *Econometrica* 63 (1995), No. 5, 1161-1180.
- [4] M. Ayoagi, A Collusion Mechanism in Dynamic Bertrand Oligopoly under Private Observation of Correlated Demand Shocks, University of Pittsburgh, mimeo, 1997.
- [5] Bhaskar V. and E. van Damme, Moral Hazard and Private Monitoring, University of St. Andrews and CentER Tilburg, mimeo, 1997.
- [6] O. Compte, On failing to Cooperate when Monitoring is Private, CERAS-ENPC, mimeo, 1996.
- [7] D. Fudenberg, B. Holmstrom and P. Milgrom, Short-term Contracts and Long-term Agency Relationships, *J. of Econ. Theory* 51 (1990), 1-31.

- [8] D. Fudenberg, D. Levine and E. Maskin, The Folk Theorem with Imperfect Public Information, *Econometrica* 62 (1994), No. 5, 997-1039.
- [9] D. Fudenberg and E. Maskin, The Folk Theorem with discounting and with Incomplete Information, *Econometrica* 54 (1986), 533-554.
- [10] E. Green and R. Porter, Noncooperative collusion under imperfect price formation, *Econometrica* 52 (1984), 87-100.
- [11] J.C. Harsanyi, Games of incomplete information played by Bayesian players, Parts I, II, III, *Management Science* 14 (1967-68), 159-182, 320-334, 486-502.
- [12] M. Kandori and H. Matsushima, Private Observation, Communication and Collusion, *Econometrica* 66 (1998), 627-653.
- [13] G. Mailath and S. Morris, Repeated Games with Imperfect Private Monitoring: Notes on a Coordination Perspective, CARESS Working Paper 98-07, 1998.
- [14] H. Matsushima, On the theory of repeated games with private information, Parts I, II, *Economic Letters* 35 (1991), 253-256 and 257-261.
- [15] J.-F. Mertens, Repeated Games, in "Proceedings of the International Congress of Mathematicians (Berkeley, 1986)", pp. 1528-1577, American Mathematical Society, 1987.
- [16] J.-F. Mertens and T. Parthasarathy, Non zero sum stochastic games, in "Stochastic Games and Related Topics" (T.E.S. Raghavan et al. eds.), Kluwer Academic Publishers, pp. 145-148, 1991.
- [17] J.-F. Mertens and S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory* 14 (1985), 1-29.
- [18] P.R. Milgrom and R.J. Weber, Distributional Strategies for Games with Incomplete Information, *Mathematics of Operation Research* 10 (1985), 619-32.
- [19] T. Sekiguchi, Efficiency in Repeated Prisoner's Dilemma with Private Monitoring, *University of Tokyo, J. Econ. Theory* 76 (1997), 345-361.
- [20] G. Stigler, A Theory of Oligopoly", *Journal of Political Economy* 72 (1964), 44-61.
- [21] N. L. Stokey and R.J. Lucas, *Recursive Methods in Economic Dynamics*, Harvard University Press, 1989.