

Recurrent Hyperin^oations and Learning

Albert Marcet

Universitat Pompeu Fabra (Barcelona), CREI and CEPR
and

Juan Pablo Nicolini[□]

Universidad Torcuato di Tella (Buenos Aires)

November 23, 2001

Abstract

This paper uses a model of boundedly rational learning to account for the observations of recurrent hyperin^oations in the last decade. We study a standard monetary model where the fully rational expectations assumption is replaced by a formal de⁻inition of quasi-rational learning. The model under learning is able to match remarkably well some crucial stylized facts observed during the recurrent hyperin^oations experienced by several countries in the 80's. We argue that, despite being a small departure from rational expectations, quasi-rational learning does not preclude falsi⁻ability of the model, it does not violate reasonable rationality requirements and it can be used for policy evaluation.

Keywords: Hyperin^oations, convertibility, stabilization plans, quasi-rationality.

JEL classi⁻cation: D83, E17, E31.

[□]We thank Tony Braun, Jim Bullard, George Evans, Seppo Honkapohja, Rodi Manuelli, Ramon Marimon, Tom Sargent, Stacia Sowerby, Harald Uhlig, Neil Wallace and Carlos Zarazaga for helpful conversations and Marcelo Delajara and Ignacio Ponce Ocampo for research assistance. All errors are our own. Part of this work was done when both authors were visiting the Federal Reserve Bank of Minneapolis and, specially, when Marcet was visiting CEMFI. Research support from DGICYT, CIRIT, HCM and CREI is greatly appreciated.

1 Introduction

The goal of this paper is to develop a model that accounts for the main features of the hyperinflation of the 80's and to study the policy recommendations that arise from it. The model is standard, except for the assumption of quasi-rational learning. Modern macroeconomics has been reluctant to use boundedly rational expectations models to match empirical observations. It is commonly believed that such models are not falsifiable and expectations are not consistent with the model. This view is stated clearly in the following quotation from Sargent (1993): "... the literature on adaptive decision processes seems to me to fall far short of providing a secure foundation for a good theory of real-time transition dynamics. There are problems of arbitrariness and the need for prompting, with a concomitant sensitivity of outcomes to details of adaptive algorithms". A side contribution of the paper is to show with an example that, contrary to Sargent's statement, if certain rationality requirements are imposed, learning models can be useful to understand real-time transition dynamics.

The long run relationship between money and prices is a well understood phenomenon. The price level and the nominal quantity of money over real output hold an almost proportional relationship so that the inflation rate is essentially equal to the growth rate of money supply minus the growth rate of output. There is widespread consensus in the profession that successfully stopping inflation involves substantial reductions in money growth rates. On the other hand, long periods of high money growth rates are associated with large seignorage collection required to finance government deficits. A simple story about hyperinflation is often told: when the government is unable to either reduce its fiscal deficit or finance it through the capital market, high seignorage is required and high inflation rates are unavoidable. This is the logic behind the IMF advice to countries experiencing high inflation rates. Cross country evidence very strongly supports this story. Hyperinflation have occurred in countries with high seignorage, and many countries that successfully stopped inflation did so by eliminating the fiscal imbalance that required high seignorage.

However, this simple story fails when we closely look at time series of inflation and seignorage for very high inflation countries. Countries that undergo very rapid price increases typically exhibit periods of relatively high but stable inflation rates, followed by a sudden explosion in the rate of inflation; this often happens without any important change in the level of seignorage. We observe inflation rates multiplying by 8 or 10 in a couple of months while seignorage remains roughly the same or even decreases. This could challenge the validity of the IMF advice to hyperinflationary countries to decrease their

seignorage.

In this paper we develop a model that accounts for this and other crucial observations that occurred during the hyperin^oations of the 80's. These episodes involve very high in^oation rates (for instance, in^oation in Argentina in June 89 peaked at 200% a month) and all we know about the welfare e^oects of in^oation suggest that they are very costly.

Sargent and Wallace (1987) explained these hyperin^oations as bubble equilibria. Their model generates a standard La^oer curve with two stationary rational expectations equilibria; hyperin^oations could occur as speculative equilibria converging to the high-in^oation steady state. Their paper explains how in^oation can grow even though seignorage is stable; but it fails to explain other facts observed in the hyperin^oationary episodes. Our work builds upon Sargent and Wallace's by introducing learning; we show that, with this modi^ocation, the model matches observations much better. Our model is consistent with the very high hyperin^oations, their recurrence, the fact that exchange rate rules temporarily stop hyperin^oations, the cross country correlation of in^oation and seignorage, and the lack of serial correlation of seignorage and in^oation in hyperin^oationary countries.

The last decade has witnessed a renewed interest in learning models in macroeconomics. This literature focussed on limiting properties, studying convergence of learning to rational expectations¹. This literature has made enormous progress, and convergence of learning models to rational expectations can now be studied in very general setups. Nevertheless, few attempts have been made to explain observed economic facts with models of boundedly rational learning: among others, Arifovic, Bullard and Du^oy (1997) and Evans and Honkapohja (1993) have compared the overall behavior of learning models with some general features of the data; Timmermann (1993,1996) and Chung (1990) have attempted to match concrete facts in set ups very di^oerent from ours. However, with the partial exception of Evans and Honkapohja (see our discussion following De^oinition 3), none of them formally addressed the critique to boundedly rational models that is commonplace in today's macro literature and that is clearly stated in the above quote from Sargent. This critique is based on the nowadays standard view in macroeconomics that using models of boundedly rational learning would entail problems similar to those found in models of adaptive expectations of the pre-rational-expectations era, namely: i) there are too many degrees of freedom available to the economist so that the model is not falsi^oable, ii) agents' expectations are inconsistent with the model, and iii) the model does not predict how expectation formation will change if there is a change in policy.

¹See Sargent (1993), Marimon (1997) and Evans and Honkapohja (1999, 2001) for reviews.

We address these criticisms by restricting the learning mechanisms to produce good forecasts within the model. We only consider learning mechanisms that produce small departures from rationality within the model, in a way that is precisely defined in the paper. We show that the model has empirical content and that expectations are endogenous to policy.²

Some papers have presented models that explain some of the facts we consider. Eckstein and Leiderman (1992) and Bental and Eckstein (1996) explain the very large inflation rates in Israel with an ever increasing Laffer curve. Zarazaga (1993) develops a model of endogenous seignorage, where spikes in inflation can happen because of moral hazard in the demands for revenue of several branches of government. These papers account for some of (but not all) the facts we describe in the paper. Their stories could be combined with the story of the current paper.

The paper is organized as follows. Section 2 presents the stylized facts and provides supporting evidence. Section 3 presents the model and characterizes rational expectations equilibria. Section 4 discusses the lower bounds in rationality in a general setup. Section 5 discusses the behavior of the model under the lower bounds on rationality. The paper ends with some concluding remarks.

2 Evidence on Recurrent Hyperinflation

A number of countries, including Argentina, Bolivia, Brazil and Peru, experienced during the eighties the highest average inflation rates of their history. While the duration and severity of the hyperinflation and the policy experiments differ substantially, there are several stylized facts that are common to those experiences (and, to some extent, to those of some European countries after the first world war, and those of East European countries after the end of the cold war). These stylized facts are

1. Recurrence of hyperinflationary episodes. Time series show relatively long periods of moderate and steady inflation, and a few short periods of extremely high inflation rates.
2. Exchange rate rules (ERR) stop hyperinflation. In most circumstances, however, these plans only lower inflation temporarily, and new hyperinflation eventually occur.

²Recent literature imposing consistency requirements in learning models are Evans and Honkapohja (1993), Kurz (1994), Fudenberg and Levine (1995) and Hommes and Sorger (1998).

3. For a given country where hyperinflation occurs, the contemporaneous correlation across time between seignorage and inflation is low.
4. Average inflation and seignorage are strongly correlated across countries. Hyperinflation only occurs in countries where seignorage is high on average.

Points 2 and 4 can be combined to state the following observation on monetary policy: stabilization plans based on ERR -"heterodox" policy - that do not permanently reduce average seignorage, may be successful in substantially reducing the inflation rate only in the short run. Some stabilization plans not only relied on the fixing of the exchange rate but also permanently reduced the deficit -"orthodox" policy- and the need for seignorage. It is now relatively well accepted that this combination of both orthodox and heterodox ingredients has been successful at stopping hyperinflation permanently. To our knowledge, ours is the first economic model that satisfactorily explains the above facts and is consistent with this policy recommendation.

Our summary of stylized facts should be uncontroversial³, but first-hand evidence to support them is provided in Figure 1, which presents data on the recent inflationary experiences of Argentina, Bolivia, Brasil and Peru. Inflation rates were computed from IFS consumer price indices. Periods when an explicit fixed ERR was in place are indicated by shaded areas. The end of the shading indicates the date in which the ERR was explicitly abandoned. Figure 1 illustrates quite clearly stylized facts 1 and 2.

Figure 2 depicts the quarterly inflation rate for Argentina together with the seignorage as a share of GDP for the period 1982 to 1990⁴. Note that while seignorage is between two and eight percent of GDP, inflation ranges from almost zero to 300% a quarter. The Figure shows, for instance, that the level of seignorage leading to the spectacular hyperinflation at the end of 1989 (more than 250% a quarter) is very similar to the one of the second quarter of 1984, with subsequent inflation rates that were around 60%. Also, note that in the second half of 1984 seignorage and inflation were going in opposite directions. This documents fact 3.⁵

³See Bruno et al. (1988) and (1991).

⁴Given the banking regulation during the period, there is some debate regarding the right monetary aggregate to compute the resources raised by the government. We report seignorage computed using the monetary base, but a similar picture arises if we use M1. Figure 2 reports, for both series, four-period averages (using current and the three previous quarters) to eliminate the strong seasonal movements that obscure the picture.

⁵A closer look at Figure 2 points to some interesting facts that merit a more careful empirical investigation. Note, in particular, that seignorage appears to lead the hyperinflationary bursts. Also, there is some correlation between inflation and seignorage in the

3 The Model

3.1 Economic Fundamentals

The assumptions in this subsection are standard. The model consists of a portfolio equation for the demand of real money balances, a government budget constraint relating money creation and changes in reserves, and a rule for establishing fixed exchange rates.⁶

Money demand

The demand for real balances is given by

$$\frac{M_t^d}{P_t} = \begin{cases} \bar{A} \left(\frac{P_{t+1}^e}{P_t} \right)^{\alpha} & \text{if } 1 - \alpha \left(\frac{P_{t+1}^e}{P_t} \right) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $\bar{A} > 0$ are parameters, $P_t; M_t^d$ are price level and nominal demand of money; P_{t+1}^e is the price level that agents expect for next period.

Money supply

We assume that money creation is driven by the need to finance seignorage. On the other hand, government's concern about current levels of inflation prompts the adoption of ERR when inflation gets out of hand or to restore equilibrium.

In periods that the ERR is not in place, the government budget constraint is given by

$$M_t = M_{t-1} + d_t P_t \quad (2)$$

Seignorage is given by an exogenous i.i.d. stochastic process d_t with mean $E(d_t)$ and variance σ_d^2 , and it is the only source of uncertainty in the model⁷. Equations (1) and (2) plus a hypothesis of expectations formation, determine the equilibrium values for $M_t; P_t$ in periods of floating exchange rates.

Exchange Rate Rules

sub samples periods when inflation was not too high. Both of these features are consistent with our model but they are not studied carefully in this version of the paper.

⁶Appendix 1 shows that the following equations can be rationalized as the equilibrium conditions of an OLG monetary model of a small open economy.

⁷The i.i.d. assumption is made for simplicity. For example, if d_t were a Markov process, P_{t+1}^e would have to depend on d_t for the learning scheme to satisfy the lower bounds on rationality, and agents would have to learn about at least two parameters. It would be interesting to generalize the model to this case, specially since seignorage is, indeed, serially correlated in the data. We conjecture that the main results of the paper would go through with serially correlated seignorage, but some analytical results would be harder to prove.

In periods of ERR, the government pegs the nominal exchange rate by buying or selling foreign reserves at an exchange rate e_t satisfying

$$\frac{P_t^f}{P_{t-1}^f} \frac{e_t}{e_{t-1}} = \bar{\pi};$$

where $\bar{\pi}$ is the targeted inflation rate, and P_t^f is the price level abroad. Assuming full mobility of goods, purchasing power parity implies

$$\frac{P_t}{P_{t-1}} = \bar{\pi} \quad (3)$$

and the targeted inflation rate is achieved. In the case that targeted inflation $\bar{\pi}$ is the same as foreign inflation, the government announces a fixed exchange rate. Otherwise, a crawling peg is followed.

Under ERR, equilibrium price level is determined by (3). Given this price level and an expectations hypothesis, (1) determines money demand. In general, this money demand will not match money supply as determined by (2). As it is standard in fixed exchange rate models, international reserves (denoted below R_t) adjust so the right level of money balances is achieved. Thus, instead of (2) the following equation holds in periods of ERR:

$$M_t = M_{t-1} + d_t P_t + e_t (R_t - R_{t-1}); \quad (4)$$

Finally, we impose the rule that government acts to satisfy

$$\frac{P_t}{P_{t-1}} \leq \bar{\pi}^u; \quad (5)$$

where $\bar{\pi}^u$ is the maximum inflation tolerated. ERR is only imposed in periods when inflation would otherwise violate this bound or in periods where no positive price level clears the market if $R_t = R_{t-1}$.⁸

Our model makes the implicit assumption that ERR can always be enforced. In fact, governments may run out of foreign reserves, and they may be unable to enforce ERR for a sufficiently long period. Hence, we are making the implicit assumption that the non-negativity constraint on foreign reserves is never binding. Since we will choose the target inflation rate $\bar{\pi}$ to be the lower stationary rational expectations equilibrium steady state inflation, the loss of reserves is likely to be small. Modelling reserve accumulation formally is unlikely to change our main results, but it opens up a host of interesting

⁸It may be that no positive price level clears the market, for example, if perceived inflation is too high. See Section 5.2 for details.

issues. For example, the government may run out of reserves during a hyperinflation, so that "orthodox" measures can not be avoided, a feature that is consistent with our model. Alternatively, by increasing the length of the ERR after a hyperinflation the monetary authority could accumulate reserves since the real value of the money stock is increased after the stabilization⁹.

We have modelled policy in this way because it mimics the broad features of policies followed by South-American countries during the 80's. The issue of why these countries followed this kind of policy is not addressed formally in this paper, but we can advance three possible justifications for using this rule in our model. First, the fact that ERR has been established only after some periods of high inflation is justified because then the value of foreign reserves is high, and a large part of the domestic money can be backed with existing reserves¹⁰. Second, in principle, any reduction in the government deficit of $e_t(R_t - R_{t-1})$ units would also fix the inflation to $\bar{\pi}$ in periods of ERR. In fact, the reduction in seignorage that is needed to achieve an inflation equal to $\bar{\pi}$ is often quite moderate, which raises the issue of why governments have used ERR instead of lowering the fiscal deficit (and seignorage) sufficiently. One possible answer is that the exact value of $e_t(R_t - R_{t-1})$ can only be inferred from knowledge of the true model and all the parameter values, including those that determine the (boundedly rational) expectations P_{t+1}^e ; and all the shocks. By contrast, an ERR can be implemented only with knowledge of the foreign price level and the policy parameters $(\bar{d}; \bar{u})$. A third advantage of establishing ERR for real governments would be the existence of institutions that can implement this measure quickly, while lowering government expenditures or increasing taxes often takes a long time.

An important policy decision is how long to maintain the ERR. Obviously, the longer the ERR is maintained, the closer expected inflation will be to $\bar{\pi}$: In our simulations, we hold the ERR till expected inflation is close to $\bar{\pi}$ in a sense to be made precise below.

In summary, the government in our model sets money supply to finance exogenous seignorage; if inflation is too high, the government establishes ERR. The parameters determining government policy are \bar{d} , \bar{u} and the distribution of d_t :

⁹For instance, Central Bank reserves grew, in Argentina, from 1991 (year in which the Convertibility plan was launched) to 1994 from 500 million dollars to more than 12 billion.

¹⁰This interpretation would suggest that the burst in inflation at the beginning of 1991 in Argentina was crucial for the success of the Convertibility Plan launched in April of the same year, because it substantially reduced the value of the money stock to a point where, at a one dollar=one peso exchange rate, the government could back the whole money stock.

3.2 The Model under Rational Expectations.

If we assume that agents form expectations rationally, the model is very similar to that of Sargent and Wallace (1987) (SW from now on). As long as seignorage is not too high, the model has two stationary equilibria with constant expected inflation levels (called low- and high-inflation equilibria), and a continuum of bubble equilibria that converge to the high-inflation equilibrium¹¹.

The main motivation behind the work of SW was to explain 'fact 3' in section 2 as rational bubble equilibria¹². Their original model does not allow for recurrence of hyperinflations (fact 1), but the work by Funke et al. (1994) shows that recurrence can be explained by introducing a sunspot that turns rational bubbles on and off. Even if one accepts rational sunspots as an explanation, fact 1 is not matched quantitatively: for reasonable parameter values, the magnitude of the hyperinflations that can be generated with this model is very small¹³. Fact 4 is contradicted: the long run inflation rate in any rational bubble equilibrium is lower when seignorage is higher, so the model under RE predicts that hyperinflations are less severe in countries with high seignorage.

The papers of Obstfeld and Rogoff (1983) and Nicolini (1996) introduce ERR that goes into effect if inflation goes beyond a certain level and, therefore, these papers can be used to address fact 2. Their results show that just the threat of convertibility eliminates bubble equilibria altogether and that the ERR, under rational expectations equilibria, never takes place. Thus, once ERR is introduced, the rational expectations equilibrium is inconsistent with the existence of hyperinflations, since convertibility was certainly a credible threat and hyperinflations were indeed observed in the 80's.

Marcet and Sargent (1989b) studied stability of rational expectations equilibria in the SW model under least squares learning. They found that the low-inflation equilibrium is locally stable and the high-inflation equilibrium is always unstable. Taken literally, these results would say that bubble equilibria can not be learned by agents. Therefore, none of the above facts is appropriately matched if we restrict our attention to rational expectations equilibria that are stable under learning.¹⁴

¹¹We reproduce these results in appendix 2. As our model is, contrary to SW, stochastic, some of the results are slightly different.

¹²There has been some work testing the existence of rational bubbles in the German hyperinflation of the twenties. A summary of the literature and a test of bubble versus stationary equilibria in the SW model can be found in Imrohorglu (1993).

¹³This is documented in our discussion of Figure 4 in subsection 5.5 below.

¹⁴Marcet and Sargent (1989b) is a special case of the present paper when uncertainty is eliminated, σ^u is arbitrarily high, and agents forecast P_i by regressing it on P_{i-1} . These

In the next section we propose several criteria to assess models with quasi-rational learning and to address the criticisms of learning models commonly found in the literature.

4 Learning and Lower Bounds on Rationality

Before the rational expectations revolution, economic agents' expectations were specified in macroeconomics according to ad-hoc assumptions; one popular alternative was 'adaptive expectations'. This practice was criticized because: i) it introduced too many degrees of freedom in the specification of expectations so it made the models less falsifiable and, ii) agents' expectations were inconsistent with the model; then rational agents would be likely to abandon their ad-hoc expectations after a while, and the predictions of the model would be invalid. Related to this point, the model does not say if expectations will change when policy changes. The first criticism is hyperbolized by the sentence: 'any economic model can match any observation by choosing expectations appropriately'; the second criticism is typified by the sentence 'economic agents do not make systematic mistakes'.

The rational expectations hypothesis solved these two issues: under RE, expectations are determined by the model and, after some time, agents will just realize that their beliefs are right.

In this paper we use a boundedly rational learning model to explain stylized facts, so a natural question is: are we slipping into a use of learning models that is as objectionable as, say, adaptive expectations?

The term boundedly rational learning (which, in this paper, we use as synonymous with the term learning) is used to denote learning mechanisms that place upper bounds on rationality. For example, agents are assumed not to know the exact economic model or to have bounded memory. But this admits too many models of learning. Indeed, once we rule out RE, anything can be a boundedly rational learning scheme and we could be falling back into old mistakes and the 'wilderness of irrationality'¹⁵.

Our approach is to allow for only small deviations from rationality, both along the transition and asymptotically. Given an economic model we only admit learning mechanisms that satisfy certain lower bounds on rationality

authors noted that if inflation goes beyond the high steady state it may enter an unstable region where inflation tends to grow without bound. This feature of the model with learning constitutes the core of the dynamics in the current paper.

¹⁵It might seem that Bayesian learning is a way out of this dilemma, but the literature has recognized many problems with this approach. See, for example, Bray and Kreps (1987), Easley and Rustichini (1995) and Marimon (1997) for descriptions of paradoxes and shortcomings of Bayesian learning about the model.

within this model. In section 5 we will show how this small departure from rationality generates equilibria that are quite different from RE, precisely in the direction of improving the match of empirical observations.¹⁶

4.1 A general framework and quasi-rationality

Let us now be precise about the lower bounds that we place on rationality. Assume that an economic model satisfies

$$x_t = g(x_{t-1}; x_{t+1}^e; \varepsilon_t; \gamma) \quad (6)$$

where g is a function determined by market equilibrium and agents' behavior, x_t contains all the variables in the economy, x_{t+1}^e is agents' expectation of the future value of x ; ε_t is an exogenous shock, and γ is a vector of parameters, including the parameters of government policy and the parameters that govern the distribution of ε_t . For example, in our model, x_t is inflation and real balances, ε_t is seignorage, the function g is given by the demand for money (1), the government budget constraint (2), and the ERR rule, while the vector of parameters γ includes $\rho; \bar{A}; \bar{u}$ and the parameters of the distribution of seignorage:

Assume that agents' expectations are given by

$$x_{t+1}^e = z(\bar{\omega}_t(\gamma); x_t) \quad (7)$$

where $\bar{\omega}_t(\gamma)$ is a vector of statistics inferred from past data and z is the forecast function. The statistics $\bar{\omega}$ are generated by a learning mechanism f and learning parameters γ^1 according to

$$\bar{\omega}_t(\gamma^1) = f(\bar{\omega}_{t-1}(\gamma^1); x_t; \gamma^1) \quad (8)$$

The learning mechanism f dictates how new information on x_t is incorporated into the statistics $\bar{\omega}$. The learning parameters γ^1 govern, for example, the weight that is given to recent information. For now, $(z; f; \gamma^1)$ are unrelated to the true model $(g; \gamma)$; but later in this section we will define bounds on rationality that amount to imposing restrictions on the space of $(z; f; \gamma^1)$ given a model $(g; \gamma)$.

In the context of our model in section 3, the function z will be defined as

$$P_{t+1}^e = \bar{\omega}_t P_t \quad (9)$$

where $\bar{\omega}_t$ is expected inflation, estimated somehow from past data.

¹⁶Easley and Rustichini (1995) and Marimon (1997) also argue that learning can be used for more than a stability criterion.

Equations (6), (7) and (8) determine the equilibrium sequence for given learning parameters θ . Obviously, the process for x_t depends on the parameters θ . This dependence will be left implicit in most of the paper, and we will write x_t only if we want to make the dependence explicit.

Let $\mu_{\theta}^{2;T}$ be the probability that the perceived errors in a sample of T periods will be within $\epsilon > 0$ of the conditional expectation error. Formally:

$$\mu_{\theta}^{2;T} = P \left(\frac{1}{T} \sum_{t=1}^T (x_{t+1} - E_t^1(x_{t+1}))^2 < \frac{1}{T} \sum_{t=1}^T [x_{t+1} - E_t^1(x_{t+1})]^2 + \epsilon^2 \right) \quad (10)$$

where E_t^1 is the true conditional expectation under the learning model.

The first lower bound on rationality we propose is:

Definition 1 Asymptotic Rationality (AR): the expectations given by $(z; \theta; \theta)$ satisfy AR in the model $(g; \theta)$ if, for all $\epsilon > 0$;

$$\mu_{\theta}^{2;T} \rightarrow 1 \text{ as } T \rightarrow \infty;$$

This requires the perceived forecast to be asymptotically at least as good as the forecast from the conditional expectation in terms of sample mean square prediction error. In this case, agents would not have any incentive to change their learning scheme after they have been using it for a sufficiently long time.

AR can be viewed as a minimal requirement in the sense that it only rules out behavior that is inconsistent forever. It rules out, for example, learning mechanisms where a relevant state variable is excluded from the forecasting rule z (this feature would exclude adaptive expectations, for example, if d_t were serially correlated). It is satisfied by least squares learning mechanisms in models where this mechanism converges to RE and certain continuity assumptions are satisfied.¹⁷ Similar concepts can be found in the literature¹⁸.

However, AR admits learning mechanisms that generate very bad forecasts along the transition for very long periods. For example, OLS in a model with recurrent hyperinflation would generate very bad forecasts every time a hyperinflation starts, and their forecast would be worse for each new hyperinflation, because least squares learning gives less and less importance to

¹⁷Perhaps surprisingly, AR excludes many 'rational equilibria' in the terminology of Kurz (1994), which allows for agents to make systematic mistakes forever, as long as these mistakes are not contemplated in the prior distribution.

¹⁸This requirement was implicitly imposed in the literature on stability of RE under learning, where least squares learning was optimal in the limit. Also, AR is related to the $(\epsilon; \delta)$ consistency of Fudenberg and Levine (1995), where agents in a game are required to only accept small deviations from best response asymptotically.

recent events as time goes by, so it would take longer and longer for agents to realize that a hyperinflation is starting.

To use only learning mechanisms that generate good forecasts along the transition we impose the next two lower bounds.

Definition 2 Epsilon-Delta Rationality (EDR): the expectations given by $(z; f; \theta^1)$ satisfy EDR for $(\epsilon; \delta; T)$ in the model $(g; \gamma)$ if:

$$\frac{1}{T} \sum_{t=1}^T |e_t| \leq \epsilon$$

If EDR is satisfied for small $\epsilon; \delta > 0$, agents are unlikely to switch to another learning scheme after period T , even if they were told "the whole truth"¹⁹.

It is only interesting to study EDR for moderately high values of T : the sample mean of the prediction error has no chance to settle down for very low T ; and EDR is unlikely to be satisfied for low $(\epsilon; \delta)$ because of large sampling error. Also, for learning mechanisms that are asymptotically rational, EDR is satisfied for T large enough even if agents made large mistakes along the transition. The precise empirical application that the researcher has in mind should suggest an interesting value for T . For example, in our application below, we choose $T = 10$ years, which is the length of the hyperinflationary period in many of the countries studied.

AR is unambiguously satisfied (there is a yes or no answer), but EDR can only be satisfied in a quantitative way, for certain ϵ and δ .

The next bound on rationality requires agents to use learning parameters θ^1 that are nearly optimal within the learning mechanism f : Denote by $e_t(\theta^1; \theta^0)$ the forecast produced by the learning parameter θ^0 when all agents are using the parameter θ^1 : Formally,

$$e_t(\theta^1; \theta^0) = f(e_{t-1}(\theta^1; \theta^0); x_t^1; \theta^0);$$

Definition 3 Internal Consistency (IC): Given $(g; \gamma)$; the expectations given by $(z; f; \theta^1)$ satisfy IC for (T, ϵ) if

$$E \frac{1}{T} \sum_{t=1}^T |x_{t+1}^1 - z_t(\theta^1; x_t^1)|^2 \leq \epsilon$$

$$\cdot \min_{\theta^0} E \frac{1}{T} \sum_{t=1}^T |x_{t+1}^1 - z_t(\theta^0; x_t^1)|^2 + \epsilon \quad (11)$$

¹⁹Bray and Savin (1986) study whether the learning model rejects the hypothesis of serially uncorrelated prediction errors by assuming that agents run a Durbin and Watson test. That exercise carries the flavor of EDR.

Thus, if IC is satisfied, agents are doing almost as well as possible within the learning mechanism specified after T periods, so that they are likely to stay with θ^1 .²⁰

IC is, in general, more restrictive than AR. In particular, any θ^1 satisfying AR also verifies internal consistency for T large enough. As in the case of EDR, it only makes sense to study IC in the context of 'moderately high' T .

The first two bounds compare the performance of the learning mechanism used by agents relative to an external agent who knows the best prediction that can be computed from knowledge of $(f; \theta^1; z; g; \gamma)$. The bound IC, instead, compares the learning mechanism with forecasts that use the same family of mechanisms f but are allowed to pick alternative parameter values θ^1 : This last bound contains some of the intuition of rational expectations, in the sense of looking for an approximate fixed point in which agents' expectations minimize the errors within the mechanism f . Notice that this restriction will, in general, imply that agents under different policy environments use different learning parameters θ^1 ; so that the learning parameter that satisfies IC is endogenous to the model and to government policy. For example, in our model, agents in high seignorage countries (say, Argentina in the 80's) will use a different learning parameter from agents in low seignorage countries (say, Switzerland). These Definitions can be generalized to more complicated models or to objective functions other than the average prediction error.

Imposing these lower bounds on rationality is our way of relaxing rational expectations while maintaining the requirement that agents do not make mistakes forever. Certainly, agents have a certain amount of forward-looking capabilities under Definitions 2 and 3 but far less than under rational expectations.

Rational expectations can be interpreted as imposing extreme versions of the second and third bounds. Obviously, RE satisfies AR. It would appear that requiring EDR for all $\theta^2; \pm > 0$ and all T is the same as imposing rational expectations, but a careful proof should be worked out. Also, if the REE is recursive, if the appropriate state variables are included in z ; if z is a dense class of functions (for example, polynomials), then imposing IC for any $\theta^2; T$ is then same as rational expectations.

5 Learning Equilibrium

In this section, we propose a learning mechanism f that combines least squares learning with tracking, and we show that it satisfies the three lower

²⁰Evans and Honkapohja (1993) developed a very similar criterion in a different context.

bounds on rationality defined in the previous section.

5.1 The Learning Mechanism

In the model of section 3 with expectations given by (9), we assume that the learning mechanism is given by

$$\bar{y}_t = \bar{y}_{t-1} + \frac{1}{\alpha_t} \sum_{i=1}^n \frac{P_{t,i}}{P_{t,i-1}} (\bar{y}_{t-1} - y_{t-1,i}) \quad (12)$$

for given \bar{y}_0 : That is, perceived inflation \bar{y}_t is updated by a term that depends on the last prediction error²¹ weighted by the gain sequence $1/\alpha_t$: This is a simple version of stochastic approximation algorithm. Equation (12) together with the evolution of the gains $1/\alpha_t$ determines the learning mechanism f in equation (8).

One common assumption for the gain sequence is

$$\alpha_t = \alpha_{t-1} + 1 \quad (13)$$

for $\alpha_1 = 1$. In this case, $\alpha_t = t$; and simple algebra shows that (with $\bar{y}_0 = 0$)

$$\bar{y}_{t+1} = \frac{1}{t} \sum_{i=1}^n \frac{P_i}{P_{i-1}}$$

so that, under (13), perceived inflation is just the sample mean of past inflation or, equivalently, it is the OLS estimator of the mean of inflation.

Another common assumption for the gain sequence is $\alpha_t = \alpha > 1$. These have been termed 'tracking' or 'constant gain' algorithms.²² In this case, perceived inflation satisfies (with $\bar{y}_0 = 0$)

$$\bar{y}_{t+1} = \frac{1}{\alpha} \sum_{i=0}^n \frac{1}{\alpha^i} \sum_{j=1}^n \frac{P_{t,i,j}}{P_{t,i,j-1}}$$

so that past information is now a weighted average of past inflations, where the past is discounted at a geometric rate.²³

²¹As usual in models of learning, we make the convenient assumption that the last observation used to formulate expectations is dated at $t-1$. Including today's inflation in \bar{y}_t would make it even easier for the learning scheme to satisfy the lower bounds and to match the stylized facts, and it would not change the dynamics of the model.

²²Evans and Honkapohja (1993), Sargent (1993) and Chung (1990) also discuss the use of tracking algorithms.

²³In this simple model 'tracking' is equivalent to adaptive expectations with a delay. In a more general model tracking is different from adaptive expectations and it generates

Notice that least squares gives equal weight to all past observations, while tracking gives more importance to recent events. Tracking produces better forecasts when there is a sudden change in the environment, because it adapts more quickly while OLS is known to be a consistent estimator of the mean in stationary setups.

Both alternatives are likely to fail the lower bounds on rationality of section 4 in a model that replicates fact 1, where periods of stability are followed by hyperinflation. Tracking performs poorly in periods of stability because perceived inflation is affected by small shocks even though, in truth, the shocks are i.i.d. and they should not affect today's expected inflation: formally, tracking does not converge to RE and it does not even satisfy AR, while OLS has a chance of converging and satisfying AR.

On the other hand, least squares does not generate 'good' forecasts along a hyperinflation, because it will be extremely slow in adapting to the rapidly changing inflation level. During hyperinflations 'tracking' performs better. Least squares does not satisfy EDR or IC and its performance is likely to worsen as there are more successive hyperinflations.

We will specify a learning mechanism that mixes both alternatives: it will use OLS in stable periods and it will switch to 'tracking' when some instability is detected. This amounts to assuming that agents use an endogenous gain sequence such that, as long as agents don't make large prediction errors, θ_t follows a least squares rule, but in periods where a large prediction error is detected, θ_t becomes a fixed positive value θ_{t-1} as in 'tracking'.²⁴ Formally, the gain sequence follows

$$\begin{aligned} \theta_t &= \theta_{t-1} + 1 && \text{if } \frac{p_{t-1} - p_{t-2}}{p_{t-1}} < \epsilon \\ &= \theta_{t-1} && \text{otherwise} \end{aligned} \quad (14)$$

The learning mechanism is the same whether or not ERR is enforced in a given period. The conventional wisdom that the importance of an ERR is the effect it has on expectations is consistent with the model, since the exchange rate rule has an impact on expectations by its effect on the current price level and by setting the gain factor to its base value θ .

In summary, the gain sequence is assumed to be updated according to OLS in periods of stability, but it uses constant gain (or tracking) in periods

better forecasts. For example, if seignorage is autoregressive of order 1, expected inflation would have to depend on current seignorage in order to satisfy any of the lower bounds on rationality. In that case, tracking would be fundamentally different from adaptive expectations.

²⁴Evans and Ramey (1998) also analyze the properties of a learning mechanism that responds endogenously to the performance of the predictions within the model and within the realization.

of instability. The learning mechanism f is fully described by equations (12) and (14), and the learning parameters β are given by (10, 11):

5.2 Learning and Stylized Facts

The variables we need to solve for are $\frac{P_t}{P_{t-1}}$; π_t ; π_t^e : Simple algebra shows that equilibrium inflation satisfies

$$\frac{P_t}{P_{t-1}} = H(\pi_t; \pi_{t-1}; d_t) \quad (15)$$

where²⁵

$$H(\pi_t; \pi_{t-1}; d_t) = \begin{cases} \frac{1 - \pi_t^e - \pi_{t-1}}{1 - \pi_t^e - \pi_{t-1} - d_t} & \text{if } 0 < \frac{1 - \pi_t^e - \pi_{t-1}}{1 - \pi_t^e - \pi_{t-1} - d_t} < 1 - \pi_t^e \text{ and } 1 - \pi_t^e - \pi_{t-1} > 0 \\ \pi_t^e & \text{otherwise,} \end{cases} \quad (16)$$

Equations (12), (14) and (15) define a system of stochastic, second-order difference equations. Characterizing the solution analytically is unfeasible since the system is highly non-linear.

We now provide some intuition on the behavior of inflation. Let $h(\pi; d) \equiv H(\pi; \pi; d)$: Notice that if $\pi_t = \pi_{t-1}$; then $P_t = P_{t-1} \cdot h(\pi_t; d_t)$; so that letting $S(\pi) \equiv E(h(\pi; d_t))$, we have that $S(\pi_t) = E_t(\frac{P_{t+1}}{P_t})$. Then, the graph of h in Figure 3 provides an approximation of the actual inflation rate as a function of perceived inflation and it can be used to describe the approximate dynamics of the model (see appendix 2 for the properties of $S = E(h(\pi; d_t))$).

The first graph corresponds to a realization for low average seignorage $E(d_t)$. The low rational expectations equilibrium π_{RE}^1 is locally stable under least squares learning (see appendix 2). The horizontal axis can be split into the intervals S , U and ERR .

If $\pi_t \in S$; actual inflation is on average closer to π_{RE}^1 than perceived inflation and the learning mechanism pushes perceived inflation towards π_{RE}^1 : Roughly speaking, S is the stability set of perceived inflation. On the other hand, if perceived inflation is in U ; actual inflation is on average higher than π_t ; perceived inflation tends to increase and a hyperinflation is likely to occur.

²⁵Notice that an ERR will prevail at t if one of the following (mutually exclusive) cases occur:

-Case i): $1 - \pi_t^e - \pi_{t-1} < 0$; which implies $M_{t-1} = 0$; so the budget constraint of the government is incompatible with the demand for real balances unless reserves adjust.

-Case ii): $1 - \pi_t^e - \pi_{t-1} - d_t < 0$ and $1 - \pi_t^e - \pi_{t-1} > 0$ so only a negative price level clears the market, and

-Case iii): none of the above and $\frac{1 - \pi_t^e - \pi_{t-1}}{1 - \pi_t^e - \pi_{t-1} - d_t} > 1 - \pi_t^e$, such that the market generates a level of inflation unacceptable to the government if reserves do not adjust.

Then, when the set ERR is reached, a fixed exchange rule is established and inflation is sent back to S. The economy may end up in the unstable set U due to a number of reasons: a few high shocks to seignorage when $1/\pi_t$ is not yet close to zero, initially high perceived inflation, the second-order dynamics adding momentum to increasing inflation, etc.

It is clear that the model can only generate recurrent hyperinflation if there is uncertainty. Otherwise, once the economy is in the stable set, there is no force to take it out of it. Another difference with the deterministic case is that the stable region S shrinks with higher σ_d^2 (see part 4 of proposition 3 in appendix 2). Hence, for a high variance of seignorage, the probability of hyperinflation is high for two reasons: i) a shock large enough to send π_t to the unstable region U is more likely to occur and, ii) the stable region S shrinks.²⁶

If a shock to inflation occurs agents will switch to tracking and set $1/\pi_t = 1/\pi$; perceived inflation will then be more heavily influenced by actual inflation and it is more likely to end up in U than under pure OLS. The presence of hyperinflation prompts agents to pay more attention to recent observations by switching to tracking, this in turn makes hyperinflation more likely to occur and predictions with tracking better, thus reinforcing the switch to tracking in periods of instability. Only if $1/\pi_0$ is very small relative to the variance of inflation and if initial inflation starts out in S (and σ is large enough), hyperinflation is impossible.

This intuition suggests that the model is consistent with stylized fact 1, since a number of hyperinflation may occur in the economy before it settles down. Also, it is clear that an ERR will end each hyperinflation temporarily, so that fact 2 is found in this model. Also, once π_t is in the set U, inflation is likely to grow even if seignorage does not, which is consistent with fact 3.

To analyze fact 4, consider the second graph of Figure 3, which corresponds to a high average level of seignorage. Now, the unstable set U is much larger. Furthermore, U is "dangerously" close to the rational expectations equilibrium π_{RE}^{-1} where the economy tends to live, and it is likely for the model to end up in U and a hyperinflation to occur even if inflation has been stable for a while. Thus, a country with a high average seignorage tends to have hyperinflationary episodes more often, and fact 4 is consistent with the model.

²⁶That a high σ_d^2 increases the probability of a hyperinflation is roughly consistent with the data, but we will not study this property of the model any further in this paper.

5.3 Asymptotic Rationality

We show that asymptotic rationality obtains by first proving convergence to RE.

Proposition 1 In addition to Assumptions 1-2 of Appendix 2, assume that average seignorage and its variance are low enough for two stationary REE to exist, that $\bar{\pi} \geq \bar{\pi}^2$ (targeted inflation belongs to the stability set) and that $\bar{\pi}$ and $\bar{\sigma}$ are large enough²⁷. Then

$$\pi_t \rightarrow \pi_{RE}^{-1} \quad \text{a.s.,}$$

and Asymptotic Rationality obtains.

Proof

We will show the theorem holds for

$$\bar{\pi} > \frac{1 - \bar{\sigma}}{K_i = \bar{A}} \quad \text{and} \quad \bar{\sigma} > \frac{\bar{\pi} - U}{\min\{1; \frac{K_i = \bar{A}}{1 - K_i = \bar{A}}\}} \quad (17)$$

In order to show that the learning mechanism stays in the OLS form in all periods $t > \bar{\pi}$, we first show that inflation is bounded below. For each t and each realization, only three cases are possible: Case i): an ERR is activated at t ; ii): an ERR is not activated at t and $1 - \bar{\sigma} - \pi_{t-1} - d_{t-1} = \bar{A} < 0$; and iii): an ERR is not activated at t and $1 - \bar{\sigma} - \pi_{t-1} - d_{t-1} = \bar{A} > 0$:

Notice that in cases ii) and iii) the first branch of (16) applies and we have $\pi_{t-1} < \bar{\sigma}^{-1}$. Note also that since π_t is a weighted average of past inflations and $-U$ is an upper bound of inflation $\pi_t < -U$.

In case i), inflation is equal to $\bar{\pi}$: In case ii)

$$\frac{P_t}{P_{t-1}} = \frac{1 - \bar{\sigma} - \pi_{t-1}}{1 - \bar{\sigma} - \pi_{t-1} - d_{t-1}} = \frac{1}{1 - \frac{\bar{\sigma}(\bar{\pi} - \pi_{t-1}) + d_{t-1}}{1 - \bar{\sigma} - \pi_{t-1}}} \quad (18)$$

²⁷The assumption on $\bar{\pi}$ can be interpreted as saying that convergence occurs if the importance given to recent news is never too high. This assumption is needed in order to obtain a lower bound of inflation in the first part of the proof. A lower bound on inflation can also be obtained for unrestricted $\bar{\pi}$ by changing the model in reasonable ways. For example, assuming that the government has the objective of avoiding devaluation and it achieves this by activating an ERR and insure that $\frac{P_t}{P_{t-1}} > 1$ at the same time that reserves increase. A lower bound in $\bar{\sigma}$ can be interpreted as saying that agents do not easily switch to tracking; a lower bound is necessary because, if $\bar{\sigma}$ is too small, even if π_t is very close to π_{RE}^{-1} , it will eventually happen that $\frac{P_{t-1}}{P_{t-2} - \pi_{t-1}} > \bar{\sigma}$; then $\bar{\pi}_t = \bar{\pi}$, perceived inflation will have positive variance, and convergence will never occur.

where the first equality follows from (16) and (12), and the second equality follows from the fact that an ERR was established at t_{i-1} so that $\frac{P_{t_{i-1}}}{P_{t_{i-2}}} = \bar{r}$. Now, using $\bar{r}_{t_{i-1}} < \bar{r}_{t_{i-2}}$; $\bar{r}_t \geq \bar{r}$, $d_t \geq K^i$ and (17) we have

$$\frac{\bar{r}_t}{\bar{r}} (\bar{r}_{t_{i-1}} - \bar{r}_{t_{i-2}}) + d_t = \bar{r} > \frac{\bar{r}}{\bar{r}} (\bar{r}_{t_{i-1}} - \bar{r}_{t_{i-2}}) + K^i = \bar{r} > 0$$

which, together with (18), implies that $\frac{P_t}{P_{t_{i-1}}} > 1$ in case ii).

In case iii), the condition on $\bar{r}_{t_{i-1}}$ and simple algebra imply

$$\frac{P_t}{P_{t_{i-1}}} > \frac{d_{t_{i-1}} = \bar{r}}{1 - \bar{r}_{t_{i-1}}} > \frac{K^i = \bar{r}}{1 - K^i = \bar{r}}$$

Therefore we find the lower bound $\frac{P_t}{P_{t_{i-1}}} \geq \min(1, \frac{K^i = \bar{r}}{1 - K^i = \bar{r}}) > 0$. Since \bar{r}_t is an average of past innovations we also have $\bar{r}_t \geq \bar{r}_L$.

For any $\epsilon > \bar{r}_L - 1$ we clearly have $\frac{P_{t_{i-1}} - \bar{r}_{t_{i-1}}}{P_{t_{i-1}}} < \epsilon$ with probability one for all $t > \bar{t}$, then $\bar{r}_t = \bar{r}_{t_{i-1}} + 1$ for all $t > \bar{t}$ and the learning mechanism stays in the OLS form in all periods.

Now, letting $C = \{ \bar{r}_t \in S \text{ i.o.g.} \}$ we want to argue that $P(C) = 1$. Consider a realization $\bar{r} \notin C$; since $\bar{r}_t \in S$ implies $\bar{r}_{t+1} \in S$, it is clear that for t large enough $\bar{r}_t \notin S$. Also, if $\bar{r}_t \in U$ for all t large enough in \bar{r}_t tends to grow and eventually goes into ERR; which would also contradict $\bar{r} \notin C$. Therefore, $\bar{r}_t \in S$ i.o. with probability one. Appendix 1 shows that, in this case, \bar{r}_t converges to \bar{r}_{RE}^{-1} almost surely.

The rest of the proof simply shows that, if the learning scheme converges to \bar{r}_{RE}^{-1} ; then the sample mean square errors converge to the best forecasts and AR obtains. For this purpose, notice first that $\bar{r}_t \rightarrow \bar{r}_{RE}^{-1}$ a.s. and the fact that H is a continuous function for \bar{r} 's in the set S imply

$$\frac{P_t}{P_{t_{i-1}}} - \frac{P_t^{RE}}{P_{t_{i-1}}^{RE}} = H(\bar{r}_{t_{i-1}}; \bar{r}_{t_{i-2}}; d_t) - H(\bar{r}_{RE}^{-1}; \bar{r}_{RE}^{-1}; d_t) \rightarrow 0 \quad \text{a.s.}$$

as $t \rightarrow \infty$, so that

$$E_{t_{i-1}} \left[\frac{P_t}{P_{t_{i-1}}} \right] - E_{t_{i-1}} \left[\frac{P_t^{RE}}{P_{t_{i-1}}^{RE}} \right] \rightarrow 0 \quad \text{a.s.}$$

by Lebesgue dominated convergence, boundedness of innovation and that by definition $\bar{r}_{RE}^{-1} = E \left[\frac{P_t^{RE}}{P_{t_{i-1}}^{RE}} \right]$. Therefore, both \bar{r}_t and $E_{t_{i-1}} \left[\frac{P_t}{P_{t_{i-1}}} \right]$ converge to \bar{r}_{RE}^{-1} ; we have

$$\frac{1}{T} \sum_{t=1}^T \left(\frac{P_t}{P_{t_{i-1}}} - \bar{r}_{RE}^{-1} \right)^2 \rightarrow 0 \quad \text{a.s.} \quad (19)$$

as $T \rightarrow \infty$, and

$$\frac{1}{T} \sum_{t=1}^T \frac{P_t}{P_{t-1}} \left(\frac{P_t}{P_{t-1}} - \mu \right) \leq \frac{1}{T} \sum_{t=1}^T \frac{P_t}{P_{t-1}} E_{t-1} \left(\frac{P_t}{P_{t-1}} - \mu \right) + o_p(1) \quad \text{a.s.}$$

as $T \rightarrow \infty$ for any $\epsilon > 0$:

Notice that AR imposes very few restrictions on the learning scheme. In particular, AR poses no restriction on the choice of the parameter θ despite the fact that this is a key parameter determining the probability of experiencing a hyperinflation. Also, even if AR is satisfied agents could be making systematic mistakes; for example, if agents used pure OLS, they would be making very large forecasting errors whenever a hyperinflation happened, since OLS does not weigh more heavily recent events.

5.4 Internal Consistency

In subsection 5.2 we explained intuitively why hyperinflations are more likely to occur with high $1-\theta$. Also, a high value of $1-\theta$ is likely to generate better forecasts during a hyperinflation. Therefore, there is potential for IC to be satisfied precisely for $1-\theta$'s that generate hyperinflations.

IC is the criterion we use to define equilibria in the paper. The variables we have to determine are the sequences of inflation, expected inflation and nominal balances, together with the parameter θ : Notice that, since θ is determined as part of the equilibrium, the θ that satisfies IC will vary as the process for d_t changes so that the learning mechanism is endogenous to government policy.

Definition 4 A stochastic process $\left\{ \frac{P_t}{P_{t-1}}, \pi_t, M_t \right\}$ together with θ is an IC equilibrium for $(\theta; T)$ if:

1. Given $\theta; \left\{ \frac{P_t}{P_{t-1}}, \pi_t, M_t \right\}$ satisfy (15), (12), (14) with $M_t \geq 0$ for all t ;
2. Given $\left\{ \frac{P_t}{P_{t-1}}, \pi_t, M_t \right\}; \theta$ satisfies

$$E \left[\frac{1}{T} \sum_{t=1}^T \frac{P_{t+1}}{P_t} \pi_t \right] \leq \min_{\theta^0} E \left[\frac{1}{T} \sum_{t=1}^T \frac{P_{t+1}}{P_t} \pi_t(\theta; \theta^0) \right]$$

Where $\pi_t(\theta; \theta^0)$ has been defined in the previous section²⁸.

²⁸The careful reader will note that we did not impose IC on the learning parameter θ in this definition or in the simulations we describe below. This was done only for simplicity.

Since the dynamics are highly non-linear, characterizing analytically the equilibrium π^0 s is impossible. We solve the model by simulation and search numerically for θ that satisfy IC in a way to be described below. This will show that IC does impose restrictions on the space of learning parameters, and that the resulting equilibria match the stylized facts of the hyperinflationary experiences remarkably well.

5.5 Characterization of the solution by simulation.

To generate simulations we must assign values to the parameters of the money demand equation (α, \bar{A}) and the distribution of d_t . We choose values $(\alpha = 0.4$ and $\bar{A} = 0.37)$ in order to replicate some patterns of the Argentinean experience during the 80's, for details see the appendix 2. We assume that seignorage is normally distributed, truncated to have positive values of seignorage, with mean that varies across experiments we perform and $\frac{1}{4}d = 0.01$ ²⁹

The parameter α was set equal to 10%. We also assumed that the government established ERR whenever expectations were such that inflation rates would be above 5000%, so that we set $\bar{u} = 50$. The ERR is enforced until expected inflation is inside the stable set S.

Since our purpose is to show that a small deviation from rational expectations can generate dynamics different from RE and more similar to the data, we choose as initial beliefs $\pi_0 = \pi_{RE}^1$ so that our simulations are biased in favor of looking like RE.³⁰ For the specified parameters, the maximum level of average seignorage in the deterministic model for which a REE exists is $E(d_t) = 0.05$. In spirit of making it difficult for the model to depart from RE, we have chosen values of the average seignorage for which a REE exists. In order to quantify the relevance of average seignorage (fact 4), we performed our calculations for four different values: $E(d_t) = 0.049; 0.047; 0.045$ and 0.043 :

First of all, we describe the typical behavior of the model. A particular realization is presented in Figure 4. That realization was obtained with $E(d_t) = 0.049$ and $\theta = 0.2$: We will show below that this value of the learning parameter satisfies IC. This graph shows the potential of the model to generate enormous inflation rates. In the same graph, we also plotted two horizontal lines, one at each of the stationary deterministic rational expectation equilibria, to show how the model under learning can generate much higher inflation rates than the rational expectations version.

²⁹The results for lower values of $\frac{1}{4}d$ were similar. Of course, hyperinflationary experiences were then less frequent.

³⁰For example, it would be trivial to generate at least one hyperinflation by choosing $\pi_0 > \pi_{RE}^2$.

This graph displays some of the stylized facts in the learning model³¹. In the first periods, the inflation rate is close to the low stationary equilibrium. When a relatively large shock occurs, it drives perceived inflation into the unstable region U and a hyperinflation episode starts. Eventually, ERR is established and the economy is brought back into the stable region. If no large shocks occurred for a long while, π_t would be revised according to the OLS rule $\pi_t = \pi_{t-1} + 1$; and the model would converge to the rational expectations equilibrium; however, since average seignorage is high for this simulation, it is likely that a new large shock will put the economy back into the unstable region and a new burst in inflation will occur. Clearly, we have recurrent hyperinflations, stopped by ERR (facts 1 and 2). Since seignorage is i.i.d., and since the graph shows some periods of sustained increases in inflation, it is clear that there is little correlation of inflation and seignorage (fact 3). In order to reduce (or eventually eliminate) the chances of having a new burst, the government must reduce the amount of seignorage collected (i.e., an "orthodox" stabilization plan) in order to increase the size of the stable set. This would separate the two horizontal lines and it would stabilize the economy permanently around the low stationary equilibrium. Establishing ERR just before a reduction in average seignorage would help stabilize the expectations of agents more quickly, so there is room for a positive effect of a 'heterodox' intervention as well.

To find the learning equilibrium we look for values of α that satisfy the lower bound criterion IC for $(\beta; T) = (0.01; 120)$. This value of T is chosen to represent 10 years, roughly the length of the hyperinflationary episodes we are studying. The value of β is just chosen to be 'small', it will be clear below how the results may change if this parameter changes.

To find numerically those values of α that satisfy IC we proceed as follows: we define a grid of $\alpha \in [0; 1.2]$ separated by intervals of length 0.1. The same grid is used both for α and the alternative learning parameters α^0 considered. We compute the mean squared errors in the right side of (11) by Monte-Carlo integration³², and we find the minimum over α^0 for each α on the grid. Figure 5 shows the result of these calculations: in the horizontal axis we plot α , while the vertical axes plots α^0 . The interval of alternative learning parameters that generate a mean square error within $\epsilon = .01$ of the minimum in each column is marked with a dark area. An IC equilibrium for

³¹The behavior of the REE in this economy is clear: for the stationary REE, inflation would be i.i.d., fluctuating around the horizontal line of π_{RE}^1 ; For bubble equilibria, inflation would grow towards the horizontal line of π_{RE}^2 :

³²More specifically, we draw 1000 realizations of $d_1; \dots; d_{120}$, find the equilibrium inflation rates for each realization, we compute the sample mean square error for each alternative α^0 in the grid, and we average over all realizations.

$(\beta; T) = (0.01; 120)$ is found when the dark area cuts the 45 degree line.

Table 1 reports the probabilities of having n hyperinations in 10 years for different values of average seignorage and for those values of β that satisfy the IC criterion.

As Figure 5 shows, for a low value $E(d_t) = 0.043$, only $\beta = 0$ and 0.1 satisfy the IC requirement. It turns out that for those two values the probability of a hyperinflation in 120 periods is zero. Therefore, if IC is imposed, this value of the average seignorage rules out hyperinflations. Since hyperinflations do not occur, giving too much importance to recent observations does not generate good forecasts, so a low β is a good choice within the model. If seignorage is increased to 0.045 ; the criterion is satisfied for all values of α between 0.5 and zero. As indicated by Table 1, for this average seignorage there are equilibria in which the probability of experiencing recurrent hyperinflations is high, so that higher alternative β 's generate good forecasts, and the hyperinflationary behavior is reinforced. Table 1 and Figure 5 show that as the mean of seignorage increases, quasi-rational learning is consistent with hyperinflations. Furthermore, hyperinflations are more likely when seignorage is high. This documents how fact 4 is present in our model.

This exercise formalizes the sense in which the equilibria with a given learning mechanism reinforces the use of the mechanism. For instance, when seignorage is 0.49 and $\beta = 0.2$; an agent using an alternative α equal to zero, which is the collective behavior that replicates the REE, will make larger MSE than the agent using $\beta = 0.2$: The reason is that in equilibrium there are many hyperinflations, and the agent that expects the REE will make bad forecasts.

Whenever equilibria with hyperinflations exist, there is multiplicity of equilibria (several β 's satisfy IC). The behavior of inflation does not change much for different equilibrium β 's.³³

The numerical solutions show that the chances of a hyperinflation during the transition to the rational expectations equilibrium depend on the size of the deficit. The lower the deficit, the lower the chances of experiencing a hyperinflation. Notice how the equilibrium learning parameters depend on the size of average seignorage: higher seignorage correspond to higher equilibrium β 's, which are more likely to generate a hyperinflation.³⁴

³³Since we chose $\beta_0 = \beta_{RE}^{-1}$ when we set $\beta = 0$ we have the REE: When initial beliefs are far apart from the REE, then $\beta = 0$ will no longer satisfy IC.

³⁴We have simulated the model under many other values for the parameters. The main results of this subsection about the behavior of inflation are observed for a wide range of the parameters.

5.6 Epsilon-Delta Rationality (EDR).

In this section we show that in the equilibria with hyperinflation discussed above, the criterion EDR is satisfied if the highest admissible inflation $\bar{\pi}^U$ is large enough, for values of β that are closely related to the probability of experiencing a hyperinflation. This is because, when a hyperinflation occurs, the conditional expectation can be arbitrarily high due to the existence of an asymptote in the mapping from perceived to actual inflation (see Figure 3) but, in fact, the actual value of inflation is unlikely to be ever so high in a given realization. Thus, for every realization when a hyperinflation occurs, the learning forecast can do better than the conditional expectation with very high probability in finite samples.

Proposition 2 Consider the model of section 3. If the regularity assumptions 1-2 of appendix 2 are satisfied, and d_t has a density, then for given parameter values of the model and any $(\beta; T)$; there is a $\bar{\pi}^U$ large enough such that

$$\forall \beta; T, \exists \bar{\pi}^U \text{ s.t. } P(\text{ERR at some } t \in T) < \epsilon$$

where $P(\text{ERR at some } t \in T)$ is the probability that the government implements ERR at some period $t = 1; \dots; T$:

Proof

Fix $\beta; T$: We first consider the case that $\bar{\pi}^U = 1$: Consider a realization \bar{t} where an ERR is established at some $t = 0; \dots; T$: Letting $\bar{t} + 1$ be the first period where this occurs, it has to be the case that $1 - \beta \pi_{\bar{t}+1}(\bar{t}) - d_{\bar{t}+1}(\bar{t}) - \bar{A} < 0$ and $1 - \beta \pi_{\bar{t}}(\bar{t}) > 0$. Clearly, we can only have the first inequality if $\bar{d} > \bar{A}(1 - \beta \pi_{\bar{t}+1}(\bar{t})) < K^+$. Therefore, since inflation is given by equation (15) and $(\pi_{\bar{t}+1}, \pi_{\bar{t}})$ are known with information available at \bar{t} , we have that

$$\begin{aligned} E_{\bar{t}} \left[\frac{P_{\bar{t}+1}}{P_{\bar{t}}} \right] &= \int_{K^-}^{K^+} H(\pi_{\bar{t}+1}(\bar{t}); \pi_{\bar{t}}(\bar{t}); \theta) dF_{d_{\bar{t}+1}}(\theta) = \\ &= \int_{K^-}^{\bar{d}} \frac{1 - \beta \pi_{\bar{t}}(\bar{t})}{1 - \beta \pi_{\bar{t}+1}(\bar{t}) - \theta - \bar{A}} dF_{d_{\bar{t}+1}}(\theta) + P[d_{\bar{t}+1} > \bar{d}] = \end{aligned} \quad (20)$$

The integral in (20) corresponds to the values of $d_{\bar{t}+1}$ for which there is a positive price level that clears the market without ERR and the first branch of (16) holds, while the second term accounts for those values of next period shock for which an exchange rate rule needs to be enforced.

Now we show that the integral in (20) is unbounded. Using arguments similar to the ones used in the appendix to show that S has an asymptote

we have

$$\int_0^{\infty} \frac{1 - e^{-\lambda t}}{\lambda} dF_{d_t}(\lambda) = (1 - e^{-\lambda t})Q(t) \int_0^{\infty} \frac{1}{x} dx = 1$$

for some finite constant $Q(t)$ and small λ :

This proves that $E_t(P_{t+1}=P_t) = 1$; therefore

$$\frac{1}{T} \sum_{t=1}^T \frac{P_{t+1}(t)}{P_t(t)} \leq \frac{P_{t+1}^e(t)}{P_t} < \frac{1}{T} \sum_{t=1}^T \frac{P_{t+1}(t)}{P_t(t)} + E_t \frac{P_{t+1}}{P_t} + \epsilon; \quad (21)$$

because the right hand side is, in fact, finite. So, (21) holds for all realizations where there is one hyperinflation and $\frac{1}{4}^{2,T} \leq P[\text{ERR at some } t \in T]$

The case of λ finite but arbitrarily large follows from observing that, with arbitrarily high probability, the sequences of the case $\lambda = 1$ are below a certain bound \bar{c} ; also, for arbitrarily high λ the conditional expectation is arbitrarily close to the one with λ finite, so that all the inequalities are maintained with arbitrarily high probability.²

Since hyperinflations occur with high probability if average seignorage is high, this proposition shows that EDR is satisfied with high \pm when seignorage is high. For example, Table 1 shows that the probability of having at least one hyperinflation is .84, .91 and .97 for average seignorage .045, .047 and .049 respectively, so EDR is satisfied for $\pm > .84, .91$ and .97.

6 Conclusion

There is some agreement by now that the hyperinflations of the 80's were caused by the high levels of seignorage in those countries, and that the cure for those hyperinflations was fiscal discipline and abstinence from seignorage. The IMF is currently imposing tight fiscal controls on the previously hyperinflationary countries that are consistent with this view. Nevertheless, to our knowledge, no currently available model justified this view and was consistent with some basic facts of hyperinflations. In particular, the fact that seignorage has gone down while inflation continued to grow in some hyperinflations makes it difficult for the IMF to argue in favor of fiscal discipline. Furthermore, some Eastern European economies are now engaging in hyperinflationary episodes similar to those of the 80's, and it seems important to have a solid model that can help judging the reasonability of the IMF recommendations.

Our model is consistent with the main stylized facts of recurrent hyperinflations and with the policy recommendations mentioned above: an exchange rate rule (ERR) may temporarily stop a hyperinflation, but average

seignorage (and also its standard deviation) must be lowered to eliminate hyperinflation permanently.

The economic fundamentals of the model are completely standard except for the use of a boundedly rational learning rule instead of rational expectations. We show that the learning rule is quasi-rational in a sense that it must perform fairly well within the model at hand. Despite abandoning rational expectations, we maintain falsifiability of the model, and the learning rule driving expectation formation is endogenous to government policy. This deviation from rational expectations is attractive because it avoids the strong requirements on rationality placed by rational expectations, and because the fit of the model improves dramatically even if the deviation is small.

On the practical side, this paper shows that hyperinflation can be stopped with a combination of heterodox and orthodox policies. The methodological contribution of the paper is to show that, as long as we carry along adequate equipment for orientation and survival, an expedition into the "wilderness of irrationality" can be quite a safe and enjoyable experience.

Appendix 1

Households:

In this appendix we solve a deterministic small open economy version of a standard overlapping generations model with equilibrium conditions as in section 3³⁵. It is easy to generalize the model in various directions and maintain the equivalence with the main text.

Each cohort has a continuum of agents living two periods. There is one type of consumption good in the world. Preferences are given by $(\ln c_t^y + \beta \ln c_{t+1}^o)$ where c_t^y is consumption of young agents at time t and c_{t+1}^o is consumption of old agents at time $t + 1$: Agents are endowed with $(\omega^y; \omega^o)$ units of consumption when young and old respectively, where $\omega^y > \omega^o > 0$:

Asset Markets:

There are two assets in the economy: domestic and foreign currency. In our hyperinflationary equilibria, domestic currency will be return-dominated by foreign currency. To ensure that money demand is positive we will impose a cash-in-advance constraint for local currency on net purchases of consumption³⁶.

$$M_t \leq P_{t+1}(c_{t+1}^o; \omega^o)$$

³⁵This appendix extends the closed economy results of Sargent and Wallace (1987).

³⁶Freeman and Kydland (2000) describe a model where a fraction of goods are always bought using currency, even though it is return dominated, and foreign currency is also used. Modelling all those details is well beyond the scope of this appendix and the paper.

for $t \geq 1$: This condition makes foreign currency value-less for households. Therefore, we can write the constraints for the household as

$$\begin{aligned} P_t y_t &= P_t c_t^y + M_t \\ M_t &= P_{t+1} (c_{t+1}^o - \beta^o) \\ M_t &\geq 0 \end{aligned}$$

Households' optimization implies

$$\begin{aligned} c_t^y &= \frac{P_t y_t + P_{t+1} \beta^o}{(1+\beta^o)P_t} & , & \quad \frac{M_t}{P_t} = \frac{\beta^o y_t}{(1+\beta^o)} + \frac{P_{t+1} - \beta^o}{P_t (1+\beta^o)} & \quad \text{if } \frac{\beta^o y_t}{(1+\beta^o)} \geq \frac{P_{t+1}}{P_t} \\ &= \beta^o y_t & ; & \quad = 0 & \quad \text{otherwise} \end{aligned} \quad (22)$$

This gives a microfoundation to equation (1).

Foreign sector:

The world is inhabited by large wholesale firms that can buy and sell goods in any country without transaction costs and are not subject to cash in advance constraints. If we let X_t^j (which can be negative) be the net number of units of the consumption good bought domestically and sold abroad by firm j , profits are given by

$$\pi_t^j = X_t^j P_t^f e_t - X_t^j P_t$$

where e_t is the nominal exchange rate and P_t^f the price of the consumption good abroad. Free entry into the business implies that profits must be zero, therefore

$$P_t^f e_t = P_t \quad (23)$$

If we let TB_t be the trade balance in units of consumption, market clearing implies that

$$\beta^o y_t + \beta^o = c_t^y + c_t^o + d_t + TB_t$$

where $d_t \geq 0$ is exogenously given government consumption at time t .

The government budget constraint:

We assume the government does not tax agents³⁷, it only generates income by seignorage and, occasionally, by changing its stock of foreign currency R_t .

³⁷Taxes and government debt are easy to introduce by reinterpreting d_t and the endowments β^o : all equations are consistent with β^o denoting endowments net of age-dependent, constant, lump-sum taxes, and with d_t being the primary deficit of the government. Debt can be introduced, for example, if we assume that government debt is constant (perhaps because the government is debt constrained) and d_t represents interest payments on debt plus primary deficit.

The budget constraint of the government is therefore given by

$$\frac{M_t - M_{t-1}}{P_t} = d_t + (R_t - R_{t-1}) \frac{e_t}{P_t} \quad (24)$$

Equilibrium in all markets implies $(R_t - R_{t-1})e_t = TB_t P_t$

Government policy:

Government policy must set money supply and reserves to satisfy (24). Reserves can be set according to two regimes:

A Floating Regime

in this regime the government does not change its position on foreign currency. Then, all the government expenditure is financed by means of money creation, so that $TB_t = 0$ and

$$M_t = d_t P_t + M_{t-1}$$

which together with the money demand (22) solve for the equilibrium sequences of M_t and P_t : The nominal exchange rate is given by equation (23):

A fixed ERR regime

In this regime the government buys or sells foreign currency at a given exchange rate. Given $P_t^f; P_{t-1}^f; e_{t-1}$ and a desired level of inflation $\bar{\pi}$, the desired exchange rate is

$$e_t = \bar{\pi} e_{t-1} \frac{P_{t-1}^f}{P_t^f} \quad (25)$$

Equation (23) implies that with this policy the government achieves $\bar{\pi} = \frac{P_t}{P_{t-1}}$. The money demand (22) determines the level of nominal money demand consistent with the nominal exchange rate target. Given this level of money supply and d_t , foreign reserves and, consequently, the trade balance adjust so as to satisfy (24). Of course, to the extent that there are constraints on the evolution of the government foreign asset position, ERR may not be feasible.

We assume that the first regime is used if inflation achieves an acceptable level less than $\bar{\pi}^U$; the ERR regime is followed otherwise.

The equilibrium is therefore given by equations (22), (23) and (24) which are deterministic versions of equations (1) to (4) in the paper. The analogy between this deterministic version and the stochastic one in the paper is only exact up to a linear approximation, a usual simplification in macroeconomic models under learning.

Appendix 2

In this appendix we characterize the set of stationary REE of the model with uncertainty; we discuss how the sets U and S are affected by the process of d_t ; and we show that least squares learning converges to the lower stationary rational expectations equilibrium.

Assume that expectations about inflation are given by (9) with constant perceived inflation $\bar{\pi}_t = \bar{\pi}$. Then, $\bar{\pi}_{RE}$ is a stationary REE ⁱ, when prices are generated by $\bar{\pi}_t = \bar{\pi}_{RE}$; then $E_t(P_{t+1} | I_t) = \bar{\pi}_{RE} P_t$: Let us make some assumptions on the model³⁸:

Assumption 1 the support of d_t is the interval $[K^-, K^+]$, where $K^- > 0$ and $K^+ < \bar{A}$:

Assumption 2 d_t has a continuous density f_{d_t} and $f_{d_t}(K^+) > 0$

If $\bar{\pi}_t = \bar{\pi}$ equation (15) implies $P_{t+1} = h(\bar{\pi}; d_{t+1})P_t$; where h is as defined in section 5.2. Letting

$$S(\bar{\pi}) = E(h(\bar{\pi}; d_{t+1}));$$

S is interpreted as the mapping from perceived to actual expectations on inflation.

In the next proposition we characterize the properties of S .

Proposition 3 There can not be a stationary rational expectations equilibrium with perceived inflation $\bar{\pi} > (1 - \beta K^+ / \bar{A}) = \bar{\pi}^*$.

The set of stationary rational expectations equilibria coincides with the fixed points of the mapping $S : [0; (1 - \beta K^+ / \bar{A}) = \bar{\pi}^*] \rightarrow \mathbb{R}$:

S has the following properties:

1. In the set $[0; (1 - \beta K^+ / \bar{A}) = \bar{\pi}^*]$; the mapping S is increasing and convex. If $\bar{\pi}^U = 1$, then S has an asymptote at $\bar{\pi} = (1 - \beta K^+ / \bar{A}) = \bar{\pi}^*$:
2. S has at most two fixed points denoted $\bar{\pi}_{RE}^1 < \bar{\pi}_{RE}^2$. For a distribution where d_t is low enough, and for $\bar{\pi}^U$ large enough two equilibria exist. For a distribution where d_t is large enough no equilibrium exists.
3. When two fixed points exist, $S^0(\bar{\pi}_{RE}^1) < 1$
4. The stability set S is smaller under uncertainty. More precisely, let $\bar{\pi}_{RE}^1 < \bar{\pi}_{RE}^2$ be the rational expectations equilibria without uncertainty (when $\beta_d = 0$ and $E(d_t) = \bar{d}_t$). Assume that two fixed points of S exist.

Then $\bar{\pi}_{RE}^1 < \bar{\pi}_{RE}^1 < \bar{\pi}_{RE}^2 < \bar{\pi}_{RE}^2$.

³⁸All the theorems would also work under the assumption that d_t had infinitely many possible realizations, all of them positive.

Proof

Under the assumptions of the model, once an ERR is established, it is maintained until expectations are back in the set S . If a stationary REE existed then $E_t \frac{P_{t+1}}{P_t} = \bar{r}_{RE}$ for all t ; if an ERR was possible, it could occur that an ERR is established at period \bar{t} ; it is maintained at $\bar{t} + 1$, and $E_{\bar{t}} \frac{P_{\bar{t}+1}}{P_{\bar{t}}} = \bar{r}$: Since, in general, $\bar{r}_{RE} \neq \bar{r}$ this is impossible. So, for a stationary REE it must be the case that ERR never occur. Now, if $\bar{r}_{RE} > (1 - \frac{K^+}{K^+ + \bar{A}}) = \bar{r}$ the event $1 - \frac{K^+}{K^+ + \bar{A}} > \bar{r}_{RE}$ would have positive probability, which contradicts an ERR never taking place. Stationary REE are, by definition, fixed points of S .

1. Using the definition of S we have

$$S^0(\bar{r}) = E \frac{h(\bar{r}; d_t)}{\bar{r}} = E \frac{1 - \frac{K^+}{K^+ + \bar{A}}}{(1 - \frac{K^+}{K^+ + \bar{A}} - \bar{r}_{RE})^2} \quad \text{and}$$

$$S^0(\bar{r}) = E \frac{2 \frac{K^+}{K^+ + \bar{A}}}{(1 - \frac{K^+}{K^+ + \bar{A}} - \bar{r}_{RE})^3};$$

since the expressions inside the expectation are non-negative, this proves that $S^0, S^0 > 0$.

To prove the existence of an asymptote; note that

$$S((1 - \frac{K^+}{K^+ + \bar{A}}) = \bar{r}) = \int_{K^+}^{\infty} \frac{K^+}{K^+ + \bar{A}} f_{d_t}(d) d > \int_{K^+}^{\infty} \frac{K^+}{K^+ + \bar{A}} f_{d_t}(d) d; \quad (26)$$

for small $\bar{r} > 0$. According to assumption 2, we can choose \bar{r} small enough such that, if $d > K^+ + \bar{r}$, then $f_{d_t}(d) > 0$; this implies the inequality in

$$S((1 - \frac{K^+}{K^+ + \bar{A}}) = \bar{r}) > \int_{K^+}^{\infty} \frac{K^+}{K^+ + \bar{A}} d = \int_0^{\infty} \frac{K^+}{x} dx = 1;$$

the first equality follows from a trivial change of variables, and the last equality because the integral of a hyperbola at zero is infinite. This shows that $S((1 - \frac{K^+}{K^+ + \bar{A}}) = \bar{r}) = 1$:

2. That we have at most two fixed points follows immediately from convexity of S . Increasing the probability mass of d_t near zero means that values of d_t close to zero determine the expectation that gives $S(\bar{r})$. Since, for any given \bar{r} ; $h(\bar{r}; d) \rightarrow 1$ as $d \rightarrow 0$; and since inflation is bounded we have that $S(\bar{r})$ becomes arbitrarily close to 1

and $S^0(\bar{r})$ arbitrarily close to zero. This means that as f_{d_t} puts more probability mass near zero S has to cross the 45° line from above. The fact that S has an asymptote (part 1 of this proposition) implies that there is a second fixed point if \bar{r}^U is large enough. The fact that S is concave implies that there are at most two fixed points. Since S is increasing in d_t , f_{d_t} puts more probability mass at large values no equilibrium exists.

3. Clearly, $S(0) > 0$: Therefore, S cuts the 45° line from above, and $S^0(\bar{r}_{RE}^1) < 1$.
4. Notice that \bar{r}_{RE}^1 and \bar{r}_{RE}^2 are the fixed points of $h(\bar{r}; E(d_t))$. Since h is a convex function of d_t ; Jensen's inequality implies that $S(\bar{r}) > h(\bar{r}; E(d_t))$.

Now, we argue that the least squares learning mechanism converges to the lower rational expectations equilibrium. This is a routine application of the framework of Marcat and Sargent (1989a), so the details are omitted. The associated differential equation in the set $[0, (1 + K^+ \bar{A})^{-1}]$ is given by

$$\dot{\bar{r}} = S(\bar{r}) - \bar{r} \quad (27)$$

and we know that, under least squares learning (the case that $\bar{r}_t = \bar{r}$), the system converges if and only if the differential equation is globally stable in a set D where the beliefs lie infinitely often. That stability of the differential equation is necessary and sufficient follows from the results in Ljung (1977).

Now, $S^0(\bar{r}_{RE}^1) < 1$ implies that (27) is locally stable at \bar{r}_{RE}^1 ; the basin of attraction of \bar{r}_{RE}^1 is the set $[0, \bar{r}_{RE}^2)$: In the proof to proposition 1 we have shown that \bar{r}_t visits the stable set infinitely often, therefore least squares learning converges to the rational expectations equilibrium \bar{r}_{RE}^1 a.s.

Appendix 3

In this appendix we explain the choice of parameter values for the demand for money used in the numerical solution of section 5. The money demand equation (1) is linear with respect to expected inflation. It is well known, though, that the linear functional form does not perform very well empirically. However, departing from linearity would make the analysis of the model impossible to deal with. While we do maintain linearity, we want to use parameter values that are not clearly at odds with the observations. Since we are interested in the public finance aspect of inflation, we use observations from empirical Laffer curves to calibrate the two parameters. In

particular, as one empirical implication of our model is that "high" average deficits increase the probability of a hyperinflation, we need to have a benchmark to discuss what high means. Thus, a natural restriction to impose to our numbers is that the implied maximum deficit is close to what casual observation of the data suggest. We also restrict the inflation rate that maximizes seignorage in our model to be consistent with the observations.

We use quarterly data on inflation rates and seignorage as a share of GNP for Argentina ³⁹ from 1980 to 1990 from Ahumada, Canavese, Sanguinetti y Sosa (1993) to fit an empirical Laffer curve. While there is a lot of dispersion, the maximum feasible seignorage is around 5% of GNP, and the inflation rate that maximizes seignorage is close to 60%. These figures are roughly consistent with the findings in Kiguel and Neumeyer (1992) and other studies. The parameters of the money demand ϕ and \bar{A} , are uniquely determined by the two numbers above. Note that the money demand function (1) implies a stationary Laffer curve equal to

$$\frac{\pi}{1 + \pi} m = \frac{\pi}{1 + \pi} \bar{A} (1 - i - \phi(1 + \pi)) \quad (28)$$

where m is the real quantity of money and π is the inflation rate. Thus, the inflation rate that maximizes seignorage is

$$\pi^* = \frac{r}{\phi} \frac{1}{i + 1}$$

which, setting $\pi^* = 60\%$; implies $\phi = 0.4$: Using this figure in (28), and making the maximum revenue equal to 0.05, we obtain $\bar{A} = 0.37$:

References

- [1] Ahumada, H., A. Canavese, P. Sanguinetti and W. Sosa Escudero, "Efectos distributivos del impuesto inflacionario: una estimación del caso argentino" *Economía Mexicana*, pp. 329-385, (1993)
- [2] Arifovic, J, J. Bullard and J. DuBois, "The Transition from Stagnation to Growth: An Adaptive Learning Approach", *Journal of Economic Growth*, 2:185-209, (1997).

³⁹The choice of country is arbitrary. We chose Argentina because we were more familiar with the data.

- [3] B. Bental and Z. Eckstein "On the Fit of a Neoclassical Monetary Model in High Inflation: Israel 1972-1990". *Journal of Money Credit and Banking*, (1998), 29, 725-752.
- [4] Bray, M.M. and D. Kreps, "Rational Learning and rational expectations", in *Arrow and the Ascent of Modern Economic Theory*, George Feiwel ed. 597-625, University Press, New York (1997).
- [5] Bray, M. M. and E. Savin "Rational Expectations Equilibria, Learning, and Model Specification", *Econometrica*, 54, pp. 1129-60 (1986).
- [6] Bruno, M., G. Di Tella, R. Dornbusch and S. Fisher. *Inflation Stabilization. The experience of Israel, Argentina, Brasil and Mexico* MIT Press, (1988)
- [7] Bruno, M., S. Fisher, E. Helpman and N. Liviatan. *Lessons of Economic Stabilization and its Aftermath*, MIT Press, (1991).
- [8] Chung, Hee-Taik. "Did Policy Makers Really Believe in the Phillips Curve?. An Econometric Test". Ph. D. dissertation, University of Minnesota, (1990).
- [9] Easley, D. and A. Rustichini. "Choice without belief", working paper, CORE, Université Catholique de Lovaine, (1995). Forthcoming in *Econometrica*.
- [10] Eckstein, Z. and L. Leiderman, "Seignorage and the Welfare Cost of Inflation", *Journal of Monetary Economics*, no. 29, pp. 389-410, (1992).
- [11] Evans, G.W. and S. Honkapohja. "Adaptive Forecasts, Hysteresis, and Endogenous Fluctuations", *Economic Review*, 1, Federal Reserve Bank of San Francisco, (1993).
- [12] Evans, G.W. and S. Honkapohja. "Learning Dynamics" chapter 7 of the *Handbook of Macroeconomics*, Vol.1, eds. J. Taylor and M. Woodford, North-Holland, pp.449-542, (1999).
- [13] Evans, G.W. and S. Honkapohja. *Learning and Expectations in Macroeconomics*, Princeton University Press (2001).
- [14] Evans, G.W. and G. Ramey, "Calculation, Adaptation and Rational Expectations", *Macroeconomic Dynamics*, Vol. 2, pp. 156-182, (1998).
- [15] Fudenberg, D. and D. Levine. "Universal consistency and cautious fictitious play", *Journal of Economic Dynamics and Control*, 19, pp. 1065-89 (1995).

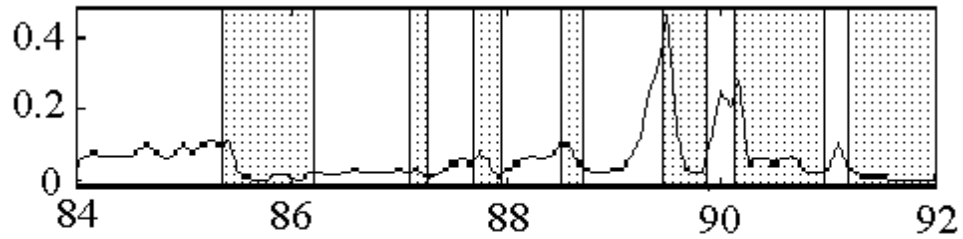
- [16] Funke, M., S. Hall and M. Solp. " Rational bubbles during Poland's hyperin°ation: implications and empirical evidence", *European Economic Review*, 38, pp 1257-76, (1994).
- [17] Hommes, C. and G. Sorger. "Consistent Equilibria", *Macroeconomic Dynamics*, Vol. 2, pp 287-321 (1998).
- [18] Imrohroglu, S. "Testing for Sunspot Equilibria in the German Hyperin°ation". *Journal of Economic Dynamics and Control*, 17, (1993).
- [19] Kiguel, M. and P. A. Neumeyer, "In°ation and Seignorage: The Case of Argentina," *Journal of Money Credit and Banking*, 27:3, 672-82, August (1995).
- [20] Kurz, M. "On rational equilibria", *Economic Theory*, 4, pp. 859-876, (1994).
- [21] Ljung, L. "Analysis of Recursive Stochastic Algorithms", *IEEE Trans. Automat. Control* AC-22, pp 551-557 (1977),
- [22] Marcet, A. and T.J. Sargent. \Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models", *Journal of Economic Theory*, no. 48, pp. 337-368, (1989a).
- [23] Marcet, A. and T.J. Sargent. \Least Squares Learning and the Dynamics of Hyperin°ation", in *International Symposia in Economic Theory and Econometrics*, vol. 4, edited by W. Barnett, J.Geweke and K. Shell, Cambridge University Press (1989b).
- [24] Marimon, R. \Learning from learning in economics", in *Advances in Economics and Econometrics: theory and applications, invited sessions at the 7th World Congress of the Econometric Society*, Cambridge University Press (1997).
- [25] Nicolini, J. "Ruling out Speculative Hyperin°ations: the Role of the Government", *Journal of Economic Dynamics and Control*, 20, pp 791-809, (1996).
- [26] Obstfeld, M and K. Rogo®. "Speculative Hyperin°ations in Maximizing Models: can we Rule them Out?", *Journal of Political Economy*, 91, pp. 675-687, (1983).
- [27] Sargent, T.J. *Bounded Rationality in Macroeconomics*, Oxford University Press, (1993).

- [28] Sargent, T.J. and N. Wallace. "Inflation and the Government Budget Constraint", in: A. Razin and E. Sadka, eds., *Economics in theory and practice*, McMillan, New York, (1987).
- [29] Timmermann, A. "How learning in financial markets generates excess volatility and predictability of excess returns" *Quarterly Journal of Economics*, 108 1135-1145, (1993).
- [30] Timmermann, A. "Excess volatility and predictability of stock market prices in autoregressive dividend models with learning" *Review of Economic Studies*, 63 523-557, (1996)
- [31] Zarazaga, C. "Hyperinflation and moral hazard in the appropriation of seignorage", Working Paper 93-26, Federal Reserve Bank of Philadelphia, (1993).

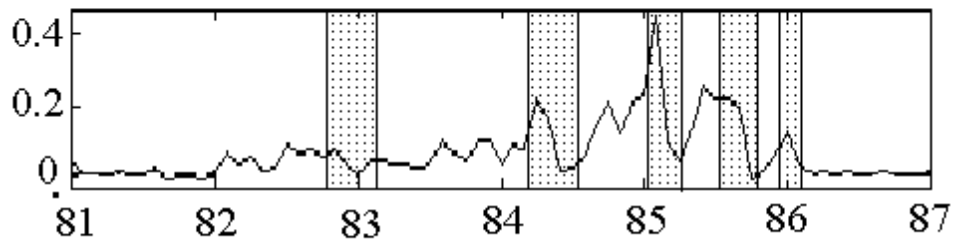
FIGURE I

Monthly Inflation Rate

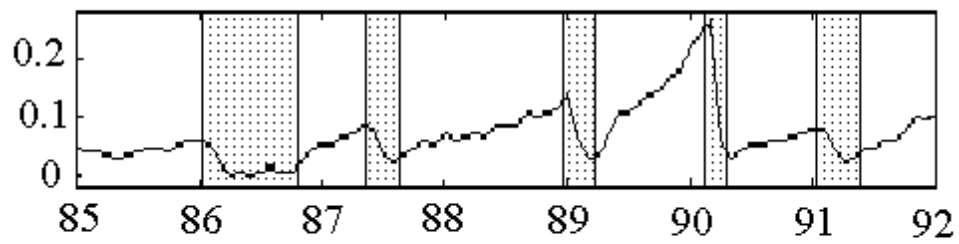
Argentina



Bolivia



Brasil



Peru

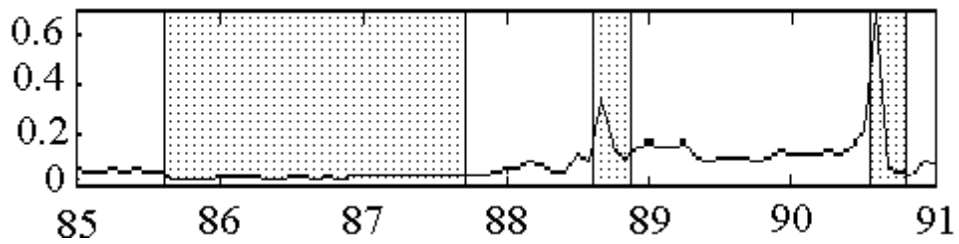


Figure 2
Time Series of Inflation Rate and Seignorage

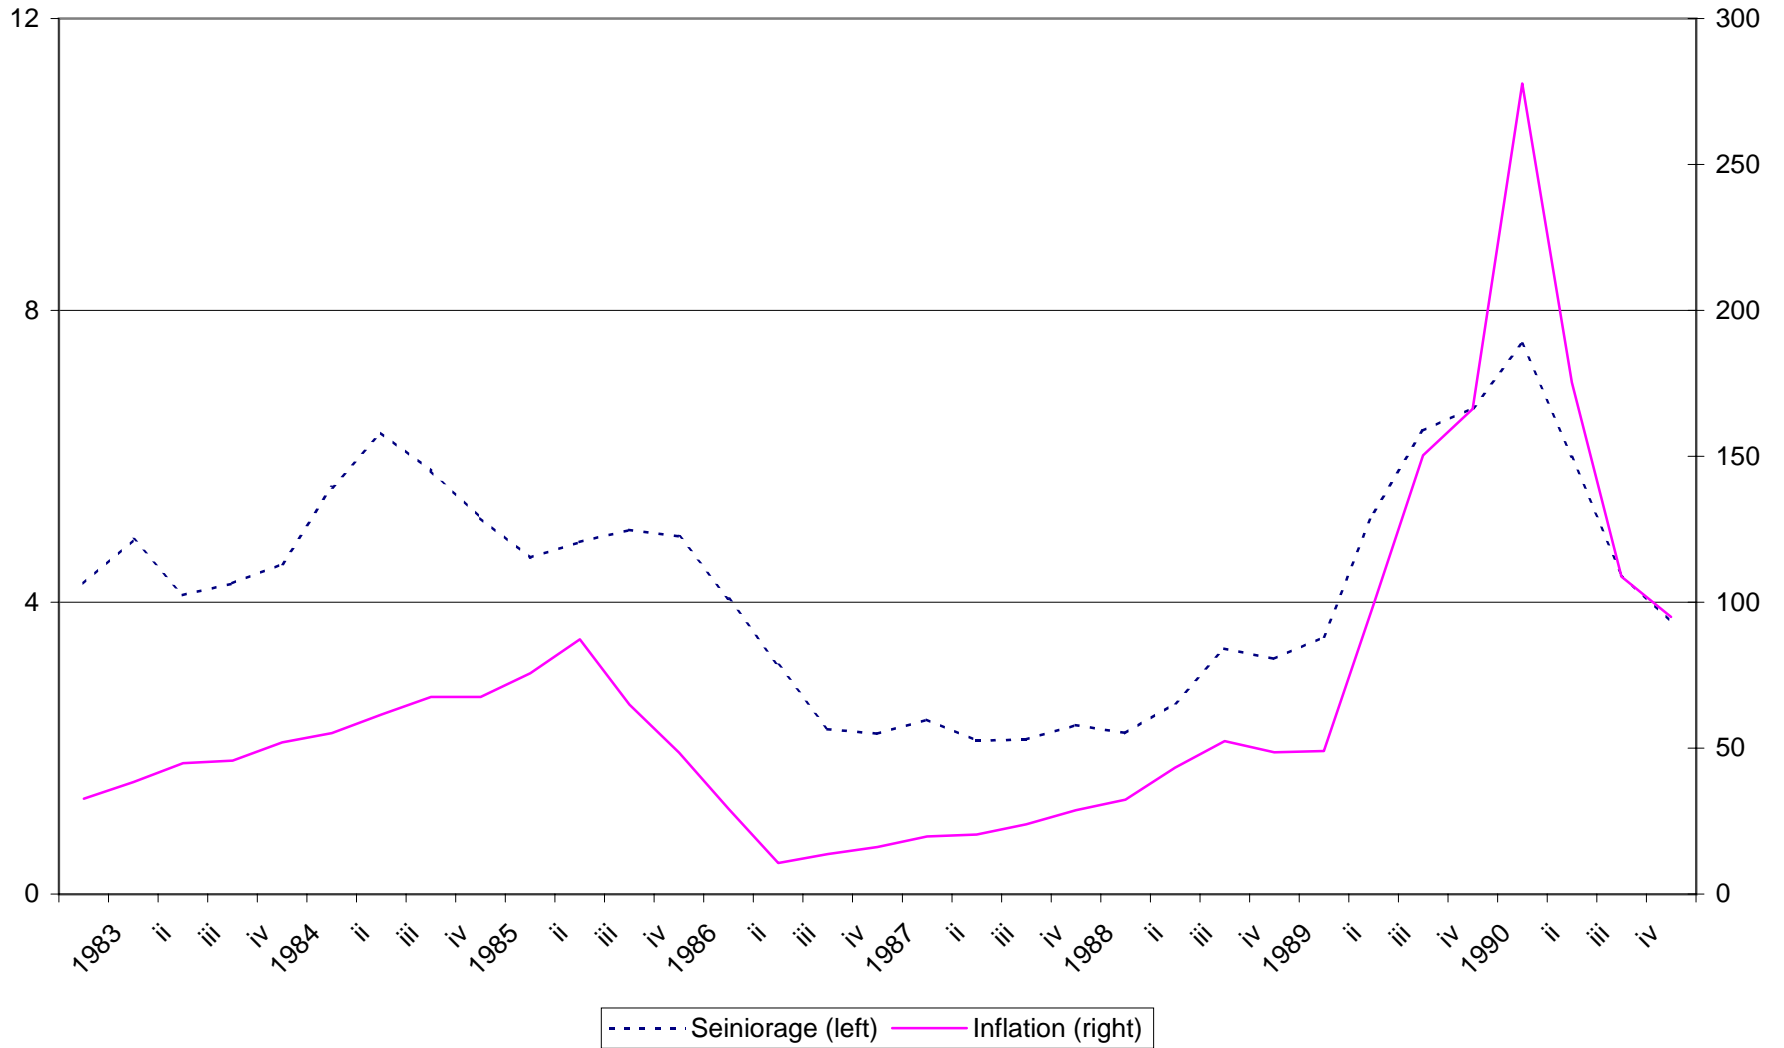


Figure 3

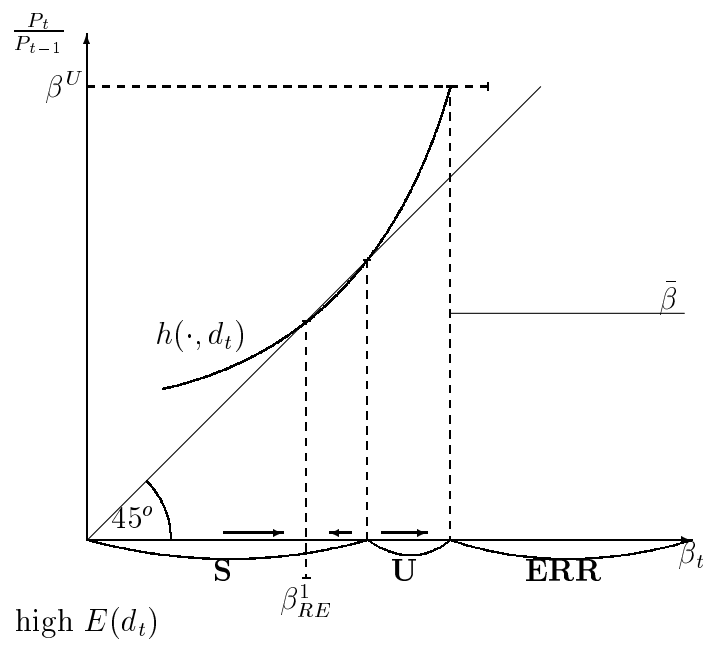
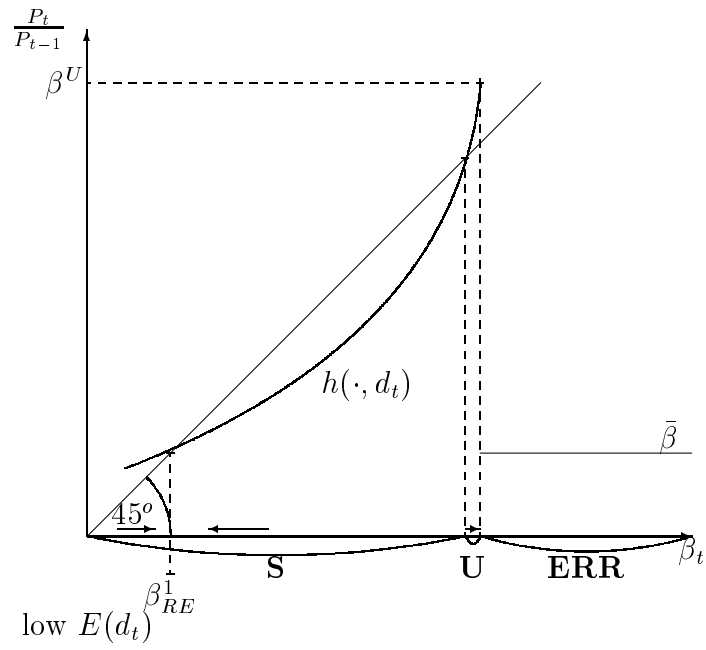


Figure 4. Monthly Inflation Rates
Inflation = $\log(P_{t+1}) - \log(P_t)$

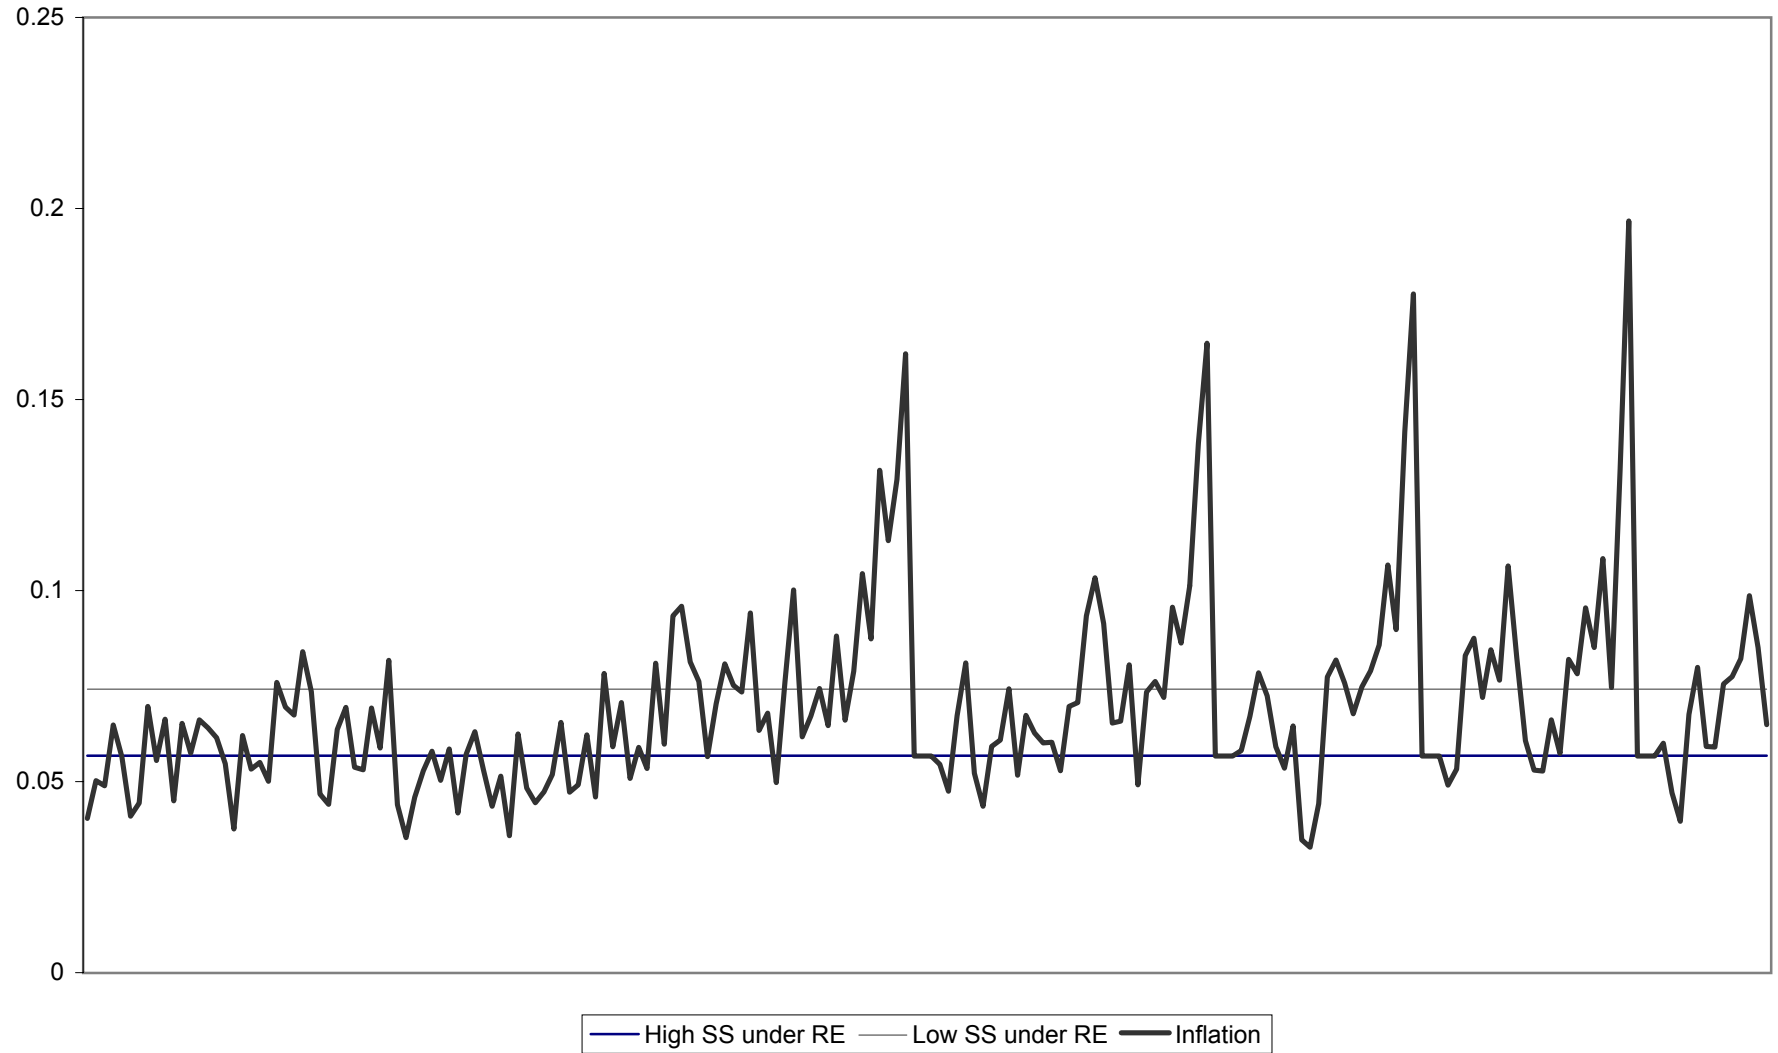


Figure 5: Efficient Values of $\bar{\alpha}$

Columns represent possible values for $1/\bar{\alpha}$ actually used by agents.

Rows depict alternative values for $1/\bar{\alpha}'$.

Light gray cells indicate the 45 degree line

Dark gray cells indicate that the value for $1/\bar{\alpha}$ is efficient.

Black cells indicate fixed points on $1/\bar{\alpha}$.

Average Seiniorage = 4.3%

0.9						Dark Gray				Light Gray
0.8						Dark Gray			Light Gray	
0.7								Light Gray		
0.6								Light Gray		
0.5								Light Gray		
0.4					Light Gray					
0.3				Light Gray						
0.2										
0.1		Black	Dark Gray	Dark Gray	Dark Gray			Dark Gray		Dark Gray
0.0	Black	Dark Gray	Dark Gray	Dark Gray	Dark Gray		Dark Gray		Dark Gray	
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9

Average Seiniorage = 4.5%

0.9					Dark Gray	Dark Gray				Light Gray
0.8					Dark Gray	Dark Gray			Light Gray	
0.7								Light Gray		
0.6								Light Gray		
0.5						Black				
0.4					Dark Gray	Dark Gray				
0.3				Black	Dark Gray	Dark Gray				
0.2		Dark Gray	Dark Gray	Dark Gray	Dark Gray	Dark Gray				
0.1		Black	Dark Gray	Dark Gray	Dark Gray		Dark Gray			
0.0	Black	Dark Gray	Dark Gray	Dark Gray	Dark Gray					
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9

Average Seiniorage = 4.7%

0.9				Dark Gray	Dark Gray	Dark Gray	Dark Gray			Light Gray
0.8				Dark Gray	Dark Gray	Dark Gray	Dark Gray			Light Gray
0.7								Light Gray		
0.6								Light Gray		
0.5								Light Gray		
0.4					Dark Gray	Dark Gray				
0.3				Dark Gray	Dark Gray	Dark Gray				
0.2		Dark Gray	Dark Gray	Dark Gray	Dark Gray	Dark Gray				
0.1	Dark Gray	Dark Gray	Dark Gray	Dark Gray	Dark Gray					
0.0	Black	Dark Gray	Dark Gray	Dark Gray	Dark Gray			Dark Gray	Dark Gray	Dark Gray
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9

Average Seiniorage = 4.9%

0.9			Dark Gray	Dark Gray	Dark Gray					Light Gray
0.8			Dark Gray	Dark Gray	Dark Gray					Light Gray
0.7								Light Gray		
0.6								Light Gray		
0.5								Light Gray		
0.4					Light Gray					
0.3				Light Gray						
0.2		Dark Gray	Dark Gray	Dark Gray	Dark Gray					
0.1	Dark Gray	Dark Gray	Dark Gray	Dark Gray	Dark Gray					Dark Gray
0.0	Black	Dark Gray	Dark Gray	Dark Gray	Dark Gray			Dark Gray	Dark Gray	Dark Gray
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9

TABLE 1
Deficit = 4.5 %

Alpha	Probability of No Hyperinflations	Probability of one Hyperinflation	Probability of two Hyperinflations	Probability of three Hyperinflations	Probability of more than three Hyperinflations
0.5	0.16	0.34	0.28	0.16	0.06
0.4	0.55	0.34	0.09	0.01	0
0.3	0.90	0.10	0	0	0
0.2	0.99	0.01	0	0	0
0.1	1	0	0	0	0
0	1	0	0	0	0

TABLE 2
Deficit = 4.7%

Alpha	Probability of No Hyperinflations	Probability of one Hyperinflation	Probability of Two Hyperinflations	Probability of three Hyperinflations	Probability of more than three Hyperinflations
0.4	0.09	0.26	0.30	0.22	0.13
0.3	0.45	0.37	0.15	0.03	0
0.2	0.82	0.14	0.04	0	0
0.1	1	0	0	0	0
0	1	0	0	0	0

TABLE 3
Deficit = 4.9%

Alpha	Probability of No Hyperinflations	Probability of one Hyperinflation	Probability of Two Hyperinflations	Probability of three Hyperinflations	Probability of more than three Hyperinflations
0.2	0.23	0.40	0.27	0.09	0.02
0.1	0.73	0.26	0.01	0	0
0	1	0	0	0	0