JOURNAL OF ECONOMIC THEORY 67, 497–519 (1995)

Limited Horizon Forecast in Repeated Alternate Games

PHILIPPE JÉHIEL*

C.E.R.A.S.–E.N.P.C., 28 rue des Saints-Pères, 75007 Paris, France

Received December 10, 1992; revised September 20, 1994

In two-player infinite-horizon alternating-move games, a limited forecast (n_1, n_2)-equilibrium is such that (1) player i chooses actions according to his n_i-length forecasts so as to maximize the average payoff over the forthcoming n_i periods, and (2) players' equilibrium forecasts are correct. With finite action spaces, (n_1, n_2)-solutions always exist and are cyclical, and the memory capacity of the players has no influence on the set of solutions. A solution is hyperstable if it is an (n_1, n_2)-solution for all n_1, n_2 sufficiently large. Hyperstable solutions are shown to exist and are characterized for generic repeated alternate-move 2×2 games. *Journal of Economic Literature* Classification Numbers: C72, D81.

I. INTRODUCTION

It is now commonplace to solve extensive form games of complete and perfect information, like chess, by the concept of subgame perfect nash equilibrium. However, as Simon [9] has pointed out, it is clear that no chess player is able to compute the equilibrium path (see also Simon and Schaeffer [10]). Stated differently, chess players are only boundedly rational. There are obviously many ways to think of bounded rationality (see, for example, Rubinstein [7, 8]), but in this context a natural approach is to view chess players as trying to formulate, at each stage where they must move, predictions about the forthcoming n moves, say, as a function of their current move (and the pre-play position of the game). In other words, players have a limited ability to forecast the future, and presumably, the longer the length of foresight, the better the player. However, a major difficulty in the game of chess is to understand what a reasonable value function at the end of the horizon of foresight may be. The treatment of

* The comments of an anonymous associate editor and anonymous referees greatly improved the quality of the paper. I am grateful to Abhijit Banerjee, Olivier Compte, Eddie Dekel, Jim Friedman, Ehud Kalai, Ariel Rubinstein, Sylvain Sorin, Jacques Thisse, Jean Tirole, and Marie-Odile Yanelle for helpful remarks on earlier drafts of this paper. I also thank seminar participants in Bonn, at "Summer in Tel Aviv 1992," at SITE 1993, at ESEM 1993, and at a conference on epistemic logic and the theory of games in Luminy 1994 for their comments.

0022-0531/95 $12.00

that point is extremely complex. We will avoid such a difficulty by considering *repeated alternate-move games* for which we propose a natural criterion to make comparisons based only on limited forecasts.

The objective of this paper is first to define a game-theoretic solution concept that has the feature of limited forecasting and, second, to analyze the properties of such solutions in a simple class of games, the class of repeated alternate-move 2×2 games. The general class of games we consider is as follows. There are two players who move sequentially. At each period t, the current payoff to a player is a function of the period t action of the player who moves currently and the previous period action of the other player.[1] The action spaces of the players are finite and remain the same throughout the play. The horizon is infinite and, for simplicity, we assume that players do not discount the future: streams of payoffs are evaluated according to the overtaking criterion.[2]

Players are assumed to have a *limited ability to forecast* the future. Player i is characterized by the length of his foresight n_i. At period t, player i formulates predictions for the forthcoming n_i moves after his own move. Then he must make his choice of current action on the basis of his limited forecasts only. A natural criterion for player i to compare actions on the basis of his limited forecasts is the average payoff obtained over the length of foresight. This is because (1) player i cannot build his criterion on what will come after n_i periods, since he makes no prediction about (or he has no idea of) it, and (2) given the stationarity of the game, the average payoff over the length of foresight may be perceived as a good approximation of the true objective function.

What is a plausible outcome when players 1 and 2 play such a repeated alternate-move game? This paper proposes as solution concept, the (n_1, n_2)-solution. Two preliminary notions are required: (1) A strategy for player i is *justified* by a sequence of forecasts if the strategy only prescribes actions that maximize the average payoff obtained over the length of foresight (as given by the forecasts), and (2) a sequence of forecasts for player i is *consistent* with a strategy profile if the forecasts coincide with the truncation to the first n_i actions of the respective continuation paths induced by the strategy profile. A (n_1, n_2)-solution is then defined as a strategy profile that can be *justified* by *consistent* sequences of forecasts for

[1] The alternating-move paradigm is more appropriate than the simultaneous-move paradigm to study the effect of limited forecasting because it avoids folk-theorem-like arguments, and allows one to have unambiguously different outcomes with different lengths of foresight.

[2] We could alternatively assume that players do discount the future, but that they are patient.

players 1 and 2. In other words, at an (n_1, n_2)-solution, (1) current actions are chosen so as to maximize the average payoff over the length of foresight, and (2) at any period where player i must move, his forecasts for the forthcoming n_i actions as a function of his current action are *correct*. It should be noted that the predictions for the forthcoming n_i actions made by player i include his *own* actions and that the equilibrium forecasts about all these actions are assumed to be correct whatever his current action and not only for the action on the equilibrium path.

We show that an (n_1, n_2)-solution always exists, and that the period t forecasts associated with (n_1, n_2)-solutions repeat cyclically as t varies. We also observe that players may be better off with shorter lengths of foresight. To be more precise, in the following analysis, we allow the limited forecasts made by the players to depend on history. In addition to his limited ability to forecast the future, player i is assumed to have a bounded recall. The solution concept with bounded recall allows player i's forecasts to depend on the past N_i actions as well. We show that the set of these solutions coincides exactly with the set of (n_1, n_2) solutions. In other words, the equilibrium forecasts associated with such solutions do not depend on history, and the memory capacity of the players has no impact on the set of solutions as long as it is finite.

We next define the concept of *hyperstability*. Hyperstable solutions are strategy profiles that are (n_1, n_2)-solutions for *all* n_1, n_2 sufficiently large. The requirement for hyperstability is, in general, extremely strong. However, if hyperstable solutions happen to exist, they have highly desirable properties since they are robust to a great deal of behavioral changes. A somewhat surprising result is that for generic repeated alternate-move 2×2 games, a hyperstable solution always exists. A full characterization of hyperstable solutions is provided for this case.

The remainder of the paper is organized as follows. In Section 2, we describe the model and define the (n_1, n_2)-solution concept. In Section 3, we define hyperstability and characterize hyperstable solutions of repeated alternate-move 2×2 games. Section 4 provides concluding remarks.

2. The Model

2.1. *Repeated Alternate-Move Games*

We consider two players indexed by $i = 1, 2$. Player i chooses actions a_i from a finite action space A_i. Players act in discrete time, and the horizon is infinite. Periods are indexed by t $(t = 1, 2, 3 \ldots)$. At time t, player i's

TABLE I

UL	UR	DL	DR
$(u_1 = a, u_2 = a')$	$(u_1 = b, u_2 = c')$	$(u_1 = c, u_2 = b')$	$(u_1 = d, u_2 = d')$

single-period payoff is a function of the current actions a_i^t of the two players $i = 1, 2$, but not of time: $u_i = u_i(a_1^t, a_2^t)$. We assume that players do not discount the future. Streams of payoffs are evaluated according to the overtaking criterion.

Players move sequentially and player 1 moves first. At each odd period $t = 2k - 1$ ($t = 1, 3, 5, \ldots$), player 1 chooses an action that remains the same for the two periods t and $t + 1$: $a_1^{2k} = a_1^{2k-1}$ for all k. Similarly, player 2 moves at each even period $t = 2k$ ($t = 2, 4, 6, \ldots$) and $a_2^{2k+1} = a_2^{2k}$.

Games with the above features are referred to as *repeated alternate-move games*. When the cardinality of the action spaces is two, the latter are denoted by $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$. Players' single-period payoffs as a function of pairs of actions are then given as in Table I.

A stream of action profiles $\{q_i^t\}_{t=1}^{\infty} = \{q_1^{2k-1}, q_2^{2k}\}_{k=1}^{\infty}$, where $q_1^{2k-1} \in A_1$ and $q_2^{2k} \in A_2$ is referred to as a path and is denoted by Q. Since players may only change actions every other period, a move at period t affects payoffs both at periods t and $t + 1$. In path Q, each action q_2^{2k} (resp. q_1^{2k+1}) of player 2 (resp. 1) is thus combined both with the previous action, q_1^{2k-1} (resp q_2^{2k}), and the next action, q_1^{2k+1} (resp. q_2^{2k+2}), of player 1 (resp. 2): At periods $2k$ and $2k + 1$, the *current* payoffs to player i induced by path Q are $u_i(q_1^{2k-1}, q_2^{2k})$ and $u_i(q_1^{2k+1}, q_2^{2k})$, respectively.[3] We first introduce some preliminary and standard notation.

Notation. (1) Let R_n denote an arbitrary n-length stream of alternate actions. $v_i(R_n)$ denotes the sum of the per period payoffs to player i induced by R_n, where each action of R_n is combined both with the previous (except for the first one) and the next (except for the last one) action of R_n. For example, in the 2×2 case, the 4-length stream $R_4 = (U, L, D, R)$ induces $v_i(R_4) = v_i(ULDR) = u_i(U, L) + u_i(D, L) + u_i(D, R)$.

(2) $[Q]_n$ denotes the truncation of path, $Q = \{q_i^t\}_{t=1}^{\infty}$, to the *first n* actions: $[Q]_n = \{q_i^t\}_{t=1}^{n}$.

(3) $[q]^N$ denotes the truncation of $q = \{q_i^t\}_{t=v}^{w}$ to the *last N* actions: $[q]^N = \{q_i^t\}_{t=w-N+1}^{w}$.[4]

[3] Single-period payoffs start at period 2.

[4] This expression is valid if $N \leqslant w - v$. When $N > w - v$, $[q]^N$ is identified with q.

(4) (q, q') denotes the concatenation of $q = \{q_i^t\}_{t=r}^{s}$ with $q' = \{q_i^t\}_{t=s+1}^{w}$: $(q, q') = \{q_i^t\}_{t=r}^{w}$.

2.2. *The Solution Concept with Limited Horizon Forecast*

Players are assumed to have a *limited* ability to forecast the *future* and a *bounded recall*.[5] The idea is that players who have some units of brain power at their disposal allocate them partly to the study of the future and partly to the analysis of the past. Player i has a two-dimensional ability that is represented by two integers: n_i, his length of foresight and N_i, his memory capacity. At each period where player i must move, this player makes *limited* predictions for the forthcoming moves: his forecasts are restricted to the forthcoming n_i moves after his own move. Since player i has a bounded recall, his limited forecasts may only depend on the last N_i actions and the current time period. We then assume that player i chooses actions on the basis of his limited forecasts. Restricting attention to pure strategies and predictions, we now introduce some definitions and notation.

Definitions and Notation. Let $\mathscr{H}(N_i)$ denote the set of N_i-length histories of alternate actions, the last action of which is an element of A_j $(j \neq i)$ and let h denote an arbitrary element of $\mathscr{H}(N_i)$.

(1) An n_i-length (pure) prediction for player i is a stream of alternate actions of length n_i starting with an action in A_j $(j \neq i)$. The set of n_i-length predictions is denoted by P_{n_i}: $P_{n_i} = (A_j \times A_i)^{n_i/2}$ if n_i is even and $P_{n_i} = A_j \times (A_i \times A_j)^{(n_i - 1)/2}$ if n_i is odd.

(2) An n_i-length forecast for player i at a period t where this player has to move is denoted by f_i^t. It maps, for every N_i-length history $h \in \mathscr{H}(N_i)$, the set of actions A_i to be currently chosen into the set of predictions P_{n_i}. Formally, $f_i^t = \{f_i^t(\cdot \mid h)\}_h$, where $\forall h \in \mathscr{H}(N_i)$, $f_i^t(\cdot \mid h)$: $A_i \to P_{n_i}$; $f_i^t(a_i \mid h)$ is the prediction about the forthcoming n_i actions made by player i at period t if he currently chooses a_i given the last N_i actions $h \in \mathscr{H}(N_i)$.

(3) $f_i = \{f_i^t\}_t$ denotes an arbitrary sequence of forecasts f_i^t for every period t where player i must move. The set of f_i is denoted $\mathscr{F}_i$. A pair $(f_1, f_2) \in \mathscr{F}_1 \times \mathscr{F}_2$ is denoted by f and the set of f is denoted $\mathscr{F}$.

(4) A pure strategy for player i is denoted by σ_i. It is a sequence of functions σ_i^t, one for each period t where player i must move. The function at period t, σ_i^t, is the behavior strategy of player i at that period. It determines player i's action at period t as a function of the last N_i actions.

[5] Still we assume that players can identify the current time period.

Formally, $\sigma_i^t \colon \mathscr{H}(N_i) \to A_i$.[6] The set of player i's strategies is denoted by Σ_i. A strategy profile (σ_1, σ_2) is denoted by σ, and the set of strategy profiles $\Sigma_1 \times \Sigma_2$ is denoted Σ.

Any strategy profile $\sigma \in \Sigma$ generates a path denoted by $Q(\sigma) = \{q_i^t(\sigma)\}_t$, $i = 1$ (resp. 2) if t is odd (resp. even).[7] Let $\mathscr{H}^t$ denote the set of histories of alternate actions of length t. Let h^* be an arbitrary history of length $t-1$, i.e., $h^* \in \mathscr{H}^{t-1}$. The strategy profile and the path induces by σ on the subgame following h^* are denoted by $\sigma|_{h^*}$ and $Q(\sigma|_{h^*})$, respectively. Given $h^* \in \mathscr{H}^{t-1}$ and the action $a_i \in A_i$ at period t, the continuation path induced by σ after (h^*, a_i) is thus $Q(\sigma|_{h^* a_i})$. In the following, we will consider the set of all such period t continuation paths. This set is referred to as the period t continuation set and is denoted by $Q^t(\sigma)$, where $Q^t(\sigma) = \{(a_i, Q(\sigma|_{h^* a_i}))\}_{h^*, a_i}$.[8] The sequence of continuation sets, $Q^t(\sigma)$, $t = 1, 2, \ldots$, is denoted by $\hat{Q}(\sigma) = \{Q^t(\sigma)\}_t$. Note that given a sequence of continuation sets, we can easily construct the associated strategy profile. Hence, a strategy profile can equivalently be described in terms of σ or $\hat{Q}(\sigma)$.

The basic idea of this paper is that players' strategies (i.e., choices of actions) are to be based on their limited forecasts only. Hence, in order to define the solution concept, we still must (1) specify a criterion based on limited forecasts (the criterion will induce choices of actions given forecasts) and (2) say how equilibrium forecasts are related to equilibrium strategies.

We assume that players use the average per period payoff over the length of foresight as their criterion. Such a criterion is natural, since the environment faced by the players is stationary and, by assumption, player i has no idea of (or does not consider) what the stream of actions after n_i periods will be. In other words, the outcomes over the length of foresight are viewed as a fair sample of all future outcomes, and we assume that the associated average payoff is perceived by player i as a good approximation of the true objective function. Formally:

DEFINITION 1. A strategy $\sigma_i \in \Sigma_i$ is *justified* by a sequence of forecasts $f_i = \{f_i^t\}_t \in \mathscr{F}_i$ if

[6] The limitation to the last N_i periods is due to the bounded recall of player i. At period t, $t \leqslant n_i$, there are less than N_i past actions, and all previous actions must be considered: for such periods t, we identify $\mathscr{H}(N_i)$ with the set of histories of length $t-1$.

[7] This path is defined inductively as follows: $q_1^1(\sigma) = \sigma_1^1$, and for $t = 2, \ldots, q_i^{t+1}(\sigma) = \sigma_i^{t+1}([q_1^1(\sigma), \ldots, q_j^t \sigma)]^{N_i})$.

[8] For notational convenience current actions are included in continuation sets.

(1) $\forall t$, $\forall h \in \mathscr{H}(N_i)$, $\forall a_i \in A_i$, $\sigma_i^t(h) \in \arg_{a_i} \max v_i(ha_i f_i^t(a_i \mid h))$, and[9, 10]

(2) $\forall t, t'$, $\forall h \in \mathscr{H}(N_i)$, $\forall h' \in \mathscr{H}(N_i)$, if $(\forall a_i \in A_i,\ v_i([h]^1 a_i f_i^t(a_i \mid h)) = v_i([h']^1 a_i f_i^{t'}(a_i \mid h')))$, then $\sigma_i^t(h) = \sigma_i^{t'}(h')$.

While condition (1) defines the criterion, condition (2) is a tie-breaking rule, meaning that if all current actions yield the same value of the criterion in two different situations, the same action should be chosen. Observe first that condition (2) is relevant only when several actions yield the same (maximum) average payoff over the length of foresight. (Otherwise, condition (2) is implied by condition (1).) Second, such a tie-breaking rule is natural since it requires the choice of current action to depend exclusively on the values of the criterion (which result from the values of the forecasts), and we have assumed that the choices of current actions should be based on limited forecasts only. Third, if we had introduced discounting, then for arbitrary limited predictions and truncated histories, two different actions would have generically yielded different average payoffs over the length of foresight. Thus, condition (2) would have been generically irrelevant.[11]

We next assume that player i's equilibrium forecasts are related to equilibrium strategies by a consistency relationship, which is defined as follows. Given any history of length $t-1$, h^*, and any current period t action a_i, $Q(\sigma|_{h^* a_i})$ is the continuation path induced by σ after (h^*, a_i). At period t, the N_i last actions are $h = [h^*]^{N_i}$. Consistency requires that for every (h^*, a_i), the prediction $f_i^t(a_i \mid h)$ coincides with the truncation to the first n_i actions of the continuation path induced by σ, $[Q(\sigma|_{h^* a_i})]_{n_i}$. In other words, consistency means that forecasts are correct on and off the equilibrium path. Formally:

Definition 2. $f_i = \{f_i^t\}_t \in \mathscr{F}_i$ is *consistent* with $\sigma \in \Sigma$ if for every period t where player i must move: $\forall a_i \in A_i$, $\forall h^* \in \mathscr{H}^{t-1}$, $f_i^t(a_i \mid h) = [Q(\sigma|_{h^* a_i})]_{n_i}$ with $h = [h^*]^{N_i}$.

[9] The last action of h matters because the game is alternate. Note that we could equivalently consider as criterion $v_i([h]^1 a_i f_i^t(a_i \mid h))$, where only the last action of h, i.e., $[h]^1$, matters.

[10] When $n_i = \infty$, we use the overtaking criterion; i.e., player i strictly prefers path Q to path Q' if $\exists T^*$ s.t. $\forall T > T^*$, $v_i([Q]_T) - v_i([Q']_T) > 0$.

[11] By genericity, we mean that there is no rational combination of the single-period payoffs (different from the null combination) that is equal to zero. With no discounting, indifferences may generically occur if, for example, in the 2×2 case $n_2 = 3$, $f_2^{2k+1}(L \mid h) = DRU$, $f_2^{2k+1}(R \mid h) = DLU$, and $[h]^1 = U$. This is because $v_2(ULDRU) = v_2(URDLU)$. However, with discount factor δ (and discounted average as criterion), we generically have $u_2(U, L) + \delta u_2(D, L) + \delta^2 u_2(D, R) + \delta^3 u_2(U, R) \neq u_2(U, R) + \delta u_2(D, R) + \delta^2 u_2(D, L) + \delta^3 u_2(U, L)$. Hence, there is no indifference.

We can now define the solution concept: An $(n_1, N_1; n_2, N_2)$-solution is a strategy profile that can be *justified* by *consistent* forecasts for players 1 and 2, i.e., a strategy profile that is associated with sequences of forecasts such that (1) players choose their actions in order to maximize the average payoff over the length of foresight, and (2) player i's forecasts for the forthcoming n_i moves after his own move (given the past N_i actions) are correct on and off the equilibrium path. Formally;

DEFINITION 3 (THE SOLUTION CONCEPT). A strategy profile $\sigma = (\sigma_1, \sigma_2) \in \Sigma$ is an $(n_1, N_1; n_2, N_2)$-solution if and only if there exists sequences of forecasts $f = (f_1, f_2) \in \mathcal{F}$ such that for $i = 1, 2$,

(1) σ_i is *justified* by f_i and

(2) f_i is *consistent* with σ.

Note that we have given no justification for why forecasts should be correct in equilibrium. This issue is addressed in Jéhiel [4], where I discuss a learning process based on limited predictions such that players eventually learn to have correct forecasts. Hence, players eventually behave as in an $(n_1, N_1; n_2, N_2)$-solution.

2.3. *Backward Construction*

Consider an arbitrary forecast f_i^T for player i at period T. We show that if f_i^T is associated with an $(n_1, N_1; n_2, N_2)$-solution σ, then the period $T-1$ equilibrium forecast associated with σ can be derived backward on the sole basis of f_i^T.

An example in the 2×2 case follows. Both players have the same length of foresight $n_1 = n_2 = 1$, and the same memory capacity $N_1 = N_2 = 2$. We assume that there exists an $(n_1, N_1; n_2, N_2)$-solution σ such that the associated player 1's forecast at period $2k+1$ is given by $f_1^{2k+1}(U \mid UL) = L$, $f_1^{2k+1}(U \mid UR) = R$, $f_1^{2k+1}(U \mid DL) = R$, $f_1^{2k+1}(U \mid DR) = L$, and $f_1^{2k+1}(D \mid UL) = R$, $f_1^{2k+1}(D \mid UR) = R$, $f_1^{2k+1}(D \mid DL) = L$, $f_1^{2k+1}(D \mid DR) = L$. Note that these forecasts do depend on the last two actions. We now construct player 2's forecast at period $2k$. To fix ideas, we consider the following single-period payoffs:

EXAMPLE 1.

UL	UR	DL	DR
(3, 2)	(0, 3)	(2, 1)	(2, 0)

Consider period $2k$ with an arbitrary $(2k-1)$-period history h^*. The last action of h^* is U. If player 2 chooses to play L at period $2k$, then the history at period $2k+1$ is (h^*, L), and the truncation to be considered at period $2k+1$ is $h = [h^*, L]^{N_2} = (U, L)$. Using condition (1) of Definition 1, we get that player 1 must choose U rather than D at period $2k+1$ because $v_1(hUf_1^{2k+1}(U \mid h)) > v_1(hDf_1^{2k+1}(D \mid h))$ $(3+3+3 > 3+2+2)$. Consequently, since (by consistency) the equilibrium predictions made by player 2 at period $2k$ must be correct, player 2 must forecast at period $2k$ that if he chooses L, player 1 will choose U at period $2k+1$: $\forall a_2 \in A_2$, $f_2^{2k}(L \mid (a_2, U)) = U$. By the same argument, we find that $\forall a_2 \in A_2$, $f_2^{2k}(R \mid (a_2, U)) = D$, $f_2^{2k}(L \mid (a_2, D)) = D$, and $f_2^{2k}(R \mid (a_2, D)) = D$. We have constructed the forecast at period $2k$. Observe that player 2's predictions at period $2k$ only depend on the last action (even though they could *a priori* depend on the last two actions since $N_2 = 2$).

2.4 *First Properties*

The backward construction introduced in Subsection 2.3 is now used for the general analysis of $(n_1, N_1; n_2, N_2)$-solutions. Starting from any possible forecast at period T, f_i^T, we can construct the equilibrium forecast at period $T-1$, f_j^{T-1}, *assuming* that f_i^T is the period T equilibrium forecast. This, in turn, restricts the set of possible equilibrium forecasts at period $T-1$. Starting from f_j^{T-1}, we next construct the equilibrium forecast at period $T-2$. Continuing in this way, we find that prior to period $T - \text{Max}(N_1, N_2)$, equilibrium forecasts necessarily have two properties: (1) they are history independent, i.e. for $t < T - \text{Max}(N_1, N_2)$, $f_i^t(\cdot \mid h)$ is independent of $h \in \mathcal{H}(N_i)$, and (2) sequences $\{f_i^t\}_{t < T - \text{Max}(N_1, N_2)}$ are cyclical, i.e., $\exists s$ s.t. $\forall t < T - \text{Max}(N_1, N_2), f_i^{t-s} = f_i^t$. Since T can be chosen arbitrarily large, these two properties hold for the entire sequences of equilibrium forecasts f_i. Moreover, we can use the cycles of forecasts as derived in the backward constructions to obtain the existence of $(n_1, N_1; n_2, N_2)$-solutions.

PROPOSITION 1. *For arbitrary repeated alternate-move games with finite action spaces A_i,*

(1) *there always exists an $(n_1, N_1; n_2, N_2)$-solution,*

(2) *equilibrium forecasts associated with $(n_1, N_1; n_2, N_2)$-solutions are history independent, and*

(3) *sequences of equilibrium forecasts (f_1, f_2) are necessarily cyclical,* i.e., $\exists s$ s.t. $\forall t, f_i^{t+s} = f_i^t$ *for* $i = 1, 2$.

Proof. Without loss of generality assume that $n_1 \geqslant n_2$.

(i) We first prove that the equilibrium forecasts are history independent. Let $N = \text{Max}(N_1, N_2)$. Since all $(n_1, N_1; n_2, N_2)$-solutions are $(n_1, N; n_2, N)$-solutions, we can restrict our attention to $N_1 = N_2 = N$. Let σ be an $(n_1, N; n_2, N)$-solution, and f denote the associated sequence of forecasts. Let f_1^{2k+1} be the forecast of player 1 at period $2k+1$. As in Example 1, we derive (backward) the equilibrium values of the forecast at period $2k$, i.e., $f_2^{2k}(a_2 \mid h)$ for every $h \in \mathscr{H}(N)$, $a_2 \in A_2$. Because the last $(N-1)$ actions of h, i.e., $[h]^{N-1}$, combined with the current action, a_2, of player 2 at period $2k$ fully determines an N-length history (i.e., $([h]^{N-1}, a_2)$) at period $2k+1$, the action of player 1 at period $2k+1$ can be derived on the sole basis of the last $N-1$ actions of h using player 1's forecast at period $2k+1$: That action must satisfy condition (1) of Definition 1. Since player 1's forecast at period $2k+1$ is correct, player 1's action at period $2k+1$ yields the true stream of actions up to period $2k+n_1+1$. Since player 2's forecast at period $2k$ is correct, and since the past $(N-1)$ actions fully determine the stream of future actions up to period $2k+n_1+1$ as a function of his current action, it follows that player 2's forecast at period $2k$ may only depend on the past $(N-1)$ actions. Continuing the backward construction for earlier periods, we find by induction that, at all periods before $2k-N+1$, forecasts must be history independent (we use condition (2) of Definition 1 for cases of indifferences). Since k can be chosen arbitrarily large, we have proven (2).

(ii) We next prove the existence of $(n_1, N_1; n_2, N_2)$-solutions. For k large enough, consider an arbitrary forecast for player 1 at period $2k+1, f_1^{2k+1}$, and construct the forecasts before $2k+1$ by backward induction as in 2.3. Since there is only a finite number of possible n_1-length forecasts for player 1, there exist $k', k'' \neq k'$ s.t. $f_1^{2k'+1} = f_1^{2k''+1}$. The cyclical sequence of forecasts where a basic cycle is defined by the chain of those forecasts between periods $2k'+1$ and $2k''+1$ defines implicitly an $(n_1, N_1; n_2, N_2)$-solution.

(iii) Finally, we prove the cyclicity of $(n_1, N_1; n_2, N_2)$-solutions. Since there is only a finite number of possible n_1-length forecasts, at least one forecast of player 1 must occur an infinite number of times in any $(n_1, N_1; n_2, N_2)$-solution. Hence, there exist $\{k_m\}_{m=1}^{\infty}$ s.t. $\forall m, m'$, $f_1^{2k_m+1} = f^{2k_{m'}+1}$. By condition (2) of Definition 1 (or (1) if there is no indifference), using the backward construction we obtain that $f_2^{2k_m} = f_2^{2k_{m'}}$. Continuing in this way, we can conclude that the sequence of forecasts is necessarily cyclical, where a basic cycle is given by the chain of forecasts between $2k_m+1$ and $2k_{m'}+1$. Q.E.D

EXAMPLE 1 (CONTINUATION). We illustrate Proposition 1 through Example 1. Given the forecast at period $2k$ constructed above, we proceed to period $2k-1$. Using the backward construction, we find that the forecast of player 1 at period $2k-1$ is history independent. We obtain that $\forall h_1 \in \mathscr{H}(N_1)$, $f_1^{2k-1}(U \mid h_1) = L$ and $f_1^{2k-1}(D \mid h_1) = L$. Continuing in this way, we next determine the forecasts of players 1 and 2 at all periods before $2k-1$. We find that $\forall h_1 \in \mathscr{H}(N_1)$, $\forall h_2 \in \mathscr{H}(N_2)$, $\forall k' \leqslant k$, $f_1^{2k'-1}(U \mid h_1) = f_1^{2k'-1}(D \mid h_1) = L$, and $f_2^{2k'-2}(L \mid h_2) = U$, $f_2^{2k'-2}(R \mid h_2) = D$.

It follows that the forecasts at the periods preceding $2k$ do form a cyclical sequence of history-independent forecasts. We conclude that the original forecast of player 1 at period $2k+1$ (see Subsection 2.3) cannot be part of a $(1, 2; 1, 2)$-solution, since such a forecast is never reached in the above cycle. However, the above cycle allows us to construct a $(1, 2; 1, 2)$-solution. Consider $\sigma \in \Sigma$ and $f \in \mathscr{F}$ defined by $\forall h_2 \in \mathscr{H}(N_2)$, $\forall h_1 \in \mathscr{H}(N_1)$, $\forall k$, $f_1^{2k-1}(U \mid h_1) = f_1^{2k-1}(D \mid h_1) = L$, $f_2^{2k}(L \mid h_2) = U$, $f_2^{2k}(R \mid h_2) = D$, and $\sigma_1^1 = U$, $\sigma_2^{2k}(h_2) = L$, $\sigma_1^{2k+1}(UL) = \sigma_1^{2k+1}(DL) = U$, $\sigma_1^{2k+1}(UR) = \sigma_1^{2k+1}(DR) = D$. It is readily verified that σ is a $(1, 2; 1, 2)$-solution associated with f.

Observe that (1) for any $h^* \in \mathscr{H}^{2k-2}$, the continuation paths induced by σ after (h^*, U) and (h^*, D) are the same and given by $Q(\sigma|_{h^*U}) = Q(\sigma|_{h^*D}) = LU\underline{LU}\ldots$ and (2) for any $h^* \in \mathscr{H}^{2k-1}$, the paths induced by σ after (h^*, L), (h^*, R) are $Q(\sigma|_{h^*L}) = UL\underline{UL}\ldots$, $Q(\sigma|_{h^*R}) = DL\underline{UL}\ldots$, respectively.

Remark 1. Note that the cyclicity of $(n_1, N_1; n_2, N_2)$-solutions refers to cycles in the sequence of forecasts, f_i^t, and not only in the actions actually played in equilibrium. Observe that *all* $(n_1, N_1; n_2, N_2)$-solutions can be constructed by backward induction. For $n_1 \geqslant n_2$, consider successively all possible forecasts f_1^{2k+1} for player 1 and proceed to the backward construction. The associated cycles (implicitly) define all possible $(n_1, N_1; n_2, N_2)$-solutions.

Remark 2. One might infer from Remark 1 that the number of $(n_1, N_1; n_2, N_2)$-solutions is an increasing function of the lengths of foresight since the number of possible forecasts increases as the lengths of foresight get larger. However, this intuition is not correct in general, and one can find parameter values in the 2×2 case for which the number of $(n_1, N_1; n_2, N_2)$-solutions is not a monotonic function of n_i.

We now use Proposition 1 to simplify notation. Proposition 1 shows that the memory capacity has no effect on the set of solutions. Hence, we may drop the parameters N_1 and N_2. An $(n_1, N_1; n_2, N_2)$-solution is called an (n_1, n_2)-solution. Similarly, we drop the parameter h in forecasts: $f_i^t(a_i \mid h)$ is replaced by $f_i^t(a_i)$. The history independence of forecasts together with condition (2) of Definition 1 also implies that the behavior strategy of player i at period t following history $h^* \in \mathscr{H}^{t-1}$ only depends on t and the last action of h^*, i.e., the action chosen by player j at period $t-1$. This, in turn, yields that the continuation path following $(h^*, a_i) \in \mathscr{H}^{t-1} \times A_i$, i.e., $Q(\sigma|_{h^*a_i})$, may only depend on t and player i's action at period t, a_i: This path is more simply denoted $Q^t(a_i)$.[12] The continuation set at period t is then $Q^t = \{(a_i, Q^t(a_i))\}_{a_i \in A_i}$. As an immediate consequence of these observations we get:

COROLLARY 1. *Let σ be an (n_1, n_2)-solution and $Q^t(a_i)$ be the associated continuation path following action a_i at period t. Assume that the θth action of $Q^t(a_i)$ (i.e., the action at period $t+\theta$) is the same whatever $a_i \in A_i$. Then for all $\phi \geq \theta$, the ϕth action of $Q^t(a_i)$ (i.e., the action at period $t+\phi$) is the same whatever $a_i \in A_i$.*

Proof. Since for arbitrary s, the action at period $t+s+1$ only depends on the action chosen at period $t+s$, the result is derived by induction. Q.E.D

Corollary 1 allows us to simplify further the notation of equilibrium continuation sets. Since equilibrium forecasts are consistent, $f_i^t(a_i) = [Q^t(a_i)]_{n_i}$ for all a_i and t (see Definition 2). Hence, equilibrium forecasts can be constructed from equilibrium continuation sets. Consider now an (n_1, n_2)-solution and suppose that, at period t, the equilibrium value of the θth action of $Q^t(a_i)$ is the same for all a_i. Without loss of information with respect to the choice of current action, the continuation set at period t, Q^t, can be *reduced* to the first θ actions. This is because Corollary 1 guarantees that the information contained in the reduced continuation set is sufficient to compare the original n_i-length predictions of player i on the basis of player i's criterion, even if $\theta < n_i$.[13]

Simplified Notation. (1) From now on, an (n_1, n_2)-solution σ will be described in terms of the sequence of continuation sets, $\hat{Q} = \{Q^t\}$, that σ generates.

(2) The original representation of continuation sets is identified with the reduced one.

(3) In the 2×2 case, continuation sets Q^{2k-1} and Q^{2k} are denoted by $\binom{UQ^{2k-1}(U)}{DQ^{2k-1}(D)}$ and $\binom{LQ^{2k}(L)}{RQ^{2k}(R)}$, respectively.

To illustrate the reduction process, consider the (1, 1) (originally (1, 2; 1, 2))-solution constructed in Example 1. The continuation paths at period $2k-1$ satisfy $Q^{2k-1}(U) = Q^{2k-1}(D) = LU\underline{LU}\ldots$. Because the first action of $Q^{2k-1}(U)$ is the same as that of $Q^{2k-1}(D)$, the continuation set at period $2k-1$ can be reduced to the first action: It simplifies into $Q^{2k-1} = \binom{UL}{DL}$. Similarly, because the second action of $Q^{2k}(L) = UL\underline{UL}\ldots$ is the same as that of $Q^{2k}(R) = DL\underline{UL}\ldots$, the continuation set at period $2k$ can be reduced to the first two actions: It simplifies into $Q^{2k} = \binom{LUL}{RDL}$.

[12] The reference to σ is omitted for simplicity.

[13] In such a situation, we have $\mathrm{Arg}_{a_i}\max(ha_i[Q^t(a_i)]_{n_i}) = \mathrm{Arg}_{a_i}\max v_i(ha_i(Q^t[Q^t(a_i)]_\theta)$.

(4) To summarize, the continuation sets associated with the (1, 1)-solution of Example 1 are $Q^{2k-1} = \binom{UL}{DL}$ and $Q^{2k} = \binom{LUL}{RDL}$ for all k. We say that this (1, 1)-solution is defined by the cycle: $Q^t = \binom{UL}{DL} \leftarrow Q^{t-1} = \binom{LUL}{RDL} \leftarrow Q^{t-2} = \binom{UL}{DL}$. (Note that the arrow indicates that we proceed to the next period.)

2.5. *When a Player Is Better Off with a Shorter Length of Foresight*

In an example, we compare the payoffs obtained with different lengths of foresight. The example is chosen so that for every (n_1, n_2), all equilibrium outcomes lead to the same payoffs. At first glance, one might think that for a given length of foresight of player j, the longer the foresight of player i, the better for player i. This intuition is led by the idea that an increase in n_i makes the criterion of player i closer to his true objective function. This loose argument is not correct in general, since an increase in player i's length of foresight may adversely affect the behavior of player j, even if player j keeps the same length of foresight n_j. The following example shows that the payoff to player i may not be an increasing function of player i's length of foresight n_i. More, it shows that player i may be better off not being fully rational.

EXAMPLE 2. Player 2's length of foresight is 1, and the single-period payoffs are:

UL	UR	DL	DR
$(a = 1, a' = 1)$	$(b = 2, c' = 2)$	$(c = 3, b' = 5)$	$(d = 1, d' = 6)$

Assume first that player 1's length of foresight is 0. All (0, 1)-solutions lead to the cycle $\binom{UL}{DL} \leftarrow \binom{LDL}{RUL} \leftarrow \binom{UL}{DL}$. This is because $c > a$, $b > d$, $a' + b' > 2c'$, $2b' > c' + d'$, and $\binom{LDX}{RUY} \leftarrow \binom{UL}{DL}$ whatever X and Y. The sequence of actions actually played at any (0, 1)-solution is $DLDL\ldots$, which leads to player 1's average payoff $c = 3$.

Assume now that player 1 is rational: His length of foresight is $n_1 = \infty$. One can check that all $(\infty, 1)$-solutions lead to the cycle $\binom{UL}{DL} \leftarrow \binom{LD}{RD} \leftarrow \binom{UR}{DR} \leftarrow \binom{LDR}{RUR} \leftarrow \binom{UL}{DL}$. This is because $c > a$, $b > d$, $c + d > a + b$, $a' + b' > 2c'$, $2b' > c' + d'$, $c' + d' > a' + b'$, $d' > b'$, $c' > a'$, which also guarantees that $\binom{LDX}{RUY} \leftarrow \binom{UL}{DL}$, $\binom{LU}{RU} \leftarrow \binom{UR}{DR}$, and $\binom{LUX}{RDY} \leftarrow \binom{UR}{DR}$ whatever X, Y. The sequence of actions actually played at any $(\infty, 1)$-solution is $DLDRDLDR\ldots$, which leads to player 1's average payoff $\frac{1}{2}(c + d) = 2 < 3$. Player 1 gets more by having a myopic behavior $(n_1 = 0)$ rather than by being fully rational $(n_1 = \infty)$.

3. Hyperstability

3.1. *The Concept*

This section proposes a concept of stability, called hyperstability, in connection with the concept of an (n_1, n_2)-solution.

Definition 4. A strategy profile σ is *hyperstable* if there exist n_1^* and n_2^* such that σ is an (n_1, n_2)-solution for *all* (n_1, n_2) with $n_1 > n_1^*$ and $n_2 > n_2^*$.

Hyperstable solutions have highly desirable properties. First, they are robust to increases of the lengths of foresight. Assume in a learning context (see Jéhiel [4]) that player i, once he is able to correctly forecast the forthcoming n_i moves after his own move, decides to forecast the forthcoming $n_i + 1$ moves. Provided the lengths of foresight of players 1 and 2 are sufficiently large, a hyperstable solution is immune against such behavioral changes. Second, hyperstable solutions are also robust to reductions of the lengths of foresight. Assume that at some point of the learning process, player i finds it unnecessary to be very sophisticated. He may then decide to reduce his length of foresight. As long as his length of foresight n_i remains larger than n_i^*, the players may keep the behavior strategies induced by the original hyperstable solution. Third, in an evolutionary context (see Maynard Smith [6]), a hyperstable solution is a solution that is immune against many types of invaders. More precisely, whatever his type (length of foresight) player i will not change his behavior strategy if his opponent (i.e., player j) is replaced by any invader with length of foresight $n_j > n_j^*$. Finally, note that hyperstable solutions are rational solutions since the case of rational players $(n_1 = n_2 = \infty)$ is covered by Definition 4.

As far as the existence of hyperstable solutions is concerned, one might think at first glance that it is not problematic: Proposition 1 ensures that (n_1, n_2)-solutions are defined by cyclical sequences of forecasts. If, for (n_1, n_2) sufficiently large, we can find an (n_1, n_2)-solution σ such that the length of the cycle induced by σ is smaller than n_1 and n_2, then cyclicity ensures that no new information about future behavior is obtained when the lengths of foresight increase. Such a strategy profile might well be hyperstable. However, a uniform upper bound on the length of the cycle induced by (n_1, n_2)-solutions does not necessarily exist. The point is that by increasing $n = \mathrm{Max}(n_1, n_2)$, the set of possible n-length forecasts is also enlarged: When the action spaces have the same cardinality $\# A$, the best upper bound on the length of the cycle induced by an (n_1, n_2)-solution is then $(\# A)^{n+1}$. Consequently, it is not *a priori* possible to ensure (for large n_1 and n_2) that there exists an (n_1, n_2)-solution with a length of cycle smaller than n_1 and n_2.

3.2. *Results in the* 2×2 *Case*

A somewhat surprising result is that for generic repeated alternate-move 2×2 games there exists at least one hyperstable solution (3.2.1). Hyperstable solutions are completely characterized in the 2×2 case (3.2.2), and an example is provided where the set of hyperstable solutions bears no relationship to the set of Markov perfect equilibria (3.2.3).

3.2.1. *Existence*

THEOREM 1. *For generic*[14] *repeated alternate-move* 2×2 *games, there exists at least one hyperstable solution.*

The proof of Theorem 1 heavily relies on the specificity of the 2×2 case. It is relegated to the Appendix. The proof is constructive. Take the continuation set at period k (k large) to be any continuation set reduced to the first action (i.e., the continuation paths from period $k+1$ on are independent of the period k action): $Q^k = (\substack{UL \\ DL}), (\substack{UR \\ DR})$ if k is odd or $(\substack{LU \\ RU}), (\substack{LD \\ RD})$ if k is even. Assuming that the lengths of foresight are infinite ($n_1 = n_2 = \infty$), we use the backward construction as introduced in Subsection 2.3 to derive the sequence of continuation sets $Q^{k'}$ at every period $k', k' < k$. Using genericity arguments (see Lemma 1 in the Appendix), we find that it is always possible to construct $Q^{k'}$ so that as many $Q^{k'}$ as we wish can be reduced to the first action.[15] Since there are only four possible continuation sets reduced to the first action, we conclude that, in this sequence, there are necessarily two identical continuation sets reduced to the first action. Define $\hat{Q} = \{Q^t\}_{t=1}^{\infty}$ to be a cyclical sequence of continuation sets where one cycle is defined by the chain of $Q^{k'}$ between the two identical reduced continuation sets as constructed above. Clearly, $\hat{Q}$ is an (∞, ∞)-solution. Moreover, if T is the length of the cycle induced by $\hat{Q}$, it is readily verified that $\hat{Q}$ is an (n_1, n_2)-solution for all $n_1, n_2 \geqslant T-1$. Hence, $\hat{Q}$ is a hyperstable solution. Note that in the above construction, all continuation sets Q^t can be reduced to a *finite* number of actions, that is, they are such that for some *finite* θ, the associated stream of actions from period $t+\theta$ on is independent of the period t action.

[14] We need that $u_i(U, R) \neq u_i(D, L)$, $u_i(U, R) + u_i(D, L) \neq u_i(D, R) + u_i(U, L)$ for $i = 1, 2$ and all similar inequalities obtained by permuting U and D, and/or L and R. Those are only a *finite* number of inequalities on the components of the single-period payoff matrix which are generically satisfied.

[15] In the 2×2 case, Lemma 1 imposes so much structure that we are able to derive such a property from the backward construction.

3.2.2. *Characterization*

DEFINITION 5. Hyperstable solutions $\hat{Q} = \{Q^t\}_{t=1}^{\infty}$ such that all continuation sets Q^t can be reduced to a *finite* number of actions are called *simple hyperstable solutions.*

The proof of Theorem 1 as sketched above shows that for generic repeated alternate-move 2×2 games, there exists at least one hyperstable solution such that all continuation sets can be reduced to a finite number of actions, i.e., a simple hyperstable solution. We now show how to construct *all* simple hyperstable solutions. Observe first that in the 2×2 case, if at some period t the continuation set Q^t can (minimally) be reduced to a finite number θ of actions, then at period $t+\theta-1$, the continuation set $Q^{t+\theta-1}$ can necessarily be reduced to the first action. Consider next the sequence of predecessors of Q^t obtained in the backward construction with $n_1 = n_2 = \infty$, where Q^t can be any continuation set reduced to the *first* action. Such backward constructions necessarily lead to *cycles* in the sequence of continuation sets (see the proof of Theorem 1). It is readily verified that the set of these cycles define all simple hyperstable solutions.

Remark. It can be shown that the length of the cycle (in the sequence of continuation sets) induced by a simple hyperstable solution can be arbitrarily large. Moreover, for a given per period payoff matrix, simple hyperstable solutions may at most lead to four different cycles that correspond to the four continuation sets reduced to the first action.

We next investigate whether there may be hyperstable solutions that are not simple. To this end, we first note that the sequence of continuation sets, $\hat{Q} = \{Q^t\}_{t=1}^{\infty}$, where $Q^{2k-1} = \left(\begin{smallmatrix} ULDR\underline{ULDR}\ldots \\ DRUL\underline{DRUL}\ldots \end{smallmatrix}\right)$, $Q^{2k} = \left(\begin{smallmatrix} LDRU\underline{LDRU}\ldots \\ RULD\underline{RULD}\ldots \end{smallmatrix}\right)$ for all k, is hyperstable (with $n_1^* = n_2^* = 0$) if $b > a$, $c > d$, $c + d > 2a$, $a + b > 2d$, $a + c > 2d$, $b + d > 2a$, and $a' > c'$, $d' > b'$, $b' + d' > 2c'$, $a' + c' > 2b'$, $a' + b' > 2c'$, $c' + d' > 2b'$. Since such conditions can be met, and since no $Q^t \in \hat{Q}$ can be reduced to a finite number of actions, we conclude that some hyperstable solutions may not be simple.

DEFINITION 6. The sequence of continuation sets $\hat{Q} = \{Q^t\}_{t=1}^{\infty}$ is completely alternate if it is of the form $Q^{2k-1} = \left(\begin{smallmatrix} ULDR\underline{ULDR}\ldots \\ DRUL\underline{DRUL}\ldots \end{smallmatrix}\right)$, $Q^{2k} = \left(\begin{smallmatrix} LDRU\underline{LDRU}\ldots \\ RULD\underline{RULD}\ldots \end{smallmatrix}\right)$ for all k (or the similar sequence obtained by permuting U and D, say).

There are only two completely alternate solutions, and we have shown above the conditions on the per period payoff matrix that guarantee the existence of completely alternate hyperstable solutions (for the other completely alternate solution, use permutations). In the Appendix we prove that:

THEOREM 2. *Hyperstable solutions of generic repeated alternate-move 2×2 games are either simple or they are completely alternate.*

3.2.3. *Comparison with Markov Perfect Equilibria*

Because action spaces are finite, discounted repeated alternate-move games are known to possess at least one stationary equilibrium, i.e., a Markov perfect equilibrium for which the current state is minimally defined to be the action chosen in the previous period (see Fudenberg and Tirole [2]). Several authors justify the Markov perfect equilibrium (MPE) on the grounds that it is *simple*. The concept of hyperstability is clearly another approach to simplicity. For a specific repeated alternate-move 2×2 game, we show that the two notions of simplicity may lead to very different outcomes.

EXAMPLE 3.

UL	UR	DL	DR
(2, 2)	(7, 1)	(1, 1)	(4, 10)

Hyperstable solutions are all simple. They are strict (no indifferences) and all correspond to the cycle $\binom{LU}{RU} \leftarrow \binom{ULU}{DRU} \leftarrow \binom{LD}{RD} \leftarrow \binom{UR}{DR} \leftarrow \binom{LU}{RU}$. The sequence of actions actually played at any hyperstable solution is $RDRU\underline{RDRU}\ldots$ The average payoff to player i is defined by $\frac{1}{2}[u_i(U, R) + u_i(D, R)]$: Each player receives 5.5.

Consider now the discounted game, where player i's criterion is the discounted means over his length of foresight. For discount factors sufficiently close to one, because hyperstable solutions without discounting are strict, they are still hyperstable even with discounting. Clearly, the sequence of actions played at the hyperstable solution cannot correspond to a MPE, since after action R of player 2, player 1 chooses alternatively U and D. Actually, the difference between hyperstable solutions and MPE is even more severe. Consider the limit as the discount factor goes to one of the MPE of the discounted repeated alternate-move game. One can first observe that an MPE in pure strategies fails to exist. (Gurvich [3] also provides a Markov game with finite action spaces and no MPE in pure strategies.) Moreover, the limit as the discount factor goes to one of the unique MPE is such that player 1 chooses D when he sees L; he plays U with probability 4/9 and D with probability 5/9 when he sees R. Player 2 chooses R when he sees D; he plays L with probability 3/4 and R with probability 1/4 when he sees U. Such reaction functions cannot be related to the sequence of actions played at the hyperstable solution. (They lead to average payoffs 5 and 4 to players 1 and 2, respectively.) We conclude that there is no relationship between hyperstable solutions and MPE.[16]

4. Concluding Remarks

Before discussing the solution concept in view of the limited horizon forecast interpretation, observe that at least two alternative interpretations of the (n_1, n_2)-solution concept can be given.[17] First an (n_1, n_2)-solution is *formally* equivalent to the *consistent* outcome of the sequential game with the same structure as described in Subsection 2.1, where players have now a perfect forecast of the entire future but have tastes which vary along the play path (see Strotz [11]): player i's current taste should be then given by the average payoff obtained over the next n_i periods.

Second, the (n_1, n_2)-solution concept can be given an overlapping generations interpretation. Consider two countries $i = 1, 2$. Each country is composed of overlapping generations. The representative (or median voter) of country 1 (resp. 2) chooses an action (or policy) $a_1 \in A_1$ (resp. $a_2 \in A_2$) at every odd (resp. even) period. That action remains the same for the next period (because of political inertia, say). The representatives of countries $i = 1, 2$ change every other period. At every period t, the representative of country i who is in charge of the current policy is assumed to have an exogenously given age. The country i policy maker at period t dies at period $t + n_i + 1$: He is thus only concerned with the forthcoming n_i periods. Assume that individuals do not discount the future as long as they are alive, and that, given the current policies (a_1^t, a_2^t), the period t payoff to an individual in country i (who is alive at period t) is $u_i(a_1^t, a_2^t)$. It is readily verified that an (n_1, n_2)-solution is a subgame perfect Nash equilibrium of that multi-player game. Moreover, hyperstable solutions are solutions which are robust to variations of the lengths of lifetimes in countries 1 and 2.

Coming back to the limited forecast interpretation, we note that the (n_1, n_2)-solution concept can hardly be justified with an eductive approach (see Binmore [1]'s terminology). To see this, let $n_1 = n_2 = n$, for simplicity. If player i were able by introspection to correctly forecast the forthcoming n moves after his own move, he would be able to understand the reaction function of player j for the next period. But, in order to find out his current action at the following period, player j needs forecasts for *his* forthcoming n actions. From the original viewpoint of player i, this represents the

[16] If we restrict ourselves to cases where there exists an MPE in pure strategies, we still find that the relationship between hyperstable solutions and MPE is ambiguous.

[17] I am grateful to the associate editor for this observation.

forthcoming $n+1$ actions. If this argument is pursued, we see that with the eductive approach, the correct forecasting of the next n periods requires a full understanding of all subsequent actions, that is, full rationality. It may then be questionable why players only consider a limited number of (and not all) forthcoming actions.

In Jéhiel [4], I provide a learning justification for the (n_1, n_2)-solution concept that somewhat follows the line of Kalai and Lehrer [5]. Player i initially has a (private) belief over infinite sequences of forecasts f_i (as introduced in Subsection 2.2), chooses current actions so as to maximize the expected average payoff over the forthcoming n_i periods given his current belief, and updates his beliefs according to Bayes' rule. By allowing the possibility of trembles, it is shown that if the initial beliefs of the players assign positive weight to sufficiently many sequences of forecasts, players eventually learn to make correct forecasts and the play of the game is asymptotically (for small trembles) that of an (n_1, n_2)-solution.

Appendix

Proof of Theorem 1. Take player 1's continuation set to be $Q^{-1} = \binom{UL}{DL}$. Assuming that $n_1 = n_2 = \infty$ (see Footnote 10), we next proceed to the backward construction (as introduced in Subsection 2.3) starting from Q^{-1}. We inductively define $Q^{-(k+1)}$ to be the continuation set that precedes Q^{-k} for every k in that construction. This yields an infinite sequence of continuation sets $\{Q^{-k}\}_{k=1}^{\infty}$.[18]

Case 1. There exist an infinite number of k s.t. Q^{-k} can be reduced to the first action. Since there is a finite number (4) of continuation sets that are reduced to the first action, we are sure that there exist k' and k'' such that $Q^{-k'}$ and $Q^{-k''}$ can be reduced to the first action, and $Q^{-k'} = Q^{-k''}$. As explained in the main text, we now define the sequence $\hat{Q} = \{Q^{2k-1}, Q^{2k}\}_{k=1}^{\infty}$ to be cyclical where a basic cycle is implicitly given by the chain of Q^{-k} between $Q^{-k'}$ and $Q^{-k''}$: $\hat{Q}$ is a hyperstable solution.

Case 2. There is a finite number of k. s.t. Q^{-k} can be reduced to the first action. Consider the smallest k such that Q^{-k} can be reduced to the first action. Without loss of generality (w.l.o.g.), we assume that it corresponds to $Q^{-1} = \binom{UL}{DL}$. The predecessor of Q^{-1} is either $Q^{-2} = \binom{LDL}{RUL}$ or $Q^{-2} = \binom{LUL}{RDL}$ (since the two other possible continuation sets can be

[18] Since an infinite sequence is considered, we cannot assume that Q^{-1} occurs at some given finite period. (Otherwise, Q^{-1} would only have a finite number of predecessors.) However, the backward construction can still be used for the derivation of the infinite sequence: It is as if Q^{-1} occurred at infinity.

reduced to the first action). Assume w.l.o.g. that $Q^{-2}=\binom{LDL}{RUL}$. Since Q^{-3} cannot be reduced to the first action, either $Q^{-3}=\binom{URUL}{DLDL}$ or $Q^{-3}=\binom{ULDL}{DRUL}$. We will show that generically only $Q^{-3}=\binom{ULDL}{DRUL}$ is possible. Similarly, generically only $Q^{-4}=\binom{LDRUL}{RULDL}$ and $Q^{-5}=\binom{ULDRUL}{DRULDL}$ will be possible. Given the form of Q^{-5}, we will be able to explicitly construct a hyperstable solution in this case. The fact that generically $Q^{-3}=\binom{ULDL}{DRUL}$ relies on the following lemma which is proven below.

LEMMA 1. *For generic repeated alternate-move* 2×2 *games, let* $\hat{Q}=\{Q^t\}_{t=1}^{\infty}$ *denote the sequence of continuation sets associated with an* (∞,∞)*-solution. Then, there is no period t such that the continuation set* Q^t *is of the form* $\binom{b_i a_i a_j a_i x}{a_i b_i b_j b_i y}$, *where* x, y *are any finite lists of (alternate) actions,* $(a_i, b_i)\in A_i\times A_i$, $(a_j, b_j)\in A_j\times A_j$ $(j\neq i)$, *and* b_i *(resp.* b_j*) stands for the action in* A_i *(resp.* A_j*) other than* a_i *(resp.* a_j*).*

Assume now (by contradiction) that $Q^{-3}=\binom{URUL}{DLDL}$. From Lemma 1 applied to $a_i=U$ $(b_i=D)$, $a_j=R$ $(b_j=L)$, $x=y=L$, we know that $Q^{-4}\neq\binom{LURUL}{RDLDL}$. Since Q^{-4} cannot be reduced to the first action (by assumption in Case 2), we conclude that $Q^{-4}=\binom{LDLDL}{RURUL}$. Again from Lemma 1 applied to $a_i=L$, $a_j=D$, $x=DL$, $y=UL$, we know that $Q^{-5}\neq\binom{ULDLDL}{DRURUL}$. Since Q^{-5} cannot be reduced to the first action, we have $Q^{-5}=\binom{URURUL}{DLDLDL}$. Continuing in this way (using Lemma 1), we find that $Q^{-(2k+1)}=\binom{(UR)_k UL}{(DL)_k DL}$ and $Q^{-2k}=\binom{(LD)_k L}{(RU)_k L}$ for all k where $(X)_k$ designates the repetition of X, k times. However, *generically* the matching (U,R) and (D,L) do not give the same payoff to players 1 and 2, in particular to player 1. Assume, for instance, that $u_1(U,R)>u_1(D,L)$. We are sure then that there exists k^* such that $v_1(L(UR)_{k^*}UL)>v_1(L(DL)_{k^*}DL)$ and $v_1(R(UR)_{k^*}UL)>v_1(R(DL)_{k^*}DL)$. Consider the step $-(2k^*+1)$ for which we have shown that $Q^{-(2k^*+1)}=\binom{(UR)_{k^*}UL}{(DL)_{k^*}DL}$. The above inequalities can be used for the derivation of Q^{-2k^*-2} in the backward construction starting at $Q^{-(2k^*+1)}$. They imply that $Q^{-2k^*-2}(L)=Q^{-2k^*-2}(R)=(UR)_{k^*}UL$. We conclude that the continuation set Q^{-2k^*-2} can be reduced to the first action so that $Q^{-2k^*-2}=\binom{LU}{RU}$, which contradicts the premise that Q^{-2k^*-2} cannot be reduced to the first action. Hence, $Q^{-3}\neq\binom{URUL}{DLDL}$, and $Q^{-3}=\binom{ULDL}{DRUL}$.

Similarly, we obtain that $Q^{-4}=\binom{LDRUL}{RULDL}$ and $Q^{-5}=\binom{ULDRUL}{DRULDL}$. This configuration is possible. However, we now show that, for such a configuration, $Q^{-5}=\binom{UL}{DL}$ is also possible. Because, in the latter situation, $Q^{-1}=Q^{-5}=\binom{UL}{DL}$, we can construct a hyperstable solution as in Case 1. It is implicitly given by the cyclical sequence of continuation sets where a basic cycle is $\binom{UL}{DL}\leftarrow\binom{LDL}{RUL}\leftarrow\binom{ULDL}{DRUL}\leftarrow\binom{LDRUL}{RULDL}\leftarrow\binom{UL}{DL}$.

To see that $Q^{-5}=\binom{UL}{DL}$ is possible whenever $Q^{-5}=\binom{ULDRUL}{DRULDL}$ is possible, let us have a closer look at the derivation of $Q^{-5}(D)=RULDL$. We must compare

$$v_2(DLDRUL)=u_2(D,L)+u_2(D,L)+u_2(D,R)+u_2(U,R)+u_2(U,L)$$

and

$$v_2(DRULDL)=u_2(D,R)+u_2(U,R)+u_2(U,L)+u_2(D,L)+u_2(D,L).$$

These two expressions are the same. Hence, for such a configuration, both $Q^{-5}(D)=RULDL$ and $Q^{-5}(D)=LDRUL$ are possible. In the former case, we find that $Q^{-5}=\binom{ULDRUL}{DRULDL}$ while in the latter case, we find that $Q^{-5}=\binom{UL}{DL}$. (Note that, for such a configuration, the hyperstable solution that we have constructed above is not strict; that is, at some nodes of the game tree, there are some indifferences and these are generic.)

Since we have exhausted all possible cases and shown that in each case there exists a hyperstable solution, we have shown Theorem 1. (Observe that a hyperstable solution as in Case 1 always exists.) Q.E.D

Before we prove Lemma 1, we prove:

LEMMA 2. *For generic repeated altenate-move 2×2 games, let $\hat{Q}$ be a hyperstable solution. Then, there is no period t such that the continuation set Q^t is of the form $\binom{b_j a_i a_j a_i x}{a_i b_i b_j b_i y}$, where x, y are any lists of (alternate) actions, $(a_i, b_i)\in A_i\times A_i$, $(a_j, b_j)\in A_j\times A_j$ $(j\neq i)$, and b_i (resp. b_j) stands for the action in A_i (resp. A_j) other than a_i (resp. a_j).*

Proof. Consider a hyperstable solution $\hat{Q}=\{Q^t\}_{t=1}^{\infty}$ and suppose by contradiction that there is k s.t. $Q^{2k}=\binom{LURUX}{RDLDY}$. This also implies that $Q^{2k+2}=\binom{LDY}{RUX}$. Since $\hat{Q}$ is hyperstable, it is also an (n_1+1, n_2)-solution and an (n_1+3, n_2)-solution for n_1 sufficiently large. Consider player 1's length of foresight to be n_1+3 at period $2k+1$ and n_1+1 at period $2k+3$, where n_1 is sufficiently large. X and Y are then relevant only up to their n_1 first actions. We consider the truncations to the n_1 first actions of X and Y. For notational convenience, we still denote these truncations by X and Y, respectively. Given that $Q^{2k+2}=\binom{LDY}{RUX}$ is part of a (n_1+1, n_2)-solution, we have

$$v_1(LDY)\geqslant v_1(LUX) \tag{1}$$

and

$$v_1(RUX)\geqslant v_1(RDY). \tag{2}$$

Similarly, given that $Q^{2k}=\binom{LURUX}{RDLDY}$ is part of an (n_1+3, n_3)-solution, we have

$$v_1(LURUX)\geqslant v_1(LDLDY) \tag{3}$$

$$v_1(RDLDY)\geqslant v_1(RURUX). \tag{4}$$

By adding up (1), (2), (3), and (4), we obtain that $0\geqslant 0$.

This implies that there are only equalities in (1), (2), (3), and (4). By adding up (1) and (2) now, we obtain that $u_1(U, R) + u_1(D, L) = u_1(D, R) + u_1(U, L)$. Since such a requirement is generically false, we have shown Lemma 2. Q.E.D

Proof of Lemma 1. Consider, by contradiction, an (∞, ∞)-solution, $\hat{Q} = \{Q^t\}_{t=1}^{\infty}$, such that for some k, $Q^{2k} = \binom{LURUX}{RDLDY}$ where X, Y are *finite* lists of actions. (This implies that $Q^{2k+2} = \binom{LDY}{RUX}$ and both Q^{2k} and Q^{2k+2} can be reduced to a finite number of actions.) At period $2k$ (resp. $2k+2$), the criterion of player 1 that is defined by $n_1 = \infty$ is equivalent to the criterion defined by any $n_1 + 3$ (resp. $n_1 + 1$) greater than the number of actions to which Q^{2k} (resp. Q^{2k+2}) is reduced. By considering such an n_1, we can replicate the proof of Lemma 2 to show Lemma 1. Q.E.D

Proof of Theorem 2. Assume the sequence $\hat{Q} = \{Q^{2k-1}, Q^{2k}\}_{k=1}^{\infty}$ is hyperstable.

Case 1. An infinite number of Q^k can be reduced to the first action. Then, there exist two stages k' and k'' such that $Q^{k'} = Q^{k''} = Q$, where Q can be reduced to the first action. We now inductively use the fact that if limited forecasts are the same at periods t' and t'', the backward construction (for $n_i > n_i^*$) should yield the same limited forecasts at periods $t' - 1$ and $t'' - 1$. This implies that the sequence $\hat{Q} = \{Q^{2k-1}, Q^{2k}\}_{k=1}^{\infty}$ is cyclical where a basic cycle is defined by the chain of Q^t between $Q^{k'}$ and $Q^{k''}$. Hence, it is a simple hyperstable solution.

Case 2. A finite number of Q^k can be reduced to the first action. Then, there is a stage s such that from this stage on, continuation sets cannot be reduced to a finite number of actions. Suppose that the sequence of continuation sets from stage s on is not completely alternate. From Lemma 2, we can deduce that there is a stage k^* where the continuation set is of the form $\binom{(UL)_k\underline{DRUL}\ldots}{(DR)_k\underline{ULDR}\ldots}$. (If this were not the case, we would have, say, $\binom{(UL)_k DRDLW}{(DR)_k ULURZ}$, which implies that there is a period for which the continuation set is $\binom{LDRDLW}{RULURZ}$: This contradicts Lemma 2 applied to $a_i = D$, $a_j = R$, $x = LW$, and $y = RZ$). This implies that the continuation set at period $k^* + 2k$ is $\binom{ULDR\underline{ULDR}\ldots}{DRUL\underline{DRUL}\ldots}$. Given the cyclicity of the continuation sets after stage $k^* + 2k$ (and condition (2′) of Definition 2), we can conclude that the sequence has to be completely alternate. Hence, when a hyperstable solution is not simple, it is necessarily completely alternate. Q.E.D

REFERENCES

1. K. BINMORE, Modeling rational players: Part I, *Econ. Phil.* **3** (1987), 179–214.
2. D. FUDENBERG AND J. TIROLE, "Game Theory," MIT Press, Cambridge, MT, 1991.

3. V. GURVICH, A stochastic game with complete information and without equilibrium situations in pure stationary strategies, *Comm. Moscow Math. Soc.* (1986), 171–172.
4. P. JÉHIEL, Learning to play limited forecast equilibria, mimeo C.E.R.A.S. (1995).
5. E. KALAI AND E. LEHRER, Rational learning leads To Nash equilibrium," *Econometrica* **61** (1993), 1019–1045.
6. J. MAYNARD SMITH, Evolution and the theory of games, Cambridge Univ. Press, Cambridge, UK, 1982.
7. A. RUBINSTEIN, Finite automata play the repeated prisoner's dilemma, *J. Econ. Theory* **39** (1986), 83–96.
8. A. RUBINSTEIN, 7 lectures on bounded rationality, mimeo, University of Pennsylvania, 1989.
9. H. A. SIMON, On how to decide what to do, *Bell J. Econ.* **9** (1978), 494–507.
10. H.A. SIMON AND J. SCHAEFFER, The game of chess, *in* "Handbook of Game Theory with Economic Applications" (R. Aumann and S. Hart, Eds.), pp. 1–16, North Holland, Amsterdam, 1992.
11. R. STROTZ, Myopia and dynamic inconsistency in dynamic utility maximization, *Rev. Econ. Stud.* **23** (1955), 165–180.