

Recurrent Hyperinflations and Learning

By ALBERT MARCET AND JUAN P. NICOLINI*

We use a model of boundedly rational learning to account for the observations of recurrent hyperinflations in the 1980's. In a standard monetary model we replace the assumption of full rational expectations by a formal definition of quasi-rational learning. The model under learning matches some crucial stylized facts observed during the recurrent hyperinflations experienced by several countries in the 1980's remarkably well. We argue that, despite being a small departure from rational expectations, quasi-rational learning does not preclude falsifiability of the model, it does not violate reasonable rationality requirements, and it can be used for policy evaluation. (JEL D83, E17, E31)

The goal of this paper is to develop a model that accounts for the main features of the hyperinflations of the 1980's and to study the policy recommendations that arise from it. The model is standard, except for the assumption of quasi-rational learning. Modern macroeconomics has been reluctant to use boundedly rational expectations models to match empirical observations. It is commonly believed that such models are not falsifiable and expectations are not consistent with the model. This view is stated clearly in the following quotation from Thomas J. Sargent (1993): "... the literature on adaptive decision processes seems to me to fall far short of providing a secure foundation for a good theory of real-time transition dynamics. There are problems of arbitrariness and the need for

prompting, with a concomitant sensitivity of outcomes to details of adaptive algorithms." A side contribution of the paper is to show with an example that, contrary to Sargent's statement, if certain rationality requirements are imposed, learning models can be useful to understand real-time transition dynamics.

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

However, this simple story fails when we closely look at time series of inflation and sei-

* Marcet: Department of Economics and Business, Universitat Pompeu Fabra, C/Ramon Trias Fargas, 23-25, 08005, Barcelona, Spain, CREI, and CEPR (e-mail: albert.marcet@upf.edu); Nicolini: Universidad Torcuato Di Tella, C/Miñones 2177, C1428ATG Buenos Aires, Argentina (e-mail: juanpa@utdt.edu). We thank Tony Braun, Jim Bullard, George Evans, Seppo Honkapohja, Rodi Manuelli, Ramon Marimon, Tom Sargent, Stacia Sowerby, Harald Uhlig, Neil Wallace, and Carlos Zarazaga for helpful conversations and Marcelo Delajara and Ignacio Ponce Ocampo for research assistance. Any errors are our own. Part of this work was done when both authors were visiting the Federal Reserve Bank of Minneapolis. Most of the work by Marcet was done when visiting CEMFI, Madrid. Research support from DGICYT (Spanish Ministry of Science and Education), CIRIT (Conselleria d'Ensenyament Generalitat), CREI, the Barcelona Economics Program of CREA, ANCT (Argentina, PICT 98-02-03543), and HCM (European Union) is greatly appreciated. Many of the proofs can be found in Appendix C, available from the authors or in the working paper version of this article.

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

In this paper we develop a model that accounts for this and other crucial observations that occurred during the hyperinflations of the 1980's. These episodes involve very high inflation rates (for instance, inflation in Argentina in June 1989 peaked at 200 percent a month) and all we know about the welfare effects of inflation suggest that they are very costly.

Sargent and Neil Wallace (1987) explained these hyperinflations as bubble equilibria. Their model generates a standard Laffer curve with two stationary rational expectations equilibria; hyperinflations could occur as speculative equilibria converging 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. Our work builds upon Sargent and Wallace's 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.

The last decade has witnessed a renewed interest in learning models in macroeconomics, mostly focusing on issues of convergence to rational expectations.¹ This literature has made enormous progress, and convergence of learning models to rational expectations can now be studied in very general setups. But few attempts have been made to explain observed economic

facts with models of boundedly rational learning. Some exceptions are Evans and Honkapohja (1993), Allan Timmermann (1993, 1996), and Jasmina Arifovic et al. (1997). Some papers looking at policy implications are Hee-Taik Chung (1990) and Evans et al. (2001). However, with the partial exception of Evans and Honkapohja (see our discussion following Definition 3), none of them formally addressed the critique to boundedly rational models that is commonplace in today's macro literature and that is clearly stated in the above quote from Sargent. This critique says that using models of boundedly rational learning would entail problems similar to those found in models of adaptive expectations of the prerational-expectations era, namely: (i) too many degrees of freedom are available to the economist, so the model is not falsifiable; (ii) agents' expectations are inconsistent with the model, so rational agents would eventually abandon their ad hoc expectations; and (iii) the model does not predict how expectation formation will change if there is a change in policy.

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

Quite a few papers have presented models that explain some of the facts we consider, among others, Zvi Eckstein and Leonardo Leiderman (1992) and Benjamin Bental and Eckstein (1997) explain the very large inflation rates in Israel with an ever-increasing Laffer curve, and Carlos Zarazaga (1993) develops a model of endogenous seignorage. These papers 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 I presents the stylized facts and provides supporting evidence. Section II presents the model and characterizes rational expectations equilibria. Section III discusses the lower bounds in

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

² Recent literature imposing consistency requirements in learning models are Evans and Honkapohja (1993); Mordecai Kurz (1994); Drew Fudenberg and David K. Levine (1995); and Cars Hommes and Gerhard Sorger (1998).

rationality in a general setup. Section IV discusses the behavior of the model under the lower bounds on rationality. The paper ends with some concluding remarks.

I. Evidence on Recurrent Hyperinflations

A number of countries, including Argentina, Bolivia, Brazil, and Peru experienced during the 1980's the highest average inflation rates of their history. 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:

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. But often an ERR only lowers inflation temporarily, and new hyperinflations eventually occur.
3. During a hyperinflation, seignorage and inflation are not highly correlated.
4. Average inflation and seignorage are strongly correlated across countries. Hyperinflations only occur in countries where seignorage is high on average.

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

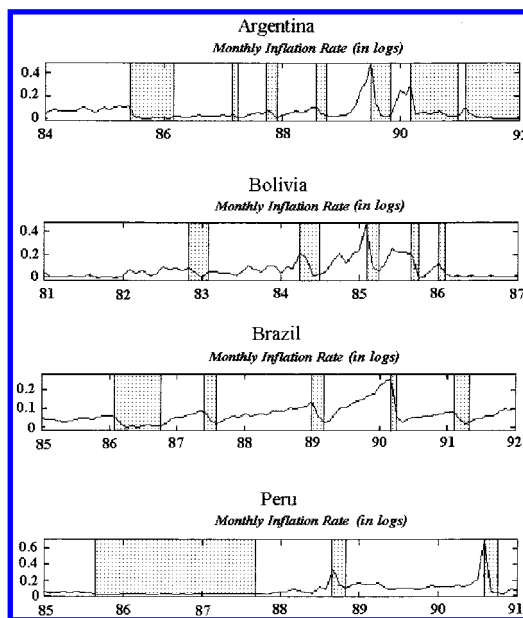


FIGURE 1. MONTHLY INFLATION RATE (IN LOGS)

Our summary of stylized facts should be uncontroversial.³ Facts 1 and 2 are clearly shown in Figure 1, which presents data on the inflationary experiences of Argentina, Bolivia, Brazil, and Peru in the 1980's.⁴ Periods when an explicit fixed ERR was in place are indicated by shaded areas. The end of the shading indicates the date in which the ERR was explicitly abandoned.

Fact 3 has been well documented in the literature and it has motivated quite a bit of research including, for example, that of Sargent and Wallace (1987). Hyperinflations did not happen together with high peaks in seignorage; very different levels of seignorage were present during different hyperinflations in a given country and, in some instances, seignorage even decreased while a hyperinflation was taking place.⁵

³ See Michael Bruno et al. (1988) and (1991).

⁴ Inflation rates were computed from International Financial Statistics consumer price indices.

⁵ See a plot of inflation and seignorage for Argentina in the 1980's in the working paper version of this paper (Marcet and Nicolini, 1998).

II. The Model

A. Money Demand and Money Supply

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

The demand for real balances is given by

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

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

Money supply is driven by the need to finance seignorage. On the other hand, government’s concern about current levels of inflation prompts the adoption of ERR when inflation gets out of hand or to restore equilibrium.

If no ERR is in place at t , the government budget constraint is given by

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

Seignorage is given by an exogenous independently and identically distributed (i.i.d.) stochastic process $\{d_t\}_{t=0}^\infty$ and it is the only source of uncertainty in the model.⁷ Equations (1) and (2) plus a hypothesis of expectations formation, determine the equilibrium values for $\{M_t, P_t\}_{t=0}^\infty$ in periods of floating exchange rates.

If an ERR is in place at t , the government

⁶ Appendix A shows that the following equations can be rationalized as the equilibrium conditions of an overlapping generations (OLG) monetary model of a small open economy.

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

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. Assuming full mobility of goods, purchasing power parity implies

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

and the targeted inflation rate is achieved. In the case that targeted inflation $\bar{\beta}$ 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). Given this price level and an expectations hypothesis, (1) determines money demand. In general, this money demand will not match money supply as determined by (2). As it is standard in fixed exchange rate models, international reserves (R_t) adjust so the right level of money balances is achieved. Thus, instead of (2) money supply is now given by:

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

Finally, we impose the rule that government acts to satisfy

$$(5) \quad \frac{P_t}{P_{t-1}} \leq \beta^U,$$

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 positive price level clears the market if $R_t = R_{t-1}$.

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 nonnegativity constraint on foreign reserves is never binding. Since we will choose the target inflation rate $\bar{\beta}$ to be the lower stationary rational expectations equilibrium steady-state inflation, the loss

of reserves is small in our simulations and it is likely to be small for most parameter values. Modeling reserve accumulation formally is unlikely to change our main results, but it opens up a host of interesting issues. For example, the government may run out of reserves during a hyperinflation, so that “orthodox” measures cannot be avoided, a feature that is consistent with our model. Alternatively, by increasing the length of the ERR after a hyperinflation the monetary authority could accumulate reserves since the real value of the money stock is increased after the stabilization.⁸

We have modeled policy in this way because it mimics the broad features of policies followed by South-American countries during the 1980’s. The issue of why these countries followed this kind of policy is not addressed formally in this paper, but we can advance three possible justifications for using this rule within our model. First, the fact that ERR has been established only after some periods of high inflation is justified because then the value of foreign reserves is high, and a large part of the domestic money can be backed with existing reserves.⁹ Second, in principle, any reduction in the government deficit of $e_t(R_t - R_{t-1})$ units would also fix the inflation to $\bar{\beta}$ in periods of ERR. In fact, the reduction in seignorage that is needed to achieve an inflation equal to $\bar{\beta}$ 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 the exact value of $e_t(R_t - R_{t-1})$, can only be inferred from knowledge of the true model and all the parameter values, including those that determine the (boundedly rational) expectations P_{t+1}^e , and all the shocks. By contrast, an ERR can be implemented only with knowledge of the foreign price level and the policy parameters ($\bar{\beta}$, β^U). A third advantage of establishing ERR for

real governments would be the existence of institutions that can implement this measure quickly, while lowering government expenditures or increasing taxes often takes a long time.

An important policy decision is how long to maintain the ERR. Obviously, the longer the ERR is maintained, the closer expected inflation will be to $\bar{\beta}$. In our simulations, we hold the ERR until expected inflation is close to $\bar{\beta}$ 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 $\bar{\beta}$, β^U , and the distribution of d_t .

B. Rational Expectations

If we assume that agents form expectations rationally, the model is very similar to that of Sargent and Wallace (1987) (henceforth, SW). As long as seignorage is not too high, the model has two stationary equilibria with constant expected inflation levels (called low- and high-inflation equilibria), and a continuum of bubble equilibria that converge to the high-inflation equilibrium.¹⁰

The main motivation behind the work of SW was to explain fact 3 in Section I as rational bubble equilibria.¹¹ Their original model does not allow for recurrence of hyperinflations (fact 1), but the work by Funke et al. (1994) shows that recurrence can be explained by introducing a sunspot that turns rational bubbles on and off. Even if one accepts rational sunspots as an explanation, fact 1 is not matched quantitatively: for reasonable parameter values, the magnitude of the hyperinflations that can be generated with this model is very small.¹² Fact 4 is contradicted: the long-run inflation rate in any rational bubble equilibrium is *lower* when seignorage is *higher*, so the model under ratio-

⁸ For instance, Central Bank reserves grew, in Argentina, from 1991 (year in which the Convertibility Plan was launched) to 1994 from 500 million dollars to more than 12 billion.

⁹ This interpretation would suggest that the burst in inflation at the beginning of 1991 in Argentina was crucial for the success of the Convertibility Plan launched in April of the same year, because it substantially reduced the value of the money stock to a point where, at a one dollar = one peso exchange rate, the government could back the whole money stock.

¹⁰ We reproduce some of these results for our model in Appendix 2, available from the authors or in the working paper version of the paper.

¹¹ There has been some work testing the existence of rational bubbles in the German hyperinflation of the 1920’s. A summary of the literature and a test of bubble versus stationary equilibria in the SW model can be found in Selahattin İmrohoroğlu (1993).

¹² This is documented in our discussion of Figure 3 in Section IV, subsection E.

nal expectations (RE) predicts that hyperinflation is *less* severe in countries with high seignorage.

The papers of Maurice Obstfeld and Kenneth Rogoff (1983) and Nicolini (1996) maintain rational expectations and introduce ERR that goes into effect if inflation goes beyond a certain level and, therefore, these papers can be used to address fact 2. Their results show that just the threat of convertibility eliminates bubble equilibria altogether and that the ERR, under rational expectations equilibria, never takes place. Thus, once ERR is a credible threat the rational expectations equilibrium is inconsistent with the existence of hyperinflation. But this was certainly a credible threat in these countries, since ERR did take place.

Marcet and Sargent (1989) studied stability of rational expectations equilibria in the SW model under least-squares learning. They found that the low-inflation equilibrium is *locally* stable and the high-inflation equilibrium is always unstable. Taken literally, these results would say that bubble equilibria cannot be learned by agents. Therefore, none of the above facts is appropriately matched if we restrict our attention to rational expectations equilibria that are stable under learning.¹³

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

III. Learning and Lower Bounds on Rationality

Before the “rational expectations (RE) revolution,” economic agents’ expectations were specified in macroeconomics according to ad hoc assumptions; one popular alternative was “adaptive expectations.” We explained in the introduction that this was criticized because it leads to: (i) too many degrees of freedom, (ii) irrational expectations, and (iii) expectations that are exogenous to the model. The first crit-

icism 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.” RE became the way to overcome these criticisms.

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

The term *boundedly rational learning* (which, in this paper, we use as synonymous with the term *learning*) is used to denote learning mechanisms that place *upper* bounds on rationality. For example, agents are assumed not to know the exact economic model or to have bounded memory. But this admits too many models of learning. Indeed, once we rule out RE, anything can be a boundedly rational learning scheme and we could be falling back into old mistakes and the “wilderness of irrationality.”¹⁴

Our approach is to allow for only *small* deviations from rationality, both along the transition and asymptotically. Given an economic model we only admit learning mechanisms that satisfy certain *lower* bounds on rationality within this model. In Section IV we will show how this small departure from rationality generates equilibria in the model of Section II that are quite different from RE, precisely in the direction of improving the match of empirical observations.¹⁵

We now set up a general framework and define lower bounds that we place on rationality. Assume that an economic model satisfies

$$(6) \quad \mathbf{x}_t = g(\mathbf{x}_{t-1}, \mathbf{x}_{t+1}^e, \boldsymbol{\xi}_t, \boldsymbol{\eta})$$

where g is a function determined by market equilibrium and agents’ behavior, \mathbf{x}_t contains all the variables in the economy, \mathbf{x}_{t+1}^e is agents’ expectation of the future value of \mathbf{x} , $\boldsymbol{\xi}_t$ is an exogenous shock, and $\boldsymbol{\eta}$ is a vector of parameters,

¹³ Marcet and Sargent (1989) 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} . These authors noted that if inflation goes beyond the high steady state it may enter an unstable region where inflation tends to grow without bound. This feature of the model with learning constitutes the core of the dynamics in the current paper.

¹⁴ It might seem that Bayesian learning is a way out of this dilemma, but the literature has described several paradoxes and shortcomings of this approach. See, for example, Margaret M. Bray and David M. Kreps (1987), Marimon (1997), and David Easley and Aldo Rustichini (1999).

¹⁵ Marimon (1997) and Easley and Rustichini (1999) also argue that learning can be used for more than a stability criterion.

including the parameters of government policy and the parameters that govern the distribution of ξ_t . For example, in our model, \mathbf{x}_t is inflation and real balances, 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 $\boldsymbol{\eta}$ includes $\gamma, \phi, \beta, \beta^U$, and the parameters of the distribution of seignorage.

Assume that agents' expectations are given by

$$(7) \quad \mathbf{x}_{t+1}^e = z(\boldsymbol{\beta}_t(\boldsymbol{\mu}), \mathbf{x}_t)$$

where $\boldsymbol{\beta}_t(\boldsymbol{\mu})$ is a vector of statistics inferred from past data and z is the forecast function. The statistics $\boldsymbol{\beta}$ are generated by a learning mechanism f and learning parameters $\boldsymbol{\mu}$ according to

$$(8) \quad \boldsymbol{\beta}_t(\boldsymbol{\mu}) = f(\boldsymbol{\beta}_{t-1}(\boldsymbol{\mu}), \mathbf{x}_t, \boldsymbol{\mu}).$$

The learning mechanism f dictates how new information on \mathbf{x}_t is incorporated into the statistics $\boldsymbol{\beta}$. The learning parameters $\boldsymbol{\mu}$ govern, for example, the weight that is given to recent information. For now, $(z, f, \boldsymbol{\mu})$ are unrelated to the true model $(g, \boldsymbol{\eta})$, but later in this section we will define bounds on rationality that amount to imposing restrictions on the space of $(z, f, \boldsymbol{\mu})$ given a model $(g, \boldsymbol{\eta})$.

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

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

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

Equations (6), (7), and (8) determine the equilibrium sequence for given learning parameters $\boldsymbol{\mu}$. Obviously, the process for \mathbf{x}_t depends on the parameters $\boldsymbol{\mu}$. This dependence will be left implicit in most of the paper, and we will write \mathbf{x}_t^μ only 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. Formally:

$$(10) \quad \pi^{\epsilon, T} \equiv \mathcal{P}\left(\frac{1}{T} \sum_{t=1}^T |\mathbf{x}_{t+1} - \mathbf{x}_{t+1}^e|^2 < \frac{1}{T} \sum_{t=1}^T |\mathbf{x}_{t+1} - E_t^\mu(\mathbf{x}_{t+1})|^2 + \epsilon\right)$$

where E_t^μ is the true conditional expectation under the learning model and \mathcal{P} is probability.

The first lower bound on rationality we propose is:

Definition 1 Asymptotic Rationality (AR): The expectations given by $(z, f, \boldsymbol{\mu})$ satisfy AR in the model $(g, \boldsymbol{\eta})$ if, for all $\epsilon > 0$,

$$\pi^{\epsilon, T} \rightarrow 1 \text{ as } T \rightarrow \infty.$$

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

AR can be viewed as a minimal requirement in the sense that it only rules out behavior that is inconsistent forever. It rules out, for example, learning mechanisms where a relevant state variable is excluded from the forecasting rule z (this feature would exclude adaptive expectations, for example, if d_t were serially correlated). It is satisfied by least-squares learning mechanisms in models where this mechanism converges to RE and certain continuity assumptions are satisfied.¹⁶ Similar concepts of consistency can be found in the literature.¹⁷

However, AR admits learning mechanisms that generate very bad forecasts along the transition for very long periods. For example, ordinary least squares (OLS) in a model with recurrent hyperinflations would generate very bad forecasts every time a hyperinflation starts, because least-squares learning gives less and less importance to recent events as time goes by, so it would take longer and longer for agents to realize that a hyperinflation is starting. Even worse, under OLS the agents' expectations would adjust more slowly for each subsequent outburst in inflation.

¹⁶ Perhaps surprisingly, AR excludes many "rational equilibria" in the terminology of Kurz (1994), which allows for agents to make systematic mistakes forever, as long as these mistakes are not contemplated in the prior distribution.

¹⁷ This requirement was implicitly imposed in the literature on stability of RE under learning, where the use of least squares was often justified because of its optimality 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.

To restrict learning mechanisms so that they generate good forecasts along the transition we impose the next two lower bounds.

Definition 2 Epsilon-Delta Rationality (EDR): The expectations given by (z, f, μ) satisfy *EDR* for (ϵ, δ, T) in the model (g, η) if:

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

If *EDR* is satisfied for small $\epsilon, \delta > 0$, agents are unlikely to switch to another learning scheme after period T , even if they were told “the whole truth.”¹⁸

It is only interesting to study *EDR* for “moderately high” values of T . If T is too low the sample means have no chance to settle down. If T is large enough and *AR* is satisfied, *EDR* is also satisfied. The precise empirical application that the researcher has in mind should suggest an interesting value for T . For example, in our application below, we choose $T = 10$ years, which is the length of the hyperinflationary period in many of the countries studied.

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

The next bound on rationality requires agents to use learning parameters 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 . Formally,

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

Definition 3 Internal Consistency (IC): Given (g, η) , the expectations given by (z, f, μ) satisfy *IC* for (ϵ, T) if

$$(11) \quad E \left(\frac{1}{T} \sum_{t=1}^T |\mathbf{x}_{t+1}^\mu - z(\beta_t(\mu), \mathbf{x}_t^\mu)|^2 \right) \leq \min_{\mu'} E \left(\frac{1}{T} \sum_{t=1}^T |\mathbf{x}_{t+1}^\mu - z(\tilde{\beta}_t(\mu, \mu'), \mathbf{x}_t^\mu)|^2 \right) + \epsilon.$$

¹⁸ Bray and Nathan E. Savin (1986) study whether the learning model rejects the hypothesis of serially uncorrelated prediction errors by assuming that agents run a Durbin and Watson test. That exercise carries the flavor of *EDR*.

Thus, if *IC* is satisfied, agents are doing almost as well as possible within the learning mechanism specified after T periods, so that they are likely to stay with μ .¹⁹

IC is, in general, more restrictive than *AR*, since it requires that good forecasts are generated along the transition, not only at the limit. As in the case of *EDR*, it only makes sense to study *IC* in the context of “moderately high” T .

The first two bounds compare the performance of the learning mechanism used by agents relative to an external agent who knows the best prediction that can be computed from knowledge of (f, μ, z, g, η) . The bound *IC*, instead, compares the learning mechanism with forecasts that use the same family of mechanisms f but are allowed to pick alternative parameter values μ . This last bound contains some of the intuition of rational expectations, in the sense of looking for an approximate fixed point in which agents’ expectations minimize the errors within the mechanism f . Notice that this restriction will, in general, imply that agents under different policy environments use different learning parameters μ , so that the learning parameter that satisfies *IC* is endogenous to the model and to government policy. For example, in our model, agents in high seignorage countries (say, Argentina in the 1980’s) will use a different learning parameter from agents in low seignorage countries (say, Switzerland). These Definitions can be readily generalized to more complicated models or to objective functions other than the average prediction error.

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

Rational expectations can be interpreted as imposing extreme versions of the second and third bounds. Obviously, RE satisfies *AR*. It would appear that requiring *EDR* for all $\epsilon, \delta > 0$, and all T is the same as imposing rational expectations, but a careful proof should be worked out. Also, if the RE equilibrium (REE) is recursive, if the appropriate state variables are

¹⁹ Evans and Honkapohja (1993) developed a very similar criterion in a different context.

included in z , if z is a dense class of functions (for example, polynomials or splines), imposing IC for any ϵ , T is the same as rational expectations.

IV. Learning Equilibrium

In this section, we propose a learning mechanism f that combines least-squares learning with tracking and we show that, in the model of Section II, it satisfies the three lower bounds on rationality defined in the previous section.

A. The Learning Mechanism

In the model of Section II with expectations given by (9), we assume that the learning mechanism is given by the stochastic approximation algorithm

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

for given β_0 . That is, perceived inflation β_t is updated by a term that depends on the last prediction error²⁰ weighted by the gain sequence $1/\alpha_t$. Equation (12) together with the evolution of the gains $1/\alpha_t$ determines the learning mechanism f in (8).

One common assumption for the gain sequence is $\alpha_t = \alpha_{t-1} + 1$, for $\alpha_0 = 1$. In this case, $\alpha_t = t$, and simple algebra shows that (with $\beta_0 = 0$)

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

So, in this case, perceived inflation is just the sample mean of past inflations or, equivalently, it is the OLS estimator of the mean of inflation.

Another common assumption for the gain sequence is $\alpha_t = \tilde{\alpha} > 1$. These have been termed “tracking” or “constant gain” algorithms.²¹ In this case, perceived inflation satisfies (with $\beta_0 = 0$)

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

so that past information is now a weighted average of past inflations, where the past is discounted at a geometric rate.²²

Least squares gives equal weight to all past observations, while tracking gives more importance to recent events. Tracking produces better forecasts when there is a sudden change in the environment but it does not converge. OLS is known to be a consistent estimator in stationary setups but it reacts slowly to sudden changes.

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

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

We will specify a learning mechanism that mixes both alternatives: it will use OLS in stable periods and it will switch to “tracking” when some instability is detected. This amounts to assuming that agents use an endogenous gain sequence such that, as long as agents don’t make large prediction errors, α_t follows a least-squares rule, but in periods where a large prediction error is detected, α_t becomes a fixed

²⁰ As usual in models of learning, we make the convenient assumption that the last observation used to formulate expectations is dated at $t - 1$. Including today’s inflation in β_t would make it even easier for the learning scheme to satisfy the lower bounds and to match the stylized facts, and it would not change the dynamics of the model.

²¹ Chung (1990), Evans and Honkapohja (1993), and Sargent (1993) also discuss tracking algorithms.

²² 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 seignorage is autoregressive of order 1, expected inflation would have to depend on current seignorage in order to satisfy any of the lower bounds on rationality. In that case, tracking would be fundamentally different from adaptive expectations.

positive value $\bar{\alpha} \geq 1$ as in “tracking.”²³ Formally, the gain sequence follows

$$(13) \quad \alpha_t = \alpha_{t-1} + 1 \quad \text{if} \quad \left| \frac{\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}}{\beta_{t-1}} \right| < \nu$$

$$= \bar{\alpha} \quad \text{otherwise.}$$

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

The learning mechanism f is fully described by equations (12) and (13). The learning parameters $\mu \equiv (\nu, \bar{\alpha})$ and the statistics $\beta_t \equiv (\beta_t, \alpha_t)$.

B. Learning and Stylized Facts

The variables we need to solve for are $\{P_t, \beta_t, \alpha_t\}$. Simple algebra implies

$$(14) \quad \frac{P_t}{P_{t-1}} = H(\beta_t, \beta_{t-1}, d_t)$$

where²⁴

$$(15) \quad H(\beta_t, \beta_{t-1}, d_t) = \frac{1 - \gamma\beta_{t-1}}{1 - \gamma\beta_t - d_t/\phi} \quad \text{if } 0 < \frac{1 - \gamma\beta_{t-1}}{1 - \gamma\beta_t - d_t/\phi} < \beta^u$$

$$= \bar{\beta} \quad \text{and } 1 - \gamma\beta_{t-1} > 0$$

$$= \bar{\beta} \quad \text{otherwise.}$$

²³ Evans and Garey Ramey (1998) also analyze the properties of a learning mechanism that responds endogenously to the performance of the predictions within the model and within the realization.

²⁴ Notice that the second part of this equation, where ERR will prevail, applies if one of the following (mutually exclusive) cases occurs:

Case (i): $1 - \gamma\beta_{t-1} < 0$, which implies $M_{t-1} = 0$, so the budget constraint of the government is incompatible with the demand for real balances unless reserves adjust.

Case (ii): $1 - \gamma\beta_t - d_t/\phi < 0$ and $1 - \gamma\beta_{t-1} > 0$ so only a negative price level clears the market, and

Case (iii): None of the above and $\frac{1 - \gamma\beta_{t-1}}{1 - \gamma\beta_t - d_t/\phi} > \beta^u$, such that the market generates a level of inflation unacceptable to the government if reserves do not adjust.




FIGURE 2. ACTUAL INFLATION AS A FUNCTION OF PERCEIVED INFLATION

Equations (12), (13), and (14) define a system of stochastic, second-order difference equations. Characterizing the solution analytically is unfeasible since the system is highly nonlinear.

We now provide some intuition on the behavior of inflation. Let $h(\beta, d) \equiv H(\beta, \beta, d)$. Notice that if $\beta_t \approx \beta_{t-1}$, then $P_t/P_{t-1} \approx h(\beta_t, d_t)$, so that the graph of h in Figure 2 provides an approximation to the actual inflation as a function of perceived inflation and it can be used to describe the approximate dynamics of the model.

The first graph corresponds to a low d_t . The low rational expectations equilibrium β_{RE}^1 is

locally stable under least squares learning.²⁵ The horizontal axis can be split into the intervals **S**, **U**, and **ERR**.

If $\beta_t \in \mathbf{S}$, actual inflation is on average closer to β_{RE}^1 than perceived inflation and the learning mechanism pushes perceived inflation towards 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 on average higher than β_t , perceived inflation tends to increase and a hyperinflation is likely to occur. Then, when the set **ERR** is reached, a fixed exchange rule is established and inflation is sent back to **S**. The economy may end up in the unstable set **U** due to a number of reasons: a few high shocks to seignorage when $1/\alpha_t$ is not yet close to zero, initially high perceived inflation, the second-order dynamics adding momentum to increasing inflation, etc.

If a shock to inflation occurs agents are likely to switch to tracking and set $1/\alpha_t = 1/\bar{\alpha}$, perceived inflation will then be more heavily influenced by actual inflation and it is more likely to end up in **U** than under pure OLS. Hyperinflations prompt agents to switch to tracking and to pay more attention to recent observations; this in turn makes hyperinflations more likely to occur and predictions with tracking better, thus reinforcing the switch to tracking in periods of instability. Only if $1/\alpha_0$ is very small relative to the variance of inflation and if initial inflation starts out in **S** (and ν is large enough), hyperinflations are impossible.

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 is likely to grow even if seignorage does not, which is consistent with fact 3.

To analyze fact 4, consider the second graph of Figure 2, for a high d_t . Now, the unstable set **U** is much larger. Furthermore, **U** is “dangerously” close to the rational expectations equilibrium β_{RE}^1 where the economy tends to live, and it is likely for the model to end up in **U** and

a hyperinflation to occur even if inflation has been stable for a while. Thus, a country with a high average seignorage tends to have hyperinflationary episodes more often, and fact 4 is consistent with the model.

C. Asymptotic Rationality (AR)

To prove convergence to RE and that AR is satisfied we need:

ASSUMPTION 1: *The support of d_t is the interval $[K^-, K^+]$, where $K^- > 0$, $K^+ < \phi$.*

ASSUMPTION 2: *d_t has a continuous density f_{d_t} and $f_{d_t}(K^+) \equiv \Pi > 0$.*

Letting $S(\beta) \equiv E(h(\beta, d_t))$, Appendix C [available from the authors, and in the working paper version of the paper (Marcet and Nicolini, 1998)] proves that

- S is increasing and convex;
- as $\beta^U \rightarrow \infty$, S has an asymptote at $\beta = (1 - K^+/\phi)/\gamma$;
- S has at most two fixed points $\beta_{RE}^1, \beta_{RE}^2$;
- $\beta_{RE}^1, \beta_{RE}^2$ are the stationary rational expectations equilibria;
- two fixed points exist if f_{d_t} is close enough to zero;
- no fixed point exists if f_{d_t} is large enough;
- as we consider larger distributions for d_t (as f_{d_t} shifts to the right) $\beta_{RE}^1, \beta_{RE}^2$ get closer together. Therefore, for larger d_t the stable set **S** shrinks and β_{RE}^1 is closer to the unstable set **U**;
- if two fixed points exist, least-squares learning converges to β_{RE}^1 a.s.

PROPOSITION 1: *In addition to Assumptions 1–2, assume that average seignorage and its variance are low enough for two stationary REE to exist, that $\bar{\beta} \in \mathbf{S}$ (targeted inflation belongs to the stability set) and that $\bar{\alpha}$ and ν are large enough.²⁶ Then $\beta_t \rightarrow \beta_{RE}^1$ a.s. and Asymptotic Rationality obtains.*

²⁵ This discussion assumes that $S(\beta)$ is close to $h(\beta, d_t)$, which is approximately correct if d_t is close to its expectation. The proofs of stability of least-squares learning and of the properties of S are in an Appendix available from the authors and in the working paper version of the paper.

²⁶ The assumption on $\bar{\alpha}$ can be interpreted as saying that convergence occurs if the importance given to recent news is never too high. This assumption is needed in order to obtain a lower bound of inflation in the first part of the proof. A lower bound on inflation can also be obtained for unrestricted $\bar{\alpha}$ by changing the model in reasonable ways.

PROOF:

The theorem holds for

$$(16) \quad \bar{\alpha} > \frac{1 - \gamma\bar{\beta}}{K^-/\phi}$$

$$\text{and } \nu > \left| \frac{\beta^U}{\min\left(1, \frac{K^-/\phi}{1 - K^-/\phi}, \bar{\beta}\right)} - 1 \right|.$$

In order to show that the learning mechanism stays in the OLS form in all periods $t > \bar{\alpha}$, we first show that inflation is bounded below. For each t and each realization, only three cases are possible: Case (i): an ERR is activated at t , (ii): an ERR is not activated at t and $1 - \gamma\beta_{t-1} - d_{t-1}/\phi \leq 0$, and (iii): an ERR is not activated at t and $1 - \gamma\beta_{t-1} - d_{t-1}/\phi > 0$.

Notice that in cases (ii) and (iii) the first branch of (15) applies and we have $\beta_{t-1} < \gamma^{-1}$. Note also that β_t is a weighted average of past inflations and β^U is an upper bound of inflation so that $\beta_t < \beta^U$. We now find a $\beta^L > 0$ such that $\beta_t > \beta^L$ for all t .

In case (i), inflation is equal to $\bar{\beta}$.

In case (ii)

$$(17) \quad \frac{P_t}{P_{t-1}} = \frac{1 - \gamma\beta_{t-1}}{1 - \gamma\beta_{t-1} - \frac{\gamma}{\alpha_t} \left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) - d_t/\phi} = \frac{1}{1 - \frac{\gamma}{\alpha_t} (\bar{\beta} - \beta_{t-1}) + d_t/\phi}$$

where the first equality follows from (15) and

(12), and the second equality follows from the fact that an ERR was established at $t - 1$ so that $P_{t-1}/P_{t-2} = \bar{\beta}$. Now, using $\beta_{t-1} < \gamma^{-1}$, $\alpha_t \geq \bar{\alpha}$, $d_t \geq K^-$ and (16) we have

$$\frac{\gamma}{\alpha_t} (\bar{\beta} - \beta_{t-1}) + \frac{d_t}{\phi} > \frac{\gamma}{\bar{\alpha}} (\bar{\beta} - \gamma^{-1}) + \frac{K^-}{\phi} > 0$$

which, together with (17), implies that $P_t/P_{t-1} > 1$.

In case (iii), the condition on β_{t-1} and simple algebra imply

$$\frac{P_t}{P_{t-1}} > \frac{d_{t-1}/\phi}{1 - \gamma\beta_{t-1} - d_t/\phi} > \frac{K^-/\phi}{1 - K^-/\phi}.$$

Therefore we find the lower bound $P_t/P_{t-1} \geq \min[1, (K^-/\phi)/(1 - K^-/\phi), \bar{\beta}] \equiv \beta^L > 0$. Since β_t is an average of past inflations we also have $\beta_t > \beta^L$ for $t > \bar{\alpha}$ large enough such that the effect of the initial condition has disappeared.

For any $\nu > |\beta^U/\beta^L - 1|$ we clearly have $|(P_{t-1}/P_{t-2} - \beta_{t-1})/\beta_{t-1}| < \nu$ with probability one for all $t > \bar{\alpha}$, then $\alpha_t = \alpha_{t-1} + 1$ for all t large enough and the learning mechanism stays in the OLS form.

Now, let $C^S \equiv \{\omega \in \Omega : \beta_t(\omega) \in \mathbf{S} \text{ i.o.}\}$, $C^U \equiv \{\omega \in \Omega : \beta_t(\omega) \in \mathbf{U} \text{ i.o.}\}$, and $C^{\text{ERR}} \equiv \{\omega \in \Omega : \beta_t(\omega) \in \mathbf{ERR} \text{ i.o.}\}$ we want to argue that $\mathcal{P}(C^S) = 1$. Clearly, any realization ω belongs to either C^S , C^U , or C^{ERR} . Consider a $\bar{\omega} \in C^{\text{ERR}}$; for any t such that $\beta_t(\bar{\omega}) \in \mathbf{ERR}$, an ERR is then enforced sufficiently long for the beliefs to go back to the stable set, so that $\beta_{t+j}(\bar{\omega}) = \bar{\beta} \in \mathbf{S}$ for some j ; therefore, it is clear that also $\bar{\omega} \in C^S$. Therefore $C^{\text{ERR}} \subset C^S$. Now consider $\bar{\omega} \in C^U$, the theorem of Ljung (1977) implies that the differential equation $\dot{\beta} = S(\beta) - \beta$ in Appendix C (available from the authors or in the working paper version of this article) governs the dynamics of β_t , therefore for t large enough inflation tends to grow and eventually goes into \mathbf{ERR} , which would also imply $\bar{\omega} \in C^S$. Therefore $\omega \in C^S$ for all ω , and $\beta_t \in \mathbf{S}$ i.o. with probability one. In Appendix C we apply the o.d.e. approach for convergence of learning schemes in dynamic models to show that this implies that β_t converges to β_{RE}^1 almost surely.

The rest of the proof simply shows that, if the learning scheme converges to β_{RE}^1 , then the sample mean square errors converge to the best

For example, assuming that the government has the objective of avoiding deflation and it achieves this by activating an ERR and insures that $P_t/P_{t-1} \geq 1$ at the same time that reserves increase. A lower bound in ν can be interpreted as saying that agents do not easily switch to tracking; a lower bound is necessary because, if ν is too small, even if β_t is very close to β_{RE}^1 , it will eventually happen that $\left| \frac{P_{t-1}}{P_{t-2}\beta_{t-1}} - 1 \right| > \nu$, then $\alpha_t = \bar{\alpha}$, perceived inflation will have positive variance, and convergence will never occur.

forecasts and AR obtains. Notice that $\beta_t \rightarrow \beta_{RE}^1$ a.s. and the fact that H is continuous at β_{RE}^1 imply

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

as $t \rightarrow \infty$, where we used the fact that, by definition, $P_t^{RE}/P_{t-1}^{RE} = H(\beta_{RE}^1, \beta_{RE}^1, d_t)$. Now

$$\begin{aligned} & \left| E_{t-1} \left(\frac{P_t}{P_{t-1}} \right) - \beta_{RE}^1 \right| \\ &= \left| E_{t-1} \left(\frac{P_t}{P_{t-1}} \right) - E_{t-1} \left(\frac{P_t^{RE}}{P_{t-1}^{RE}} \right) \right| \rightarrow 0 \text{ a.s.} \end{aligned}$$

where $\beta_{RE}^1 = E(P_t^{RE}/P_{t-1}^{RE})$ is by definition, and convergence follows by Lebesgue-dominated convergence and boundedness of inflation. Therefore, both β_t and $E_{t-1}(P_t/P_{t-1})$ converge to β_{RE}^1 . This, together with boundedness of prices and β implies that

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

as $T \rightarrow \infty$, so that

$$\begin{aligned} \pi^{\epsilon, T} &\equiv \mathcal{P} \left(\frac{1}{T} \sum_{i=1}^T \left[\frac{P_i}{P_{i-1}} - \left(\frac{P_i}{P_{i-1}} \right)^{\epsilon} \right]^2 \right. \\ & \left. < \frac{1}{T} \sum_{i=1}^T \left[\frac{P_i}{P_{i-1}} - E_{i-1} \left(\frac{P_i}{P_{i-1}} \right) \right]^2 + \epsilon \right) \rightarrow 1 \text{ a.s.} \end{aligned}$$

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

Notice that AR imposes very few restrictions on the learning scheme. In particular AR holds for many $\bar{\alpha}$'s. Even if AR is satisfied, agents could be making systematic mistakes; for example, in periods where α_i is updated according to OLS, agents could be making very large

forecasting errors whenever a hyperinflation happened, since OLS does not weigh recent events more heavily.

D. Internal Consistency (IC)

In Section IV, subsection B, we explained intuitively why hyperinflations are more likely to occur with high $1/\bar{\alpha}$. Also, a high value of $1/\bar{\alpha}$ is likely to generate better forecasts during a hyperinflation. Therefore, there is potential for IC to be satisfied precisely for the $1/\bar{\alpha}$'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 $\bar{\alpha}$. Notice that, since $\bar{\alpha}$ is determined as part of the equilibrium, the $\bar{\alpha}$ that satisfies IC will vary as the process for d_t changes so that the learning mechanism is endogenous to government policy.

Definition 4: A stochastic process $\{P_t, \beta_t, M_t\}$ together with $\bar{\alpha}$ is an IC equilibrium for (ϵ, T) if:

1. Given $\bar{\alpha}$, $\{P_t, \beta_t, M_t\}$ satisfy (7), (12), (13), (14) for all t .
2. $\bar{\alpha}$ satisfies IC for (ϵ, T) .²⁷

Since the dynamics are highly nonlinear, characterizing analytically the equilibrium $\bar{\alpha}$'s is impossible. We solve the model by simulation 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.

E. Characterization of the Solution by Simulation

To generate simulations we must assign values to the parameters of the money demand equation (γ, ϕ) and the distribution of d_t . We choose values $(\gamma = 0.4$ and $\phi = 0.37)$ in order to replicate some patterns of the Argentinean

²⁷ The careful reader will note that we did not impose IC on the learning parameter ν in this definition or in the simulations we describe below. This was done only for simplicity.

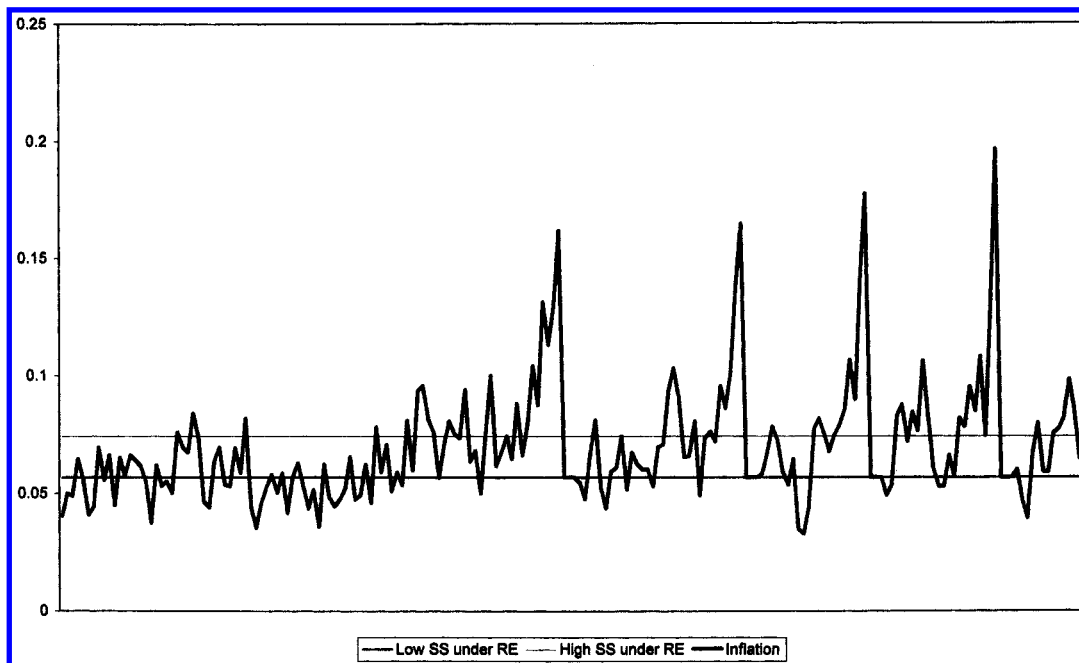


FIGURE 3. SIMULATION OF THE (LOG) INFLATION

experience during the 1980's; for details see Appendix B. We assume that seignorage is normally distributed, truncated to have positive values of seignorage, with mean that varies across experiments we perform and $\sigma_d = 0.01$.²⁸

The parameter ν was set equal to 10 percent. We also assumed that the government established ERR whenever expectations were such that inflation rates would be above 5,000 percent, so that we set $\beta^U = 50$. The ERR is enforced until expected inflation is inside the stable set S .

Since our purpose is to show that a small deviation from rational expectations can generate dynamics quite different from and closer to the data than RE, we choose as initial beliefs $\beta_0 = \beta_{RE}^1$ so that our simulations are biased in favor of looking like RE.²⁹ For the specified parameters, the maximum level of average seignorage in the deterministic model for which a REE exists is $E(d_t) = 0.05$. In the spirit of

making it difficult for the model to depart from RE, we have chosen values of the average seignorage for which a REE exists. In order to quantify the relevance of average seignorage (fact 4), we performed our calculations for four different values: $E(d_t) = 0.049, 0.047, 0.045,$ and 0.043 .

First of all, we describe the typical behavior of the model. A particular realization is presented in Figure 3. That realization was obtained with $E(d_t) = 0.049$ and $1/\bar{\alpha} = 0.2$. 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 deterministic rational expectation equilibria, to show how the model under learning can generate much higher inflation rates than the rational expectations version.

This graph displays some of the stylized facts in the learning model.³⁰ In the first periods, the

²⁸ The results for lower values of σ_d^2 were similar. Of course, hyperinflations were then less frequent.

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

³⁰ 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_{RE}^1 . For bubble equilibria, inflation would grow towards the horizontal line of β_{RE}^2 .

inflation rate is close to the low stationary equilibrium. When a relatively large shock occurs, it drives perceived inflation into the unstable region \bar{U} and a hyperinflation episode starts. Eventually, ERR is established and the economy is brought back into the stable region. If no large shocks occur 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, β_1^{RE} is close to the unstable set (see Figure 2) and it is likely that a new large shock will put the economy back into the unstable region and a new burst in inflation will occur. Clearly, we have recurrent hyperinflations, stopped by ERR (facts 1 and 2). Since seignorage is i.i.d., and since the graph shows some periods of sustained increases in inflation, it is clear that there is little correlation of inflation and seignorage (fact 3). In order to reduce (or eventually eliminate) the chances of having a new burst, the government must reduce the amount of seignorage collected (i.e., an “orthodox” stabilization plan) in order to increase the size of the stable set. This would separate the two horizontal lines, it would place β_1^{RE} far from the unstable set and it would stabilize the economy permanently around the low stationary equilibrium. Establishing ERR just before a reduction in average seignorage would help stabilize the expectations of agents more quickly, so there is room for a positive effect of a “heterodox” intervention as well.

To find the learning equilibrium we look for values of $\bar{\alpha}$ that satisfy the lower bound criterion IC for $(\epsilon, T) = (0.01, 120)$. This value of T is chosen to represent ten years, roughly 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.

To find numerically those values of $\bar{\alpha}$ that satisfy IC we proceed as follows: we define a grid of $1/\bar{\alpha} \in [0, 1.2]$ separated by intervals of length 0.1. The same grid is used both for $1/\bar{\alpha}$ 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,³¹ and we find the minimum over

$1/\bar{\alpha}'$ for each $1/\bar{\alpha}$ on the grid. Figure 4 shows the result of these calculations: in the horizontal axis we plot $1/\bar{\alpha}$, while the vertical axis plots $1/\bar{\alpha}'$. The interval of alternative learning parameters that generate a mean square error within $\epsilon = 0.01$ 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.

Table 1 reports the probabilities of having n hyperinflations in ten years for different values of average seignorage and for those values of $1/\bar{\alpha}$ that satisfy the IC criterion.

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

This exercise formalizes the sense in which the equilibria with a given learning mechanism reinforces the use of the mechanism. For instance, when seignorage is 0.49 and $1/\bar{\alpha} = 0.2$, 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 $1/\bar{\alpha} = 0.2$. The reason is that in equilibrium there are many hyperinflations, and the agent that expects the REE will make bad forecasts.

³¹ More specifically, we draw 1,000 realizations of $\{d_1, \dots, d_{120}\}$, find the equilibrium inflation rates for each

realization, we compute the sample mean square error for each alternative $1/\bar{\alpha}'$ in the grid, and we average over all realizations.

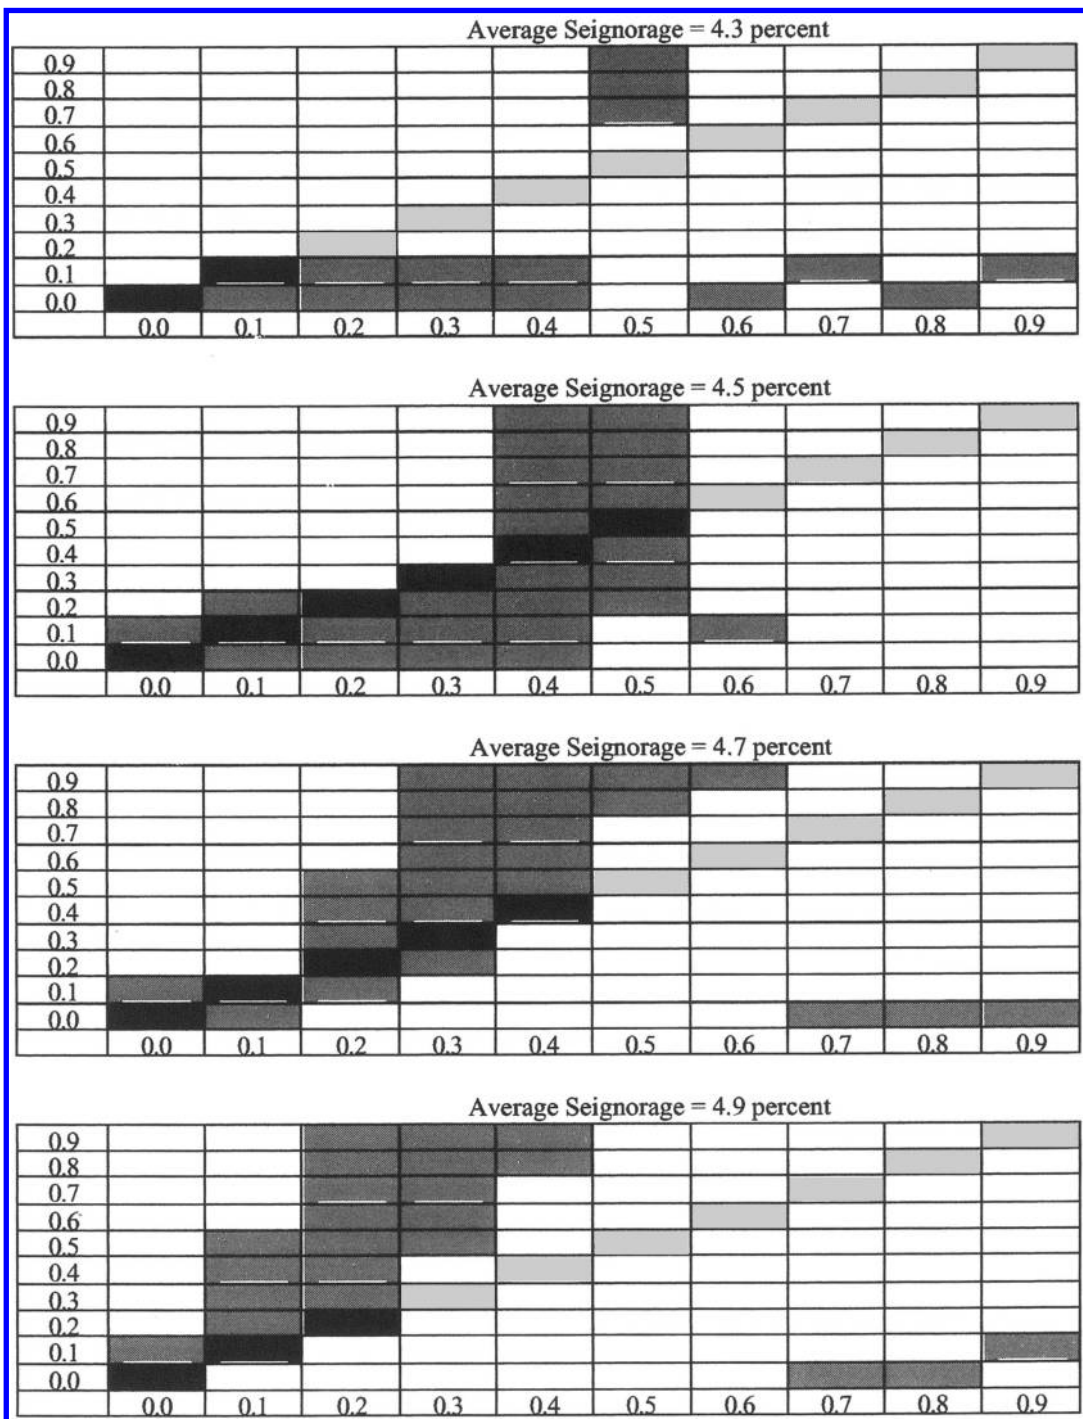


FIGURE 4. INTERNAL CONSISTENCY

Notes: Columns represent possible values for $1/\bar{\alpha}$ actually used by agents. Rows depict alternative values for $1/\bar{\alpha}'$. Light gray cells indicate the 45-degree line. Dark gray cells indicate that the value for $1/\bar{\alpha}$ is efficient. Black cells indicate fixed points on $1/\bar{\alpha}$.

TABLE 1—PROBABILITIES OF HYPERINFLATIONS OCCURRING IN TEN YEARS FOR DIFFERENT DEFICIT MEANS AND LEARNING PARAMETERS

Deficit mean = 4.5 percent						
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	
Deficit mean = 4.7 percent						
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	
Deficit mean = 4.9 percent						
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	

Whenever equilibria with hyperinflations exist, there is multiplicity of equilibria (several $1/\bar{\alpha}$'s satisfy *IC*). The behavior of inflation does not change much for different equilibrium $1/\bar{\alpha}$'s.³²

The numerical solutions show that the chances of a hyperinflation during the transition to the rational expectations equilibrium depend on the size of the deficit. The lower the deficit, the lower the chances of experiencing a hyperinflation. Notice how the equilibrium learning parameters depend on the size of average seignorage: higher seignorage corresponds to higher equilibrium $\bar{\alpha}$'s, which are more likely to generate a hyperinflation.³³

³² Since we chose $\beta_0 = \beta_{RE}^1$ when we set $1/\bar{\alpha} = 0$ we have the REE. When initial beliefs are far apart from the REE, then $1/\bar{\alpha} = 0$ will no longer satisfy *IC*.

³³ We have simulated the model under many other values for the parameters. The main results of this subsection about the behavior of inflation are observed for a wide range of the parameters.

F. Epsilon-Delta Rationality (EDR)

In this subsection 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, when a hyperinflation occurs, the conditional expectation can be arbitrarily high due to the existence of an asymptote in the mapping from perceived to actual inflation (see Figure 2) but, in fact, the actual value of inflation is unlikely to be ever so high in a given realization. Thus, for every realization when a hyperinflation occurs, the learning forecast can do better than the conditional expectation with very high probability in finite samples.

PROPOSITION 2: *Consider the model of Section II. If Assumptions 1–2 are satisfied, then for given parameter values of the model and any (ϵ, T) , there is a β^U large enough such that*

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

where $\mathcal{P}(\text{ERR at some } t \leq T)$ is the probability that the government implements ERR at some period $t = 1, \dots, T$.

PROOF:

Fix ϵ, T . We first consider the case that $\beta^U = \infty$. Consider a realization $\bar{\omega}$ where an ERR is established at some $t = 0, \dots, T$. Letting $\bar{t} + 1$ be the *first* period where this occurs, it has to be the case that $1 - \gamma\beta_{\bar{t}+1}(\bar{\omega}) - d_{\bar{t}+1}(\bar{\omega})/\phi < 0$ and $1 - \gamma\beta_{\bar{t}}(\bar{\omega}) > 0$. Clearly, we can only have the first inequality if $\bar{d} \equiv \phi(1 - \gamma\beta_{\bar{t}+1}(\bar{\omega})) < K^+$. Therefore, since inflation is given by equation (14) and $(\beta_{\bar{t}+1}, \beta_{\bar{t}})$ are known with information available at \bar{t} , we have that

$$\begin{aligned} (18) \quad E_{\bar{t}}\left(\frac{P_{\bar{t}+1}}{P_{\bar{t}}}\right)(\bar{\omega}) &= \int_{K^-}^{K^+} H(\beta_{\bar{t}+1}(\bar{\omega}), \beta_{\bar{t}}(\bar{\omega}), \bar{d}) dF_{d_{\bar{t}+1}}(\bar{d}) \\ &= \int_{K^-}^{\bar{d}} \frac{1 - \gamma\beta_{\bar{t}}(\bar{\omega})}{1 - \gamma\beta_{\bar{t}+1}(\bar{\omega}) - \bar{d}/\phi} dF_{d_{\bar{t}+1}}(\bar{d}) \\ &\quad + \mathcal{P}[d_{\bar{t}+1} > \bar{d}]\bar{\beta}. \end{aligned}$$

The integral in (18) corresponds to the values of $d_{\bar{t}+1}$ for which there is a positive price level that clears the market without ERR and the first branch of (15) holds, while the second term accounts for those values of next period shock for which an exchange rate rule needs to be enforced.

Now we show that the integral in (18) is unbounded. Using arguments similar to the ones used in Appendix C (available from the authors or in the working paper version of this article) to show that S has an asymptote we have

$$\begin{aligned} &\int_{K^-}^{\bar{d}} \frac{1 - \gamma\beta_{\bar{t}}(\bar{\omega})}{1 - \gamma\beta_{\bar{t}+1}(\bar{\omega}) - \bar{d}/\phi} dF_{d_{\bar{t}+1}}(\bar{d}) \\ &\geq (1 - \gamma\beta_{\bar{t}}(\bar{\omega}))Q(\bar{\omega}) \int_0^\eta \frac{1}{x} dx = \infty \end{aligned}$$

for some finite constant $Q(\bar{\omega})$ and small η .

This proves that $E_{\bar{t}}(P_{\bar{t}+1}/P_{\bar{t}})(\bar{\omega}) = \infty$, therefore

$$\begin{aligned} (19) \quad &\frac{1}{T} \sum_{t=1}^T \left[\frac{P_{t+1}(\bar{\omega})}{P_t(\bar{\omega})} - \frac{P_{t+1}^e}{P_t}(\bar{\omega}) \right]^2 \\ &< \frac{1}{T} \sum_{t=1}^T \left[\frac{P_{t+1}(\bar{\omega})}{P_t(\bar{\omega})} - E_t\left(\frac{P_{t+1}}{P_t}\right)(\bar{\omega}) \right]^2 + \epsilon, \end{aligned}$$

because the right-hand side is, in fact, infinite. So, (19) holds for all realizations where there is one hyperinflation and $\pi^{\epsilon, T} \geq \mathcal{P}[\text{ERR at some } t \leq T]$.

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 $\underline{\beta}$; also, for arbitrarily high β^U the conditional expectation is arbitrarily close to the one with β^U infinite, so that all the inequalities are maintained with arbitrarily high probability.

Since hyperinflations occur with high probability if average seignorage is high, this proposition shows that EDR is satisfied with high δ when seignorage is high. For example, Table 1 shows that the probability of having at least one hyperinflation is 0.84, 0.91, and 0.97 for average seignorage 0.045, 0.047, and 0.049, respectively, so this proposition implies that EDR is satisfied for $\delta > 0.84, 0.91, \text{ and } 0.97$.

V. Conclusion

There is some agreement by now that the hyperinflations of the 1980's were caused by the high levels of seignorage in those countries, and that the cure for those hyperinflations was fiscal discipline and abstinence from seignorage. The IMF is currently imposing tight fiscal controls on the previously hyperinflationary countries that are consistent with this view. Nevertheless, to our knowledge, no currently available model justified this view and was consistent with some basic facts of hyperinflations. In particular, the fact that seignorage has gone down during some hyperinflations makes it difficult for the IMF to argue in favor of fiscal discipline.

Our model is consistent with the main stylized facts of recurrent hyperinflations and with

the policy recommendations mentioned above: an exchange rate rule (ERR) may temporarily stop a hyperinflation, but average seignorage must be lowered to eliminate hyperinflations permanently.

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

On the practical side, this paper shows that hyperinflations can be stopped with a combination of heterodox and orthodox policies. We have been working on this paper in the second half of the 1990's; at that time it might have seemed that hyperinflations were a purely academic issue: South American countries seemed to have solid fiscal stances and hyperinflations were a thing of the past. Unfortunately, the recent events in Argentina and the experience of some Eastern European countries have lent some immediate interest to the policy conclusions of this paper. It still seems important to have a solid model that can help judging the reasonability of the IMF recommendations. The methodological contribution of the paper is to show that, with adequate equipment for orientation and survival, an expedition into the "wilderness of irrationality" can be quite safe and productive.

APPENDIX A

Households:

To provide some microfoundations for the model in Section II, subsection A, we solve a deterministic small open economy version of a standard overlapping generations model.³⁴

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

Asset markets:

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

$$M_t \geq P_{t+1}(c_{t+1}^o - \omega^o)$$

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

$$P_t \omega^y = P_t c_t^y + M_t$$

$$M_t = P_{t+1}(c_{t+1}^o - \omega^o)$$

$$M_t \geq 0.$$

Household's optimization implies

³⁴ This Appendix extends the closed economy results of Sargent and Wallace (1987).

$$(A1) \quad c_t^y = \frac{P_t \omega^y + P_{t+1} \omega^o}{(1 + \lambda) P_t}, \quad \frac{M_t}{P_t} = \frac{\omega^y \lambda}{(1 + \lambda)} - \frac{P_{t+1}}{P_t} \frac{\omega^o}{(1 + \lambda)} \quad \text{if } \frac{\omega^y \lambda}{\omega^o} \geq \frac{P_{t+1}}{P_t}$$

$$c_t^y = \omega^y, \quad \frac{M_t}{P_t} = 0 \quad \text{otherwise.}$$

This gives a microfoundation to equation (1), with $\phi = \frac{\omega^y \lambda}{1 + \lambda}$ and $\gamma = \frac{\omega^o}{\omega^y \lambda}$.

Foreign sector:

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

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

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

$$(A2) \quad P_t^f e_t = P_t.$$

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

$$\omega^y + \omega^o = c_t^y + c_t^o + d_t + TB_t$$

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

The government budget constraint:

We assume the government does not tax agents,³⁵ it only generates income by seignorage and, occasionally, by changing its stock of foreign currency R_t . The budget constraint of the government is therefore given by

$$(A3) \quad \frac{M_t - M_{t-1}}{P_t} = d_t + (R_t - R_{t-1}) \frac{e_t}{P_t}.$$

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

Government policy:

Government policy must set money supply and reserves to satisfy (A3). Reserves can be set according to two regimes:

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

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

which together with the money demand (A1) solve for the equilibrium sequences of M_t and P_t . The nominal exchange rate is given by equation (A2).

³⁵ Taxes and government debt are easy to introduce by reinterpreting d_t and the endowments ω : all equations are consistent with ω denoting endowments net of age-dependent, constant, lump-sum taxes, and with d_t being the primary deficit of the government. Debt can be introduced, for example, if we assume that government debt is constant (perhaps because the government is debt constrained) and d_t represents interest payments on debt plus primary deficit.

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

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

Equation (A2) implies that with this policy the government achieves $\bar{\beta} = P_t/P_{t-1}$. The money demand (A1) determines the level of nominal money demand consistent with the nominal exchange rate target. Given this level of money supply and d_t , foreign reserves and, consequently, the trade balance adjust so as to satisfy (A3). Of course, ERR is only feasible if the constraints on the government asset position is never binding.

We assume that the first regime is used if inflation achieves an acceptable level less than β^U ; the ERR regime is followed otherwise.

The equilibrium is therefore given by equations (A1), (A2), and (A3) which are deterministic versions of equations (1) to (4) in the paper. The analogy between this deterministic version and the stochastic one in the paper is only exact up to a linear approximation, a usual simplification in macroeconomic models under learning.

APPENDIX B

In this Appendix we explain the choice of parameter values for the demand for money used in the numerical solution of Section IV. The money demand equation (1) is linear with respect to expected inflation. It is well known, though, that the linear functional form does not perform very well empirically. However, departing from linearity would make the analysis of the model impossible to deal with. While we do maintain linearity, we want to use parameter values that are not clearly at odds with the observations. Since we are interested in the public finance aspect of inflation, we use observations from empirical Laffer curves to calibrate the two parameters. In particular, as one empirical implication of our model is that “high” average deficits increase the probability of a hyperinflation, we need to have a benchmark to discuss what high means. Thus, a natural restriction to impose on our numbers is that the implied maximum deficit is close to what casual observation of the data suggests. 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 from 1980 to 1990 from Hidelgort Ahumada et al. (1993) to fit an empirical Laffer curve. While there is a lot of dispersion, the maximum feasible seignorage is around 5 percent of GNP, and the inflation rate that maximizes seignorage is close to 60 percent. These figures are roughly consistent with the findings in Miguel Kiguel and Pablo A. Neumeyer (1995) and other studies. 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

$$(B1) \quad \frac{\pi}{1 + \pi} m = \frac{\pi}{1 + \pi} \phi(1 - \gamma(1 + \pi))$$

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

$$\pi^* = \sqrt{\frac{1}{\gamma}} - 1$$

which, setting $\pi^* = 60$ percent, implies $\gamma = 0.4$. Using this figure in (B1), and making the maximum revenue equal to 0.05, we obtain $\phi = 0.37$.

REFERENCES

- Ahumada, Hidalgort; Canavese, Alfredo; Sanguinetti, Pablo and Sosa, Walter.** "Efectos Distributivos del Impuesto Inflacionario: Una Estimación del Caso Argentino." *Economía Mexicana*, July–December 1993, II(2), pp. 329–85.
- Arifovic, Jasmina; Bullard, James and Duffy, John.** "The Transition from Stagnation to Growth: An Adaptive Learning Approach." *Journal of Economic Growth*, June 1997, 2(2), pp. 185–209.
- Bental, Benjamin and Eckstein, Zvi.** "On the Fit of a Neoclassical Monetary Model in High Inflation: Israel 1972–1990." *Journal of Money, Credit, and Banking*, November 1997, Pt. 2, 29(4), pp. 725–52.
- Bray, Margaret M. and Kreps, David M.** "Rational Learning and Rational Expectations," in George Feiwel, ed., *Arrow and the ascent of modern economic theory*. New York: New York University Press, 1987, pp. 597–625.
- Bray, Margaret M. and Savin, Nathan E.** "Rational Expectations Equilibria, Learning, and Model Specification." September 1986, *Econometrica*, 54(5), pp. 1129–60.
- Bruno, Michael; Di Tella, Guido; Dornbusch, Rudiger and Fischer, Stanley.** *Inflation stabilization: The experience of Israel, Argentina, Brasil and Mexico*. Cambridge, MA: MIT Press, 1988.
- Bruno, Michael; Fischer, Stanley; Helpman, Elhanan and Liviatan, Nissan.** *Lessons of economic stabilization and its aftermath*. Cambridge, MA: MIT Press, 1991.
- Chung, Hee-Taik.** "Did Policy Makers Really Believe in the Phillips Curve? An Econometric Test." Ph.D. dissertation, University of Minnesota, 1990.
- Easley, David and Rustichini, Aldo.** "Choice without beliefs." *Econometrica*, September 1999, 67(5), pp. 1157–84.
- Eckstein, Zvi and Leiderman, Leonardo.** "Seigniorage and the Welfare Cost of Inflation: Evidence from an Intertemporal Model of Money and Consumption." *Journal of Monetary Economics*, June 1992, 29(3), pp. 389–410.
- Evans, George W. and Honkapohja, Seppo.** "Adaptive Forecasts, Hysteresis, and Endogenous Fluctuations." *Federal Reserve Bank of San Francisco Economic Review*, 1993, (1), pp. 3–13.
- _____. "Learning Dynamics," in John Taylor and Michael Woodford, eds., *Handbook of macroeconomics*, Vol. 1A. Amsterdam: North-Holland, 1999, pp. 449–542.
- _____. *Learning and expectations in macroeconomics*. Princeton, NJ: Princeton University Press, 2001.
- Evans, George W.; Honkapohja, Seppo and Marimon, Ramon.** "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules." *Macroeconomic Dynamics*, February 2001, 5(1), pp. 1–31.
- Evans, George W. and Ramey, Garey.** "Calculation, Adaptation and Rational Expectations." *Macroeconomic Dynamics*, June 1998, 2(2), pp. 156–82.
- Fudenberg, Drew and Levine, David K.** "Consistency and Cautious Fictitious Play." *Journal of Economic Dynamics and Control*, July–September 1995, 19(5–7), pp. 1065–89.
- Funke, Michael; Hall, Stephen and Solà, Martín.** "Rational Bubbles during Poland's Hyperinflation: Implications and Empirical Evidence." *European Economic Review*, June 1994, 38(6), pp. 1257–76.
- Hommes, Cars and Sorger, Gerhard.** "Consistent Expectations Equilibria." *Macroeconomic Dynamics*, September 1998, 2(3), pp. 287–321.
- İmrohoroğlu, Selahattin.** "Testing for Sunspot Equilibria in the German Hyperinflation." *Journal of Economic Dynamics and Control*, January–March 1993, 17(1–2), pp. 289–317.
- Kiguel, Miguel and Neumeyer, Pablo A.** "Seigniorage and Inflation: The Case of Argentina." *Journal of Money, Credit, and Banking*, August 1995, 27(3), pp. 672–82.
- Kurz, Mordecai.** "On Rational Belief Equilibria." *Economic Theory*, October 1994, 4(6), pp. 859–76.
- Ljung, Lennart.** "Analysis of Recursive Stochastic Algorithms." *IEEE Trans. Automatic Control AC-22*, August 1977, pp. 551–57.
- Marcet, Albert and Nicolini, Juan Pablo.** "Recurrent Hyperinflation and Learning." Centre for Economic Policy Research Discussion Paper No. 1875, May 1998.
- Marcet, Albert and Sargent, Thomas J.** "Least-Squares Learning and the Dynamics of Hyperinflation," in William Barnett, John Geweke, and Karl Shell, eds., *International*

- symposia in economic theory and econometrics*. Cambridge: Cambridge University Press, 1989, pp. 119–37.
- Marimon, Ramon.** “Learning from Learning in Economics,” in David M. Kreps and Kenneth F. Wallis, eds., *Advances in economics and econometrics: Theory and applications*, Vol. 1. Cambridge: Cambridge University Press, 1997, pp. 278–315.
- Nicolini, Juan Pablo.** “Ruling Out Speculative Hyperinflations: The Role of the Government.” *Journal of Economic Dynamics and Control*, May 1996, 20(5), pp. 791–809.
- Obstfeld, Maurice and Rogoff, Kenneth.** “Speculative Hyperinflations in Maximizing Models: Can We Rule Them Out?” *Journal of Political Economy*, August 1983, 91(4), pp. 675–87.
- Sargent, Thomas J.** *Bounded rationality in macroeconomics*. Oxford: Oxford University Press, 1993.
- Sargent, Thomas J. and Wallace, Neil.** “Inflation and the Government Budget Constraint,” in Assaf Razin and Efraim Sadka, eds., *Economic policy in theory and practice*. New York: St. Martin’s Press, 1987, pp. 170–200.
- Timmermann, Allan.** “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices.” *Quarterly Journal of Economics*, November 1993, 108(4), pp. 1135–45.
- . “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning.” *Review of Economic Studies*, October 1996, 63(4), pp. 523–57.
- Zarazaga, Carlos.** “Hyperinflations and Moral Hazard in the Appropriation of Seigniorage.” Federal Reserve Bank of Philadelphia Research Working Paper No. 93–26, November 1993.

This article has been cited by:

1. Josef Hollmayr, Christian Matthes. 2015. Learning about fiscal policy and the effects of policy uncertainty. *Journal of Economic Dynamics and Control* **59**, 142-162. [[CrossRef](#)]
2. George A. Waters Careful Price Level Targeting 29-40. [[CrossRef](#)]
3. Francesco Caprioli. 2015. Optimal fiscal policy under learning. *Journal of Economic Dynamics and Control* . [[CrossRef](#)]
4. Fernando de Holanda Barbosa, Tito Nícias Teixeira da Silva Filho. 2015. Bubble, weak and strong hyperinflation: Theory and empirical evidence. *Economia* . [[CrossRef](#)]
5. Thomas W.L. Norman. 2015. Learning, hypothesis testing, and rational-expectations equilibrium. *Games and Economic Behavior* **90**, 93-105. [[CrossRef](#)]
6. ARTURO ORMEÑO, KRISZTINA MOLNÁR. 2015. Using Survey Data of Inflation Expectations in the Estimation of Learning and Rational Expectations Models. *Journal of Money, Credit and Banking* **47**:4, 673. [[CrossRef](#)]
7. Christos Axioglou, Spyros Skouras. 2015. Asset pricing with flexible beliefs. *Studies in Nonlinear Dynamics & Econometrics* **19** . [[CrossRef](#)]
8. Fabio Milani. 2014. Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control* **47**, 94-114. [[CrossRef](#)]
9. Krisztina Molnár, Sergio Santoro. 2013. Optimal monetary policy when agents are learning. *European Economic Review* . [[CrossRef](#)]
10. Eric Gaus. 2013. Robust Stability of Monetary Policy Rules under Adaptive Learning. *Southern Economic Journal* **80**, 439-453. [[CrossRef](#)]
11. Kaushik Mitra, George W. Evans, Seppo Honkapohja. 2013. Policy change and learning in the RBC model. *Journal of Economic Dynamics and Control* **37**, 1947-1971. [[CrossRef](#)]
12. Paul De Grauwe, Agnieszka Markiewicz. 2013. Learning to forecast the exchange rate: Two competing approaches. *Journal of International Money and Finance* **32**, 42-76. [[CrossRef](#)]
13. Olena Kostyshyna. 2012. APPLICATION OF AN ADAPTIVE STEP-SIZE ALGORITHM IN MODELS OF HYPERINFLATION. *Macroeconomic Dynamics* **16**, 355-375. [[CrossRef](#)]
14. Fabio Milani The Modeling of Expectations in Empirical DSGE Models: A Survey 3-38. [[CrossRef](#)]
15. Francisco Buera, Gaston Navarro, Juan Pablo Nicolini. 2011. The Argentine Economy after Two Centuries. *Latin American Journal of Economics* **48**:10.7764/LAJE.48.2, 133-156. [[CrossRef](#)]
16. Sergey Slobodyan, Raf Wouters. 2011. Learning in an estimated medium-scale DSGE model. *Journal of Economic Dynamics and Control* . [[CrossRef](#)]
17. William A. Branch,, George W. Evans. 2011. Learning about Risk and Return: A Simple Model of Bubbles and Crashes. *American Economic Journal: Macroeconomics* **3**:3, 159-191. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
18. Alexandre Sokic. 2011. The Monetary Analysis of Hyperinflation and the Appropriate Specification of the Demand for Money. *German Economic Review* no-no. [[CrossRef](#)]
19. Klaus Adam, Albert Marcet. 2011. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory* **146**, 1224-1252. [[CrossRef](#)]
20. Michael Wegener, Frank Westerhoff. 2011. Evolutionary competition between prediction rules and the emergence of business cycles within Metzler's inventory model. *Journal of Evolutionary Economics* . [[CrossRef](#)]
21. STEFANO EUSEPI. 2010. Central Bank Communication and the Liquidity Trap. *Journal of Money, Credit and Banking* **42**:10.1111/jmcb.2010.42.issue-2-3, 373-397. [[CrossRef](#)]

22. George W. Evans, Seppo Honkapohja, Noah Williams. 2010. GENERALIZED STOCHASTIC GRADIENT LEARNING. *International Economic Review* 51:10.1111/iere.2010.51.issue-1, 237-262. [[CrossRef](#)]
23. Lars Peter Hansen, Thomas J. Sargent. Wanting Robustness in Macroeconomics 1097-1157. [[CrossRef](#)]
24. Vitor Gaspar, Frank Smets, David Vestin. Inflation Expectations, Adaptive Learning and Optimal Monetary Policy 1055-1095. [[CrossRef](#)]
25. Ching-Hsue Cheng, Liang-Ying Wei. 2009. Data spread-based entropy clustering method using adaptive learning. *Expert Systems with Applications* 36, 12357-12361. [[CrossRef](#)]
26. George W. Evans, Seppo Honkapohja, Kaushik Mitra. 2009. Anticipated fiscal policy and adaptive learning#. *Journal of Monetary Economics* 56, 930-953. [[CrossRef](#)]
27. George W. Evans, Seppo Honkapohja. 2009. Learning and Macroeconomics. *Annual Review of Economics* 1, 421-449. [[CrossRef](#)]
28. George A. Waters. 2009. LEARNING, COMMITMENT, AND MONETARY POLICY. *Macroeconomic Dynamics* 13, 421. [[CrossRef](#)]
29. FABIO MILANI. 2009. Adaptive Learning and Macroeconomic Inertia in the Euro Area. *JCMS: Journal of Common Market Studies* 47:10.1111/jcms.2009.47.issue-3, 579-599. [[CrossRef](#)]
30. Peter Bernholz, Peter Kugler. 2009. The Success of Currency Reforms to End Great Inflations: An Empirical Analysis of 34 High Inflations. *German Economic Review* 10:10.1111/geer.2009.10.issue-2, 165-175. [[CrossRef](#)]
31. Ching-Hsue Cheng, Liang-Ying Wei. 2009. Volatility model based on multi-stock index for TAIEX forecasting. *Expert Systems with Applications* 36, 6187-6191. [[CrossRef](#)]
32. Ricardo Nunes. 2009. LEARNING THE INFLATION TARGET. *Macroeconomic Dynamics* 13, 167. [[CrossRef](#)]
33. Kevin X.D. Huang, Zheng Liu, Tao Zha. 2009. Learning, Adaptive Expectations and Technology Shocks. *The Economic Journal* 119:10.1111/eoj.2009.119.issue-536, 377-405. [[CrossRef](#)]
34. Roc Armenter, Martin Bodenstein. 2008. CAN THE U.S. MONETARY POLICY FALL (AGAIN) IN AN EXPECTATION TRAP?. *Macroeconomic Dynamics* 12, 664. [[CrossRef](#)]
35. ROC ARMENTER. 2008. A General Theory (and Some Evidence) of Expectation Traps in Monetary Policy. *Journal of Money, Credit and Banking* 40:10.1111/jmcb.2008.40.issue-5, 867-895. [[CrossRef](#)]
36. IN-KOO CHO, KENNETH KASA. 2008. LEARNING DYNAMICS AND ENDOGENOUS CURRENCY CRISES. *Macroeconomic Dynamics* 12. . [[CrossRef](#)]
37. Avik Chakraborty, George W. Evans. 2008. Can perpetual learning explain the forward-premium puzzle?#. *Journal of Monetary Economics* 55, 477-490. [[CrossRef](#)]
38. Thomas J. Sargent. 2008. Evolution and Intelligent Design. *American Economic Review* 98:1, 5-37. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
39. Teck H. Ho, Xin Wang, Colin F. Camerer. 2008. Individual Differences in EWA Learning with Partial Payoff Information. *The Economic Journal* 118, 37-59. [[CrossRef](#)]
40. F MILANI. 2007. Expectations, learning and macroeconomic persistence#. *Journal of Monetary Economics* 54, 2065-2082. [[CrossRef](#)]
41. George A. Waters. 2007. Regime changes, learning and monetary policy. *Journal of Macroeconomics* 29, 255-282. [[CrossRef](#)]
42. K ADAM, G EVANS, S HONKAPOHJA. 2006. Are hyperinflation paths learnable?. *Journal of Economic Dynamics and Control* 30, 2725-2748. [[CrossRef](#)]

43. Fernando H. Barbosa, Alexandre B. Cunha, Elvia Mureb Sallum. 2006. Competitive equilibrium hyperinflation under rational expectations. *Economic Theory* **29**, 181-195. [[CrossRef](#)]
44. Gilles Saint-paul. 2005. Some Evolutionary Foundations for Price Level Rigidity. *American Economic Review* **95**:3, 765-779. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
45. KLAUS ADAM. 2005. LEARNING TO FORECAST AND CYCLICAL BEHAVIOR OF OUTPUT AND INFLATION. *Macroeconomic Dynamics* **9**. . [[CrossRef](#)]
46. Max Gillman, Anton Nakov. 2004. Granger causality of the inflation-growth mirror in accession countries*. *The Economics of Transition* **12**:10.1111/ecot.2004.12.issue-4, 653-681. [[CrossRef](#)]