EUROPEAN UNIVERSITY INSTITUTE, FLORENCE

DEPARTMENT OF ECONOMICS

EUI WORKING PAPER No. 88/331

CONVERGENCE TO EQUILIBRIUM IN A SEQUENCE OF GAMES WITH LEARNING.

by

David CANNING*

* Pembroke College, Cambridge.

This paper was written while the author was a Jean Monnet Fellow at the European University Institute.

BADIA FIESOLANA, SAN DOMENICO (FI)

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

All rights reserved.
No part of this paper may be
reproduced in any form without
permission of the author.

(C) David Canning
Printed in Italy in February 1988
European University Institute
Badia Fiesolana
- 50016 San Domenico (Fi) -
Italy

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

CONVERGENCE TO EQUILIBRIUM IN A SEQUENCE OF GAMES WITH LEARNING

David Canning
European University Institute and
Pembroke College, Cambridge

December 1987

Abstract.

A group of agents repeatedly play a two player game against each other, the pairing of agents in each period being random. Agents assume that they face a stationary probability distribution over which strategy their opponent will choose and estimate this distribution by Bayesian updating on an arbitrary prior. A perfect belief equilibrium is defined for normal form games as an equilibrium which is stable against small mistakes in players' beliefs; it is shown to be a strict refinement of perfect equilibrium for two player games. Conditions are found under which the beliefs, and mixed strategy choices, of the agents converge. Under these conditions the limit beliefs (and mixed strategy choices) are shown to form a Perfect Belief equilibrium of the underlying one-shot game with complete information and common knowledge of the players' payoff functions.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

1. Introduction.

People's actions depend on their beliefs, and their beliefs change over time depending on their observations. In many situations of economic interest their observations will be the actions of others, depending, in turn, on their beliefs. The beliefs of any one agent will be dominated by chance; the original prior belief may be fairly arbitrary and their observations, from which they learn, may represent a biased sample of the actions being taken by others. While it is difficult to say anything about the individual in such circumstances we may be able to say something about the long run average behaviour of such a system, if the number of agents is large, and the agents' learning is based on observations which are not too biased. The aim of this paper is to use such an approach to define an equilibrium concept for games in normal form. The equilibria will be defined as limits to such a learning process.

Most equilibrium concepts in game theory are based on the terrifyingly strict assumption of common knowledge. All agents playing the game share a pool of basic information about the strategic features of the game. Brandenburger and Dekel (1986) demonstrate the relationship between the different solution concepts which have been proposed and the degree of common knowledge which is assumed in the game.

Harsanyi (1967-68, 1980) examines games where the agents are not completely informed and shows how the game can be changed from one of incomplete information to one of complete but imperfect information. Agents subjective probability distributions can be modeled as if they where the result of a known objective randomization. However, this makes the solution of the game a depend entirely on agents beliefs and the

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

theory has no predictive power; almost any outcome can be sustained as an equilibrium providing agents begin with the appropriate beliefs (see Ledyard (1986)).

While the solution concepts used in game theory are based on a very strict assumption of common knowledge there is support for them from a different source. Maynard-Smith (1982) considers a population of creatures who play a particular strategy in a game by instinct. The number of descendants of each creature is assumed to be the creature's payoff in the game; creatures with successful strategies (given the strategies being chosen by others) reproduce more prodigiously. If the young inherit the characteristics of their parent (its strategy choice) we have a dynamic system and can find the equilibrium points (equilibrium proportions of creatures with different strategies) and investigate the stability of these equilibria; which equilibria are stable against an invasion by a "mutant" group with non-equilibrium strategy instincts.

The surprising result is that the equilibria and stable equilibria of such a process correspond closely with the game theoretic equilibrium concepts put forward for rational players maximizing their expected utility with complete information and common knowledge. Evolution provides for the survival of the fittest and the fittest are maximizers.

While these results add to the case for the existing solution concepts the applicability of the inheritance of characteristics model to games played by people in an economic context seems limited. The foundation of economic theory is that actions are a means to an end and not an end in themselves, so while agents' desires remain unchanged

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

their actions adapt with changing circumstance; they have an element of choice.

The following model is proposed for economic situations. Each agent is informed that he will play a sequence of games with randomly chosen anonymous opponents and knows his payoff function over strategy choices. Agents attempt to maximize their expected payoff in each period given their probability distribution over which strategy their opponent will play. The probability distribution over which strategy their opponent will play begins arbitrarily but it is assumed that agents believe they face a stationary distribution and that they update their beliefs on what this stationary distribution is by Bayes rule.

In this environment it seems natural to think of a Nash equilibrium as a rational expectations solution of the system. The existence of such rational expectations equilibria in games where agents have only one period memory is studied by Rosenthal (1979). The obvious question to ask is does the system converge to a rational expectations solution from any set of initial beliefs? The problem is similar to that studied by Bray and Savin (1986) in the context of rational expectations equilibrium in a simple market. Agents believe they face a stationary probability distribution and attempt to estimate its parameters by Bayesian updating. Here agents follow a similar course of action, assuming that the proportions in which the oppositions' strategies are played are constant. The difference in this paper is that there is no market wide price signal which is commonly observed; each agent sees only the strategy chosen by his or her opponent, and so beliefs may depend on the particular outcome of the matching process.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

The main result is that under the simple learning rule the subjective probability distributions of the agents over the strategies they face converge, and these limit distributions are the same for each agent.

Not only do agents' beliefs converge but the objective probability distribution faced by agents in each period converges to a stationary distribution and the strategies observed by agents are, in the limit, stationary and statistically independent. This provides some justification for the simple learning rule we have postulated.

To prove the result three restrictive assumptions are required. Firstly, agents' initial beliefs are restricted to a finite dimensional space which is invariant under Bayesian updating. This set, while quite large, rules out some plausible priors. Secondly, we assume that when agents calculate their probability assignments over what strategy they will face in a particular period they make a small random error; this implies that strategy choices are continuous in beliefs. Finally, it is assumed that the number of agents is large (in fact countably infinite).

Aumann (1986) has argued that equilibria in games should be regarded as equilibria in the players' beliefs rather than in strategy choices. The concept of a Perfect Belief Equilibrium is defined to capture the idea of an equilibrium point of a normal form game which is stable against small mistakes in beliefs. The existence of such equilibria is demonstrated for N player games in normal form and it is shown to be a strict refinement of Nash equilibrium (and a refinement of perfect equilibrium in two player games). Further, we can ensure that for our sequence of games, in which agents learn by Bayesian updating, beliefs converge to such a Perfect Belief equilibrium.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

2. The Model.

Agents repeatedly play a 2 player, finite strategy, game. At each time t each of the players of type one is matched randomly, with equal probabilities, with one of the players of type two in the opposing pool. Each player i of type one chooses a strategy s_1 from the set S_1, similarly players of type two choose from the set S_2. The payoffs to player i from playing against agent j (selected from the opposing pool) who chooses the strategy s_1, for one play of the game, are given by $U_1(s_1, s_2)$. Player j gets the payoff $U_2(s_1, s_2)$. Each player is assumed to know his own payoff if a particular strategy pair is played, he need not know the payoff of his opponent.

Let M_k denote the set of mixed strategies which can be generated from the finite strategy set S_k. Clearly, M_k is a compact subset of a finite dimensional Euclidean space. Usually in game theory we think of agents choosing mixed strategies. Here agents may or may not resort to mixed strategies, the important point is that they believe they face a stationary distribution of strategy choices by their opponents. They may either believe that the agents in the opposing pool are randomizing, or that they individually play pure strategies, but the random matching generates a mixed strategy for the pool. In this context it is natural to think of M_k as the space of possible probability measures agents of type $k' \neq k$ can have over which strategy players of type k will choose. It is useful to make M_k into a normed space by setting

$$|m_k| = \max\{m_k(s_k): s_k \in S_k\}$$

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

where $m_k(s_k)$ is the probability weight given to strategy s_k in the mixed strategy m_k. Hereafter we will assume that agent i comes from pool 1, the appropriate definitions for agents from pool 2 follow by interchanging the symbols 1 and 2 where they appear.

(i) Beliefs.

The problem for each agent in our framework is that they are unsure as to what probability assignment to use over the strategies to be played by the opponent. We suppose that the agents believe they face a stationary distribution but are is unsure of its parameters. We shall assume they act in a Bayesian manner, starting with a prior (which incorporates all relevant initial information) over the set of possible mixed strategies they may face, and updating this prior by Bayes rule in the light of their observations.

In principle the set of possible beliefs of the agent is very large; any probability measure on M_2 is a possible belief for agents of type one. This set is infinite dimensional which raises technical problems in the analysis.

The following simplifying assumption is made. The initial prior belief of each agent is a Dirichlet function; that is, the initial beliefs of each agent over the probability assignments $m_2 = (p_1, \quad ,p_n)$ taking $p_j = m_2(s_j)$ for the n pure strategies of the opposition, can be written in the form

$$b_{i0}(p_1, p_2, \ldots ,p_n) = \frac{\Gamma(\alpha_1+\alpha_2+ \ldots \quad +\alpha_n)}{\Gamma(\alpha_1)\Gamma(\alpha_2)\ldots\Gamma(\alpha_n)} p_1^{\alpha_1} p_2^{\alpha_2} \ldots p_n^{\alpha_n}$$

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

with parameters $\alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_n > 0$. The agent's initial beliefs can be parameterized by n numbers (different of course for each agent).

The advantage of restricting initial beliefs to this family is that it forms a set of conjugate functions for the multinomial distribution which the agents' are trying to estimate; updating a Dirichlet function by Bayes rule produces a new Dirichlet function (see De Groot (1970), Winkler (1972) discusses the restrictiveness of this assumption for the case n = 2 in detail).

It is easy to construct a one to one invertable mapping from the positive orthant of R^n into the interior of a n dimensional unit cube. Therefore each agents belief can be represented by a point in the n dimensional unit cube. The natural mapping to use is

$$q_j = \alpha_j/H \qquad j=1, 2, \ldots, n-1.$$

and $$q_n = \alpha_1(H - \alpha_1)/[H^2(H+1)] \quad \text{where } H = \sum_{j=1}^{n} \alpha_j.$$

It is easy to check that the range of the mapping is the unit cube and that the map is invertable. This mapping is natural because q_1 to q_{n-1} are the probability assignments associated with the belief (e.g. $q_j = \alpha_j/H = \int_{M_2} b(m_2)\, m_2(s_j)\, dm_2$ if the belief b is a Dirichlet function with parameters $(\alpha_1, \ldots, \alpha_n)$) and q_n is the variance of the agents belief over the lst assignment.

Let B_1 be the set of beliefs for agents in pool one (a unit cube of dimensional equal to the number of strategies available to agents in

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

pool two) while B_2 is the set of possible beliefs for agents in pool two. We put a norm on the beliefs spaces by taking

$$|b| = \max \{q_j,\ j=1, \ldots, n\}$$

We define Bayesian updating of a belief b_{it} to a belief b_{it+1} given the observation $s_2 \in S_2$ at time t by

$$b_{it+1}(m_2) = \frac{b_{it}(m_2)\ m_2(s_2)}{\int_{M_2} b_{it}(m_2)\ m_2(s_2)\ dm_2}$$

It is easy to show that given our assumption that if b_{it} is a Dirichlet function with coefficients $(\alpha_1, \ldots, \alpha_n)$ and the agent observes the strategy $s_j \in S_2$ then b_{it+1} is a Dirichlet function with the same parameters as b_{it} except for α_{jt+1} which becomes $\alpha_{jt} + 1$.

Given our assumption that the initial parameters in each agent's prior are positive the bottom line in the Bayesian updating equation is positive and updating is well defined. For any point in the interior of our belief spaces, the unit cubes, B_1 and B_2 we can construct a Dirichlet function, update this function, and return to the space. However, the the interior of the space B_1 is not closed. It follows that the agents can construct a sequence of beliefs by Bayesian updating, each a Dirichlet function, which tend to a limit Belief which is not a Dirichlet function. We complete the Belief space by adding the definitions that;

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

for the set $\{b_{it} \in B_1 : q_n > 0, q_j = 0\}$ Bayesian updating is defined by for the observation s_j by taking $\alpha_j = 1$ and leaving the other parameters unchanged in the Dirichlet function corresponding to b_{it}, for the set $\{b_{it} \in B_1 : q_n = 0\}$ $b_{it+1} = b_{it}$ for any observation. The first case represents the edge of the unit cube when the agent assigns probability zero to an event. Such points correspond to Dirichlet functions with parameter $\alpha_j = 0$, and are never reached by our agent, but we include an updating rule for the agent at such a point for completeness.

The more interesting case is $q_n = 0$; this corresponds to the agent having an infinite number of observations and being sure what mixed strategy the opposing pool are playing; it is a probability measure with an atom of weight one at this point and zero elsewhere. The central question is whether or not the agents beliefs converge to such a point. We can define a norm on the set B_1 by taking $|b| = \max \{q_1, \ldots, q_n\}$.

Lemma 1. The mapping $\beta : B_1 \to B_1$, given by the completed updating rule, is continuous for a fixed observation $s_j \in S_2$.

proof. For $q_n > 0$ the continuity of β is straightforward, the mapping from B_1 to the Dirichlet functions is continuous, the updating mechanism is continuous in the Dirichlet functions, and the return mapping to B_1 is continuous. For $b \in B_1$ with $\alpha_n = 0$ the updating rule gives $\beta(b) = b$. Now consider b' with nth argument q_n'. It is easy to show that

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

$|b' - \beta(b')|$ can be made as small as we please by making q_n' close to zero. Hence $|\beta(b) - \beta(b')| \leq |\beta(b) - b'| + |b' - \beta(b')|$

$$= |b - b'| + |b' - \beta(b')|$$

can be made arbitrarily small by making $|b - b'|$ small enough and β is also continuous at b with $q_n = 0$.

Q.E.D.

(ii) Strategy Choices.

Suppose that player i believes he faces a mixed strategy $m_i \in M_2$. That is, m_i is a probability measure over S_2. We shall assume that he plays a strategy $s_i \in S_1$, so as to maximize

$$\sum_{S_2} U_1(s_i, s_2)\, m_i(s_2)$$

Players are assumed to act so as to maximize their expected payoff in the current period. We can think of players who completely discount the future or of players who believe that future payoffs are independent of their current actions.

For the set of probability distributions where the choice is not unique we assume that the agent chooses between these best replies in a manner that is arbitrary but stationary over time. That is, for any given probability assignment $m_i \in M_2$ over opposing strategies the agent has a unique action, either a pure strategy choice or a well defined mixed strategy. We further assume that this function from probability assignments to a probability for a particular strategy being played is a behavioural strategy in the sense of Aumann (1964), and so is measurable with respect to generalized Lebesgue measure. This is quite a weak

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

assumption. Let ψ denote this best reply mapping for agents of type one;

$$\psi : M_2 \rightarrow M_1.$$

Given their beliefs at each time agents will form a probability assignment as to what strategy will be played by their opponent. We shall assume that while agents update correctly they make a small error in calculating their probability assignments from their beliefs.

Let the beliefs of agent i at time t be denoted by b_{it}. Probability assignments for agent i at time t over the strategies s_2 the opposing pool can play are given by

$$m_{it}(s_2) = P(s_2 \mid b_{it}) = \int_{M_2} b_{it}(m_2)\, m_2(s_2)\, dm_2$$

We denote this mapping by φ: $B_1 \rightarrow M_2$. That is, for each belief the mapping generates a probability assignment over the opponent's strategies in this period. Beliefs are continuous functions and the probability weights on a particular strategy s_2 are certainly continuous over M_2 the function is integrable. It would be quite natural to combine this probability assignment m_{it} with the strategy choice mechanism ψ to give strategy choices as a function of beliefs.

However, we assume that the agent makes a small mistake in the calculation of the probability assignment. That is we have a mapping

$$\mathcal{M} : M_2 \rightarrow C(M_2)$$

so that, $\mathcal{M}(m_{it}) = g_{it} \in C(M_2)$, the set of continuous functions over M_2, where we assume

(a) g_{it} can be regarded as a probability density function on M_2.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

(b) $g(m_2) > 0 \iff |m_2 - m_{it}| < \varepsilon$ for some $\varepsilon > 0$.

(c) η is a continuous function on M_2.

and we define a norm on $C(M_2)$ by $|g| = \sup \{g(m_2),\ m_2 \in M_2\}$.

We can think of a small (continuously distributed) tremble, of size bounded above by ε, when agents are calculating their probability assignments. Given the probability assignment which results from this error process the agent chooses a best reply strategy. We can construct a mapping from the space of beliefs to the space of mixed strategy choices of the agent given by

$$\tau : B_1 \to M_1$$

where $\tau(b_{it}) = \int_{M_2} \eta\varphi(b_{it})\psi(m_2)\,\lambda\,dm_2$.

Since η generates a continuous function and ψ is measurable with respect to the generalized Lebesgue measure λ the integral is well defined.

Lemma 2. τ is continuous over B_1.

proof. φ is a bounded linear operator and hence is continuous. η is continuous by assumption so $\eta\varphi$ is continuous. The mapping from $C(M_2)$ to M_1 given by $\int_{M_2} g\ \psi(m_2)\lambda\,dm_2$ for $g \in C(M_2)$ is a bounded linear operator and so is continuous. τ is the composition of this continuous mapping with the continuous function $\eta\varphi$ and so is continuous.

Q.E.D.

The advantage of assuming agents make mistakes in calculating their probability assignments is that it makes their strategy choices continuous in beliefs. Without such random mistakes there is a sharp

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

dividing line between beliefs for which each strategy is a best reply. Allowing mistakes means that we can regard the agents as choosing a mixed strategy as a function of their beliefs, the mistakes in probability assignments being the mechanism which does the randomizing.

Note that the functions ψ and η are assumed to be the same for each agent i. In principle we could make the mistakes of each agent have a different distribution and allow different agents to choose differently at points where they are indifferent between a set of strategies. Doing this does not change our results; however, the assumption that all agents in the same pool are are identical, except for their beliefs, greatly simplifies the analysis. For almost all games the set of points at which agents are indifferent between strategies is of measure zero; it follows that for a large class of games γ is defined independently of the choice of ψ.

The structure of the model is set out in figure one. We can regard the agents beliefs as the state space of the game. Given their beliefs each agent calculates a probability assignment for the strategy they face. Given this probability assignment (including an element of error) they choose a best reply. They are matched with an agent from the other pool and learn what strategy they actually face and their payoff. They then update their beliefs in the face of this new information.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

3. Convergence to Equilibrium.

The problem we now address is the convergence of the system set out in section 2 for arbitrary initial beliefs of the agents (restricted to Dirichlet functions). Let

$$X = B_1 \times B_2$$

represent the state space of the system made of of the beliefs of two arbitrary agents, one from each pool. We suppose these two agents meet at time t and investigate what we can say about their beliefs at time t+1 given some probability distribution over their beliefs at time t. It is useful to put a norm on the space X by setting

$$|x| = \max \{|b_1|, |b_2|\}$$

An event at time t assigns an observation of a strategy from the opposing strategy set to each of the two agents. Let E be the set of all possible events. Since each agent can observe only a finite number of possible strategies E is clearly a finite set (it is simply the set of pairs $\{(s_2, s_1)\}$. We take E* to be the set of all subsets of E, which is again a finite set. Since E* is finite we can define a measure on E by the number of elements in any set in E*.

Let $Q: X \times E^* \rightarrow [0, 1]$ be the function induced by the decision rule for the agents given by τ in section two. The intrepretation of Q is that for $x \in X$ and $F \in E^*$ the value Q(x, F) is the probability that if the beliefs of the agents are given by x then an event in the set F occurs. For $e \in E$ the probability Q(x, e) is simple to calculate; each agent has a belief which generates a probability distribution over strategy choices, these are then assigned as observations to the other agent.

Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.
© The Author(s). European University Institute.

Lemma 3. For fixed x, Q is a probability measure on E*.

Proof: For fixed x it is easy to show that Q is a probability measure on E*, since every set is measurable and we have $Q(x, E) = 1$, $Q(x, \emptyset) = 0$, and for $F, F' \in E^*$ such that $F \cap F' = \emptyset$ we have $Q(x, F \cap F') = Q(x, F) + Q(x, F')$.

Q.E.D.

Take X* to be the Borel sets of X, that is X* is the smallest σ-algebra containing the open sets of X; (X, X*) is a measurable space. Consider the mapping h: X x E -> X from agents beliefs, and an event E, to new beliefs for the agents, given by the completed Bayesian updating rule.

Lemma 4. h is continuous in X for fixed $e \in E$.

h is measurable in E for fixed $x \in X$.

Proof. The continuity of h follows immediately from lemma 1. h is measurable in E for fixed x provided that for each $A \subseteq X$, $A \in X^*$, the set $\{e \in E: h(x, e) \in A\}$ is measurable. Since every subset of E is measurable this is automatic.

Q.E.D.

Note that even if we knew the initial beliefs of the agents the evolution of the system would still be stochastic. The strategy each agent observes depends on the outcome of the stochastic matching technology and the chance mistakes of the opponent he is matched with.

Now let $Y \subseteq X \times E$. Define $Y_x = \{e \in E : (x, e) \in Y\}$. Define the function $\pi : X \times X^* \to [0, 1]$ by

$$\pi(x, A) = Q(x, (h^{-1}(A))_x)$$

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

This is the probability that if beliefs start in a state x they end in a state in the set A after agents have observed their opponent's strategy choice. The set $(h^{-1}(A))_x$ is the set of events in E which move the system from x to a state in A; the probability of ending in A is just the probability of one of these events occurring given the system starts in x.

Theorem 5. π is a transition probability on (X, X*)

That is : (i) for fixed x in X, $\pi(x, .)$ is a probability measure on X*.

(ii) for fixed A in X*, $\pi(., A)$ is measurable in X.

proof .

(i) For fixed $x \in X$ the set the every updating takes us to a new point in X, so $(x, \phi) = 0$ and $(x, X) = 1$. If A and B are disjoint elements of X* the events which map into them from x are disjoint, E_A and E_B say, and we have

$$\pi(x, A) + \pi(x, B) = Q(x, E_A) + Q(x, E_B) = Q(x, E_A \cap E_B) = \pi(x, A \cap B)$$

so π is a countable additive set function on X* and so is a probability measure.

(ii) Consider the set $\{x: Q(x, F) \leq p\}$ for some fixed $F \in E^*$ and $p \in [0, 1]$. Consider any sequence (x_n) in this set which converges to a limit x*. By lemma two the mixed strategies played by agents are continuous in their beliefs so if x_n converges the associated mixed strategies, $\gamma(x_n)$, also converge. It follows that the probability of a particular set of events converges since an event is just a

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

permutation of the agents' strategy choices. Since $Q(x_n, F) \leq p$ for each x_n in the sequence we have $Q(x^*, F) \leq P$ and the set $\{x: Q(x, F) \leq p\}$ is closed. Closed sets are Borel sets, hence Q is measurable in X*.

We can write $\pi(x, A) = \sum_E Q(x, e)\, \chi_A(h(x, e))$

where χ_A is the characteristic function of the set A, that is:

$$\chi_A(x) = \begin{cases} 1 \text{ if } x \in A \\ 0 \text{ if } x \notin A \end{cases}$$

If A is a Borel set χ_A is obviously measurable, h is continuous in X (by lemma 4) so $\chi_A h$ is measurable regarded as a function over X for a fixed $e \in E$. . Hence for fixed e we have that $Q(x, e)\chi_A(h(x, e))$ is measurable over X in X* since the product of two measurable functions is measurable. Since π can be written as a sum of a finite number of these measurable functions (one for each $e \in E$) we have that π is a measurable function on X (measurable in X*) for fixed $A \in X^*$.

Q.E.D.

We can define a measure space (X, X^*, λ) by considering λ to be Lebesque measure generalized to the higher dimensional space X. Let $L_1(X, X^*, \lambda)$ be the set of all real valued functions on X which are measurable with respect to X* and whose absolute integrals

$$\int_X |f(x)|\, \lambda\, dx$$

are finite. For $f(x) \geq 0$ and $\int_X f(x)\, \lambda\, dx = 1$ we can regard f(x) as a probability density function, giving the probability that the system is in state x. For each f we can define a new mapping on X by

$$Tf(y) = \int_X f(x)\, \pi(x, y)\, \lambda\, dx \qquad \text{for all } y \in X.$$

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

By theorem 5 $\pi(x, y)$ is measurable over $x \in X$ for fixed y so this integral is well defined. Since π is a transition probability T is a Markov operator. T takes a probability distribution over the state space into another by the event and updating mechanisms implicit in π. To check that $Tf \in L_1$ all that is necessary is to show that T is null preserving, that is that Tf gives zero weight to sets of Lebesgue measure zero. It is easy to show that for each x there are at most a finite set of points (one for each event) which map to x under π. For fixed e it is clear that h^{-1} is a continuous function and so maps sets of measure zero into sets of measure zero. The finite union of sets of measure zero (one for each e) has zero measure so any set of measure zero has zero weight in Tf.

If f_0 is our initial probability distribution over X then after the the random matching and updating we have the distribution Tf_0.

We now assume that the initial beliefs of the agents are given by independent, identically distributed, random variables on B_1 and B_2 with distributions f_0^1 and f_0^2 (each a Lebesgue measurable function) so that our initial distribution over the beliefs of a randomly matched pair is $f_0 = f_0^1 \times f_0^2 \in L_1(X, X^*, \lambda)$. Even if we knew each agents' initial belief for sure the initial random matching generates a probability distribution over the beliefs of the agents in any pairing (note it is the same for any pair).

The state space of our system is in fact X^∞, a belief for each agent of each type. However, let us label the pairs that are matched each

Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.
© The Author(s). European University Institute.

period as 1, 2, 3 ,... , and follow the beliefs of the agents who happen to be matched in pair one. For each initial pairing our distribution of the beliefs of the two agents is f_0 and it follows that after they have observed their opponents strategy we have a distribution Tf_0 over their beliefs. The actual outcome will be different for each pairing but the distribution Tf_0 is the same for each, since for each pairing we have the same initial distribution f_0. In addition, the actual outcome is statistically independent across pairings since the mistakes of each agent are statistically independent.

The two agents in our pairing labeled no. 1 are now replaced by agents drawn randomly from the entire population. The probability distribution on agents beliefs in the replication we are studying is now the average of the probability distributions of the different populations. But these are all the same (our distribution over each pair of agents' beliefs is identical), so after the reallocation of agents the probability distribution over agent types in the case we are studying is still Tf_0.

This distribution is identical for each replication. After matching and updating again the probability distribution over beliefs is given by T^2f_0, again the same for each replication. It follows by the above argument that this distribution is once again unchanged by reassigning agents between replications. We can continue this process indefinitely so that the probability distribution for the replication we are studying is, at time t, given by $f_t = T^tf_0$.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Clearly $L_1(X, X^*, \lambda)$ is a vector space under the usual definitions of function addition and scaler multiplication and we can define a norm on $L_1(X, X^*, \lambda)$ by $|f| = \int_X |f(x)| \lambda dx$. We can then define a norm on operators over $L_1(X, X^*, \lambda)$ in the usual way

$$|T| = \sup \{|Tf| : f \in L_1(X, X^*, \lambda), |f| = 1\}$$

Lemma 6. T is a positive contraction on $L_1(X, X^*, \lambda)$.

proof. We require that if $f(x) \geqslant 0$ for all x then $Tf(x) \geqslant 0$ for all x, and that $|T| \leqslant 1$. Both these follow immediately from the definition of T and the fact that π is a transition probability.

Q.E.D.

In fact it is easy to show that $|T| = 1$. Now let

$$S_N f = \sum_{t=0}^{N-1} T^t f$$

Since T is a positive contraction we can perform the Hopf decomposition of X, constructing two disjoint sets, a conservative set C and a dissipative set D defined by

For all $f \in L_1^+$, $S_\infty f = \infty$ on $C \cap \{x \in X: S_\infty f(x) > 0\}$ and

for all $f \in L_1^+$, $S_\infty f < \infty$ on D.

where L_1^+ is the set of functions in $L_1(X, X^*, \lambda)$ which are non-negative everywhere on X and $S_\infty f = \lim_{N \to \infty} S_N f$. Intuitively what this says is that if we begin with a positive weight on all points in X then, after an arbitrarily long time, we still give positive weight to points in C but give zero weight to points in D.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Lemma 7. The Hopf decomposition of X under the operator T is

$$C = \{x = (b_1, b_2) \in X: q_n = 0 \text{ for both } b_1 \text{ and } b_2\}, \quad D = X - C.$$

proof. Let $x = (b_1, b_2)$ be such that $q_n = 0$ for both b_1 and b_2. It is clear that $Tf(x) \geq f(x)$ and $S_N f(x) \geq Nf(x)$. For such an x we have $x \in C$. Now suppose that $x = (b_1, b_2)$ with $q_n > 0$ for at least one of the beliefs. It is easy to see that $H_t = \sum_j \alpha_{jt}$ for the agents belief at time t is bounded below by t (recall that updating adds one to one of the α_j each time). Hence

$$\begin{aligned} q_{nt} &= \alpha_1(H - \alpha_1)/[H^2(H+1)] \\ &\leq 1/[4(H+1)] \\ &\leq 1/[4(t+1)] \text{ since } H \geq t. \end{aligned}$$

Hence for any intial belief by the agent we have an upper bound on q_n at time t and it follows that we must have $f_t(x) = T^t f(x) = 0$ for all $t > t^* = 1/[4q_n]$ for any f if x is such that one of the agents beliefs has $q_n > 0$. It follows that $S_N f(x)$ is bounded above and $x \in D$.

Q.E.D.

Consider any $f^* \in L_1(X, X^*, \lambda)$ such that $\{x \in X: f^*(x) > 0\} = C$. Clearly such a f^* invariant under T since updating on beliefs in C leaves the agents beliefs unchanged. Now define

$$A_N f = S_N f/N$$

so $A_N f$ is the time average of our probability distribution over the state of the system.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Theorem 8. For all $f \in L_1(X, X^*, \lambda)$, $A_N f$ converges almost everywhere to an invariant distribution f^* such that $\{x \in X: f^*(x) > 0\} \subseteq C$.

proof. T is a positive contraction in L_1 and there exists an invariant f^* which is strictly positive on C. The result follows immediately by applying theorem 3.12 of Krengel (1985).

Q.E.D.

Theorem 8 tells us that there exists an f^* (not necessarily in $L_1(X, X^*, \lambda)$ since this space is not complete) such that $A_N f(x) \to f^*(x)$ for almost all $x \in X$ (that is except on a set of generalized Lebesque measure zero). The sequence of probability distributions $A_N f$ have an associated sequence of probability measures u_n, say, given by

$$u_n(Y) = \int_Y A_N f(x)\, \lambda\, dx \qquad \text{for } Y \in X^*.$$

We say that the sequence of probability measures u_n converge weakly to u^* if for every continuous function $g: X \to R$ we have that

$$\int_X g(x)\, u_n dx \to \int_X g(x)\, u^* dx$$

Theorem 9. The probability measures associated with $A_N f$ converge weakly.

proof. Let $g \in C(X)$ and consider the sequence of random variables defined by g_n on R with distribution given by the probability measures P_n

$P_n(I) = \lambda\{x \in X: g(x)A_N f(x) \in I\}$ for each $I \in R^*$, the Borel sets of R, (normalizing so that $\lambda(X) = 1$) . But

$$|g(x)A_N f(x) - g(x)f^*(x)| \leq |g(x)||A_N f(x) - f^*(x)|$$

$$\leq M\, |A_N f(x) - f^*(x)|$$

Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository. © The Author(s). European University Institute.

where M is the supremum of g(x) over X (which is bounded since g is continuous and X is compact). Hence for any $\delta > 0$ and $\delta^* = \delta/M$ we have

$$\{x \in X: |g(x)A_N f(x) - g(x)f^*(x)| > \delta\} \subseteq \{x \in X: |A_N f(x) - f^*(x)| > \delta^*\}$$

and so

$$\lambda\{x \in X: |g(x)A_N f(x) - g(x)f^*(x)| > \delta\} \leq \lambda\{x \in X: |A_N f(x) - f^*(x)| > \delta^*\}$$

but by theorem 8 $A_N f$ converges to f^* almost everywhere, hence we have that $g(x)A_N f$ converges in probability to $g(x)f^*(x)$, that is

$$\lambda\{x \in X: |g(x)A_N f(x) - g(x)f^*(x)| > \delta\} \to 0 \text{ as } N \text{ becomes large.}$$

Billingsley (1968) shows that convergence in probability implies convergence in distribution and weak convergence of the underlying probability measures on X provided X is separable. This is certainly the case here, so P_n converges weakly. Let $z: R \to R$ be continuous then

$$\int_R z(r)\, P_n dr = \int_X z(g(x)A_N f(x))\, \lambda\, dx$$

converges. Taking $z(r) = r$ we have that

$$\int_X g(x)A_N f(x)\, \lambda\, dx$$

converges. Since the choice of g was arbitrary the sequence $A_N f$ converges weakly.

Q.E.D.

Consider the number of times an agent of type one plays a particular strategy s_1 as a proportion of their his strategy choices. If N rounds have elapsed we can denote this ratio as

$$R_N(s_1) = \sum_{t=0}^{T-1} J_{it}(s_1)/\, N$$

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

where $J_{it}(s_l)$ is an indicator function taking the value one if he plays s_l at time t and zero otherwise. We write $P_{it}(s_l)$ the probability he plays strategy s_i in round t as

$$P_{it}(s_l) = \int_{B_l} \tau(b_l)\, f_{it}(b_l)\, \lambda\, db_l$$

where we take f_{it} and $A_{iN}f$ to be the marginal distributions of f_t and $A_N f$ over B_l.

Theorem 10. $E[R_N(s_l)]$ converges to a limit, p_l^* say, as N becomes large.

proof

$$E[R_N(s_i)] = \lim_{N\to\infty} \sum_{t=0}^{N-1} P_{it}(s_l) / N = \lim_{N\to\infty} \sum_{t=0}^{N-1} \int_{B_l} \tau(b_l)\, f_{it}(b_l)\, db_l /N$$

changing the order of summation and putting the sum of the integrals equal to the integral of the sum gives

$$= \lim_{N\to\infty} \int_{B_l} \tau(b_l)\Big(\sum_{t=0}^{N-1} f_{it}(b_l)/N \Big)\lambda\, db_l = \lim_{N\to\infty} \int_{B_l} \tau(b_l)\, A_{iN}f(b_l)\, db_l$$

By theorem 9 the $A_{iN}f$ converge weakly. By lemma 2 τ is continuous in B_l hence the limit of the integral is equal to the integral of the limit. Hence we have that the limit is exists. So we can write

$$\lim_{N\to\infty} E[R_N(s_l)] = p_l^*.$$

Q.E.D.

The expected value of the proportion of times a particular strategy s_l is played converges to a limit p_l^*. It is easy to show that these limit probabilities p_k^* for $s_k \in S_k$ form an element of M_k.

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Since the limit probability measures u* is the same for each pair of agents in the pools this expected proportion is the same for every agent. The key to this result is that the strategy choices are continuous in beliefs. If this were not the case, that is, if agent did not make small mistakes in their probability assignments, the result is not valid. The limit distribution f* might have an atom at a point where an agent was indifferent between strategies and choose different strategies arbitrarily close to this point. In this case even if f_t converges to f* agents can be jumping arbitrarily between strategies in such a way that the proportions do not converge.

Since this result holds for each pool and an agent has an equal chance of meeting each agent in the opposing pool the expected proportion of times an agent of type one faces a particular strategy s_2 also converges as N becomes large. While this expected value converges agents only observe a sample sequence drawn from this distribution.

Let $J^1_{it}(s_2)$ be an indicator function taking the value one if agent i of type one observes the strategy $s_2 \in S_2$ at time t and zero otherwise.

Define $R^1_N(s_j) = \sum_{t=0}^{N-1} J^1_{it}(s_2)/N$

and let p^*_2 be expected value of the proportion of times agents of type two play s_2 (by applying theorem 10 to agents of type two, this exists).

Theorem 11. $R^1_N(s_2)$ converges to p^*_2 with probability one.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

proof. We can regard each J'_{it} as a random variable whose probability depends on f_t. Consider the conditional distributions $(f_t \mid J'_{i0}, J'_{i1}, J'_{it-1})$ The assumption that there are an infinite number of agents to choose from in replacing the current set in the pool implies that this conditional distribution is simply f_t; the information on the finite set of outcomes gives no assistance in predicting the next outcome since the new agent we are paired with has been chosen randomly for an infinite set. Therefore the random variables in the sequence J'_{it} are independent. It is easy to check that the mean and variance of each J'_{it} is bounded so by the strong law of large numbers (see Loeve 1977) we have

$$P(\lim_{N\to\infty} R'_N(s_2) = \lim_{N\to\infty} E(R'_N(s_2))) = 1$$

But the $\lim_{N\to\infty} E(R'_N(s_2)) = \lim_{N\to\infty} E(R_N(s_2)) = p_2^*$ since each agent i has equal probability of observing each strategy played by the agents in the opposing pool. Hence

$$P(\lim_{N\to\infty} R'_N(s_2) = p_2^*) = 1$$

Q.E.D.

The assumption that the observations are independent is necessary to apply the strong law of large numbers in its strict form to ensure that the average of the sample sequence observed by each agent converges to the expected value of the average. With only a finite number of agents the outcome each period will be correlated and the average of an observed sequence need not converge to the ex ante population average. However, independence is an unnecessarily strict requirement. All we really need is that the average covariance between observation goes to zero as the time elapsed becomes large (see Whittle (1976)); intuitively

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

this means that the covariance between the current outcome and outcomes in the distant future become small relatively quickly.

Theorem 12. The beliefs of all the agents in any pool 1 converge with probability one to put an atom of weight one on the same probability assignment $m_2^* = (p_1^*, p_2^*, \quad , p_n^*) \in M_2$ and the objective probability distribution faced by the agents in pool 1 also converges to this limit.

Proof: By theorem 11 the sample proportions observed by the agents converge with probability one. It is straightforward to show that the beliefs of the agents under Bayesian updating converge to put weight one on these limit sample proportions (our restriction on priors rules out any probability assignment having a prior probability of zero). Each agent's strategy choice (over mixed strategies) is continuous in their beliefs so this also converges (to $\gamma(m_K^*)$). Similarly the beliefs and mixed strategies played by agents in pool two also converge, and are stationary in the limit and so must have the same limit proportions as estimated by the Bayesian updaters in the other pool.

Q.E.D.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

4. Characterization of the Equilibrium.

The strategies played by the agents converge to a limit mixed strategy. We now characterize this equilibrium point of the process.

Consider an N-player, finite strategy, game in normal form. Let m_i be a mixed strategy for player i and let m_{-i} be the vector of mixed strategies played by the other N-1 players. M_{-i} is the set of such mixed strategies. Player i has a payoff function $U_i(s_i, s_{-i})$ where S_i be players i's pure strategy set and S_{-i} is the set of vectors of pure strategies of the other N-1 players. We denote by $C(m_i)$ the set of carriers of m_i, that is, the set of pure strategies which have strictly positive probability in the mixed strategy m_i. Let $m_i(s_i)$ denote the probability attached to strategy s_i in the mixed strategy m_i. We have

$$U_i(s_i, m_{-i}) = \sum_{S_{-i}} U_i(s_i, s_{-i})\, m_{-i}(s_{-i}).$$

Define a norm on player i's mixed strategy set by the relation

$$|m_i| = \max \{m_i(s_i) : s_i \in S_i\}$$

For any vector of mixed strategies for the N agents we define

$$|(m_1, \ \ldots \ , m_n)| = \max \{|m_1|, \ \ldots \ , |m_n|\}.$$

We now define an equilibrium concept.

Definition: The mixed strategy n-tuple $(m_1^*, m_2^*, \ \ , m_n^*)$ is an ε-perfect belief equilibrium if and only if for each $s_i^* \in C(m_i^*)$ there exists a set $G(s_i^*) \subseteq M_{-i}$ of strictly positive (generalized Lebesque)

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

measure such that for all m'_{-i} in $G(s_i^*)$ we have $|m'_{-i} - m^*_{-i}| < \varepsilon$ and $U_i(s_i^*, m'_{-i}) \geq U_i(s_i, m'_{-i})$ for all $s_i \in S_i$.

Most refinements of the Nash equilibrium concept rely on agents making errors, with small probability, in their strategy choices. Examples of this approach are Selten's (1975) perfect equilibrium and Myerson's (1978) proper equilibrium. Here we consider the case where agents choose their strategy perfectly, but make small mistakes in their estimate of the mixed strategy being chosen by the other players.

By our definition if a pure strategy is being played with positive probability it must be a best reply to some set of beliefs, close to the true belief, which has positive measure. The condition that the set must have positive measure is justified by the idea that a (continuously distributed) random mistake will take us to a set of measure zero with zero probability. If we regard the agents as choosing a mixed strategy as a function of his beliefs, this mixed strategy being generated by a continuous probability density function on mistakes in his probability assignment, only strategies which are best replies on a set of positive measure can have positive weight in the mixed strategy.

Lemma 13. The limit mixed strategy (m_1^*, m_2^*) pair given by theorem 12 is an ε-perfect belief equilibrium of the underlying two-person where ε is the maximum size of the error made in the agents probability assignment.

proof. In the limit every strategy played by agents of type one is a best reply to a probability assignment generated by the random error mechanism η which must lie within ε of m_2^*. Since the error is

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

continuously distributed the only strategies in $C(m_1^*)$ are best replies to assignments which have positive measure in this neighbourhood. Similarly for players of type two.

Q.E.D.

Theorem 14. For fixed $\varepsilon > 0$ every N-player, finite strategy, game in normal form has an ε-perfect belief equilibrium.

proof: Take our fixed ε to be the limit of the size of the mistake each player can make in his probability assignment as set out in section two. By theorem 12 the system converges to limit mixed strategies for each agent and the beliefs of each agent converge to the average of these mixed strategies in the other pool. These limit strategies form a ε-perfect belief equilibrium of the game by lemma 13. This proves the theorem for the case N = 2. However, the results of section 3 can be generalized to N player games simply by increasing the belief space of agents to M_{-i}. the convergence results go through unchanged; hence we can construct an ε-perfect belief equilibrium for any N player game.

Q.E.D.

Definition: A n-tuple of mixed strategies $(m_1^*, \quad , m_n^*)$ in an N player game is a perfect belief equilibrium if there exists a sequence of ε-perfect belief equilibria such that ε_k goes to zero and the associated ε_k-perfect belief equilibria converge to $(m_1^*, \quad , m_n^*)$.

Theorem 15. Every N-player, finite strategy, game in normal form has a perfect belief equilibrium.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

proof: Construct a sequence of positive ε_k's which converges to zero. For each ε_k we can find an associated ε_k-perfect belief equilibrium of the Game by theorem 14. The n-tuple of mixed strategies is a compact set so this sequence of ε_k-perfect belief equilibria has a convergent subsequence. Take the limit of this subsequence.

Q.E.D.

Theorem 16. Every perfect belief equilibrium is a Nash equilibrium.

proof. Best replies are upper hemi continuous in beliefs, so for each m_{-i} there exists $\delta > 0$ such that for all m_i' within δ of this point the best rely set contains only best relies to m_{-i}. It follows that at the limit point of the sequence of ε-perfect belief equilibrium as ε goes to zero each agent is playing only best replies to the strategies chosen by the others. Hence this limit point is a Nash Equilibrium.

Q.E.D.

The Perfect Belief Equilibrium concept is a strict refinement of Nash Equilibrium. To show this we prove the following.

Theorem 17. In two player, finite strategy, normal form games every Perfect Belief Equilibrium is a Perfect Equilibrium.

proof: Every perfect belief equilibrium is a Nash Equilibrium by theorem 16. Dominated strategies are best replies only on a space of dimension lower than the opponent's mixed strategy space. This is a set of measure zero; hence dominated strategies are never played in a ε-perfect belief equilibrium. It follows that dominated strategies have zero weight in a perfect belief equilibrium. Van Damme (1983) proves that an equilibrium

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

in a two person game is perfect if and only if it is undominated; hence the perfect belief equilibrium is perfect.

Q.E.D.

To see that the perfect belief equilibrium is is a refinement of perfect equilibrium in two player games consider the following example

Example

		Column Player	
		a	b
	A	1, 1	-1,-1
Row Player	B	0, 0	0, 0
	C	-1,-1	1, 1

The first number gives the payoff of the row player while the second gives the payoff of the column player. The equilibrium (B, 1/2a+1/2b) is both perfect and proper, however, it is not a perfect belief equilibrium. The row player will play B only if his probability assignment is (1/2, 1/2) a point which has measure zero in the space of possible probability assignments.

The definition of perfect belief equilibrium cannot therefore be indentified with either perfect or proper equilibria. However, it seems to be similar to the P-stable equilibria of games where agents have incomplete information about the payoffs of the game they are playing (see Harsanyi (1973) and Van Damme (1983)). Having discussed the nature of the Perfect Belief Equilibrium we now show

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Theorem 18. For ε sufficiently small every ε-Perfect Belief Equilibrium is close to a Perfect Belief Equilibrium.

proof: Suppose not, that is, we can construct a sequence of ε_k's converging to zero, such that each of the associated ε_k-perfect belief equilibria are bounded away from the set of perfect belief equilibria. This infinite sequence has a convergent subsequence which, by the argument of theorem 14, is a perfect belief equilibrium, contradicting the hypothesis.

Q.E.D.

By theorem 12 and lemma 13 the dynamic system converges with probability one to a ε-perfect belief equilibrium of the one shot game and by theorem 18 this equilibrium can be made as close as we please to a perfect belief equilibrium by choosing ε, the maximum size of the agents' errors in their probability assignments, small enough.

5. Conclusion.

We have demonstrated that, under suitable assumptions, in a sequence of games, the strategy choices of agents who form beliefs by Bayesian updating will converge, with probability one, to an equilibrium of the game they are playing. This equilibrium will be close to a perfect belief equilibrium of the game played with complete information where these are defined as Nash equilibria which are stable to small perturbations in agents' beliefs.

While the existence of such perfect belief equilibria is shown for finite strategy games in normal form the usual problem of multiplicity of equilibria persists. However which equilibrium the system converges

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

to under the learning process is determined uniquely by the distribution of the agents' initial beliefs.

References

Aumann R.J. (1964), "Mixed and Behaviour strategies in infinite extensive games, in M. Dresher, L.S. Shapely and A.W. Tucker (eds.), Contributions to the Theory of Games, Vol. IV, Annals of Mathematics Studies 40, Princeton University Press, Princeton.

Aumann R.J. (1986), "Correlated Equilibrium as an Expression of Bayesian Rationality", Econometrica, Vol. 55, pp. 1-18.

Billingsley P. (1968), "Convergence of Probability Measures", J. Wiley and Sons Inc., New York.

Brandenburger A. and Dekel E. (1986), "On the Axiomatic Approach to Refinements of the Nash Equilibrium", University of Cambridge, Economic Theory discussion Paper, No. 104.

Bray M.M. and Savin N.E. (1986), "Rational Expectations Equilibria, Learning, and Model Specification", Econometrica, Vol. 54, pp. 1129-1160.

Damme E. van (1983), "Refinements of the Nash Equilibrium Concept", Lecture Notes in Economics and Mathematical Systems, No. 219, Springer Verlang, New York.

De Groot M.H. (1970), "Optimal Statistical Decisions", McGraw Hill, New York.

Harsanyi J.C. (1967-68), "Games with Incomplete Information Played by Bayesian Players, I, II, and III", Management Science, Vol. 14, pp. 159-182, 320-334, 486-502.

© The Author(s). European University Institute. Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Harsanyi J.C. (1973), "Games with Randomly Distributed Payoffs: a New Rational for Mixed Strategy Equilibrium Points", International Journal of Game Theory, Vol. 2, pp.1-23.

Harsanyi J.C. (1982), "Uses of Bayesian Probability Models in Game Theory", in D. H. Mellor (ed.) Science, Belief and Behaviour: Essays in Honour of R. B. Braithwaite, Cambridge University Press, Cambridge, England.

Krengel U. (1985), "Ergodic Theorems", De Grunter studies in mathematics; 6, de Grunter, Berlin.

Ledyard J.O. (1986), "The scope of the Hypothesis of Bayesian Equilibrium", Journal of Economic Theory, Vol. 39, pp. 54-82.

Loeve M. (1977), "Probability Theory I", 4th Edition, Springer Verlang, New York.

Maynard-Smith J. (1982), "Evolution and the Theory of Games", Cambridge University Press, Cambridge, England.

Myerson R. B. (1978), "Refinements of the Nash Equilibrium Concept", International Journal of Game Theory, Vol. 7, 73-80.

Rosenthal R.W. (1979), "Sequences of Games with Varying Opponents". Econometrica, Vol 49, pp. 1353-1366.

Selten R. (1975), "A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games", International Journal of Game Theory, Vol. 4, pp 25-55.

Whittle P. (1976), "Probability", T. Wiley and Sons Ltd., London.

Winkler R.L. (1972), "Introduction to Bayesian Inference and Decision", Holt, Rinehart and Winston Inc., New York.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

Figure 1. Schematic Representation of the Model.

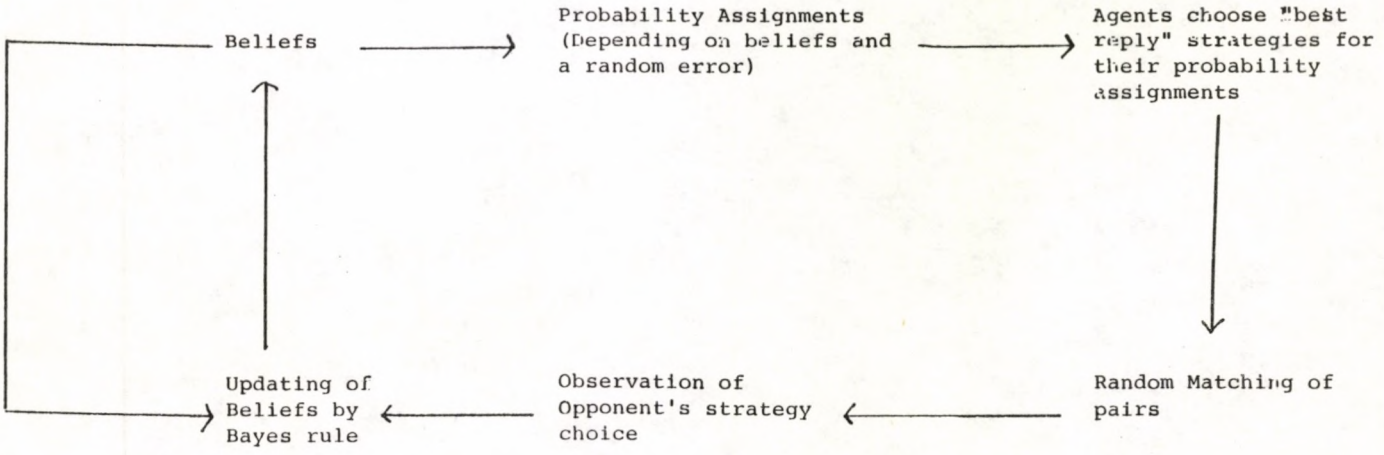

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

WORKING PAPERS ECONOMICS DEPARTMENT

85/155: François DUCHENE	Beyond the First C.A.P.
85/156: Domenico Mario NUTI	Political and Economic Fluctuations in the Socialist System
85/157: Christophe DEISSENBERG	On the Determination of Macroeconomic Policies with Robust Outcome
85/161: Domenico Mario NUTI	A Critique of Orwell's Oligarchic Collectivism as an Economic System
85/162: Will BARTLETT	Optimal Employment and Investment Policies in Self-Financed Producer Cooperatives
85/169: Jean JASKOLD GABSZEWICZ Paolo GARELLA	Asymmetric International Trade
85/170: Jean JASKOLD GABSZEWICZ Paolo GARELLA	Subjective Price Search and Price Competition
85/173: Berc RUSTEM Kumaraswamy VELUPILLAI	On Rationalizing Expectations
85/178: Dwight M. JAFFEE	Term Structure Intermediation by Depository Institutions
85/179: Gerd WEINRICH	Price and Wage Dynamics in a Simple Macroeconomic Model with Stochastic Rationing
85/180: Domenico Mario NUTI	Economic Planning in Market Economies: Scope, Instruments, Institutions
85/181: Will BARTLETT	Enterprise Investment and Public Consumption in a Self-Managed Economy
85/186: Will BARTLETT Gerd WEINRICH	Instability and Indexation in a Labour-Managed Economy - A General Equilibrium Quantity Rationing Approach
85/187: Jesper JESPERSEN	Some Reflexions on the Longer Term Consequences of a Mounting Public Debt
85/188: Jean JASKOLD GABSZEWICZ Paolo GARELLA	Scattered Sellers and Ill-Informed Buyers: A Model of Price Dispersion
85/194: Domenico Mario NUTI	The Share Economy: Plausibility and Viability of Weitzman's Model
85/195: Pierre DEHEZ Jean-Paul FITOUSSI	Wage Indexation and Macroeconomic Fluctuations

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

85/196: Werner HILDENBRAND	A Problem in Demand Aggregation: Per Capita Demand as a Function of Per Capita Expenditure
85/198: Will BARTLETT Milica UVALIC	Bibliography on Labour-Managed Firms and Employee Participation
85/200: Domenico Mario NUTI	Hidden and Repressed Inflation in Soviet-Type Economies: Definitions, Measurements and Stabilisation
85/201: Ernesto SCREPANTI	A Model of the Political-Economic Cycle in Centrally Planned Economies
86/206: Volker DEVILLE	Bibliography on The European Monetary System and the European Currency Unit.
86/212: Emil CLAASSEN Melvyn KRAUSS	Budget Deficits and the Exchange Rate
86/214: Alberto CHILOSI	The Right to Employment Principle and Self-Managed Market Socialism: A Historical Account and an Analytical Appraisal of some Old Ideas
86/218: Emil CLAASSEN	The Optimum Monetary Constitution: Monetary Integration and Monetary Stability
86/222: Edmund S. PHELPS	Economic Equilibrium and Other Economic Concepts: A "New Palgrave" Quartet
86/223: Giuliano FERRARI BRAVO	Economic Diplomacy. The Keynes-Cuno Affair
86/224: Jean-Michel GRANDMONT	Stabilizing Competitive Business Cycles
86/225: Donald A.R. GEORGE	Wage-earners' Investment Funds: theory, simulation and policy
86/227: Domenico Mario NUTI	Michal Kalecki's Contributions to the Theory and Practice of Socialist Planning
86/228: Domenico Mario NUTI	Codetermination, Profit-Sharing and Full Employment
86/229: Marcello DE CECCO	Currency, Coinage and the Gold Standard
86/230: Rosemarie FEITHEN	Determinants of Labour Migration in an Enlarged European Community
86/232: Saul ESTRIN Derek C. JONES	Are There Life Cycles in Labor-Managed Firms? Evidence for France

Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.
© The Author(s). European University Institute.

86/236: Will BARTLETT Milica UVALIC	Labour Managed Firms, Employee Participation and Profit Sharing - Theoretical Perspectives and European Experience.
86/240: Domenico Mario NUTI	Information, Expectations and Economic Planning
86/241: Donald D. HESTER	Time, Jurisdiction and Sovereign Risk
86/242: Marcello DE CECCO	Financial Innovations and Monetary Theory
86/243: Pierre DEHEZ Jacques DREZE	Competitive Equilibria with Increasing Returns
86/244: Jacques PECK Karl SHELL	Market Uncertainty: Correlated Equilibrium and Sunspot Equilibrium in Market Games
86/245: Domenico Mario NUTI	Profit-Sharing and Employment: Claims and Overclaims
86/246: Karol Attila SOOS	Informal Pressures, Mobilization, and Campaigns in the Management of Centrally Planned Economies
86/247: Tamas BAUER	Reforming or Perfecting the Economic Mechanism in Eastern Europe
86/257: Luigi MONTRUCCHIO	Lipschitz Continuous Policy Functions for Strongly Concave Optimization Problems
87/264: Pietro REICHLIN	Endogenous Fluctuations in a Two-Sector Overlapping Generations Economy
87/265: Bernard CORNET	The Second Welfare Theorem in Nonconvex Economies
87/267: Edmund PHELPS	Recent Studies of Speculative Markets in the Controversy over Rational Expectations
87/268: Pierre DEHEZ Jacques DREZE	Distributive Production Sets and Equilibria with Increasing Returns
87/269: Marcello CLARICH	The German Banking System: Legal Foundations and Recent Trends
87/270: Egbert DIERKER Wilhelm NEUEFEIND	Quantity Guided Price Setting
87/276: Paul MARER	Can Joint Ventures in Hungary Serve as a "Bridge" to the CMEA Market?

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

87/277: Felix FITZROY	Efficiency Wage Contracts, Unemployment, and Worksharing
87/279: Darrell DUFFIE Wayne SHAFER	Equilibrium and the Role of the Firm in Incomplete Markets
87/280: Martin SHUBIK	A Game Theoretic Approach to the Theory of Money and Financial Institutions
87/283: Leslie T. OXLEY Donald A.R. GEORGE	Perfect Foresight, Non-Linearity and Hyperinflation
87/284: Saul ESTRIN Derek C. JONES	The Determinants of Workers' Participation and Productivity in Producer Cooperatives
87/285: Domenico Mario NUTI	Financial Innovation under Market Socialism
87/286: Felix FITZROY	Unemployment and the Share Economy: A Sceptical Note
87/287: Paul HARE	Supply Multipliers in a Centrally Planned Economy with a Private Sector
87/288: Roberto TAMBORINI	The Stock Approach to the Exchange Rate: An Exposition and a Critical Appraisal
87/289: Corrado BENASSI	Asymmetric Information and Financial Markets: from Financial Intermediation to Credit Rationing
87/296: Gianna GIANNELLI	On Labour Market Theories
87/297: Domenica TROPEANO	The Riddle of Foreign Exchanges: A Swedish-German Debate (1917-1919)
87/305: G. VAN DER LAAN A.J.J. TALMAN	Computing Economic Equilibria by Variable Dimension Algorithms: State of the Art
87/306: Paolo GARELLA	Adverse Selection and Intermediation
87/307: Jean-Michel GRANDMONT	Local Bifurcations and Stationary Sunspots
87/308: Birgit GRODAL Werner HILDENBRAND	Income Distributions and the Axiom of Revealed Preference
87/309: Eric PEREE Alfred STEINHERR	Exchange Rate Uncertainty and Foreign Trade
87/312: Pietro REICHLIN	Output-Inflation Cycles in an Economy with Staggered Wage Setting

Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.
© The Author(s). European University Institute.

87/319: Peter RAPPOPORT Lucrezia REICHLIN	Segmented Trends and Nonstationary Time Series
87/320: Douglas GALE	A Strategic Model of Labor Markets with Incomplete Information
87/321: Gianna GIANNELLI	A Monopoly Union Model of the Italian Labour Market: 1970-1984
87/322: Keith PILBEAM	Sterilization and the Profitability of UK Intervention 1973-86
87/323: Alan KIRMAN	The Intrinsic Limits of Modern Economic Theory
87/324: Andreu MAS-COLELL	An Equivalence Theorem for a Bargaining Set
88/329: Dalia MARIN	Assessing Structural Change: the Case of Austria
88/330: Milica UVALIC	"Shareholding" in Yugoslav Theory and Practice
88/331: David CANNING	Convergence to Equilibrium in a Sequence of Games with Learning
88/332: Dalia MARIN	Trade and Scale Economies. A causality test for the US, Japan, Germany and the UK.
88/333: Keith PILBEAM	Fixed versus Floating Exchange Rates Revisited.

Spare copies of these working papers and/or a complete list of all working papers that have appeared in the Economics Department series can be obtained from the Secretariat of the Economics Department.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

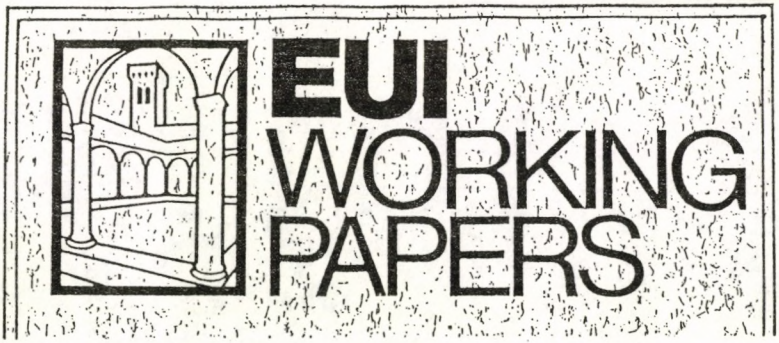

EUI Working Papers are published and distributed by the European University Institute, Florence.

A complete list and copies of Working Papers can be obtained free of charge -- depending on the availability of stocks -- from:

The Publications Officer

European University Institute

Badia Fiesolana

I-50016 San Domenico di Fiesole (FI)

Italy

Please use order form overleaf

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

PUBLICATIONS OF THE EUROPEAN UNIVERSITY INSTITUTE

To The Publications Officer
European University Institute
Badia Fiesolana
I-50016 San Domenico di Fiesole (FI)
Italy

From Name.....
Address.....
.....
.....
.....

Please send me: ☐ a complete list of EUI Working Papers
☐ the following EUI Working Paper(s):

No.:.....

Author, title:.....
.....
.....
.....
.....

Date:.....

Signature: .....

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

PUBLICATIONS OF THE EUROPEAN UNIVERSITY INSTITUTE DECEMBER 1987

87/284:	Saul ESTRIN and Derek JONES	The Determinants of Workers' Participation and Productivity in Producer Cooperatives
87/285:	Domenico Mario NUTI	Financial Innovation under Market Socialism
87/286:	Felix FITZROY	Unemployment and the Share Economy: A Sceptical Note
87/287:	Paul HARE	Supply Multipliers in a Centrally Planned Economy with a Private Sector
87/288:	Roberto TAMBORINI	The Stock Approach to the Exchange Rate: an Exposition and a Critical Appraisal
87/289:	Corrado BENASSI	Asymmetric Information and Financial Markets: from Financial Intermediation to Credit Rationing *
87/290:	Johan BARNARD	The European Parliament and Article 173 of the EEC Treaty
87/291:	Gisela BOCK	History, Women's History, Gender History
87/292:	Frank PROCHASKA	A Mother's Country: Mothers' Meetings and Family Welfare in Britain, 1850 - 1950
87/293:	Karen OFFEN	Women and the Politics of Motherhood in France, 1920 - 1940
87/294:	Gunther TEUBNER	Enterprise Corporatism
87/295:	Luciano BARDI	Preference Voting and Intra-Party Competition in Euro-Elections
87/296:	Gianna GIANNELLI	On Labour Market Theories
87/297:	Domenica TROPEANO	The Riddle of Foreign Exchanges: A Swedish-German Debate
87/298:	B. THOM, M.BLOM T. VAN DEN BERG, C. STERK, C. KAPLAN	Pathways to Drug Abuse Amongst Girls in Britain and Holland
87/299:	V. MAQUIEIRA, J.C. LAGREE, P. LEW FAI, M. De WAAL	Teenage Lifestyles and Criminality in Spain, France and Holland

* :Working Paper out of print

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

PUBLICATIONS OF THE EUROPEAN UNIVERSITY INSTITUTE DECEMBER 1987

87/300:	A. ELZINGA, P. NABER, R. CIPPOLLINI, F. FACCIOLI, T. PITCH	Decision-Making About Girls by the Criminal Justice System in Holland and Italy
87/301:	S. LEES, J. SHAW, K. REISBY	Aspects of School Culture and the Social Control of Girls
87/302:	Eleanor MILLER, Rosa ANDRIEU-SANZ and Carmen VAZQUEZ ANTON	Becoming a Teenage Prostitute in Spain and the U.S.A.
87/303:	Mary EATON and Lode WALGRAVE	A comparison of crime and its treatment amongst girls in Britain and Belgium
87/304:	Annie HUDSON Edna OPPENHEIMER	Towards an effective policy for delinquent girls
87/305:	G. VAN DER LAAN and A.J.J. TALMAN	Computing, Economic Equilibria by Variable Dimension Algorithms: State of the Art
87/306:	Paolo C. GARELLA	Adverse Selection and Intermediation
87/307:	Jean-Michel GRANDMONT	Local Bifurcations and Stationary Sunspots
87/308:	Birgit GRODAL/Werner HILDENBRAND	Income Distributions and the Axiom of Revealed Preference
87/309:	Eric PEREE/Alfred STEINHERR	Exchange Rate Uncertainty and Foreign Trade
87/310:	Giampaolo VALDEVIT	American Policy in the Mediterranean: The Operational Codes, 1945-1952
87/311:	Federico ROMERO	United States Policy for Postwar European Reconstruction: The Role of American Trade Unions
87/312:	Pietro REICHLIN	Output-Inflation Cycles in an Economy with staggered wage setting
87/313:	Neil KAY, Jean-Philippe ROBE and Patrizia ZAGNOLI	An Approach to the Analysis of Joint Ventures
87/314:	Jane LEWIS	Models of Equality for Women: The Case of State Support for Children in 20th Century Britain

* :Working Paper out of print

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

PUBLICATIONS OF THE EUROPEAN UNIVERSITY INSTITUTE FEBRUARY 1988

87/315: Serge NOIRET	Nuovi motivi per studiare i meccanismi delle leggi elettorali. Una riflessione metodologica a proposito della legge del 1919 in Italia
87/316: Alain GOUSSOT	Les sources internationales de la culture socialiste italienne à la fin du 19e siècle et au début du 20e siècle. Problèmes de la composition de l'idéologie du PSI et ses rapports avec la circulation des idées en Europe
87/317: Eamonn NOONAN	Württtemberg's exporters and German protection, 1931-36
87/318: Jean-Pierre CAVAILLE	Theatrum Mundi. Notes sur la théatralité du Monde Baroque.
87/319: Peter RAPPOPORT and Lucrezia REICHLIN	Segmented Trends and Nonstationary Time Series
87/320: Douglas GALE	A Strategic Model of Labor Markets with Incomplete Information
87/321: Gianna GIANNELLI	A Monopoly Union Model of the Italian Labour Market
87/322: Keith PILBEAM	Sterilization and the Profitability of UK Intervention 1973-86
87/323: Alan KIRMAN	The Intrinsic Limits of Modern Economic Theory
87/324: Andreu MAS-COLELL	An Equivalence Theorem for a Bargaining Set
88/325: A. GROPPI	"La classe la plus nombreuse, la plus utile et la plus précieuse". Organizzazione del lavoro e conflitti nella Parigi rivoluzionaria.
88/326: Bernd MARIN	Qu'est-ce que c'est "Le Patronat"? Quelques enjeux théoriques et observations empiriques
88/327: Jean BLONDEL	Decision-Making Processes, Conflicts, and Cabinet Government
88/328: Ida KOPPEN	The European Community's Environment Policy. From the Summit in Paris, 1972, to the Single European Act, 1987
88/329: Dalia MARIN	Assessing Structural Change: The Case

* :Working Paper out of print

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

	of Austria
88/330: Milica UVALIC	"Shareholding" in Yugoslav Theory and Practice
88/331: David CANNING	Convergence to Equilibrium in a Sequence of Games with Learning
88/332: Dalia MARIN	Trade and Scale Economies. A causality test for the U.S., Japan, Germany and the UK
88/333: Keith PILBEAM	Fixed versus Floating Exchange Rates Rivisited
88/334: Hans Ulrich Jessurun d'OLIVEIRA	Die EWG und die Versalzung des Rheins

* :Working Paper out of print

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.

© The Author(s). European University Institute.
Digitised version produced by the EUI Library in 2020. Available Open Access on Cadmus, European University Institute Research Repository.