

Limited Foresight May Force Cooperation

PHILIPPE JEHIEL

CERAS, Ecole Nationale des Ponts et Chaussées and University College London

First version received November 1998; final version accepted May 2000 (Eds.)

This paper considers discounted repeated games with boundedly rational players. In each period, player i chooses his current action on the basis of his forecast about the forthcoming n_i action profiles; his assessment of the payoffs he will obtain next depends on his state of mind, which is non-deterministic. A limited forecast equilibrium is such that after every history the limited horizon forecasts formed by the players are correct. The set of all limited forecast equilibria is characterized and constructed. Application to the repeated prisoner's dilemma shows that limited foresight may sometimes induce purely cooperative paths while purely non-cooperative paths cannot arise.

1. INTRODUCTION

Repeated game theory assumes that players are fully rational. In long horizon interactions, this implies that players should be able to form (rational) predictions about what will come in a very distant future as a function of what they currently plan to do (and possibly the history of play). Such an assumption seems extremely strong, and we wish to investigate the more plausible paradigm in which players only try to make *limited* predictions about what will happen over the horizon of interaction.

Several forms of limitation in the players' predictions may *a priori* be considered. A limitation of particular interest seems to be that each player i when he must move only tries to predict what will come in the next n_i periods (possibly as a function of what happened in the past). Regarding the outcome in later periods beyond the *horizon of foresight*, player i is viewed as forming no prediction about it (or as having no idea about it).

Such a view on agents' forecasting abilities seems in accordance with the behaviour of many human beings in long-horizon interaction contexts (see, for example, the functioning of chess players, Simon and Schaeffer (1992)). We wish to investigate the implications of such modes of behaviour in strategic interaction contexts.

We need first to define a solution concept in order to describe how agents with *limited foresight* will interact. The solution concept we propose has two parts. First, it specifies how a player makes his choice of action on the basis of his limited forecast about the future. Second, it imposes some constraints on the limited forecasts formed by the players in equilibrium.

In every period each player i has a limited forecast, which is restricted to the forthcoming n_i action profiles. Formally, a limited forecast is a mapping from the set of player i 's current actions into the set of possible distributions of truncated paths over the forthcoming n_i periods. In general, those mappings may also depend on history. The criterion used by player i to choose his current action is a convex combination of two terms: The first term stands for the average discounted payoff over the forthcoming n_i periods as

given by his limited forecast. The second term stands for player i 's assessment of the expected payoff he will obtain in subsequent periods beyond his horizon of foresight. The relative weight between the two terms is that of the two time periods as given by the discount factor δ and the length n_i of the horizon of foresight.¹

Player i 's assessment of what is to come beyond his horizon of foresight is assumed to be the realization of some exogenously given random variable $\tilde{\varepsilon}_i$. The realization of the random variable—which may be interpreted as his current *state of mind*—may possibly assign different values to the various possible current actions (because player i 's state of mind need not be such that he believes his current action will not affect what is to come beyond his horizon of foresight). The realizations of the state of mind may also vary from period to period. All realizations are required to be feasible payoffs given the stage game payoff matrix.

The stochastic nature of the state of mind is to be related to an important strand of literature in psychology, which goes back to Thurstone (1927). In order to explain some results in psychological experiments, Thurstone proposed a model based on the idea that a given stimulus provokes a “psychological state” that is the *realization* of a *random variable*, and the response of an individual having to compare two stimuli results in the *comparison of the realization of two random variables* that represent the sensations provoked by the two stimuli. Our formulation of states of mind follows that view.² Note that the state of mind in our approach bears only on what the player does not forecast.³

A sequence of n_i length limited forecasts for each period of interaction will be referred to as a forecasting rule for player i and denoted by f_i . We will say that a strategy σ_i for player i is *justified* by a forecasting rule f_i whenever in every period and for every history the criterion implicitly defined above (as a function of f_i and ε_i) induces player i to adopt the behavioural strategy given by σ_i .

In equilibrium, we impose some consistency constraints on the limited forecasts formed by the players. Specifically, we assume that in every period and for every history the prediction over the forthcoming n_i periods made by player i is *correct* for every action player i may choose. More precisely, consider a history of play h at time t , and suppose the behavioural strategies used by the players in the forthcoming n_i periods induce some distribution over player i 's horizon of foresight when player i plays action a_i in the current period. In equilibrium we require that player i 's prediction over the forthcoming n_i periods were a_i to be played in the current period coincides with that distribution. A forecasting rule f_i will be said to be *consistent* with a strategy profile σ whenever that correctness property is met in every period, for every history of play and for every current action player i may choose.⁴

To sum up, consider a repeated game where, for $i = 1, 2$, player i 's length of foresight is given by n_i , and player i 's distribution of state of mind is (in every period) given by $\tilde{\varepsilon}_i$.⁵

1. That is, the weight of the former term is $(1 + \delta + \dots + \delta^{n_i-1})(1 - \delta)$ and the weight of the latter is δ^{n_i} .

2. That the realizations of states of mind may vary from period to period is also in line with modern views of Thurstone's model, which consider the psychological state as momentary (see Edgell and Geisler (1980, p. 266).

3. This should be contrasted with the models developed in Anderson *et al.* (1992) or in Chen–Friedman–Thisse (1997) in which (implicitly) agents make perfect predictions and the psychological state affects the whole preference.

4. The consistency requirement imposed on equilibrium forecasts (which might seem strong at first glance) should be thought of as resulting from a learning process (not from introspection or calculation). I believe the learning approach also suggests to rule out solution concepts within which players would base their choices of current actions on plans of actions (as resulting from their limited forecasts) that they would not follow afterwards (a time-inconsistency like problem, see Rubinstein (1998, ch. 7) and Jehiel (1998a, Subsection 6.5)).

5. For simplicity, we will assume that the distributions of states of mind are independent from period to period.

A limited forecast equilibrium of this game is a strategy profile σ such that for some forecasting rule f_i of player $i = 1, 2$ (1) player i 's strategy σ_i is justified by f_i and (2) player i 's forecasting rule f_i is consistent with σ .⁶

The paper next proceeds to analyse the set of limited forecast equilibria in the general setting of repeated games with finite action spaces. We first show that a limited forecast equilibrium necessarily exists whatever the stage game payoffs, the lengths of foresight of the players and their distributions of states of mind (Proposition 1).

We next provide a general (and simple) algorithm for constructing all limited forecast equilibria whether in mixed or in pure strategies. The idea underlying the algorithm is as follows. Let m be no smaller than the maximum length of foresight among the players minus one. Consider an arbitrary set X of distributions over m -length paths (*i.e.* over m consecutive action profiles). The antecedent of X , $W(X)$, is a set of distributions over m -length paths constructed from X as follows. An element of $W(X)$ is a distribution of action profiles from period t to period $t + m - 1$ obtained when (1) for every period t action profile, the true distribution of action profiles from period $t + 1$ to period $t + m$ (as induced by the strategy profile) is one of the distributions in X , (2) the period t limited forecasts of the players are correct (in the sense defined above) and (3) the players use the criterion based on limited forecasts as defined above.⁷ Starting from the set of all possible distributions over m -length paths, the iterative application of the antecedent operator is shown to converge to the set of all possible m -length paths from period 1 to period m that can be sustained in limited forecast equilibria (Theorems 1–2). The same analysis is carried out for limited forecast equilibria in pure strategies (Theorem 3). In the case of limited forecast equilibria in pure strategies the algorithm is extremely simple and converges in a finite number of steps (because there is a finite number of possible m -length pure paths).⁸

Finally, the paper contains an application section. We first consider repeated 2×2 coordination games. We observe that, for intermediate values of the discount factors, it may happen that the repetition of the risk-dominant stage game Nash equilibrium can be sustained as a limited forecast equilibrium while the repetition of the other stage game pure Nash equilibrium cannot (see Harsanyi and Selten (1988) or Subsection 5.1 for the definition of risk-dominance).

We next consider the repeated prisoner's dilemma. The most interesting finding is that for intermediate discount factors, a purely cooperative path can be sustained as a limited forecast equilibrium outcome while a purely non-cooperative path cannot. Furthermore, we apply the general construction techniques to a numerical example (of the repeated prisoner's dilemma). We show (in the example) that the worst (in the Pareto sense) limited forecast equilibrium in pure strategies induces an equilibrium path made of a limited number of periods in which both players play non-cooperatively and followed by an infinite repetition of the cooperative action profile. In this sense, limited foresight may force cooperation.⁹

6. It should be noted that we make no consistency requirement on the random assessment made by the players beyond their horizon of foresight, which constitutes the essence of bounded rationality in the solution concept.

7. The complexity of constructing $W(X)$ is that of deriving a fixed point over a mixed strategy in the stage game action space.

8. Limited forecast equilibria in pure strategies may sometimes fail to exist, but when they exist they have the (desirable) feature that they are independent of the distributions of state of mind as long as the support of the latter remains unchanged.

9. There may be limited forecast equilibria in mixed strategies in which (after some histories) the actual choice of action depends on the realization of the state of mind. However, these equilibria have the unappealing feature that they are sensitive to the specific form of the distribution of state of mind.

The reason why limited foresight may force cooperation is as follows. A deviation from the stationary strategies underlying the repetition of the stage game non-cooperative Nash equilibrium imposes a cost in the current period, but no punishment in later periods (because the maximal punishment has already been reached). In contrast, a deviation from a cooperative behaviour is typically followed by a punishment phase, which—if within the horizon of foresight—is taken into account by the players. In some cases, the incentives to comply with the cooperative behaviour are then sufficient to outweigh possibly unfavourable states of mind (regarding what is to come beyond the horizon of foresight), while the incentives to comply with a stationary pure non-cooperative behaviour are insufficient. Cooperation emerges rather than non-cooperation because it allows to better reduce the effect of the ignorance of the players about what is to come beyond their horizon of foresight.

The rest of the paper is organized as follows. Section 2 presents the model and the solution concept. Section 3 reports some preliminary analysis, in particular about the relationship between Subgame Perfect Nash Equilibria and limited forecast equilibria as the lengths of foresight of the players increase to infinity. Section 4 provides a general algorithm for constructing all limited forecast equilibria whether in mixed or in pure strategies. Section 5 is the application section. In addition to repeated coordination games and repeated prisoner's dilemmas, we briefly consider repeated battles of the sexes. Section 6 mentions the relationship with the literature. Section 7 concludes. All proofs are gathered in the Appendix.

2. THE LIMITED FORECAST EQUILIBRIUM

Consider a stage game G with two players $i = 1, 2$. Each player i has a finite action space A_i . Player i derives utility $u_i(a)$ when the action profile is $a \in A_i \times A_{-i}$ and $-i$ stands for the player other than i . We will consider the discounted repetition of the stage game G in which δ is the common discount factor of players 1 and 2. The repeated game will be denoted by G^δ .

We wish to define a solution concept for the game G^δ when each player i is assumed to have a limited foresight of length n_i for $i = 1, 2$.

Some notation is required before the solution concept can be defined. We let h denote a history at time t ; it is a $(t-1)$ -tuple of action profile one for each period $k = 1, \dots, t-1$. We let H denote the set of h . We denote by p_i a (pure) prediction for player i . It is a sequence of action profiles for the forthcoming n_i periods including the current one. The set of player i 's pure predictions p_i is denoted P_i (we have $P_i = A^{n_i}$).

A key notation is that of the limited forecast. At period t , given the history of play $h \in H$ and for each possible action $a_i \in A_i$ player i may choose, player i forms a forecast over the forthcoming n_i action profiles. That forecast may be stochastic and therefore it is an element of the convex hull of P_i , which is denoted by ΔP_i . Formally, a period t limited forecast for player i is denoted by $f_i^t = (f_i^t(\cdot|h))_h$ where $f_i^t(\cdot|h): A_i \rightarrow \Delta P_i$ and $f_i^t(a_i|h)$ is player i 's stochastic prediction (or belief) about the action profiles a^{t+k} for $k = 0, \dots, n_i-1$ when the history of play is h and player i considers playing a_i in the current period t . For every $p_i \in P_i$, $f_i^t(a_i|h)[p_i]$ will denote the weight of the pure prediction $p_i \in P_i$ in the belief $f_i^t(a_i|h)$. That is, according to $f_i^t(a_i|h)$ player i believes that given h , the forthcoming n_i action profiles from period t to period $t+n_i-1$ if he plays a_i at period t are given by p_i with probability $f_i^t(a_i|h)[p_i]$. We also denote by $f_i = (f_i^t)_t$, which is referred to as player i 's forecasting rule. F_i denotes the set of player i 's forecasting rules. A forecasting rule profile (f_1, f_2) is denoted by f , and the set of f is denoted F .

Finally, σ_i^t denotes the behavioural strategy of player i in period t ; it maps the set of histories H onto the set of possibly mixed actions ΔA_i , where $\sigma_i^t(h)[a_i]$ denotes the probability that player i plays a_i in period t given $h \in H$; $\sigma_i = (\sigma_i^t)_t$ denotes a strategy for player i , and the set of σ_i is denoted by Σ_i . A strategy profile $(\sigma_1, \sigma_2) \in \Sigma_2$ is denoted by σ , and the set of σ is denoted by Σ .

The solution concept has two parts. First, it makes explicit how each player i makes his choice of current action on the basis of his limited forecasts. Formally, this will result in a relationship between the strategy σ_i and the forecasting rule f_i for each player i . Second, some consistency of the forecasting rules will be required relating each forecasting rule f_i with the strategy profile σ .

We first define the criterion used by each player i to select current actions. The criterion is a convex combination of two terms, where the first term is the average expected payoff obtained over the horizon of foresight as given by his forecast and the second is the realization of some random variable which stands for the discounted average flow of player i 's payoffs after the horizon of foresight as given by his current *state of mind*. Formally, let $\varepsilon_i = (\varepsilon_i(a_i))_{a_i}$ denote player i 's state of mind regarding the unpredicted period outcome. It is the realization of a vector of exogenously given random variables $\tilde{\varepsilon}_i$ with density $\gamma(\cdot)$ and full support on $(u_i^{\text{inf}}, u_i^{\text{sup}})^{\#A_i}$, where $u_i^{\text{inf}} \cong \min_{a \in A} u_i(a)$, $u_i^{\text{sup}} \cong \max_{a \in A} u_i(a)$ and $\#A_i$ denotes the cardinality of the action space A_i . The realization $\varepsilon_i(a_i)$ should be interpreted as player i 's period t assessment of the average expected per period payoff from period $t + n_i$ onwards when he plays a_i in period t (as given by his current state of mind).¹⁰ We will assume that the draws of states of mind are independent from period to period (extension to the correlated case would raise no conceptual difficulties). Note that a realization of state of mind ε_i allows for possibly different assessments $\varepsilon_i(a_i)$ for each action $a_i \in A_i$.

Given the period t forecast f_i^t , the history of plays $h \in H$ and the vector of draws $\varepsilon_i = (\varepsilon_i(a_i))_{a_i}$, player i 's assessment of action a_i is given by

$$c_i(a_i | f_i^t, h, \varepsilon_i) = (1 - \delta)v_i(f_i^t(a_i | h)) + \delta^{n_i}\varepsilon_i(a_i),$$

where

$$v_i(f_i^t(a_i | h)) = \sum_{p_i \in P_i} f_i^t(a_i | h)[p_i] \cdot w_i(p_i),$$

and

$$w_i(p_i) = \sum_{k=0, \dots, n_i-1} \delta^k u_i(a^{(k)}),$$

with

$$p_i = (a^{(k)})_{k=0, \dots, n_i-1}.$$

Note that $v_i(f_i^t(a_i | h))$ is the expected discounted sum of player i 's payoffs over his horizon of foresight as given by his forecast $f_i^t(a_i | h)$, and thus $c_i(a_i | f_i^t, h, \varepsilon_i)$ is the exact expected per period payoff player i would obtain by playing a_i if the distribution of action profile from period t to period $t + n_i - 1$ were given by $f_i^t(a_i | h)$ and the average expected per period payoff from period $t + n_i$ onwards were given by $\varepsilon_i(a_i)$.

In order to define the solution concept it will be useful to introduce the set

$$\mathcal{R}_i(a_i | f_i^t(a_i | h)) = \{\varepsilon_i | c_i(a_i | f_i^t, h, \varepsilon_i) > c_i(a'_i | f_i^t, h, \varepsilon_i) \text{ for all } a'_i \neq a_i\},$$

10. We require that $u_i^{\text{inf}} \cong \min_{a \in A} u_i(a)$, and $u_i^{\text{sup}} \cong \max_{a \in A} u_i(a)$ because realizations of states of mind must correspond to (feasible) stage game payoffs.

and the probability

$$\pi_i(a_i | f_i^t(a_i | h)) = \int_{\varepsilon_i} \mathbf{1}_{\{\varepsilon_i \in R_i(a_i | f_i^t(a_i | h))\}} \gamma(\varepsilon_i) d\varepsilon_i.$$

$R_i(a_i | f_i^t(a_i | h))$ is the set of player i 's states of mind such that given the history h and the period t limited forecast f_i^t , player i prefers action a_i to any other action a_i' (according to his criterion c_i); $\pi_i(a_i | f_i^t(a_i | h))$ is the associated *ex ante* probability that player i chooses a_i prior to the realization of his state of mind.¹¹

To summarize, the behavioural strategy of player i with forecasting rule f_i can be described as:

Definition 1. (The criterion). A strategy $\sigma_i \in \Sigma_i$ is *justified* by a forecasting rule $f_i \in F_i$ if and only if for all periods t , for all $h \in H$, and for all $a_i \in A_i$, $\sigma_i^t(h)[a_i] = \pi_i(a_i | f_i^t(a_i | h))$.

The second requirement is that the limited forecasts formed by the players coincide with the true distribution of action profiles over the horizon of foresight as induced by the strategy profile. In other words, the forecasts formed by the players over their horizon of foresight in equilibrium are correct on and off the equilibrium path. Formally, for each strategy profile $\sigma \in \Sigma$, we denote by $\sigma^t|_h$ the strategy profile induced by σ from period t onwards after history h . We also denote by $q^t(\sigma^t|_h, a_i)$ the distribution over the action profiles from period t onwards generated by σ when the action profiles from periods 1 to period $t-1$ are given by h and player i 's action in period t is a_i . Finally, we denote by $[q^t(\sigma^t|_h, a_i)]_{n_i}$ the marginal distribution of $q^t(\sigma^t|_h, a_i)$ over the action profiles from period t to period $t+n_i-1$. Consistency is defined as:

Definition 2. (Consistency). A forecasting rule $f_i \in F_i$ is *consistent* with the strategy profile $\sigma \in \Sigma$ if for all periods t , for all $h \in H$ and for all $a_i \in A_i$, $f_i^t(a_i | h) = [q^t(\sigma^t|_h, a_i)]_{n_i}$.

We can now define the solution concept: It is a strategy profile, *i.e.* a pattern of behaviour, such that (i) it results from forecasting rules as summarized in Definition 1 and (ii) players' limited forecasts are correct on and off the equilibrium path for every history of play as expressed in Definition 2.

Definition 3. (The solution concept). A strategy profile $\sigma = (\sigma_1, \sigma_2) \in \Sigma$ is an (n_1, n_2) limited forecast equilibrium if and only if there exists a forecasting rule profile $f = (f_1, f_2) \in F$ such that for $i = 1, 2$ (i) σ_i is *justified* by f_i and (ii) f_i is *consistent* with σ .

Several comments are in order. First it should be emphasized that the solution concept takes as exogenous the lengths n_i of foresight of the players and their distributions of states of mind $\tilde{\varepsilon}_i$. The limited forecasts f_i are endogenized (given n_i and $\tilde{\varepsilon}_i$) through the consistency requirement. Note that the correctness of the limited forecasts should not—in my view—be interpreted as resulting from introspection on the part of the players.¹²

11. The treatment of indifference is immaterial because they have probability 0.

12. Indeed if player i were able by introspection to make the correct forecast over the forthcoming n_i moves, I fail to see any major reason why he would not be able to go further and make a complete forecast about the entire future (unless calculation takes time and the player faces time-constraints, but this is not the point of view adopted in this paper).

Rather it should be viewed as resulting from a (successful) learning process (see Jehiel (1998a) for a learning process yielding asymptotically the consistency in a repeated alternate-move context with deterministic criterion). Despite the apparent complexity of the description of solution concept, observe that the learning required for a limited forecast equilibrium is less demanding than the learning required for a Subgame Perfect Nash Equilibrium, since it bears on more limited objects (distributions over truncated paths rather than distributions over infinite paths).¹³

A feature of the criterion used by the players is that it is not very sophisticated regarding the assessment of the payoffs beyond the horizon of foresight. (The distributions of states of mind $\tilde{\varepsilon}_i$ are left exogenous, and satisfy no consistency requirement of any kind.) The lack of sophistication on this part is the essence of the boundedly rational character of the players.¹⁴ We have chosen to model player i 's assessment of the continuation value beyond his horizon of foresight in a stochastic fashion. However, note that whenever player i must make a decision, he has a deterministic assessment of this continuation value, which is summarized by the draw ε_i (called state of mind). The stochastic representation of the state of mind is meant to capture the (psychologically plausible) idea that sometimes players may be optimistic and sometimes pessimistic regarding what they do not forecast (they have no clear idea about it, see Chapter 2 of Anderson, de Palma and Thisse (1992) and the references therein).^{15,16} So from the point of view of an outside observer (which is the one adopted by the players when they look for regularities over truncated paths), it is *as if* players used stochastic continuation values. In the same vein, note that we have assumed that the realizations of the terms $\varepsilon_i(\cdot)$ depend on the action a_i player i may currently choose. Whether the true payoffs beyond the horizon of foresight are affected by a_i or not depends on the true strategy profile σ used by the players. However player i is unaware of the true strategy profile σ , and therefore of the correct dependence. Whether player i considers his current action will affect or not his payoffs beyond his horizon of foresight is described by his state of mind ε_i .

A significant part of the following analysis will be devoted to limited forecast equilibria which have the property that (despite the randomness of the distributions of states of mind) the behaviour of the players is independent of the realization of their state of mind. Thus, from the viewpoint of an outside observer the players are viewed as employing pure strategies—we will refer to such equilibria as limited forecast equilibria in pure strategies.

13. The solution concept could easily be adapted to contexts in which each player i restricts himself to make limited forecasts that may depend only on the past N_i action profiles (rather than on the whole history h and the time period t). In such contexts, the learning would bear only on a finite number of objects (*i.e.* how the past N_i action profiles together with the current action a_i affects the distribution of action profiles over the forthcoming n_i moves), thus making successful learning even more likely. In some applications, we will observe that all limited forecast equilibria in pure strategies are of this form.

14. If the state of mind ε_i were allowed to depend on history h and on the limited forecast $f_i^t(\cdot|h)$ and if ε_i were required to satisfy a consistency requirement (*i.e.* that it corresponds for each realization of events over P_i to the true expected continuation payoff beyond the horizon of foresight), then the above Definition 3 would coincide with that of Subgame Perfect Nash Equilibrium.

15. One might argue that sometimes players will extrapolate what they forecast onto what they do not forecast to assess the continuation payoffs beyond their horizon of foresight. However, players would still remain unsure as to the effect of their actions beyond their horizon of foresight thus making the randomness of the assessment of the continuation values plausible. In this more general case, the realizations of the state of mind should also depend on the limited forecasts f_i . However no additional insights would be gained, and this would make the construction of the solutions in mixed strategies a bit more cumbersome.

16. In applications we will briefly discuss the case where players have perfect foresight when assessing actions on the equilibrium path and have a limited horizon foresight when assessing actions off the equilibrium path (with an associated stochastic state of mind). This appears to be very similar to assuming that players use extrapolations to assess continuation values associated with actions on the equilibrium path (and not for the actions off the equilibrium path).

A nice property of limited forecast equilibria in pure strategies is that they remain equilibria for all distributions of states of mind $\tilde{\varepsilon}_i$ as long as the support $(u_i^{\text{inf}}, u_i^{\text{sup}})^{\#A_i}$ of $\tilde{\varepsilon}_i$ remains unchanged. They are thus robust to variations on the specific distributional assumptions of states of mind.¹⁷ The special case where $u_i^{\text{inf}} = \min_{a \in A} u_i(a)$ and $u_i^{\text{sup}} = \max_{a \in A} u_i(a)$ is of particular interest, since it covers all possible formulations of the states of mind (the continuation value has to be feasible). In some applications, we will discuss how the set of limited forecast equilibria in pure strategies varies with the degree of randomness as measured by $\Delta_i = u_i^{\text{sup}} - u_i^{\text{inf}}$.

3. PRELIMINARY RESULTS

We start with some preliminary results and we remind the reader that all proofs are gathered in the Appendix. First we prove the existence of (n_1, n_2) -limited forecast equilibria for arbitrary stage game payoffs, discount factors and distributions of states of mind.

Proposition 1. There always exists at least one (n_1, n_2) -limited forecast equilibrium.

We next explore how (n_1, n_2) -limited forecast equilibria relate to Subgame Perfect Nash Equilibria (SPNE) when the lengths of foresight n_1, n_2 increase to infinity, keeping the distributions of states of mind $\tilde{\varepsilon}_i, i = 1, 2$, fixed. We make two statements: Proposition 2 shows that the limit as the lengths of foresight increase to infinity of (a subclass of) limited forecast equilibria must lie in the set of SPNE; Proposition 3 shows that (a subclass of) SPNE are (n_1, n_2) -limited forecast equilibria for sufficiently large lengths of foresight.

For the first result, we will restrict attention to time-independent strategies with bounded recall:

Definition 4. A strategy profile σ is time-independent and has bounded recall of size N if at any time t and for every history h the behavioural strategy $\sigma_i^t(h)$ of player i depends only on the last N action profiles of h (i.e. it does not depend on t nor on the action profiles of h in periods $1, \dots, t - N - 1$).¹⁸

We have:

Proposition 2. Let $(\sigma^{n_1, n_2})_r$ be a sequence of time-independent (n_1^r, n_2^r) -limited forecast equilibria of G^{δ} with bounded recall of size N such that $n_i^r, i = 1$ and 2 , tends to infinity as r tends to infinity.¹⁹ Then the sequence has an accumulation point,²⁰ and all accumulation points of $(\sigma^{n_1, n_2})_r$ are Subgame Perfect Nash Equilibria with time-independent strategies and bounded recall of size N .

For the next result, we will restrict attention to a subclass of SPNE in pure strategies in which at any node the incentives to comply with the equilibrium strategy are strict and uniformly bounded by a strictly positive number. Specifically, we denote by $\tilde{u}_i(q)$ the average discounted per period expected payoff obtained by player i when the path q is

17. It should also be clear that this robustness idea extends to the case where the distribution of state of mind may also depend on history, the time period and the limited forecast.

18. If h contains less than N action profiles, then the behavioural strategy $\sigma_i^t(h)$ may depend on the whole h .

19. The proof of Proposition 1 shows that such limited forecast equilibria exist for all r .

20. The topology of the set of time-independent strategies with bounded recall of size N is unambiguously defined, see Appendix.

being played, and (with some abuse of notation) we denote by $\sigma_i^t(h)$ the pure action prescribed by player i 's pure strategy σ_i in period t after history h . Recalling that $q^t(\sigma^t|_h, a_i)$ is the path induced by σ from period t onwards when the history up to period t is h and player i plays a_i at t , we define:

Definition 5. A SPNE of G^δ with locally strict uniform incentives is a strategy profile σ in pure strategies such that there exists $\alpha > 0$, $\forall t, \forall h \in H, \forall i = 1, 2, \forall a_i \neq \sigma_i^t(h)$

$$\tilde{u}_i(q^t(\sigma^t|_h, \sigma_i^t(h))) > \tilde{u}_i(q^t(\sigma^t|_h, a_i)) + \alpha.$$

Observe that if the stage game G admits a strict Nash equilibrium (in pure strategies) then the repetition of this stage game Nash equilibrium constitutes a SPNE with locally strict uniform incentives. More generally, SPNE which employ a finite number of states (in either the sense of Abreu and Rubinstein (1988) or that of Kalai and Stanford (1988)) and such that no player faces any indifference in any state is a SPNE with locally strict uniform incentives.²¹ We have:

Proposition 3. Consider any SPNE σ with locally strict uniform incentives. Then there exists a threshold length of foresight n^* such that for any $n_i > n^*$ for $i = 1, 2$, σ is a (n_1, n_2) -limited forecast equilibrium.

Finally, we derive some robustness properties of limited forecast equilibria in pure strategies when the supports of states of mind cover all possible payoffs:

Proposition 4. Assume that for $i = 1, 2$, the distribution $\tilde{\epsilon}_i$ has full support on $(u_i^{\text{inf}}, u_i^{\text{sup}})^{\#A_i}$ with $u_i^{\text{inf}} = \min_{a \in A} u_i(a)$ and $u_i^{\text{sup}} = \max_{a \in A} u_i(a)$. Let σ be a (n_1, n_2) -limited forecast equilibrium in pure strategies. Then σ is a SPNE of G^δ . Besides, σ is a (n'_1, n'_2) -limited forecast equilibrium for all (n'_1, n'_2) , $n'_1 \geq n_1, n'_2 \geq n_2$.

Proposition 4 has interesting implications (especially in view of a learning model). Suppose each player i with length of foresight n_i has managed to make correct n_i -length forecasts. Suppose further that the associated pattern of behaviour is independent of the realization of states of mind so that a (n_1, n_2) -limited forecast equilibrium in pure strategies σ is being played. Suppose then that after a while, player 2 decides to increase his length of foresight from n_2 to $n_2 + 1$. Provided the distributions of states of mind of each player i covers all possible values, the play of the game will remain unchanged, since (by Proposition 4) σ is also a $(n_1, n_2 + 1)$ -limited forecast equilibrium.

4. CONSTRUCTION OF (n_1, n_2) -LIMITED FORECAST EQUILIBRIA

This section provides the main tool for the analysis of (n_1, n_2) -limited forecast equilibria. Specifically, we provide a general algorithm for constructing all (n_1, n_2) -limited forecast equilibria whether in mixed (Subsection 4.1) or in pure (Subsection 4.2) strategies.

21. Because there is a finite number of states, the equilibrium conditions are characterized by a finite number of inequalities. Because every such inequality must be strict (since there is no indifference), the lower bound on the differences $\tilde{u}_i(q^t(\sigma^t|_h, \sigma_i^t(h))) - \tilde{u}_i(q^t(\sigma^t|_h, a_i))$ may be chosen to be strictly positive, thus strictly larger than some $\alpha > 0$.

4.1. Limited forecast equilibria in mixed strategies

In the construction we make use of an operator—called the antecedent operator—that is defined for every subset of distributions over paths of given length m , $m \geq n_i - 1$ for $i = 1, 2$.

Before we can define the antecedent operator, some additional notation is required. We denote by Ψ_m the set of all distributions over paths of length m for any integer m . Let x be a distribution over paths of length m , i.e. $x \in \Psi_m$. We denote by $[x]_k$ the marginal distribution over the first k action profiles induced by x . For each action profile $a \in A$, we denote by (a, x) the distribution over paths of length $(m+1)$ induced by the action profile a and followed by the distribution x .

Consider a mapping $x(\cdot)$ from the set of action profiles A into the set Ψ_m of distributions of length m . For each probability distribution $\beta(\cdot)$ on A (i.e. $\beta(\cdot) \in \Psi_1$), and each function $x(\cdot)$, we denote by $\langle \beta(\cdot), x(\cdot) \rangle$ the distribution over paths of length $m+1$, which assigns weight $\beta(a)$ to the distribution $(a, x(a))$.²²

For each probability distribution $\beta_{-i}(\cdot)$ on A_{-i} , we denote by $\langle a_i; \beta_{-i}(\cdot), x(\cdot) \rangle$ the distribution over paths of length $m+1$, which assigns weight $\beta_{-i}(a_{-i})$ to the distribution $(a, x(a))$ where $a = (a_i, a_{-i})$. Finally, $\langle a_i; \beta_{-i}(\cdot), x(\cdot) \rangle_{n_i}$ will denote the marginal distribution over the first n_i action profiles induced by the distribution $\langle a_i; \beta_{-i}(\cdot), x(\cdot) \rangle$.

Definition 6. (The antecedent operator). Let X be a subset of Ψ_m . A distribution $\langle \beta(\cdot), y(\cdot) \rangle$ where $\beta(\cdot) \in \Delta A$ and $y(\cdot): A \rightarrow \Psi_{m-1}$ over paths of length m belongs to $W(X)$ —the antecedent of X —if and only if

- (1) There exists a mapping $x(\cdot): A \rightarrow X$ such that, for all $a \in A$, $y(a) = [x(a)]_{m-1}$ and
- (2) There exists $\beta_i \in \Delta A_i$ for $i = 1, 2$ such that
 - (2.1) For all $a = (a_1, a_2) \in A$, $\beta(a) = \beta_1(a_1)\beta_2(a_2)$ and
 - (2.2) For all $a_i \in A_i$ and $i = 1, 2$, $\beta_i(a_i) = \pi_i(a_i | f_i^t(a_i | h))$ where $f_i^t(a_i | h) = \langle a_i; \beta_{-i}(\cdot), x(\cdot) \rangle_{n_i}$ (see definition of $\pi_i(\cdot | \cdot)$ in Section 2).

In other words, the antecedent of X is the set of distributions over paths of lengths m from period t to period $t+m-1$ say, that would result if for every possible action profile in the current period t the *true* ensuing distribution over action profiles from period $t+1$ to period $t+m$ were to lie in the set X , assuming that (1) the period t limited forecasts of the players are *correct* and (2) the players behave according to the criterion as defined in Section 2. The mapping $x(\cdot)$ in the above definition is said to *sustain* the distribution $\langle \beta(\cdot), y(\cdot) \rangle \in W(X)$.

We first derive some preliminary properties of the antecedent operator, which are gathered in the following lemma:

Lemma 1. (i) (*Existence*) For all $X \subseteq \Psi_m$, the set $W(X)$ is non-empty. (ii) (*Monotonicity*) For all $X \subseteq X'$, $W(X) \subseteq W(X')$. (iii) (*Superadditivity*) $W(X) \cup W(X') \subseteq W(X \cup X')$. (iv) (*Continuity*) If X is a closed subset of Ψ_m , so is $W(X)$. For all closed subsets X of Ψ_m , $\lim_{X' \rightarrow X} W(X') \subseteq W(X)$.

The following theorem relates the antecedent operator to the set of distributions over paths of length m that can be sustained in (n_1, n_2) -limited forecast equilibria:

²² Observe that any distribution over paths of length $m+1$ can be described that way for some probability function $\beta(\cdot)$ and some distribution function $x(\cdot) \in \Psi_m$.

Theorem 1. (*Self-Generation*). Let S be the set of all possible distributions over paths of action profiles from period 1 to period m induced by (n_1, n_2) -limited forecast equilibria. S is the maximal set $X \subseteq \Psi_m$ such that $X \subseteq W(X)$.²³

We now provide a general algorithm for constructing all possible distributions over m -length paths induced by limited forecast equilibria.

Theorem 2. (*Algorithm*). The set S of all possible (n_1, n_2) -limited forecast equilibrium distributions over m length paths satisfies:

$$S = \lim_{k \rightarrow +\infty} X_k,$$

where $X_{k+1} = W(X_k)$ for all $k = 0, 1, \dots$ and $X_0 = \Psi_m$.

The above theorem provides an algorithm for the construction of the set S of all possible distributions over paths of action profiles from period 1 to period m induced by (n_1, n_2) -limited forecast equilibria. From the set S , we now explain how to construct all (n_1, n_2) -limited forecast equilibria. For every distribution $s \in S$, define $X(s) \equiv \{x(\cdot) : A \rightarrow S \text{ such that } x(\cdot) \text{ sustains } s \in W(S)\}$. Consider some $x^{(1)}(\cdot) \in X(s)$ for some $s \in S$. Consider next for each k and for each history h^k up to period k (i.e. h^k is of the form $(a^1, \dots, a^{k-1}) \in A^{k-1}$) some $x^{(k)}(\cdot | h^k) \in X(x^{(k-1)}(a^{k-1} | h^{k-1}))$ where h^{k-1} is the truncation of h^k to the history up to period $k-1$ (i.e. $h^{k-1} = (a^1, \dots, a^{k-2})$).²⁴ (The sequence $x^{(k)}(\cdot | h^k)$ is defined by forward induction on histories.) Finally, let $x^{(t-1)}(a^{t-1} | h^{t-1}) = \langle \beta(\cdot), y(\cdot) \rangle$ with $\beta(a) = \beta_1(a_1)\beta_2(a_2)$. Consider the strategy profile $\sigma = (\sigma_1, \sigma_2)$ (defined by forward induction on histories) such that for every history h^t at time t , and for every $a_i \in A_i$, $\sigma_i^t(h^t)[a_i] = \beta_i(a_i)$ as just defined. Then σ is a (n_1, n_2) -limited forecast equilibrium and every (n_1, n_2) -limited forecast equilibrium is of this form.

Comment. The complexity of the algorithm is that of determining $W(X)$ given $X \subseteq \Psi_m$, which corresponds to the complexity of deriving a fixed point over ΔA (i.e. a fixed point over (β_1, β_2) for each possible mapping $x(\cdot) : A \rightarrow X$).

4.2. Limited forecast equilibria in pure strategies

The construction of (n_1, n_2) -limited forecast equilibria in pure strategies parallels that in mixed strategies.

Let $m, m \geq n_i - 1$ for $i = 1$ and 2 . Let x^* be a m length (pure) path, i.e. $x^* \in A^m$. We denote by $[x^*]_k$ the truncation of x^* to the first k action profiles. For each action profile $a \in A$ and each $x^* \in A^m$, we denote by (a, x^*) the concatenation of a and x^* .

Definition 7. (The *antecedent operator). Let X^* be a subset of the set A^m of all pure m length paths. A m length path (a^*, y^*) where $a^* = (a_1^*, a_2^*) \in A$ and $y^* \in A^{m-1}$ belongs to $W^*(X^*)$ if and only if

- (1) There exists a mapping $x^*(\cdot) : A \rightarrow X^*$ such that $y^* = [x^*(a^*)]_{m-1}$ and
- (2) For $i = 1, 2$ and for all $a_i \in A_i$, $a_i \neq a_i^*$, where $a = (a_i, a_i^*)$ (see definition of $w_i(\cdot)$ in Section 2)

$$(1 - \delta)w_i((a^*, x^*(a^*))_{m_i}) + \delta^{n_i}u_i^{\text{inf}} \geq (1 - \delta)w_i((a, x^*(a))_{m_i}) + \delta^{n_i}u_i^{\text{sup}}. \quad (*)$$

23. Maximality is defined with respect to the inclusion relation.

24. For $k = 2$, h^0 should be interpreted as the null history and $x^{(1)}(a^1 | h^0)$ should be identified with $x^{(1)}(a^1)$ as just defined.

Observe that inequality (*) guarantees that if player i believes that by playing a_i (resp. $a_i \neq a_i^*$) in the current period t the stream of action profiles from period t to period $t + n_i - 1$ will be $(a^*, x^*(a^*))_{n_i}$ (resp. $(a, x^*(a))_{n_i}$ with $a = (a_i, a_i^*)$) then player i chooses to play a_i at period t irrespective of his state of mind (the most pessimistic (resp. optimistic) draw was taken for a_i^* (resp. a_i)). It follows that the *antecedent operator has the same interpretation as the antecedent operator except for the restriction to pure behavioural strategies in the current behaviour (a^* rather than $\beta(\cdot)$) and the restriction to deterministic m length path distributions for subsequent action profiles.

The mapping $x^*(\cdot): A \rightarrow X^*$ in the above definition will be referred to as *sustaining* $(a^*, y^*) \in W^*(X^*)$. It should be noted that the mapping $x^*(\cdot)$ can be chosen without loss of generality to satisfy $x^*(a) = \arg \min_{x^* \in X^*} w_i((a, x^*)_{n_i})$ for all $a = (a_i, a_i^*)$, $a_i \neq a_i^*$. (Note that this x^* may be chosen to be independent of a for $a = (a_i, a_i^*)$, $a_i \neq a_i^*$.) This is so because if there is some mapping sustaining (a^*, y^*) so does the induced mapping obtained by replacing $x^*(a)$ for $a = (a_i, a_i^*)$, $a_i \neq a_i^*$ by $\arg \min_{x^* \in X^*} w_i((a, x^*)_{n_i})$. (This is in spirit of Abreu (1988), see below the related literature.)

In contrast to Lemma 1, the set $W^*(X^*)$ may sometimes be empty. (To see this, consider discount factors δ sufficiently close to 1 with $u_i^{\text{sup}} > u_i^{\text{inf}}$. Then whatever the subset X^* , the mapping $x^*(\cdot): A \rightarrow X^*$, and the m length path (a^*, y^*) there is no way to satisfy (*).) When the set $W^*(X^*)$ is non-empty however, constructing $W^*(X^*)$ requires a finite number of checks.

We now provide the general tool for constructing all (n_1, n_2) -limited forecast equilibria in pure strategies (if any).

Theorem 3. *Let S^* be the set of all possible m length paths of action profiles from period 1 to period m induced by (n_1, n_2) -limited forecast equilibria in pure strategies. (i) (Self-generation) S^* is the maximal set $X^* \subseteq A^m$ such that $X^* \subseteq W^*(X^*)$. (ii) (Algorithm) $S^* = \lim_{k \rightarrow +\infty} X_k^*$, where $X_{k+1}^* = W^*(X_k^*)$ for all $k = 0, 1, \dots$ and $X_0^* = A^m$.²⁵*

From the set S^* , we now construct all (n_1, n_2) -limited forecast equilibria in pure strategies (if there are any). For every m length path $s^* \in S^*$, define $X^*(s^*) \equiv \{x^*(\cdot): A \rightarrow S^*$ such that $x^*(\cdot)$ sustains $s^* \in W^*(S^*)\}$. Consider some $x^{*(1)}(\cdot) \in X^*(s^*)$ for some $s^* \in S^*$. Consider next for each k and for each history h^k up to period k (i.e. h^k is of the form $(a^1, \dots, a^{k-1}) \in A^{k-1}$) some $x^{*(k)}(\cdot | h^k) \in X^*(x^{*(k-1)}(a^{k-1} | h^{k-1}))$ where h^{k-1} is the truncation of h^k to the history up to period $k-1$ (i.e. $h^{k-1} = (a^1, \dots, a^{k-2})$). (The sequence $x^{*(k)}(\cdot | h^k)$ is defined by forward induction on histories.) Finally, let $x^{*(t-1)}(a^{t-1} | h^{t-1}) = (a^*, y)$ with $a^* = (a_i^*, a_i^*)$. Consider the strategy profile $\sigma = (\sigma_1, \sigma_2)$ defined by forward induction on histories such that for every history h^t at time t , $\sigma_i^t(h^t)$ prescribes player i to play a_i^* in period t after history h^t (as just defined). Then σ is a (n_1, n_2) -limited forecast equilibrium in pure strategies and every (n_1, n_2) -limited forecast equilibrium in pure strategies is of this form.

5. APPLICATIONS

This section aims at analysing in various repeated game contexts the implications of limited forecast equilibria. We first consider repeated 2×2 coordination games. We then

25. There exists a (n_1, n_2) -limited forecast equilibrium in pure strategies whenever $\lim_{k \rightarrow +\infty} X_k^*$ is non-empty.

move to the analysis of repeated prisoner’s dilemmas. Finally, we briefly consider repeated battles of the sexes.

5.1. *Repeated coordination games*

Consider the repeated coordination game whose stage game payoffs are given by:

	<i>L</i>	<i>R</i>
<i>U</i>	<i>a, a</i>	<i>b, c</i>
<i>D</i>	<i>c, b</i>	<i>d, d</i>

with $a > c$ and $d > b$. We have $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$. Note that the stage game has two pure Nash equilibria (U, L) and (D, R) . When $a > d$, (U, L) Pareto-dominates (D, R) . When $a - c > d - b$, (U, L) is said to risk-dominate (D, R) (because there is a greater range of beliefs about the opponent’s action that justifies playing according to (U, L) rather than according to (D, R) , in a static version of the game, see Harsanyi and Selten (1988)).²⁶ Finally, both players have the same discount factor δ .

Let $n_i \geq 1$ be the length of foresight of player i . Player i ’s distribution of state of mind $\tilde{\epsilon}_i$ has full support on $[u_i^{\text{inf}}, u_i^{\text{sup}}]^2$ with $\min(a, b, c, d) \leq u_i^{\text{inf}} \leq u_i^{\text{sup}} \leq \max(a, b, c, d)$. We let $\Delta_i = u_i^{\text{sup}} - u_i^{\text{inf}}$, and for simplicity we assume that the supports are player-independent, *i.e.* for $i = 1, 2$, $u_i^{\text{inf}} = u^{\text{inf}}$, $u_i^{\text{sup}} = u^{\text{sup}}$ and $\Delta_i = \Delta$.

We wish to analyse when stationary pure strategies can be sustained as limited forecast equilibria. It should be clear that a stationary strategy profile that employs pure strategies *cannot* be sustained as a (n_1, n_2) -limited forecast equilibrium if it does *not* prescribe the players to behave (in each period and whatever history) according to one of the stage game Nash equilibria.²⁷ However, it need not be that the repetition of (U, L) or of (D, R) constitutes a (n_1, n_2) -limited forecast equilibrium. The following result makes clear the conditions under which each of these stationary strategy profiles is a (n_1, n_2) -limited forecast equilibrium.

Result

1. *The stationary strategy profile (U, L) is a (n_1, n_2) -limited forecast equilibrium if and only if*

$$(1 - \delta)(a - c) - \delta^{\min(n_1, n_2)} \Delta \geq 0. \tag{1}$$

2. *The stationary strategy profile (D, R) is a (n_1, n_2) -limited forecast equilibrium if and only if*

$$(1 - \delta)(d - b) - \delta^{\min(n_1, n_2)} \Delta \geq 0. \tag{2}$$

A few comments are in order. First, for n_1, n_2 and $\Delta > 0$ fixed, observe that none of conditions (1) or (2) are satisfied when δ is close enough to 1. The reason is that the weight of the unpredicted component—which is noisy (because $\Delta > 0$)—then dominates in the criterion of the players, and therefore it is impossible to sustain pure behaviours as (n_1, n_2) -limited forecast equilibrium outcomes. For sufficiently low values of δ , however,

26. See Carlsson–van Damme (1993), Kandori–Mailath–Rob (1993), Morris–Rob–Shin (1995) for various arguments in favour of the risk-dominance criterion.

27. To see this, it suffices to consider the behaviour of player i with state of mind ϵ_i such that $\epsilon_i(a_i)$ coincides for all actions $a_i \in A_i$.

both conditions (1) and (2) are satisfied. The reason is that then the weight of the unpredicted component is very small so that the static incentives drive the behaviour of the players. More interesting is the case of intermediate discount factors. Then it may well be that only one of conditions (1) or (2) is satisfied. Which condition is more likely to be satisfied depends on the comparison between $a - c$ and $d - b$. Specifically, if $a - c > d - b$, the stationary strategy profile (U, L) (UL to make it short) is a (n_1, n_2) -limited forecast equilibrium whenever DR is a (n_1, n_2) -limited forecast equilibrium; sometimes UL is a (n_1, n_2) -limited forecast equilibrium and not DR . Increasing $\min(n_1, n_2)$ while keeping all other parameters of the model fixed yields the following:²⁸ For intermediate values of $\min(n_1, n_2)$, UL is a (n_1, n_2) -limited forecast equilibrium and not DR whenever $a - c > d - b$ (and condition (2) is not satisfied with $\min(n_1, n_2) = 1$).

The above considerations thus suggest that, in repeated coordination games, the approach of limited foresight gives some support to the selection of the repetition of the risk-dominant stage game Nash equilibrium over the repetition of the other stage game Nash equilibrium (in particular, the Pareto-dominant stage game Nash equilibrium if they differ). The logic for this result is however very different from the logic developed by Harsanyi and Selten in static versions of coordination games (their logic in terms of risk has been briefly summarized above). In the present context, players are certain (because they have learned it, say) of the behaviour of their opponent within a given horizon of foresight, and they are ignorant of what will come next. The incentives to comply with stationary pure strategies are driven by aggregating the current stage incentives to follow the assumed strategy and the (random) state of mind that stands for the assessment of what is to come beyond the horizon of foresight (as a function of the current action). The repetition of the risk-dominant stage game Nash equilibrium is more easily sustainable as an (n_1, n_2) -limited forecast equilibrium than the repetition of the other stage game Nash equilibrium because the static incentives to comply with the former are greater than those of the latter, and the magnitude of the noise beyond the horizon of foresight is assumed to be the same in both cases.

Comment. An interesting extension of the limited foresight framework proposed in Section 2 is one where we allow the players to have different horizons of foresight when considering actions on and off the equilibrium path.²⁹ Anticipating on future work that allows for such an extension, assume that for actions *on* the equilibrium path, players have *perfect* foresight while for actions *off* the equilibrium path the length of foresight of player i is *limited* and equal to n_i ; player i 's assessment of the effect of the action off the equilibrium path beyond the horizon of foresight is described by player i 's state of mind, which is stochastic and takes values in $[u^{\text{inf}}, u^{\text{sup}}]$. Consider again the two stationary strategy profiles UL and DR . The repetition of (U, L) is an (n_1, n_2) *-limited forecast equilibrium (in this new sense) if and only if

$$(1 - \delta)(a - c) - \delta^{\min(n_1, n_2)}(u^{\text{sup}} - a) \geq 0. \quad (3)$$

Similarly, the repetition of (D, R) is an (n_1, n_2) *-limited forecast equilibrium if and only if

$$(1 - \delta)(d - b) - \delta^{\min(n_1, n_2)}(u^{\text{sup}} - d) \geq 0. \quad (4)$$

28. Observe that the mere effect of increasing $\min(n_1, n_2)$ is to reduce the weight of the unpredicted component, since we are considering stationary strategies.

29. This can be motivated by the idea that learning the effect of on the equilibrium path actions does not require exogenous experimentation, and is thus easier.

Thus in this modified framework, the repetition of the Pareto-dominant stage game Nash equilibrium is always an $(n_1, n_2)^*$ -limited forecast equilibrium (because $\max(a, d) \geq u^{\text{sup}}$ so that even optimistic assessments of the continuation value associated with the off the equilibrium path action are not enough to destabilize the Pareto-dominant strategy profile) while the repetition of the other stage game Nash equilibrium need not be an $(n_1, n_2)^*$ -limited forecast equilibrium (for some parameter values). When players have limited horizons of foresight for actions both on and off the equilibrium path, but have a longer horizon of foresight for actions on the equilibrium path, then the stationary strategy profile that can most easily be sustained as a limited forecast equilibrium will be the one that maximizes some weighted average between the Pareto-dominance criterion and the risk-dominance criterion.

5.2. *Repeated prisoner's dilemma*

We consider repeated symmetric prisoner's dilemmas whose stage game payoffs are given by:

	<i>C</i>	<i>D</i>
<i>C</i>	1, 1	- <i>l</i> , 1 + <i>g</i>
<i>D</i>	1 + <i>g</i> , - <i>l</i>	0, 0

with $g > 0$ and $l > 0$. We have $A_i = \{C, D\}$ for $i = 1, 2$. Both players have the same discount factor δ . The length of foresight of player i is denoted by n_i . Player i 's distribution of state of mind $\tilde{\epsilon}_i$ has full support on $[u_i^{\text{inf}}, u_i^{\text{sup}}]^2$ with $-l \leq u_i^{\text{inf}} \leq u_i^{\text{sup}} \leq 1 + g$. We let $\Delta_i = u_i^{\text{sup}} - u_i^{\text{inf}}$, and for simplicity we assume that for $i = 1, 2$, $u_i^{\text{inf}} = u^{\text{inf}}$, $u_i^{\text{sup}} = u^{\text{sup}}$ and $\Delta_i = \Delta$.

5.2.1. *Stationary vs. non-stationary strategies.* For $n_i \geq 1$, $i = 1, 2$, the repetition of the stage game Nash equilibrium in which each player $i = 1, 2$ plays D irrespective of the past is a (n_1, n_2) -limited forecast equilibrium if and only if

$$(1 - \delta)l - \delta^{\min(n_1, n_2)} \Delta \geq 0. \tag{5}$$

For $n_1, n_2, \Delta > 0$ fixed, condition (5) defines a threshold discount factor $\delta(n_1, n_2)$ such that if $\delta \leq \delta(n_1, n_2)$ the defect strategy profile constitutes a (n_1, n_2) -limited forecast equilibrium, and if $\delta > \delta(n_1, n_2)$ it does not.

We wish to explore here whether there may exist (n_1, n_2) -limited forecast equilibria in pure strategies when the defect strategy profile fails to be a (n_1, n_2) -limited forecast equilibrium. To this end, we will assume that for $i = 1, 2$, player i 's length of foresight n_i is no smaller than 2.³⁰ We will consider the following strategy σ_i^* for player i . Player i 's behavioural strategies depend only on the action profile in the previous period: When the last period action profile is (C, C) or (D, D) player i plays C ; when the last period action profile is (C, D) or (D, C) player i plays D (in the first period, player i plays C , say).

Observe that the path generated by $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is the infinite repetition (C, C) . If one player deviates, say player 1 plays D , then in the next stage (D, D) is played after which (C, C) is played again for ever.

30. If $n_i \leq 1$ for $i = 1$ or 2 , then one of the players makes no prediction about the future and only the defect strategies can possibly be sustained (because no punishment is forecast by at least one player).

For $n_i \geq 2$, $i = 1, 2$, the strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a (n_1, n_2) -limited forecast equilibrium if and only if³¹

$$(1 - \delta)(-g + \delta) - \delta^{\min(n_1, n_2)} \Delta \geq 0. \quad (6)$$

It is instructive to compare conditions (5) and (6). Keeping $n_1, n_2, \Delta > 0$ fixed, none of conditions (5) and (6) are satisfied for δ close to 1, since again the unpredicted component (which is driven by the noisy state of mind) dominates players' criterion. For δ close to 0, only condition (5) is met, since then players do not care about the future and thus only the repetition of the stage game Nash equilibrium can be sustained.

For intermediate values of δ , it may well be thought that condition (6) is met and not condition (5). In such cases, a pure cooperative path can be obtained as a (n_1, n_2) -limited forecast equilibrium while a purely non-cooperative path cannot.

The reason for this seemingly surprising result is as follows. Sustaining the repetition of a purely non-cooperative behaviour, on the one hand, is not possible when $\delta > \delta(n_1, n_2)$ because of the magnitude of the noisy character of the assessment of what is to come beyond the horizon of foresight. The difficulty in sustaining the non-cooperative behaviour lies in the fact that the current period incentives to comply with the defect strategies are not high enough relative to the noisy character of the state of mind of the players: If player i were to forecast that player $-i$ is to play D in the current period and that whatever the action profile today the action profile within his horizon of foresight is (D, D) , he would sometimes play C today for those states of mind which are optimistic when he assesses the beyond horizon of foresight effect of C and pessimistic when he assesses D , i.e. $\varepsilon_i(C) \approx u_i^{\text{sup}}$ and $\varepsilon_i(D) \approx u_i^{\text{inf}}$, which in turn would yield a contradiction to the consistency of forecasts (because this implies that sometimes in the next period action C is played).

The strategy profile σ^* , on the other hand, includes a punishment phase of one period (one more period of (D, D) instead of (C, C)) in case of deviation from the assumed equilibrium behaviour. Since the players do forecast what is to come in the next period (their horizon of foresight is no smaller than 2), they do perceive the 1 period punishment phase in case of deviation. When condition (6) is satisfied, this one period punishment is sufficient to outweigh the effect of the state of mind (regarding the assessment of what is to come beyond the horizon of foresight), and thus a purely cooperative path can be sustained as a (n_1, n_2) -limited forecast equilibrium.³²

Comments:

1. The main insight from the above considerations is that sometimes pure non-stationary strategies are more easily sustainable as (n_1, n_2) -limited forecast equilibria than pure stationary strategies are. The reason is that by playing on the length of the punishment phase (as long as it remains within the horizon of foresight) non-stationary strategies may increase the incentives to comply with the assumed equilibrium behaviour within the horizon of foresight, which in turn may allow to better neutralize the noisy character of the states of mind of the players (regarding what they do not forecast).

31. This condition is obtained when checking players' incentives to play C after (C, C) or (D, D) whatever their state of mind. The incentives to play D after (C, D) or (D, C) are then automatically satisfied (i.e. condition (6) implies that $(1 - \delta)(I + \delta) - \delta^{\min(n_1, n_2)} \Delta \geq 0$).

32. Of course, incentives have to be checked in every subgame, and that is why cooperation cannot be sustained using trigger strategies whenever condition (5) is not met. The strategy profile σ^* allows for forgiveness after the punishment phase.

2. We have briefly mentioned above (in the application to coordination games) the extension to the case where the horizon of foresight is longer for actions on the equilibrium path than for actions off the equilibrium path. When each player i has perfect foresight for actions on the equilibrium path and has a n_i -length horizon of foresight for actions off the equilibrium path, the repetition of the defect strategies is a $(n_1, n_2)^*$ -limited forecast equilibrium if and only if

$$(1 - \delta)l - \delta^{\min(n_1, n_2)} u^{\sup} \geq 0, \tag{7}$$

where u^{\sup} denotes the supremum of players' states of mind regarding the per period payoff induced by the off the equilibrium path action beyond the horizon of foresight. Similarly, σ^* is a $(n_1, n_2)^*$ -limited forecast equilibrium if and only if

$$(1 - \delta)(-g + \delta) - \delta^{\min(n_1, n_2)} (u^{\sup} - 1) \geq 0. \tag{8}$$

There are even more parameter values for which condition (8) is satisfied and not (7), thus reinforcing the insight that non-stationary strategies may be easier to sustain than stationary strategies when players have limited foresight.

5.2.2. A complete construction. In this part, we apply the general construction techniques of Section 4 to the repeated symmetric prisoner's dilemma described above with $g = 0.05$, $l = 0.2$. We will consider the case where each player has a length of foresight equal to 2, i.e. $n_i = 2$ for $i = 1, 2$ and the support of each player i 's state of mind satisfies $u^{\inf} = -l$, $u^{\sup} = 1 + g$, yielding $\Delta = 1 + g + l = 1.25$.

We apply the general technique of Subsection 4.2. The analysis bears over m length paths where $m \geq \min(n_1, n_2) - 1$. We choose the minimum such m so that we set $m = 1$. We refer to a^{CC} , a^{CD} , a^{DC} , a^{DD} as the action profiles (C, C) , (C, D) , (D, C) , and (D, D) , respectively.

The general algorithm technique of Subsection 4.2 yields:

$$S^*(\delta) = \begin{cases} \{a^{DD}\}, & \text{for } \delta < \delta_1, \\ \{a^{DD}, a^{CC}\}, & \text{for } \delta_1 \leq \delta < \delta_2, \\ \{a^{DD}, a^{CD}, a^{DC}, a^{CC}\}, & \text{for } \delta_2 \leq \delta \leq \delta_3, \\ \{a^{DD}, a^{CC}\}, & \text{for } \delta_3 < \delta \leq \delta_4, \\ \emptyset, & \text{for } \delta_4 < \delta, \end{cases}$$

where $\delta_1 \approx 5.3828 \times 10^{-2}$,³³ $\delta_2 \approx 0.22597$,³⁴ $\delta_3 \approx 0.35403$,³⁵ and $\delta_4 \approx 0.41284$.³⁶

Comment. One may wonder whether, for other parameter values of g , l , and Δ , limit sets S^* other than those displayed for the example may arise. For every g , l , and Δ , we obtain that if S^* is a singleton it must be $\{a^{DD}\}$ (this results from the standard insight that in order to sustain something other than the repetition of the stage game Nash equilibrium there must be several possible equilibrium paths). We also obtain that it cannot be that S^* coincides with $\{a^{CC}, a^{CD}, a^{DC}\}$ (this is because if $\{a^{CC}, a^{CD}, a^{DC}\} \subseteq$

33. δ_1 is the smallest solution to $(1 - \delta)(\delta - g) - \delta^2(1 + g + l) = 0$ (allowing to sustain a^{CC} in S^*).
 34. δ_2 is the smallest solution to $(1 - \delta)(\delta - l + \delta(g + l)) - \delta^2(1 + g + l) = 0$ (allowing to sustain a^{CD} and a^{DC} in S^*).
 35. δ_3 is the largest solution to $(1 - \delta)(\delta - l + \delta(g + l)) - \delta^2(1 + g + l) = 0$ (so that for a larger discount factor a^{CD} and a^{DC} can no longer be sustained in S^*).
 36. δ_4 is the largest solution to $(1 - \delta)(\delta - g) - \delta^2(1 + g + l) = 0$ so that for larger discount factors a^{CC} and a^{DD} can no longer be sustained in S^* .

$W^*\{a^{CC}, a^{CD}, a^{DC}\}$ then necessarily $a^{DD} \in W^*\{a^{CC}, a^{CD}, a^{DC}\}$). Other than that, every possibility can arise depending on the parameter sets.

Description of (2, 2)-equilibrium strategies for $\delta_3 < \delta \leq \delta_4$. The description of (n_1, n_2) -limited forecast equilibria in pure strategies for every value of δ appears in Jehiel (1998b). We only report here the results when $\delta_3 < \delta \leq \delta_4$ for which $S^*(\delta) = \{a^{DD}, a^{CC}\}$. Sustaining a^{CC} can only be done using a mapping $x_2^*(\cdot)$ such that $x_2^*(a^{CC}) = a^{CC}$ and $x_1^*(a^{CD}) = x_1^*(a^{DC}) = a^{DD}$. Sustaining a^{DD} can only be done using a mapping $x_1^*(\cdot)$ such that $x_1^*(a^{DD}) = a^{CC}$ and $x_1^*(a^{CD}) = x_1^*(a^{DC}) = a^{DD}$. (Observe that a^{DD} cannot be sustained using the mapping $x_0^*(\cdot)$ such that for every $a \in A$, $x_0^*(a) = a^{DD}$ because $\delta > \delta(2, 2) \approx 0.32792$.)

Thus there are two (2, 2)-limited forecast equilibrium pure paths, which are $(a^{CC}, a^{CC}, a^{CC}, \dots)$ and $(a^{DD}, a^{CC}, a^{CC}, \dots)$. That is, either there is an infinite repetition of a^{CC} or there is a first stage with a^{DD} followed by an infinite repetition of a^{CC} . The corresponding strategy profile employs strategies with bounded recall of size one. That is, the behaviour strategy depends only on the play in the previous period (it does not depend on the time period t either). Player i plays C whenever in the previous stage the action profile is either (C, C) or (D, D) . He plays D whenever in the previous stage the action profile is either (C, D) or (D, C) .³⁷ The two pure paths correspond to the set of paths generated by these strategies after any possible history.

Discussion. For $\delta_3 < \delta \leq \delta_4$ the above analysis shows that there is essentially a unique (2, 2)-limited forecast equilibrium in pure strategies.³⁸ Accordingly, the lowest payoff induced by any (2, 2)-limited forecast equilibrium in pure strategies corresponds to the path $(a^{DD}, a^{CC}, a^{CC}, \dots)$. Insofar as we are interested in limited forecast equilibria in pure strategies (because they are robust to changes in the distributions of states of mind, see above Sections 2 and 3), limited foresight forces cooperation for these parameter values.

Considering other lengths of foresight yields the following. Increasing the length of foresight has several effects. On the one hand, it allows the players to forecast a punishment phase in case of deviation over a longer horizon, which may facilitate in some cases the possibility of cooperation. On the other hand, it reduces the weight of the unpredicted component (the state of mind), which may also facilitate the sustainability of the defect strategies.³⁹

One may be interested in how the possibility of cooperation is affected when the players have different lengths of foresight, say $n_1 = 2$ for player 1 and $n_2 = 3$ for player 2. For the above numerical example, we derive that σ^* is still the unique (2, 3)-limited forecast equilibrium in pure strategies for $\delta_3 < \delta \leq \delta_4$.⁴⁰

Finally, going through the characterization of the set of limited forecast equilibria in mixed strategies will in general require some heavy calculations (*i.e.* iterative application of the algorithm shown in Theorem 2). It is worth noting though that in some cases it is possible to sustain the repetition of the stage game cooperative action profile even though

37. This corresponds to the strategy profile σ^* introduced in Subsection 5.2.1.

38. Strictly speaking, there are two of them depending on the state of the strategy profile σ^* from which period 1 starts.

39. In our numerical example, we derive that the defect strategies constitute a (3, 3)-limited forecast equilibrium (*i.e.* $n_i = 3$ for $i = 1, 2$) whenever $\delta \leq \delta(3, 3) \approx 0.4459$. We also observe that a cooperative path can be sustained for $5.0166 \times 10^{-2} \leq \delta \leq 0.5638$. Moreover, for $0.4459 < \delta \leq 0.5638$, we obtain $S^* = \{(a^{DD}, a^{CC}), (a^{CC}, a^{CC}), (a^{DD}, a^{DD}), (a^{CC}, a^{CC}), \dots, (a^{CC}, \dots)\}$.

40. More generally, it can be shown that, in the repeated symmetric prisoner's dilemma, the set of symmetric limited forecast equilibria in pure strategies is determined by $\min(n_1, n_2)$.

there is no limited forecast equilibrium in pure strategies (i.e. $(a^{CC}, a^{CC}, \dots, a^{CC}) \in S$ is compatible with $S^* = \emptyset$).

5.3. Other applications

Other applications of interest include the repeated battle of the sexes whose stage game payoffs are given by:

	L	R
U	$1+k, 1$	$0, 0$
D	$0, 0$	$1, 1+k$

with $k > 0$. Because of space constraints, the complete analysis will not be provided here. However, in the same vein as the insights developed for the repeated prisoner's dilemma, we would obtain that for some intermediate values of the discount factor and of the lengths of foresight it is not possible to sustain the infinite repetition of either of the stage game pure Nash equilibrium action profile $((U, L), \dots (U, L) \dots$ or $(D, R), \dots (D, R) \dots)$ as a limited forecast equilibrium path while it is possible to sustain the alternation between the two stage game Nash equilibrium action profiles $((U, L), (D, R), \dots (U, L), (D, R) \dots)$.

6. RELATIONSHIP WITH THE LITERATURE

On a technical level, the general construction method proposed in Section 4 bears some similarity with the techniques used in Abreu (1988) and Abreu *et al.* (1990) for the case of perfectly rational players. An essential difference with the work of Abreu *et al.* (1990) is that here it is not possible to make the analysis over continuation values because the horizon of the players is *rolling*. We therefore had to make the recursive analysis using continuation distributions over truncated paths rather than continuation values. Another difference is that mixed behavioural strategies were not ruled out from our analysis. Also, we obtained a very simple way to construct all (n_1, n_2) -limited forecast equilibria in pure strategies if any (the algorithm of Theorem 3 ends in a finite number of steps), which should be contrasted with the possibly great complexity of finding the optimal penal codes in Abreu (1988).

It may be instructive to relate the results of Subsection 5.2 to the other existing results on the repeated prisoner's dilemma. In standard paradigms where the players are assumed to be perfectly rational, a cooperative path can be sustained as a Subgame Perfect Nash Equilibrium when the players are sufficiently patient (see Friedman (1971), Aumann and Shapley (1976), Rubinstein (1979), Fudenberg and Maskin (1986)). However, even when the cooperative outcome can be sustained as a SPNE the repetition of the stage game Nash equilibrium is also a SPNE, and therefore that theory does not ensure that a cooperative outcome will indeed obtain (it only shows that a cooperative behaviour cannot be ruled out). This should be contrasted with our finding (Subsection 5.2) that sometimes only cooperative paths can be sustained as equilibrium paths of limited forecast equilibria in pure strategies.

We next relate our approach to the literature on bounded rationality in repeated games. The most studied approach to bounded rationality in repeated games is that of automaton theory. Neyman (1985) takes as given the complexity of the strategies that the players may employ and shows how in the finitely repeated prisoner's dilemma game such a constraint may allow the players to cooperate (again cooperation is a possibility but

need not arise in equilibrium). Rubinstein (1986) and Abreu and Rubinstein (1988) analyse how the equilibrium strategies are constrained when the players care both about the complexity of the automaton they use to implement their strategies and the payoff they get (see also Kalai and Stanford (1988)).⁴¹

The links between the finite automaton approach and the limited foresight approach are not immediate. Of course, in a given environment, the equilibrium strategies of a limited forecast equilibrium can (in general) be described as finite automata, but keeping player i 's length of foresight fixed while changing the environment (either through the stage game payoff or through the discount factor or through the length of foresight of the other player n_{-i}) will in general result in equilibrium strategies for player i that need not have the same complexity in the sense of automaton theory.

Finally, limited foresight in repeated alternate-move games has been previously considered in Jehiel (1995). Apart from the restriction to alternate-move games, a key difference with the present paper bears on the modelling of continuation values beyond the horizon of foresight (implicitly they are deterministic and/or the same for all actions in Jehiel (1995)). Jehiel (1998a) proposes a learning model to justify the consistency requirement imposed on equilibrium limited forecasts in Jehiel (1995).

In his discussion of limited foresight, Rubinstein (1998), Chapter 7 (together with the concept proposed in Jehiel (1995)) suggests an alternative solution concept (see also Jehiel (1998a, Subsection 6.5)). In that concept, players know the reaction function of their opponent within their horizon of foresight, and decide on the current action by solving a finite horizon game as induced by the horizon of foresight and the associated opponent's reaction function. The finite path generated by the resolution of this finite horizon game can naturally be interpreted as a plan of actions (within the horizon of foresight) made by the player. A drawback that results from this concept is that players base their choices of current actions on plans of actions that they need not follow afterwards—a time-inconsistency like problem that does not arise with the solution concept considered in this paper.

7. CONCLUSION

This paper has shown how limited foresight considerations may change our understanding of strategic interactions even in simple games like repeated coordination games or repeated prisoner's dilemmas. It would clearly be of interest now to investigate the implications limited foresight might have in more elaborate real life strategic interaction contexts.

APPENDIX

Proof of Proposition 1.

We prove the existence of a stationary (n_1, n_2) -limited forecast equilibrium by standard application of Brouwer's fixed point theorem. Consider a stationary strategy profile $\sigma = (\sigma_1, \sigma_2)$, and (with some abuse of notation) denote by $\sigma_i[a_i]$ the probability that player i plays a_i in every period t according to σ_i . Denote by $\tilde{\sigma}_i(\sigma_{-i})$ the stationary player i 's strategy defined from σ_{-i} by

$$\tilde{\sigma}_i(\sigma_{-i})[a_i] = \Pr \left\{ (1 - \delta) \sum_{a_{-i}} u_i(a_i, a_{-i}) \sigma_{-i}[a_{-i}] + \delta^n \varepsilon_i(a_i) > (1 - \delta) \sum_{a'_{-i}} u_i(a'_i, a_{-i}) \sigma_{-i}[a'_{-i}] + \delta^n \varepsilon_i(a'_i) \quad \text{for all } a'_i \neq a_i \right\}.$$

41. The restrictions imposed by the Abreu–Rubinstein equilibrium payoffs follow a logic that is very different from the one highlighted in Section 5.2. Within our framework, the restriction follows from robustness considerations regarding what players do not forecast; in Abreu–Rubinstein the restriction follows from optimization considerations on the complexity of the strategies.

In the class of stationary strategies, the mappings $\sigma_{-i} \rightarrow \tilde{\sigma}_i(\sigma_{-i})$ are continuous (because the distributions of $\tilde{\varepsilon}_i$, ε_{-i} are assumed to have no mass points). By application of Brouwer (or Nash) theorem, the continuity ensures that there exists a stationary strategy profile $\sigma = (\sigma_1, \sigma_2)$ such that $\tilde{\sigma}_1(\sigma_2) = \sigma_1$ and $\tilde{\sigma}_2(\sigma_1) = \sigma_2$. It is readily verified that such a σ is a (n_1, n_2) -limited forecast equilibrium. ||

Proof of Proposition 2.

The topology of the set of time-independent strategy profiles with bounded recall of size N is that of \mathfrak{R}^M with

$$M = [(\#A_1 \#A_2)^N]^{A_1 + A_2 - 2},$$

(the number of histories of size N is $(\#A_1 \#A_2)^N$ and for every history of size N the behavioural strategy of each player i must assign a probability vector on A_i). That $(\sigma^{r_i, r'_i})_r$ admits an accumulation point follows from well-known topological results (a sequence of bounded elements of a compact set of \mathfrak{R}^M must admit an accumulation point). Consider an arbitrary accumulation point σ of the sequence $(\sigma^{r_i, r'_i})_r$. First, σ is a time-independent strategy profile with bounded recall of size N (because the set of such strategy profiles is compact). Second, σ is a SPNE: To see this, consider an arbitrary N -length history h , and assume that player i 's strategy σ_i assigns him to play a_i with positive probability, i.e. $\sigma_i(h)[a_i] > 0$. By definition of an accumulation point it must be that for an infinite sequence of r there are draws of states of mind ε_i^r such that $c_i(a_i | f_i^r, h, \varepsilon_i^r) > c_i(a'_i | f_i^r, h, \varepsilon_i^r)$ for all $a'_i \neq a_i$ where f_i^r is player i 's forecasting rule associated with σ^{r_i, r'_i} (because of consistency f_i^r must also be time-independent with bounded recall of size N). For any draw ε_i and for every action $a'_i \in A_i$ the associated sequence $c_i(a'_i | f_i^r, h, \varepsilon_i)$ must converge to $\tilde{u}_i(q(\sigma|_h, a'_i))$ where $\tilde{u}_i(q)$ stands for the discounted sum of payoffs induced by the infinite path q and $q(\sigma|_h, a'_i)$ stands for the infinite path generated by σ after the N -length history h when player i plays a'_i . It thus follows that $\tilde{u}_i(q(\sigma|_h, a_i)) \geq \tilde{u}_i(q(\sigma|_h, a'_i))$ for all $a'_i \neq a_i$. Hence σ is a SPNE. ||

Proof of Proposition 3.

Consider a SPNE σ with locally strict uniform incentives in which the incentives are bounded by α as in Definition 5. Take n^* to be such that

$$2\delta^{n^*} \left(\max_{a \in A} u_i(a) - \min_{a \in A} u_i(a) \right) > \alpha,$$

for $i = 1, 2$. Then it is readily verified that σ is a (n_1, n_2) -limited forecast equilibrium for all $n_1, n_2 > n^*$. (The difference between the true assessment and the proxy assessment for player i is at most $\max_{a \in A} u_i(a) - \min_{a \in A} u_i(a)$ for each current action player i may choose.) ||

Proof of Proposition 4.

Let σ be a (n_1, n_2) -limited forecast equilibrium in pure strategies. For each possible action a_i player i may choose and for each history h , the true continuation value beyond player i 's horizon of foresight must lie in $(u_i^{\text{inf}}, u_i^{\text{sup}})$. Thus, player i will optimally decide to play according to σ even if he has a correct assessment of the continuation value beyond his horizon of foresight. Similarly, σ is a (n'_1, n'_2) -limited forecast equilibrium (in pure strategies) for all (n'_1, n'_2) , $n'_1 \geq n_1$, $n'_2 \geq n_2$ because the weighted average between the true (i.e. according to σ) expected average payoff obtained within $n'_i - n_i$ periods and the subsequent periods beyond the n'_i horizon of foresight must lie in $(u_i^{\text{inf}}, u_i^{\text{sup}})$, and is therefore a possible realization of the state of mind. ||

Proof of Lemma 1.

(i) Consider an element x^* of X . Let σ_i, σ_{-i} denote the behavior strategies induced (in every period) by the stationary (n_1, n_2) -limited forecast equilibrium constructed in Proposition 1. Let $\beta_j \in \Delta A_j$ be such that $\beta_j(a_j) = \sigma_j[a_j]$ for all a_j . Let $\beta \in \Delta A$ and $\gamma(\cdot) \in \Psi_{m-1}$ be such that $\beta(a) = \beta_1(a_1)\beta_2(a_2)$ and $\gamma(a) = [x^*]_{m-1}$ for all $a = (a_1, a_2) \in A$. It is readily verified that the distribution $\langle \beta(\cdot), \gamma(\cdot) \rangle$ belongs to $W(X)$, and therefore $W(X)$ is non-empty. The monotonicity (ii) and superadditivity (iii) properties follow trivially from the very definition of the antecedent operator. (iv) That $W(X)$ is closed when X is closed follows from the continuity of the function $\pi_i(\cdot | \cdot)$ with respect to its second argument. We now show that $\lim_{X^k \rightarrow X} W(X^k) \subseteq W(X)$ when X is closed. Consider a sequence $\langle \beta^{(k)}(\cdot), \gamma^{(k)}(\cdot) \rangle \in W(X^{(k)})$ where $X^{(k)}$ converges to X as k tends to infinity. Any accumulation point of this sequence is an element of $W(X)$ (because X is closed so that the mappings $x^{(k)}(\cdot) : A \rightarrow X^{(k)}$ have accumulation points in mappings from A to X and the function $\pi_i(\cdot | \cdot)$ is continuous with respect to its second argument). ||

Proof of Theorem 1.

(0) As a preliminary step, observe that there exists a unique maximal set \bar{X} such that $\bar{X} \subseteq W(\bar{X})$ because of the superadditivity and the continuity property ($X \subseteq W(X)$ and $X' \subseteq W(X')$ imply that $X \cup X' \subseteq W(X \cup X')$; together with the continuity property it shows that the closure of the union of all sets X such that $X \subseteq W(X)$ is the only maximal such set).

(1) We first establish that $S \subseteq \bar{X}$. To this end, note that $S \subseteq W(S)$ follows from the observation that the distributions over m length paths from period 2 to period $m+1$ induced by limited forecast equilibria must belong to the set S whatever the first period action profile (the strategy profile induced by a limited forecast equilibrium after an arbitrary first period action profile is a limited forecast equilibrium). We thus have $S \subseteq \bar{X}$ by definition of \bar{X} .

(2) We next establish that $\bar{X} \subseteq S$. The proof is constructive. For every distribution $x \in \bar{X}$ let $\bar{x}(\cdot | x)$ be a mapping from A to \bar{X} that allows to sustain $x \in W(\bar{X})$ ($\supseteq \bar{X}$) (see definition of $W(\bar{X})$). Consider a distribution x in \bar{X} . Let $h = (a^1, \dots, a^{t-1})$ be an arbitrary history at time t . For every $k = 2, \dots, t$, we denote by h^k the corresponding history of play up to time k , i.e. $h^k = (a^1, \dots, a^{k-1})$. (The history at time 1 is denoted by $h^1 = \emptyset$.) First we define $x^1(\cdot) \equiv \bar{x}(\cdot | x)$. We next define inductively $x^k(\cdot | h^k) \equiv \bar{x}(\cdot | x^{k-1}(a^{k-1} | h^{k-1}))$ for each $k = 2, \dots, t-1$. Finally, we let $x^{t-1}(a^{t-1} | h^{t-1}) = \langle \beta(\cdot), y(\cdot) \rangle$ with $\beta(a) = \beta_1(a_1)\beta_2(a_2)$ (remember that $x^{t-1}(a^{t-1} | h^{t-1}) \in \bar{X} \subseteq W(\bar{X})$ and see definition of $W(\bar{X})$). Consider the strategy profile $\sigma = (\sigma_1, \sigma_2)$ such that for every history h at time t , and for every $a_i \in A_i$, $\sigma_i^t(h)[a_i] = \beta_i(a_i)$ as just defined. Then it is readily verified that the distribution induced by σ over the first m action profiles is x , and σ is a (n_1, n_2) -limited forecast equilibrium. Thus $\bar{X} \subseteq S$ as desired. ||

Proof of Theorem 2

(i) We first establish by induction that for all k , $S \subseteq X_k$: $S \subseteq X_0$ by definition of X_0 ; if $S \subseteq X_k$ then by Lemma 1 it follows that $W(S) \subseteq W(X_k)$, which writes $W(S) \subseteq X_{k+1}$ by definition of X_{k+1} and thus $S \subseteq X_{k+1}$ since $S \subseteq W(S)$ by Theorem 1. (ii) We next establish by induction that the sequence $(X_k)_k$ is nested in the sense that $X_0 \supseteq X_1 \supseteq \dots \supseteq X_k$ for all k : If $X_{k+1} \supseteq X_k$ then by Lemma 1 it follows that $W(X_{k+1}) \subseteq W(X_k)$, which writes $X_{k+2} \subseteq X_{k+1}$. (iii) Each X_k is closed (by iterative application of Lemma 1(iii)), and the sequence $(X_k)_k$ is nested. Denote by X_∞ the set of all accumulation points in the sequence $(X_k)_k$. By application of standard topology arguments we have that X_∞ is closed and $X_\infty = \lim_{k \rightarrow \infty} X_k$, which implies that $\lim_{k \rightarrow \infty} W(X_k) \subseteq W(X_\infty)$ (by the continuity property) or equivalently $X_\infty \subseteq W(X_\infty)$ since $\lim_{k \rightarrow \infty} W(X_k) = \lim_{k \rightarrow \infty} X_{k+1} = X_\infty$. Thus, by Theorem 1, $X_\infty \subseteq S$ since S is the maximal set \bar{X} such that $\bar{X} \subseteq W(\bar{X})$. Using (i) (which implies that $S \subseteq X_\infty$), we may conclude that $S = X_\infty$ as required. ||

Proof of Theorem 3.

The proof follows from easy adaptations of the proof of Theorem 2. It is therefore omitted. ||

Acknowledgements. I would like to thank seminar participants in Paris (ENS and Roy), Gerzensee ESSET 1997, EEA 1997, UCL, ECARE, Tel Aviv, Haifa (Technion), Toulouse IDEI, Cambridge, ESRC game theory workshop, Northwestern University, MIT-Harvard as well as Hyun Shin and two anonymous referees for helpful comments.

REFERENCES

- ABREU, D. (1988), "Toward a Theory of Discounted Repeated Games", *Econometrica*, 56, 383–396.
 ABREU, D., PEARCE, D. and STACCHETTI, E. (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", *Econometrica*, 58, 1041–1064.
 ABREU, D. and RUBINSTEIN, A. (1988), "The Structure of Nash Equilibria in Repeated Games with Finite Automata", *Econometrica*, 56, 1259–1282.
 ANDERSON, S., DE PALMA, A. and THISSE, J. (1992), *Discrete Choice Theory of Product Differentiation* (Cambridge: MIT Press).
 CARLSSON, H. and VAN DAMME, E. (1993), "Global Games and Equilibrium Selection", *Econometrica*, 61, 989–1018.
 CHEN, H. C., FRIEDMAN, J. and THISSE, J. (1997), "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach", *Games and Economic Behavior*, 18, 32–54.
 EDGELL, S. E., and GEISLER, W. S. (1980), "A Set-Theoretic Random Utility Model of Choice Behavior", *Journal of Mathematical Psychology*, 21, 265–278.
 FRIEDMAN, J. (1971), "A Non-Cooperative Equilibrium for Supergames", *Review of Economic Studies*, 38, 1–12.
 FUDENBERG, D. and MASKIN, E. (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information", *Econometrica*, 54, 533–556.

- JEHIEL, P. (1995). "Limited Horizon Forecast in Repeated Alternate Games", *Journal of Economic Theory*, 67, 497–519.
- JEHIEL, P. (1998a). "Learning to Play Limited Forecast Equilibria", *Games and Economic Behavior*, 22, 274–298.
- JEHIEL, P. (1998b). "Limited Foresight May Force Cooperation" (CERAS Discussion paper 98-05).
- KALAI, E. and STANFORD, W. (1988). "Finite Rationality and Interpersonal Complexity in Repeated Games", *Econometrica*, 61, 29–56.
- KANDORI, M., MAILATH, G. and ROB, R. (1993). "Learning, Mutation and Long-Run Equilibria in Games", *Econometrica*, 61, 397–410.
- MORRIS, S., ROB, R. and SHIN, H. S. (1995). " p -Dominance and Belief Potential", *Econometrica*, 63, 145–157.
- NEYMAN, A. (1985). "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma", *Economics Letters*, 19, 227–229.
- RUBINSTEIN, A. (1979). "Equilibrium in Supergames with the Overtaking Criterion", *Journal of Economic Theory*, 21, 1–9.
- RUBINSTEIN, A. (1986). "Finite Automata Play the Repeated Prisoner's Dilemma", *Journal of Economic Theory*, 39, 83–96.
- RUBINSTEIN, A. (1998) *Modeling Bounded Rationality* (Cambridge: MIT Press).
- SIMON, H. A. and SCHAEFFER, J. (1992). "The Game of Chess" in Aumann, R. and Hart, S. (eds.), *Handbook of Game Theory with Economic Applications* (Amsterdam: North-Holland), 1–16.
- THURSTONE, L. L. (1927). "A Law of Comparative Judgement", *Psychological Review*, 34, 273–286.