EUI Working Paper ECO No. 96/12

Learning from Learning in Economics

RAMON MARIMON

EUI WORKING PAPERS

EUROPEAN UNIVERSITY INSTITUTE

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE

ECONOMICS DEPARTMENT

EUI Working Paper **ECO** No. 96/12

Learning from Learning in Economics

RAMON MARIMON

BADIA FIESOLANA, SAN DOMENICO (FI)

Printed in Italy in May 1996
European University Institute
Badia Fiesolana
I – 50016 San Domenico (FI)
Italy

Learning from learning in economics

Ramon Marimon
European University Institute
and
Universitat Pompeu Fabra*

January 1996

"The process of learning,
of the growth of subjective knowledge,
is always fundamentally the same.
It is *imaginative criticism*."
(Karl Popper (1979), p. 148)

Abstract

In this paper I summarize some of the recent literature on learning in games and macroeconomic models. I emphasize adaptive learning: the current efforts to attain a behavioral foundation for a broad class of adaptive learning rules; the alternative definitions of weaker —subjective— equilibrium notions needed to characterize the asymptotic outcome of learning processes; the selection of equilibria when learning rules satisfy certain properties (e.g., experimentation or perturbations); the ability of learning models to explain observed economic phenomena which are not properly accounted for by existing equilibrium theories (in particular, financial and macroeconomic data), and, the possibility of using learning models as normative models to help the design of economic policies and of political and economic institutions.

*Symposia paper of the *7th World Congress of the Econometric Society.* To appear in David Kreps' ed. *Advances in Economic Theory*, Cambridge University Press.I would like to thank Buz Brock, Drew Fudenberg, Ed Green, Seppo Honkapohja, Larry Jones, Ehud Kalai, Ehud Lehrer, David Levine, Albert Marcet, Ellen McGrattan, John Nachbar, Hamid Sabourian, Thomas Sargent, and, specially, Giorgia Giovannetti for their comments on previous drafts, David Kreps, for his patience and encouragement, and all of those who have sent me their current work. Of course, the opinions and remaining misunderstandings can only be attributed to me.

1 Introduction

Learning and evolutionary theory in economics are two related research fields that have experienced exponential growth in the last five years. In a sense, this renewed interest on learning theory seems unjustified: the main questions being addressed are not new. Nash himself wondered how agents will reach the equilibrium he proposed; Muth thought of the rational expectations hypothesis as an extreme counterpart to concurrent macroeconomic models with naïve expectations' formation, to mention two pioneer examples. In this paper I review some of the recent contributions on learning theory in games and macroeconomic models. Such a closer look reveals that even if the old questions have not yet been fully answered and it remains difficult —if not impossible— to set a dividing line between rational and adaptive behavior, a *theory of the learnable in economics* is on its way.

To the question *why do we study learning in economics?* four main, non mutually exclusive, answers have been put forward: *i) Bounded rationality*; *ii) Equilibrium justification*; *iii) Equilibrium selection*, and *iv) Observed non–equilibrium behavior*. Let me briefly comment on each of them.

i) Bounded rationality. It is often argued that "our rationality assumptions are too strong, that our models must take into account human weaknesses such as limited ability to compute, etc." In fact, recent developments in *decision theory* (see, Dekel and Gul contribution to this volume) show how the theory of choice, and equilibrium, can be founded in weaker rationality, and common knowledge, assumptions. Furthermore, as some recent work shows (see Section 3), it is also possible to develop a theory based on "behavioral axioms" as a foundation for adaptive learning models. Nevertheless, such a *choice theoretical model of learning* would be of very limited interest if it only contributed in making our assumptions more realistic. As it has been recognized by many, "realism" can not be the object of economic modeling.

ii) Equilibrium justification. This was Nash's concern. Can we attain equilibria, not as a result of a fixed point argument (i.e., with the help of Adam Smith's "invisible hand"), but as the asymptotic outcome of a decentralized learning process? This remains a focal question and, as we will see in Sections 4 and 5, substantial progress has been done in the last few years. Answers to this question link with applied econometric work: an *observed* social outcome must be *learnable* by agents. That is, positive answers to the above question can help to integrate theoretical and applied work. But, in this sense, learning theory plays a secondary or transitory role: once we have shown that we can *justify* certain equilibrium outcome as the outcome of a learning process, we might as well leave learning aside and concentrate on equilibrium theory. In other words, for learning theory in economics to stand on its own we must look into other explanations.

iii) Equilibrium selection. If not all equilibrium outcomes are equally *justified* by learning, then learning theory can provide a much needed selection criteria. Due to multiplicity —often, indeterminacy— of equilibria many of our theories have little predictable power. In contrast, experimental evidence shows that in many environments some equilibria are more likely to emerge than others, suggesting that our theories need more structure to capture these features. For example, in games, not all Nash equilibria are equally played. A simple benchmark, that anyone can try, is the *Battle of the sexes* game: the *mixed* Nash equilibrium typically does not emerge. Similarly, competitive rational expectations models (REE) with incomplete market structure have, in general, a continuum of equilibria and experimental evidence shows that not all are equally likely to be observed. In fact, the outcomes observed in these different experiments seem to be better characterized by the stability properties of learning algorithms.[1]

The natural step is to replace Nash equilibrium and Rational Expectations Equilibrium for some suitable refinement. But, as we will see in Sections 4 and 5, two problems arise. First, if some *stability* criteria are imposed to select among equilibria, then these criteria are not independent of the class of learning algorithms under consideration. Second, some games try to model relatively sophisticated agents (e.g., CEO's managerial decisions, etc.) and some proposed refinement concepts are based on high levels of deduction by the players involved. Whether such levels of sophistication are achieved is, again, an empirical question, but learning models should allow for these forms of behavior when agents gain enough experience; that is, the "bound" on rationality should be displaced away as agents learn through experience.

Learning theory can help to systematically explore and explain evidence on *equilibrium selection* by providing more structure to the map that defines how agents form and coordinate their expectations, and actions, based on their experience, and capabilities. In particular, alternative consistency and behavioral conditions on learning rules should "make falsifiable" the set of observed paths and, in turn, the set of *learnable equilibria*.

[1]For an overview of experimental evidence see Kagel and Roth (1995). For games, see also Crawford's contribution to this volume, and for evidence on macroeconomic models see Section 5.

iv) Observed non–equilibrium behavior. While existing equilibrium theories must be "refined" to account for some economic facts, they should also be "extended" to account for others. While equilibrium conditions must be stringent to preclude certain unobserved multiplicities, at the same time, may have to be weakened to account for observed persistent patterns that seem to show the complexity of attaining certain equilibrium outcomes. The evidence comes from different sources. The experimental lab provides a useful tool to isolate these phenomena. There are, at least, three types of experimental facts that should be considered and a learning theory may help to explain. First, the "sensitivity to marginal payoffs," that is, typically agent's actions are affected by the relative performance of their actions, something not taken into account by a simple –marginalist– maximization point of view, which only considers that the best action is taken. Second, the "experience or expectations across environments," seems to affect agents' behavior. In a theoretical model, agents' behavior (e.g., in which information, and how, should they condition their decisions?) is usually predetermined. Nevertheless, experience —and/or expectations about— other environments may result in well defined "non equilibrium" patterns by agents that try to act rationally across different environments. Third, some equilibria seem to be "stable" in the large but "locally unstable". For example, equilibrium selection may occur for low frequency data, while high frequency data may show persistent volatility. As a result, observed patterns around the equilibrium are more complicated than standard theory prescribes; or may be misinterpreted by unproperly randomizing the model.

Macroeconomic and, in particular, financial data also provide evidence which is difficult to reconcile with existing dynamic rational expectations equilibrium models, but may be explained by taking the learning process into account. For example asset market data shows that there is persitency of returns' volatility and of trading volume (and cross–correlation between absolute returns and trading volume). Another example, from macroeconomics, is the existence of inflationary spells which can not be accounted as non–stationary rational expectations equilibrium paths (see, Section 5).

This possibility of explaining some *non–equilibrium behavior* raises an obvious question. Does learning theory account for any possible pattern? That is, does it *justify "everything goes"*?. Unfortunately, there are examples that seem to point out in this direction (e.g., convergence to fairly arbitrary "aspiration levels" or to indeterminate REE, etc.). However, a closer look at these examples shows that agents must *coordinate* in fairly arbitrary learning rules to attain such outcomes. As we will see, when the learning process is considered as a decentralized process in which agents' learning rules can be arbitrarily chosen from a large class (satisfying basic behavioral and/or consistency conditions). then it does not justify *everything goes.* As we will see in Section 4 and 5, on the one hand, notions of equilibrium, which are weaker than Nash, or REE, may be defined by only requiring, for example, that agents' subjective beliefs are self–confirmed in equilibrium. With this notion, non–Nash (or REE) equilibrium patterns are possible, but not all patterns are possible. In particular, in "simple" environments only Nash (or REE) are equilibria. On the other hand, for large classes of adaptive learning rules it is possible to dicriminate among Nash, or REE, equilibria.

The last two answers, which I would rank as more relevant, can be summarized by saying that the development of a *theory of the learnable in economics* can provide our equilibrium theories with more predictable power, by helping to reject some equilibria and by helping to account for other non–equilibrium patterns. But then, learning theory plays can also play other roles in economics:

v) The study of complex economic environments. If a class of learning models provides consistent results in well understood environments, then such a class can be used as a tool to explore more complicated environments that we do not know how to characterize *ex ante.* For example, if the learning model converges, then it is usually possible to characterize *ex post* the resulting outcome; e.g., see whether it is a particular type of equilibrium. Some models with Artificially Intelligent Agents (AIA) have been successfully applied in this way. In fact, in the last few years a growing number of computational (and estimation) algorithms based on learning and evolutionary principles have been developed. These algorithms have been used both as "theorist tool" to study, for example, nonlinear stochastic rational expectations models or as an "applied economist, or financial analyst, tool" for estimation and prediction. The development of learning theory can help to characterize the "learnable solutions" that such computational and estimation procedures may attain (see, for example, Sargent 1993 and White 1992).

vi) As a normative theory. This is an area of research which has been little explored, but learning can contribute to our normative theories in different ways (see Section 5 for some examples). First, a disturbing feature of models with multiplicity of equilibria is that the welfare implications of an economic policy usually are different across equilibria. If, for a large class of learning algorithms and initial conditions, the process converges to a particular equilibrium, then one can make policy prescriptions with high confidence. Second, new policies, or institutions, can be designed taking into account the fact that agents must learn from their experience; say, about the effects of different incentives or tax schemes. For example, many economic policies suffer from some form of indeterminacy since the effect of their announcement depends on agents' expectations. While there is an extend literature on "credibility problems" as incentive problems, reputation (inference from policy announcements given past experience) is also a learning problem and a better understanding of it can help policy design. Similarly, market and organizational forms can be designed taking into account how institutional arrangements affect the stability properties of equilibria and, therefore, affect welfare when the process of convergence to equilibrium is also taken into account. Third, the designer, or planner, may have to take into account that he, or the designed organization, has also to learn. In other words,the problem may not of designing a contract or organization that it is efficient from a given period zero and is simply executed as uncertainties unfold, but the problem may be to design a contract that will adapt well to "unforeseen contingencies;" in a similar fashion that living organisms can perform well, but not because they are executing God's plan at period zero.

As it can be seen, the scope of a *theory of the learnable in economics* goes well beyond the standard —and Nash's original question— of whether we can justify an equilibrium concept as being the result of a learning process. My aim in the remaining of the paper is to summarize some of the recent results that, I think, are helping in building up such a theory. Given the rapid growth of this literature, I only report on some of the recent contributions on individual learning (Section 3), learning in games (Section 4) and learning in dynamic macroeconomic models (Section 5).[2] In the next Section, I introduce some notation and describe a basic framework.

2 Basic framework

A large class of economic, and game theoretical, models can be casted in a relatively simple general framework. There is a *set of agents* I, time is discrete $(0, 1, \ldots)$. At each period of time there are actions sets for agents, A_i, a public outcome set X, and a set of states of nature S. In period t, agent i's one-period *ex-post* payoff, are represented by $u_i(a_{i,t}, x_{t+1})$, and *ex-ante* present value payoffs by $(1-\delta)E\sum_{n=0}^{\infty}\delta^n Eu(a_{i,t+n}, x_{t+1+n})$, where $\delta \in [0, 1)$ (notice

[2]See Kandori's contribution to this volume for a review of the, closely related, *evolutionary* and *social learning* models.

that $\delta = 0$ corresponds to *myopic* behavior and the $\lim_{\delta \to 1}$ is given by the long run average expected payoff: $\lim_{T\to\infty} \frac{1}{T} \sum_{n=0}^{T+t-1} Eu(a_{t+n}, x_{t+1+n}))$. The public outcome evolves according to: $x_t = \phi(g(a_t, \ldots, a_{t-m}), x_{t-1}, \ldots, x_{t-m}, s_t)$, where g aggregates the actions of all the agents, and $s_t \in S$ is an exogenous shock.[3]

Within this framework we can consider several environments. *i) Single agent without intertemporal feedback*: $I = \{i\}$ and $x_t = s_t$. This is the standard case of individual decision theory with uncertainty.*ii) Single agent with intertemporal feedback*: $I = \{i\}$ and $x_t = \phi(g(a_{i,t-1}, \ldots, a_{i,t-m}), x_{t-1}, \ldots, x_{t-m}, s_t)$. Individual investment problems take this form and, as we will see in Section 3, most learning difficulties, usually associated with multiagent problems, already appear in these environments. *iii) Multiagent without intertemporal feedback*: $I = \{1, \ldots, n\}$, and $x_t =, \phi(g(a_{t-1}), s_t)$. The standard example is a *multistage game*[3]: $x_t = a_{-i,t-1}$ (notice that for every player a "public outcome" is defined). *iv) Competitive with intertemporal feedback*: $I = [0, 1]$ (i.e., a continuum of agents) and $x_t = \phi(g(a_t, a_{t-1}), x_{t-1}, s_t)$. For example, in a deterministic *temporary equilibrium model*, such as the standard *overlapping generations model*, $x_t = g(a_t, a_{t-1})$.

In most of the above environments, the actions of agent i do not affect the public outcome. Then, from the perspective of agent i, given the actions of other agents $\{a_{-i,t}\}$, the public outcome can be thought of as "the state of the outside world." The realizations of the stochastic process $\{x_t\}_{t=0}^{\infty}$ are governed by a probability measure ν on $(X_\infty, \mathcal{F}_\infty)$; where $X_t = \prod_{t=0}^{t} X$ and $\mathcal{F}_t$ is the σ-field generated by $(x_0, \ldots, x_t)$, the corresponding empirical distribution –as it is perceived by agent i– is denoted by m_t (with subindex only when needed); *i.e.*, m_t is a probability measure on $(X_t, \mathcal{F}_t)$. Even when agent i's actions affect the public outcome, as long as he does not take account of this dependence, he may still consider the public outcome as "the exogenous outside world."

Some further notation is needed to denote agent's decisions. At any point in time –say, t– the agent must take a (mixed) action, $\alpha_t \in \Delta(A)$, based on all his past information, represented by $h_t \in H_t$ (h_t may include $(x_0, \ldots, x_t), (s_0, \ldots, s_{t-1})(a_0, \ldots, a_{t-1})$, as well as all the information that the agent has about (A, X, S, u, and ν)). A *behavioral rule* (or strategy), σ maps histories into (mixed) actions (i.e., $\sigma : H_\infty \to \Delta(A)$, σ_t is $\mathcal{H}_t$ measurable).

---

[3]I use the following standard notation: if a_t is denotes the actions of all the agents in period t, then $a_t = (a_{i,t}, a_{-i,t})$, where $a_{-i,t}$ denotes the actions of "all agents but i". Furthermore, $a^t = (a_0, \ldots, a_t)$; similarly for other variables. Also, when there is no confusion, $u(\alpha, m)$ denotes expected utility with respect to the mixed action α and the distribution m.

In the standard framework, A, X, S, u, and ν are assumed to be known, and the Savage or Anscombe and Aumann (1963) axioms prescribe that the agent will choose the path of actions that maximizes expected utility. In this case, it is irrelevant for the solution of the problem whether the agent is facing a "once and for all" problem or has accumulated "experience." The problem, even in its infinite-horizon formulation is essentially a static problem that has a resolution in period zero by choosing the optimal behavioral rule. In the general equilibrium framework, this corresponds to solving for Arrow-Debreu contingent plans in period zero. As it is well known, even in this context, the characterization, and computation, of optimal behavioral rules is greatly simplified if they have a *recursive structure*. That is, whether there is an "optimal policy" R, such that optimal σ^* satisfy $\sigma_t^*(h_t) = R(x_t, \theta_t)$, where θ_t is some vector statistic from h_t which follows a pre-specified law of motion $\theta_t = \psi(x_{t-1}, \theta_{t-1})$ (with initial conditions).

In a learning process, *experience* is essential and so it is some form of recursivity that explicitly defines how experience is accumulated. As we will see, when $\sigma_t^*(h_t) = R(x_t, \theta_t)$, R may be a fairly complicated object, which may include a selection procedure among simpler recursive rules, an internal accounting system, etc. In many economic problems, however, a relatively simple behavioral rule R is assumed. For example, $\sigma_t^*(h_t) = BR(x_{t+1}^e)$, where $BR(x_{t+1}^e)$ is –a selection from– agent's *best reply* to his expectations about the public outcome x_{t+1}^e, and $x_{t+1}^e = f(x_t, \theta_t)$, where f is a forecasting rule. For example, $\theta_t = (x_t^e, \alpha_t)$, $f(x_t, x_t^e, \alpha_t) = x_t^e + \alpha_t(x_t - x_t^e)$ and $\alpha_t = \frac{\alpha_{t-1}}{\alpha_{t-1}+1}$; that is, if $\alpha_0 = 1$ then x_{t+1}^e is the empirical mean. Notice that, in this case, $R \equiv BR \cdot f$ and $\psi(x_t, \theta_t) \equiv (f((x_t, x_t^e, \alpha_t), \alpha_t + 1)$. More generally, one can consider R as an element of a class of learning rules $\mathcal{R}$. In this case, the agent may "choose how to learn," that is, select which learning rule $R \in \mathcal{R}$ fits better his learning needs. For example, in the last formulation, the class of learning rules, $\mathcal{R}$, may be defined by two parameters (a, α_0) characterizing the *step size* sequence: $\alpha_t = \frac{\alpha_{t-1}}{\alpha_{t-1}+a}$ (e.g., if $a = (1 - \alpha_0)$ then $\alpha_t = \alpha_0$ for all t), or by having a finite memory with some parametric form.

It is standard to separate decision models, according to the number of agents, that is whether it is an individual choice model, a finite agent, or competitive equilibrium model. I will follow this classical separation in the following sections. However, from the point of view of learning theory, more than the number of agents, what matters is how agents actions *feedback* into, or are correlated with, the public outcome. In a sense, *learnable equilibria are solutions to multiagent problems achieved without recursion to a fixed point argument.*

For an equilibrium to be the asymptotic outcome of a learning process (i.e., a *learnable equilibrium*) it is required that for every individual agent his expectations about the public outcome are *self-confirmed*, and that his resulting actions do not disturb such expectations. Different forms of *feedback* –together with the *complexity* of the "outside world" and of the decision process– make such convergence more or less difficult to achieve. This will be a recurrent theme through the following sections.

3 Learning to choose

Ad-hoc learning economic models have been around for long time without building up to a theory of adaptive learning. In contrast, Bayesian statistical methods have given foundation to a well developed theory of *Bayesian learning*. While the basic elements of the theory date back to the development of the corresponding statistical theory (see, for example, de Groot 1970), the theory was "put in use" in the seventies and eighties as a learning foundation of rational expectations equilibria by Jordan and Blume and Easley, among others, and in the early nineties as a learning foundation of Nash and Correlated equilibria by Jordan (1991), Kalai and Lehrer (1993a, 1993b, and 1995) and Nyarko (1994), among others.

3.1 Rational learning?

Bayesian learning has been labeled "rational learning" since *within* the Bayesian framework, the Bayesian learner satisfies the standard rationality axioms. In particular, along the learning process, the agent follows "optimal statistical procedures" and his views of the world "can not be contradicted," if anything, they become "more accurate." However, as it is well known, the quotation marks can not be dismissed. In relatively complex environments the optimal procedures are optimal only in reference to the agent's simplified view of the environment, not in relation to the actual environment; the agent's simplified view maybe highly counterfactual, although the agent may not perceive this and, as a result, the agent's predictions may be far from converging to the right predictions. It is for these reasons that Bayesian learning will only be a first short stop in our search for a theory of the learnable in economics (see Blume and Easley 1995 for a more complete account).

A Bayesian learner summarizes his uncertainty about the aspects of the economy that are unknown to him with an appropriate prior distribution. That is, the *agent knows what he does not know*. For example, if payoffs are unknown the prior must also be defined over possible payoff functions. To make the problem manageable, most Bayesian learning model reduce learning to a forecasting problem. The Bayesian learner starts with a prior μ on $(X_\infty, \mathcal{F}_\infty)$, and follows a process of Bayesian updating $\mu(x_{t+1}|x^t)$. The question is whether this updated beliefs converges to the real distribution $\nu(x_{t+1}|x^t)$. If such convergence is achieved then forecasted beliefs and the objective distribution are said to *strongly merge*. The basic result in which this literature builds is a theorem by Blackwell and Dubins (1963) which says that if ν is *absolutely continuous* with respect to μ (denoted $\nu << \mu$), then the forecasts of the corresponding conditional distributions *strongly merge* ($\nu << \mu$ if for any subset D of X_∞, $\nu(D) > 0$ implies $\mu(D) > 0$). That is, the forecaster can not place *ex-ante* zero probability to paths that have positive probability in the environment (see Kalai and Lehrer 1993a and Nyarko 1994).

Blackwell and Dubins' theorem exemplifies the strengths and weaknesses of Bayesian learning theory. The theorem, as many analytic learning results, is based on the martingale convergence theorem (and the Radon-Nikodyn theorem) and it provides a very strong tool for convergence theorems. However, it also also shows how the learner must "show his cards" *ex-ante* by committing himself to a prior μ and to follow the Bayesian rules of the game. To see this, consider a simple environment where $X = \{0, 1\}$. Let ν define a sequence of Bernoulli trials (i.e., Prob$\{x_t = 1\} = p \in (0, 1)$ for all $t \geq 1$) and let the Bayesian's prior distribution on $(0, 1)$, $\mu(p_0)$, be a *Beta* distribution with parameters $(q_0, 1 - q_0)$. Then the Bayesian process of updating posteriors $\mu(p_{t+1}|x^t)$ is –almost– that of a *frequentalist* in the sense that, $p^e_{t+1}, = q_t$ and, for $t > 0$, $q_t = q_{t-1} + \frac{1}{t+1}(x_t - q_{t-1})$ and $\mu_t \rightarrow \delta_p$. That is, his forecasts are accurate in the limit.[4]

Two small variations of the above endowment are enough to question the "rationality" of the Bayesian learner. First, suppose that $X = \{0, 1, 2\}$, but that the agent thinks that $X = \{0, 1\}$ and behaves as before. The unexpected event $x_t = 2$ will be ignored by the dogmatic Bayesian while it will be easily accounted for by the frequentalist. Second, consider that $X = \{0, 1\}$, but

[4]Here δ_p denotes the (Dirac) distribution with point-mass at p. Notice that a *frequentalist* will have *point-expectations* $p^e_{t+1} = p^e_t + \frac{1}{t}(x_t - p^e_t)$ for $t > 0$. That is, the mean forecasts of a Bayesian with an appropriate Beta prior distribution are those of a *frequentalist* with a prior q_0 (see de Groot 1970).

that ν defines a deterministic cycling sequence $(0, 1, 0, 1, \ldots, 0, 1, \ldots)$. Both, the Bayesian and the frequentalists described above will converge to the belief $p^e = 1/2$, disregarding the "obvious" cycling *pattern*.

These last examples show the importance of some form of *absolute continuity* condition on the prior beliefs with respect to the "true" environment in order to achieve accurate forecasts (*strong merging*). But, as Nachbar (1995) has recently shown, even when such absolute continuity conditions are satisfied, there is another limitation on the "rational learning" approach that must be taken into account. An idea underlying "rational learning" is that if an agent knows how to *best reply* to his forecast, then *strong merging* of beliefs results in *optimal* outcomes against *accurate predictions*. This is correct if one considers stationary environments in which the actions of the individual agent do not affect the distribution of the public outcome, but when such distribution is affected by the agent's actions then this property of "rational learning" may be too much to ask for.

To better understand this difficulty consider the following example. Let $X = \{0, 1\}$, $A = \{a^1, a^2\}$, $u(a^1, 1) = u(a^2, 0) = 0$, $u(a^1, 0) = u(a^2, 1) = 1$. Suppose that a bayesian learner starts with a prior μ satisfying the absolute continuity condition with respect to a class of envirnonments $\nu \in \mathcal{N}$. His prior and the process of bayesian updating will result in a *best reply* behavioral strategy, σ, for *any* environment $\nu \in \mathcal{N}$. Suppose that such strategy is a pure strategy, that is $\forall h_t$, $\sigma_t(h_t) \in A$. Now consider a class of "miscoordination environments" defined by $\text{Prob}\{x_t = 0|x^{t-1}\} > 1/2$ if $\sigma_t(x^{t-1}) = a^2$ and $\text{Prob}\{x_t = 1|x^{t-1}\} > 1/2$ if $\sigma_t(x^{t-1}) = a^1$. If these environments are in $\mathcal{N}$ then the agent will, eventually learn to forecast, but then σ will no longer be his optimal strategy. Associated with the new *best reply* strategy there are other "miscoordination environments," etc. The problem is that, in general, one can not "close" this process. That is, in general, one can not have an optimal strategy for all environments in $\mathcal{N}$ and, at the same time, remain within this class $\mathcal{N}$ when the the *feedback* from the optimal strategy is taken into account. For example, if, moving ahead, we consider the public outcome as the resulting action of a second player in a *Battle the Sexes* game, then the lack of "closedness" means that if $\mathcal{N}$ is "rich enough" it can not simultaneously be the set of *plausible* strategies of the opponent and the set to which best reply strategies belong.

Nachbar's warning is in the spirit of the *impossibility theorems* that, since Göedel's theorem, have been common in the "logicians approach" to learning theory and, as Binmore (1988) has shown, are also present in our context. Far from being "impossibility of learning" results are warnings against us making excessive "rationality and completeness" demands.[5] With this lesson of humility, a first direction to pursue is to be less demanding about what is meant by a learnable outcome.

3.2 Rational, calibrated beliefs and other consistency conditions

Strong merging of beliefs requires that even tail events are forecasted correctly. This is a very *strong consistency* condition. A weaker, and reasonable, condition is to require only accurate forecasts over arbitrarily long but finite future horizons. If this property is satisfied then it is said that forecasts *weakly merge* with the true environment (see, Lehrer and Smordinsky 1993). Therefore, in non-stationary environments, tail events may prevent forecasts from being accurate and, in multiple agent environments, may prevent agents from asymptotically agreeing.

Similar ideas have been formalized by Kurz (1994a, 1994b) as a basis for his concept of *rational belief equilibria.* A rational belief is one that it is "consistent" with the observed long-run empirical distribution. That is, even if ν does not define a stationary process, as long as the long-run empirical distribution is well defined –say, m– then the environment is considered to be *stable.* *Rational beliefs,* μ, are those which are "compatible with the data," in the sense of generating the same empirical distribution m, which must also be absolutely continuous with respect to the beliefs (i.e., $m << \mu$; see, Kurz (1994a) for details). Since many measures on $(X_\infty, \mathcal{F}_\infty)$ are compatible with the same empirical distribution, heterogeneity of beliefs (at the tail!) is allowed. A rational believer can be thought as a bayesian learner that is born with experience, and forced to use the empirical distribution as a prior, and for which forecast *weakly merge* with the environment.

[5] In the context of evolutionary repeated games, Anderlini and Sobourian (1995) provide a global convergence result for computable environments which contrasts with Nachbar's negative results. We can translate an informal version of their results to our context. Now instead of having a single behavioral rule, the agent will be considering a large number of them and, as in an evolutionary process, give more weight to those rules that are performing better. Anderlini and Sabourian show that asymptotically an optimal rule will be selected, which, in the context of games, means there is *global convergence* to a Nash equilibrium. Anderlini and Sabourian' result does not contradict Nachbar's since the selected rule will depend on the environment, which does not mean that such specific rule could not be tricked in another environment.

A similar -and weaker- concept is that of a "well calibrated forecaster" (see, Dawid (1982)). A well calibrated forecaster is one whose forecasts -asymptotically- conform with the empirical distribution. For example, the stock price should go up 2/3 of the periods that it is forecasted that it will go up with probability 2/3. Foster and Vohra (1995a and 1995b) have recently further developed this idea and applied to study convergence of calibrated learning schemes to correlated equilibria (see Section 4).

It should be clear that "calibrated forecasts" may be far from accurate. Consider again the cyclic environment $X = \{0, 1\}$ with ν generating $(0, 1, 0, 1, \ldots, 0, 1, \ldots)$. A forecast of $x_t = 1$ with probability one half will be well calibrated while an accurate forecast for ν^1 will make a prediction of x_t conditioned on x_{t-1}; or conditioned on whether t is odd or even.

As in statistical inference, alternative *consistency conditions* provide different convergence tests. In line with statistical tests, Fudenberg and Levine (1995a) propose a weak asymptotic consistency test to assess the "fitness" of a behavioral rule. A behavioral rule σ is said to be $(\epsilon - \delta)$ *consistent* if there exist a T such that for *any i.i.d. environment* ν, and for $t \geq T$ with *confidence* δ the realized long-run average payoff (up to t) is, within an *error* ϵ, at least as good as the optimal -one period- mixed strategy against the empirical distribution m_t.[6]

In other words, Fudenberg and Levine take the empirical distribution -not the "true environment"- as reference and require that the behavioral rule performs almost as well as optimal play against the empirical distribution. Immediately comes to mind a simple generalization of the *frequentalist forecaster*. The well known *fictitious play* rule: compute the empirical distribution of the opponents' play and best reply against such empirical distribution. The problem is that while the fictitious player in *i.i.d* environments may attain accurate forecasts and almost maximal payoffs, one can also construct "miscoordination" (non i.i.d.) environments to trick a fictitious player. However, Fudenberg and Levine (1995) show that an exponential version of fictitious play is $(\epsilon - \delta)$-consistent for *any environment*; they call such property *universal consistency*.

To base consistency tests on the *empirical* distribution, instead than on the *true* distribution, is a necessary step for a *theory of the learnable*. However, two difficulties must be considered. First, the fact that *the agent perceives the*

[6]That is, $\nu\{h_t | \frac{1}{t} \sum_{n=0}^{t} u(a_n, x_n) + \epsilon \geq \max_\alpha u(\alpha, m_t)\} \geq (1 - \delta)$. Valiant's *Probably Approximately Correct (PAC)* learning theory relates the $(\epsilon - \delta)$ and T consistency requirements with the *complexity* of the class of concepts to be learned (*sample complexity*) and with the *computational complexity* of the learning algorithm (see, Natarajan (1991)).

marginal empirical distribution and in environments with *feedback* the actual –correlated– distribution may be difficult to recover from such a partial picture. Second, the problem, already mentioned, of *pattern recognition* of the empirical distribution, that is, of finding the right conditioning of the data. The following example illustrates these difficulties.

Example 1 There are three actions and three possible states. Payoffs are given by

	x^1	x^2	x^3
a^1	0	2	1
a^2	1	0	2
a^3	2	1	0

I consider two environments, in the first environment, private actions and public outcomes are correlated in the following form: $\text{Prob}\{x_t = x^1|a_t = a^1\} = \text{Prob}\{x_t = x^2|a_t = a^2\} = \text{Prob}\{x_t = x^3|a_t = a^3\} = 0$, and $\text{Prob}\{x_t = x^2|a_t = a^1\} = \text{Prob}\{x_t = x^3|a_t = a^1\} = \text{Prob}\{x_t = x^1|a_t = a^2\} = \text{Prob}\{x_t = x^3|a_t = a^2\} = \text{Prob}\{x_t = x^1|a_t = a^3\} = \text{Prob}\{x_t = x^2|a_t = a^3\} = 1/2$. Notice that if the agent plays the three actions with (almost) equal probability, the empirical distribution, m, can converge to the uniform distribution and well calibrated (universally consistent) beliefs will be *self confirmed*, and the agent will receive an average payoff close to 1. Nevertheless, if the agent conditions on his own actions, or if simply sticks to one action, he can receive a payoff of almost 1.5 and with respect to such (conditioned) empirical distribution be a well calibrated forecaster.

In the second environment (a one agent version of Shapley's (1964) example), there is the following *intertemporal feedback*: $\text{Prob}\{x_t = x^1|a_{t-1} = a^2\} = \text{Prob}\{x_t = x^2|a_{t-1} = a^3\} = \text{Prob}\{x_t = x^3|a_{t-1} = a^1\} = 1$. If the agent behaves as a well calibrated (unconditional) forecaster, the history of play will cycle and the perceived distribution can approximate the uniform distribution. Again, almost equal weight to all three actions can be a consistent behavior. But then, the agent does not take into account two important elements. First, along any (countable) history, the agent infinitely often assigns positive probability to zero probability events (see Jordan 1993 for a similar remark). Second, at the beginning of period t, there is a much better sufficient statistic to predict the future outcome, x_t: the last action of the agent a_{t-1}. That is, with the same information –and different conditioning– the agent can receive an average payoff of almost 2.

In a sense, Fudenberg and Levine' *universal consistency* condition tests a behavioral rule against an econometrician who knows how to adopt the optimal action with respect to the empirical distribution. That is, there still is an external element of reference. However, it may be that either the agent is constrained to a set of behavioral rules, $\mathcal{R}$, or that the problem is complex enough that the econometrician is not in a better position than the learning agent. In these contexts, one may have to consider an *internal consistency* condition with respect to a class of behavioral rules, $\mathcal{R}$. This is, for example, the approach taken by Evans and Ramey (1994), Brock and Hommes (1995) and by Nicolini and Marcet (1995) (see Section 5). Brock and Hommes consider learning agents as consumers of forecasting services who are willing to tradeoff accuracy for lower forecasting fees. Nicolini and Marcet consider agents that choose, among a class of adaptive learning rules, the one that better tracks the possible nonstationarities of their economic environment. Consistency is then defined in relative terms within such a class. In these models that agents "choose how to learn" consistency is not defined only asymptotically, since agents keep testing different rules and which one is the best rule may keep changing in a nonstationary environment.

We have been moving away from the set frame of Bayesian learning. In considering classes of rules and, particularly, in studying their "consistency," opens two obvious questions (see Kreps 1990): Can we characterize a broad class of adaptive learning algorithms based on few behavioral assumptions? Does such behavior result in "standard" optimal, or equilibrium, choices in relatively simple environments?

3.3 Learnable choice

In the context of games, Milgrom and Roberts (1991) have provided a first characterization of a learning class based on a weak monotonicity condition: "an adaptive rule must only play strategies that, in relation to the observed finite history, are undominated." Marimon and McGrattan (1995) consider a general class of learning rules based on a stronger monotonicity condition (together with *experimentation* and *inertia*): an adaptive learning rule must move in the direction of the –possibly, unknown– best reply map , that is, must follow the evolutionary principle of assigning more weight to actions that have been shown to work better in the past.

Recently, Easley and Rustichini (1995) have "rationalized" replicator type dynamics. They consider a class of finite memory simple rules –say, a pure strategy– and adaptive rules $\mathcal{R}$ which select –or weight– different simple rules. They impose basic "axioms" on the selection procedure $\mathcal{R}$: *i)* a *monotonicity* condition: increase the weight on rules that have relatively high payoffs; *ii)* a *symmetry* condition: selection procedures should not be affected by relabeling, and *iii)* an *independence* (of irrelevant alternatives) condition: the effect of a payoff on the weight of a simple rule should be independent of other payoffs. In the context of single agent problems without intertemporal feedback, they show that: first, an adaptive rule $\mathcal{R}$ satisfying these three axioms asymptotically selects rules which are objective expected utility maximizers, among the set of simple rules; second, any expected utility maximizer rule can be selected by a procedure satisfying the axioms, and, third the class of selection procedures, $\mathcal{R}$, satisfying these axioms are, at least asymptotically, strict monotone transforms of replicator dynamics.

In particular, the family of *exponential fictitious play* rules, which, as we have seen before, Fudenberg and Levine (1995a) show is $(\epsilon - \delta)$-*universally consistent*, is representative of the class of rules satisfying Easley and Rustichini' axioms. In other words, the universal consistency condition is satisfied for a class of behavioral rules characterized by few "behavioral axioms." Furthermore, a representative of such a class can be constructed.

Unfortunately, in more complex environments, these basic axioms may not be enough for learning. First, when either the learning problem is not just a forecasting problem, and the agent must also learn about his payoffs or the set of actions at his disposal, or in non stationary environments, *experimentation* is an essential feature of the learning process. Second, in environments with *intertemporal feedback* from individual agents, there may be a problem of *overreactions* resulting in agents never learning the final consequences of their actions. *Inertia*, that is to stick to the same action with some probability, is a form of *building stationarity* into the environment on the part of the agent. In summary, in addition to a *strong monotonicity condition* –say, of the replicator dynamics type– *experimentation* and *inertia* characterize a general class of learning rules and a more general theory of *learnable choice* must take them into account (Marimon and McGrattan 1995; see also Fudenberg and Levine 1995a and Kaniovski and Young 1995).

3.4 Explicit behavioral rules

Learnable outcomes must be achieved by explicit behavioral rules. Experimental economics, applied decision theory, and behavioral sciences may help us understand how economic subjects actually learn from their experience, but in the same way that we need explicit functional forms representing preferences and technologies to test and develop our economic theories and econometricians need explicit estimation programs, a theory of the *learnable in economics* needs explicit behavioral rules: recursive rules suitable for computational experimentation. *Fictitious play*, and its variants, are examples of explicit adaptive rules. They are examples of more general classes of rules. First, as a recursive estimation procedures, they are examples of *stochastic approximation algorithms*.

Second, as *recursive meta-rules* that choose among simple rules or actions, they are examples of *artificially intelligent agents*. I finish this section with a brief discussion of these general classes of explicit behavioral rules (see, Sargent 1993, for other economic examples).

Adaptive learning as a stochastic approximation algorithm

The frequentalist methods of computing the sample mean or the empirical distribution, are simple examples of stochastic approximation methods. The following example (a variant of *fictitious play*) illustrates the strength of these algorithms (see, Arthur 1993, Easley and Rustichini 1995 and Marimon and McGrattan 1995).

Example 2 The agent only knows his set of actions $A = (a^1, \ldots, a^n)$ and only observes his realized payoffs, i.e., does not observe $\{x_t\}$. We assume that $\{x_t\}$ is (strictly) stationary and that payoffs are strictly positive and bounded ($\bar{u} \geq u(a, x) \geq \underline{u} > 0$). A Bayesian learner would have to form a prior over possible payoff matrices, in addition to the prior on μ, etc. The adaptive agent, instead defines his behavioral strategy by randomly choosing among actions according to some measure of relative strength. Let S_t^k be the strength attached to action a^k in period t and $S = \sum_{k=1}^{n} S_t^k$. Fix $S_1 \in \mathcal{R}_{++}(i.e., S_1^k > 0 \forall k)$ (this is the extent in which our adaptive agent has a prior). Then strengths evolve according to the rule

$$S_{t+1}^k = S_t^k + \begin{cases} u(a^k, x_t) & \text{if } a^k \text{ is used at } t \\ 0 & \text{otherwise} \end{cases}$$

then actions are randomly selected according to $\alpha_t(a^k) = \frac{S_t^k}{S_t}$ To express this behavioral rule in stochastic approximation form, let $z_t^k \equiv \alpha_t(a^k)$, and $v(a^k, x_t)$ be a random variable that takes value 0 with probability $(1 - z_t^k)$ and $u(a^k, x)$ with probability $z_t^k \nu_t(x)$. Then, we have that mixed actions evolve according to

$$z_{t+1}^k = z_t^k + \frac{1}{S_{t+1}}[v(a^k, x_t) - (S_{t+1} - S_t)z_t^k]$$

This is an example of a stochastic process of the form

$$z_{t+1} = z_t + \alpha_t[h(a_t, z_t)] \tag{1}$$

where z is the object to be learned and $h(a_t, z_t)$ is a random variable, which depends on the parameter z and, possibly, on the action a taken by the agent. If "the gain sequence" (for simplification taken to be deterministic) $\{\gamma_t\}$ satisfies a "decay condition" ($\sum_{t=1}^{\infty} \alpha_t = \infty$, $\sum_{t=1}^{\infty} \alpha_t^p < \infty$ for some $p > 1$), and z_t is a stationary process, then the steady states of the system of stochastic difference equations (1) are given by the solutions of $E[h(a, z_t)] = 0$.[7] A specially useful result is that the corresponding asymptotic behavior of $\{z_t\}$ are characterized by the solutions of the *ordinary difference equation*

$$\frac{dz}{d\tau}(\tau) = E[h(a, z(\tau))] \tag{2}$$

Returning to our Example 2, the corresponding ODE is,

$$\frac{dz^k}{d\tau}(\tau) = (1 - z^k(t))z^k(t)\left[E[u(a^k, x)] - E\sum_{j\neq k} z^j(t)[u(a^j, x)]\right]$$

This is a form of the *replicator dynamics* and the stable stationary point of the ODE correspond to assigning all the weight to the actions with highest expected payoff. That is, with very little information, and fairly simple behavioral rules, the agent learns to choose actions that maximize the objective expected utility of the one-period decision problem. Although, this may take a long time! As Arthur (1993) shows comparing his artificial agents with the results with human subjects reported by Busch and Mosteller (1955), the rates of convergence, and convergence itself, is very sensitive to the degree of disparity of payoffs.

As we have seen, a frequentalist who computes the empirical distribution of a process $\{x_t\}$, $x_t \in \{0, 1\}$ can completely miss a well defined cycling pattern. This is a general problem in defining learning processes as Stochastic Approximation Algorithms (or any type of algorithm). The underlying environment prescribes which algorithm is better suited for learning (in terms of a *consistency condition*). For example, in the cycling case would be enough to properly condition the algorithm (e.g., have a forecast for odd periods and another for even periods). Similarly, one may want to adjust the "gain sequence" to the environment (see, for example, Benveniste *et al.* 1990). It is well understood that there is a tradeoff between *tracking* and *accuracy*: a decrease of the gain α reduces variance of the error (allows for more accuracy), but, at the same time, may increase the bias of the forecast since it tends to underestimate the possible drift or nonstationarity of the underlying process (may not "track" well the process). This, which is a standard problem in statistical inference, shows

[7] See, Benveniste *et al.* (1990) and Ljung *et al.* (1992) for an account of Stochastic Approximation Theory. In particular, when the process has a fixed α, then it is a Markovian process, and, with enough perturbations (experimentations), the process is ergodic. That is, there is an ergodic distribution $m(\alpha)$ characterizing the asymptotic behavior of the process $\{z_t\}$. We can then take the limit of these distributions $m(\alpha)$ as $\alpha \to 0$ at the right rate. This is the approach taken by *simulated annealing* methods which have been widely applied in *evolutionary models* (see Kandori's contribution to this volume). This shows that, as long as adaptation and evolution are governed by similar rules, differences in asymptotic behavior can only be attributed to different form in which perturbations and limits of α interact.

the type of tradeoffs that are common in learning dynamics and that a learning theory must account for.

Artificially intelligent agents

Example 2 of an adaptive decision maker selecting rules according to their relative strength, can be seen as a very simplified brain of an *artificially intelligent agent* using Holland's *classifier systems* (CS) (see, for example, Holland 1995). Classifier systems, based on genetic algorithms, were designed as an "all purpose adaptive algorithm." That is, in contrast with *experts systems* specially designed for a task or problem, and in contrast with learning machines, such as *neural networks* (NN), that even if there were not designed for a particular task needed external training with an environment before they could appropriately perform, CS could –in principle– adapt to alternative tasks and learn from the beginning. In fact, some of the CS principles –mainly, selection of rules by genetic algorithm operations– have been incorporated in *neural network* architecture. *Recurrent neural networks* (RNN) share many things in common with Classifier Systems.

It is not possible to describe here these different classes of recursive algorithms (see, for example, Holland 1995, Sargent 1993, Cho and Sargent 1995 and White 1994), but only to point out some features that characterize their "higher adaptability:" *i*) Decisions, as languages, are decomposed into basic components ("if ... then" statements, neurons, etc.); *ii)* the "builiding blocks" structure means that there is "parallel learning" of different components (*schemata* in CS); *iii)* an *internal accounting system* allows aggregation and internal transmission of information, which may result in building up "internal models" (e.g., some forms of endogenous pattern recognition); *iv)* even if certain "tuning to the environment" is needed, stochastic search can result in "discoveries," that is, "unforeseen contingencies" may be accounted for. These features of *artificial intelligent agents* allows them to learn through relatively complex environments generating interesting –non linear– feedback. For this reason, an economy with *feedback* populated by *artificial intelligent agents* is labeled a *Complex Adaptive Systems*. At the same time, such sophistication makes their analytical study more difficult and computational simulations are an integral part of the development of these models.

4 Learning in games

From the perspective of learning theory, game theory provides a set of models where *feedback* effects may be specially perverse, except that they are governed by the norms of behavior that characterize individual learning. I discuss some of the recent contributions to the theory of learning in games (see also Fudenberg and Levine, 1995b). Broadly speaking, I proceed from weaker behavioral assumptions and solutions to stronger ones.

Monotonicity and rationalizable solutions

As we have seen, adaptive behavior can be characterized by a monotonicity condition. Milgrom and Roberts (1991) condition of only playing strategies which in relation to a finite past history of play are not strictly dominated, is a fairly weak monotonicity condition. Nevertheless, as they have shown, it is enough to guarantee that, in the long run only *rationalizable strategies* will be played. That is only strategies that survive an iterated process of elimination of strictly dominated strategies. This result has an immediate corollary, if agents behavior satisfies their monotonicity condition and play converges to a single strategy profile, σ^*, then it must be a pure strategy Nash equilibrium.

If an agent is not fully aware of his feedback on the environment, or not all actions are observable and experimentation is costly enough, then he may converge to play an action which does not maximize objective expected utility, but maximizes his subjective expected utility. In a sense, Milgrom and Roberts result is "too strong" since if their monotonicity condition is expressed in terms of subjective beliefs, then only *subjective rationalizability* may be achieved.

4.1 Subjective equilibrium notions

A process of iterated elimination of strategies which are not a best reply to subjective beliefs may converge to a situation where every player i, adopts a behavioral strategy σ_i that is optimal, given his subjective beliefs μ_i of other agents' play (of the environment), and his beliefs are not contradicted. That is, μ_i coincides with ν_i, where ν_i is the objective marginal distribution induced on player i's play paths (i.e., $\nu = (\sigma_1 \times \cdots \times \sigma_I)$). In other words, agent i's beliefs may not coincide with the objective distribution outside his play paths, as it is required in a Nash equilibrium.

Subjective equilibrium notions in economics go back to Hayek (1945). Hahn (1973) defined a *conjectural equilibrium* as a situation where agents' subjective beliefs (agents' *models*) about the economy were self-fulfilled and agents had no incentive to change their policies. This notion has been formalized, in the context of games, by Battigalli (see Battigalli *et al* 1994). In the context of games, more recently, Fudenberg and Levine (1993) and Fudenberg and Kreps (1995) have proposed the notion of *self-confirming equilibrium* and –subsequently and independently– Kalai and Lehrer (1993b, 1995) the notion of *subjective equilibrium*. In a *subjective equilibrium* agents may misperceive their *feedback* on the social outcome. For example, in an oligopolistic market, with a large number of firms, a competitive equilibrium would correspond to a *subjective equilibrium* where firms do not take into account their individual effect on prices (see also Brousseau and Kirman, 1995).

Even if agents observe the actions of their opponents and know how to *best reply* to them, non-objective beliefs can be *self-confirmed*. For example, in an extensive game with limited experimentation, two players can have inconsistent beliefs about a third player and never realize this since the information that would show such inconsistency is never revealed (see Fudenberg and Kreps 1995).

Kalai and Lehrer (1993, 1995), Fudenberg and Levine (1993) and Fudenberg and Kreps (1995 and 1994) provide conditions under which such subjective equilibria are, in fact, Nash equilibria. For instance, in a one stage game *self-confirming* equilibria are Nash equilibria since all information sets are reached. The convergence of bayesian learning in repeated games to Nash equilibria has also been established by Nyarko (see Nyarko 1994) and, recently, Sandroni (1995) has extended these results to finitely repeated games.

Subjectivism, marginal best replies and correlated equilibria

As seen in Section 3, if the distribution of outcomes, x_t, is correlated with the agent's actions, a_t, then learning can converge to different subjective solutions, depending on how the agent conditions his forecasts. In a repeated game, correlations are likely to occur since agents respond to each other past actions. The question is whether, such correlations result in a distribution of play which is a –possibly, subjective– *correlated equilibrium*. In a *subjective correlated equilibrium*, agents' subjective beliefs may contemplate correlated play by other players and these beliefs can be different while playing different pure strategies (non unitary beliefs). That is, agents best replies can be subjective *conditional* best replies.

There are several reasons why correlated equilibrium, and not Nash equilibrium, seems to emerge as the central reference equilibrium concept in learning theory. First, as it has been said, correlations naturally arise in repeated play. Second, as Aumann (1987) has argued, Bayesian rationality in a multiagent context (without common priors) is equivalent to a *subjective correlated equilibrium*. Third, as Hart and Schmeidler (1989) have shown, in a correlated equilibrium every agent can be viewed as playing a *zero sum game* against the opponents. This has two related implications. The first is that the existence of correlated equilibria can be derived using standard separation arguments (i.e., without the recall to a fixed point argument); the second is that, as we will see, the simple zero sum game, such as that between a player and the environment, it has been known to be learnable. The question is whether players condition properly, to guarantee that the joint solution of this learning process results in a correlated equilibrium.

In the context of repeated games, Kalai and Lehrer (1995) and, independently, Nyarko (1994) have shown that bayesian learning leads, eventually, to approximate subjective correlated equilibria when beliefs satisfy an absolute continuity condition. They also provide conditions that guarantee that such convergence is to an objective correlated equilibrium. That is, *strong merging* (or *weak merging*) of beliefs translate into objective (subjective) equilibrium outcomes. But, as we have seen, these consistency conditions require strong absolute continuity assumptions. The question is whether weaker consistency conditions result in some form of equilibrium outcome.

Recently, Foster and Vohra (1995), in the context of a two players repeated stage game (with myopic payoffs), have proved a remarkable result. They show that if both players behave as *well calibrated forecasters*, then the asymptotic distribution of forecasts, or *calibrated beliefs* is a correlated equilibrium, and that –for almost every game– any correlated equilibrium is learnable by well calibrated players. Unfortunately, their existence proof of well calibrated players, based in the zero-sum characterization of correlated equilibria, is not constructive.

As it has been seen in Example 1, if the actions of players are correlated, players may correctly perceive the marginal distribution, but this does not mean that the product of these marginals distributions is the joint distribution (a basic and well known fact about joint distributions). Example 1 translates into the famous Shapley example. That is,

Example 1' (Shapley)

	b^1	b^2	b^3
a^1	0,0	2,1	1,2
a^2	1,2	0,0	2,1
a^3	2,0	1,2	0,0

Fudenberg and Levine (1995a) show that the long run play of *universally consistent* behavioral rules are (infinitely often) correlated distributions that have a marginal best response property. That is. agents –subjectively– best reply to the empirical marginal distributions. For instance. in Example 1'. it is known that fictitious play cycles along the strategy vectors with positive payoffs and that these cycling patterns slow down. The asymptotic distributions have the *marginal best reply* property, but the joint distribution may not be a correlated equilibrium.

The convergence of learning processes to subjective non Nash. or correlated, equilibria (or, just *marginal best reply* solutions) is directly related to the *complexity* of the game and the potential benefits that players may get from reaching "objective" beliefs. Three possible sources of *subjectivism* are: *i)* the lack of enough *experimentation*; *ii)* the existence of correlations that may occur through the game, and *iii)* the "misspecification" of the model. While the experimentation costs are easy to measure (for example, "enough experimentation" in an extensive form game means that all nodes are "tested infinitely often"; Fudenberg and Kreps (1994, 1995), Hendon *et al.* (1995)), it is less obvious how to measure the "complexity" of correlated patterns or of possible model misspecifications. Such "complexity" measures could help us to understand the gap between subjective and objective equilibrium outcomes.

4.2 The learnability of Nash equilibria

As it has been seen, under relatively weak monotonicity assumptions, *if the learning process converges, it will converge* to a Nash equilibrium. Also, under the appropriate absolute continuity (and independence) assumptions bayesian learning converges to Nash equilibrium in repeated games (see also footnote 5). Furthermore, it has been known for some time that, for some classes of games, convergence to Nash equilibria is guaranteed for certain classes of learning algorithms. For instance, in addition to Milgrom and Roberts' adaptive play in (strict) dominance solvable games, convergence to Nash has been shown for fictitious play in zero-sum games (Robinson 1951 and Brown 1951) and in 2×2 games (with a prespecified "breaking ties" rule; Miyasawa, 1961) . These results (as well as Foster and Vohra' (1995) convergence to correlated equilibrium) are in terms of *beliefs*. It is well known that convergence of beliefs does not imply convergence of the empirical distributions of play (see, for example, Jordan 1993). As we have also seen, Fudenberg and Levine (1995a) show that if players' behavioral rules satisfy their asymptotic *universal consistency* condition, then players can only be playing infinitely often strategies that satisfy the *marginal best reply* property.

Fudenberg and Kreps (1993) provide a –sharper– global convergence result for fictitious play in 2×2 games (with a unique totally mixed Nash equilibrium) by considering an augmented game in which payoffs are perturbed. This work has been extended by Benaïm and Hirsch (1994) and, recently, Kaniowski and Young (1995) have shown that, in the same context of 2×2 games, if both agents learning rules are of fictitious play type –with perturbations– then the learning process converges almost surely to a *stable Nash equilibrium*, either mixed or pure (in fact, arbitrarily close to a *dynamically stable equilibiurm*). For example, in the *Battle of the Sexes* game the learning process converges to one of the two pure strategy equilibria with probability one. All these global convergence results for 2×2 games use stochastic approximation techniques and, as Kaniowski and Young point out, they can also be interpreted as the output of an evolutionary process with two populations where agents are fictitious players with random sampling (as in Young 1993). These results, however, do not generalize to more than two actions games since Shapley's example (Example 1') remains a counterexample to global convergence of play.

With respect to *local stability* results, it has also been known that *strict Nash equilibria* (that is, Nash equilibria with single valued best replies) are *learnable equilibria* in the sense that are asymptotically locally stable to some process of adaptive learning –say, satisfying Milgrom and Roberts monotonicity condition (see, for example, Fudenberg and Kreps, 1993). As we have seen, Foster and Vohra (1995) also provide a local stability result for correlated equilibrium (i.e., any correlated equilibrium can be *forecasted* by calibrated forecasters. Fudenberg and Kreps (1993) provide a local stability result for Nash equilibrium. They consider "convergence of behavioral strategies" instead of the previously used "convergence of beliefs (or empirical frequencies)." More specifically, they consider a strategy profile, σ^* *locally stable* if the following *weak asymptotic consistency* conditions are satisfied: *i)* the beliefs (assessments) of the all the players converge to the empirical distribution of play; *ii)* if, for every player, the mixed strategy played at t, $\sigma_t^i(h_t)$ (not necessarily all the pure strategies in its support) is within an $\epsilon_t > 0$ of being the best reply, where $\epsilon_t \to 0$, and *iii)* for every $\delta > 0$, there is a t such that $\text{Prob}\{h_t | \lim_{t' \to \infty} \sigma_t'(h_t') = \sigma^*\} > 1 - \delta$. They show that every Nash equilibrium is locally stable and that any non Nash equilibrium strategy profile is not locally stable (in particular, correlated equilibrium that are not Nash equilibrium are not locally stable).

Fudenberg and Kreps (1993) results can be viewed as a *learning justification* of Nash equilibria. Loosely speaking, in repeated play finite games, "every Nash equilibrium is learnable and the only learnable strategy profiles are Nash equilibria." As I remarked, however, their asymptotic consistency condition –that is, their *learnability* concept– is very weak, and it must be since, for example, it prescribes that, Nash equilibria where agents play *weakly dominated* strategies, or the mixed Nash equilibrium in the *Battle of the Sexes*, are *learnable*. However, for most learning algorithms that have any chance to learn in relatively complicated environments such equilibria do not emerge as the outcome of the learning process. For example, as we have just seen, in 2×2 games, Kaniowski and Young' "perturbated learners" learn such –non *dynamically stable*– equilibria with zero probability. Of course, the same remark applies to the local belief stability concept of Foster and Vohra (1995).

4.3 The "support" of attraction

As we have seen, for a large class of learning algorithms, *when play converges* it must be to a Nash equilibrium and, when agents learn with experimentation (or other forms of perturbations), it must converge to a refinement of Nash equilibrium. We have also seen some global convergence results for 2×2 games. A characterization for general strategic form games is available if only the "support of play" is considered. This is the approach taken by Marimon and McGrattan (1995) and by Hurkens (1994). The basic idea is that if a learning algorithm satisfies a monotonicity condition, then "learning processes move in the direction of the best reply map." Furthermore, if agents *experiment*, individual best replies are perturbated. As a result, the support of the attractors of large classes of learning rules can be characterized.

Hurkens (1994) considers a learning model where players are members of classes and at each period of time a player of a class is randomly selected to play (as in Young 1993). These players have fixed finite memory and satisfy Milgrom and Roberts monotonicity condition. Given the finiteness assumptions the learning process has a Markovian structure. He then shows that, if memory is long enough, play eventually sets into a *curb set.* A *curb set* is one that it is closed under best replies and a curb* set is one that it is closed under undominated best replies.[8] Curb sets may include Nash equilibria with dominated strategies. However, Hurkens extends the same result to minimal curb* sets, by strengthening the monotonicity condition.

Marimon and McGrattan (1995) (see Section 3) consider variations of *persistent retracts.* A persistent retract is minimal with respect to the property of being closed under perturbed best replies[8]. A *robust equilibrium* is a singleton persistent retract. For example, in a *matching pennies* game all the strategies define a persistent retract, but in the *battle of the sexes*, only the pure strategy Nash equilibria are robust equilibria (and curb* sets; but there are robust equilibria which are not curb* sets). They also extend these concepts to correlated equilibria, to capture correlated cycles, such as the one of Shapley's example. Long run dynamics of these learning algorithms either converge to a *robust equilibrium* or *the support* of the asymptotic play is a *robust correlated cycle.* In 2×2 games, Posh (1995) shows how "perturbated learning" can result in well defined cycles. That is, a cycle (of the joint distribution of play) with support in a *robust correlated cycle*, as in the *matching pennies game.*

As we see, when agents use learning algorithms satisfying basic behavioral assumptions (monotonicity, inertia and experimentation) then play converges –if it does– to a refinement of Nash equilibria (robust equilibria). The problem arises when agents' feedback translate into correlated cycling behavior. Such patterns may not result in well defined cycles and then it is not clear how agents condition their actions on the observed play. This problem does not arise in –say, *supermodular*– games where convergence is monotone. In such cases, *pattern recognition* of the asymptotic play does not disrupt the convergence process. Recently, Sonsino (1995) has extended adaptive behavior to allow for certain degree of pattern recognition. He shows how agents may learn and sustain a cycle, such as "take turns in playing pure strategy equilibria" in the *Battle of the sexes*, but he also shows that unbounded memory is needed to *pattern recognize and converge* to a mixed strategy profile.

[8]Formally, let $PB_i(\sigma_{-i})$ denote the set of non dominated pure strategies which are i's best replies to σ_{-i}, and $PB(\sigma)$ the corresponding joint best reply correspondence. A set $D = \prod_I D_i$ is a *curb set* if $PB(\prod_I \Delta(D_i)) \subset D$. Such a set is called a *minimal curb set* if it does not properly contain a curb set. Curb* sets are similarly defined by considering only undominated best replies. Furthermore, D is a persistent retract if it is minimal with respect to the property: $PB(\prod_I \Delta(D_i^\epsilon)) \subset D$, where D_i^ϵ is an open neighborhood of D_i.

4.4 Neuralnets rediscover the Folk Theorem

There is one more lesson to learn from the recent literature on learning in games. In all the above discussion, learning agents did not have a final goal. The learning problem, however, can be greatly simplified if agents "know what they want." For example, at any point in time, an agent can *discriminate* whether his "aspiration" level has been achieved or not. Cho (1994, 1995) develops a model in which two artificial intelligent agents (AIA), with *neural network* capabilities, play repeated games. He recovers the *Folk Theorem*. That is, any individually rational payoff can be achieved as the long-run payoff of the AIA game (with or without discount). He also extends this result (without discounting) to games with imperfect monitoring.

5 Learning dynamic rational expectations equilibria

In competitive environments an individual agent does not affect the social outcome and, therefore, he does not create strange correlations out of his optimal actions and mistakes. Furthermore, most learning models only address the problem of forecasting a public outcome, such as prices. These features simplify the learning problem. In competitive environments, however, we typically study intertemporal problems with continuous action sets. This feature complicates again the learning problem.

5.1 Learnable Rational Expectations Equilibria

Most macroeconomic models are examples of the general competitive model with intertemporal feedback, where there is a continuum of agents –say, $I = [0, 1]$– and a public outcome evolves according to $x_t = \phi(g(a_t, a_{t-1}), x_{t-1}, s_t)$ (see Section 2). Furthermore, in many models, $a_{i,t} = BR_i(x^{e_i}_{t+1})$, where $x^{e_i}_{t+1}$ is agent i's expectation of x_{t+1}, at t. In this case, there is a well defined mapping $\Gamma : X_{\infty}^{[0,1]} \to X_{\infty}$, such that $\{x_t\} = \Gamma\left(\{x^e_{i,t}\}_{i\in I}\right)$ Rational expectations equilibria are fixed points of Γ, in the sense that $\{x^*_t\} = \Gamma\left(\{x^*_t\}_{i\in I}\right)$. That is, agents *agree in the right forecast*, even when this means agreeing on a non stationary paths. Most existing learning models make three important shortcuts. First, they assume a *representative agent*; second, the model is reduced to a "temporary equilibrium" model of the form $x_t = \gamma(x^e_{t+1})$, where $\gamma = \phi \cdot g \cdot BR$

(see, for example, Grandmont, 1994) and, third, the learnability question is limited to the fixed points of the temporary equilibrium map, $x^* = \gamma(x^*)$, that is to the stationary fixed points of $\Gamma(\cdot)$. In addition to a a fixed point there can be the stationary fixed cycles, which are stationary rational expectations equilibria (SREE).[9] Of these "short cuts", probably the most "distorting" one is the *representative agent.* It violates a basic *learnability principle*: agents may imitate others (social learning), but "agents learn independently." In a REE different agents must share the same beliefs about public outcomes, but learning is supposed to study how such a coordination of beliefs takes place. The *representative agent* introduces *feedback* effects (e.g., his mistakes are not smoothed out by others) that are not present in historical competitive economies.

Subjectivist and rational belief equilibria

As in games, self-fulfilling expectations may not imply rational expectations. In fact, as I mentioned in Section 3, some "subjectivist" concepts, such as Kurz's "rational beliefs," have been postulated as a –weaker than REE– solution to dynamic competitive models. Nevertheless, some of the sources for "subjectivism," discussed in Section 4, disappear in competitive environments. In particular, the agent should not be mislead about his effects on public outcomes, which are nil, nor be concerned about individual correlations. This, however, does not mean that *subjective competitive equilibria* can not exist. For example, in Kurz's (1994a and 1994b) *rational belief* equilibria, it is only required that all agents beliefs must be "compatible with the data." That is, they must satisfy the absolute continuity property of only assigning zero probability to unobserved events. In his examples, non REE *rational belief equilibria* exist by having agents disagreeing with the objective REE distribution –ν– in its nonstationary component. In other words, agents agree in "pattern recognizing" stationary public events, but fail to grasp the complexity of the underlying –possibly, nonstationary– economy.

One way to get disagreement between REE and *rational belief* equilibria is by having some nonstationarity built into the environment –say, an important tail event. But then, it is not clear that there should be any loss of efficiency from not reaching REE. In fact, continuity of infinitely lived agents' preferences implies that tail events should have no effect on prices, otherwise (without continuity) the same existence of REE is at stake. This suggests an interesting relationship between *learnability of equilibria* and *market completeness.*

[9] A fixed cycle –say, of period k– is a vector $(x_1^*, \ldots, x_k^*)$ such that, $x_1^* = \gamma(x_2^*), \ldots, x_{k-1}^* = \gamma(x_k^*)$, $x_k^* = \gamma(x_1^*)$.

Araujo and Sandroni (1994) show, that, in an exchange economy with *dynamically complete markets*, bayesian learners *always* converge to rational expectations; a remarkable result. As it is standard in this literature, they appeal to Blackwell-Dubins theorem (see, Section 3) and the facts that agents' preferences are continuous and all public information –that is, prices– is endogenous. Nevertheless, it has been known, since, for example, the work of Townsend (1983) that in stationary economies –with incomplete markets– the learning process may result in generating non stationary prices, which presumably could result in rational belief equilibria which are not REE.

There is another source of *subjectivism*, which has already been mentioned, of which we have a number of examples: *model misspecification*. An example, can be that all agents use a linear forecasting model when REE prices have non linearities and the economy converges to a "linear rational belief equilibrium." Most of these models, however, maintain a *representative agent* assumption or that agents share similar learning rules (models). Both assumptions impose a high degree of "coordinated subjectivism".

E-stability

The *learnability* of the SREE with fixed learning rules has been well studied. In particular, the following classes of learning rules have been studied: *i)* finite memory forecasting rules $x^e_{t+1} = f(x_{t-1}, \ldots, x_{t-m})$, where the class $f(\cdot)$ satisfies certain basic assumptions, such as the ability to forecast stationary paths (i.e., $\bar{x} = f(\bar{x}, \ldots, \bar{x})$ (see, for example, Grandmont and Laroque, 1991 and Grandmont 1994); *ii)* Cagan's adaptive learning rules: $x^e_{t+1} = x^e_t + \alpha(x_{t-1} - x^e_t)$, with a *constant gain* $\alpha \in (0, 1)$ (see, for example, Guesnerie and Woodford 1991, and Evans and Honkapohja, 1993a, 1993b) and, *iii)* learning rules with an stochastic approximation representation of the form $x^e_{t+1} = x^e_t + \alpha_t(x_{t-1} - x^e_t)$, such as recursive least square rules (the *gain sequence* satisfies the "decay condition," mentioned in Section 3.4, $\{\alpha_t\}$, $\alpha_t \in (0, 1)$ $\sum_{t=0}^{\infty} \alpha_t = +\infty$ and $\sum_{t=0}^{\infty} \alpha_t^p < +\infty$, $p > 1$)[10] (see Bray 1983, Marcet and Sargent 1989a, 1989b,

[10]Evans and Honkapohja (1995a), call an E-stable equilibrium that is stable to overparametrizations of the learning rule *strongly* E-stable, otherwise they call it *weakly* E-stable (Overaparmetrizations examples are: in a k cycle, consider $k \cdot n$ cycles; in an ARMA(k, n) model, consider increasing (k, n), etc.) In their examples where a "representative agent" learns an indeterminate equilibrium with an ARMA rule, such equilibrium is only *weakly* E-stable. In Duffy's example the "representative agent" only reacts to current events and his behavior immediately *feedback* into the economy: his forecasting rule is $\pi^e_{t+1} = [1 + (b-1)\pi_t]/b$ where $b > 1$ parameterizes the continuum of REE!

Woodford 1990, Evans and Honkapohja 1994, 1995a, 1995b, 1995c, among others). Some of this work studies the the relation between the "determinacy" (local uniqueness) of the the SREE, or of the corresponding REE sunspot cycle, and the *learnability* of the corresponding equilibrium using adaptive learning rules (see, for example, Guesnerie and Woodford 1991).

By using specific classes of forecasting rules the concept of *learnable equilibrium* is well defined. The learnability of a SREE, $x^* = \gamma(x^*)$, is given by the interaction of the γ map and the forecasting map (assuming a representative agent). That is, for finite memory rules *(i)* or for Cagan rules *(ii)*, one can use standard stability theory (i.e., characterizing the asymptotic stability by linearization around the steady state, etc.), and for rules of the form *(iii)*, stochastic approximation theory provides the necessary characterization. In particular, if the corresponding ODE (recall Section 3), $\frac{dx}{d\tau} = E\left[\gamma(x(\tau)) - x(\tau)\right]$, is locally asymptotically stable at x^*, then the SREE is called *E-stable*. In models with constant gain, *(ii)*, the asymptotic stability properties also depend on α being small enough. That is, if agents place enough weight on current events (α close to one) the SREE may not be learnable, even if the ODE is locally asymptotically stable and, therefore, *learnable* with rules with decreasing gain. These conditions generalize to "determinate" SREE cycles.

Woodford (1990) was the first to show that sunspot cycles could be *learnable* equilibria. In a sense, this result created skepticism on the idea that learning can help to "select" among REE. In fact, there are examples where adaptive learning agents' beliefs converge to *indeterminate* REE (Evans and Honkapohja, 1994, Duffy, 1994). Does this mean that *any REE is learnable*?

As we have seen in relation to Fudenberg and Kreps (1993)' results (*any Nash equilibrium is learnable*), the answer to the above question depends on the definition that we use of learnability. For instance, a closer look at the "convergence to indeterminate equilibria" examples reveals that a very weak concept of learnability is being used: they are either not robust to overparametrizations of the learning rule or require –usually, a representative– "overreactive" agent with important *feedback* effects [34]. Similarly, a closer look at the "convergence to cycles" models shows, that learnability of –say, a k– cycle requires that agents follow rules of the form $x^e_{t+k} = x^e_t + \alpha_t(x_t - x^e_t)$. In other words, the cycle must have been "pattern recognized" before it can be learned. For example, Evans, Honkapohja and Sargent (1993) study a model of REE sunspots k-cycles populated by a fraction of perfect foresight agents and a fraction of agents that do not recognize cycles (e.g., they treat price fluctuations like noise). They show that, if the fraction of non "perfect foresighters" is large enough, all the cycles of period $k \geq 2$ disappear. In other words, cycles may exist only if enough agents are simultaneously "tuned" to them.

Ultimately, however, *learnability* is an empirical matter. The experimental evidence shows well defined regularities. In particular, non stationary REE and non-stable SREE have never been observed (see, for example, Marimon and Sunder 1995). Similarly, with respect to *sunspot* cycles, the experimental evidence is revealing: purely belief driven cycles have never been observed, however, when agents share a common experience with a real cycle, perfectly correlated with a sunspot cycle, the E-stable sunspot cycle may persist after the real shock, driving the real cycle, disappears (Marimon, Spear and Sunder, 1993).

5.2 Global stability and local instability?

The macro experimental evidence is consistent with the theoretical (and experimental) results in games. When learning is not "tuned to an equilibrium" but agents must experiment and coordinate based on past experience, then many REE dynamics may de disregarded. In particular, *dynamic stability* criteria for large classes of learning rules, provide a good global characterization of observed paths (e.g., inflation paths tend to cluster around a *E-stable* SREE; as in the *Battle of the Sexes* agents converge to play pure strategy equilibria). Nevertheless, experimental evidence also shows that local dynamics, around a E-stable SREE, may be fairly complicated and not well predicted by the local stability properties of simple learning algorithms. This phenomena is particularly true when there are fluctuations around the steady state (e.g., there are complex roots). Subjects tend to react to such fluctuations (see, Marimon and Sunder 1995), as in games players tend to react to cycling patterns. As a result, the local convergence properties depend on how agents "pattern recognize" such high frequency fluctuations and how their behavior "feedback" into the economy. This dichotomy between the predictability of *low frequency* data and the difficulty to predict *high frequency* data is, in turn, consistent with macroeconomic and financial data and, therefore, learning models can not only help to explain it, but also help to design economic policies.

Stabilization policies with learning

Agents expectations always play an important role in the successful implementation of economic policies, and this is particularly true of stabilization policies. Unfortunately, most macroeconomic models have either made an *ad-hoc* treatment of expectations precluding agents from anticipating the effects of policies or, as a healthy reaction to such models, assumed *rational expectations* precluding agents from *learning* the possible consequences of policies. Recently learning models have started to fill this gap.

Evans, Honkapohja and Marimon (1995) show how introducing a (credible) bound on the amount of deficit to GDP ratio that can be financed by seignorage can be a powerful instrument for stabilization policy, resulting in *global stability*, in an Overlapping Generations model with learning by heterogeneous agents. Such a policy, however, only makes more acute the problem of indeterminacy of REE equilibria. In other words, a policy which is often proposed (for example, in the Maastricht Treaty) is an advisable policy only when learning is taken into account. They also provide experimental data showing how subjects learn better, than simple adaptive learning algorithms, to foresee the effects of policies, but in their model this results in making the stabilization polices even more effective.

Similarly, Marcet and Nicolini (1995) show that recurrent inflationary episodes interrupted by *pegging exchange rates policies*, such as the ones experienced in several South American countries in the last twenty years, can be explained by an adaptive learning model, where agents endogenously change the weight place in current events (the "tracking" parameter α discussed in Section 3). Again, incorporating learning not only allows to explain some macro data, but also provides a different evaluation of an often proposed stabilization policy.

Volatility, persistent correlations and complicated dynamics around the steady state

Learning models are increasingly used to explain patterns which do not seem to satisfy REE restrictions. From long run development trends (Arifovic *et al.* 1995) to short term trade and price volatility in financial and exchange markets (see, for example, Hussman, 1992, Brock and LeBaron, 1995, de Fontnouvelle, 1995). Some of these "high frequency" models exploit the fact that around the steady state variability may be high, and correlations persistent, in an adaptive learning model. In fact, it is well understood that the *speed of convergence* near equilibrium depends on the sophistication of learning rules (again, the tradeoff between *tracking* and *accuracy* discussed in Section 3).

A claim, often made, is that final convergence to a competitive equilibrium can be difficult to achieve since "near equilibrium competitive forces may soften." This claim goes back to the problem studied by Grossman and Stiglitz. among others, of who has an incentive to purchase information if equilibrium prices are fully revealing. Brock and Hommes (1995) have, recently, developed an interesting model that addresses this issue. A simple example of their economies goes as follows. Agents can either forecast using their, relatively myopic, forecasting rules or purchase the services of a more sophisticated forecaster. All forecasting rules are of the stochastic approximation type, however the corresponding ODE is unstable if agents use their "free" rules, while it is stable if they use the –more accurate– costly rule. Far from the steady state most agents buy the forecasting services, pushing the economy towards the steady state, however, near the steady state it is not worth it to pay for forecasting services. This implies local instability (as in Grandmont and Laroque 1991) and, at the same time, non divergence from a neighborhood of the steady state. As a result adaptive paths may be fairly complicated; they prove the existence of *homoclinic orbits* and of *strange chaotic attractors* for a class of predictor rules.

In the experimental lab, we have observed similar "complicated dynamics" around the steady state. There are, however, two distinct underlying phenomena. One is that, as in Brock and Hommes, agents' loose the incentives to sharpen their decisions rules as to finally converge to equilibrium, and learning paths only "cluster around" the E-stable SREE (see Marimon and Sunder, 1993). The other is that around the SREE there may be fluctuations and even agents that can appropriately guess the SREE (e.g., is the announced policy) try to follow the fluctuations in order to capture possible short-run rents (Marimon and Sunder, 1995). In other words, a trader in the stock market may not be interested in "fundamentals."

5.3 Social planners as AIA

As in games, there are macroeconomic problems where it seems appropriate to think that the learner has the "right pattern and equilibrium in mind." These are planner's problems, in which the planner (or principal) is concerned that the agent follows a contract as planned. The resulting optimal strategies in these type of problems are fairly complicated history dependent contracts. Nevertheless, building on the work of Cho on *neural nets* playing repeated games (Section 4), Cho and Sargent (1995), show how relatively sophisticated planner's problems (e.g., an optimal capital accumulation problem with private information) can be implemented by *artificially intelligent planners* with neural nets capabilities.

6 Concluding remark

As I mentioned in the Introduction, to build a *theory of the learnable in economics* goes way beyond some standard justifications for the study of learning in economics. According to this view, I have placed the emphasis in this abridged survey in several lines of inquiry: *i)* the characterization of *adaptive learning*, of broad classes of rules, based on behavioral assumptions and resulting in learning process which are "well behaved" (i.e., satisfy certain *consistency condition*) in well understood environments; *ii)* parallel to the definition of alternative *consistency conditions*, the definition of *subjective forms of equilibria*; *iii)* the characterization of the asymptotic properties of certain classes of learning algorithms and their ability to *select among equilibria*; *iv)* the ability of learning models *to explain observed data* which is not properly accounted for by existing equilibrium theories, and *v)* the use of learning models as *normative models* to help the design of economic policy and of political and economic institutions.

Although –possibly, for sociological reasons– research in *learning in games* and in *learning in macro* have been conducted fairly separately, as sometimes has been the case between *learning theory* and *experimental evidence*, I have tried to bring these different lines together, showing how progress in all the above lines of inquiry "crosses field lines." In summary, I hope this paper will help others, as it has helped me, to go to the original sources and *learn from learning in economics.*

REFERENCES[11]

Anderlini, Luca and Sabourian, Hamid 1995b. "The Evolution of Algorithmic Learning Rules: A Global Stability Result," mimeo, Cambridge University.

Anscombe, F. and Aumann, Robert 1963. "A Definition of Subjective Probability," *Annals of Mathematics and Statistics*, **34**, 199-205.

Arifovic, Jasmina, Bullard, James and Duffy, John 1995. "Learning in a Model of Economic Growth and Development." (mimeo) Department of Economics, Universiry of Pittsburgh.

[11]Unfortunately, space limitations have prevented me from properly quoting many interesting contributions. The interested reader can find a more complete list of references in my www page at <http://www.iue.it/ECO/marimon.html>.

Arthur, Brian W. 1993. "On Designing Economic Agents that Behave Like Human Agents." *Journal of Evolutionary Economics*, **3**, 1–22.

Arthur, Brian W. 1995. "Complexity in Economic and Financial Markets." *Journal of Complexity*, **1**.

Aumann, Robert J. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* **55**, 1-18.

Battigalli, P., Gilli, M and Molinari, M.C. 1994. "Learning Convergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey." *Ricerche Economiche.*

Benhaïm, M. and Hirsch, M.W. 1994. "Learning Processes, Mixed Equilibria and Dynamical Systems Arising from Repeated Games." (mimeo) Department of Mathematics, University of California at Berkeley.

Benveniste, Albert, Métivier, Michel and Priouret, Pierre 1990. *Adaptive Algorithms and Stochastic Approximations*, Berlin: Springer-Verlag.

Binmore, K. 1988. "Modeling Rational Players: Part II." *Economics and Philosophy*. **4**, 9–55.

Blume, Lawrence E. and and Easley, David 1995. "What has the Rational Learning Literature Taught Us?" In Kirman, A. and Salmon, M. (eds.) 12–39.

Bray, Margaret M. 1982. "Learning Estimation and the Stability of Rational Expectations," *Journal of Economic Theory* **26**, 318–339.

Bray, Margaret M. and Kreps, David 1987. "Rational Learning and Rational Expectations," in Feiwel, George (ed.) *Arrow and the Ascent of Modern Economic Theory.* New York: New York University Press. 597–625.

Brock, William A. and Hommes, Cars H. 1995. "Rational Routes to Randomness," SSRI wp # 9506, Department of Economics, University of Wisconsin.

Brock, William A. and LeBaron, Blake 1995. "A Dynamic Structural Model for Stock Return Volatility and Trading Volume." NBER wp # 4988.

Brousseau, Vincent and Kirman, Alan 1995. "The Dynamics of Learning in N-Person Games with the Wrong N." in Kirman, A. *et al.* (eds.)

Brown, G. 1951. "Iterated Solution of Games by Fictitious Play," in Koopmans, T.C. (ed.) Activity Analysis of Production and Allocation, Wiley: New York, 374–376.

Bullard, James 1994. "Learning Equilibria," *Journal of Economic Theory* **64** 2, 468–485.

Busch, R.R. and Mosteller, F. 1955. *Stochastic Models of Learning.* New York: Wiley.

Canning, David 1992. "Average Behavior in Learning Models." *Journal of Economic Theory*, **57**, 442–472.

Cho, In-Koo 1994. "Bounded Rationality, Neural Networks and Folk Theorem in Repeated Games with Discounting," *Economic Theory* **4** 6, 935–957.

Cho, In-Koo 1995. "Perceptrons Play Repeated Games with Imperfect Monitoring," mimeo, University of Chicago.

Cho, In-Koo and Sargent, Thomas 1995. "Neural Networks for Encoding and Adapting in Dynamic Economies," in *Handbook of Computational Economics.*

Dawid, A. P. 1982. "The Well Calibrated Bayesian," *Journal of The American Statistical Association,* **77**, 605–613.

de Fontnouvelle, Patrick 1995. "Informational Strategies in Financial Markets: the Implications for Volatility and Trading Volume Dynamics." (mimeo) Iowa State University.

de Groot, Morris H. 1970. *Optimal Statistical Decisions.* New York: McGraw-Hill.

Duffy, John 1994. "On Learning and the Nonuniqueness of Equilibrium in an Overlapping Generations Model with Fiat Money," *Journal of Economic Theory* **64** 2, 541–553.

Easley, David and Rustichini, Aldo 1995. "Choice without Belief" mimeo, Cornell University.

Evans, George W. and Honkapohja, Seppo 1994. "Convergence of Least Squares Learning to a Non-Stationary Equilibrium," *Economic Letters*, **46**, 131–136.

Evans, George W. and Honkapohja, Seppo 1995a. "Adaptive Learning and Expectational Stability: An Introduction", in Kirman, A. and Salmon, M. (eds.) 102–126.

Evans, George W. and Honkapohja, Seppo 1995b. "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models," *Econometrica* **63** 1, 195-206.

Evans, George W., Honkapohja, Seppo and Marimon, Ramon 1995. "Convergence in Monetary Models with Heterogeneous Learning Rules," mimeo, European University Institute.

Evans, George W., Honkapohja, Seppo and Sargent, Thomas 1993. "On the Preservation of Deterministic Cycles when Some Agents Perceive them to be Random Fluctuations," *Journal of Economic Dynamics and Control*, **17**, 705–721.

Evans, George W. and Ramey, Gary 1994. "Expectation Calculation, Hyperinflation and Currency Collapse", forthcoming in *The New Macroeconomics: Imperfect Markets and Policy Effectiveness*, eds. H. Dixon and N. Rankin, Cambridge: Cambridge University Press.

Foster, Dean and Vohra, Rakesh V. 1995. "Calibrated Learning and Correlated Equilibrium," mimeo, University of Pennsylvania.

Fudenberg, Drew and Kreps, David 1993. "Learning Mixed Equilibria" *Games and Economic Behavior.*

Fudenberg, Drew and Kreps, David 1994. "Learning in Extensive Games, II: Experimentation and Nash Equilibrium," Economic Theory Disc. Paper #20 Harvard Insitute for Economic Research, Harvard University.

Fudenberg, Drew and Kreps, David 1995. "Learning in Extensive Games, I: Self Confirming Equilibrium," *Games and Economic Behavior* **8**, 20–55.

Fudenberg, Drew and Levine, David 1993. "Self Confirming Equilibria" *Econometrica* **61** 3, 523–546.

Fudenberg, Drew and Levine, David 1995a. "Universal Consistency and Cautious Fictitious Play," *Journal of Economic Dynamics and Control* (forthcoming).

Fudenberg, Drew and Levine, David 1995b. *Theory of Learning in Games* (manuscript) UCLA.

Grandmont, Jean-Michael 1994. "Expectations Formation and Stability of Large Socioeconomic Systems," WP No. 9424, CEPREMAP, Paris.

Grandmont, Jean-Michael, and Laroque, Guy 1991. "Economic Dynamics with Learning: Some Instability Examples," in *Equilibrium Theory and Applications, Proceedings of the Sixth International Symposium in Economic Theory and Econometrics*, ed. by W. A. Barnett *et al.*. Cambridge, Cambridge University Press, 247–273.

Guesnerie, Roger and Woodford, Michael 1991. "Stability of Cycles with Adaptive Learning Rules," in Barnett, *et al.* (eds.) *Equilibrium Theory and Applications.* Cambridge: Cambridge University Press.

Hahn, Frank 1973. *On the Notion of Equilibrium in Economics: An Inaugural Lecture.* Cambridge: Cambridge University Press.

Hart, Sergiu and Schmeidler, David 1989. "Existence of Correlated Equilibria," *Mathematics of Operations Research*, **14**.

Hayek, Frederik A. 1945. "The Use of Knowledge in Society," *American Economic Review,* **35**, 519–530.

Hendon, Ebbe, Jacobsen Hans J. and Sloth, Birgitte 1995. "Adaptive Learning in Extensive Form Games and Sequential Equilibrium," mimeo, Institute of Economics, University of Copenhagen.

Holland, John 1995. *Hidden Order: How Adaptation Builds Complexity.* Menlo Park, CA: Addison-Wesley.

Howitt, Peter 1992. Interest Rate Control and Nonconvergence of Rational Expectations," *Journal of Political Economy*, **100**, 776–800.

Hussman, John 1992. "Market Efficiency and Inefficiency in Rational Expectations Equilibria," *Journal of Economic Dynamics and Control*, **16**, 655–680.

Hurkens, Sjaak 1994. "Learning by Forgetful Players: From Primitive Formations to Persistent Retracts," WP No. 9437, CentER, Tilburg University.

Jordan, James S. 1991. "Bayesian Learning in Normal Form Games," *Games and Economic Behavior*, **3**, 60–81.

Jordan, James S.1993. "Three Problems in Learning Mixed–Strategy Nash Equilibria," *Games and Economic Behavior* **5**, 368–386.

Kagel, John H. and Roth, Alvin E. 1995. *The Handbook of Experimental Economics.* Princeton, NJ. Princeton University Press.

Kalai, Ehud and Lehrer, Ehud 1993a. "Rational Learning Leads to Nash Equilibrium," *Econometrica* **61** 5, 1019–1046.

Kalai, Ehud and Lehrer, Ehud 1993b. "Subjective Equilibria in Repeated Games," *Econometrica* **61**, 1231–1240.

Kalai, Ehud and Lehrer, Ehud 1995. "Subjective Games and Equilibria," *Games and Economic Behavior* **8**, 123–163.

Kaniovski, Yuri M. and Young, H. Peyton 1995. "Learning Dynamics in Games with Stochastic Perturbations," *Games and Economic Behavior* **11**, 330–363.

Kirman, Alan and Salmon, Mark (eds.) 1995. *Learning and Rationality in Economics.* Blackwell: Oxford, UK.

Kreps, David 1990. *Game Theory and Economic Modeling.* Oxford: Clarendon Press.

Kurz, Mordecai 1994a. "On the Structure and Diversity of Rational Beliefs," *Economic Theory* **4** 6, 877–900.

Kurz, Mordecai 1994b. "On Rational Belief Equilibria," *Economic Theory* **4** 6, 859–876.

Ljung, Lennart, Pflug, Georg and Walk, Harro 1992. *Stochastic Approximation and Optimization of Random Systems*, Basel: Birkhauser.

Marcet, Albert and Nicolini, Juan P. 1995. "Recurrent Hyperinflations and Learning," mimeo, Universitat Pompeu Fabra.

Marcet, Albert, and Sargent, Thomas 1989a. "Convergence of Least Squares Learning Mechanisms in Self Referential, Linear Stochastic Models," *Journal of Economic Theory* 48, 337–368.

Marcet, Albert, and Sargent, Thomas 1989b. "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information", *Journal of Political Economy*, 97, 1306-1322.

Marimon, Ramon and McGrattan, Ellen 1995. "On Adaptive Learning in Strategic Games," in Kirman, A. and Salmon, M. (eds.) 63–101.

Marimon, Ramon, McGrattan, Ellen and Sargent, Thomas 1990. "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents," *J. Economic Dynamics and Control*, **47**, 282-366.

Marimon, Ramon and Sunder, Shyam 1993. "Indeterminacy of Equilibria in an Hyperinflationary World: Experimental Evidence", *Econometrica* **61** 5, 1073–1108. missing

Marimon, Ramon and Sunder, Shyam 1995. "Does a Constant Money Growth Rule Help Stabilize Inflation?: Experimental Evidence." (mimeo) European University Institute, *Journal of Monetary Economics*.

Marimon, Ramon, Spear, Stephen and Sunder, Shyam 1993. Expectationally-driven Market Volatility: An Experimental Study. *Journal of Economic Theory*.

Milgrom, Paul and Roberts, John 1991. "Adaptive and Sophisticated Learning in Normal Form Games," *Games and Economic Behavior*, **3**, 82–100.

Miyasawa, K. 1961. "On the Convergence of the Learning Process in a 2×2 Non-zero-sum Game," Econometric Research Program, Research Memorandum No 33, Princeton University.

Muth, John F. 1960. "Optimal Properties of Exponentially Weighted Forecasts." *Journal of the American Statistical Association* **55**, 299–306.

Nachbar, John H. 1995. "Prediction, Optimization and Rational Learning in Games," mimeo, Washington University, St. Louis.

Nash, John F. 1950. "Equilibrium Points in n-Person Games," *Proc. Nat. Acad. Sci. USA*, **36**, 48–49.

Natarajan, Balas K. 1991. *Machine Learning*, Morgan Kaufmann Publishers, Inc.: San Mateo, California.

Nyarko, Yaw 1994. "Bayesian Learning Leads to Correlated Equilibria in Normal Form Games, " *Economic Theory* **4** 6, 821–842.

Popper, Karl R. 1979. *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.

Posch, Martin 1995. "Cycling in Stochastic Learning Algorithms for Normal Form Games," mimeo. University of Vienna.

Sandroni, Alvaro 1995. "Does Rational Learning Lead to Nash Equilibrium in Finitely Repeated Games?" mimeo, University of Pennsylvania.

Sargent, Thomas J. 1993. *Bounded Rationality in Macroeconomics.* Oxford: Claredon –Oxford University– Press.

Shapley, Lloyd 1964. "Some Topics in Two–Person Games," *In Advances in Game Theory, Annals of Mathematical Studies*, **5**, 1–28.

Sonsino, Doron 1995. "Learning to Learn, Pattern Recognition, and Nash Equilibrium," mimeo, GSB Stanford University.

Swinkels, Jeroen M. 1993. "Adjustment Dynamics and Rational Play in Games," *Games and Economic Behavior* **5**, 455–484.

Timmermann, Alan 1993. "How Learning in Financial Markets Generates Excess Volatility and Predictability of Excess Returns," *Quarterly Journal of Economics*, **108** 1135–1145.

Townsend, Robert M. 1983. "Forecasting the Forecasts of Others". *Journal of Political Economy*, 91, 546-588.

van Huyck, John, Battalio, Raymond and Cook, Joseph 1994. "Selection Dynamics, Asymptotic Stability, and Adaptive Behavior," *Journal of Political Economy*, **102**, 975–1005.

White, Halbert 1992. *Artificial Neural Networks. Approximation and Learning Theory.* Blackwell: Oxford, UK.

Woodford, Michael 1990. "Learning to Believe in Sunspots," *Econometrica* 58, 2, 277-308.

Young, H. Peyton 1993. "The Evolution of Conventions," *Econometrica* **61** 1, 57–84.

EUI WORKING PAPERS

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

Copies 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

Publications of the European University Institute

To The Publications Officer
European University Institute
Badia Fiesolana
I-50016 San Domenico di Fiesole (FI) – Italy
Telefax No: +39/55/4685 636
E-mail: publish@datacomm.iue.it

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

❐ Please send me a complete list of EUI Working Papers
❐ Please send me a complete list of EUI book publications
❐ Please send me the EUI brochure Academic Year 1996/97

Please send me the following EUI Working Paper(s):

No, Author . . . .
Title: . . . .
No, Author . . . .
Title: . . . .
No, Author . . . .
Title: . . . .
No, Author . . . .
Title: . . . .

Date . . . .

Signature . . . .

Working Papers of the Department of Economics Published since 1994

*out of print

ECO No. 94/20
Andrzej BANIAK/Louis PHLIPS
La Pléiade and Exchange Rate Pass-Through

ECO No. 94/21
Mark SALMON
Bounded Rationality and Learning; Procedural Learning

ECO No. 94/22
Isabelle MARET
Heterogeneity and Dynamics of Temporary Equilibria: Short-Run Versus Long-Run Stability

ECO No. 94/23
Nikolaos GEORGANTZIS
Short-Run and Long-Run Cournot Equilibria in Multiproduct Industries

ECO No. 94/24
Alexander SCHRADER
Vertical Mergers and Market Foreclosure: Comment

ECO No. 94/25
Jeroen HINLOOPEN
Subsidising Cooperative and Non-Cooperative R&D in Duopoly with Spillovers

ECO No. 94/26
Debora DI GIOACCHINO
The Evolution of Cooperation: Robustness to Mistakes and Mutation

ECO No. 94/27
Kristina KOSTIAL
The Role of the Signal-Noise Ratio in Cointegrated Systems

ECO No. 94/28
Agustín MARAVALL/Víctor GÓMEZ
Program SEATS "Signal Extraction in ARIMA Time Series" - Instructions for the User

ECO No. 94/29
Luigi ERMINI
A Discrete-Time Consumption-CAP Model under Durability of Goods, Habit Formation and Temporal Aggregation

ECO No. 94/30
Debora DI GIOACCHINO
Learning to Drink Beer by Mistake

ECO No. 94/31
Víctor GÓMEZ/Agustín MARAVALL
Program TRAMO "Time Series Regression with ARIMA Noise, Missing Observations, and Outliers" - Instructions for the User

ECO No. 94/32
Ákos VALENTINYI
How Financial Development and Inflation may Affect Growth

ECO No. 94/33
Stephen MARTIN
European Community Food Processing Industries

ECO No. 94/34
Agustín MARAVALL/Christophe PLANAS
Estimation Error and the Specification of Unobserved Component Models

ECO No. 94/35
Robbin HERRING
The "Divergent Beliefs" Hypothesis and the "Contract Zone" in Final Offer Arbitration

ECO No. 94/36
Robbin HERRING
Hiring Quality Labour

ECO No. 94/37
Angel J. UBIDE
Is there Consumption Risk Sharing in the EEC?

ECO No. 94/38
Berthold HERRENDORF
Credible Purchases of Credibility Through Exchange Rate Pegging: An Optimal Taxation Framework

ECO No. 94/39
Enrique ALBEROLA ILA
How Long Can a Honeymoon Last? Institutional and Fundamental Beliefs in the Collapse of a Target Zone

ECO No. 94/40
Robert WALDMANN
Inequality, Economic Growth and the Debt Crisis

*out of print

ECO No. 94/41
John MICKLEWRIGHT/
Gyula NAGY
Flows to and from Insured
Unemployment in Hungary

ECO No. 94/42
Barbara BOEHNLEIN
The Soda-ash Market in Europe:
Collusive and Competitive Equilibria
With and Without Foreign Entry

ECO No. 94/43
Hans-Theo NORMANN
Stackelberg Warfare as an Equilibrium
Choice in a Game with Reputation Effects

ECO No. 94/44
Giorgio CALZOLARI/Gabriele
FIORENTINI
Conditional Heteroskedasticity in
Nonlinear Simultaneous Equations

ECO No. 94/45
Frank CRITCHLEY/Paul MARRIOTT/
Mark SALMON
On the Differential Geometry of the Wald
Test with Nonlinear Restrictions

ECO No. 94/46
Renzo G. AVESANI/Giampiero M.
GALLO/Mark SALMON
On the Evolution of Credibility and
Flexible Exchange Rate Target Zones *

***

ECO No. 95/1
Paul PEZANIS-CHRISTOU
Experimental Results in Asymmetric
Auctions - The 'Low-Ball' Effect

ECO No. 95/2
Jeroen HINLOOPEN/Rien
WAGENVOORT
Robust Estimation: An Example

ECO No. 95/3
Giampiero M. GALLO/Barbara PACINI
Risk-related Asymmetries in Foreign
Exchange Markets

ECO No. 95/4
Santanu ROY/Rien WAGENVOORT
Risk Preference and Indirect Utility in
Portfolio Choice Problems

ECO No. 95/5
Giovanni NERO
Third Package and Noncooperative
Collusion in the European Airline
Industry *

ECO No. 95/6
Renzo G. AVESANI/Giampiero M.
GALLO/Mark SALMON
On the Nature of Commitment in Flexible
Target Zones and the Measurement of
Credibility: The 1993 ERM Crisis *

ECO No. 95/7
John MICKLEWRIGHT/Gyula NAGY
Unemployment Insurance and Incentives
in Hungary

ECO No. 95/8
Kristina KOSTIAL
The Fully Modified OLS Estimator as a
System Estimator: A Monte-Carlo
Analysis

ECO No. 95/9
Günther REHME
Redistribution, Wealth Tax Competition
and Capital Flight in Growing
Economies

ECO No. 95/10
Grayham E. MIZON
Progressive Modelling of
Macroeconomic Time Series: The LSE
Methodology *

ECO No. 95/11
Pierre CAHUC/Hubert KEMPF
Alternative Time Patterns of Decisions
and Dynamic Strategic Interactions

ECO No. 95/12
Tito BOERI
Is Job Turnover Countercyclical?

ECO No. 95/13
Luisa ZANFORLIN
Growth Effects from Trade and
Technology

ECO No. 95/14
Miguel JIMÉNEZ/Domenico
MARCHETTI, jr.
Thick-Market Externalities in U.S.
Manufacturing: A Dynamic Study with
Panel Data

*out of print

ECO No. 95/15
Berthold HERRENDORF
Exchange Rate Pegging, Transparency, and Imports of Credibility

ECO No. 95/16
Günther REHME
Redistribution, Income cum Investment Subsidy Tax Competition and Capital Flight in Growing Economies

ECO No. 95/17
Tito BOERI/Stefano SCARPETTA
Regional Dimensions of Unemployment in Central and Eastern Europe and Social Barriers to Restructuring

ECO No. 95/18
Bernhard WINKLER
Reputation for EMU - An Economic Defence of the Maastricht Criteria

ECO No. 95/19
Ed HOPKINS
Learning, Matching and Aggregation

ECO No. 95/20
Dorte VERNER
Can the Variables in an Extended Solow Model be Treated as Exogenous? Learning from International Comparisons Across Decades

ECO No. 95/21
Enrique ALBEROLA-ILA
Optimal Exchange Rate Targets and Macroeconomic Stabilization

ECO No. 95/22
Robert WALDMANN
Predicting the Signs of Forecast Errors *

ECO No. 95/23
Robert WALDMANN
The Infant Mortality Rate is Higher where the Rich are Richer

ECO No. 95/24
Michael J. ARTIS/Zenon G. KONTOLEMIS/Denise R. OSBORN
Classical Business Cycles for G7 and European Countries

ECO No. 95/25
Jeroen HINLOOPEN/Charles VAN MARREWIJK
On the Limits and Possibilities of the Principle of Minimum Differentiation

ECO No. 95/26
Jeroen HINLOOPEN
Cooperative R&D Versus R&D-Subsidies: Cournot and Bertrand Duopolies

ECO No. 95/27
Giampiero M. GALLO/Hubert KEMPF
Cointegration, Codependence and Economic Fluctuations

ECO No. 95/28
Anna PETTINI/Stefano NARDELLI
Progressive Taxation, Quality, and Redistribution in Kind

ECO No. 95/29
Ákos VALENTINYI
Rules of Thumb and Local Interaction *

ECO No. 95/30
Robert WALDMANN
Democracy, Demography and Growth

ECO No. 95/31
Alessandra PELLONI
Nominal Rigidities and Increasing Returns

ECO No. 95/32
Alessandra PELLONI/Robert WALDMANN
Indeterminacy and Welfare Increasing Taxes in a Growth Model with Elastic Labour Supply

ECO No. 95/33
Jeroen HINLOOPEN/Stephen MARTIN
Comment on Estimation and Interpretation of Empirical Studies in Industrial Economics

ECO No. 95/34
M.J. ARTIS/W. ZHANG
International Business Cycles and the ERM: Is there a European Business Cycle?

ECO No. 95/35
Louis PHLIPS
On the Detection of Collusion and Predation

ECO No. 95/36
Paolo GUARDA/Mark SALMON
On the Detection of Nonlinearity in Foreign Exchange Data

*out of print

ECO No. 95/37
Chiara MONFARDINI
Simulation-Based Encompassing for Non-Nested Models: A Monte Carlo Study of Alternative Simulated Cox Test Statistics

ECO No. 95/38
Tito BOERI
On the Job Search and Unemployment Duration

ECO No. 95/39
Massimiliano MARCELLINO
Temporal Aggregation of a VARIMAX Process

ECO No. 95/40
Massimiliano MARCELLINO
Some Consequences of Temporal Aggregation of a VARIMA Process

ECO No. 95/41
Giovanni NERO
Spatial Multiproduct Duopoly Pricing

ECO No. 95/42
Giovanni NERO
Spatial Multiproduct Pricing: Empirical Evidence on Intra-European Duopoly Airline Markets

ECO No. 95/43
Robert WALDMANN
Rational Stubbornness?

ECO No. 95/44
Tilman EHRBECK/Robert WALDMANN
Is Honesty Always the Best Policy?

ECO No. 95/45
Giampiero M. GALLO/Barbara PACINI
Time-varying/Sign-switching Risk Perception on Foreign Exchange Markets

ECO No. 95/46
Víctor GÓMEZ/Agustín MARAVALL
Programs TRAMO and SEATS
Update: December 1995

***

ECO No. 96/1
Ana Rute CARDOSO
Earnings Inequality in Portugal: High and Rising?

ECO No. 96/2
Ana Rute CARDOSO
Workers or Employers: Who is Shaping Wage Inequality?

ECO No. 96/3
David F. HENDRY/Grayham E. MIZON
The Influence of A.W.H. Phillips on Econometrics

ECO No. 96/4
Andrzej BANIAK
The Multimarket Labour-Managed Firm and the Effects of Devaluation

ECO No. 96/5
Luca ANDERLINI/Hamid SABOURIAN
The Evolution of Algorithmic Learning: A Global Stability Result

ECO No. 96/6
James DOW
Arbitrage, Hedging, and Financial Innovation

ECO No. 96/7
Marion KOHLER
Coalitions in International Monetary Policy Games

ECO No. 96/8
John MICKLEWRIGHT/ Gyula NAGY
A Follow-Up Survey of Unemployment Insurance Exhausters in Hungary

ECO No. 96/9
Alastair McAULEY/John MICKLEWRIGHT/Aline COUDOUEL
Transfers and Exchange Between Households in Central Asia

ECO No. 96/10
Christian BELZIL/Xuelin ZHANG
Young Children and the Search Costs of Unemployed Females

ECO No. 96/11
Christian BELZIL
Contiguous Duration Dependence and Nonstationarity in Job Search: Some Reduced-Form Estimates

*out of print

ECO No. 96/12
Ramon MARIMON
Learning from Learning in Economics

*out of print