LEARNING HYPERINFLATIONS

Atanas Christev

Department of Economics, Heriot-Watt University
Edinburgh EH14 4AS, UK e-mail: a.christev@hw.ac.uk

PRELIMINARY, PLEASE DO NOT QUOTE!

This Version: December 2005

Abstract:

Emprical studies of hyperinflations reveal that the rational expectations hypothesis fails to hold. To address this issue, we study a model of hyperinflation and learning in an attempt to better understand the volatility in movements of expectations, money, and prices. The findings surprisingly imply that the dynamics under neural network learning appear to support the outcome achieved under least squares learning reported in the earlier literature. Relaxing the assumption that inflationary expectations are rational, however, is essential since it improves the fit of the model to actual data from episodes of severe hyperinflation. Simulations provide ample evidence that if equilibrium in the model exists, then the inflation rate converges to the low inflation rational expectations equilibrium. This suggests a classical result: a permanent increase in the government deficit raises the stationary inflation rate (Marcet and Sargent, 1989).

JEL classification: C62 E63 E65

Key Words: Hyperinflation, Learning, Rational Expectations Equlibria, Neural Networks

1. The Main Questions

In the United States and Western Europe, it is generally accepted that new money creation (seigniorage) is not typically resorted to as a source of public finance (Plosser and King, 1985). This is in stark contrast with the recent experience in some East European and Former Soviet Union (FSU) countries (Miller and Zhang, 1997). With the collapse of the command system and with the deterioration of the balance sheets in state-owned enterprises, governments were forced to seek new sources of revenue. Reform delays in most countries and slow progress toward privatization in some countries created further obstacles to improvement in the fiscal position of governments. In periods of such upheaval, the absence of strong independence of the central bank meant that weak governments could easily resort to extensive money creation. The results have been very high and variable rates of inflation, especially in the initial stages of transition. In some instances, hyperinflation has occurred. Such episodes are the subject of this paper.

Hyperinflations are economic anomalies in the sense that variations in nominal and monetary variables almost completely outweigh variations in real variables (Cagan, 1956, p. 25). Thus, such episodes provide opportunities for testing of long-standing hypotheses in monetary theory, especially regarding the linkage between money and prices. They thus also provide natural laboratories for studying inflationary expectations and policy regime changes. The economic theory of hyperinflation was stimulated by the observations of the dramatic historical episodes witnessed over the past century.[1] There is little disagreement that movements in expected inflation account principally for the

[1] These include the classic hyperinflations of the interwar period: Austria, Germany, Hungary, Poland, and Russia. Other examples involve Hungary and China during and after WWII; some Latin American countries in the 1980s, viz., Argentina, Bolivia, Brazil, Peru; Israel in the mid-1980s; and most recently Yugoslavia, Russia, Ukraine, Georgia and Bulgaria.

changes in demand for money during hyperinflations. However, the process of stabilizing inflation is more controversial. A traditional view holds that stopping inflation necessarily requires that money creation be severely restrained. The classic essay by Cagan (1956) elaborates on this point of view.[2] Sargent (1993, p. 45) extends these earlier ideas to include the need for a *credible* fiscal restraint that becomes "sufficiently binding to be widely believed."[3] His main message implies that hyperinflations end only when accompanied by sudden, credible policy change by the government. When the government is unable to finance the deficit through means other than money creation, high seigniorage is exacted and inflation soars. Notwithstanding these contributions, the enduring questions remain: what causes hyperinflations; what stops them; are expectations rational during episodes of dramatic increases in money and prices, and during regime shifts; can agents in hyperinflationary environments "learn to stabilize?"

We use a standard overlapping generations model with money in which the deficit is financed through seigniorage to study some of these questions. Much like Cagan's findings (1956), the model results in two equilibria with stationary inflation levels (high and low). We next introduce two alternative learning mechanisms to explain better the stylized facts of economies during hyperinflation. This work follows, among others,

---

[2] Also of note is Bresciani-Turroni (1937, p. 355), who calls stabilization a "miraculous event," observing "The stabilisation of the German exchange showed, as did that of the Austrian crown, this characteristic: The exchange was stabilized *before* there existed conditions (above all the equilibrium of the Reich Budget) which alone could assure a *lasting* recovery of the monetary situation." Keynes' eloquent *A Tract on Monetary Reform* (1923. p. 40 and 47), which in effect pre-dates actual stabilizations of the time, makes a subtler point: "Has the public in the last resort no remedy, no means of protecting itself against these ingenious depredations? It has only one remedy—to change its habits in the use of money. The initial assumption on which our argument rested was that the community did *not* change its habits in the use of money. ….. a minimum is reached eventually from which the least favorable circumstance will cause a sharp recovery."

[3] Sargent's original essay was published in 1982.

Sargent and Wallace (1987), Marcet and Sargent (1989) and Marcet and Nicolini (2003). It captures two main ideas: the workings of a credible regime shift that is typically required to end a hyperinflation, and the study of the dynamics of adjustment under learning and allowing for a small departure from rational expectations (to be precisely defined below). We employ the internal consistency (IC) equilibrium requirement in Marcet and Nicolini (2003) to compare and numerically contrast our results. We show that, under both least squares and neural network learning, the large size of fiscal shocks and higher levels of average seiniorage in early transition economies account for the dynamics and explosive patterns of hyperinflation we observe in this period and countries.

2. Motivation

There is widespread agreement that inflationary expectations play a central role in hyperinflations. What is still little understood is how these expectations are formed. Since Cagan's pioneering work, Sargent and Wallace (1973), Sargent (1977) and Sargent and Salemi (1979) have formulated the model under rational expectations to explain governments' reliance on deficit finance through money creation observed in most hyperinflations. This result points to the presence of a feedback from inflation to future rates of money creation, i.e., inflation Granger-causes money creation. In at least several of the hyperinflations that Cagan studied, empirical results in Sargent and Wallace (1973, p. 418-419) provide evidence that inflation influences subsequent rates of money growth. This link is stronger than the weaker effects going in the other direction. Thus, money creation is endogenous, governed by the need to finance fiscal deficits.

This happens, Sargent (1977, p. 431) claims, because the system during hyperinflation operates under a particular money supply rule which predetermines the outcome. If, however, the monetary regime changes, so will the causality in the money creation-inflation process. Furthermore, Sargent (1977) explains, due to this simultaneity bias in Cagan's estimator, a "paradox" emerges which may easily be resolved: the actual rates of inflation in these hyperinflationary episodes far exceeded the rates that would yield the maximum sustainable revenue from money creation available to governments.[4]

To address this and other regularities exhibited in the data, such as the tendency for the real balances to sharply decline while inflation soars, Sargent and Wallace (1981, 1987) use a Cagan-type, portfolio equation for the demand for real money balances together with a government budget constraint relating seigniorage and money creation. Under rational expectations (RE), they show that the model exhibits a continuum of equilibria converging to two stationary rational expectations ones. The model implies that the economy may end up on "the slippery side of the Laffer Curve." These dual equilibria, studied also in Bruno and Fisher (1990), Benthal and Eckstein (1990) and Lee and Ratti (1993) though under a slightly different (continuous-time) set-up, suggest that the same amount of seigniorage (the revenue from money creation) may be extracted at either the high "inflation trap" or the lower rate of inflation consistent with RE. Bruno and Fisher (1990, p. 353 and 373) conclude that this may be the direct result of the operation rules (regime) the government selects for its monetary and fiscal policy. How and which equilibrium path is to be selected remains an open question. However, Sargent and Wallace (1987) also provide little explanation of the way in which agents

[4] Companion empirical work shows this to be the case in several East European hyperinflations. See also Salemi and Sargent (1979).

come to acquire RE and knowledge about the economy experiencing hyperinflation (Fischer, 1987). The model with rational expectations cannot account for rapid and increasing rates of inflation that end with a sudden "restoration of normality."

Laboratory evidence and a survey of expectations (in the case of the Bulgarian hyperinflation) also suggest that significant adjustment of expectations occurred only in the period immediately preceding policy change (Marimon and Sunder (1993), Carlson and Valev (2001)). Burdekin and Whited (2001, p. 78), however, point to some differences in this regard in stabilizing Taiwanese hyperinflation (1945-1953). The reform package of 1949 had little influence on the public's future expectations of inflation, and stabilization occurred only gradually and only in the presence of multiple regime shifts.

Marcet and Sargent (1989c) present a model of learning and hyperinflation, in which either the economy converges to the low inflation steady state or no equilibrium exists under least squares learning.[5] However, the result is classical in the sense that a higher deficit is associated with a higher, stable stationary inflation rate. Under RE, an increase in deficits results in a lower stationary inflation rate. These results appear to be not completely satisfactory, and as Bullard (1994) demonstrates, under an alternative preference map, least squares learning converges to a periodic equilibrium. Agents are also not allowed to re-specify their perceived law of motion of the system in response to any structural breaks, or regime shifts.

---

[5] By least squares learning, in this context, we mean an adaptive learning rule, which is a mapping of past inflation observations into future forecasts whereby agents rely on misspecified economic models and "run" least squares regressions that converge asymptotically to the actual ones. See Marcet and Sargent (1989a,b). For earlier insight and analysis, see DeCanio (1979), Bray (1982), and Bray and Savin (1986). For convergence, these authors rely on methods different than the ones presented in the above papers (Ljung's ordinary differential equation method). See also Woodford (1990) for exposition of Ljung's (1977) main theorems. Margaritis (1987) is the first application of Ljung's results in economics.

Arifovic (1995) and Bullard and Duffy (1998) study the model under genetic algorithms, interpreting it as a model of learning and emulation in which heterogeneous agents interact and evolve to understand how to forecast the future. The results in their artificial economies indicate that the lower of the two is a stationary equilibrium, while at the same time a self-fulfilling hyperinflation may be reached and sustained. Duffy and Bullard (1998) also show non-convergence to any of the equilibria.

Adaptive learning of expectations addresses a number of issues raised by the above inconsistencies in modeling expectations and hyperinflation episodes in particular. One set of questions involves the examination of transitional dynamics in equilibrium selection, where the experience gained from transition economies is particularly useful to motivate its application. Following some type of structural change or policy regime shift, economies necessarily grope for ways to reach RE. Some RE are learnable while others are not, as summarized in Evans and Honkapohja (2001) and Arifovic and Bullard (2001).

This paper uses a model developed by Marcet and Nicolini (2003) to test, via simulations, the stability and the empirical validity of the model in a hyperinflationary environment. It also examines the behavior of the model under two alternative learning rules and addresses the stylized facts of recent hyperinflation experiences.[6] We reflect, in general, on the convergence of learning processes based on neural networks for which no analytical results are yet available.[7] It is of considerable interest to understand whether the stability conditions for this particular self-referential model are the same under

[6] In the same spirit, though different in approach, Sargent (1999) and Lansing (2001) provide some explanation for the "Great Inflation" in the US.
[7] See White, (1989, p. 1003): "…investigations of network learning methods have so far proceeded by Monte Carlo simulation; a rigorous and general analysis of network learning has not yet been carried out."

alternative (reasonable) econometric learning rules. Are these (in)stability conditions the same as the frequently encountered "expectational stability" or E-stability conditions? If not, why not, and does this vindicate the model if it can explain stylized facts in hyperinflations? Section 3 discusses the stylized facts. Section 4 finds motivation in the existing literature on learning and hyperinflation and addresses some of its deficiencies. Neural networks and learning are the subject of Section 5. Section 6 describes the model. Learning is introduced, motivated and defined in Section 7. Section 8 studies the model under neural network learning. Simulations and results are presented in Section 9.

3. Hyperinflations: The Stylized Facts and Stabilizations

Table 1 reports selected hyperinflations over the last century. It is evident that the hyperinflation phenomenon is relatively short in duration and characterized by highly unstable and explosive paths of inflation and money growth (Bruno, 1991). In most instances, there was a complete collapse of the currency and the monetary system. Most of those hyperinflations are well documented. The important common feature that emerges is the fiscal difficulties that accompanied each episode. In some sense, hyperinflations have been largely fiscal in origin. The way these economies regained stability was only after a credible and abrupt change in fiscal and monetary regimes. This attests to the importance of inflationary expectations in successful stabilizations. It is only after a drastic revision of the public's expectations of future government policies that inflation is promptly curtailed (Sargent, 1993). In Germany, inflation was finally brought down suddenly after the implementation of a new fiscal regime in November

1923.[8] Bolivia managed to end its hyperinflation by unifying and stabilizing the exchange rate while at the same committing to adjusting the fundamental fiscal imbalances in the economy (Sachs, 1987).

It is well established in the literature that these types of stabilizations work successfully through influencing inflationary expectations in such a way that the public understands the reform packages carried out by governments (*transparency*) and anticipates that the announced policies will be pursued (*credibility*). This is achieved through a combination of "orthodox" and "heterodox" policy. Stabilizations that occur through incomes policies are called heterodox, and those that do not orthodox. Orthodox policies aim to reduce the deficit permanently through both monetary and fiscal reforms. In addition, those programs that stabilize inflation using the exchange rate as the nominal anchor for the economy are *exchange–rate based* stabilizations.[9] In *money-based* stabilizations, the central bank targets the money stock. Table 2 summarizes some recent stabilization experiences. With the exception of Brazil, the hyperinflations of Table 1 all ended with the implementation of some type of an orthodox program. Thus, in what follows, we model this feature explicitly to discuss the end of hyperinflations and the role of learning.

Figures 1 and 2 exhibit the quarterly, and where available the monthly, inflation rate and seigniorage as a percent of GDP for each of the hyperinflations we use to

---

[8] Vegh (1992) describes the experiences of effectively stopping hyperinflations in 1920s Europe and 1980s Latin America. His evidence suggests that the exchange rate served as the nominal anchor, and hyperinflations were stopped "almost overnight" without significant loss in output. Webb (1986) also asserts that inflationary expectations were driven in an important way during the German hyperinflation by fiscal news, i.e., the steady increase and anticipation of budget deficits.

[9] For more on the implications of money-based and exchange rate-based stabilization, see Bofinger (1996). He examines the experience of three former Soviet Union countries. In addition, see also Dornbusch et. al. (1990) and Bruno (1991).

motivate this study. For these countries, the data show that seigniorage rises as the deficit increases and inflation remains high.[10] Notice also that the levels of seigniorage that lead to hyperinflation in all cases are similar to those that occur at much lower rates of inflation. In other words, bursts of explosive inflation occurred in periods in which seigniorage was low and even declined. For example, in Ukraine, a five percent seigniorage level corresponds to both the highest rate of inflation in the fourth quarter of 1993 and the lower inflation rate of the first quarter of 1995. This suggests that a low contemporaneous relationship exists between these two variables in hyperinflations. Not reported here, the cross correlograms of these two series reveal a well-defined spike in the third and fourth lags and then descrease to zero.[11] The cross correlations are positive since an increase in seigniorage produces a subsequent rise in inflation. It also appears that seigniorage leads the periods of accelerating inflation.

Finally, following Marcet and Nicolini (2003), we may detail the stylized facts of hyperinflations in the following way (and as our examination of recent hyperinflations in the 1990s reveals):

1. Hyperinflations occur in countries with high, on average, levels of seigniorage.
2. There exists a low contemporaneous correlation between seigniorage and inflation, for a given hyperinflation country.

---

[10] See also Fischer, et. al. (2002). In addition, we estimate the following non-linear relationship to establish this in our case: $seigniorage = \beta_0 + \beta_1 * \inf lation + \beta_2 * \inf l^2 + \varepsilon$. Controlling for fixed effects for each country, we find β_1 to be positive and significant and β_2 negative and significant but of smaller magnitude. This may indicate the presence of a Laffer curve effect.

[11] The estimated cross-correlation coefficient is defined as: $\hat{r}_j^{S\pi} = \dfrac{\sum_{t=j+1}^{T}(S_t - \bar{S})(\pi_{t-j} - \bar{\pi})}{\sum_{t=1}^{T}(S_t - \bar{S})\sum_{t=1}^{T}(\pi_t - \bar{\pi})} \quad for\ \ j = 0 \ldots T-1.$

3. A *credible* regime shift (an exchange rate rule) stops hyperinflations successfully.

These stylized facts appear to be uncontroversial in the literature; see Bruno (1991, 1993), Dornbusch, Sturzenegger, and Wolf (1990) and Fisher, Sahay, and Vegh (2002).

4. Why Learning Hyperinflations?

Hyperinflationary episodes are usually accompanied by a sharp fall in real balances and chaotic behavior of money growth and prices. Cagan (1956) attempted to isolate a stable money demand function even in these extreme circumstances. In periods of such short but drastic changes and upheaval, he explained that demand for real balances is not a function of income or the real interest rate but rather, more generally, of the rate of return on holding money. The opportunity cost of holding money is the rate of inflation, as even anecdotal evidence suggests. However, what matters is the rate of inflation that is expected to prevail in the immediate future and not past values of inflation.

$$\frac{M_t}{P_t} = F({\pi_t}^e) = e^{-\alpha {\pi_t}^e} \qquad (1)$$

Cagan also postulated that adjustment occurred as a fraction of the forecast error of the previous period. His adaptive expectations formula is as follows:

$${\pi_{t+1}}^e - {\pi_t}^e = \beta(\pi_t - {\pi_t}^e) \qquad (2)$$

β is the speed of adjustment. As discussed earlier, there are a number of problems with this approach.[12] Furthermore, Cagan (1956, p. 77) points out that sudden shifts in expectations will not occur if expected future inflation is determined only by past events.

These concerns were first addressed in Sargent and Wallace (1973) who formulated the model with rational expectations and followed by a series of articles by Sargent and others. The issues that were mainly addressed concern the inertia evident in Cagan's adaptive expectations and the sudden and abrupt stabilization that accompany most hyperinflations. This point was fully developed in Sargent's "The End of Four Big Inflations" (1982). In addition, it appeared that the estimates of alpha in Equation (1) did not accord well with the actual average inflation rates experienced. Governments inflated too fast to maximize revenue through seigniorage. There is some empirical evidence to this effect, especially in recent transition countries' experience.

Sargent and Wallace (1981 and 1987) formulated a model with a Cagan-type money demand with rational expectations in which they specified explicitly the endogeneity of the money creation process, motivated by the need to finance a given level of seigniorage. They discover (confirm) a duality of rational expectations equilibria, with high inflation stationary equilibrium being the stable solution and the attractor of a continuum of equilbria. Given the stylized facts of hyperinflations, this is not completely satisfactory. In some sense, it even implies that the economy may end up in high inflation equilibria indefinitely. In addition, raising the level of seigniorage,

[12] Cagan (1956, 72) also notes this deficiency eloquently: "…In the light of the sharp rise in the balances when a reform of the currency approaches, any diminution in the rate at which notes were issued would likely alter the prevailing expectation of a certain future inflation to one of a less rapid rate, whether the reaction index were greater than unity or not. If so, the balances would rise at once if the policies of the note-issuing authorities justified more confidence in the future value of the currency. These sudden revisions in expectations cannot be accounted for in a model that predicts future prices on the basis of past changes in prices and money alone."

reduces inflation. This, however, is counterintuitive and does not also accord with actual observations. This is illustrated in Figures 3 and 4 of Appendix 3 in the paper.

Marcet and Sargent (1989b) analyze least squares learning in the context of this model. They find that the model either converges to the low inflation equilibrium, which is Pareto superior, or no equilibrium exists, for high enough seigniorage. This result is classical in the sense that higher deficits financed through seigniorage obtain stable stationary inflation rate. These findings are the basis of the dynamics of the model developed by Marcet and Nicolini (2003) and the one we study under two alternative learning procedures in the rest of the paper.

As described in Marcet and Sargent (1988, 1989a), adaptive learning arises in the following manner: agents stipulate a mapping of the equilibrium dynamics of the state variable in the model economy; then they estimate an econometric model that fits this relationship. This econometric model describes the agent’s perceived law of motion (PLM). Agents thus use these estimates to evaluate forecasts of future values of state variables. In other words, “agents act like econometricians.” These forecasts are then used in the model to obtain solutions for the actual path of the variables of interest. This is the agent’s actual law of motion (ALM). This process is updated each period and repeated recursively as more data become available. The main interest in the analysis is the conditions under which PLM converges to the ALM and the choice of the updating (decision) rule that governs the agents’ behavior. Evans and Honkapohja (2001) discuss the notion of expectational stability (E-stability) that defines the convergence of learning dynamics to rational expectations equilibria. In other words, the conditions for the asymptotic stability of the rational expectations equilibrium under least squares learning, for example, are described by the mapping from PLM to ALM.[13] In this instance, for the model adopted below, Marcet and Nicolini (2003) describe the stability properties of the mapping T.

Marcet and Nicolini's (2003) work also suggests that rational expectations and adaptive learning may be seen as non-competing approaches. If rational expectations obtain, then those may still be learnable in the long run. Thus the adaptive learning approach to expectations formation allows us to explore the learnability and stability of steady state solutions in the seigniorage model (Evans and Hokapohja, 2001 and Lettau and Van Zandt, 2002). Recent contributions are Barucci (2001) and Heinemann (2000). Bullard and Butler (1993), discussing policy implications of these types of models, declare: "A learning hypothesis is attractive, since it is not clear that perfect foresight is even a feasible assumption when equilibrium dynamics become complicated or chaotic, because then agents in the model predict perfectly even when observers outside the model see only white noise. These issues remain open. When deviations from baseline assumptions are incorporated, *stabilization policy* can be both desirable and efficacious. This is therefore an important avenue for future research." (p. 866)

5. Hidden layer neural networks and learning

Neural networks are a class of nonlinear models developed by cognitive scientists to mimic the architecture of the human brain. Implicitly, they are designed to 'learn' the

[13] Assume a mapping $\phi \to T(\phi)$, from PLM to ALM. Assume the rational expectations equilibrium (REE) is a fixed point of T, i.e., $T(\phi^*) = \phi^*$ under some regularity conditions. Then a given REE is said to be E-stable if the differential equation: $\frac{d\phi}{d\tau} = T(\phi) - \phi$ is locally asymptotically stable at ϕ^* (see also Evans and Honkapohja, 2001, p. 140).

model parameters through a complex interaction with their ‘environment’ (patterns of weights and connections) in a process that can be viewed as a recursive estimation procedure (the method of back-propagation). In this way, these networks discover an unknown functional relationship between inputs and outputs. They represent an input-output system in which a mapping of inputs $x_1, \ldots, x_p$ into outputs $y_1, \ldots, y_n$ occurs. For our purposes, they express functional forms that are approximating unknown relationships among relevant variables. The single hidden-layer feedforward network has the following form:

$$y = F(x;\delta) = a(\sum_{j=0}^{q} \beta_j a(\sum_{i=0}^{m} \lambda_{ji} x_i)) \tag{3}$$

We can view Equation (3) as a nonlinear function, in which only one layer of hidden units exists; the model is “feedforward” because signals flow in only one direction, more complicated networks may have both more hidden units and allow for certain “feedbacks” as well. a is the activation function, usually of sigmoid shape; β and λ are the unknown weights.

The usefulness of this functional form is its ability to approximate a set of mappings from input to output spaces as indicated by Hornik, Stinchcombe and White (1989) and White (1989). The theorem is formally stated in Judd (1998, p. 246). However, performance is sensitive to specification and one should be careful in applying these models to specific empirical phenomenon involving inputs and outputs. In this sense, NN coefficients are chosen to minimize the distance between the target and the output match. Determination of those coefficients is done through stochastic approximation methods that model a dynamic process of learning and updating (Judd, 1998).

The recursive learning procedure that is used in solving the minimization problem is called back-propagation, and is implemented through a local gradient descent (see Kosko, 1992 and White, 1989 for details). It is expressed in the following way:

$$\theta_t = \theta_{t-1} + \eta \nabla f(x_t, \theta_{t-1})(y_t - f(x_t, \theta_{t-1})) \quad for \forall t = 1, 2... \tag{4}$$

The back propagation starts with initial random weights θ and updates occur through (4). These values are a function of the βs and λs in (3). η is the learning rate, and the gradient with respect to the weights is ∇f. The weights continue to adjust in response to the errors of reaching the target output, y_t. The error, in this sense, is propagated back. If neural network training is done "on-line" as a recursive scheme, it closely resembles adaptive learning as discussed above. It simply represents an alternative statistical learning rule. η can be either decreasing to accommodate the effects of shocks that die out in the approximation or it can be a constant where it is appropriate for learning in evolving systems and tracking is necessary for abrupt changes (White, 1990). In this sense, the learning rate controls the influences innovations have on updating the current estimate. We attempt to exploit this feature of the learning algorithm below, where vanishing learning rates may fail to recognize a sudden shift in the system and we resort to some optimal fixed learning rates to track evolving phenomenon better as in the case of hyperinflations. To our knowledge, this application of the back propagation algorithm in neural networks, in which the learning rate "switches" endogenously, appears for the first time in the literature.

The introduction of neural network learning in the seigniorage model, as described below, helps us redress the following issues in hyperinflation episodes:

1. Is modeling learning important for understanding stabilizations? Can agents learn the rational expectations equilibrium over time, as indicated by Cagan (1991) and Bullard and Butler (1993)?

2. How does the convergence in neural network vs. least squares learning occur and at what values? Gradient descent learning has been shown to converge under E-stability correspondence (Evans and Honkapohja (1998) and Barucci and Landi (1997), what about the neural network learning variant? As suggested by Chen and White (1998), is it preferable that we use other nonparametric techniques, i.e., PLM be specified non-parametrically?

3. What do we learn about the behavior of the model and does it indeed match the stylized facts of hyperinflations?

6. A Model of Seigniorage and Learning

6.1. The Economy

The underlying economic model is one of inflation with government financing the deficit by seigniorage.[14] We consider a basic overlapping generations (OLG) economy with fiat money, where agents are two-period lived. Trade and consumption of a single perishable good take place in periods $t \in \{0,1,\ldots.\}$. Every agent of generation t lives only over t and t+1 (young and old, respectively) and consumes in these periods c_{1t} and

[14] An extensive description of the model can be found in Ljunqvuist and Sargent (2001, Chapter 8), and Lettau and Van Zandt (2002), among others. Sargent (1993, p. 90-95) provides a stochastic version of the model.

c_{2t}. Each agent has a unique endowment (e_1, e_2), $e_1 > e_2$, and the same utility function U(.). There is also a storable good, "money", used in exchange that has no consumption value. The initial stock of money is $m_{-1}>0$, supplied to the initial generation that lives only one period with endowment of e_0. In addition, we consider a government that finances the deficit through seigniorage only; this implies money does have value. Let p_t be the price of the perishable good at period t. Let $\pi_{t+1}=p_{t+1}/p_t$ be the (gross) inflation rate in period t+1. Young agents trade and form expectations about the price to prevail in the next period. Given the current price and an expected future price, p^e_{t+1}, agents in generation t choose consumption and money to:

$$\begin{aligned} &\max_{c_{1t}, c_{2t}, m_t} U(.) = \ln c_{1t} + \beta \ln c_{2t} \\ &\textit{subject to}: \quad c_{1t} \le e_1 - m_t / p_t \\ &\qquad\qquad\qquad c_{2t} \le e_2 + m_t / p^e_{t+1}. \end{aligned} \tag{5}$$

We assume the problem has a unique solution and the constraints are binding in equilibrium.[15] The solution depends in an obvious way only on the expected inflation. It is well known that the solution to this problem yields a Cagan-type demand (savings function) for real balances given by:

$$m_t = \gamma_1 - \gamma_2 \pi^e_{t+1} \tag{6}$$

where the γ's are functions of the discount factor β in the utility function and endowments. Let $\gamma_1 > \gamma_2 > 0$.

As mentioned above, money creation is only driven by the need to finance seigniorage. The government augments the initial money stock in each period t in the following way (this is its budget constraint):

[15] For further derivations of the results that follow, see the Appendices.

$$h_t = h_{t-1} / \pi_t + d_t \tag{7}$$

where $\{d_t\}_{t=0}^{\infty}$ is an exogenous iid stochastic process and represents seigniorage, h_t is the money supply.[16] Using the equilibrium condition in the money markets and equations (6) and (7), we readily obtain an equilibrium map for the economy:

$$F(\pi^e_{t+1}, \pi^e_t, \pi_t) = 0\,.^{17}$$

In particular, the equilibrium dynamics of the economy is described by the actual inflation as a function of expected inflation this (current) period and expected inflation in the following period. Furthermore, given our particular choice of the utility function in (5), we can write:

$$\pi_t = \phi(\pi^e_{t+1}, \pi^e_t) = \frac{\gamma_1 - \gamma_2 \pi^e_t}{\gamma_1 - \gamma_2 \pi^e_{t+1} - d_t} \tag{8}$$

The rational expectations "dual" equilibria (REE) to this model are well known and thus the main results are relegated to Appendix 2.[18]

We augment this formulation of the model along the lines of Marcet and Nicolini (1998, p. 8), to capture the behavior of governments and economies experiencing hyperinflation (see Table 1). Equation (8) describes the solution in periods of no exchange rate rules (or other similar rules that reduce governments' incentives to collect seigniorage through money creation). In periods of a "fixed exchange rate rule" (ERR), the government's concern about the current levels of inflation prompts it to peg the

---

[16] A more straightforward, intuitive way to express (7) is the following: $m_t = m_{t-1} + p_t d_t$.

[17] Stationary solutions satisfy $F(\overline{\pi}, \overline{\pi}, \overline{\pi}) = 0$. There are two stationary rational expectations solutions, as discussed in the introduction, (π^L, π^H), low and high inflation states. These are shown in Appendix 2.

[18] Detailed description of the behavior of these solutions when $\beta^U = \infty$, or very high, and there is no uncertainty introduced through assuming random seigniorage d_t (or $\sigma_d^2 = 0$) exists in Sargent and Wallace (1981), Marcet and Sargent (1989) and Arifovic (1995). However, Proposition 3 in Marcet and Nicolini (2003) verifies the properties of the equilibrium map with these additional extensions.

foreign exchange rate thorough foreign reserves operation in the market (some of these rules are easily instituted through a currency board, for example Argentina (1991) and Bulgaria (1997)). Then the targeted inflation rate is $\frac{p^f{}_t}{p^f{}_{t-1}}\frac{e_t}{e_{t-1}}=\bar{\theta}$, where e_t is the exchange rate in terms of the foreign currency and p^f_t is the foreign price level. Assuming purchasing power parity and arbitrage in foreign exchange markets,[19] (9) obtains:

$$\frac{p_t}{p_{t-1}}=\pi_t=\bar{\theta} \tag{9}$$

To implement this policy, the government needs to know only the foreign price level and the exchange rate. The government budget constraint will not hold under the new regime; it must be supplanted with the stock of foreign exchange reserves to acknowledge the adjustment the government makes to enforce ERR.[20] Since we do not intend to formally model reserves in this model, we impose a rule on how the government acts in times of increasing inflation. In a sense, they enforce the rule whenever inflation reaches an upper limit:

$$\frac{p_t}{p_{t-1}}=\pi_t\leq\theta^U \tag{10}$$

where θ^U is the maximum “tolerated” inflation rate, after which ERR is imposed.

Hence, the model can be expressed as:

[19] Some evidence of PPP in high inflation transition economies is provided in Christev and Noorbakhsh (2000).

[20] This implies as in Marcet and Nicolini (2003): $H_t=H_{t-1}+d_t p_t+e_t(R_t-R_{t-1})$, where R_t is the level of foreign exchange reserves and H_t the nominal level of money stock.

$$\pi_t = \phi(\pi^e{}_{t+1}, \pi^e{}_t) = \frac{\gamma_1 - \gamma_2 \pi^e{}_t}{\gamma_1 - \gamma_2 \pi^e{}_{t+1} - d_t} \quad if \ 0 < \frac{\gamma_1 - \gamma_2 \pi^e{}_t}{\gamma_1 - \gamma_2 \pi^e{}_{t+1} - d_t} < \theta^U$$

$$= \bar{\theta} \qquad\qquad otherwise \tag{11}$$

6.2. Adaptive Learning and Lower Bounds to Rationality

Assume a self-referential model of the type discussed in Marcet and Sargent (1988, 1989a,b) that satisfies:[21]

$$x_t = H(x_{t-1}, x^e{}_{t+1}, \xi_t, \lambda)\ , \tag{12}$$

where H is determined by (6) and (7) and market equilibrium. x_t is inflation and money supply, x^e_{t+1} is agents' expectations of the future value of x_t and ξ_t is seigniorage. λ contains a vector of parameters relating to the behavior of the public and government policy: γ_1, γ_2, $\bar{\theta}$ *and* θ^U .

Agents form expectations in the following way:

$$x^e{}_{t+1} = z(\theta_t(\mu), x_t) \tag{13}$$

$\theta_t(\mu)$ is estimated from past data, and z is the forecast function that depends on the state variable x_t. The θs are governed by a learning rule *f*, where μ's explicitly model how past data is to be used in estimation of the forecasts:

$$\theta_t(\mu) = f(\theta_{t-1}(\mu), \mu, x_t). \tag{14}$$

[21] See also Evans and Honkapohja (1995) for additional description and stability results for this type of models.

Let us assume that for this model, z is $\pi^{e}{}_{t+1} = \theta_t$. Below, by imposing lower bounds to rationality of agents in the model, due to Marcet and Nicolini (2003), we relate and restrict the behavior of the above parameters via the learning rule, *f*.

Adaptive learning places *upper* bounds on rationality by assuming agents do not know the exact economic model under which they operate and have bounded memory. There may be a variety of learning rules that meet these bounds, and as Sims (1980) has warned us, our models may simply slip into the "wilderness of irrationality" and not be falsifiable. To avoid this dilemma, Marcet and Nicolini (2003) suggest that we allow for small departure (in transition and asymptotically) from rational expectations. This is proposed so that, given an economic model and some empirical observations, we may introduce learning rules that satisfy certain *lower* bounds while at the same time confront the model with the observed behavior of the economy. Much like Marcet and Nicolini (2003, p. 13 and 25), the simulations discussed in Section 9 provide evidence that the model may generate equilibria quite different from the rational expectations (RE) ones and match the data better.

Marcet and Nicolini (2003, p. 15-17) propose three *lower* bounds to rationality that we adopt in what follows and study two alternative learning rules. Let $\pi^{\varepsilon,T}$ be the probability that, in a sample of T periods, the perceived prediction errors will be within ε>0 of the true conditional error:

$$\pi^{\varepsilon,T} \equiv P(\frac{1}{T}\sum_{t=1}^{T}\left[x_{t+1} - x_{t+1}{}^{e}\right]^{2} < \frac{1}{T}\sum_{t=1}^{T}\left[x_{t+1} - \mathrm{E}_t(x_{t+1})\right]^{2} + \varepsilon) \tag{15}$$

where E_t is the true conditional expectation under the learning rule and x's are defined as above. This definition implies that after T periods the prediction error from the agents in the model is within ε of the true prediction error. The convergence of this probability asymptotically limits the behavior of our learning rules and allows the weights (learning steps), with which agents in the model update their forecasts, to be determined endogenously in what follows. The first bound to rationality in Marcet and Nicolini (2003) is as follows:[22]

Definition 1: *Asymptotic Rationality (AR): the learning rule satisfies asymptotic rationality in the model of Section 6.1 if:* $\pi^{\varepsilon,T} \to 1 \quad as \quad T \to \infty \quad for \ \forall \varepsilon > 0$.

This allows agents to make forecasting errors for some time in the future but not do it consistently forever.

Definition 2: *Epsilon-Delta Rationality (EDR): the learning rule satisfies epsilon-delta rationality for (ε, T, δ) in the model of Section 6.1 if:* $\pi^{e,T} \geq 1-\delta$.

Definition 2 requires that agents make optimal or near optimal forecasts within a given learning rule. Denote $\tilde{\theta}(\mu,\mu')$ to be the forecast given by the learning parameter μ' when in fact all agents in the model are using the parameter μ:

$$\tilde{\theta}_t(\mu,\mu') = f(\tilde{\theta}_{t-1}(\mu,\mu'), x_t^{\mu}, \mu'),$$

where f is the learning rule under study.

Definition 3: *Internal Consistency (IC): Given the model in Section 6.1, the learning rule satisfies internal consistency for (ε, T)* if :

$$\mathrm{E}\left(\frac{1}{T}\sum_{t=1}^{T}(x_{t+1}{}^{\mu} - z(\theta_t(\mu), x_t{}^{\mu}))^2\right) \leq \min_{\mu'} \mathrm{E}\left(\frac{1}{T}\sum_{t=1}^{T}(x_{t+1}{}^{\mu} - z(\tilde{\theta}_t(\mu,\mu'), x_t{}^{\mu}))^2\right) + \varepsilon, \quad (16)$$

[22] As these authors acknowledge, they are not the first to impose these requirements in the literature on stability of RE under adaptive learning. See Kurtz (1994) and also Fudenburg and Levine (1995).

where z(.) is defined as the expectations of the agents in the model. This requirement, in fact, helps us select an efficient learning parameter from a set of alternatives that identifies an equilibrium in the model given the learning rule used by the agents. While it is not straightforward to show that *Epsilon-Delta Rationality* satisfies a given learning rule, it implies that learning forecasts might in effect be better at detecting a regime change and once in place the agents are not willing to change to a different learning mechanism within the model.[23] We only employ the IC criteria of Eq. (16) to select the learning step in the back-propagation algorithm and compare the results to the least squares learning as defined below.

7. Learning Equilibria with Least Squares

Now suppose, we depart from the rational expectations hypothesis used in the model above. Instead, let us assume expectations are given by: ${\pi^e}_{t+1} = \theta_t$. Under certain conditions, not all equilibria could be the outcomes of adaptive learning processes described in Section 6.2. The learning rule could change adaptively in a way that the parameters are altered in response to forecast errors. Agents form expectations of inflation in the next period as a weighted average of the previous inflation expectations and updates on the last prediction error. One way to express this learning mechanism is:

$$\theta_t = \theta_{t-1} + \frac{1}{\alpha_t}(\pi_{t-1} - \theta_{t-1}) \tag{17}$$

[23] See Marcet and Nicoli (2003, p.25-26) for a proof that ERD obtains when least squares learning is used by the agents under uncertainty.

given initial values θ_0. This is a well-known adaptive expectations formula, a particular expression of (13), where the coefficient ("gain" parameter) specifying the revision of the forecast error $(\frac{1}{\alpha_t})$ decreases over time at rate t.[24] The evolution of the learning step and (17) determine the learning mechanism *f* in (13).

The gain sequence is often expressed as:

$$\alpha_t = \alpha_{t-1} + 1 , \tag{18}$$

with initial α_1=1. Thus, when $\alpha_t = t$ (commonly used assumption), the forecast rule is simply a regression of the mean of past inflations.

Alternatively, when there is a structural change in the economy, and agents have to update their perceived inflation accordingly, the approach that appears more robust is to assume that α_t is not decreasing infinitely but rather is equal to some fixed value $\bar{\alpha}$ >1. This is called a "constant-gain" algorithm and is studied in Benveniste (et. al., 1990, ch.1 and 4 part I). Unlike (18), where new information is given proportionately less weight, the constant-gain step allows for abrupt changes in the economy to be followed more rapidly, allowing higher weights on the last available information. The optimal choice will depend on the magnitude and frequency of the structural change.[25]

Marcet and Nicolini (2003) propose to combine the two approaches, and have least squares updates with a decreasing gain to occur in periods of stability and "tracking"

[24] See, for example, Evans and Honkapohja (1995) and Marcet and Sargent (1989a). Lettau and Van Zandt (2002) study the stability of this learning rule for the same model of hyperinflation as in this paper.
[25] As shown in Ljung, Pflug and Walk (1992), there is a trade-off between bias and accuracy when a constant-gain algorithm is used but it is clearly preferred to detecting abrupt changes in the structure of the economy, evolving over time. These are called algorithms with "tracking."

when sudden and abrupt change or instability is observed.[26] In this way, we allow the agents in the model to use an endogenous gain sequence in which agents use least squares when they do not make large mistakes and switch to fixed values of alpha when prediction errors increase above a certain threshold ν, say more than 10%, in cases of hyperinflation. Therefore, the gain sequence in (17) becomes:

$$\begin{aligned} \alpha_t &= \alpha_{t-1} + 1 \qquad\qquad if \ \left| \frac{\pi_{t-1}}{\theta_{t-1}} - 1 \right| < \upsilon \\ &= \bar{\alpha} \qquad\qquad\qquad otherwise \end{aligned} \tag{19}$$

Hence, the model in (11) with least squares learning becomes:

$$\begin{aligned} \phi(\theta_t(\alpha), \theta_{t-1}(\alpha), d_t) = \pi_t &= \frac{\gamma_1 - \gamma_2\theta_{t-1}}{\gamma_1 - \gamma_2\theta_t - d_t} \quad if \ 0 < \frac{\gamma_1 - \gamma_2\theta_{t-1}}{\gamma_1 - \gamma_2\theta_t - d_t} < \theta^U \\ &= \bar{\theta} \qquad\qquad\qquad otherwise \end{aligned} \tag{20}$$

Provided the learning mechanism in (17) and (19), Equation (20) is highly nonlinear and thus we describe its solution via simulations in Section 9 below. Marcet and Nicolini (2003) study the stability properties of the mapping from PLM to ALM when seigniorages are stochastic (see Proposition 3 of their paper, p. 30). This element of uncertainty determines the equilibrium path of the model. They prove that the function is increasing and convex and possesses two stationary rational expectations equilibria (fixed points of the mapping T): low inflation rate and high inflation rate (see also Figures 5 and 6). They also show convergence to the lower of the two stationary equilibria and that, for high enough seigniorage, no equilibrium exists. Figure 6, in addition, illustrates that the stable set (S set) decreases as seigniorage increases. A

[26] This is indeed also suggested in Evans and Honkapohja (1993).

succession of high seigniorage shocks increases perceived inflation in the unstable set (U set) and the system tends to unravel in an ever-increasing inflation. When inflation becomes intolerable, the government sets up a credible exchange rate rule (ERR) and agents switch to tracking as more recent events have become increasingly more important. This sends inflation back into S set. When the unstable U set is closer to the lower stable rational expectations equilibrium, hyperinflations become more likely since seigniorage is higher. We will use this intuition of the model to explain some of the stylized facts discussed earlier.

8. Learning Equilibria with Neural Networks

We introduce neural network expectations and learning in the model of Section 6. Thus, economic agents in this model will continually employ neural networks to learn and approximate their perceived (expectations of) inflation to the actual, true conditional expectations of the future inflation rate (Salmon, 1995). This process of trial and error (as explained above and represented by the back-propagation algorithm), of adjusting to “reality” as new data become available, proceeds until a trade-off margin is reached and it is no longer optimal for agents to adjust. In a sense, further accuracy is no longer necessary, given economic conditions (Salmon, 1995, p. 250). As mentioned, this stems from the powerful approximation properties of the hidden layer neural networks (NN). (See White 1989, 1992 and Kosko 1992). The measure of accuracy is the sum of the mean of squared prediction errors. Back-propagation learning (BPL) converges to a set of weights[27] (model parameters) that could provide locally (mean-square) optimal predictions of future expected inflation. (White (1989) and Kuan and White (1994, p.40 - 56) provide formal treatment).

In this nonparametric sense, without specifying the structure of expectations formation, the back-propagation algorithm and the NN that defines it may indeed have the potential of learning the true rational expectations equilibrium of the model (Salmon, 1995 and Heinemann, 2000). These studies show that learning, even 'approximate,' rational expectations equilibria, is not straightforward. The recursive, iterative nature of the algorithm used corresponds closely to adaptive learning mechanisms of Section 6.2, and mimics "the solution to forward-looking rational expectations." (Salmon 1995, p. 267) In addition, providing asymptotic convergence results for NN estimates, Kuan and White (1994, p.40-41, Theorem II.2.1) show that the solution trajectory of the back-propagation algorithm does not cycle between two stable equilibria. In what follows we examine whether (and if not, why not) the time varying, nonlinear inflation process implies a θ that converges and compare the results in quantitative sense to the least squares learning algorithm. It is of interest to understand how and under what set of initial conditions and parameter values the behavior of the system changes and evolves over time. Again, the goal is to explain some of the stylized facts of hyperinflations and contrast the results under different learning rules. In the process, the small departure, defined in the precise sense of Definition 3, away from rational expectations, is under scrutiny.

[27] These are the model parameters that we estimate through a back-propagation algorithm and represent future expected inflation. That is, they provide an approximation to the function (mapping) relating perceived to actual inflation in the model, i.e., the law of motion in the economy.

We construct a single hidden layer NN with two inputs: a constant and the current inflation rate. Let NN expectation of future inflation be:

$$\pi_{t+1}{}^{e} = F(\pi_t, \delta_t, cons\tan t) \tag{21}$$

Using the model equations (5) and (6), we obtain the current inflation generation process:

$$\begin{aligned} \pi_t &= \frac{\gamma_1 - \gamma_2 F(\pi_{t-1}, \delta_{t-1})}{\gamma_1 - \gamma_2 F(\pi_t, \delta_t) - d_t} \quad if \; 0 < \frac{\gamma_1 - \gamma_2 F(\pi_{t-1}, \delta_{t-1})}{\gamma_1 - \gamma_2 F(\pi_t, \delta_t) - d_t} < \theta^U \\ &= \overline{\theta} \qquad\qquad\qquad\qquad\qquad otherwise\,, \end{aligned} \tag{22}$$

where $\overline{\theta}$, θ^U and d_t are defined as before. The network expectation function $g(\pi, \delta)$ defined in Eq. (23) is obtained in the process described above. The inputs (inflation this period and a bias) are sent to intermediate (hidden) unit, which is activated through the sigmoid function, and sent forward to the output level. The hidden layer (indexed by j=1…q) yields an activation function $a_{ij} = G(\pi_i' \lambda_j + a_{i-1}' \mu_j)$, where G is the sigmoid function and i=1…T. The network output can be expressed as:

$$g(\pi, \delta) = F(\pi_t, \delta_t) = \beta_0 + \sum_{j=}^{q} a_{ij}\beta_j \;, \tag{23}$$

where $\delta = (\beta_0 \ldots \beta_q, \lambda_1' \ldots \lambda_q', \mu_q' \ldots \mu_q')'$. These are the weights, or connection strengths, and the parameters to be learned through a recursive back-propagation learning algorithm. The agents are estimating their updates of future expected inflation in a loop that allows them to target the previous period iteration's value as an approximation they learn based on current and past inflation rates. The learning rule in the model is:

$$\delta_t = \delta_{t-1} + \frac{1}{\alpha_t} \nabla F_\delta (\hat{\pi}_{t-1} - F(\pi_t, \delta_t))\,, \tag{24}$$

with a gain sequence α_t defined as in the least squares learning:

$$
\begin{aligned}
\alpha_t &= \alpha_{t-1} + 1 \qquad if \quad \left| \frac{\pi_{t-1}}{F(\pi_t, \delta_t)} - 1 \right| \le \nu \\
&= \bar{\alpha} \qquad\qquad\quad otherwise.
\end{aligned} \tag{25}
$$

The advantages of using this algorithm in the model are twofold. First, multiple passes through the data allow for reaching a global optimum. In addition, the algorithm represents a simple updating estimation procedure to solving complex models of unknown functions without much prior knowledge of the particular relationship among the variables involved. If the estimation converges, we expect it to converge to approximate rational expectations equilibria, thereby minimizing the mean square prediction error (Heinemann, 2000). However, apart from misspecification, problems may arise related to the feedback feature of the algorithm, and may prevent convergence. In addition, there is a dearth of analytic results for general nonlinear economic models using NN to approximate expectations. In simulations below, we show the behavior of the system regarding convergence to rational expectations solutions.

9. Model Solutions by Simulation: Comparing Different Learning Equilibria

IC criteria of Eq. (16) in Definition 3 is what we use to define equilibria in the model with both learning rules; thus the variables we determine provide sequences of current inflation, expected (perceived) inflation and the parameter $\bar{\alpha}$. Hence, these sequences should satisfy (20 and 22), (17 and 24) and (19 and 25) for all periods, given $\bar{\alpha}$, and $\bar{\alpha}$ should satisfy (16) for (ε, T) appropriately chosen. We numerically search for those alphas that fulfill IC and are efficient in the sense that they impose restrictions on the behavior of agents. First, we select $\frac{1}{\overline{\alpha}}$ over the grid [0,1.1]. The interval is discrete with a length of 0.1. We use the same values for the alternative learning parameter within a given learning rule, $\frac{1}{\overline{\alpha}'}$. We next compute the left-hand side of Equation (16) by Monte Carlo integration. We draw one thousand realizations of the random seigniorage d_t for t=1…T. In the model, we find equilibrium inflation rates for each of the realizations. Simultaneously, we find predictions of inflation for each alternative value of $\frac{1}{\overline{\alpha}'}$. We then calculate the mean square error for each of these alternative forecasts and average them over all realizations.

These Monte Carlo results are presented in Table 4. We seek the minimum across rows ($\frac{1}{\overline{\alpha}'}$) for each column ($\frac{1}{\overline{\alpha}}$). Those values over the grid that form a mean square error (MSE) within ε = 0.01 of the minimum in each column and lie on the diagonal of Table 4 are called the efficient values of the learning parameter α.[28] This is what is required in Equation (16). This implies that agents using an efficient value of alpha have little incentive to change their learning parameter given the choice of other agents with the same learning rule.

To simulate the model, we need to choose values for the parameters in the money demand function (6) and the distribution of d_t. We model the seigniorage as a truncated, normally distributed, variable with different means and standard deviation of 0.01. Since some evidence exists to support a Laffer curve in the experiences we study, we note that

[28] In this exercise, we closely followed the one done in Marcet and Nicolini (2003). However, for the neural network learning in this paper, we report Monte Carlo MSEs for only one high level of seigniorage since it is more computationally intensive. Numerous other runs were performed with different levels of seigniorage but only with 250 draws of the random variable.

the money demand function in the OLG model implies a stationary Laffer curve.[29] However, unlike Marcet and Nicolini (2003), who primarily want to explain the Argentine experience and used only one set of parameters, we use an interval of inflation rates that maximize seigniorage and a measure of maximum revenue from seigniorage to construct a number of different parameter values to use in the simulations. Inflation varied between 20 and 60 percent and the revenue from 5 to 40 percent. Simulations of the model with learning through least squares and neural networks were performed with eight different sets of parameter values. Table 3 presents those and gives the stationary rational expectations equilibria inflation rates (π^*_{REE1}, π^*_{REE2}) associated with these model economies.

We start with the parameter values used in Marcet and Sargent (1989b). The results are presented in Set 2 of Table 3. Both least squares and neural network learning failed to converge to the lower stationary inflation rate in this instance. This may be due to the difference in initial beliefs (though starting values in each case are given fairly close to π^*_{REE1}) or different convergence of the algorithms at values of seigniorage close to the maximum sustainable levels under rational expectations. Parameter Set 4 (used in Marcet and Nicolini) points to an interesting result: the relevance of the average level of seigniorage for the stability of (hyper) inflationary processes. While at the lower levels of average seigniorage (0.045) the inflation rate tends to stabilize around the lower rational expectations equilibrium value, at higher seigniorage levels (0.049) no convergence occurs using both learning rules. Increased average seigniorage leads to

[29] See Fisher, et. al. (2002). Azariadis (1993, p. 397) provides the formula for a stationary Laffer curve one may use to compute the range of parameters in the model. See also Kiguel and Neumeyer (1992) and Miller and Zhang (1997). The latter authors provide the range of parameters they used in their analysis of the Ukrainian hyperinflation.

higher and unsustainable inflation. What is surprising, however, is that, under different values of the parameters (Sets 1, 3, 5-8), both least squares and neural network learning algorithms behave strikingly similarly. Inflation converges to the lower rational expectations equilibrium. This result suggests that, given certain conditions, the agents in our model are able to learn the rational expectations outcome through either of the two alternative learning mechanisms.

Hence, simulations of the seigniorage model with a wide range of parameters implies that there may exist a correspondence between the stability (E-stability) conditions for the convergence of the NN learning and the least squares learning. However, this finding is in contrast to Salmon (1995), who suggests that the domain of attraction of equilibria is much higher and wider for neural networks than least squares regressions. Furthermore, Heinemann (2000, p. 2009) claims that neural network learning may give rise to expectations that do not minimize the mean square prediction error. Certainly, our results do not indicate that this is the case. Since in the form used here,[30] NN learning is similar to a stochastic gradient algorithm, it is worth pointing out that Barucci and Landi (1997) prove convergence to REE under conditions that turn out to be quite different than the ones generally established for least squares learning. Their results, however, depend crucially on assuming a decreasing learning step η_t. We, on the other hand, adopt a modified version of the back propagation learning to account for a sudden regime shift that effectively stop hyperinflations. In this way, by choosing an optimal η_t within a given learning rule, we not only provide justification for rational expectations equilibria as illustrated, but also show that learning introduces dynamics that

[30] Sometimes also referred to as the *generalized delta rule*.

are important and significant in their own right and track evolving expectations better than the rational expectation version of the model.

Figures 7-9 plot the results of simulations for Sets 6 and 8 and a particular, efficient value of $\frac{1}{\bar{\alpha}}$=0.2. It is clear that the model under learning creates much higher rates of inflation (hyperinflation) than the stationary rational expectations version. It also captures the sudden and abrupt stabilizations that occurred in most hyperinflations and exemplifies the importance of inflationary expectations in the process. Since sustained increases in inflation appear to be generated by normally distributed (iid) random d_t, there is little contemporaneous correlation between inflation and seigniorage. Introducing ERR, for example, successfully stabilizes inflation, returning it back to the stable S set around the low REE, for a time, but unless levels of seigniorage are reduced permanently, the system may be cast back into the unstable U set as a result of a large shock and another hyperinflation occurs (see Fig. 9-10). These figures clearly display some the stylized facts analyzed earlier. Hyperinflations occurred, on average, in high seigniorage countries.

10. Conclusion

The findings of this analysis are threefold. First, in the seigniorage model of hyperinflation with learning dynamics, the REE are learnable. Second, results from simulations suggest that, for a certain range of parameters, the behavior of the neural network learning with endogenously changing learning step is very similar to the one often exhibited by least squares learning. Third, the stylized facts of hyperinflationary episodes seem to be better explained with the introduction of learning dynamics than using only the rational expectations version of the model.

In addition, one may augment the model presented here and study the situation in which we replace the assumption of real government spending with nominal money growth rule, thus make government endogenous so that the government budget constraint is satisfied. In this case, introducing the type of learning described here may result in different equilibria.

REFERENCES:

Arifovic, J. (1995). Genetic algorithms and inflationary economies. Journal of Monetary Economics, 36, 219-243.

Aslund, A., Boone, P., Johnston, S., (1996). How to Stabilize: Lessons from Post-Communist Countries. Brookings Papers on Economic Activity, 1, 217-313.

Azariadis, C., (1993). Intertemporal Macroeconomics. Blackwell.

Barucci, E., (2001). Fading Memory learning in a class of forward-looking models with an application to hyperinflation dynamics. Economic Modelling, 18, 233-252.

Barucci, E. and Landi, L., (1997). Least Mean Squares Learning in Self-Referential Stochastic Models. Economics Letters, 57, 313-317.

Benveniste, A., Metivier, M. and Priouret, P. (1990) Adaptive Algorithms and Stochastic Approximations. New York: Springer-Verlag.

Bofinger, P., (1996). The Economics of orthodox money-based stabilizations (OMBS): The recent experience of Kazakhstan, Russia and Ukraine. European Economic Review, 40, 663-671.

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

Bresciani-Turroni, C. (1937). The Economics of Inflation. London: Allen and Unwin.

Bruno, M. and Fischer, S. (1990). Seigniorages, operating rules, and the high inflation trap. Quarterly Journal of Economics, 105, 353-374.

Bruno, M., (1991). High Inflation and the Nominal Anchors of an Open Economy. Essays in International Finance, 183, June. Princeton University

Bruno, M., (1993a). Stabilization and the Macroeconomics of Transition- How Different Eastern Europe? Economics of Transition,

Bruno, M., (1993b). Crisis, Stabilizations, and Economic Reform: Therapy by Consensus. Claredon Press, Oxofrd.

Bullard, J., Butler, A., (1993). Nonlinearity and Chaos in Economic Models: Implications for Policy Decisions. The Economic Journal, 103 (419), Jul., 849-867.

Bullard, J. (1994). Learning Equilibria. Journal of Economic Theory, 64, 468-485.

Cagan, P. (1956). The monetary dynamics of hyperinflation. In M. Friedman (Ed.), Studies in the Quantity Theory of Money. Chicago: University of Chicago Press.

Carlson, J. and Valev, N. (2001). Credibility of a new monetary regime: the currency board in Bulgaria. Journal of Monetary Economics, 47, 581-594.

Dornbusch, R., Sturzenergger, F., and Wolf, H. (1990). Extreme inflation: Dynamics and stabilization. Brookings Papers on Economic Activity, 2.

Duffy, J. (1994). On learning and the nonuniqueness of equilibrium in an overlapping generations model with fiat money. Journal of Economic Theory, 64, 541-553.

Evans, G. W. (1995). Expectational Stability and the multiple equilibria problem in linear rational expectations models. Quarterly Journal of Economics, 100, 1217-1233.

Evans, G. W. and Honkapohja, S (1993). Adaptive Forecasts, Hysterisis, and Endogenous Fluctuations. Economic Review, 1, Federal Reserve Bank of San Fransisco, 3-13.

Evans, G. W. and Honkapohja, S (1995). Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models. Econometrica, 63, 195-206.

Evans, G. W. and Honkapohja, S. (1998). Economic Dynamics with Learning: New Stability Results. Review of Economic Studies, 65, 23-44.

Evans, G. W. and Honkapohja, S. (1998). Stochastic Gradient Learning in the Cobweb Model. Economics Letters, 61, 333-337.

Evans, G. W. and Honkapohja, S. (2000a). Convergence for difference equations with vanishing time-dependece, with applications to adaptive learning. Economic Theory, 15, 727-725.

Evans, G. W. and Honkapohja, S. (2000b). Learning dynamics. In J. B. Taylor and M. Woodford (Eds.), Handbook of Macroeconomics, Volume 1. Amsterdam: Elsevier Science.

Evans, G. W. and Honkapohja, S. (2001). Learning and Expectations in Macroeconomics. Princeton: Princeton University Press.

Fischer, S., Sahay, R., Vegh, C., (2002). Modern Hyper- and High Inflations. Journal of Economic Literature, Vol. XL, Sept., 837-880.

Heinemann, M., (2000). Adaptive Learning of rational expectations using neural networks. Journal of Economic Dynamics and Control, 24, 1007-1026.

Hornik, K., Stinchcombe, M., White, H., (1989). Multi-layer Feedforward Networks Are Universal Approximators. Neural Networks, 2.

Judd, K., (1998). Numerical Methods in Economics. The MIT Press.

Keynes, J. M. (1923) A Tract on Monetary Reform. In The Collected Writings of John Maynard Keynes. London: MacMillan Press, 1971

Kiguel, M., Neumeyer, P., (1995). Seigniorage and Inflation: The Case of Argentina. Journal of Money, Credit and Banking, 27 (3), 673-682.

Kosko, B. (1992). Neural Networks and Fuzzy Systems: A Dynamical Systems Approach to Machine Intelegence. Prentice Hall.

Kuan, C-M., White, H., (1994). Artificial Neural Networks: An Econometric Perspective. Econometric Reviews, 13 (1), 1-91.

Kuan, C-M., White, H., (1994). Adaptive Learning with Nonlinear Dynamics Driven by Dependent Processes. Econometrica, 62, 1087-1114.

Lettau, M. and Van Zandt, T. (2002). Robustness of adaptive expectations as an equilibrium selection device. Forthcoming Macroeconomic Dynamics

Ljung, L., Pflung, G. and Walk, H. (1992). Stochastic Approximation and Optimaization of Random Systems. Berlin: Birkhauser Verlag.

Marcet, A. and Nicolini, J. P. (2003). Recurrent Hyperinflations and Learning. The American Economic Review. 93, 5, 1476-1498, December.

Marcet, A. and Sargent, T. (1988). The fate of systems with “adaptive” expectations. American Economic Review, Papers and Proceedings, 78, 168-172, May.

Marcet, A. and Sargent, T (1989a). Convergence of least squares learning mechanism in self-referential models. Journal of Economic Theory, 48, 337-368.

Marcet, A. and Sargent, T. (1989b). Convergence of least-squares learning in environments with hidden state variables and private information. Journal of Political Economy, 97, 6, 1306-1322.

Marcet, A. and Sargent, T. (1989c). Least squares learning and the dynamics of hyperinflation. In W. Barnett, J. Geweke, and K. Shell (Eds.), Chaos, Sunspots, Bubbles and Nonlinearities. Cambridge: Cambridge University Press.

Marimon, R. and Sunder, S. (1993). Indeterminacy of equilibria in a hyperinflationary world: experimental evidence. Econometrica, 61, 5, 1073-1107.

Miller, M. and Zhang, L. (1997). Hyperinflation and Stabilization: Cagan Revisited. The Economic Journal, 107, 441-454.

Sachs, J., (1987). The Bolivian Hyperinflation and Stabilization. The American Economic Review. 77, 2, Papers and Proceedings, May, 279-283.

Salemi, M. and Sargent, T. (1979). The demand for money during Hyperinflation under Rational Expectations: II. International Economic Review, 20 (3), 741-758.

Salmon, M. (1995). Bounded Rationality and Learning; Procedural Learning. In A. Kirman and M. Salmon (Eds.), Learning and Rationality in Economics. Blackwell.
Vegh, C., (1992). Stopping High Inflation: an Analytical Overview. IMF Staff Papers, 39, 3, Sept., 626-695.

Sargent, T. (1977). The Demand for Money during Hyperinflations under Rational Expectations. International Economic Review, 18 (1),

Sargent, T. (1993). Bounded Rationality in Macroeconomics. Oxford: Oxford University Press.

Sargent, T. and Wallace, N. (1973). Rational Expectations and the Dynamics of Hyperinfaltion. International Economic Review, 14 (2),

Sargent T. and Wallace, N. (1981). Some unpleasant monetarist arithmetic. Federal Reserve Bank of Minneapolis Quarterly Review, Fall, pp. 1-17. Also published in Rational Expectations and Inflation, T. Sargent, ed., New York: Harper Collins College Publishers, 1993.

Sargent, T. and Wallace, N. (1987). Inflation and the government budget constraint. In A. Razin and E. Sadka (Eds.), Economic Policy in Theory and Practice. London: MacMillan Press.

Webb, S., (1986). Fiscal News and Inflationary Expectations in Germany After World War I. Journal of Economic History, XLVI, 3, Sept., 769-794.

White, H., (1989). Some Asymptotic Results for Learning in Single Hidden-Layer Feedforward Network Models. Journal of the American Statistical Association, 84, Dec., 1003-1013.

White, H., (1990). Connectionist Nonparametric Regression: Multilayer Feedforward Networks Can Learn Arbitrary Mappings. Neural Networks, 3, 535-549

Woodford, M. (1990). Learning to Believe in Sunpsots. Econometrica, 58, 277-307.

Appendix 1:

TABLE 1. SELECTED HYPERINFLATIONS: KEY FACTS

Country	Average Monthly Inflation Rate	Peak Monthly Inflation Rate (Date)	Number of Months with Inflation >50% (>25%)
Austria	47.1	134	4
		(8/1922)	(10)
Germany	322	32400	11
		(10/1923)	(20)
Poland	81.4	275	9
		(10/1923)	(16)
Russia	57.0	213	10
		(1/1924)	(18)
Bolivia	48.1	182	9
		(2/1985)	(16)
Argentina	66	200	3
		(7/1989)	(16)
Brazil	19.7	73	3
		(1/1990)	(16)
Ukraine	30.6	90.8	6
		(12/1993)	(5)
Georgia	44.1	211.1	9
		(9/1994)	(4)
Bulgaria	24.5	242.8	1
		(2/1997)	(4)

Sources: Own calculations, Bruno (1990), Cagan (1956)

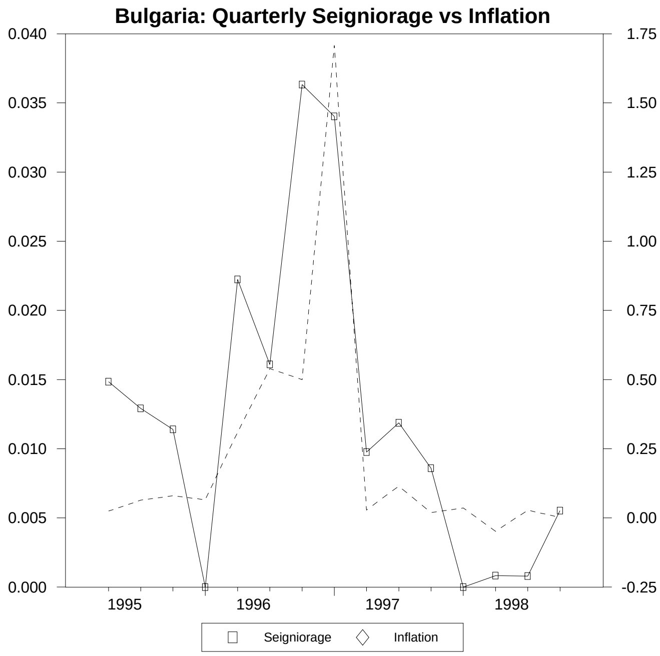

Figure 1: Bulgaria and Russia: Seigniorage and Inflation in the 1990s. $\frac{MB_t - MB_{t-1}}{GDP_t}$ is the change in the monetary base as a percentage of GDP (left axis), and $\log(P_t/P_{t-1})$ is inflation (right axis)

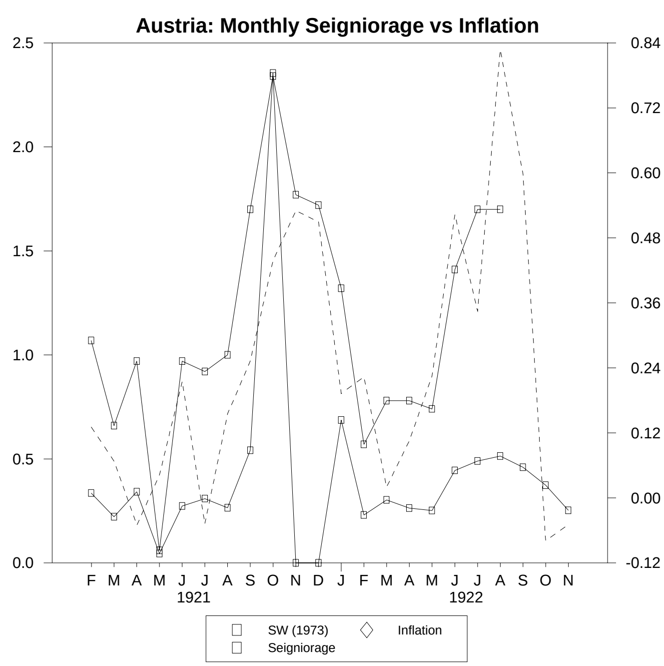

Figure 2: Ukraine in the 1990s and Austria in the 1920s

SW (1973) stands for seigniorage series found in Sargent and Wallace (1973, Table 6). This quantity is computed as $\frac{M_t - M_{t-1}}{0.5(P_t + P_{t-1})}$ (left axis) and $\Delta M/P$ as seigniorage (right axis).

TABLE 2. CLASSIFICATION OF STABILIZATION PROGRAMS

	Orthodox	Heterodox
Money-based	**Russia (1994/1995)** **Ukraine (1994/1995)**	Latvia (1992) Lithuania (1993) Peru (1990) Bulgaria (1991/1992) Romania (1991/1992) Slovenia (1992)
Exchange-Rate Based	Argentina (1978) Chile (1978) Uruguay (1978) Argentina (1991) Bolivia (1985) **Bulgaria (1997)**	Argentina (1985) Brazil (1986, 1987, 1989) Czechoslovakia (1991) Estonia (1992) Israel (1985) Mexico (1987) Peru (1985) Poland (1990)

Sources: Own and Bofinger (1996)

Appendix 2:

First, note that if a stationary rational expectations equilibrium (SREE) exists, then ERR is never established. Assuming $\pi^{e}_{t+1} = \pi_{t+1}$, i.e., rational expectations, and given $\pi_t \in (0, \frac{\gamma_1}{\gamma_2})$ to rule out zero or negative real money balances, we can solve Equation. (4) to obtain:

$$\pi_{t+1} = \frac{\gamma_1}{\gamma_2} + 1 - \frac{d_t}{\gamma_2} - \frac{\gamma_1}{\gamma_2}\frac{1}{\pi_t} \tag{A1-1}$$

Suppose d is constant, there is no uncertainty, and $d < d^{\max} = \gamma_2(1 + \frac{\gamma_1}{\gamma_2} - 2\sqrt{\frac{\gamma_1}{\gamma_2}})$, then

the two SREE solutions are: $\pi^{*}_{RRE1/2} = \frac{1}{2}\left[(\frac{\gamma_1}{\gamma_2} + 1 - \frac{\bar{d}}{\gamma_2}) \pm \sqrt{(\frac{\gamma_1}{\gamma_2} + 1 - \frac{\bar{d}}{\gamma_2})^2 - 4\frac{\gamma_1}{\gamma_2}}\right]$.

As discussed in the introduction and shown in both Sargent and Wallace (1981) and Marcet and Sargent (1989), the high stationary REE is stable and this implies that raising the seigniorage level causes the low REE to increase and high RRE decrease (see Fig. 3 and 4 below in Appendix 3). Note that when d=0, there is no inflation at the lower SREE and π^{*}_{RRE2} is bound by $\frac{\gamma_1}{\gamma_2}$. Marcet and Nicolini (2003) show in addition that these SRRE are outside the interval or close to the rational expectation solutions of the model when $\sigma_d \neq 0$.

Appendix 3:

Figure 3: Equilibrium Dynamics of Inflation with steady state rational expectations solutions: high π_{REE2} and low π_{REE1}.

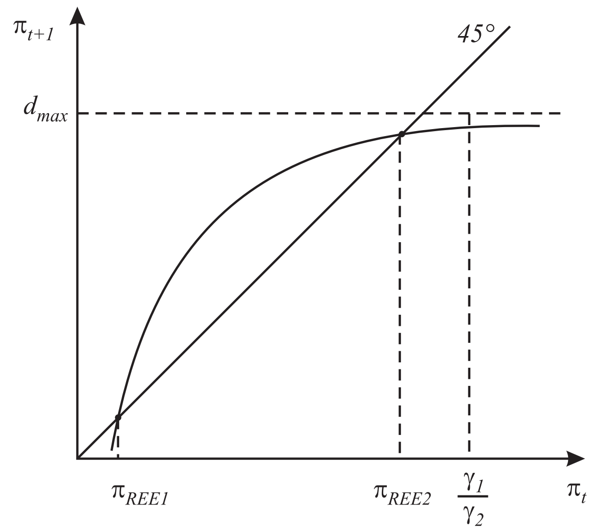

Figure 4: Equilibrium Dynamics of Inflation with high average seigniorage

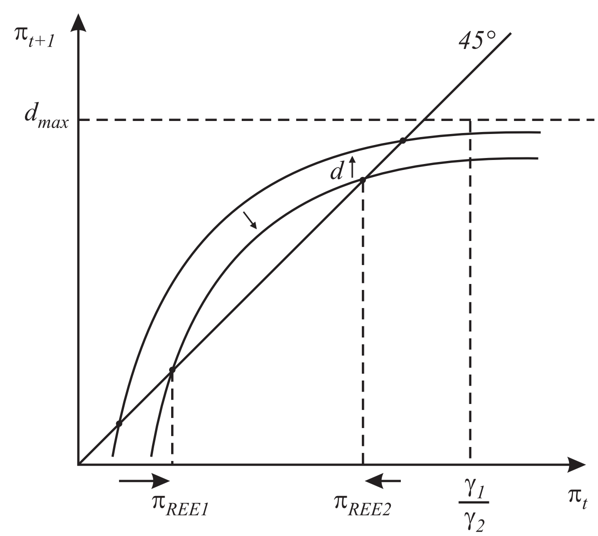

Figure 5: Equilibrium Dynamics of the model in Section 2

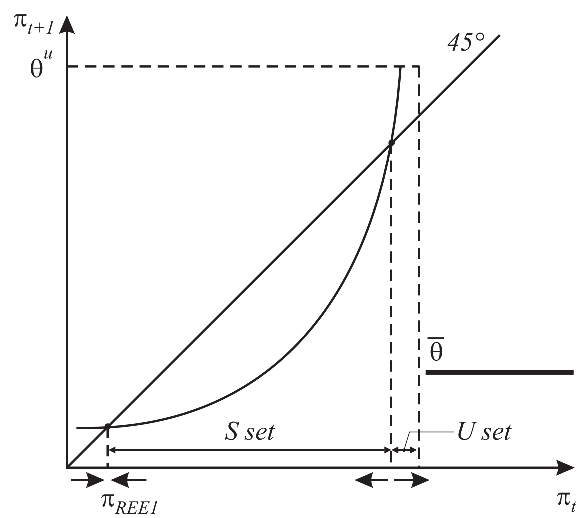

Figure 6: Equilibrium Dynamics of Inflation when seigniorage is high

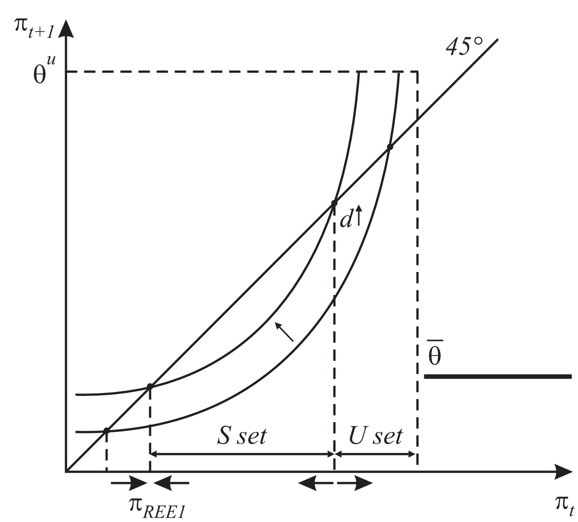

Appendix 4:

TABLE 3: SET OF PARAMETERS FOR SIMULATIONS: OLG ECONOMIES WITH VARIABLE REAL DEFICIT FINANCED THROUGH SEIGNIORAGE

Parameters	1	2	3	4
γ_1	5	1	100	0.37
γ_2	2	0.9	50	0.148
$\overline{d}$	0.01	0.0019	0.05	0.045
π^*_{REE1}	1.0033	1.026	1.001	1.37
π^*_{REE2}	2.4917	1.0838	1.998	1.829
π^{LS}	**1.003**	1.4189*	**1.0010**	**1.3757**
π^{NN}	**1.003**	1.4042*	**1.0010**	**1.4134**
γ_1	5	1	100	0.37
γ_2	2	0.9	50	0.148
$\overline{d}$	0.5	0.0024	0.075	0.049
π^*_{REE1}	1.25	1.0377	1.0015	1.48
π^*_{REE2}	2.00	1.0707	1.997	1.69
π^{LS}	**1.249**	1.3566*	**1.0015**	1.984*
π^{NN}	**1.250**	1.3863*	**1.0015**	1.855*
Parameters	5	6	7	8
γ_1	2.81	0.898	3.144	0.71
γ_2	1.67	0.399	1.396	0.28
$\overline{d}$	0.06	0.095	0.01	0.045
π^*_{REE1}	1.06	1.37	1.0058	1.1289
π^*_{REE2}	1.59	1.64	2.239	2.246
π^{LS}	**1.0618**	**1.3652**	**1.007**	**1.1321**
π^{NN}	**1.0695**	**1.375**	**1.0107**	**1.1292**
γ_1	2.81	0.898	3.144	0.71
γ_2	1.67	0.399	1.396	0.28
$\overline{d}$	0.145	0.098	0.349	0.097
π^*_{REE1}	1.25	1.42	1.47	1.51
π^*_{REE2}	1.34	1.586	1.53	1.679
π^{LS}	**1.2387**	**1.416**	**1.4070**	**1.52**
π^{NN}	1.3874*	**1.4224**	**1.445**	**1.5090**

Note: * indicates divergence or no equilibrium; bold type shows convergence.

TABLE 4: EFFICIENT VALUES OF $\frac{1}{\overline{\alpha}}$ IN EQUATION (11): *MONTE CARLO* MEAN SQUARE ERRORS (MSE) - γ_1 =0.898 AND γ_2 =0.399

LEAST SQUARES LEARNING: AVERAGE SEIGNIORAGE= 0.098

	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.1
1.1	0.00	0.01	0.01	0.07	0.05	1.88	20.16	17.26	6.85	4.58	2.29	9.15
1.0	0.01	0.01	0.01	0.06	0.05	1.73	18.26	15.69	6.30	4.20	2.11	8.24
0.9	0.01	0.01	0.01	0.06	0.05	1.60	16.68	14.31	5.85	3.89	1.95	7.46
0.8	0.00	0.01	0.01	0.06	0.05	1.49	15.35	13.11	5.47	3.62	1.81	6.80
0.7	0.00	0.00	0.01	0.06	0.05	1.40	14.24	12.06	5.14	3.38	1.68	6.21
0.6	0.00	0.00	0.01	0.06	0.05	1.31	13.30	11.13	4.82	3.15	1.54	5.71
0.5	0.00	0.00	0.01	0.06	0.04	1.24	12.47	10.31	4.50	2.94	1.41	5.25
0.4	0.00	0.00	0.01	0.05	0.04	1.17	11.73	9.58	4.17	2.73	1.30	4.84
0.3	0.00	0.00	0.00	0.05	0.04	1.10	11.04	8.93	3.84	2.54	1.19	4.47
0.2	0.00	0.00	**0.00**	0.05	0.04	1.03	10.36	8.35	3.51	2.35	1.10	4.14
0.1	0.00	**0.00**	0.01	0.04	0.03	0.94	9.69	7.86	3.20	2.16	1.02	3.85
0	**0.00**	0.15	0.33	0.20	0.59	1.40	3.41	4.68	3.64	4.37	2.60	7.51

Note: Rows represent the alternative values for $\frac{1}{\overline{\alpha}'}$; columns give the possible values for $\frac{1}{\overline{\alpha}}$ used in the model. Bold type indicates the range of efficient values of $\frac{1}{\overline{\alpha}}$.

NEURAL NETWORK (BACK-PROPAGATION) LEARNING: AVERAGE SEIGNIORAGE= 0.098

	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.1
1.1	0.00	0.01	0.00	0.01	0.14	0.57	1.49	5.03	15.70	6.32	21.83	4.97
1.0	0.00	0.01	0.00	0.01	0.13	0.51	1.34	4.43	14.09	5.70	19.56	4.29
0.9	0.00	0.01	0.00	0.01	0.11	0.48	1.21	4.02	12.77	5.21	17.69	3.76
0.8	0.00	0.00	0.00	0.01	0.10	0.46	1.12	3.66	11.70	4.79	16.15	3.33
0.7	0.00	0.00	0.00	0.01	0.10	0.43	1.04	3.66	11.68	4.44	14.83	3.28
0.6	0.00	0.00	0.00	0.01	0.10	0.39	0.98	3.32	11.58	4.34	15.92	2.67
0.5	0.00	0.00	0.00	0.01	0.09	0.37	0.92	3.25	10.33	4.24	14.08	2.55
0.4	0.00	0.00	0.00	0.01	0.08	0.35	0.87	2.85	9.19	3.78	12.41	2.43
0.3	0.00	0.00	0.00	**0.01**	0.08	0.34	0.82	2.66	8.62	3.44	11.28	2.08
0.2	0.00	0.00	**0.00**	0.01	0.07	0.31	0.77	2.47	7.83	3.26	10.67	1.97
0.1	0.00	**0.00**	0.00	0.01	0.07	0.28	0.70	2.32	7.36	3.00	10.02	1.82
0.0	0.01	0.06	0.31	1.10	2.21	1.54	2.02	1.63	6.53	2.48	7.32	4.27

Note: Rows represent the alternative values for $\frac{1}{\overline{\alpha}'}$; columns give the possible values for $\frac{1}{\overline{\alpha}}$ used in the model. Bold type indicates the range of efficient values of $\frac{1}{\overline{\alpha}}$.

Figure 7: Inflation Rate: $\ln(\frac{p_{t+1}}{p_t})$. OLG economy set 6 of Parameters in Table 1 with average seigniorage 0.098 and $\frac{1}{\overline{\alpha}} = 0.2$ – least squares learning

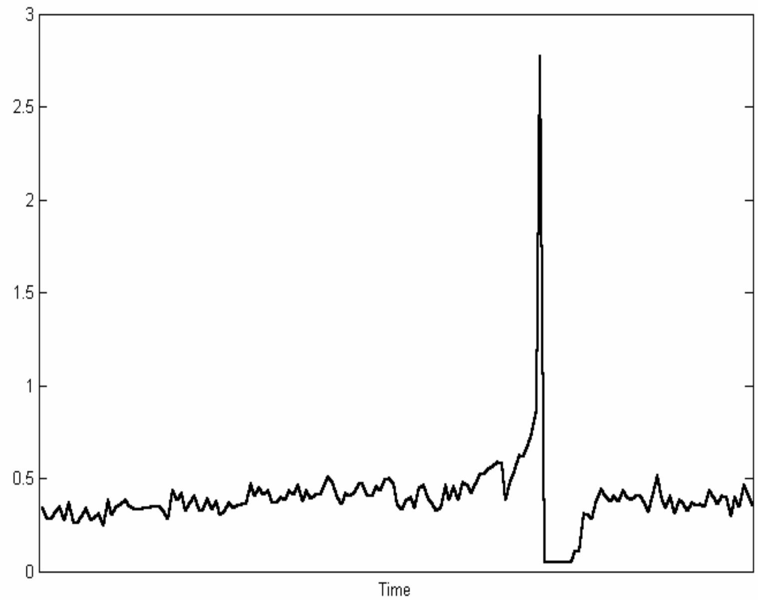

Figure 8: Inflation Rate: $\ln(\frac{p_{t+1}}{p_t})$. OLG economy set 6 of Parameters in Table 1 with average seigniorage 0.098 and $\frac{1}{\overline{\alpha}} = 0.2$ – neural network learning

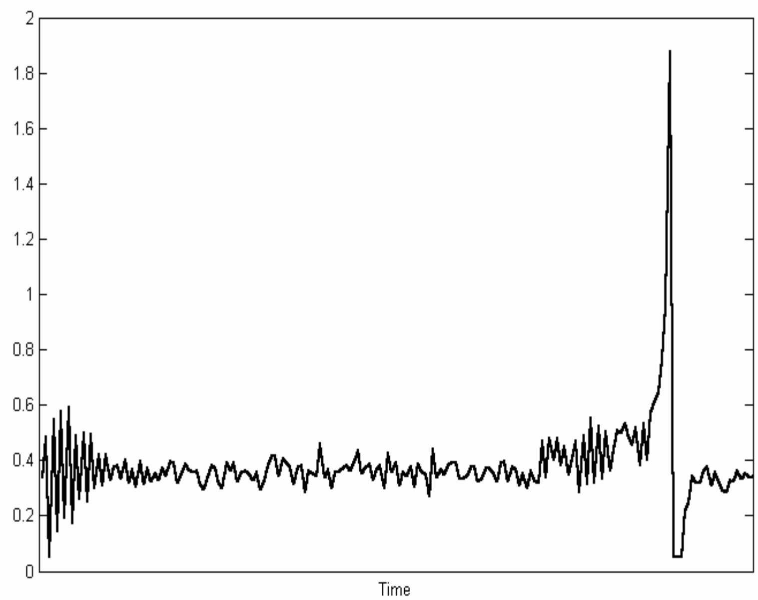

Figure 9: Inflation Rate: $\ln(\frac{p_{t+1}}{p_t})$. OLG economy set 8 of Parameters in Table 1 with average seigniorage 0.097 and $\frac{1}{\overline{\alpha}} = 0.2$ – least squares learning

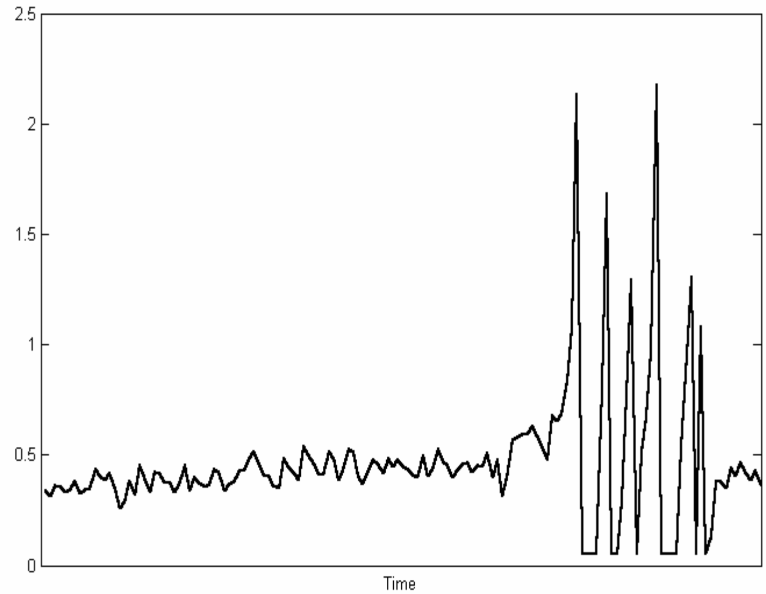

Figure 10: Inflation Rate: $\ln(\frac{p_{t+1}}{p_t})$. OLG economy set 8 of Parameters in Table 1 with average seigniorage 0.097 and $\frac{1}{\overline{\alpha}} = 0.2$ – neural network learning