No. 1875

RECURRENT HYPERINFLATIONS AND LEARNING

Albert Marcet and Juan Pablo Nicolini

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

ISSN 0265-8003

RECURRENT HYPERINFLATIONS AND LEARNING

Albert Marcet and Juan Pablo Nicolini

Discussion Paper No. 1875
May 1998

Centre for Economic Policy Research
90–98 Goswell Rd
London EC1V 7DB
Tel: (44 171) 878 2900
Fax: (44 171) 878 2999
Email: cepr@cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **International Macroeconomics**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Albert Marcet and Juan Pablo Nicolini

CEPR Discussion Paper No. 1875

May 1998

ABSTRACT

Recurrent Hyperinflations and Learning*

This paper uses a model of boundedly rational learning to account for the observations of recurrent hyperinflations in the last decade. We study a standard monetary model, where the full rational expectations assumption is replaced by a formal definition of quasi-rational learning. The model under learning is able to match remarkably well some crucial stylized facts, observed during the recurrent hyperinflations experienced by several countries in the 1980s. We argue that, despite being a small departure from rational expectations, quasi-rational learning does not preclude falsifiability of the model and it does not violate reasonable rationality requirements.

JEL Classification: D83, E17, E31
Keywords: hyperinflations, convertibility, stabilization plans, quasi-rationality

Albert Marcet
Departament d'Economia i Empresa
Universitat Pompeu Fabra
Ramon Trias Fargas 23–25
08005 Barcelona
SPAIN
Tel: (34 1) 429 0551
Fax: (34 1) 429 1056
Email: marcet@cemfi.es

Juan Pablo Nicolini
Departamento de Economia
Universidad Torcuato di Tella
Miàones 2159
1428 Capital Federal
ARGENTINA
Tel: (54 1) 784 0081
Fax: (54 1) 784 0089
Email: juanpa@utdt.edu.ar

*The authors thank Tony Braun, Jim Bullard, George Evans, Seppo Honkapohja, Rodi Manuelli, Ramon Marimon, Tom Sargent, Harald Uhlig, Neil Wallace and Carlos Zarazaga for helpful conversations and Marcelo Delajara and Ignacio Ponce Ocampo for research assistance. All errors are the authors' own. Part of this work was done when both authors were visiting the Federal Reserve Bank of Minneapolis. Research support from DGICYT, CRIT and HCM is greatly appreciated.

Submitted 9 March 1998

NON-TECHNICAL SUMMARY

Hyperinflations are very spectacular economic events that have occurred at different times throughout this century. Some of the best known occurred in Central European countries in the inter-war period and in South American countries in the 1980s. Recently, several Eastern European countries have experienced (or have been on the verge of experiencing) such hyperinflations.

Certain facts are very characteristic of the South American hyperinflations of the 1980s: inflation reached very high levels (e.g. annualized 5000% in Argentina in January 1989); hyperinflations were recurrent; there was a succession of inflationary bursts (about seven in Argentina between 1984 and 1991); each of these bursts ended by establishing convertibility of the local currency with the dollar (these are the so-called 'heterodox measures' for stopping hyperinflations); the convertibility was successful in lowering inflation only temporarily.

Many economists have associated the presence of hyperinflations to high levels of seignorage; that is, to the government using money creation heavily in order to finance its fiscal deficit. A simple story about hyperinflations could therefore be 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 one of the central messages of Sargent (1986). It is also the logic behind IMF advice to hyperinflationary countries of lowering seignorage (the so-called 'orthodox measures').

Cross country evidence very strongly supports this story. Hyperinflations have occurred in countries with high seignorage and the countries that successfully stopped inflation did so by eliminating the fiscal imbalance that leads to high seignorage. This simple story fails, however, when we look closely at time series of inflation and seignorage. We can observe periods when inflation is extremely high *and* increasing while, at the same time, seignorage is decreasing (e.g. during Argentina's hyperinflation that ended in June 1985 and the last burst that preceeded the convertibility plan in 1991).

Furthermore, most of the models available in the literature (based on the assumption of rational expectations) cannot explain the observed facts and/or are inconsistent with IMF advice. What is the empirical evidence and the theoretical support for IMF actions? Can we justify the use of both orthodox and heterodox measures that, in practice, have stopped most hyperinflations?

We show that, in a model with a very standard demand and supply for money, the above facts can be explained much better if we replace the standard assumption of rational expectations by quasi-rational learning. We also show that the hyperinflations disappear if seignorage is low on average but, that if seignorage is high on average hyperinflations can occur; once inflation accelerates it keeps going up, even if there is a temporary decrease in seignorage. Also, the combination of policy measures that has worked in practice would work in our model of learning.

Our model generates recurrent hyperinflations because there is an unstable set for expected inflation such that, if expected inflation lies in this set, a hyperinflation occurs. Lowering seignorage later on may not be enough to expel expected inflation from the unstable set. Therefore, agents' fear of a high inflation is self-justified: because agents fear that a hyperinflation may occur expected inflation is heavily affected by temporary shocks, thereby increasing the probability of entering the unstable set. This reproduces what has been termed, somewhat loosely, 'expectation-driven hyperinflations'. Also, we show that the unstable set is larger in countries with high seignorage, which is consistent with the cross-country evidence we mention above.

Our use of a learning model in order to explain observed facts is controversial. Certainly, economists have become accustomed to the assumption of rational expectations, which implies that agents are assumed to know enough about the economy to make the best possible forecasts. Rational expectations have become the norm because, supposedly, one could always find some learning scheme that justified any observed fact. Indeed, this would be a good reason to dismiss models of learning since any given model could not be rejected;. Presumably, if economists abandoned rational expectations, one can use any learning scheme, then economists would find themselves lost in the 'wilderness of irrationality'. We show in this paper how to formulate learning schemes that satisfy certain quasi-rationality requirements; agents in our model do not know everything (as they would under rational expectations) but, because they learn in a fairly rational way, they are not making systematic mistakes; rather, they make quite good predictions within the model at hand. Once a learning scheme is restricted to be quasi-rational within the model, it is no longer true that any observed fact could be explained. Therefore, our way of formulating learning schemes is testable and the model can be rejected.

On the practical side, this paper shows that hyperinflations can be stopped with a combination of heterodox and orthodox policies, thus justifying many of the IMF interventions in hyperinflationary countries. The methodological contribution of the paper is to show that, as long as we carry along adequate equipment for orientation and survival (i.e. as long as we use quasi-rational learning), an expedition into the 'wilderness of irrationality' can be quite a safe and enjoyable experience.

1 Introduction

The goal of this paper is to develop a model that accounts for the main features of the hyperinflations of last decade and to study the policy recomendations that arise from it. The model is standard, except for the assumption of quasi-rational learning. A side contribution of the paper is to show that, if certain rationality requirements are imposed, learning models can be made consistent and falsifiable.

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 hyperinflations could therefore be 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 one of the central messages of Sargent (1986). It is also the logic behind the IMF advice to countries experiencing high inflation rates.

Cross country evidence very strongly supports this story. Hyperinflations have occurred in countries with high seignorage, and the 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 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 would question the validity of the IMF advice to hyperinflationary countries to decrease their seignorage.

In this paper we develop a model that accounts for these observations. These episodes involve very high inflation rates (for instance, in July 1989, monthly inflation for Argentina peaked at 200%) and all we know about the welfare effects of inflation suggest that they are very costly. At the same time, evidence shows that they do not involve an improvement in the fiscal side so they can be considered pure waste.

Sargent and Wallace (1987) explained these hyperinflations as bubble equilibria; their model had a standard Laffer curve with two stationary rational expectations equilibria; hyperinflations could occur as speculative equilibria going from the low-inflation to the high-inflation steady state. Their paper explains how inflation can grow even though seignorage is stable; but it fails to explain other facts observed in the hyperinflationary episodes. Furthermore, bubble equilibria in their model are locally unstable under learning[1]. Our paper builds upon Sargent and Wallace's model by introducing learning; we show that, with this modification, the model matches observations much better. Our model is consistent with the very high hyperinflations, their recurrence, the fact that exchange rate rules temporarily stop hyperinflations, the cross country correlation of inflation and seignorage, and the lack of serial correlation of seignorage and inflation in hyperinflationary countries.

We assume that agents forecast the relevant variables by adopting standard learning schemes that converge to the rational expectations equilibrium; the spikes in the inflation rate can occur as a transition phenomenon due to the presence of an unstable set governing the dynamics of inflation. As the amount of seignorage increases, the rational expectations equilibria are harder to learn, recurrent hyperinflations occur, and this behavior reinforces the use of this particular learning rule.

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 expectations[2]. Indeed, this literature has made enormous progress, and convergence of learning models to rational expectations can now be studied in very general setups. Nevertheless, almost no attempt has been made to explain observed economic facts with models of boundedly rational learning[3]. It is commonly believed that this would entail problems similar to those found in models of adaptive expectations of the pre-rational-expectations era, namely, that there are too many degrees of freedom available to the economist so that the model is not falsifiable, and that expectations are inconsistent with the model[4]. We address these two criticisms by restricting our study to learning mechanisms that produce good forecasts within the model. As our choice of learning mechanism is restricted by the model itself, the model is falsifiable. In addition, since the resulting equilibrium reinforces the use of the learning mechanism (because

[1]See Marcet and Sargent (1989b).

[2]Some examples are Bray (1982), Marcet and Sargent (1989a,b), Evans and Honkapohja (1993) and Woodford (1990). See Sargent (1993), Marimon (1997) and Evans and Honkapohja (1997) for reviews.

[3]Chung (1990) is the only exception we know

[4]The conclusion of Sargent (1993) contains a clear statement of the standard view on the problems that arise in using learning to account for empirical observations.

good forecasts are generated along the equilibrium), agents' expectations are not inconsistent with the model[5].

We do this by defining small upper bounds on the "mistakes" that an agent can make along the learning equilibria. Thus, we only accept small departures from rationality in a way that is precisely defined in the paper. We then show how to construct equilibria in our model that satisfy the small departures in rationality, we show that the model has empirical content and that it replicates the facts we are after. In particular, our model explains why recurrent hyperinflations only occur in countries that collect a high average amount of seignorage, even if there is no apparent relationship between seignorage and inflation over time.

Some papers have presented models that explain some of the facts we are after. 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 seignorage can happen because of moral hazard in the demands for revenue of several branches of government. These papers are interesting additions to the literature, they account for some (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 we are after, providing supporting evidence, and presents the existing literature. Section 3 presents the model and describes the learning mechanism. Section 4 discusses the lower bounds in rationality in a general setup. Section 5 discusses the equilibria in the model of this paper and how the lower bounds on rationality apply to this model. The paper ends with some concluding remarks.

2 Evidence on Recurrent Hyperinflations

A number of countries, including Argentina, Bolivia, Brasil and Perú experienced during the eighties the highest average inflation rates of their history. Stopping inflation was then, almost the only item in the policy agenda of these countries. While the duration and severity of the hyperinflations 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

[5]Recent literature imposing consistency requirements in learning models are Kurz (1994), Fudenberg and Levine (1995) and Hommes and Sorge (1997).

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 hyperinflations. In most circumstances, however, these plans only lower inflation temporarily, and new hyperinflations eventually occur.

3. For a given country where hyperinflations occur, there is a low contemporaneous correlation across time between seignorage and inflation.

4. Across countries there is a clear relation between average inflation and seignorage, namely, hyperinflations only occur in countries where inflation rate is high on average.

Points 2 and 4 can be combined to state the following observation on monetary policy: stabilization plans that do *not* make a permanent fiscal effort (i.e., that do not reduce the average deficit and average seignorage) may be successful in substantially reducing the inflation rate *only* in the short run. Stabilization attempts that focused only on fixing the exchange rate, sometimes with additional price controls, are called *"heterodox"* plans; when the focus is on the fiscal adjustment required to reduce government deficit, they are called *"orthodox"* plans. Most stabilization plans that were successful in reducing inflation substantially and permanently, relied on the fixing of the exchange rate but they also made a severe fiscal adjustment to permanently eliminate the deficit 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 hyperinflations permanently.

Our summary of stylized facts should be uncontroversial[6], but first-hand evidence to support them is provided in figures 1 to 5, which present data on the recent inflationary experiences of Argentina, Bolivia, Brasil and Peru. Inflation rates for selected periods were computed from IFS consumer price indices. These periods have been selected so as to show the main stabilization efforts carried out by each country and the effect they had on the evolution of inflation. Periods when an explicit fixed exchange rate rule was in place are indicated by shaded areas; the end of the shading indicates the date in which convertibility was explicitly abandoned. Figures 1 to 4 illustrate quite clearly stylized facts 1 and 2.

[6]For instance, see Bruno et al. (1988) and (1991).

Figure 5 depicts the evolution of the quarterly[7] inflation rate for Argentina together with the evolution of the seignorage for the period 1983 to 1990. The left hand side vertical axis measures seignorage as a percentage of GNP, while inflation, measured as the log (P_t/P_{t-1}), is on the right hand side vertical axis. Notice that seignorage goes down in certain periods of rapidly increasing inflation, while in other periods the opposite occurs. Also, notice that the level of seignorage that led the spectacular hyperinflation of the second quarter of 1989 (more than 200%) is the same as the one of the first quarter of 1984, with subsequent inflation rates that were below 60%. This documents fact 3.[8]

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 seignorage, money creation, and changes in reserves, and a rule for establishing fixed exchange rates.

Money demand

The demand for real balances is given by

$$P_t = \frac{1}{\phi} M_t^d + \gamma P_{t+1}^e \tag{1}$$

where γ and ϕ 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. It is well known that this equation is consistent with utility maximization and general equilibrium in the context of an overlapping generations model.

[7]The data has been taken from Ahumada, Canavese, Sanguinetti y Sosa Escudero (1993). We use quarterly data for this Figure because the seignorage is typically expressed as a share of GNP.

[8]A closer look at Figure 5, however, points to some interesting facts that merit a more careful empirical investigation. Note, in particular, that seigniorage appears to lead the hyperinflationary bursts. Also, there is some correlation between inflation and seigniorage in the sub samples periods when inflation was not too high; for example, in the periods 80.I-82.IV and 86.II-88.IV. Both of these features are consistent with our model but they are not studied carefully in this version of the paper.

Money supply

We assume government policy rules that mimic those used by governments with hyperinflationary experiences in the last decade. Money creation is driven by the need to finance seignorage; on the other hand, government's concern about current levels of inflation prompts the government to establish a fixed exchange rate rule (ERR) when inflation gets out of hand. Seignorage is given by an exogenous i.i.d. stochastic process $\{d_t\}_{t=0}^{\infty}$ with mean $\overline{d}$ and variance σ_d^2, and it is the only source of uncertainty in the model[9].

In periods with no ERR, the government budget constraint is given by

$$M_t = M_{t-1} + d_t P_t \tag{2}$$

which determines money supply M_t.

Exchange Rate Rules

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

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

where $\bar{\beta}$ is the targeted inflation rate, and P_t^f is the price level abroad. Arbitrage in the international currency market implies that

$$\frac{P_t}{P_{t-1}} = \bar{\beta} \tag{3}$$

and the targeted inflation rate is achieved. In order to implement this policy, the government only needs to know past values of exchange rate and foreign price levels. In the case that targeted inflation $\overline{\beta}$ 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). This price level and equation (1) determine the demand for nominal money. In general, this money demand will not match money supply as determined by (2), so that some variable needs to be introduced in order to satisfy the government budget constraint: the stock of international reserves is the variable that

[9]The i.i.d. assumption is made for simplicity. Most of our results would go through with serially correlated seignorage, but some analytical results would be a little harder to prove.

makes the adjustment, and the government will enforce the ERR by adjusting its reserves. Therefore, the following equation holds in periods of ERR:

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

where R_t denotes the level of international reserves.

Then, we impose the rule that government acts to satisfy

$$\frac{P_t}{P_{t-1}} < \beta^U, \tag{5}$$

where β^U is the maximum inflation tolerated. ERR is *only* imposed in periods when inflation would otherwise violate this bound or in periods where no price level clears the market[10].

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. We have chosen $\overline{\beta}$ close to the lower stationary steady state mean inflation; we then check in the simulations that the loss in reserves is small, which is not surprising, since the model is close to equilibrium. This suggests that, in our model, the policy is sustainable even with a moderate accumulation of foreign reserves in periods outside ERR.

Modelling reserve accumulation formally is beyond the scope of this paper, and it is unlikely to change our results, but it opens up a host of interesting issues. For example, one could study under what conditions the government runs out of reserves during a hyperinflation, so that "orthodox" measures can not be avoided[11]. Alternatively, the accumulation of reserves can also be achieved by maintaining the ERR while the real value of the money stock is increased after the stabilization[12].

[10]Since both the demand and supply of money depend positively on the price level, no equilibrium price exists for high enough β_t. See Marcet and Sargent (1989b) for a detailed description.

[11]One could also argue that a more reasonable policy is to have a permanent ERR, so that equation (3) determines the inflation rate and there can never be hyperinflations. Obviously, from the explanation in this paragraph, the stock of international reserves will have little value, a small shock can create a balance of payment crisis, causing the ERR to be abandoned, and a hyperinflation could start. But once the real value of the money stock is low enough, a new ERR could be established to stop the hyperinflation. Thus, the qualitative nature of the equilibrium would be very similar with this alternative policy. In fact, some of the episodes could be described with this balance of payment-devaluation-hyperinflation cycle. For an early explanation along this lines, see Rodríguez (1980).

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

We have modelled policy in this way because it mimicks the policies followed by South-American countries during the 80's. The issue of *why* these countries followed this kind of policy is interesting, but it is not addressed formally in this paper. We can advance, however, three reasons why this kind of policy rule would be a good rule in our model. First, the fact that ERR has been established only after some periods of high inflation is justified because the value of foreign reserves is high, and a large part of the domestic money supply is backed by those reserves[13]. Second, in principle, any reduction in the government deficit of $e_t(R_t - R_{t-1})$ units would also fix the inflation to $\overline{\beta}$ in periods of ERR. In fact, the reduction in seignorage that is needed to achieve an inflation equal to $\overline{\beta}$ 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 lowering seignorage by the exact amount requires much more information: it can only be implemented when the government knows exactly the model and all the parameter values, including those that determine the (boundedly rational) expectations P^e_{t+1}, and all the shocks. By contrast, an ERR can be implemented only with knowledge of $\bar{\beta}$, the foreign price level, and β^U. The fact that ERR seems to have been the choice of governments under hyperinflationary experiences is further evidence that governments live in a world where agents' expectations and the model generating inflation are not easily determined. A third advantage of establishing ERR for real governments would be that, for institutional reasons, it can be implemented quickly, while lowering government expenses or increasing taxes may take a long time.

An important policy decision is how long to mantain the ERR. Obviously, the longer the ERR is mantained, the closer expected inflation will be to $\overline{\beta}$. In fact, in our simulations, we hold the ERR till expected inflation is close to $\overline{\beta}$ 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 $\overline{\beta}$, β^U and the process for d_t.

[13]This interpretation would suggest that the burst in inflation at the begining 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 Equilibria with Rational Expectations and ERR:

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 below a certain maximal level, the model has two stationary equilibria with constant inflation levels (called low- and high-inflation equilibria), and a continuum of bubble equilibria that converge to the high-inflation equilibrium[14]. These bubble equilibria can be interpreted as hyperinflations.

The main motivation behind the work of SW was precisely to explain 'fact 3' in section 2; indeed, their bubble equilibria explain this fact qualitatively[15]. 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 hyperinflations on and off. Even if one accepts rational sunspots (where agents coordinate perfectly on a particular non-fundamental variable) 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 small[16]. Fact 4 is contradicted: the long run inflation rate in any rational bubble equilibrium is *lower* when seignorage is *higher*; therefore, the model under RE predicts that hyperinflations are *less* severe in countries with high seignorage. Fact 2 is addressed by Obstfeld and Rogoff (1983) and Nicolini (1996); they introduce ERR that goes into effect if inflation goes beyond a certain level. Their results show that the threat of convertibility eliminates bubble equilibria. Thus, once ERR is introduced, the model is inconsistent with the existence of hyperinflations; since convertibility was actually introduced in the Latin-American countries that we are studying, one would think that convertibility became a credible threat, and the fact that hyperinflations emerged again is inconsistent with RE.

Marcet and Sargent (1989b) studied stability of rational expectations equilibria in the SW model under least squares learning[17]. They found that, only the low-inflation equilibrium is *locally* stable; the high-inflation equilibrium is always unstable. Taken literally, these results would say that bubble

[14]Since our model is slightly different from SW, we reproduce these results in appendix 1. There we show that, in our case, existence of a stationary equilibrium depends not only on the average value of seignorage, but also on its standard deviation

[15]A wide empirical literature tested the existence of a speculative component in the German hyperinflation of the twenties. A short summary of the literature and a test of bubble versus stationary equilibria in the SW model can be found in Imrohoroglu (1993).

[16]This is documented in our discussion of Figure 7 in subsection 5.5 below.

[17]Marcet and Sargent (1989b) is a special case of the present paper when uncertainty is eliminated, β^U is arbitrarily high, and agents forecast P_t by regressing it on P_{t-1}. In addition, they only study local stability.

equilibria can not be learned by agents. Therefore, if learning is taken seriously as a stability criterion, the model of Sargent and Wallace does not have hyperinflations and, again, none of the above facts is appropriately matched.

Note that the learning mechanisms considered by Marcet and Sargent (1989b) change the dynamics of the model in a very suggestive way, if we try to understand the hyperinflationary processes. The low steady state is *locally* stable, thus the economy can live close to it for a rather long period. However, the high steady state is unstable, so if the economy, by some reason, goes beyond the high steady state, it may enter an unstable region that can, potentially, explain the spikes in the inflation rates observed in the data. This feature of the model with learning constitutes the core of the dynamics in the current paper.

In the next section we propose several criteria to asses models with quasi-rational learning and argue that our model generates learning equilibria that are robust to the well known 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'. That expectations were ad-hoc 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; hence, rational agents would be likely to abandon their adaptive expectations after a while, and the predictions of the model would be invalid[18]. 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'. Indeed, it is a much documented and well accepted fact that 'economic agents do not make systematic mistakes'.

The rational expectations hypothesis is, nowadays, the most commonly used paradigm in macroeconomics, mainly, because it solved these two issues: under RE, expectations are determined by the model; after some time agents will just realize that they are doing the right thing, and they will never abandon their rational expectations.

[18]A careful justification of this position can be found in the conclusion of Sargent (1993).

In this paper we will show that by introducing boundedly rational learning in the model of section 3, one can match the stylized facts of section 2 much better that with the existing alternative RE models available in the literature. One could simply argue that hyperinflations are such confusing events that it is reasonable to assume non-RE behavior, but a natural question comes to mind: 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. These upper bounds often rule out RE, but it might seem that they accept too many models of learning. The dilemma is: RE is too demanding of agents' rationality; on the other hand, by moving away from RE we may just fall back into old mistakes and the 'wilderness of irrationality'. It might seem that Bayesian learning is a way out of this dilemma, but the literature has recognized many problems with this approach[19].

We take an alternative road; we only allow for *small* deviations from rationality, both along the transition and asymptotically. Our hope is that this solves the issue of falsifiability and it does not violate reasonable definitions of rationality (or quasi-rationality). In other words, given an economic model and some empirical observations, we look for learning mechanisms that satisfy certain *lower* bounds on rationality and that the model explains the observed behavior of the economy. In section 5 we will show this small departure from rationality generates equilibria that are quite different from RE, precisely in the direction of improving the match of empirical observations, even if we consider countries that were following different policies[20].

[19]First, Bayesian learning requires that agents know perfectly part of the model in order to form the likelihood function; which simply begs the question of 'how did agents learn the likelihood function?'. Second, in models with endogenous state variables (such as the model of section 3, where money, or past inflation, are state variables), Bayesian learning requires, in principle, that agents use a law of motion that changes from period to period and to remember the whole past, and it is hard to justify how agents could learn a law of motion. Finally, the literature has also accumulated a number of paradoxes generated by Bayesian learning, among them, that small mistakes in the formulation of the prior will cause agents to make very bad predictions, since errors accumulate over time. See, for example, Bolton and Rustichini (1995) and Marimon (1997) for descriptions of such paradoxes.

[20]Bolton and Rusticcini (1995) and Marimon (1997) also argue that learning can be used for more than a stability criterion.

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

$$x_t = g(x_{t-1}, x^e_{t+1}, \xi_t, \eta) \tag{6}$$

where g is determined by market equilibrium and agents' behavior and η is a vector of parameters in the economy that enters in the determination of x_t. This includes, for example, parameters of government policy. x_t is a summary of all the variables in the economy, x^e_{t+1} is agents' expectation of the future value of x_t, and ξ_t is an exogenous shock. For example, in our model, x_t is inflation and the money supply, ξ_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 η includes $\gamma, \phi, \overline{\beta}$ and β^U.

Agents' expectations are given by

$$x^e_{t+1} = z(\beta_t(\mu), x_t) \tag{7}$$

where $\beta_t(\mu)$ are certain statistics inferred from past data. The function z is the forecast function that depends on today's state and the statistics. These statistics are generated by a learning mechanism f

$$\beta_t(\mu) = f(\beta_{t-1}(\mu), x_t, \mu). \tag{8}$$

where μ are certain learning parameters that govern how past data is used into forming the statistics. The statistics are only a function of observed data, not of the true model or the true parameters.

The learning mechanism f says how new information is incorporated into the statistics; the learning parameters μ govern, for example, the weight that is given to recent information. In the following section we discuss several alternatives for f. For now, (z, f, μ) are unrelated to the true model (g, η), but later in this section we will define bounds on rationality that amount to imposing restrictions on the space of (z, f, μ) for a given model (g, η).

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

$$P^e_{t+1} = \beta_t P_t \tag{9}$$

where β_t is expected inflation, estimated somehow from past data.

Equations (6), (7) and (8) determine the equilibrium sequence for given learning parameters μ. Obviously, since the process for x_t is self-referential, it depends on the parameter μ; this dependence will be left implicit in most of the paper, and we will write x^μ_t if we want to make the dependence explicit.

Let $\pi^{\epsilon,T}$ be the probability that the perceived errors in a sample of T periods, will be within $\epsilon > 0$ of the conditional expectation error:

$$\pi^{\epsilon,T} \equiv P\left(\frac{1}{T}\sum_{t=1}^{T}\left[x_{t+1} - x_{t+1}^{e}\right]^2 < \frac{1}{T}\sum_{t=1}^{T}\left[x_{t+1} - E_t^{\mu}(x_{t+1})\right]^2 + \epsilon\right) \qquad (10)$$

where E_t^{μ} is the true conditional expectation under the learning model. For small ϵ, this is the probability that, after T periods, the (sample) prediction error made by agents in the model, is almost as small as the (sample) prediction error made with the conditional expectation.

The first lower bound on rationality we propose is:

Definition 1 Asymptotic Rationality (AR): *the learning mechanism (z, f, μ) satisfies AR in the model (g, η) if :*

$$\pi^{\epsilon,T} \to 1 \text{ as } T \to \infty \quad \text{for all } \epsilon > 0.$$

This requires the perceived forecast to be at least as good as the forecast with the conditional expectation asymptotically. In this case, agents would not have any incentive to change their learning scheme after they have been using it for an arbitrarily long time.

AR seems like a minimal requirement; it rules out behavior that is inconsistent forever. It rules out adaptive expectations for most stochastic models, or learning models where agents exclude some relevant state variables in the forecasting rule z. It is satisfied often in models of least squares learning if they converge to RE. Perhaps surprisingly, AR excludes many models that would be termed 'rational equilibria' in Kurz (1994), since this author allows for agents to make systematic mistakes, as long as these mistakes are not contemplated in the prior distribution.

Similar concepts can be found in the literature[21]. However, it is best to think of AR as only a minimum requirement, because it admits learning mechanisms that generate very bad forecasts along the transition for an arbitrarily long time. For example, if we assume that β_t is estimated by OLS in a model that generates recurrent hyperinflations, agents in the model would be adapting to sudden hyperinflations more and more slowly as time went by, due to the fact that least squares learning gives less and less importance to recent events as time goes by. This would imply that agents make worse and worse predictions in the successive hyperinflations.

[21]This requirement was implicitely 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.

For this reason, the next two additional restrictions impose that good predictions are also generated *along the transition* :

Definition 2 Epsilon-Delta Rationality (EDR): *the learning mechanism* (z, f, μ) *satisfies EDR for* (ϵ, δ, T) *in the model* (g, η) *if :*

$$\pi^{\epsilon,T} \geq 1 - \delta$$

If *EDR* is satisfied for small (ϵ, δ), agents are unlikely to switch to another learning scheme after period T, even if they were told "the whole truth"[22]. Clearly, once *AR* is satisfied, it is only interesting to study this probability for T moderately high: if T is too low, the sample mean of the prediction error has no chance to settle down, and if T is too high, the criterion is still satisfied even for learning schemes that perform very poorly, just as with *AR*. The precise application that the researcher wants to explain should suggest an interesting value for T.

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 researcher is supposed to report to the reader the probabilities $\pi^{\epsilon,T}$ for a given model and, hopefully, convince the reader that these probabilities are 'sufficiently' high for this learning scheme. *EDR* is quite stringent, but we will see that it is satisfied in our model for certain parameter values, even for very strict ϵ and δ.

The next (and last) bound on rationality requires the agent to use values of μ that are nearly optimal within the learning mechanism f. Denote by $\tilde{\beta}_t(\mu, \mu')$ the forecast produced by the learning parameter μ' when all agents are using the parameter μ :

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

Definition 3 Internal Consistency (IC): *Given* (g, η), (z, f, μ) *satisfies IC for* (T,ϵ) *if*

$$E\left(\frac{1}{T}\sum_{t=1}^{T}(x_{t+1}^{\mu} - z(\beta_t(\mu), x_t^{\mu}))^2\right) \leq$$

$$\leq \min_{\mu'}\ E\left(\frac{1}{T}\sum_{t=1}^{T}\left(x_{t+1}^{\mu} - z(\tilde{\beta}_t(\mu, \mu'), x_t^{\mu})\right)^2\right) + \epsilon. \qquad (11)$$

[22]Bray 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 paper carries the flavor of *EDR* in the sense that it requires that learning schemes are not inconsistent even along the transition.

Thus, if the mechanism satisfies this bound after T periods, agents do not perceive alternative μ's as being much better on average[23]. One would expect IC to be more restrictive than AR. In particular, any μ satisfying AR also verifies internal consistency for T large enough. Again, it only makes sense to study IC in the context of 'moderately high' T in order to allow the parameters to settle down and, also, in order for IC to be restrictive.

The first two bounds compare the performance of the agent that is learning relative to an external agent who knows the best prediction that can be computed from knowledge of (f, μ, h, g, η), i.e., the right model, the probability distributions and, in addition, the learning mechanism that all other agents are using. The bound IC, instead, compares the agents of the economy, with other agents that are forced to use the same family of mechanisms f in their forecasts, but are allowed to pick alternative parameter values μ. This last bound replicates the intuition of rational expectations, in the sense of looking for an approximate fixed point, in which the equilibrium expectations that the consumers are using, minimize the errors *within the mechanism* f. Notice that this restriction may cause agents under different environments to use different learning parameters; for example, in our model, it will cause agents in high seignorage countries (say, Argentina) to use a different learning parameter from agents in low seignorage countries (say, Switzerland). These criteria could be readily generalized to more complicated models or to objective functions other than the average prediction error.

Rational expectations can be interpreted as imposing extreme versions of the second and third bounds. Obviously, RE satisfies AR. Requiring EDR for all ϵ, T and δ is the same as imposing rational expectations. Also, if the REE is recursive, if z uses the appropriate state variables and is a dense class of functions (for example, polynomials), and we impose IC for any ϵ, T, we are left with rational expectations. In this sense, a learning mechanism that satisfies all the above bounds for small ϵ, δ can be interpreted as a small deviation from full rationality.

5 Learning Equilibrium

In this section, we propose a learning mechanism that combines least squares learning with tracking. We show that the mechanism satisfies AR for a wide rangr of learning parameter values. We argue that, even if AR is satisfied, it puts almost no restriction on the transition, so that it allows for very poor performance of the learning mechanism in the model. Then, we define an equilibrium as a set of parameters that satisfies IC and discuss how to find

[23]Evans and Honkapohja (1993) propose to use a related criterion.

such equilibria. Finally, we show that the interesting equilibrium parameters values according to *IC* also satisfy *EDR*.

We also show how the *IC* learning equilibrium is able to match the data by providing some analytical results as well as describing the outcomes of simulations.

5.1 The Learning Mechanism

We assume a learning mechanism given by

$$\beta_t = \beta_{t-1} + \frac{1}{\alpha_t}\left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}\right) \tag{12}$$

That is, perceived inflation is updated by a term that depends on the last prediction error[24] weighted by $1/\alpha_t$. This is a simple version of stochastic approximation algorithms, where the weights are often denoted the 'gain' sequence. Equation (12), together with the evolution of the gain sequence, determines the learning mechanism f in equation (8).

In stochastic approximation, the gain sequence is often specified exogenously. One common assumption is

$$\alpha_t = \alpha_{t-1} + 1 \tag{13}$$

where α_0 is set exogenously, typically equal to 0. In this case, $\alpha_t = t$, and simple algebra shows that

$$\beta_{t+1} = \frac{1}{t}\sum_{i=1}^{t} \frac{P_i}{P_{i-1}} \tag{14}$$

so that perceived inflation is just equal to the sample mean of past inflations or, equivalently, it is the result of a least squares regression of inflation on a constant.

Another exogenous gain sequence is $\alpha_t = \tilde{\alpha} > 1$; these have been termed 'tracking' or 'constant gain' algorithms. In this case, perceived inflation

[24] This formula implies that agents do not use today's inflation in order to formulate their expected inflation; the last observed inflation used to formulate expectations at t is inflation at $t-1$. This assumption is made purely for convenience, and it is often made in models of learning; it simplifies solving the model by avoiding simultaneity in the determination of perceived inflation and actual inflation. Including today's inflation in β_t would make it even easier for the learning scheme to satisfy the lower bounds, since information about prices would be used very quickly and, in a hyperinflationary world, inflation may change from one period to the next. Also, the dynamics of the model are unlikely to change. .

satisfies

$$\beta_{t+1} = \frac{1}{\bar{\alpha}}\sum_{i=0}^{t}\left(1-\frac{1}{\bar{\alpha}}\right)^{i}\frac{P_{t-i}}{P_{t-i-1}} + \beta_0\left(1-\frac{1}{\bar{\alpha}})\right)^{t+1} \qquad (15)$$

so that past information is now a weighted average of past inflations, where the past is discounted at a geometric rate. This learning scheme produces better forecasts when there is a sudden change in the environment, because it adapts more quickly to such a change. [25,26].

Notice that least squares (14) gives equal weight to all past observations, while tracking (15) gives more importance to recent events.

Unfortunately, both alternatives are likely to fail the lower bounds on rationality of section 4 in a model that exhibits recurrent hyperinflations. The reason is that with recurrent hyperinflations there are periods of stability and periods of instability. Tracking (15) performs poorly in periods of stability: since it does not converge to RE (because the RE equilibrium has a constant perceived inflation, and perceived inflation under pure tracking does not converge to a constant) it does not even satisfy our weakest requirement.

On the other hand, least squares does not generate 'good' forecasts because, along a hyperinflation, it will be extremely slow in adapting. In those periods, 'tracking' will be a better idea, and least squares does not satisfy *EDR* or *IC*.

We will specify a learning mechanism that uses OLS in stable periods and it switches 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, α_t increases over time at the same rate as in least squares, but in periods where a large prediction error is detected, α_t goes down to a fixed value $\bar{\alpha}$, mimicking the 'tracking' algorithms. Formally, the gain sequence follows

$$\begin{aligned} \alpha_t &= \alpha_{t-1}+1 \qquad && \text{if } \left|\frac{\frac{P_{t-1}}{P_{t-2}}-\beta_{t-1}}{\beta_{t-1}}\right| < \nu \\ &= \bar{\alpha} && \text{otherwise} \end{aligned} \qquad (16)$$

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

[26]In 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 better forecasts. For example, if we assumed that seignorage is autoregressive of order 1, in order to satisfy lower bounds on rationality, expected inflation would have to depend on current seignorage, the parameter multiplying seignorage could be estimated with 'tracking'. In that case, tracking would be fundamentally different from adaptive expectations, both because of its functional form and because adaptive expectations will not satisfy lower bounds on rationality.

Thus, the expectation formation mechanism is the same whether or not ERR is enforced in a given period and regardless of the parameters of the model. 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 $\bar{\alpha}$.

In summary, we assume that the gain sequence of the learning mechanism is updated according to OLS in periods of stability, but it uses constant gain (or tracking) in periods of instability. The learning mechanism f is fully described by equations (12) and (16), and the learning parameters μ are given by ν and $\bar{\alpha}$.

5.2 Learning and Stylized Facts

The variables we need to solve for are $\left\{\frac{P_t}{P_{t-1}}, \beta_t, \alpha_t\right\}$. Simple algebra shows that equilibrium inflation satisfies

$$\frac{P_t}{P_{t-1}} = H(\beta_t, \beta_{t-1}, d_t) \tag{17}$$

where

$$\begin{aligned} H(\beta_t, \beta_{t-1}, d_t) &= \frac{1-\gamma\beta_{t-1}}{1-\gamma\beta_t - d_t/\phi} && \text{if } 0 < \frac{1-\gamma\beta_{t-1}}{1-\gamma\beta_t - d_t/\phi} < \beta^U \\ &= \overline{\beta} && \text{otherwise,} \end{aligned}$$

Equations (12), (16) and (17) define a system of stochastic, second-order difference equations; characterizing the solution analytically is unfeasible due to the high non-linearities of this system.

Let $h(\beta, d) \equiv H(\beta, \beta, d)$. We can provide some intuition for the behavior of the model using the fact that, if $\beta_t \simeq \beta_{t-1}$, we have $P_t/P_{t-1} \simeq h(\beta_t, d_t)$. In this sense, the graph of $h(\cdot, d)$ in Figure 6 provides an approximation of the actual inflation rate as a function of perceived inflation. The first graph corresponds to a low average seignorage $E(d_t) = \overline{d}$. The limiting rational expectations equilibrium $\beta_{RE}^1 = S(\beta_{RE}^1)$ is close to the lower fixed point of $h(\cdot, \overline{d})$ (see appendix 1). The horizontal axis can be split into the intervals S, **U** and **ERR**.

If $\beta_t \in$ S, inflation is closer to β_{RE}^1 than perceived inflation and perceived inflation is pushed in the direction of β_{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 always higher than β_t, so that a hyperinflation will occur until the upper bound β^U would be violated, and perceived inflation falls in the set **ERR** where a fixed exchange rule will be 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 α_t is not yet close to zero, initially high perceived inflation, the second-order dynamics which add momentum to increasing inflation, etc.

It is clear that a stochastic model is required to generate recurrent hyperinflations. If there were no shocks, once the economy is at the stable set, there is no force to take it out of it. Thus, a crucial difference with the deterministic case is that even if the economy is at the stable set, the probability of having a hyperinflations is positive if seignorage is large enough. Another difference with the deterministic case is that the dynamics are governed by a function S which increases with the variance of seignorage, so that the stable region also shrinks with a higher σ_d^2 (see appendix 1). This implies that, if the variability of seignorage is high, the probability of hyperinflations is high for two reasons: *i)* given a value for today's beliefs, it is more likely to obtain a large enough shock that will send the next beliefs to the unstable region U, *ii)* the stable region S shrinks.

Notice that the economy is likely to end up in U if the gain $1/\overline{\alpha}$ is large. In that case, perceived inflation is more heavily influenced by shocks to actual inflation; if $1/\alpha_t$ is arbitrarily close to zero and initial inflation starts out in S (and if ν is large enough), hyperinflations are impossible. But if hyperinflations occur, agents will set the weight $1/\alpha_t = 1/\overline{\alpha}$, so that the presence of hyperinflations prompts agents to pay more attention to recent observations which, in turn, makes it more likely that hyperinflations occur.

This intuition suggests that the model is consistent with stylized fact 1, since a number of hyperinflations 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 β_t is in the set U, inflation will grow on average even if seignorage stays roughly constant, which is consistent with fact 3.

To analyze fact 4, consider the second graph of Figure 6, 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, where the economy tends to live; hence, 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. Our previous discussion on the effects of σ_d^2 also suggests that a high σ_d^2 increases the probability of a hyperinflation, a fact that is roughly consistent with the data, but we will not pursue this property of the model any further in this paper.

5.3 Asymptotic Rationality.

It is clear that, for ν small enough, the learning mechanism does not converge to rational expectations[27]. Therefore, we show that convergence to RE happens if ν is large enough and, in that case, AR obtains.

Proposition 1 *Under the assumptions of Appendix 1, if $\beta^U < \infty$ and $\beta_0 > 0$, if ν large enough, then*

$$\beta_t \longrightarrow \beta^1_{RE} \qquad a.s.,$$

and Asymptotic Rationality obtains.

Proof

First of all, we show that there is a ν large enough for which the learning mechanism stays in the OLS form.

Let $\tilde{\beta} \equiv \min(1, \beta_0, \overline{\beta})$. It is clear from (12) and $H(\cdot)$ defined in section 5.2 that, if β_{t-1} and $H(\beta_{t-1}, \beta_{t-2}, d_{t-1}) > \tilde{\beta}$, then $\beta_t > \tilde{\beta}$. Since $H(\beta, \beta', d) > \min(1, \overline{\beta})$ for all $\beta, \beta' \in R_+$ and all $d \in [K^-, K^+]$, it follows by induction that $\beta_t > \tilde{\beta}$ for all t with probability one. Now, letting

$$\overline{\nu} \equiv \max_{\beta,\beta' \in [\tilde{\beta},\infty),\ d \in [K^-,K^+]} \left| \frac{H(\beta, \beta', d) - \beta}{\beta} \right|,$$

since H is bounded for $\beta^U < \infty$, $\overline{\nu} < \infty$. Therefore, for any $\nu > \overline{\nu}$,

$$\left| \frac{\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}}{\beta_{t-1}} \right| = \left| \frac{H(\beta_{t-1}, \beta_{t-2}, d_t) - \beta_{t-1}}{\beta_{t-1}} \right| < \nu$$

with probability one for all t. This implies that $\alpha_t = \alpha_{t-1} + 1$ for all t and the learning mechanism is simply OLS.

Now, assume that $\beta_t > \beta^2_{RE}$ for *all* t large enough; then β_t would grow until it entered the ERR region and, if the ERR is enforced for sufficiently many periods, β_t would go back inside $[0, \beta^2_{RE} - \epsilon]$, which is a contradiction. Therefore, β_t stays in the set $[0, \beta^2_{RE}]$ infinitely often with probability one. Appendix 1 shows that, in this case, β_t converges to β^1_{RE} almost surely.

[27]This is because, even for β_t very close to the rational expectations equilibrium, it will eventually happen that $\left| \frac{\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}}{\beta_{t-1}} \right| > \nu$ and α_t will increase to $\overline{\alpha}$. When this happens, β_t reacts to the current shock, and it never converges.

The rest of the proof simply shows that, if the learning scheme converges to β^1_{RE}, then the sample mean square errors converge to the best forecasts. First of all, $\beta_t \rightarrow \beta^1_{RE}$ a.s. implies

$$\left| \frac{1}{T} \sum_{t=1}^{T} \left[\frac{P_t}{P_{t-1}} - \beta_{t-1} \right]^2 - \frac{1}{T} \sum_{t=1}^{T} \left[\frac{P_t}{P_{t-1}} - \beta^1_{RE} \right]^2 \right| \rightarrow 0 \qquad \text{a.s.} \tag{18}$$

as $T \rightarrow \infty$. Also, since H is a bounded function, $\beta_t \rightarrow \beta^1_{RE}$ a.s. implies

$$\left| \frac{P_t}{P_{t-1}} - \frac{P_t^{RE}}{P_{t-1}^{RE}} \right| = \left| H(\beta_{t-1}, \beta_{t-2}, d_t) - H(\beta^1_{RE}, \beta^1_{RE}, d_t) \right| \rightarrow 0 \qquad \text{a.s.}$$

as $t \rightarrow \infty$, so that

$$\left| E_{t-1}\left(\frac{P_t}{P_{t-1}} \right) - \beta^1_{RE} \right| = \left| E_{t-1}\left(\frac{P_t}{P_{t-1}} \right) - E_{t-1}\left(\frac{P_t^{RE}}{P_{t-1}^{RE}} \right) \right| \rightarrow 0 \qquad \text{a.s.}$$

This equation, together with (18), imply that

$$\left| \frac{1}{T} \sum^{T} \left[\frac{P_t}{P_{t-1}} - \beta_{t-1} \right]^2 - \frac{1}{T} \sum^{T} \left[\frac{P_t}{P_{t-1}} - E_{t-1}\left(\frac{P_t}{P_{t-1}} \right) \right]^2 \right| \rightarrow 0 \qquad \text{a.s.} \tag{19}$$

as $T \rightarrow \infty$, and

$$\pi^{\epsilon,T} \equiv P\left(\frac{1}{T} \sum_{t=1}^{T} \left[\frac{P_t}{P_{t-1}} - \left(\frac{P_t}{P_{t-1}} \right)^e \right]^2 < \frac{1}{T} \sum_{t=1}^{T} \left[\frac{P_t}{P_{t-1}} - E_t^{\mu}\left(\frac{P_t}{P_{t-1}} \right) \right]^2 + \epsilon \right) \rightarrow 1 \qquad \text{a.s.}$$

as $T \rightarrow \infty$ for any $\epsilon > 0$.□

This discussion shows that AR is satisfied. The problem is that AR poses no restriction on the choice of the parameter $\overline{\alpha}$. As $\overline{\alpha}$ is a key parameter determining the probability of experiencing a hyperinflation, AR is not sufficient to determine interesting lower bounds in the context of this model. For example, if agents used pure OLS, even if AR were satisfied, they would be making very large forecasting errors whenever a hyperinflation happened, since OLS tells them to give very little importance to recent events.

5.4 Internal Consistency

From the intuition given in subsection 5.2, when $1/\overline{\alpha}$ is high, hyperinflations are likely to occur. Since $1/\overline{\alpha}$ high is likely to generate good forecasts in a hyperinflation, setting $1/\alpha_t = 1/\overline{\alpha}$ is likely to generate good forecasts within the model, and there is a chance for IC to be satisfied for $1/\overline{\alpha}$'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 $\overline{\alpha}$.

Definition 4 *A sequence* $\left\{\frac{P_t}{P_{t-1}}, \beta_t, M_t\right\}$*, together with* $\overline{\alpha}$ *is an* ϵ, T *equilibrium if:*

1. *Given* $\overline{\alpha}$, $\left\{\frac{P_t}{P_{t-1}}, \beta_t, M_t\right\}$ *satisfy (17), (12), (16) at all periods.*

2. *Given* $\left\{\frac{P_t}{P_{t-1}}, \beta_t, M_t\right\}$, $\overline{\alpha}$ *satisfies*

$$E\left(\frac{1}{T}\sum_{t=1}^{T}\left(\frac{P_{t+1}}{P_t} - \beta_t(\overline{\alpha})\right)^2\right) \leq \min_{\overline{\alpha}'} E\left(\frac{1}{T}\sum_{t=1}^{T}\left(\frac{P_{t+1}}{P_t} - \tilde{\beta}_t(\bar{\alpha}, \bar{\alpha}')\right)^2\right) + \epsilon$$

here $\tilde{\beta}_t(\bar{\alpha}, \bar{\alpha}')$ is consistent with our notation in the Definition of IC and, therefore, it is the forecast of inflation obtained by a forecaster who uses $\bar{\alpha}'$ instead of $\overline{\alpha}$ in equations (12) and (16).

The solution of the model is a highly non-linear second order difference equation, so characterizing analytically the $\overline{\alpha}$'s that satisfy IC is impossible. We solve the model numerically and search numerically for $\bar{\alpha}$ 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 (γ and ϕ) and the government policy. We choose values ($\gamma = 2.7$ and $\phi = 2.56$) in order to replicate some patterns of the Argentinean experience during the 80's, for details see the appendix 2. For the standard deviation of the seignorage we used 0.01[28]. The parameter ν , which measures the error level at which the learning rule sets alpha equal to the base value 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 $\beta^U = 50$. The ERR is enforced until expected inflation is inside the stable set.

For the initial condition of the beliefs we have chosen $\beta_0 = \beta^1_{RE}$. Our purpose is to show that a small deviation from RE can generate very different results and explain better some stylized facts; so, this choice makes it more

[28] We also used a value for sigma equal to 0.005. The results were similar except that, as expected, the probabilities of hyperinflations were lower.

difficult for our model to actually generate different results under learning than under rational expectations[29]. For the specified parameters, the maximum level of average seignorage in the model for which a REE exists is 0.05. In the same spirit as with the initial condition, we have chosen values of the average seignorage for which a REE exists[30]. In order to quantify the relevance of the average seignorage (fact 4), we performed our calculations for four different values of the seignorage: $\bar{d} = 0.049$, 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 7. That realization was obtained with $\bar{d} = 0.049$ and $1/\overline{\alpha} = 0.2$. 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 rational expectation equilibria, to show how the model can generate inflation rates that are way higher than them.

This graph displays some of the stylized facts in the learning model[31]. 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. Then 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, β_t would be revised according to the OLS rule $\alpha_t = \alpha_{t-1} + 1$, 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 (facts 1, 2 and 3). In order to reduce (or eventually eliminate) the chances of having a new burst, the government must reduce the amount of seignorage collected and increase the size of the stable set (an "orthodox" stabilization plan); this would separate the two horizontal lines and it would stabilize the economy permanently around the low stationary equilibrium. Establishing ERR just before this would help stabilize the expectations of agents more

[29]For example, it would be trivial to generate at least one hyperinflation by choosing $\beta_0 > \beta^2_{RE}$.

[30]It would be trivial to generate hyperinflations if average seignorage was too high for a REE to exist.

[31]The 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 β^1_{RE}. For bubble equilibria, inflation would grow towards the horizontal line of β^2_{RE}.

quickly, so there is room for a 'heterodox' intervention as well.

An important aspect of the calibration is the choice of the learning parameter. We look for values of $\overline{\alpha}$ that satisfy the lower bound criterion IC for $(\epsilon, T) = (0.01, 120)$. This value of T is chosen to represent 10 years, the length of the hyperinflationary episodes we are studying; the value of ϵ is just chosen to be 'small', it will be clear below how the results may change if this parameter changes. According to Definition 4 (which is an application of IC to our model), a particular $\overline{\alpha}$ will be an $(\epsilon, T) = (0.01, 120)$ equilibrium if, for $T=120$, the squared sum of errors is within 0.01 of the minimum across all possible alternative values of $\bar{\alpha}'$.

To find numerically those values of $\bar{\alpha}$ that satisfy IC we proceed as follows: we define a grid of $1/\overline{\alpha} \in [0, 1.2]$ separated by intervals of length 0.1; the same grid is used both for $1/\overline{\alpha}$ and the alternative learning parameters $1/\bar{\alpha}'$ considered. We compute the mean squared errors in the right side of (11) by Monte-Carlo integration[32], and we find the minimum over $1/\bar{\alpha}'$ for each $1/\overline{\alpha}$. Figures 8 to 11 show the result of these calculations: in the horizontal axis we plot $1/\overline{\alpha}$, while the vertical axes plots $1/\bar{\alpha}'$. 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 $(\epsilon, T) = (0.01, 120)$ is found when the dark area cuts the 45 degree line: if all agents use one of these values of $1/\overline{\alpha}$, the equilibrium reinforces the use of that learning parameter, in the sense that agents could not do better (up to ϵ) by using an alternative value of the learning parameter, within this learning mechanism.

Tables 1 to 3 report the probabilities of having n hyperinflations in 10 years for different values of average seignorage and for those values of $1/\overline{\alpha}$ that satisfy the IC criterion.

Figure 8 presents the results for a low value $\bar{d} = 0.043$. In this case, only $1/\overline{\alpha} = 0$ 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. Low values of $1/\overline{\alpha}$ satisfy IC because hyperinflations do not occur; giving too much importance to recent observations does not generate good forecasts, so a low $1/\overline{\alpha}$ is a good choice within the model.

Figure 9 shows the results of increasing average seignorage to 0.045. In this case the criterion is satisfied for all values of alpha between 0.5 and

[32]More specifically, we draw 1000 realizations of $\{d_1, ..., d_{120}\}$, find the equilibrium inflation rates for each realization, we compute the sample mean square error for each alternative $1/\overline{\alpha}'$ in the grid, and we average over all realizations. Notice that Monte-Carlo is the only feasible integration procedure, since the expectations in (11) involve 120 random variables.

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 α's generate good forecasts, and the hyperinflationary behavior is reinforced. Figures 10 and 11 and tables 2 and 3 show that, as the mean of seignorage increases, pseudo-rational learning is consistent with the observation of hyperinflations. In fact, 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 $1/\overline{\alpha} = 0.2$, an agent using an alternative alpha equal to zero, which is the collective behavior that replicates the REE, will make larger MSE than the agent using $1/\overline{\alpha} = 0.2$. The reason is that in equilibrium there are many hyperinflations, and the agent that expects the REE will not make good forecasts.

Whenever there exist equilibria with hyperinflations, there is multiplicity of equilibria (several $1/\overline{\alpha}$'s satisfy IC). The results do not change much when different $1/\overline{\alpha}$'s satisfying *IC* are used.[33]

The numerical solutions show that the chances of facing a hyperinflation during the transition to the rational expectations equilibrium, depend on both the sensitivity of the learning rule with respect to changes in prices and on the size of the deficit. The lower the deficit, the lower the chances of experiencing a hyperinflation. In our model, the sensitivity of the learning rule depends on the size of the deficit. The larger the deficit, the larger will be the optimal sensitivity of the learning rule, which increases the chances of having a hyperinflation.

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

5.6 Epsilon-Delta Rationality (EDR):

In this section we show that in the equilibria with hyperinflations discussed above, the criterion *EDR* is satisfied if the highest admissible inflation β^U is large enough, for values of δ that are closely related to the probability of experiencing a hyperinflation. This is because, along equilibria with hyperinflations, the conditional expectation can be arbitrarily high due to the fact that the mapping H has an asymptote, however, the actual value of inflation is never so high. Thus, for every realization of the shocks such that

[33]The REE $1/\overline{\alpha} = 0$ is always an equilibrium; obviously, this is an artifact of having chosen $\beta_0 = \beta^1_{RE}$; when initial beliefs are far appart from the REE, then $1/\overline{\alpha} = 0$ is no longer *IC*.

a hyperinflation is experienced, the learning forecast can do better than the conditional expectation with very high probability in finite samples.

Proposition 2 *In the model of section 3, under the regularity assumptions 1, 2 and 3 of appendix 1, given any (ϵ, T), there is a β^U large enough such that*

$$\pi^{\epsilon,T} \geq P(ERR \text{ at some } t \leq T)$$

where $P(ERR$ at some $t \leq T)$ is the probability that the government implements ERR at some point before T.

Proof

Fix ϵ, T. We first show that, if $\beta^U = \infty$, given a period t and a realization such that $P(ERR \text{ at } t+1 \mid d_t, d_{t-1}, ...) > 0$, then $E_t(P_{t+1}/P_t) = \infty$. Since inflation is given by equation (17); since β_{t+1} and β_t are both in the information set at t, we have

$$E_t\left(\frac{P_{t+1}}{P_t}\right) = \int_{K^-}^{d^{ERR}} \frac{1-\gamma\beta_t}{1-\gamma\beta_{t+1}-d/\phi} \mathrm{d}F_{d_t}(d) + P(ERR \text{ at } t+1 \mid d_t, d_{t-1}, ...)\overline{\beta} \tag{20}$$

where K^- is the lower bound on the distribution of seignorage and $d^{ERR} \equiv \phi - \gamma\beta_{t+1}\phi$ is the lowest value of the shock at which ERR will have to be enforced at $t+1$. Notice how the integral in (20) corresponds to the values of d_{t+1} for which there is a price level that clears the market and the first branch of (17) holds, while the second term accounts for those values of next period shock for which an exchange rate rule needs to be enforced. The fact that d^{ERR} is a random variable known at time t is left implicit in our notation.

Now we show that the integral in (20) is unbounded. The derivation is similar to the one used in appendix 1 to show that S has an hyperbola

$$\int_{K^-}^{d^{ERR}} \frac{1-\gamma\beta_t}{1-\gamma\beta_{t+1}-d/\phi} \mathrm{d}F_{d_t}(d) \geq (1-\gamma\beta_t) Q \int_0^{\eta} \frac{1}{x} \mathrm{d}x = \infty$$

for some finite constant Q and small η.

This proves that $E_t(P_{t+1}/P_t) = \infty$ for realizations and periods where $P(ERR \text{ at } t+1 \mid d_t, d_{t-1}, ...) > 0$. Therefore, for any realization where $P(ERR \text{ at } t+1 \mid d_t, d_{t-1}, ...) > 0$ at *some* $t \leq T$,

$$\frac{1}{T}\sum_{t=1}^{T}\left[\frac{P_{t+1}}{P_t} - \frac{P_{t+1}}{P_t}^{e}\right]^2 < \frac{1}{T}\sum_{t=1}^{T}\left[\frac{P_{t+1}}{P_t} - E_t\left(\frac{P_{t+1}}{P_t}\right)\right]^2 + \epsilon, \tag{21}$$

because the right hand side is, in fact, infinite. So,

$$\begin{aligned} \pi^{\epsilon,T} &\geq P\left[P(ERR \text{ at } t+1 \mid d_t, d_{t-1}, ...) > 0 \text{ at some } t \leq T\right] \geq \\ &\geq P\left[ERR \text{ at some } t \leq T\right], \end{aligned}$$

where the first inequality follows from (21) and the definition of $\pi^{\epsilon,T}$, and the second inequality follows from elementary properties of probabilities.

The case of β^U finite but arbitrarily large follows from observing that, with arbitrarily high probability, the sequences of the case $\beta^U = \infty$ are below a certain bound $\underline{\beta}$; then, one can choose arbitrarily high β^U to make the conditional expectation arbitrarily close to the one with β^U infinite, while actual and perceived inflations are bounded by $\underline{\beta}$.□

This proposition shows that the probability that learning is better than the conditional expectation is no lower than the probability of having a hyperinflation. Obviously, for low enough values of seignorage, this probability is quite small. However, for high values of seignorage, equilibrium exhibits hyperinflations with high probability. For instance, in the computations we have in Tables 2, 4 and 6, the probability of having at least one hyperinflation is .84, .91 and .97 for average seignorage .045, .047 and .049 respectively. According to the proposition, those are lower bounds for $\pi^{\epsilon,T}$ in that model.

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 these hyperinflations is 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 these controls. 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. The policy recommendations that come out of the model are in agreement with the views we discussed in the previous paragraph: an ERR may temporarily stop a hyperinflation, but average seignorage (and also its standard deviation) must be lowered to eliminate hyperinflations permanently.

The economic fundamentals of the model are perfectly 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 is made precise in the body of the paper; despite abandoning RE, we maintain falsifiability of the model. This deviation from rational expectations is attractive because it avoids the strong requirements on rationality placed by RE, and because the fit of the model improves dramatically despite the *small* deviation from rationality.

On the practical side, this paper shows that hyperinflations 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

In this appendix we characterize the set of stationary rational expectations equilibria 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

$$P^e_{t+1} = \beta P_t, \tag{22}$$

where β is a constant. Then, β is a stationary rational expectations equilibrium iff

$$E_t(P_{t+1} \mid I_t) = \beta P_t.$$

Let us make some assumptions on the model:

Assumption 1 $\gamma, \phi > 0$;

Assumption 2 *d_t is i.i.d. with finite support* $[K^-, K^+] \subset R_+$

Define the distribution $F_{d_t}(d) \equiv P(d_t \leq d)$

Assumption 3 $\liminf_{d \nearrow K^+} \frac{F_{d_t}(d)}{K^+ - d} \equiv \Pi > 0$

These are very weak assumptions; the third assumption is satisfied, for example, if d_t has a point mass at K^+, or if d_t has a positive density at K^+.

If expectations follow (22), then equation (17) implies that

$$P_{t+1} = h(\beta, d_{t+1}) P_t, \tag{23}$$

where h is as defined in section 5.2.

Now, letting

$$S(\beta) \equiv E\left(h(\beta, d_{t+1}) \mid I_t\right),$$

S is interpreted as the mapping from perceived to actual expectations. The set of stationary rational expectations equilibria coincides with the fixed points on the mapping $S : R_+ \to R_+$. Notice that $S(\beta)$ is a constant because d_t is i.i.d.

In the next proposition we characterize the properties of S. These properties are displayed in the graph at the end of this appendix.

Proposition 3 *$S : R_+ \to R_+$ has the following properties:*

1. *In the set $[0,(1 - K^+/\phi)/\gamma)$, the mapping S is increasing and convex. If $\beta^U = \infty$, then S has an asymptote at $\beta = (1 - K^+/\phi)/\gamma$.*

2. *S has at most two fixed points denoted $\beta^1_{RE} < \beta^2_{RE}$. For $\overline{d} \equiv E(d_t)$ and σ_d low enough, and for β^U large enough two equilibria exist. For $\overline{d}$ large enough no equilibrium exists.*

3. *When a fixed point exists, $S'(\beta^1_{RE}) < 1$*

4. *Let $\widetilde{\beta}^1_{RE} < \widetilde{\beta}^2_{RE}$ be the rational expectations equilibria without uncertainty (when $\sigma_d = 0$). Assume that two fixed points of S exist.*

   Then $\widetilde{\beta}^1_{RE} < \beta^1_{RE} < \beta^2_{RE} < \widetilde{\beta}^2_{RE}$.

Proof

1. Using the definition of S we have

$$S'(\beta) = E\left(\frac{\partial\ h(\beta, d_t)}{\partial\ \beta}\right) = E\left(\frac{\gamma d_t/\phi}{(1 - \gamma\beta - d_t/\phi)^2}\right) \qquad \text{and}$$

$$S''(\beta) = E\left(2\frac{\gamma^2 d_t/\phi}{(1 - \gamma\beta - d_t/\phi)^3}\right)\ ;$$

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

To prove the existence of an asymptote; note that

$$S'((1-K^+/\phi)/\gamma) = \int_0^{K^+} \frac{\gamma d/\phi}{K^+ - d}\,\mathrm{d}\ F_{d_t}(d) > \gamma K^+ \int_{K^+-\eta}^{K^+} \frac{1}{K^+ - d}\,\mathrm{d}\ F_{d_t}(d)\,. \tag{24}$$

for any η. According to assumption 2, we can choose $\overline{\eta}$ small enough such that, if $d > K^+ - \overline{\eta}$, then $F_{d_t}(d) \geq \Pi(K^+ - d) > \Pi\overline{\eta} > 0$; this implies the inequality in

$$S'((1-K^+/\phi)/\gamma) \geq \Pi\gamma K^+ \int_{K^+-\overline{\eta}}^{K^+} \frac{1}{K^+ - d}\,\mathrm{d}\,d\ = \Pi\gamma K^+ \int_0^{\overline{\eta}} \frac{1}{x}\,\mathrm{d}\,x = \infty \qquad ,$$

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 $S'((1 - \gamma K^+)/\gamma) = \infty$.

2. That we have at most two fixed points follows immediately from convexity of S. Lowering the mean and the variance of d_t we have that all the possible values of this random variable are arbitrarily close to zero; since, for any given β, $h(\beta, d) \rightarrow 1$ as $d \rightarrow 0$, we have that $S(\beta)$ becomes arbitrarily close to 1 and $S'(\beta)$ arbitrarily close to zero. This means that there is a fixed point close to 1. The fact that S has an asymptote (part 1 od this proposition) implies that there is a second fixed point if β^U is large enough. Since S is increasing in d_t, if $\bar{d}$ is large enough no equilibrium exists.

3. Clearly, $S(0) > 0$. Therefore, at β^1_{RE}, S has to cut the 45° line from above, and $S'(\beta^1_{RE}) < 1$.

4. Notice that $\tilde{\beta}^1_{RE}$ and $\tilde{\beta}^2_{RE}$ are the fixed points of $h(\cdot, \bar{d})$. Since h is a convex function of d_t, Young's inequality, implies that $S(\beta) > h(\beta, \bar{d})$. See the graph at the end of this appendix. This implies part 4.□

Now, we argue that the least squares learning mechanism converges to the lower rational expectations equilibrium. This is a routine application of the frameword of Marcet and Sargent (1989a), so the details are omitted. The associated differential equation is given by

$$\dot{\beta} = S(\beta) - \beta \tag{25}$$

and we know that, under least squares learning (the case that $\alpha_t = t$), the sytem 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 on non-linear difference equations in Ljung (1975).

Now, $S'(\beta^1_{RE}) < 1$ implies that (25) is locally stable at β^1_{RE}; the basin of attraction of β^1_{RE} is the set $[0, \beta^2_{RE})$. In the proof to proposition 1 we have shown that β_t visits the stable set infinitely often, so that least squares learning converges to the rational expectations equilibrium β^1_{RE} a.s.

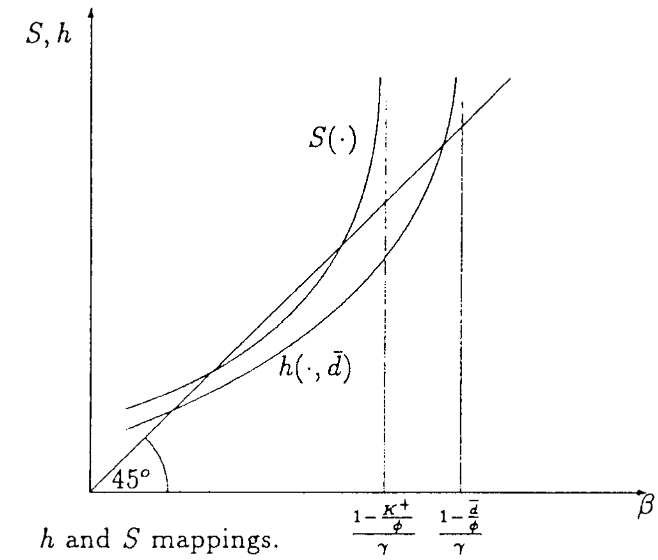

h and S mappings.

APPENDIX 2

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 analytics 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[34] 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 consistent with the findings in Fernandez and Mantel (1989), Kieguel and Newmayer (1992) and Rodríguez (1991). The parameters of the money demand γ and ϕ, are uniquely determined by the two numbers above. Note that the money demand function (1) implies a stationary Laffer curve equal to

$$\frac{\pi}{1+\pi}m = \frac{\pi}{1+\pi}\frac{1}{\gamma}\left(1 - \frac{1}{\phi}(1+\pi)\right) \tag{26}$$

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

$$\pi* = \sqrt{\phi} - 1$$

which, setting $\pi* = 60\%$, implies $\phi = 2.56$. Using this figure in (26), and making the maximum revenue equal to 0.05, implies $\gamma = 2.7$.

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

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] Bental, B. and B. Eckstein. "A Neoclassical Interpretation of Inflation and Stabilization in Israel", Working Paper 28-95, The Foerder Institu-teof Economic Research, Tel-Aviv University.

[3] Bray, M. M. "Learning, estimation and the stability of rational expectations", *Journal of Economic Theory*, 26, pp. 318-339, (1982).

[4] Bray, M. M. and E. Savin "Rational Expectations Equilibria, Learning, and Model Specification", *Econometrica*, 54, pp. 1129-60 (1986).

[5] Bruno, M., G. Di Tella, R. Dornbusch and S. Fisher. *Inflation Stabilization. The experience of Israel, Argentina, Brasil and Mexico* MIT Press, (1988)

[6] Bruno, M., S. Fisher, E. Helpman and N. Liviatan. *Lessons of Economic Stabilization and its Aftermath*, MIT Press, (1991).

[7] Chung, Heetaik. "Did Policy Makers Really Believe in the Phillips Curve?. An Econometric Test". Ph. D. dissertation, University of Minnesota, (1990).

[8] Easley, D. and A. Rustichini. "Choice without belief", working paper, CORE, Universite Catholique de Lovaine, (1995).

[9] Eckstein, Z. and L. Leiderman, "Seignorage and the Welfare Cost of Inflation", *Journal of Monetary Economics*, no. 29, pp. 389-410, (1992).

[10] Evans, G.W. and S. Honkapohja. "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models", *Econometrica*, 63, pp.195-296, (1995)

[11] Evans, G.W. and S. Honkapohja. "Adaptive Forecasts, Hysteresis, and Endogenous Fluctuations", *Economic Review*, 1, Federal Reserve Bank of San Francisco, (1993).

[12] Fernandez, R. and R. Mantel, "Fiscal lags and the problem for stabilization: Argentina's Austral plan" in *Latin-American debt and adjustment*, P. Brock, M. Connolly and C. Gonzalez-Vega, editors, New York: Praeger, (1989).

[13] Fudenberg D. and D. Levine. "Universal consistency and cautious fictitious play", *Journal of Economic Dynamics and Control*, 19, pp. 1065-89 (1995).

[14] Funke, M., S. Hall and M Sola. " Rational bubbles during Poland's hyperinflation: implications and empirical evidence", *European Economic Review*, 38, pp 1257-76, (1994).

[15] Hommes, C. and G. Sorger. "Consistent Equilibria", working paper, University of Amsterdam (1997).

[16] Imrohoroglu, S. "Testing for Sunspot Equilibria in the German Hyperinflation". *Journal of Economic Dynamics and Control*, 17, (1993).

[17] Kieguel M. and A. Newmayer, "Seignorage and inflation, the case of Argentina". Working Paper, The World Bank, (1992).

[18] Kurz, M. "On rational equilibria", *Economic Theory*, 4, pp. 859-876, (1994).

[19] Ljung, L. "Theorems for the Asymptotic Analysis of Recursive Stochastic Algorithms" Working paper, Department of Automatic Control, Lund Institute of Technology (1975).

[20] 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).

[21] Marcet A. and T.J. Sargent. "Least Squares Learning and the Dynamics of Hyperinflation", in *International Symposia in Economic Theory and Econometrics*, vol. 4, edited by W. Barnett, J.Geweke and K. Shell, Cambridge University Press (1989b).

[22] 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).

[23] Nicolini, J. "Ruling out Speculative Hyperinflations: the Role of the Government", *Journal of Economic Dynamics and Control*, 20, pp 791-809, (1996).

[24] Obstfeld, M and K. Rogoff. "Speculative Hyperinflations in Maximizing Models: can we Rule them Out?", *Journal of Political Economy*, 91, pp. 675-687,(1983).

[25] Rodríguez, C. "A stylized model of the inflation-devaluation spiral", IMF Staff Papers, (1980).

[26] Rodríguez, C. "The macroeconomics of public sector deficits, the case of Argentina"PRE Working Paper Series, # 632, Washington DC: The World Bank.

[27] Sargent, T.J. "The End of Four Big Inflations", in T.J. Sargent, *Rational Expectations and Inflation*, New York, Harper and Row, (1986).

[28] Sargent, T.J. *Bounded Rationality in Macroeconomics*, Oxford University Press, (1993).

[29] Sargent, T. 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).

[30] Woodford, M. "Learning to Believe in Sunspots", *Econometrica*, no. 58, 277-307, (1990).

[31] Zarazaga, C. "Hyperinflations and moral hazard in the appropriation of seignorage", Working Paper 93-26, Federal Reserve Bank of Philadelphia, (1993).

FIGURE I

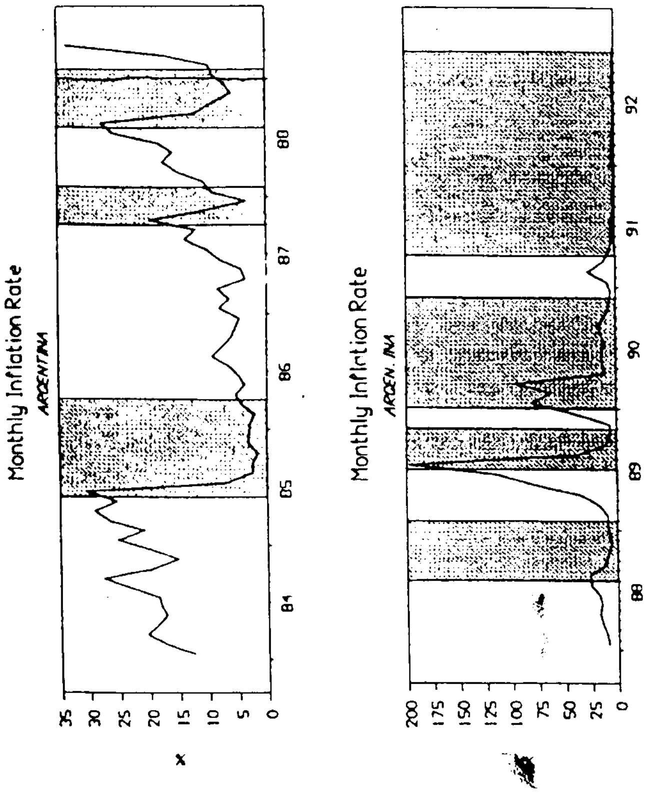

FIGURE II

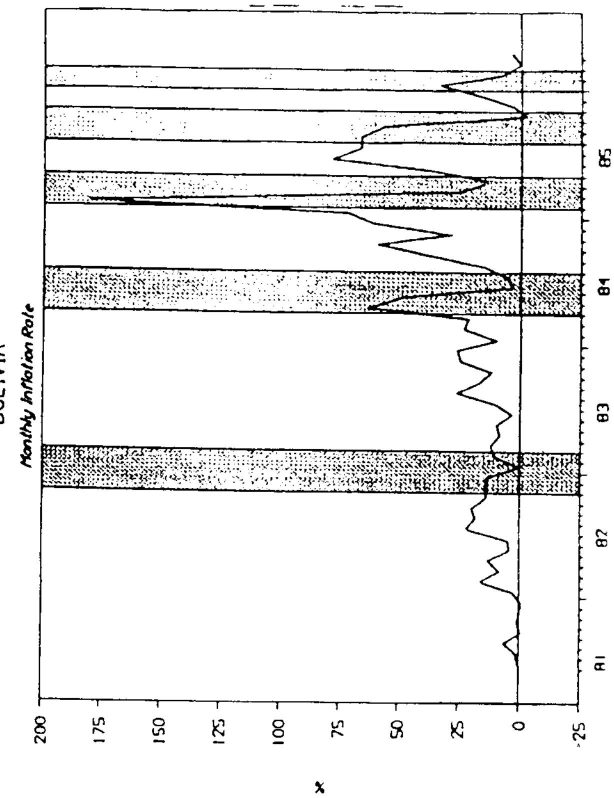

FIGURE III

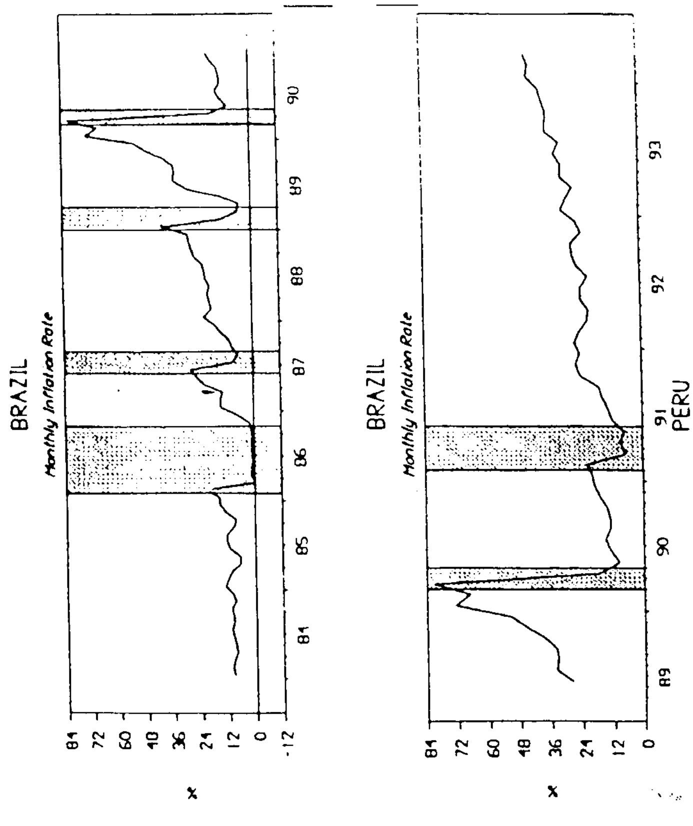

FIGURE IV

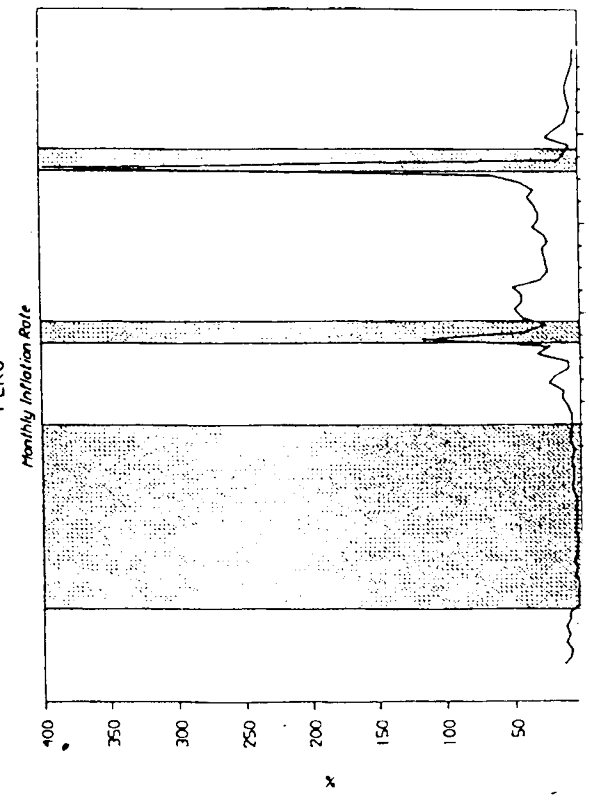

FIGURE 5

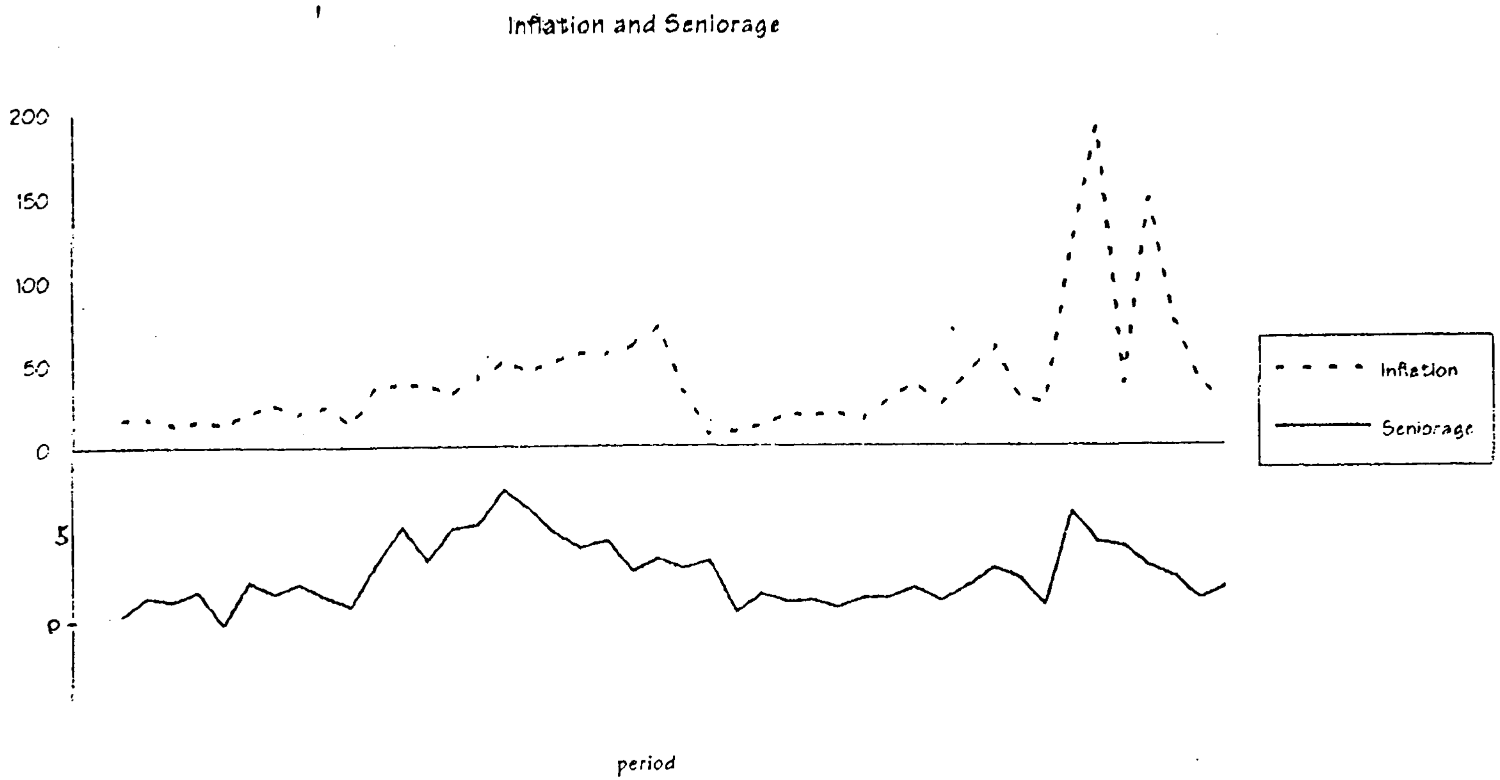
Inflation and Seniorage
200
150
100
50
0
5
0
Inflation
Seniorage
period

Figure 6

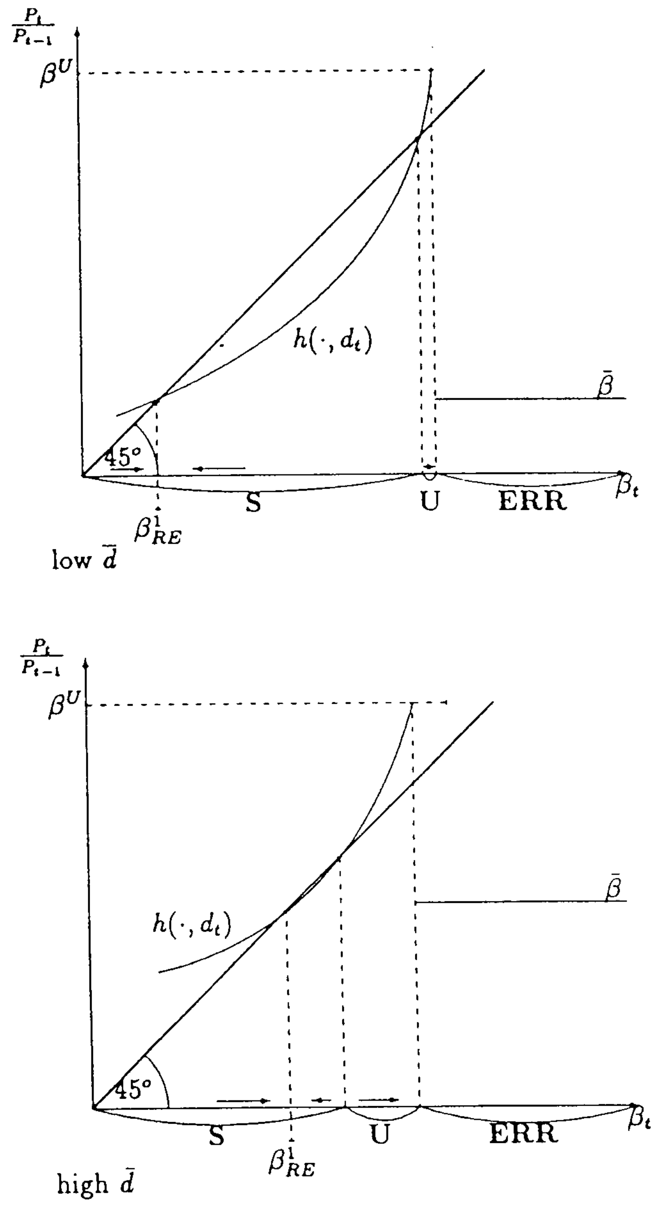

$\frac{P_t}{P_{t-1}}$
β^U
$h(\cdot, d_t)$
$\bar{\beta}$
45°
S
U
ERR
β_t
β^1_{RE}
low $\bar{d}$
$\frac{P_t}{P_{t-1}}$
β^U
$h(\cdot, d_t)$
$\bar{\beta}$
45°
S
U
ERR
β_t
β^1_{RE}
high $\bar{d}$

FIGURE 7

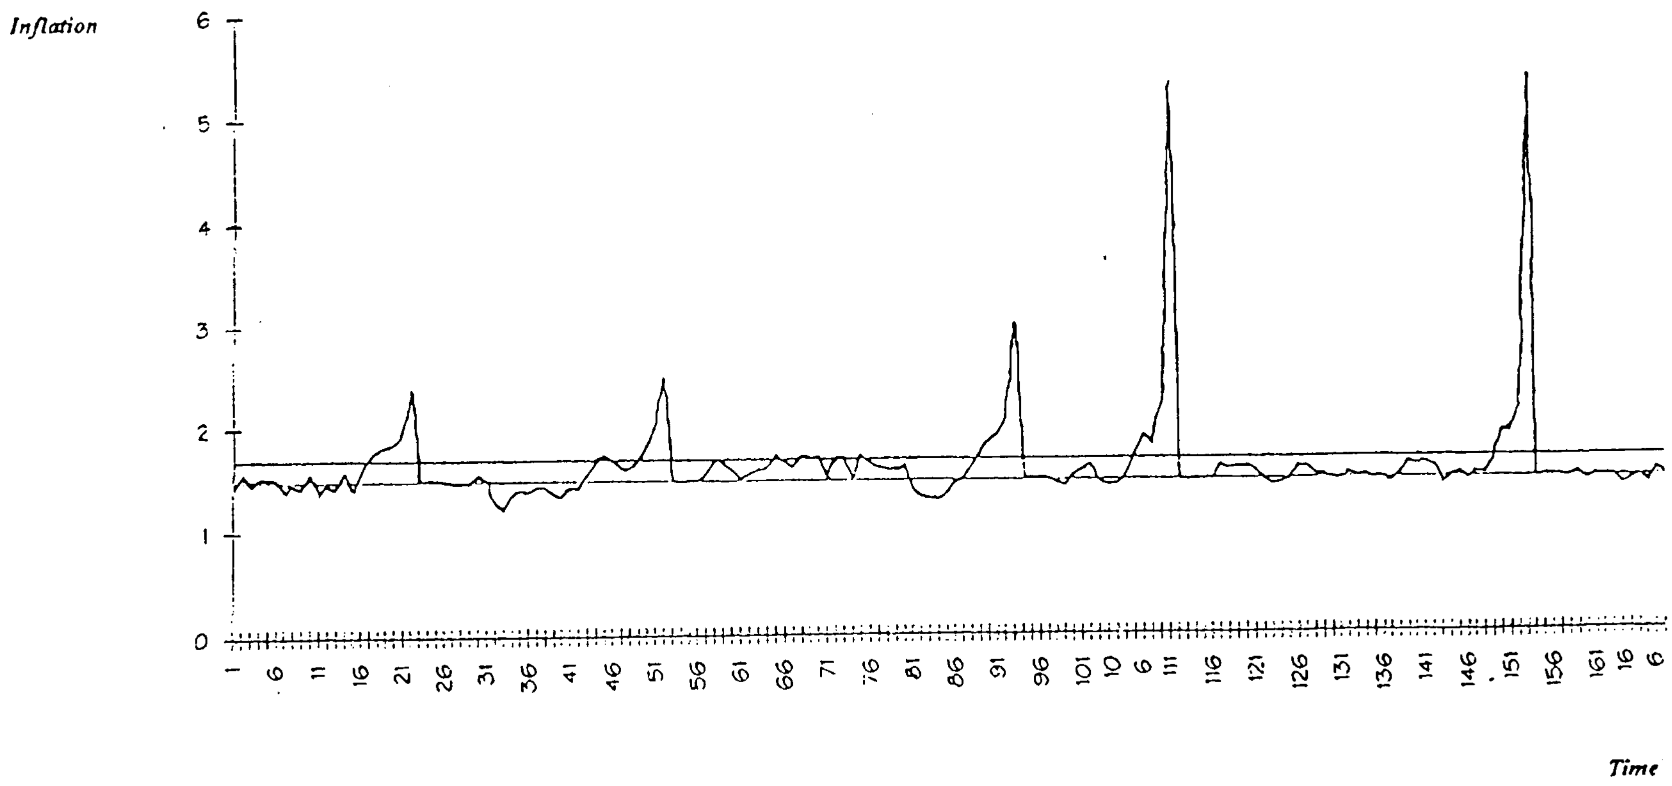

Inflation
6
5
4
3
2
1
0
1
6
11
16
21
26
31
36
41
46
51
56
61
66
71
76
81
86
91
96
101
10
6
111
116
121
126
131
136
141
146
151
156
161
16
6
Time

Efficient Values of Alpha

FIGURE 8

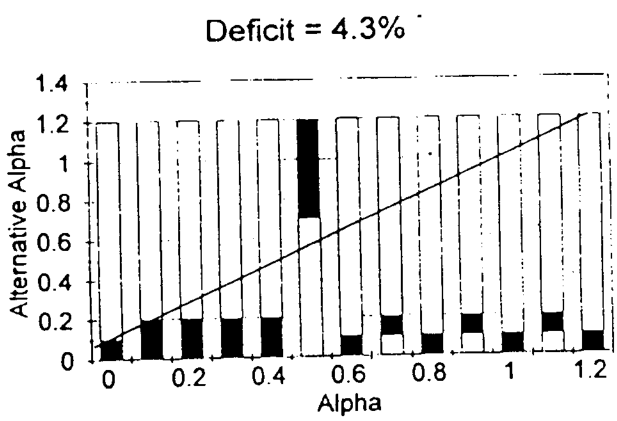

FIGURE 9

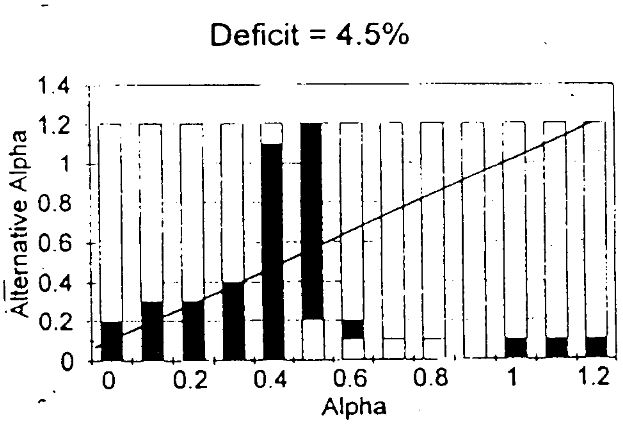

FIGURE 10

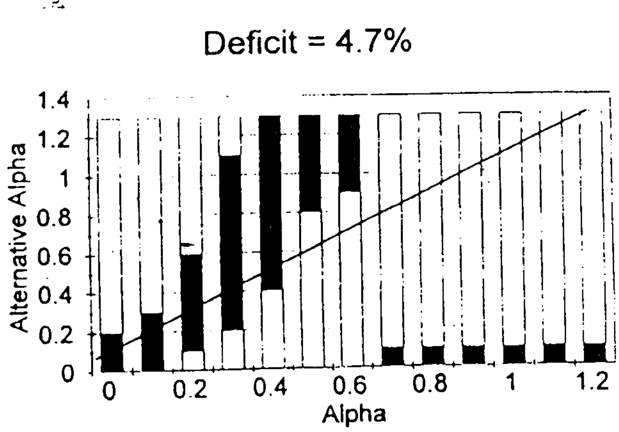

FIGURE 11

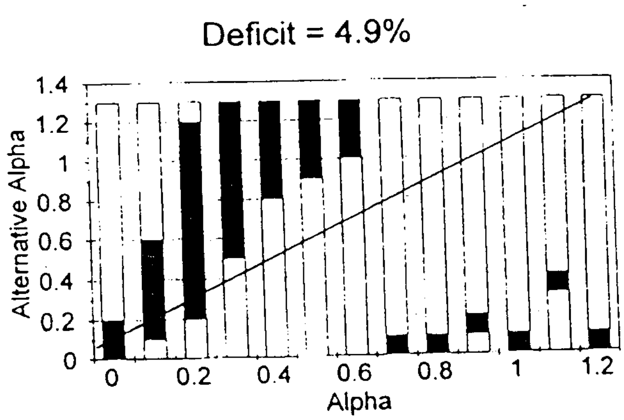

TABLE 1
Deficit = 4.5%

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