EUI WORKING PAPERS IN ECONOMICS

EUI Working Paper ECO No. 94/21

Bounded Rationality and Learning; Procedural Learning

MARK SALMON

European University Institute, Florence

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE

ECONOMICS DEPARTMENT

EUI Working Paper **ECO** No. 94/21

Bounded Rationality and Learning;
Procedural Learning

MARK SALMON

BADIA FIESOLANA, SAN DOMENICO (FI)

Printed in Italy in May 1994
European University Institute
Badia Fiesolana
I – 50016 San Domenico (FI)
Italy

Bounded Rationality and Learning; Procedural Learning

Mark Salmon
Department of Economics,
European University Institute,
Florence, Italy.
April 20, 1994

Abstract

In what follows we consider the question of boundedly rational learning and expectation formation by economic agents. Every-day observation suggests that informational constraints shape an individual's objectives, their learning activity and ultimately the type of decision rules they adopt. It also suggests that behavioural issues that arise through the interaction of an individual with their economic environment and their perception of that environment might be important when attempting to understand how economic agents actually formulate decision problems. This interaction may induce behaviour which is far from the fully rational model of *homo economicus* that provides the standard paradigm of economic theory. It also seems important to recognise that a limited knowledge of the environment may affect, not only information sets but also the *manner* by which people learn and hence behavioural theories of learning would seem to be called for. Herbert Simon, in his outstanding original contributions to the theory of boundedly rational behaviour discussed the distinction between procedural and substantive rationality and a similar distinction could perhaps be usefully drawn between procedural and substantive methods of learning.

In this chapter we explore the role of artificial neural networks in providing a conceptually simple, "non-structural", *procedural* model of how agents might learn to approximate their true but unknown conditional expectation function and hence form "boundedly rational" expectations. The minimal information required, simply a knowledge of the input and output variables, does not appear to seriously hinder the performance of the approach, either in theory or in practice. Moreover in principle a neural network approximation can evolve in complexity and hence accuracy as knowledge of the environment increases; such adaption to the environment reflects a behavioural aspect of learning which is invariably missing in the standard models of learning assumed in economics. In this chapter we examine the performance of neural network learning, first very briefly in theory and then empirically with two examples drawn from macro-economic policy issues. Our results, while positive regarding the application of neural network learning, lead us to suggest caution against drawing specific conclusions regarding economic behaviour or policy that are dependent on specific and potentially *ad hoc* assumptions as to how people learn and hence adjust their expectations. The question of how individuals learn in economic environments needs to be considered more deeply and incorporated into a fully integrated theory of boundedly rational economic behaviour.

I would like to thank Robin Wensley for several very rational discussions about bounded rationality and Alan Kirman for comments on an earlier version of this paper.

1 Introduction

Given the critical role that expectational forces have assumed in economics and the intellectual effort that has been devoted to their study it is perhaps surprising that we still have a relatively crude understanding of how expectations are *actually* formed and adjusted in the light of events. The available econometric evidence on the Rational Expectations Hypothesis, for instance, is at best ambiguous if indeed it does not clearly indicate rejection (see for example, Favero and Hendry (1992)). This undoubtedly comes as little surprise to many confirming the apparently consistently "irrational" behaviour regularly observed in experimental settings (see *inter alios* Slovic and Lichtenstein (1983), Tversky and Kahneman (1986), Tversky, Slovic and Kahneman (1991).

How should economists respond to these results? The contrast between formal optimal decision theory on the one hand and common sense and every-day observation on the other is too strong to ignore and yet it seems surprisingly difficult to give up the notion of rationality in expectation formation entirely. Baumol and Quandt (1964) nearly thirty years ago referred to " an irrational passion for dispassionate rationality" and Bray and Kreps (1987) note "a failing of economic theory in general has been that it has proved remarkably resistant to movements away from models with full rationality and consistency". Economists it would seem have been more concerned with how people *should* behave rather than with how they actually *do* behave.

We are naturally led to models of boundedly rational behaviour of some form, and the need to refine our understanding of how economic decisions are actually made if they are not always to be based on strict adherence to subjective expected utility maximisation by rational agents with perfect recall and perfect powers of inference – *homo economicus!*. The dominance of the rationality paradigm is tied most probably to an intellectual complacency that derives from the precision of the concept and the corresponding apparent lack of precision in alternative notions of bounded rationality. The means by which individuals learn, forget, process information and hence adjust expectations then becomes central to our understanding of bounded rationality and while any number of convenient assumptions can be made no generally accepted consensus or unambiguous theory as to what determines reasonable boundedly rational behaviour or learning appears to be in place. Instead the results we have regarding learning seem to relate more to the properties of particular statistical algorithms such as Least Squares or Bayes rule rather than to an underlying and more broadly justified model of economic behaviour for the construction of sensible decision rules in the face of limited knowledge of the economic environment. Why should economic agents be expected to use Least Squares, for instance, as opposed to some other method? Since both Least Squares and Bayesian learning may be seen as forms of stochastic approximation algorithms and given that there are a host of apparently equally justified alternatives both within and *outside* this class a better understanding of the behavioural basis of learning may be needed. We need to be careful that the conclusions we draw in any exercise that involves learning actually have some basis in economic behaviour and reality and are not just an artifact of the assumed statistical algorithm. Learning needs to be seen as an integral part of a boundedly rational decision problem and the manner by which we assume individuals learn justified in that context.

Common-sense also suggests that informational limitations imposed by a partial knowledge of the economic environment shape an individual's objectives, their learning activity and ultimately we would also expect the type of decision rules they adopt. This interaction between an individual and their perception of the economic environment in which they exist may induce behaviour which ranges far from the fully rational model of *homo economicus* that forms the standard paradigm of economic theory. It is important to recognise that these informational constraints may also affect the *manner* by which people learn, not only their information sets, and hence behavioural theories of learning again appear relevant. It seems strange from the behavioural point of view, for instance, to assume as is the case with the standard statistical models of "rational" learning , that agents have complete knowledge of the relevant economic structure and yet are assumed to be completely ignorant of just a subset of the parameter values within that structure[1]. The economic interactions that have taken place in the past to have left an individual in such an odd state are unspecified. A more reasonable position might be that an agent's knowledge of the structure and their learning activity evolve symbiotically and the manner by which learning takes place adapts to their increased understanding of their economic environment which in turn may grow, according to economic incentives, through deliberately increased interaction with that environment. Some flexibility within the method of learning is then needed as the agent's approximation to reality improves. Otherwise assuming that the structure itself is known and fixed from the outset may suggest one particular method of learning as optimal which may be completely inappropriate if the structure were initially unknown (and *vice-versa*). The trade off between robustness to mis-specification and efficiency is well recognised in statistics.

In this chapter we seek to discuss these issues by emphasising the interaction of the agent with the environment in which they are forced to take decisions. "Behavioural" models of decision making are, of course, far from new in Economics and were strongly emphasised in the work of Herbert Simon and others in the 1950's and 1960's (see, for instance, Simon (1955,1959,1982) and Alchian (1950)). In particu-

[1]Often for instance the "regression parameters" are assumed to be unknown but the residual variance is assumed to be known.

lar Simon stressed two aspects of boundedly rational behaviour; satisficing and the limited cognitive powers of economic agents. He also drew a distinction between procedural and substantive rationality and a similar distinction could perhaps be usefully drawn between procedural and substantive methods of learning. Substantive rationality is more concerned with the end results or choices individuals make whereas procedural rationality is concerned with the process by which these choices are made. While all learning is in one sense therefore procedural some models of learning appear to be more concerned with the properties of the end results of learning rather than with the behavioural basis or the manner by which learning takes place. Is convergence to a rational expectations equilibrium for instance a desirable or necessary objective for a model of learning under bounded rationality?

Taking these observations into account, the view of learning that we put forward below is one in which individuals construct explicitly *approximate* models of their environment which are updated as their information improves through either active or passive learning and that the decision rules they adopt reflect the fact that they know they hold mis-specified models of reality. Learning then is seen to apply *both* to the structure and any relevant parameterisation that the approximation entails. In principle the whole branch of non-parametric inference can be brought to bear on this problem although we only consider one approach based on artificial neural networks below.

Artificial neural networks provide a conceptually simple procedural model of how agents might learn to approximate their true but unknown conditional expectation function and hence form "boundedly rational" expectations. The minimal structural information required, simply a knowledge of the input and output variables, does not appear to seriously hinder the performance of the method either in theory or in practice. Moreover in principle a neural network approximation can evolve in complexity and hence accuracy as knowledge of the environment increases reflecting the behavioural aspect of learning suggested above. Such "non-parametric" learning schemes that require no *ex ante* knowledge of the structure may serve as good rules of thumb that have satisfactory operational characteristics and it is therefore of interest to explore how they perform in comparison with standard methods such as least squares learning that requires full knowledge of the structure to deliver good results.

The next section reviews aspects of boundedly rational learning, expectation formation and decision making and emphasises the procedural and behavioural characteristics of neural network learning. Two applications of the method to macro policy problems are then provided in section 3. The first considers the role of learning in the monetary policy and inflation model originally developed by Barro and

Gordon (1983) and centres on the policy maker's desire to establish credibility. The ability of a policy maker to " surprise" the private sector is clearly dependent on how the private sector forms expectations and learns and hence the potential to exploit its reputation in policy making is also critically dependent on how the private sector learns. The second application considers the Sargent/Wallace hyperinflation model used by Marcet and Sargent (1989) to investigate the impact of least squares learning. We find, for example, in this second exercise that the domain of attraction of the locally stable equilibrium is apparently much wider with neural network learning than with least squares learning although convergence is to a different stationary point in each case with a consistently lower rate of inflation being found by neural network learning across a wide range of simulations. The sensitivity of the least squares in this example is in marked contrast to the observed robustness of neural network learning.

2 Bounded Rationality and Learning

Simon (1976) suggested two forms of Bounded Rationality; Substantive and Procedural. "Behaviour is substantively rational when it is appropriate to the achievement of given goals within the limits imposed by given conditions and constraints". Substantive rationality then focuses on the outcomes or choices made by economic agents whereas procedural rationality requires that observed behaviour is the outcome from some process of appropriate reasoning and action given available knowledge and powers of computation. Simon characterises standard economic analysis as resting on two assumptions; the identification of a goal such as utility maximisation and substantive rationality. Psychologists on the other hand, have been more concerned with identifying the process or procedure that has determined behaviour which is then considered to be procedurally rational. In the context of a game of chess, Simon argues that as the complexity of the decision problem grows we are naturally forced to induce procedural rationality in place of substantive rationality which involves a shift from optimality to finding a good method of selecting moves. As in this example, but also more generally, the limited cognitive powers of economic agents forces the change from optimising to satisficing behaviour. It is perhaps also important to recognise that "rationality" has to be interpreted with care since it is not in itself absolute but only defined relative to a given loss function. Rational behaviour is then associated with action consistent with that loss function and different loss functions will imply different rational behaviour. The commonplace understanding of "rationality" is the application of correct *reasoning* to a given problem and this can clearly be applied both to evaluate the "rationality" of the outcomes of a decision problem as well as the process by which those outcomes were achieved. Bounded rationality then reflects a limited capacity to reason both because of limited information and limited powers of computation.

We should also briefly note the implications of the Theory of Cognitive Dissonance developed by Festinger (1957) which argues amongst other things that individuals strive for cognitive consistency. The observation by psychologists that individuals dislike inconsistency in their ideas and act to reduce it does not however imply full rationality, merely that individuals will adjust their subjective beliefs to achieve internal consistency which may not lead to what economists would call a rational expectations equilibrium. Since internal consistency does not necessarily coincide with the objective consistency of the correct model structure the argument that learning should lead to a rational expectations equilibrium is weakened. Since removing inconsistency does not then imply convergence to a rational expectations equilibrium, justifying a learning scheme by its ability to converge to the rational expectations equilibrium would not then necessarily seem to be supported behaviourally.

Simon's argument for focusing attention on procedural rationality given the limited cognitive powers of economic agents argues against the full structural knowledge assumptions frequently employed in rational learning schemes and against the use of standard statistical algorithms that rest on these assumptions to generate "optimal" results. As elsewhere in economics such *optimal* decisions that follow from solving some marginal or first order condition often pay little attention to the behavioural basis for carrying out the optimum action. Alternatively, so called *ad hoc* decision rules or "*rules of thumb*" can be put forward that represent actually observed behaviour or are explicitly behaviourally justified given the assumed limited information and reasoning power of the decision maker (see, for instance Baumol and Quandt (1964), Radner(1975), Day(1967), Crain et al. (1984) and Wall(1993). It is interesting to observe from these references that at least in some cases such satisficing behaviour can converge over time to optimal behaviour. These behavioural theories of action focus attention explicitly on the process by which decisions are made and therefore fall into the class of procedurally rational schemes; Radner (1975) , for instance, considered the allocation of effort by a manager between different tasks leading to behaviour which he called " putting out fires" and Baumol and Quandt (1964) and Day (1967) consider other strategies corresponding to what may be easily recognised from everyday observation as reasonable behaviour but which does not necessarily immediately deliver the optimal result. What lies at the heart of the issue is the nature of the loss function and the relevant behavioural process and constraints if we wish to formalise the agent's boundedly rational decision problem to find the "optimal imperfect decision" . When attempting to construct a general theory of boundedly rational behaviour a difficulty then lies in the identification of general restrictions that reflect *procedurally* rational behaviour. [2] Since removing constraints imposed by ignorance or the lack of information is the obvious objective of learning one step in this process is then to consider procedurally rational learning schemes and the behavioural basis for learning in economic environments.

2.1 Procedural Learning

How can we then develop procedurally rational theories of learning? Everyday life once again seems to indicate that we do not in fact *need* a complete knowledge of our environment to be able to explain behaviour. When learning to drive a car, for instance, I have no need to solve Newton's Laws of motion or solve complex differential equation systems describing the physical forces operating on the car. Not only do I in fact not know the correct model, it seems that I don't even *need* to know it in order to achieve my objective. Clearly the design of the car has embodied some of the required physical knowledge in a system which I take for granted and within which I operate. My information set is then simply the available sensual input variables along with a desired target and presumably a fairly crude method of combining them together as an approximation to the physical forces operating on the car to produce an outcome which hopefully robustly gets me to were I want to go. The "*ad hoc*" rules of thumb that we use everyday when driving a car serve in place of a precise and most probably complex "*optimal*" decision rule based on a full knowledge of the relevant environment. When driving, the error between the target and the present state is feedback into an approximate physical model to determine action that invariably has some robustness characteristics in order to insulate outcomes from the consequences of the mis-specification inherent in using the approximate model and also from potentially unanticipated shocks [3].

Both these characteristics, robustness and the use of limited information approximations in decision making, seem to be elements reflected in many everyday economic actions and can be formally justified. In a slightly different context, Salmon and Young (1979) emphasised the distinction between good, bad and optimal decision rules stressed by Rosenbrock and McMorran (1971). The argument lies in the fact that many formally "optimal" control rules are blind to mis-specification and depend for "optimality" on a precise knowledge of the environment; the parallel with the

[2] Kent Wall (1993) has recently provided a clear statement of these issues and also an approach to modelling the adaptive nature of boundedly rational decision making.

[3] Alan Kirman, who clearly has at least an unbounded memory, has recently pointed out to me a passage , in Machlup(1962), that makes several similar points to those above regarding the need to compute and the meaning of marginal conditions in economics which also uses driving a car as an analogy

learning issue would seem to be exact. The practical electrical engineer responded to the apparent failure of such "optimal" rules by implementing control strategies that explicitly included integral or error correcting action, which is in general lacking in state feedback rules. A degree of integral action of appropriate form in a decision rule may accommodate an imperfect model of the environment and ensure that the objective is robustly achieved *despite* a lack of knowledge of the true structure. The implication for our present discussion of learning would seem to be that there is no need to necessarily consider the question of convergence of a learning scheme so as to provide an exact representation of reality (even if that were possible) as long as economic agents recognise that they are using mis-specified models and adjust their decision rules accordingly to take up the slack created by their lack of knowledge. The learning decision is then intimately tied with the choice of action implemented to achieve the ultimate economic objective and the relevant loss function for learning is not then simply the criterion function that evaluates the standard statistical loss but one that reflects the true economic cost of action under ignorance.

The contrast with the way that learning is often viewed in economics is stark since these issues seem to be generally ignored and attention is often focussed instead on the debatable question of convergence to a rational expectations equilibrium. Justifying the use of a particular learning model by emphasising the properties of the end result of learning is a substantively rational argument whereas models of learning that focus more the means by which individuals acquire and subsequently process information are procedurally rational. Closely connected with procedural learning would also then seem to be the question of use of approximate models.

2.2 Least Squares and Bayesian Learning

Let us now briefly consider how least squares and Bayes' rule, stand as procedural learning devices. There are two issues that need to be separated; that of how information is accrued and then subsequently how it is processed. The role of the economic environment and the agent's approximation/perception of that environment plays a critical role in both aspects.

In the case of least squares, the implicit model would seem to be that the economic agent passively receives new signals one after the other as time advances. The sequential nature of information accrual and processing follows from the standard statistical model in which new information is viewed as arriving sequentially, and effectively from a single source although parallel algorithms of information processing may in fact be more relevant to human learning. Psychological research suggests that we bring different bundles of information together about different aspects of a problem to form a coherent picture. Information accrual may therefore take place in parallel as particular events imply new information *simultaneously* about several different aspects of a decision problem. If the economic agent employed a multivariate model then in one sense data arrive from a number of different sources simultaneously but we would then also need to recognise the obvious limitations of least squares when there are informational feedbacks and potentially several endogenous variables in the model. A more sophisticated statistical estimation method that ensures consistency may be required. Again the role of the agent's model or perception of reality is potentially crucial in determining the mechanics of their learning.

One justification for least squares is of course that within a particular class of models (such as the classical linear regression model) it provides the most efficient linear unbiased estimator and so it would be rational to use it although outside this class it may not necessarily have particularly desirable properties. Similarly assuming limited computational and cognitive powers it might behaviourally be more reasonable to model economic agents as employing suboptimal or inefficient statistical algorithms for learning. Parallel information processing algorithms may be substantially more efficient *computationally* than sequential methods although we have to be careful to discriminate between numerical and statistical efficiency. The least squares estimator retains its relative statistical efficiency independently of how it is numerically calculated. Nevertheless it can certainly be the case, as several of the simulations below indicate, that least squares can be extremely slow to converge which may be unrealistic if it is to provide a reasonable model of how economic agents learn. From the procedural point of view economic agents might in these cases take deliberate action to speed up the process of learning.

In a market context there are a number of examples, such as for instance monopolists exploring their demand curve or search theory, where deliberate action may be taken to recover information and more general models of interaction and matching provide a clear behavioural basis for learning which would seem to deviate far from the passive statistical model implied by least squares. The issue again lies with the nature of the loss function. Least squares is not itself forward looking in the sense that the loss that is measured relates solely to the cumulative sum of current actions. Active experimentation is not seen to be beneficial in that the economic value of information is not reflected in the statistical loss function and a wider economic loss function has to be considered to induce active learning, see Easley and Kiefer (1988) Kiefer (1989) for instance. Given the discussion above regarding convergence it is interesting to note that again from this point of view it can be shown that full convergence to a rational expectations equilibrium may not be justified on economic grounds when the costs of extra information accrual outweigh the benefits of greater accuracy or efficiency, see McLennan (1984). Altering the information accrual process by active experimentation again emphasises the interaction between the context within which agents perceive themselves to be operating and their method of information processing. However least squares itself is blind to these issues and simply provides a potentially efficient information processing algorithm. One obvious question is whether human learning given limited cognitive abilities can be well described by an efficient statistical algorithm or not.

Perhaps more relevant to human learning is robustness to mis-specification rather than efficiency. From this point of view least squares forms a best approximation in the least squares sense given the assumed model structure. The standard assumptions justifying both least squares and Bayesian learning rest on unattractively strong informational assumptions to generate attractive theoretical results. When the true model is not assumed to be known both approaches can produce divergence. Nyarko (1991) for instance has shown that when the true model does not lie in the support of the prior beliefs assumed by a Bayesian monopolist who is trying to learn his demand function, his beliefs and consequently actions may oscillate in a never ending cycle. Similarly, Brousseau and Kirman (1993) with a symmetric duopoly find a similar conclusion using least squares learning with the implication that the economic agents drift endlessly without ever learning to believe in anything. Such cyclical divergence has also been found , for instance, by Radner (1982), Bray (1982) and Blume and Easley (1982). Examples of explosive divergence or convergence to incorrect models can also be relatively easily found in the literature. Do we believe that such behaviour is reflected in the real world or does the problem lie with the assumed model of learning? The theory of cognitive dissonance suggests that agents would reject this instability and instead find a stable even though imperfect approximation to reality on which they could base their decisions.

Another aspect of models of learning that implicitly rests on an assumption that the correct model is being estimated is the "turn off" phenomenon. Associated with convergence and consistency, whether it be to the true parameter values or not, is the fact that the impact of the innovations in a recursive least squares formulation of learning becomes less and less over time as the variances of the estimated parameters shrink to zero. So if agents either assume that they know the correct model or at least are using a fixed approximation then the least squares learning algorithm will, in general, gradually turn itself off as it converges. This result which indicates convergence in probability is inappropriate if the agent is in fact facing a non constant data generation process as is the case of course with the feedback of learning onto the economic structure or when the agent is continuously changing his approximate model. If the agent recognises that he is dealing with a non-constant approximation to the true data generation process then it might be appropriate to instead employ a constant learning rate rather than one that decreases over time. In this way the agent could retain the flexibility to learn continuously and respond to innovations in information that arose under such a non-constant data generation process. It is important to note that least squares learning implies that economic agents respond only to the current innovations and ignore other, "dynamic" properties of the innovation sequence such as the integral of past innovations or the rate of change of the innovations. It is in this sense that the least squares model of learning is blind; as mentioned above it does not look forward but also it does not allow the agents to look back and assess their performance. The reason for this is of course that the statistical method essentially assumes that the model being estimated coincides with the data generation process. It is not unreasonable to expect economic agents to make use of mis-specification checks to assess their performance but this is, at least strictly, outside the least squares learning process. If the agent detected serial correlation in the innovations, for instance, or some other dynamic pattern then a learning method that was able to automatically respond by switching to say generalised least squares or changing the model specification might describe human learning better given that the agent will be aware *ex ante* that his assumed model is only an approximation to the data generation process. A learning method that is able to respond to any structural information in the innovation sequence that goes beyond just the current observed innovation and considers potentially dynamic functions of the innovations and does not turn off may provide a better procedural model of human learning than that given by least squares.

Given these arguments we feel the procedural basis for least squares learning seems to rest on fairly thin ground both as a general model of information accrual and as a method of processing information. The model of an economic agent as a passive sequential information processor that relies on an assumption that he holds a correct model of his economic environment seems some distance from what might be required for boundedly rational learning. Moreover it is of course still open to question that economic agents might only be able to employ relatively simple suboptimal methods of information processing rather than the efficient least squares estimator. Human frailties and cognitive constraints may be more relevant in describing how individuals learn than how an efficient statistical algorithm should compute an estimate.

Bayesian learning, because of its wider probability framework can, in principle, relax some of the reservations regarding least squares learning. In particular it appears easier to integrate the statistical and economic decision problems. Nevertheless Bayes methods also rest on strong informational assumptions to deliver good results and can yet yield inconsistency, see Diaconis and Freedman (1988). The major difficulty with Bayesian learning seems to be that there is substantial evidence that Bayes theorem lacks empirical relevance and hence its procedural justification is weak. For instance Kahneman, Slovic and Tversky (1982) show that likelihoods and preferences expressed by individuals do not satisfy the coherence properties that are necessary for the existence of subjective probabilities. Moreover experimental subjects are found to update probabilities but they don't appear to use information efficiently in that posterior probabilities differ quantitatively and systematically from those predicted by the use of Bayes rule. cf. Nelson and Winter (1982), Edwards (1968), Tversky and Kahneman (1974), Cyert and De Groot (1987) . Grether (1980), (1992) also describes experiments in which it seems clear that individuals systematically fail to take "proper" account of prior probabilities. Apart from the systematic deviations from rationality (as defined by Bayes Rule) detected in psychological research the general lack of numeracy skills in most individuals that has been well documented more generally (see Paulos(1991)) seems to suggest that the computational burden of applying Bayes rule rigorously prevents its adoption as a reasonable procedural model of learning by most economic agents.

2.3 Neural Network Learning

There are alternative approaches to modelling learning that relieve at least some of the problems with the standard methods outlined above. We concentrate here on techniques that essentially provide non-parametric estimates of the unknown quantities which may be unknown parameters or more generally conditional expectation functions representing the boundedly rational expectation generation mechanism of economic agents. No assumption about the true model structure is required to implement these procedures, only information as to the set of relevant conditioning or input variables and of course, the output or target. In this way the information requirements correspond to those employed by Fourgeaud, Gourieroux and Pradel (1986) who emphasised that in general we do not need structural information to be able to generate rational expectations in a *linear* model. An estimate of the reduced form may be constructed from simply a knowledge of the conditioning exogenous variables in the system and the output variable. The neural network technique used below differs in the rejection of least squares as an information processing tool and the relaxation of the implicit assumption that the reduced form is linear in the conditioning variables [4]. A neural network essentially provides a time varying, flexible functional form that delivers a Minimum Mean Square Error approximation to an unknown conditional expectation function or target series at each point in time.

Artificial neural networks then represent sophisticated approximation devices whose structure is inspired by biological models of how the brain functions. Assuming a boundedly rational agent wishes to approximate his conditional expectation function

[4] Hal White has recently collected together a number of his important papers on Neural Networks in White(1992). Reference should also be made to Kuan and White (1992) for applications to learning problems.

for some variable y given a conditioning information set $\{x_i, i = 1 \ldots r\}$ the neural network will determines a particular function g, such that $y = g(x, \theta)$. The biological analogy comes from referring to the inputs as neurons which send signals, the values x_i, through a network architecture to the output y. These signals can be transformed at any stage, either at the output or at any intermediate stage when there are " hidden layers" in the network. The simplest such structure arises if the output is just a linear combination of the inputs so that $y = x'\beta$ and the network $g(x, \theta) = x'\beta$ is then just a linear regression.

One physical analogy is that certain neurons may remain dormant until "fired" sending an impulse to the output. Mathematically this can be represented as

$$y = F\left(x'\beta\right)$$

where

$$F(\lambda) = \begin{cases} 1 & \text{if } \lambda \geq \lambda_0 \\ 0 & \text{otherwise} \end{cases}$$

where λ_0 is some threshold value. This threshold logic unit may be replaced by any smoothly increasing function such as a normal or logistic cumulative density function giving rise to the familiar Probit and Logit models.

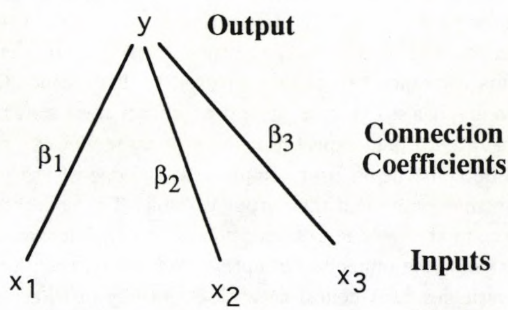

Figure 1: The simplest Neural Net: a linear regression

A further physical analogy and the really significant feature of neural networks comes through the introduction of a layer of hidden units that act effectively as unobserved state variables. In this case the inputs send their signals not directly to the output but to an intermediate level of processing units that function in exactly the same way as the output unit in transforming the signals that they receive. The parallel nature of information processing appears as each input potentially activates each hidden unit simultaneously which then in turn passes its transformed signal, in parallel, to the output. Figure 2 describes a simple single hidden layer feed forward neural network.

Such a network is termed feedforward as the information flow is in one direction from

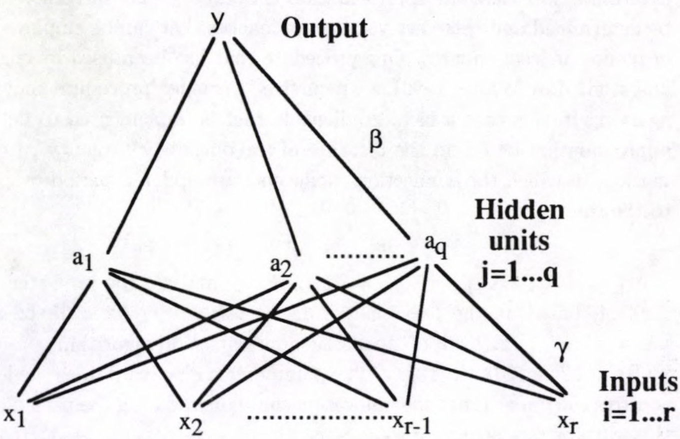

Figure 2: Neural Net with a Single Hidden Layer

the inputs towards the output. Substantially more complicated network structures with multiple layers with feedback as well as feedforward information flows can be used.

Each hidden unit produces a signal $a_j \quad j = 1 \ldots q$ where $a_j = \psi(x' \gamma_j)$ that is sent to the output as $a' \beta$. So we have

$$y = F(a'\beta) = F\left(\sum_{j=1}^{q} \psi(x'\gamma_j)\,\beta_j\right) = g(x, \theta)$$

where the parameters $\theta = (\beta', \gamma')'$ are known as connection strengths and the function $\psi(.)$ is generally taken to be some, so called, squashing function or sigmoid function such as

$$\psi(z) = (1 + \exp(-z))^{-1}$$

for each hidden unit.

Thus a neural network approximates the unknown mapping between the inputs and the target, y, by taking a potentially nonlinear combination of sigmoid functions as the basis of the approximating space. Other basis functions such as rational polynomials, Fourier series or spline functions are often used elsewhere in approximation theory. The distinguishing feature of Neural Nets is however the role of the hidden units that as state variables are defined endogenously so as to construct the best approximation to the target at each point in time given the inputs. The choice of spanning function and the number of hidden units are at the control of the agent and determine how close an approximation is obtained. The parameters θ still have to be determined and there are various approaches that can be employed including that of nonlinear least squares. One procedure that has been used by cognitive scientists and studied by White (1990) is a recursive "learning" procedure known as back propagation. In this case a local gradient descent is implemented to deliver the MMSE approximation based on the distance of the output or target, y , from the approximation, in which the connection coefficients are updated period by period according to the rule.

$$\hat{\theta}_n = \hat{\theta}_{n-1} + \eta \nabla g_n (y_n - g_n)$$

where $g_n \equiv g_n(\hat{\theta}_{n-1})$ and$\nabla g_n = \nabla g(x_n, \hat{\theta}_{n-1})$ and are the predicted output and the gradient based on the previous parameter estimates. As such the method of back propagation is one form of stochastic approximation algorithm considered originally by Robbins and Munro(1951). The parameter η is referred to as the learning rate and controls the impact that the innovations have on updating the parameter estimates. If this learning rate is fixed at some constant value the network has the continuous ability to pick-up shifts in structure whereas if it decreases as in recursive least squares with iid observations at a rate of n^{-1} then , as described above, eventually the innovations will have no impact on revising the parameter estimates and the turn off problem will appear. Back propagation is statistically inefficient relative to nonlinear least squares and susceptible to several other difficulties including converging to a local rather than a global minimum however despite these shortcomings it describes precisely how we feel economic agents who know that they have an approximate model of reality will learn. Moreover it is computationally simpler to implement than least squares and perhaps more in line with cognitive constraints. For these reasons we have used back propagation in the examples below with a fixed learning rate.

Since the neural net algorithm constructs a time varying nonlinear approximation to the unknown conditional expectation function it is quite different from the least squares or Bayesian algorithms that essentially assume there is a fixed constant structure characterised by the value of some subset of parameters which must be learnt. As the neural net does not assume there is a single invariant structure to be found it overcomes the fundamental contradiction with the use of the least squares learning. This contradiction arises as least squares rests on the assumption that the data generation process is stationary but the self referential nature of learning on the model creates nonstationarity in the data generation process even if the underlying model without learning is stationarity . The neural net approach to learning and expectation formation is free of such ambiguities.

There are also a number of potential theoretical advantages from using neural networks, in the first place they provide the ability to combine information from distinct sources (separate models) and process this information in a parallel fashion. One possible advantage is then in the speed of convergence of learning. In addition, given the non-parametric nature they can be very robust to poor information and when implemented sequentially can track with relative ease a non- constant data generation process. However, the main reason why they are of interest in the present context is the powerful theoretical results that are available demonstrating their ability to approximate unknown functions. For instance Hornik, Stinchcombe and White (1989) (1990) have recently shown, under certain relatively weak regularity conditions, that a particular neural network can approximate any measurable function to any required accuracy, see also White (1989) (1990). Thus while avoiding the need to specify the true model structure, these methods have the potential to gain convergence to the true rational expectations equilibrium! Moreover as discussed in these references the ability to approximate sufficiently accurately rests on the network structure itself evolving as information is built up. Thus economic agents who employ neural networks to learn and approximate their conditional expectation functions can continuously update their approximation to reality until the economic trade offs might indicate that further accuracy was unnecessary. Despite this powerful approximation ability we have only explored fixed structure networks in the examples below and hence have generated only partial approximations to the unknown expectation functions.

The justification of the method as a procedural learning method can be seen, quite simply, to be a sensible procedure for agents to learn in the absence of knowledge of the correct model structure. It rests on minimal assumptions and is robust through providing a time varying nonlinear approximation to the true conditional expectation function which in principle it could arbitrarily well approximate.

3 Examples

3.1 Monetary Policy and Inflation

The presumption that the private sector's expectation of government policy may critically affect the impact of that policy has been central to the study of macroeconomic policy for many years. At the heart of this issue lies the credibility of the policymaker in the eyes of the private sector since it is believed the opportunities open to a policy maker with high credibility may differ substantially from those available to the same policy maker with little credibility. Credibility itself is thought to be enhanced when expectations are confirmed by the policy maker's actions and so how the private sector learns and adjusts its expectations is central to an analysis of credibility. Invariably in this literature the private sector is assumed to be able to form fully rational expectations and the policy implications of credibility and reputation building have been based on this assumption. A number of papers (in particular, Barro and Gordon (1983), Barro (1986), Backus and Driffill (1985a,b), Cukierman and Meltzer (1986) and Basar and Salmon (1990a, b)) have considered the question in some detail[5]. In particular in Basar and Salmon (1990a) it was emphasised how the natural asymmetry in the information pattern facing the policymaker and the private sector allows the policy maker, when acting as a Stackelberg leader, to actively intervene in the learning process of the private sector and hence directly influence the formation of (rational) expectations by the private sector. In this way the government is able to choose its policy taking full account of the effect that this policy will have on the private sector and hence also on the evolution of its own credibility. The significance of the Stackelberg solution is that the policy optimisation problem facing the government must explicitly recognise the dual role of the optimal policy in that its informational impact on the private sector's uncertainty about the government's preferences must be traded off against the immediate effect on the government's welfare. In this way the Government's credibility and the resulting inflationary bias evolve from a substantially more complex, non-certainty equivalent problem than is the case with the steady state Nash (and certainty equivalent) solution provided by Cukierman and Meltzer (1986). The endogeneity of reputation building in this example, through the closing of the information loop, serves to sustain the zero inflationary bias result, which somewhat surprisingly given the Stackelberg structure, is credible when only stagewise precommitment by the policy maker is possible.

This earlier analysis rested, however, on the unrealistic assumption that both parties, the government and the aggregate private sector, hold common and correct beliefs (if not observable information) regarding the true structure of the economy. Given this full knowledge of the probability distributions describing the structure of the economy the problem facing the private sector when forming its rational expectation is simply a classic example of signal extraction. We now consider the private sector to be uncertain not only about the government's preferences but also about the structure of the process by which the government's preferences are determined. In particular we start by considering the private sector learning both in the sense of updating its estimate (or rational expectation) of the unobserved state variable (the

[5]The following example is based in part on Basar and Salmon (1990b).

government's preference parameter) and also by learning any unknown parameters in the model through which these preferences are determined[6]. As in a standard adaptive control problem with a single decision maker the fact that there are unknown parameters will usually induce nonlinearity and a non-certainty equivalent control problem for which the optimal policy will reflect the non-separability of the prediction and control aspects of the policy (see Goodwin and Sin (1984)). In what follows we ignore this aspect of active learning and therefore do not attempt to develop the optimal policy that would reflect the full impact of the private sector's parametric learning.

We employ the same model as Basar and Salmon (1990a,b) which is a finite horizon version of a dynamic "Barro and Gordon" model developed by Cukierman and Meltzer (1986). We also restrict attention here to the Nash (or discretionary) solution since it is substantially easier in this case to study the impact of parametric learning on the evolution of credibility and the inflationary bias of monetary policy.

We first derive the Nash Solution for this model under disparate beliefs without learning so that the private sector in fact holds mistaken views on the way the government's preferences are determined [7] and consider the evolution of the government's credibility and the ensuing inflationary bias. In constructing its optimal policy the government is assumed to be fully aware of this situation. We then consider the effect of the private sector recognising its ignorance and employing the two learning schemes discussed above, least squares and neural network learning. Under least squares learning the private sector updates its estimates of the unknown parameters (determining the governments preferences) while simultaneously forming rational expectations of these preferences. The neural network learning procedure approximates the entire conditional expectation function for future monetary policy given observations on past money supply.

The Nash Solution without Learning

We consider a model in which the monetary authority wishes to maximise a finite horizon objective function given by

$$J_1 = E\{\sum_{i=0}^{N} \beta^i (e_i x_i - \frac{1}{2}(m_i^p)^2)\} \tag{3}$$

[6]Apart from Basar and Salmon (1990b), Martin Cripps (1991) has also considered parameter learning in the Nash solution to a static version of the Cukierman and Meltzer model

[7]Alternatively this situation could correspond to a regime change in which the private sector's beliefs reflect the previous regime.

with respect to its choice of monetary growth m_i^p, where $e_i = m_i - \delta_i(I_i)$ represents the surprise in inflation (or alternatively monetary growth) determined by the private sector's rational expectation $E[m_i \mid I_i] = \delta_i(I_i)$ given the observable information $I_i = \{m_{i-1}, m_{i-2}, \ldots m_0\}$ and x_i is a preference parameter of the monetary authority that trades off the benefit from the growth induced by surprise inflation with the cost of a non zero level of inflation. The preferences of the monetary authority are assumed to evolve according to a first order autoregressive process given by,

$$x_i = \rho\, x_{i-1} + A\,(1-\rho) + \nu_i \tag{4}$$

where ν_i follows a normal distribution with mean zero and variance σ_ν^2. The private sector is assumed to be able to observe only the actual monetary growth rate which deviates from the planned money supply process, m_i^p, given some imperfect monetary control. In particular it observes m_i where

$$m_i = m_i^p + \psi_i \tag{5}$$

and ψ_i is a random disturbance assumed to follow a normal distribution with zero mean and variance σ_ψ^2. Since the private sector does not know the policy preference parameter, x_i, it is faced with a natural signal extraction problem when forming its expectations. We now assume that the private sector believes the government's preferences evolve according to equation (4) but parameterised by (ρ_2, A_2) whereas the true preferences are determined by (ρ_1, A_1) and, for the moment, there is no learning of the unknown parameters, ρ and A by the private sector. As shown below the optimal Nash solution to this policy problem (in which the only role of the private sector is to form rational expectations of the future course of monetary policy) is affine in the preference parameter,

$$m_i^p = M_i\, x_i + k_i \tag{6}$$

In which case the rational expectation of the private sector will be generated by the conditional expectation;

$$\delta_i^* = M_i \hat{x}_{i|i-1} + k_i \tag{7}$$

where the conditional expectation $\hat{x}_{i|i-1}$ is generated by the Kalman Filter;

$$\hat{x}_{i+1|i} = \rho_2\, \hat{x}_{i|i-1} + \rho_2\, K_i\,(m_i - M_i \hat{x}_{i|i-1} - k_i) + A_2(1-\rho_2) \tag{8i}$$

$$K_i = \frac{M_i \sum_{i|i-1}}{M_i^2 \sum_{i|i-1} + \sigma_\psi^2} \tag{8ii}$$

$$\sum\nolimits_{i|i-1} = \frac{\rho_2^2\, \sigma_\psi^2 \sum_{i|i-1}}{M_i^2 \sum_{i|i-1} + \sigma_\psi^2} + \sigma_\nu^2 \tag{8iii}$$

with

$$\sum\nolimits_{0|-1} = \sigma^2_{x_0} \tag{8iv}$$

Given this characterisation of the private sector's prediction problem the monetary authority solves for the optimal Nash decision rule by maximising J_1, given by

$$J_1 = E\{\sum_{i=0}^{N} \beta^i[(\psi_i + m_i^p)x_i - M_i\,\hat{x}_{i|i-1}x_i - k_i x_i - \frac{1}{2}(m_i^p)^2]\}$$

or

$$J_1 = E\{\sum_{i=0}^{N} \beta^i[m_i^p\,x_i - M_i\,\hat{x}_{i|i-1}x_i - \frac{1}{2}(m_i^p)^2]\}$$

where the non constant part may be rewritten as

$$E\{(m_N^p x_N - \frac{1}{2}(m_N^p)^2)\beta^N + \sum_{i=0}^{N-1} \beta^i[m_i^p\,x_i - \frac{1}{2}(m_i^p)^2\beta M_{i+1}\hat{x}_{i+1|i}x_{i+1}]\}$$

At the final stage $i = N$, the optimal policy can then simply be seen to be,

$$m_N^p = x_N$$

Moving back one period, $i = N-1$, the optimisation problem becomes

$$\max_{m_{N-1}^p} E\{m_{N-1}^p x_{N-1} - \frac{1}{2}(m_{N-1}^p)^2 - \beta c_N \hat{x}_{N|N-1}x_N \mid x_{N-1}\}$$

where

$$\begin{aligned}\hat{x}_{N|N-1} = \rho_2(1 - K_{N-1}M_{N-1})\hat{x}_{N-1|N-2} \\ + \rho_2 K_{N-1}(m_{N-1}^p - k_{N-1}) \\ + A_2(1-\rho_2) + \rho_2 K_{N-1}\psi_{N-1}\end{aligned}$$

and $c_N = M_N$

and so we find

$$m_{N-1}^p = \underbrace{(1 - \beta c_N \rho_2 \rho_1 K_{N-1}}_{M_{N-1}} x_{N-1} - \underbrace{\beta c_N \rho_2 K_{N-1} A_1(1-\rho_1)}_{k_1}$$

Proceeding in this way we find the optimal monetary policy at stage N-2 from,

$$\max_{m_{N-2}^p} E\{m_{N-2}^p x_{N-2} - \frac{1}{2}(m_{N-2}^p)^2 - \beta c_{N-1}\hat{x}_{N-1|N-2}x_{N-1} - \beta d_{N-1}\hat{x}_{N-1|N-2}\}$$

where

$$c_{N-1} = \rho_1\rho_2 c_N(1 - K_{N-1}M_{N-1})\beta + M_{N-1}$$

$$d_{N-1} = \beta\rho_2 c_N(1 - K_{N-1}M_{N-1})A_1 + (1 - \rho_1)$$

leading to an optimal policy rule for period $N - 2$ of the form;

$$m^p_{N-2} = M_{N-2}x_{N-2} + k_{N-2}$$

where

$$M_{N-2} = 1 - \beta\rho_1\rho_2 c_{N-1}K_{N-1}$$

and

$$k_{N-2} = -\beta\rho_2 A_1(1 - \rho_1)c_{N-1}K_{N-2} - \beta d_{N-1}\rho_2 K_{N-2}$$

Continuing in this way we find a general form for the Nash policy rule to be given by the following set of equations;

$$m^p_n = M_n x_n + k_n$$

where

$$M_n = (1 - \rho_1\rho_2\beta c_{n+1}K_n)$$
$$k_n = \beta\rho_2(d_{n+1} + A_1(1 - \rho_1)c_{n+1})K_n$$
$$c_n = \beta\rho_1\rho_2(1 - K_n M_n)c_{n+1} + M_n$$
$$d_n = \beta\rho_2(1 - K_n M_n)(c_{n+1}A_1(1 - \rho_1) + d_{n+1})$$

subject to

$$c_N = M_N = 1,\ \ d_N = k_N = 0$$

together with the Kalman filter equations $(8(i),.,8(iv))$ given above . Essentially what is required is the solution to

$$1 - M_n = \frac{\rho_1\rho_2 c_{n+1}M_n\Sigma_{n|n-1}}{M_n^2\Sigma_{n|n-1} + \sigma^2_\psi}$$

for M_n at each stage, n, from 1 to N. It can easily be shown that this finite horizon solution coincides in steady state and under common beliefs $(\rho_1 = \rho_2, A_1 = A_2)$ with the solution given for the infinite horizon problem by Cukierman and Meltzer (1986).

The effect of the incorrect beliefs held by the private sector in this Nash game is explored by simulation since it is not straightforward to unravel analytically the complex interactions of all the parameters on the quantities of interest. Figures 3 and 4 show the evolution of inflationary bias and credibility for the various runs described in Table One.

A simple Gauss Seidel iterative scheme was employed to ensure convergence of the policy parameters M_i, and k_i and those of the prediction or signal extraction problem, K_i, and Σ_i, over a 25 period simulation run. Credibility at each point, given

Table 1:
Simulations for Nash Solution without Learning

Run	ρ_1	ρ_2	β	σ_ν^2	σ_ψ^2	$\sigma_{x_0}^2$	A_1	A_2
1	0.9	0.9	0.95	1	1	1	1	1
2	"	0.5	"	"	"	"	"	"
3	0.5	"	"	"	"	"	"	"
4	0.1	0.1	"	"	"	"	"	"
5	"	0.9	"	"	"	"	"	"
6	0.9	0.1	"	"	"	"	"	"
7	"	0.9	"	0.1	"	"	"	"
8	"	"	"	1	5	"	"	"
9	"	"	"	"	1	0.1	"	"
10	"	"	0.1	"	"	"	"	"

Table 2:
Simulations for Nash Solution with Least Squares Learning

Run	NP	ρ_1	ρ_2	β	σ_ν^2	σ_ψ^2	$\sigma_{x_0}^2$	A_1	A_2	α
1	81	0.9	0.9	0.95	1	1	1	1	1	1
2	202	"	0.5	"	"	"	"	"	"	0.8
3	40	"	0.1	"	"	"	"	"	"	1
4	25	"	0.5	"	0.5	"	"	"	"	0.8
5	55	"	"	"	1	0.5	"	"	"	0.8
6	19	"	"	"	"	1	0.5	"	"	0.8
7	14	"	"	"	0.5	0.5	"	"	"	0.8
8	13	"	"	0.1	"	"	"	"	"	0.8

NP - The number of iterations with SOR (Successive Over Relaxation) iteration before convergence with relaxation factor α.

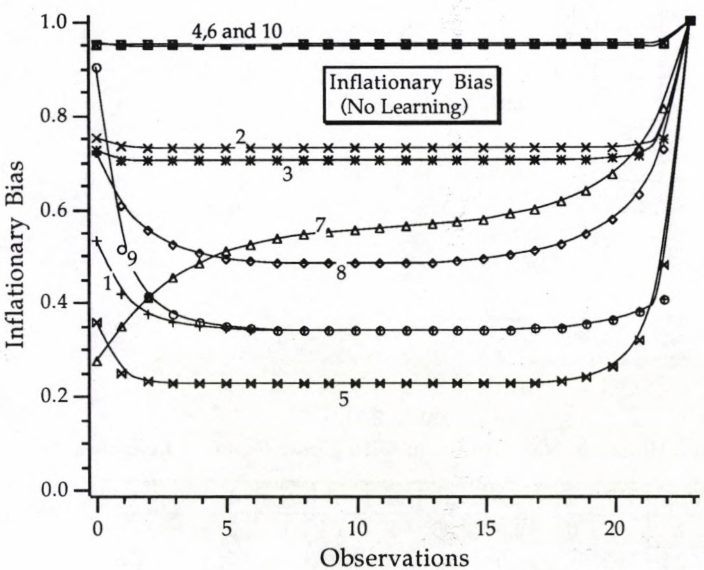

Figure 3: Inflationary Bias without LS Learning

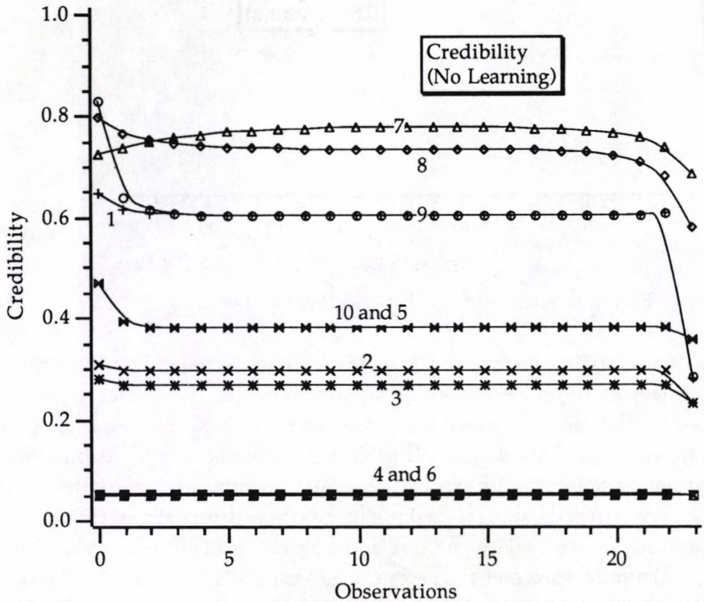

Figure 4: Credibility without Learning

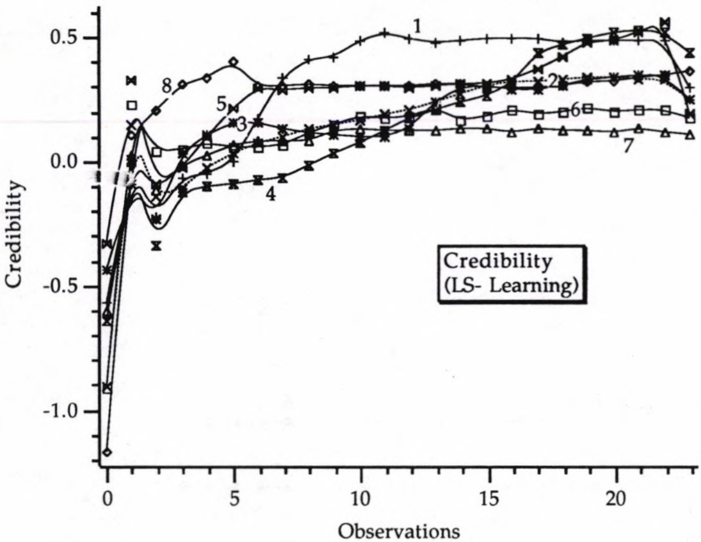

Figure 5: Credibility with Least Squares Learning

by λ_i, is defined in a similar manner to the steady state measure used by Cukierman and Meltzer where $\lambda_i = \rho_2(1 - K_i M_i)$. Basically credibility is thought of as being determined by the speed by which the private sector learns. The sequence $\{\lambda_i\}$ converges in steady state to the weight λ that the private sector uses to discount previous observations when forming its expectations. So it essentially measures the speed with which the private sector recognises shifts in government preferences, the higher λ the less important are recent developments in the formation of current expectations compared with the prior view and the lower the government's credibility[8]. Figure 4 shows how λ_i evolves over the simulation period. Inflationary bias is determined as $b_i = M_i A_2 + k_i$, being the unconditional expectation of planned monetary growth in each period.

As can be seen from these figures there is a common general pattern to the evolution of credibility and inflationary bias under most of the alternative parameter settings. Inflationary bias falls initially for several periods before attaining a long run base level before rising to unity as the end of the policy interval approaches. Similarly but less dramatically the government's credibility is typically seen to fall

[8]There are important differences between this steady state view of credibility and the transient analysis of credibility in Basar and Salmon (1990a,b)

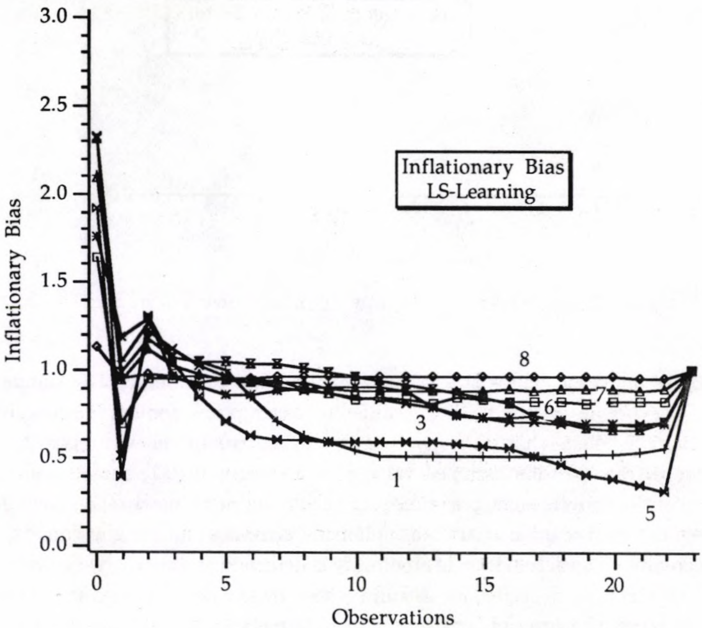

Figure 6: Inflationary Bias under Least Squares Learning

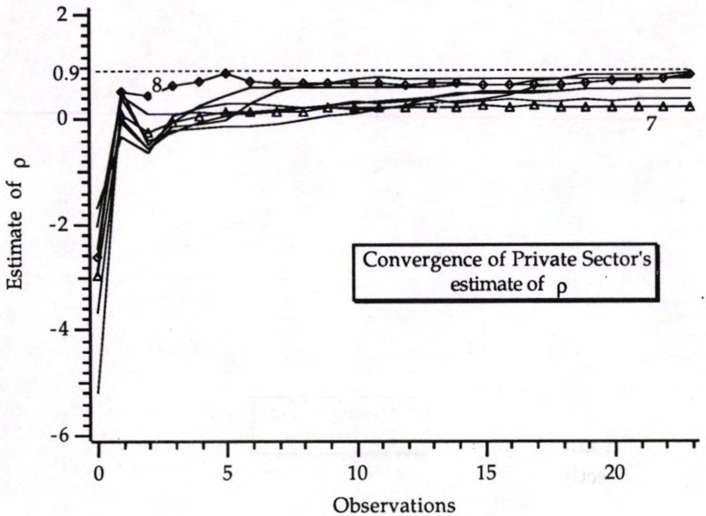

Figure 7: Convergence of private sector's estimate of ρ

initially, stabilise and then fall as the policy interval comes to an end. To examine the effects of the private sector holding a different assumption about the value of ρ, Run 1 (common beliefs that $\rho_1 = \rho_2 = 0.9$) may be compared with runs 2, 5 and 6. As the private sector's assumed value of ρ decreases to 0.1, runs 2 and 6, implying a belief that government preferences are changing more slowly than they in fact are, we see the level of inflationary bias uniformly increases and the government's credibility decreases. The actual level of credibility is determined directly by the value of ρ assumed by the private sector, ρ_2, as can be seen by the definition given above. The impact of ρ_2 on the level of inflationary bias is more subtle and results from its effect on the corresponding optimal policy rule. Run 5 reverses the manner of misperception of ρ in that the true value is 0.1 where the private sector incorrectly believes that the government's preferences are changing more rapidly ($\rho_2 = 0.9$). In this case although there is no direct impact, credibility is substantially reduced when compared with Run 1, however the level of inflationary bias is also reduced reflecting again the optimal policy when $\rho = 0.1$. Run 4 may then be used to observe the effect of a low common belief that $\rho = 0.1$ and it can be seen together with the results of Run 6 that the level of inflationary bias is critically determined by the assumed value of ρ_2. If the private sector believes the government's preferences are changing slowly, whether correctly or not, then inflationary bias will be high and the level of credibility low. Runs 7, 8 and 9 may be compared with Run 1 (all have common beliefs with $\rho = 0.9$) to examine the effect of changing the degree of uncertainty by altering the variances associated with the various stochastic terms in the model. Perhaps the most interesting observation here results from reducing the variance on the uncertainty about the government's preferences by a factor of 10 which leads to a steady growth in inflationary bias over the period but with a consistently high level of government credibility. Similarly increasing the degree of noise in observing the planned monetary growth rate leads to a relatively high level of credibility but with an average level of inflationary bias. Finally from Run 10, we can see the effect of the government discounting the future substantially more rapidly than in the base run (Run 1). In this case an average level of credibility is associated with a very high level of inflationary bias.

These results indicate the sensitivity of existing policy conclusions regarding the role of credibility to the assumption that the private sector is able to form rational expectations and has full knowledge of the economic structure and the relevant parameters. As shown in these simulations the level of credibility and the resulting inflationary bias is directly determined by the private sector's misperception of the true structure.

Least Squares Learning

We now assume that the private sector recognises the non-zero mean in the innovation of monetary growth ,that arises from its misperception of ρ, and from that infers that its view of ρ is inconsistent with the true value. The private sector is by assumption only able to observe the actual monetary growth when forming expectations and hence also when updating its estimate of ρ. It is assumed however to be aware of the structural form of the model which forms the basis of its ability to learn using least squares. In particular it knows that,

$$m_i = m_i^p + \psi_i$$

and that

$$m_i^p = M_i x_i + k_i \tag{9}$$

hence

$$m_i = M_i x_i + k_i + \psi_i \tag{10}$$

It is also assumed to be aware of the structure of the state equation generating the government's preferences (except for the value of ρ) and so equation (10) may be rewritten

$$m_i = M_i(\rho x_{i-1} + A(1-\rho) + \nu_i) + k_i + \psi_i$$

Rearranging this expression we find that

$$m_i - AM_i - k_i = \rho[M_i(x_{i-1} - A)] + M_i\nu_i + \psi_i$$

which represents a standard regression relationship of the form

$$y_i = \rho z_i + u_i \tag{11}$$

where

$$y_i = m_i - AM_i - k_i$$

and

$$z_i = M_i(x_{i-1} - A)$$

The only complication with this formulation lies in that the private sector is unable to observe the state variable x_{i-1} but they can form an estimate $\hat{x}_{i-1|i-1}$ essentially by the Kalman Filter as described by equations (8(i)-(iv)). Substituting this estimate we find an operational algorithm in terms of the constructed variable $\hat{z}_i = M_i(\hat{x}_{i-1|i-1} - A)$[9]. Applying the standard recursive least squares algorithm to (10) we have

$$\hat{\rho}_i = \hat{\rho}_{i-1} + \frac{P_{i-1}\hat{z}_i}{\sigma_u^2 + P_{i-1}\hat{z}_i^2}[y_i - \hat{z}_i\hat{\rho}_{i-1}]$$

$$P_i = P_{i-1} - \frac{P_{i-1}^2\hat{z}_i^2}{\sigma_u^2 + P_{i-1}\hat{z}_i^2}$$

where $\sigma_u^2 = M_i^2\sigma_\nu^2 + \sigma_\psi^2$

Figures 5, 6 and 7 describe a number of simulations to investigate the effect of least squares learning on the previously determined Nash policy. The plots in these figures indicate a representative stochastic run given one particular drawing from a normal random number generator for the stochastic terms in the model. In general, convergence was found to be substantially more difficult to achieve than in the previous experiments without learning. A relaxed Gauss Seidel procedure was needed in many cases to achieve convergence of the policy and prediction parameters. The dominant observation from these figures is how dramatically the policy optimisation has been affected by the dynamics of the learning behaviour of the private sector. Given this we can see the same general pattern of a low level of credibility being associated with a high level of inflationary bias. Two aspects of the learning behaviour seem to be particularly important. In the first place even when the private sector

[9] As mentioned above we are ignoring the obvious statistical feedbacks between the state estimation and parameter estimation. It is also clearly stretching rationality pretty far to assume with this least squares assumption that the private sector have sufficient knowledge to construct y_i and $\hat{z}_i$ from knowledge of m_i, M_i and k_i which are in turn based on estimates of ρ

has a correct prior on the value of ρ in terms of its mean but a large degree of initial uncertainty about this value, the learning procedure may actually drive the private sector *away* from its correct initial beliefs. This seems paradoxical but may well reflect legitimate behaviour when there is substantial initial uncertainty about government preferences. Secondly it can be seen from figure 7 that even after some 25 periods the learning scheme generally has a long way to go until convergence to the correct value of ρ. This extremely slow learning would seem to have important implications for government policy. The implicit time unit that has been used corresponds to the period between policy changes in which case it would seem to be very unwise, as is the case with standard policy analysis, for the government to treat the private sector as if it knew the true structure of the economic environment. However Figures 5 and 6 clearly show the critical impact that private sector's (parameter) learning has on the optimal policy in this model. No substantial difference is found between the different runs by the end of the period compared with the case of no learning. One effect of including learning has been to remove the majority of the difference in credibility created by the different assumptions regarding the private sectors beliefs. Ultimately the level of inflationary bias is unity with credibility falling only slightly near the end of the period. The level of inflationary bias , where comparable, seems to be higher than in the case without learning.

Neural Network Learning

To implement neural net learning and expectation formation by the private sector, a single hidden layer feed-forward architecture was set up to estimate the conditional expectation $E[m_t \mid I_t]$. Since in this model the private sector only observes actual money supply the objective was set to estimate m_t based on the first three lags of money supply. This can then be viewed as a non-linear, time-varying, autoregressive form for the conditional expectation function which, depending on interpretation, may be a solved out forward looking expectation process or a non-linear adaptive process. From the strategic point of view, it seems that the policy maker should treat the neural net expectation as exogenous since although it might be reasonable to expect that there is some feedback it is not obvious how it would itself approximate the private sector's approximation. The resulting optimal monetary policy is then determined each period by the value of the government's preference parameter, so that $m_i^p = x_i$. The neural net approximation is then in effect estimating the value of the government's preference parameter on the basis of past monetary policy. Since no structure is assumed to be known by the private sector and the Kalman filter is not employed we are not in a position to use the previous measure of credibility. One alternative is simply to measure the monetary surprise since it indicates how the private sector's expectations, however generated, have deviated from actual outcomes. The root mean squared error between the actual money supply and the respective expectation then measures the cumulative surprise. One set of results are shown in

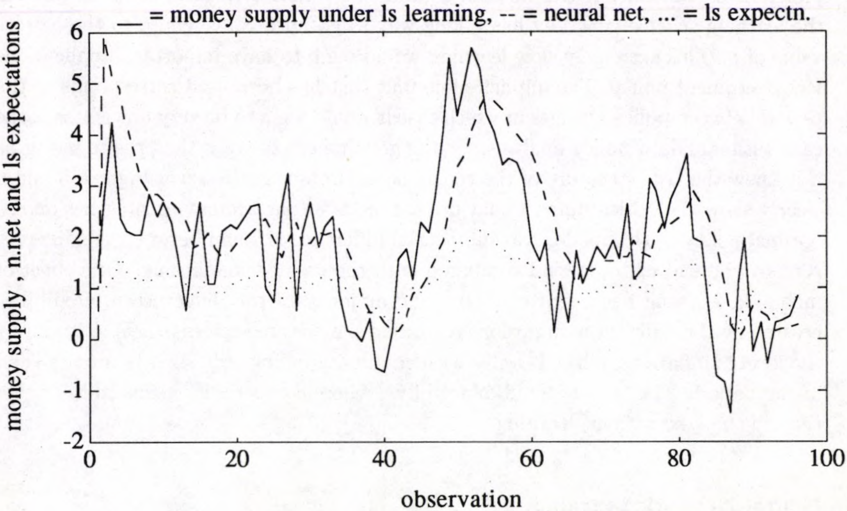

Figure 8: Least Squares and Neural Net Expectations of money supply

Figure 8 in which a 100 period simulation of the model was run with least squares learning under the "mis-specified" assumptions of Run 7 from table 2. The solid line is the actual evolution of the observed money supply determined under the least squares learning assumption; the dashed line - the neural net expectation and the dotted line, the conditional expectation estimated on the basis of the least squares learning algorithm. In this particular numerical comparison the neural net approximation delivers a root mean square error of 11.22 which is *smaller* than the least squares RMSE of 13.67 despite the fact that the least squares based expectation exploited all the (correct) knowledge about the model structure except the correct value of ρ. The least squares based expectation completely fails to capture the degree of movement in the money supply series and is generally sluggish in responding to changes in direction although near the end of the sample its relative performance improves. This simulation is obviously based on particular drawings of the random variables but the general pattern of behaviour seems to be consistent across a number of simulations that we have examined. It is of course not invariably the case that the RMSE for the neural net is lower than that for the least squares expectation, in fact the least squares estimate tends to dominate more often than not, however this example does indicate that a private sector with neural net expectations could be surprised less overall than if it had formed a "rational" least squares expectation. The implication for the scope of exploiting credibility in policy is then presumably that a policy maker should be very wary about basing his actions on any perceived opportunities to surprise the private sector until he was absolutely certain what the surprise would actually be!

The comparison given above in Figure 8 has to be interpreted carefully since the neural net expectation is calculated on the basis of a monetary policy that assumed least squares expectation formation by the private sector. Figure 9, on the other hand, shows the observed money supply following from optimal monetary policy under both neural net and least squares expectations as well as the neural net expectation calculated under the "neural net monetary policy". As can be seen from this

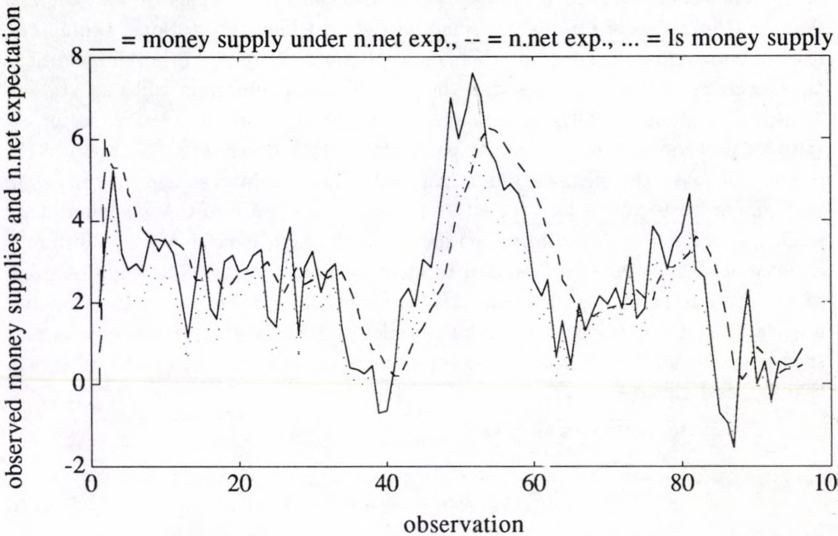

Figure 9: Optimal monetary policy under Least Squares and Neural Net Expectations

figure both policy profiles are essentially identical but differ only by a fairly constant degree with the neural net policy being more active, perhaps reflecting less ability to surprise. In fact as the sample comes closer to the end both policies tend to come together. The same shape can also therefore be seen in the neural net expectation which is now calculated on the basis of the optimal monetary policy which assumed neural net expectations. The fact that the two monetary policies are so similar in this case reflects the dynamics of the coefficients in the policy rule, eqn (9), and the dominate role, under both assumptions, of the evolution of the policy maker's preference parameter. This similarity is clearly not likely to be reflected in other problems with different structures but does indicate that the assumptions required by least squares learning may be irrelevant in practice if individuals and the private sector as a whole can act so as to approximate the government's policy mapping sufficiently accurately. The strategic implications for exploiting credibility when the private sector does not hold or learn rational expectations are perhaps more fundamental.

3.2 Learning and the Dynamics of Hyperinflation

Marcet and Sargent (1989) considered the impact of least squares learning in a model of hyperinflation analysed previously under rational expectations by Sargent and Wallace (1987). Under rational expectations the model has two stationary equilibria, one with low inflation and the other, high inflation. Only if the initial conditions are chosen to precisely coincide with the low inflation equilibrium will that stationary state be obtained. Otherwise the high inflation equilibrium is the attractor despite its perverse comparative static properties in that a permanent increase in the deficit will *lower* the stationary inflation rate. However Marcet and Sargent show that under least squares learning either there is convergence to the *low* inflationary stationary state or no equilibrium exists. Thus the high inflationary equilibrium is eliminated when agents are assumed to learn by least squares. The additional dynamics introduced through least squares learning thus completely alter the economic implications of the model; an obvious question is then how these conclusions may themselves be affected if economic agents are assumed to learn in some other way than by least squares.

The model is written as follows

$$P_t = \lambda E_t P_{t+1} + \gamma h_t \qquad 0 < \lambda < 1, \gamma > 0 \tag{12i}$$

and

$$h_t = \theta h_{t-1} + \xi P_t \qquad 1 \leq \theta \leq \lambda^{-1}, \xi > 0 \tag{12ii}$$

h_0 is given and $P_t, h_t > 0$.

P_t represents the price level, h_t per capita money and $E_t\, P_{t+1}$ the private sector's expectation of the price level at time $t+1$ based on information available at time t. Essentially equation (12ii) states that the government finances a constant per capita real deficit of ξ through money creation.

Under Rational expectations we impose

$$E_t P_{t+1} = \beta_t P_t \tag{13}$$

where $\beta_t = {}^{P_{t+1}}/_{P_t}$ is the gross inflation rate.

Substituting this expression into the model yields the relationship

$$\beta_{t+1} = (\lambda^{-1} + \theta - \xi\gamma\lambda^{-1}) - \theta\lambda^{-1}/\beta_t \tag{14}$$

and a Rational Expectations equilibrium for $\{P_t, h_t\}_{t=0}^{\infty}$ is given by the sequence $\{\beta_t\}_{t=0}^{\infty}$ that satisfies (14); there are two stationary equilibria $\beta_1^* < \beta_2^*$ provided the deficit satisfies $\xi \prec \xi_{\max} = \frac{\lambda}{\gamma}\left[\theta + \lambda^{-1} - 2(\theta\lambda^{-1})^{1/2}\right]$. Only if the initial inflation value is chosen so that $\beta_0 = \beta_1^*$, the low inflation stationary state under rational expectations, is that equilibrium found; for any other value $\beta_0 \in [\beta_1^*, \lambda^{-1})$ convergence to β_2^* follows.

Marcet and Sargent proceed by allowing the private sector to learn the value of β_t by least squares using the available price history . So that given (13)

$$\beta_t = \left[\sum_{s=0}^{t-1} P_{s-1}^2\right]^{-1}\left[\sum_{s=0}^{t-1} P_s P_{s-1}\right]$$

or when written recursively

$$\beta_t = \beta_{t-1} + \frac{1}{t}R_{t-1}^{-1}P_{t-2}\left[P_{t-1} - \beta_{t-1}P_{t-2}\right] R_t = R_{t-1} + \frac{1}{t}\left[P_{t-1}^2 - R_{t-1}\right]$$

The following dynamic relations then describe the model

$$P_t = \frac{\gamma}{1-\lambda\beta_t}h_t \tag{15}$$

$$h_t = S(\beta_t)h_{t-1} \tag{16}$$

$$P_t = \left[S\left(\beta_t\right)\frac{1-\lambda\beta_{t-1}}{1-\lambda\beta_t}\right]P_{t-1} \tag{17}$$

where the mapping

$$S(\beta_t) = \frac{(1-\lambda\beta_t)\,\theta}{(1-\lambda\beta_t-\gamma\xi)} \tag{18}$$

effectively takes the perceived law of motion for P_t into an actual law of motion for P_t. A stationary equilibrium is then given by a fixed point of $S(\beta)$ where $S(\beta)-\beta = 0$ and it can be seen that there are two solutions given again by β_1^* and β_2^*. Starting from a given set of initial conditions (β_0, R_0, h_0) we can then examine the convergence of inflation under least squares learning.

Notice immediately an inconsistency with the application of least squares learning in this problem, and indeed in general. Equation (17) is essentially the relationship to which least squares is applied once the feedback in the data generation process is recognised but it is also clearly non-constant until the fixed point itself is reached. So while learning is taking place the assumption of a constant data generation process on which least squares rests is violated.

Marcet and Sargent show both analytically and through simulation that under least squares learning;

- β_1^* is it locally stable iff $K \equiv \frac{(1-\beta_1^{*-2})\lambda\beta_1^*}{(1-\lambda\beta_1^*)} < 1$
- β_2^* is always unstable and hence β_1^* is the only candidate for a limit point.
- If β_t ever becomes larger than β_2^* then β_t will be pushed over a singularity at $\frac{(1-\lambda\xi)}{\lambda}$ into a region with no equilibrium. In other words no equilibrium exists for $\beta_t > \frac{(1-\lambda\xi)}{\lambda}$.

We introduce neural network expectation formation and learning into the model by assuming that the conditional expectation for the future price level is some unknown but also not necessarily constant function of the current price to be approximated by the neural net. Writing the neural net expectation as

$$[E_t P_{t+1}] = f(constant, P_t) \tag{19}$$

the model becomes

$$P_t = \lambda[E_t P_{t+1}] + \gamma h_t \tag{20}$$

$$h_t = \theta h_{t-1} + \xi P_t$$

and so prices will be generated by

$$P_t = \frac{\theta}{1-\gamma\xi} P_{t-1} + \frac{\lambda}{1-\gamma\xi} [[E_t P_{t+1}] - \theta[E_t P_{t+1}]_{-1}] \tag{21}$$

The neural net expectation in (20) was constructed in a similar manner as in the previous example with a single hidden layer network but with two inputs, a constant and the current price level. Unlike the previous example we have tried to mimic forward looking behaviour by the private sector by estimating the neural net approximation in a loop that enables us to use the previous iteration's value for P_{t+1} as the target for the approximation based on the current price. This process was iterated until convergence and hence consistency is found. In addition, since the neural net expectation is a function of the current price, which it also determines, a simultaneity problem exists that is again resolved by iteration until a consistent value for the current price level is employed on both sides of equation (21). This iterative procedure within each time period mirrors the solution for the forward looking rational expectation. The

Table 1: Inflation at 500 iterations under Least Squares and Neural Net Learning

β_0	ξ	Least Squares	Neural Net
1.0	0.0019	1.02519	1.0114
1.0	0.00234	1.03572	1.0133
1.0	0.0024	0.39945*	1.0135
1.0	0.003	0.72960*	1.0158
1.02	0.0019	1.02519	1.0114
1.02	0.00234	1.03572	1.0133
1.02	0.0024	0.74100*	1.0135
1.02	0.003	0.00610*	1.0158
1.0376	0.0019	1.02519	1.0114
1.0376	0.00234	1.03572	1.0133
1.0376	0.0024	1.03769	1.0135
1.0376	0.003	−0.00680*	1.0158
1.5	0.0019	1.02519*	1.0114
1.5	0.00234	0.44180*	1.0133
1.5	0.0024	0.96360*	1.0135
1.5	0.003	0.45940*	1.0158
2.0	0.0019	1.02519*	1.0114
2.0	0.00234	0.18350*	1.0133
2.0	0.0024	−0.02990*	1.0135
2.0	0.003	0.46918*	1.0158

The * indicates non-convergence and negative prices.

question then is whether or not this time varying non-linear dynamic data generation process implies an inflation rate β_t that converges and if so to what value? Since the model structure including the neural net expectation process is now different from the original model with rational expectations it is perhaps not surprising that we do not find convergence to the rational expectations equilibria β_1^*, β_2^*. Given the difficulty of studying the stationary points of the mapping under the neural network assumption analytically a number of simulations were carried out and are reported in the following table.

Several aspects of Table 1 are worth noting. In the first place it is clear that the neural net learning scheme has *always* converged regardless of the initial conditions but not to either of the stationary points found under rational expectations or least squares learning. In fact convergence is also obtained (but not shown in

Table 2: Stationary Inflation Rates under R.E. and Least Squares Learning

ξ	β_1^*	β_2^*	β_{max}
0.0019	1.02519	1.08381	1.109
0.00234	1.03572	1.07279	1.1085
0.0024	1.03769	1.07076	1.1084

the table) to the same values for a wide range of simulations starting from widely divergent initial inflation rates and deficits. The domain of attraction to the relevant equilibrium therefore seems to be substantially wider under neural network learning than under least squares learning. This robustness to changes in initial conditions offers a strong behavioural argument in favour of the neural network learning. Table 2 below shows the stationary values under least squares and rational expectations and the fact that the neural network algorithm does not converge to these values could be interpreted in several different ways; either as an indication that an element of mis-specification has been introduced into the model through the assumption of what may be considered an *ad hoc* learning rule or alternatively, if the neural net approximation is justified as a reasonable procedural assumption as to how agents learn then since the structure of the model under learning has been changed, a different but equally justified equilibrium has been robustly identified which is relevant to that new structure. Under the assumptions of bounded rationality that imply that learning is required in the first place then perhaps we should not be so concerned about failing to converge to the rational expectations equilibrium in the long run. The fact that the neural network learning rule has not converged to the rational expectation equilibrium under the original model structure should also not necessarily cause any concern following our earlier arguments since if the equilibrium found is recognised as inferior in some way agents could adopt decision rules that included integral action to reflect the divergence.

Notice also that the neural net equilibrium inflation rates are also uniformly lower than the low inflation equilibrium under least squares learning. The least squares learning rule is clearly very sensitive and its performance critically determined by the initial conditions and the value of the deficit. Column 3 of Table 1 essentially reproduces the same results found by Marcet and Sargent and the sensitivity and divergence of least squares learning is in marked contrast to the results for the neural network learning. Table 2 also shows the value of $\beta = \beta_{max}$ where there is a singularity and beyond which no equilibrium exists under least squares learning. Plots of the simulations are interesting particularly in those cases in which inflation passes through

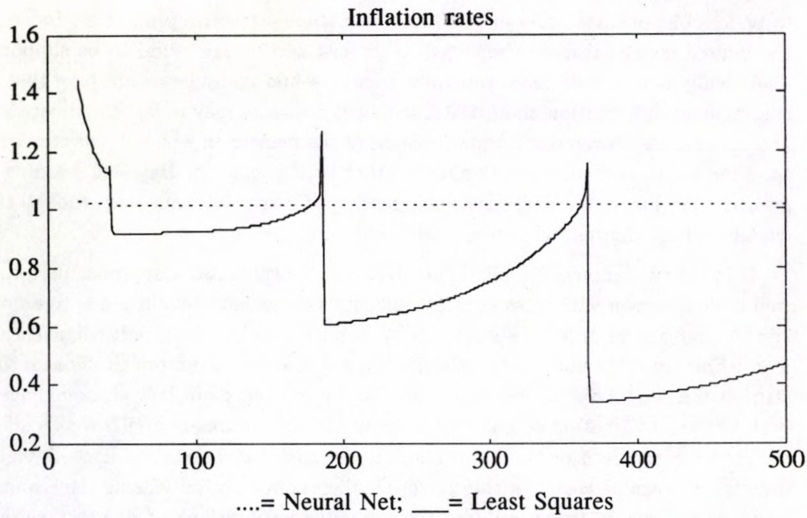

Figure 10: Least Squares and Neural Net Inflation

the singularity since although locally explosive least squares learning often appears to attempt to return towards the equilibrium value. However once having passed through the singularity inflation takes a value typically far away from the equilibrium and subsequent adjustment is extremely slow. It was also found in the simulations that the neural network was quite generally substantially faster to achieve convergence than least squares. Figure 10., which shows inflation rates corresponding to the case with $\beta_0 = 2.0$ and $\xi = 0.003$ provides one example of such behaviour as least squares passes through the singularity several times and where the robustness of neural net learning can be clearly contrasted with the sensitivity of least squares learning. Surely on behavioural grounds we would be unlikely to accept such sensitivity as reasonable in practice which would further question the value of least squares as a procedural model of learning by economic agents.

4 Conclusions

It is clear that much needs to be done to examine the question of boundedly rational learning schemes in general but from the experience we have gained and reported above the neural network approximation approach suggested here appears at least to be feasible, behaviourally justified and have some attractive properties. Indeed, if the limited empirical results reported in the last section are found to be supported analytically and to hold more generally, then it would really be astonishing that assuming more information about the true model structure may in fact be unnecessary and in some cases even detrimental because of the manner in which it was employed. In other words, seriously questioning whether least squares or Bayesian learning algorithms are likely to reflect reasonable models of human learning even though they exploit full information regarding model structure.

In the two empirical applications above the neural network approach performed well in comparison with least squares learning. One general conclusion is to emphasise the dangers of drawing specific policy implications based on particular learning algorithms given the potential sensitivity to those learning assumptions. This in turn, stresses the need for a greater understanding by economists of how economic agents with limited information actually learn about their environment and how this affects the nature of the decision rules they adopt under bounded rationality. Embedding the question of learning clearly within the original economic decision facing the economic agent would enable us to evaluate the *economic* implications of different learning schemes under the relevant loss function and may help to provide a clear behavioural basis for the manner by which learning takes place.

Market and social interaction are obvious mechanisms by which information is accrued and it does not seem clear that the standard statistical regression model adequately captures the manner by which such information is captured in the first place and subsequently processed. Behavioural models of learning under bounded rationality are being developed; Mordecai Kurz (1990),(1991),(1992), for instance, has recently emphasised the analysis of rational belief equilibria rather than rational expectations equilibria. The critical distinction is, much as argued above, that all agents are permitted to hold different beliefs or approximate models of reality which are nevertheless all consistent with the observed data. No structural assumptions need be made in this model of learning and each agent updates their beliefs in a manner which is consistent with their own perception of reality. The neural network learning procedure has several of these features since it assumes agents explicitly hold an approximate model and also that their learning mechanism is robust to misspecification since it makes few assumptions as to the true data generation process. Models of "social learning" , "herd effects" or learning from others have also been developed, see for instance Kirman(1992), Ellison and Fudenberg (1992), Topol(1991) and Vives(1992), which describe different market or social mechanisms that influence the information accrual and/or the information processing aspect of learning. Such behavioural models of learning which focus on economic incentives and the process by which learning takes place under bounded rationality may eventually provide some answers as to how economic agents actually *do* behave rather than how economists believe they *should* behave.

References

[1] **Alchian A.A.** (1950) "Uncertainty, evolution and economic theory" *Journal of Political Economy*, 58, pp. 211-21.

[2] **Arrow K.J.** (1986) "Rationality of self and others in an economic system", in *Rational Choice; the contrast between economics and psychology*, R. M. Hogarth and M.W.Reder eds.,Chicago University Press.

[3] **Backus D. and Driffill J.** (1985a) "Inflation and Reputation" *American Economic Review*, 75,1, pp. 530-538.

[4] **Backus D. and Driffill J.** (1985b) "Rational Expectations and Policy Credibility following a Change in Regime" *The Review of Economic Studies*, 52, pp. 211-21.

[5] **Barro R.** (1986) "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, 17, pp. 3–20.

[6] **Barro R. and Gordon D.** (1983) "Rules, discretion and reputation in a model of monetary policy" *Journal of Monetary Economics*, 12 pp. 101-122.

[7] **Basar T. and Salmon M.** (1988) "On the Convergence of Beliefs and Policy to a Rational Expectations Equilibrium in a Dual Policy Problem" in *Issues in Monetary Economics*, W. Marois, A. Soubeyran and D. Laussel Eds. Springer Verlag.

[8] **Basar T. and Salmon M.** (1990a) "Credibility and the value of information transmission in a model of monetary policy and inflation" *Journal of Economic Dynamics and Control*, 14, pp. 97–116.

[9] **Basar T. and Salmon M.** (1990b) "Inflation and the Evolution of Credibility with Disparate Beliefs" in *Dynamic Models and the Control of National Economics*, N. M. Christodoulakis ed, Pergamon Press.

[10] **Baumol W.J. and R.E. Quandt** (1964) "Rules of Thumb and Optimally Imperfect Decision Rules" *American Economic Review*, 54, pp. 23-46.

[11] **Blume L. and Easley D.** (1982) "Learning to Be Rational" *Journal of Economic Theory*, 26, pp. 340-51.

[12] **Bray M.** (1982) "Learning, Estimation and the Stability of Rational Expectations" *Journal of Economic Theory*, 26, pp. 318-39.

[13] **Bray M. and Kreps D. M.** (1987) "Rational Learning and Rational Expectations" in G. R. Feiwel (ed), *Arrow and the Ascent of Modern Economic Theory*, New York University Press.

[14] **Bray M. M. and Savin N. E.** (1986) "Rational Expectations Equilibria, Learning, and Model Specification" *Econometrica*, 54:5, pp. 1129-1160.

[15] **Brousseau V. and Kirman A.** (1993) "The Dynamics of Learning in N-Person Games with the Wrong N" in Binmore K.,Kirman A. and Tani P. (eds), *Frontiers of Game Theory*,MIT Press, Cambridge Massachusetts.

[16] **Crain W.M., Shughait W.F. and Tollison R.D.** (1984) "The convergence of satisficing to marginalism: an empirical test " *Journal of Economic Behaviour and Organisation*, 5, pp. 375-385.

[17] **Cripps M.** (1991) "Learning rational expectations in a policy game" *Journal of Economic Dynamics and Control*, 15, pp. 297–315.

[18] **Cukierman A. and Meltzer A.** (1986) "A theory of ambiguity, credibility, and inflation under discretion and asymmetric information" *Econometrica*, September 54:5, pp. 1009-1128.

[19] **Cyert R. M. and DeGroot M.** (1987) *Bayesian Analysis and Uncertainty in Economic Theory*, Rowman and Littlefield.

[20] **Day R.H.** (1967) "Profits, learning and the convergence of satisficing to marginalism " *Quarterly Journal of Economics*, 81, May, pp. 302-311.

[21] **Diaconis P. and Freedman D.** (1988) "On Inconsistent Bayes Estimates of Location" *The Annals of Statistics*,14,1, pp. 68-87.

[22] **Easley D. and Kiefer N.** (1988) "Controlling a Stochastic Process with unknown parameters" *Econometrica*, 56:5, pp. 1065-1084.

[23] **Edwards W.** (1968) "Conservation in Human Information Processing" in *Formal Representation of Human Judgement*, B. Kleinmuntz (ed), Wiley, New York.

[24] **Ellison G. and Fudenberg D.** (1993) "Rules of Thumb for Social Learning ", *Journal of Political Economy*, vol 101, pp. 612–43.

[25] **Favero C. and Hendry D.** (1992) "Testing the Lucas Critique: A Review", *Econometric Reviews*, 11(3), pp. 265–306.

[26] **Festinger L.** (1957) *A Theory of Cognitive Dissonance*, Stanford University Press.

[27] **Fourgeaud C., Gourieroux C. and Pradel J.** (1986) "Learning Procedures and Convergence to Rationality", *Econometrica*, 54, pp. 845-868.

[28] **Goldberg D. E.** (1989) *Genetic Algorithms, Search, Optimization and Matching Learning*, Reading, Addison-Wesley.

[29] **Goodwin G. and Sin K. S.** (1984) *Adaptive Filtering, Prediction and Control*, Prentice Hall.

[30] **Grether D.M.** (1980) "Bayes rule as a descriptive model: the representativeness heuristic", *Quarterly Journal of Economics*, 95, pp. 537-557.

[31] **Grether D.M.** (1992) "Testing Bayes rule and the representativeness heuristic: some experimental evidence", *Journal of Economic Behaviour and Organisation*, 17, pp. 31-57.

[32] **Hogarth R. M. and Reder M. W.** (1986) *Rational Choice: The Contrast Between Economics and Psychology*, University of Chicago Press.

[33] **Hornik K., Stinchcombe M. and White H.** (1989) "Multilayer Feedforward Networks are Universal Approximators", *Neural Networks*, 2, pp. 359-366.

[34] **Hornik K., Stinchcombe M. and White H.** (1990) "Universal Approximation of an Unknown Mapping and Its Derivatives Using Multilayer Feedforward Networks", *Neural Networks*, 3, pp. 551-560.

[35] **Kahneman D., Slovic P. and Tversky A** (1982) *Judgement under Uncertainty: Heuristics and Biases*, CUP.

[36] **Kiefer N.** (1989) "Optimal Collection of Information by Partially Informed Agents", *Econometric Reviews*, 7:2, pp. 113-148.

[37] **Kirman A.** (1993) "Ants, Rationality and Recuitment", *Quarterly Journal of Economics*,Feb,pp. 137-156.

[38] **Kuan Chung-Ming and H.White.** (1992) "Adaptive Learning with Nonlinear Dynamics driven by Dependent Processes ", *Working Paper, Department of Economics, University of California at San Diego.*

[39] **Kurz M.** (1990) "On the Structure and Diversity of Rational Beliefs", *Working Paper, Department of Economics, Stanford University.*

[40] **Kurz M.** (1991) "On Rational Belief Equilibria ", *Working Paper, Department of Economics, Stanford University.*

[41] **Kurz M.** (1992) "Asset prices with Rational Beliefs ", *Working Paper, Department of Economics, Stanford University.*

[42] **Machlup F.** (1967) *Essays in Economic Semantics*, Norton, New York.

[43] **Marcet A. and T. Sargent** (1989) "Least Squares Learning and the Dynamics of Hyperinflation" in W.Barnett, J.Geweke and K.Shell (eds), *Economic Complexity*, Cambridge University Press.

[44] **McLennan A.** (1984) "Price Dispersion and Incomplete Learning in the Long Run", *Journal of Economic Dynamics and Control*, 7, pp. 331-347.

[45] **Nelson R. and Winter, S.**(1982) *An Evolutionary Theory of Economic Change*, Belknap Press.

[46] **Nyarko Yaw** (1991) "Learning in Mis-specified Models and the Possibility of Cycles",*Journal of Economic Theory*,55, pp. 416–427.

[47] **Paulos J.A.** (1988) *Innumeracy; Mathematical Illiteracy and its Consequences* Viking, Penguin Books Ltd., Harmondsworth, England.

[48] **Radner R.** (1975) "A behavioural model of cost reduction", *Bell Journal of Economics*, 6, pp. 196-215.

[49] **Radner R.** (1982) "Equilibrium under Uncertainty" in K. J. Arrow and M. D. Intrilligator (eds), *Handbook of Mathematical Economics II,* Amsterdam, North Holland.

[50] **Robbins H. and Munro S.** (1951) "A Stochastic Approximation Method", *The Annals of Mathematical Statistics*, 22, pp. 400-407.

[51] **Rosenbrock H.H. and McMorran P.D.** (1971) "Good, Bad or Optimal", *IEEE Trans.Aut. Control*, vol AC. 16, no.6.

[52] **Salmon M. and Young P.** (1979) "Control Methods and Quantitative Economic Policy" in S. Holly , B. Rustem and M.Zarrop (eds), *Optimal Control for Econometric Models* , Macmillan, London.

[53] **Sargent T. and N.Wallace** (1987) "Inflation and the Government Budget Constaint " in A.Razin and E.Sadka (eds), *Economic Policy in Theory and Practice* Macmillan, London.

[54] **Simon H.** (1955) "A Behavioral Model of Rational Choice", *Quarterly Journal of Economics*, 69, pp. 99-118.

[55] **Simon H.** (1959) "Theories of Decision Making in Economics and Behavioral Science", *American Economic Review*, 49:1, pp. 253-283.

[56] **Simon H.** (1976) "From Substantive to Procedural Rationality" in *Method and Appraisal in Economics*, S. J. Latsis (ed), CUP. reprinted in *Models of Bounded Rationality* MIT Press, 1982.

[57] **Simon H.** (1982) *Models of Bounded Rationality*, MIT Press.

[58] **Slovic P. and Lichtenstein S.** (1983) "Preference reversals: a broader perspective", *American Economic Review*, 73, pp. 596–605.

[59] **Topol R.** (1991) "Bubbles and Volatility of Stock Prices: Effect of Mimetic Contagion", *The Economic Journal*, 101, pp. 786–800.

[60] **Tversky A. and Kahneman D.** (1974) "Judgement under Uncertainty " *Science* 185,pp.1124–30.

[61] **Tversky A. and Kahneman D.** (1986) "Rational Choice and the Framing of Decisions " in *Rational Choice*, edited by Hogarth and Reder (op cit).

[62] **Tversky A., Slovic P. and Kahneman D.** (1991) "The Causes of Preference Reversal", *American Economic Review*, 80:1, pp. 204-217.

[63] **Vives X.** (1992) "Learning from Others ", *Working Paper, Institut d'Analisi Economica, Unversitat Autonoma de Barcelona.*

[64] **Wall K.** (1993) "A Model of Decision Making under Bounded Rationality", *Journal of Economic Behaviour and Organisation*, 21, pp. 331-352.

[65] **White H.** (1989) "Some Asymptotic Results for Learning in Single Hidden Layer Feedforward Network Models", *Journal of the American Statistical Association*, 84:408, pp. 1003–1013, reprinted in White et al. (1992).

[66] **White H.** (1990) "Connectionist Non-Parametric Regression: Multilayer Feedforward Networks Can Learn Arbitrary Mappings", *Neural Networks*, 3, pp. 535-549, reprinted in White et al. (1992).

[67] **White H.** (1992), *Artificial Neural Networks; Approximation and Learning Theory*, Basil Blackwell, Oxford.

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

Economics Department Working Paper Series

To Department of Economics **WP**
European University Institute
Badia Fiesolana
I-50016 San Domenico di Fiesole (FI)
Italy

From Name . . . .

Address. . . . .

. . . .

. . . .

. . . .

. . . .

(Please print)

❒ Please enter/confirm my name on EUI Economics Dept. Mailing List
❒ 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 1994/95

Please send me the following EUI **ECO** 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 1993

ECO No. 93/1
Carlo GRILLENZONI
Forecasting Unstable and Non-Stationary Time Series

ECO No. 93/2
Carlo GRILLENZONI
Multilinear Models for Nonlinear Time Series

ECO No. 93/3
Ronald M. HARSTAD/Louis PHLIPS
Futures Market Contracting When You Don't Know Who the Optimists Are

ECO No. 93/4
Alan KIRMAN/Louis PHLIPS
Empirical Studies of Product Markets

ECO No. 93/5
Grayham E. MIZON
Empirical Analysis of Time Series: Illustrations with Simulated Data

ECO No. 93/6
Tilman EHRBECK
Optimally Combining Individual Forecasts From Panel Data

ECO NO. 93/7
Víctor GÓMEZ/Agustín MARAVALL
Initializing the Kalman Filter with Incompletely Specified Initial Conditions

ECO No. 93/8
Frederic PALOMINO
Informed Speculation: Small Markets Against Large Markets

ECO NO. 93/9
Stephen MARTIN
Beyond Prices Versus Quantities

ECO No. 93/10
José María LABEAGA/Angel LÓPEZ
A Flexible Demand System and VAT Simulations from Spanish Microdata

ECO No. 93/11
Maozu LU/Grayham E. MIZON
The Encompassing Principle and Specification Tests

ECO No. 93/12
Louis PHLIPS/Peter MØLLGAARD
Oil Stocks as a Squeeze Preventing Mechanism: Is Self-Regulation Possible?

ECO No. 93/13
Pieter HASEKAMP
Disinflation Policy and Credibility: The Role of Conventions

ECO No. 93/14
Louis PHLIPS
Price Leadership and Conscious Parallelism: A Survey

ECO No. 93/15
Agustín MARAVALL
Short-Term Analysis of Macroeconomic Time Series

ECO No. 93/16
Philip Hans FRANSES/Niels HALDRUP
The Effects of Additive Outliers on Tests for Unit Roots and Cointegration

ECO No. 93/17
Fabio CANOVA/Jane MARRINAN
Predicting Excess Returns in Financial Markets

ECO No. 93/18
Iñigo HERGUERA
Exchange Rate Fluctuations, Market Structure and the Pass-through Relationship

ECO No. 93/19
Agustín MARAVALL
Use and Misuse of Unobserved Components in Economic Forecasting

ECO No. 93/20
Torben HOLVAD/Jens Leth HOUGAARD
Measuring Technical Input Efficiency for Similar Production Units: A Survey of the Non-Parametric Approach

ECO No. 93/21
Stephen MARTIN/Louis PHLIPS
Product Differentiation, Market Structure and Exchange Rate Passthrough

ECO No 93/22
F. CANOVA/M. FINN/A. R. PAGAN
Evaluating a Real Business Cycle Model

ECO No 93/23
Fabio CANOVA
Statistical Inference in Calibrated Models

ECO No 93/24
Gilles TEYSSIÈRE
Matching Processes in the Labour Market in Marseilles. An Econometric Study

ECO No 93/25
Fabio CANOVA
Sources and Propagation of International Business Cycles: Common Shocks or Transmission?

ECO No. 93/26
Marco BECHT/Carlos RAMÍREZ
Financial Capitalism in Pre-World War I Germany: The Role of the Universal Banks in the Financing of German Mining Companies 1906-1912

ECO No. 93/27
Isabelle MARET
Two Parametric Models of Demand, Structure of Market Demand from Heterogeneity

ECO No. 93/28
Stephen MARTIN
Vertical Product Differentiation, Intra-industry Trade, and Infant Industry Protection

ECO No. 93/29
J. Humberto LOPEZ
Testing for Unit Roots with the k-th Autocorrelation Coefficient

ECO No. 93/30
Paola VALBONESI
Modelling Interactions Between State and Private Sector in a "Previously" Centrally Planned Economy

ECO No. 93/31
Enrique ALBEROLA ILA/J. Humberto LOPEZ/Vicente ORTS RIOS
An Application of the Kalman Filter to the Spanish Experience in a Target Zone (1989-92)

ECO No. 93/32
Fabio CANOVA/Morten O. RAVN
International Consumption Risk Sharing

ECO No. 93/33
Morten Overgaard RAVN
International Business Cycles: How much can Standard Theory Account for?

ECO No. 93/34
Agustín MARAVALL
Unobserved Components in Economic Time Series

ECO No. 93/35
Sheila MARNIE/John MICKLEWRIGHT
"Poverty in Pre-Reform Uzbekistan: What do Official Data Really Reveal?"

ECO No. 93/36
Torben HOLVAD/Jens Leth HOUGAARD
Measuring Technical Input Efficiency for Similar Production Units: 80 Danish Hospitals

ECO No. 93/37
Grayham E. MIZON
A Simple Message for Autocorrelation Correctors: DON'T

ECO No. 93/38
Barbara BOEHNLEIN
The Impact of Product Differentiation on Collusive Equilibria and Multimarket Contact

ECO No. 93/39
H. Peter MØLLGAARD
Bargaining and Efficiency in a Speculative Forward Market

***

ECO No. 94/1
Robert WALDMANN
Cooperatives With Privately Optimal Price Indexed Debt Increase Membership When Demand Increases

ECO No. 94/2
Tilman EHRBECK/Robert WALDMANN
Can Forecasters' Motives Explain Rejection of the Rational Expectations Hypothesis?

ECO No. 94/3
Alessandra PELLONI
Public Policy in a Two Sector Model of Endogenous Growth

ECO No. 94/4
David F. HENDRY
On the Interactions of Unit Roots and Exogeneity

ECO No. 94/5
Bernadette GOVAERTS/David F. HENDRY/Jean-François RICHARD
Encompassing in Stationary Linear Dynamic Models

ECO No. 94/6
Luigi ERMINI/Dongkoo CHANG
Testing the Joint Hypothesis of Rationality and Neutrality under Seasonal Cointegration: The Case of Korea

ECO No. 94/7
Gabriele FIORENTINI/Agustín MARAVALL
Unobserved Components in ARCH Models: An Application to Seasonal Adjustment

ECO No. 94/8
Niels HALDRUP/Mark SALMON
Polynomially Cointegrated Systems and their Representations: A Synthesis

ECO No. 94/9
Mariusz TAMBORSKI
Currency Option Pricing with Stochastic Interest Rates and Transaction Costs: A Theoretical Model

ECO No. 94/10
Mariusz TAMBORSKI
Are Standard Deviations Implied in Currency Option Prices Good Predictors of Future Exchange Rate Volatility?

ECO No. 94/11
John MICKLEWRIGHT/Gyula NAGY
How Does the Hungarian Unemployment Insurance System Really Work?

ECO No. 94/12
Frank CRITCHLEY/Paul MARRIOTT/Mark SALMON
An Elementary Account of Amari's Expected Geometry

ECO No. 94/13
Domenico Junior MARCHETTI
Procyclical Productivity, Externalities and Labor Hoarding: A Reexamination of Evidence from U.S. Manufacturing

ECO No. 94/14
Giovanni NERO
A Structural Model of Intra-European Airline Competition

ECO No. 94/15
Stephen MARTIN
Oligopoly Limit Pricing: Strategic Substitutes, Strategic Complements

ECO No. 94/16
Ed HOPKINS
Learning and Evolution in a Heterogeneous Population

ECO No. 94/17
Berthold HERRENDORF
Seigniorage, Optimal Taxation, and Time Consistency: A Review

ECO No. 94/18
Frederic PALOMINO
Noise Trading in Small Markets

ECO No. 94/19
Alexander SCHRADER
Vertical Foreclosure, Tax Spinning and Oil Taxation in Oligopoly

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