Policy during Hyperinflations*

Manuel Macera† Albert Marcet‡ Juan Pablo Nicolini§

February 2024

Abstract

This paper performs policy analysis in a macroeconomic model in which agents do not have rational expectations. Since this deviation is, arguably, more relevant in a changing environment, we choose as a laboratory a model of hyperinflations. We focus on policies that fix the exchange rate as a way to reduce inflation and stabilize the economy. To be successful, these policies may require access to external funds once implemented. We compute the external financing requirements and the welfare gains of alternative policies that differ in when to intervene and how fast to bring inflation down. Our results show that policy makers face a trade-off: early and fast interventions are preferable from a welfare perspective, but they require more access to external funds. Hence, if external financing is an issue, one may have to adopt gradual policies, but ideally ones that bring inflation down in two or three quarters.

*The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Marcet and Nicolini acknowledge partial support from SGR (Generalitat de Catalunya), from the European Research Council Advanced Grant Agreement No. 324048-APMPAL, and the MACFINROBODS grant.

†Universidad Torcuato Di Tella.

‡ICREA, IAE, MOVE, and Universitat Autonoma de Barcelona.

§Federal Reserve Bank of Minneapolis and Universidad Torcuato Di Tella.

1 Introduction

The goal of this paper is to perform policy analysis in a macroeconomic model in which agents do not have rational expectations (RE). Considering deviations from RE is particularly relevant in changing environments, when agents are faced with events they have not seen before. Precisely because of this, we choose as a laboratory a monetary model in which hyperinflation occurs in equilibrium. These episodes represent sudden burst of inflation following a relatively long period of high but moderate and stable inflation, being probably the most dramatic case of a changing environment that can be found on macroeconomic data.

An intended advantage of studying these episodes is that the monetary policy actions are clearly visible and identifiable in the data. The magnitudes are extreme, so even if other factors may be changing at the same time, the policy actions during these events can be taken to be almost natural experiments. Along these lines, in the conclusions of his seminal paper "The Ends of Four Big Inflations" Sargent (1982), Sargent writes:

> *"I have encountered the view that the events described here are so extreme and bizarre that they do not bear on the subject of inflation in the contemporary United States. On the contrary, it is precisely because the events were so extreme that they are relevant. The four incidents we have studied are akin to laboratory experiments in which the elemental forces that cause and can be used to stop inflation are easiest to spot."*

Therefore, there is little doubt that, during these events, monetary policy acts in the data in a similar fashion as it will in the model. That is, even if other things may vary at the same time, the magnitude of the policy changes and their effect on the evolution of inflation are so dramatic that they clearly stand out as the dominant force. Thus, it appears as the natural laboratory to discuss a methodology to study the welfare effect of alternative monetary policies in models without rational expectations.

We develop our analysis in a model in which economic agents form their expectations through learning. The key feature of the model is that agents face uncertainty regarding the value of macroeconomic variables that are relevant to their decision-making process, but use optimally the information they have available to infer their values. The model that we use has the virtue that it is the one that best describes hyperinflation episodes in the past and it has been shown to be able to replicate the data remarkably well (see Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2009)). In addition, there is a clear sense in which it encompasses rational expectations as a special case, for when the uncertainty about aggregate variables tends to zero, the solution under learning coincides with the rational expectations solution. This feature offers a precise metric to measure how much we depart from the rational expectations benchmark.

The model implies that policies that fix the exchange rate can substantially reduce inflation rates and stabilize the economy, as long as external financing is available. We compare policies that differ in terms of when and how fast to intervene the economy, and quantify both the welfare gains from stabilization and the cost in terms of the financing needs required for the success of the plan. We show that policy-makers face a trade-off since early and fast interventions are preferable from a welfare perspective, but they require more access to external funds. Hence, if external financing is an issue, one may have to adopt gradual policies, but ideally ones that bring inflation down in two or three quarters.

In our view, the usefulness of our approach lies in the fact that a wide array of policies have been implemented in past hyperinflation episodes with different degrees of success. And while most countries in the world safely and rightly consider themselves immune to these problems, some countries are not.[1] In this context, the question of how to optimally run monetary and exchange rate policy is crucial and this paper is a first attempt at providing an answer.

While we are aware that our approach brings about several challenges, we consider our methodology for policy evaluation to be as strict as under rational expectations. We provide explicit micro-foundations for how the economy works and are explicit about the belief system that allows agents to form their expectations within the model. In addition, the proposed belief system is consistent not only with empirical evidence and surveys, but also with the model outcome in the sense that agents would be unable to statistically reject their it within a few periods using data generated by the model. Thus, an additional contribution of this paper is to spell out a methodology to perform policy analysis under learning with as much discipline as under rational expectations.

[1] Specifically, both Argentina and Venezuela are now running fiscal deficits that have been recently financed by printing money, facing inflation rates that are not that different from ones that lead to hyperinflation in the past. One could also argue that if highly indebted European countries leave the Euro amid a debt crisis and with high deficits, they could also face the same problem. However, that event appears to have a very low probability nowadays.

Related Literature Policy evaluation without rational expectations is very rarely addressed in the macroeconomics literature. Quantitative optimal policy problems typically perform a wide range of robustness checks: alternative values for key parameters for which some uncertainty exists, different market structures or various functions that describe the relationship between the policy instrument and the observables that are relevant for the problem at hand are just a few examples. However, there is almost no analysis of the robustness of different policies to the alternative ways of specifying agent's expectations. Notable exceptions are, for example, Vimercati, Eichenbaum, and Guerreiro (2021) and Eusepi, Gibbs, and Preston (2022).

Relative to Marcet and Nicolini (2003), we relax two assumptions. First the fact that deficits are indeed persistent. Second, we measure the change in reserves.

Also mention Adam, Marcet, and Nicolini (2016).

2 Framework

The model consists of three equations and an assumption about how expectations are formed. The first equation is the demand for real balances:

$$\frac{M_t^d}{P_t} = \phi\left(1 - \gamma \pi_{t+1}^e\right), \tag{1}$$

where $\pi_{t+1}^e = \mathbb{E}_t\left[\frac{P_{t+1}}{P_t}\right]$ denotes the expected gross inflation rate at period t.[2] The second equation is the government budget constraint:

$$M_t^s - M_{t-1}^s = d_t P_t - s_t F_t, \tag{2}$$

where M_t^s is the money supply, d_t is the real primary deficit, F_t represents external financing, and the dummy variable $s_t \in \{0, 1\}$ indicates whether it is available ($s_t = 1$) or not ($s_t = 0$). In what follows, we first analyze the behavior of the model when $s_t = 0$, which means that the primary deficit is financed entirely through seigniorage. The third equation specifies the evolution of the primary deficit d_t over time:

$$d_t = (1-\rho)\delta + \rho d_{t-1} + \epsilon_t, \tag{3}$$

where ϵ_t denotes an *i.i.d.* perturbation term. Equation (3) generalizes Marcet and Nicolini (2003) in that it introduces positive serial correlation in the deficit process, $\rho > 0$. This feature

[2] This equation can be rationalized as part of the equilibrium conditions of an overlapping generations model monetary model of a small open economy. We briefly describe such model in Section 4.2 and refer the reader to Marcet and Nicolini (2003) for details.

is important since deficits are in fact highly serially correlated in the data. We argue below that a proper calibration of this process is crucial for our results.

2.1 The evolution of inflation

Combining the money demand equation (1) and the government budget constraint (2) delivers

$$\pi_t = \frac{\phi - \phi\gamma\pi_t^e}{\phi - \phi\gamma\pi_{t+1}^e - d_t}, \tag{4}$$

where $\pi_t \equiv P_t/P_{t-1}$ denotes the realized gross inflation rate. This equation governs the evolution of inflation in any equilibrium, regardless of how expectations are formed, and we will use it repeatedly. We start by studying the rational expectations benchmark.

2.2 Rational Expectations in the Deterministic Case

In the absence of uncertainty, imposing rational expectations amounts to require that $\pi_t^e = \pi_t$ for all t. Plugging this condition into the main equation (4) and rearranging delivers:

$$\pi_{t+1} = (1-\rho)\left(\frac{\phi + \phi\gamma - \delta}{\phi\gamma} - \frac{1}{\gamma\pi_t}\right) + \rho\left(\frac{\phi + \phi\gamma - d_{t-1}}{\phi\gamma} - \frac{1}{\gamma\pi_t}\right). \tag{5}$$

This equation governs the dynamics of inflation under rational expectations. When the deficit starts at its long run value, e.g., $d_0 = \delta$, the equilibrium is stationary. In such a case, (5) admits two stationary inflation rates, which we denote $\{\boldsymbol{\pi_1}(\delta), \boldsymbol{\pi_2}(\delta)\}$ and are obtained as the solutions to the following quadratic equation:

$$\phi\gamma\pi^2 - (\phi + \phi\gamma - \delta)\pi + \phi = 0. \tag{6}$$

As δ varies, $\{\boldsymbol{\pi_1}(\delta), \boldsymbol{\pi_2}(\delta)\}$ trace out a stationary Laffer curve, which depicts the inflation rates that allow the government to finance the long run deficit δ. From now on, we adopt the convention that $\boldsymbol{\pi_1}(\delta)$ denotes the smallest root of (6) and thus corresponds to the inflation rate in the "good" side of the Laffer curve. To ensure that the stationary inflation rates are always positive and well defined we must assume the following.

Assumption 1 $\delta \in \mathbf{D} \equiv [0, \phi(1 + \gamma - 2\gamma^{\frac{1}{2}}))$.

The upper bound of the interval $\mathbf{D}$ indicates the maximum level of primary deficit that the government can finance through seigniorage, given the primitives of the economy.

When the initial deficit d_0 differs from its long run value δ, the evolution of inflation and deficit is determined by (3) and (5). Clearly, since the environment is deterministic, d_t always reverts to its long-run mean δ. The following proposition characterizes the behavior of inflation conditional on the initial deficit d_0 under rational expectations.

Proposition 1 *Under Assumption 1, for any $d_0 \in \mathbf{D}$ there exists $\underline{\pi}(d_0)$ such that:*

1. *If $\pi_0 < \underline{\pi}(d_0)$, then $\lim_{t\to\infty} \pi_t = -\infty$.*
2. *If $\pi_0 = \underline{\pi}(d_0)$, then $\lim_{t\to\infty} \pi_t = \boldsymbol{\pi}_1(\delta)$.*
3. *If $\pi_0 > \underline{\pi}(d_0)$, then $\lim_{t\to\infty} \pi_t = \boldsymbol{\pi}_2(\delta)$.*

The proof is relegated to Appendix A. The proposition applies also when the deficit starts at its long run value, $d_0 = \delta$. In such a case, $\underline{\pi}(d_0) = \boldsymbol{\pi}_1(\delta)$ and the set of equilibria is equivalent to that corresponding to the case with no persistence, as analyzed by Sargent and Wallace (1987b) or Marcet and Nicolini (2003). In general, when d_0 differs from δ, there will be a stable inflation path that converges to the low inflation steady state and a continuum of paths that converge to the high inflation steady state.[3] Under rational expectations, the model exhibits equilibrium multiplicity, with one equilibrium driven by fundamentals and a continuum of equilibria with self-fulfilling hyperinflations. The fundamentals equilibrium, which is always on the good side of the Laffer curve, is locally unique in the sense that all other equilibria converge to the high inflation steady state, which is on the wrong side of the Laffer curve.

2.3 Rational Expectations in the Stochastic Case

To characterize the evolution of inflation under rational expectations in the stochastic case, we linearize the main equation (4) around the low inflation steady state $\boldsymbol{\pi}_1(\delta)$ and introduce a small amount of uncertainty in the deficit process. The linearization boils down to:

$$\hat{\pi}_t = \frac{\delta}{\phi - \phi\gamma\boldsymbol{\pi}_1(\delta) - \delta}\hat{d}_t \tag{7}$$

$$\hat{d}_t = \rho\hat{d}_{t-1} + \epsilon_t, \tag{8}$$

where we are using the notation $\hat{x}_t = \ln x_t - \ln \mathbf{x}$, with bold letters indicating steady state values. Thus, when shocks are sufficiently small, we can express inflation as

$$\pi_t = (1-\rho)\boldsymbol{\pi} + \rho\pi_{t-1} + \nu_t, \tag{9}$$

where $\nu_t \equiv \delta\epsilon_t/(\phi - \phi\gamma\boldsymbol{\pi}_1(\delta) - \delta)$. Hence, around the low inflation steady state, inflation behaves as an AR(1) process that inherits the persistence of the deficit process.[4]

[3] The equilibria characterized in this proposition for the case $d_0 < \delta$ is depicted in Figure 3 in Appendix E.

[4] Appendix E displays sample paths for both inflation and primary deficit according to this linearized system.

2.4 Optimal Learning in the Stochastic Case

We now relax the assumption of rational expectations, which is that agents have perfect knowledge about the function that maps the realizations of the deficit d_t to realized inflation π_t, which is captured in (4). Instead, we assume they hold subjective beliefs about how inflation evolves over time and they use optimally all the information they have available to update their beliefs and form their expectations.

Specifically, we assume that agents believe that the inflation process behaves as follows:

$$\pi_t = (1 - \rho_\pi)\, \pi_t^\star + \rho_\pi \pi_{t-1} + u_t \tag{10a}$$

$$\pi_t^\star = \pi_{t-1}^\star + \eta_t, \tag{10b}$$

where $u_t \sim N(0, \sigma_u^2)$ and $\eta_t \sim N(0, \sigma_\eta^2)$ are *i.i.d.* and independent of deficit d_t. Importantly, agents observe the realizations of inflation π_t, but not those of $\pi_t^\star$ and u_t separately.[5] The intuition behind the belief system (10) is that agents think the behavior of inflation is similar to that of the primary deficit.[6] Therefore, they think inflation behaves as an AR(1) process, although they are unsure about its long-run value $\pi_t^\star$, and they express their uncertainty about it by modeling $\pi_t^\star$ as a unit root process.

Let $\beta_t = \mathbb{E}(\pi_t^\star \mid \pi^{t-1})$ denote the prior belief, entering period t, about long run inflation. Then, the belief system implies that expected inflation evolves as follows:

$$\pi_t^e = (1 - \rho_\pi)\beta_t + \rho_\pi \pi_{t-1}. \tag{11}$$

Optimal learning involves using Bayes' inference to update β_t using realized inflation π_t. When the initial prior $\beta_0 \equiv \mathbb{E}(\pi_0^\star)$ is normally distributed, with variance $\sigma_0^2 \equiv \mathbb{E}(\pi_0^\star - \beta_0)^2$, this implies that upon observing π_t, the belief is updated as follows:

$$\beta_{t+1} = \beta_t + \frac{1}{\alpha_t}\left(\frac{\pi_t - \rho_\pi \pi_{t-1}}{1 - \rho_\pi} - \beta_t\right), \tag{12}$$

where α_t denotes the optimal Kalman gain, which in turn evolves as follows:

$$\alpha_t = 1 + \frac{\alpha_{t-1}}{1 + \alpha_{t-1}\frac{\sigma_\eta^2}{\sigma_u^2}}, \tag{13}$$

[5] In other words, agents can't tell apart transitory shocks u_t from permanent ones η_t.

[6] In principle, one could allow for inflation and primary deficit having different persistence parameters. Under rational expectations, assuming these are the same is natural since the persistence of inflation inherits that of the deficit process. While this is no longer necessarily true with learning, because of the feedback between inflation and inflation expectations, in Appendix C we show it is a remarkably good approximation, especially when persistence is high.

with initial condition $\alpha_0 = (\sigma_\eta^2 + \sigma_0^2)/(\sigma_u^2 + \sigma_\eta^2 + \sigma_0^2)$. It is possible to set the variance of the initial prior σ_0^2 to guarantee that this gain is constant and, in the rest of the paper, we focus in this case.[7] This implies, according to (13), that there is a one to one mapping between the signal to noise ratio σ_η^2/σ_u^2 and the inverse of the Kalman gain $1/\alpha$. Therefore, we can think of $1/\alpha$ as an indicator of the degree of uncertainty that agents hold about long run inflation, with a smaller $1/\alpha$ indicating less uncertainty.

2.5 Inflation Dynamics under Learning

There is a clear sense in which the learning setup encompasses the rational expectations equilibrium as a special case. To see this, suppose that agents are certain about the value of long run inflation and believe that it corresponds to the low inflation steady state. This amounts to assume $1/\alpha = 0$, $\beta_0 = \boldsymbol{\pi}_1(\delta)$, and $\sigma_0^2 = 0$. In this case, beliefs remain constant at $\boldsymbol{\pi}_1(\delta)$, and the belief system (10) coincides with the dynamics of inflation under rational expectations shown in (9), which implies that the evolution of inflation is the same in both cases. This is true regardless of how persistent the deficit process is.

Introducing small amounts of uncertainty about long run inflation amounts to increase $1/\alpha$. Importantly, this uncertainty is meant to be interpreted as due to the agents making choices in an unfamiliar environment. In a learning equilibrium, the evolution of inflation is still governed by (4). Using (11) we can write:

$$\pi_t = \frac{\phi - \phi\gamma((1-\rho_\pi)\,\beta_{t-1} + \rho_\pi \pi_{t-2})}{\phi - \phi\gamma((1-\rho_\pi)\,\beta_t + \rho_\pi \pi_{t-1}) - d_t}. \tag{14}$$

This expression makes clear that beliefs affect not only expected inflation but also realized inflation. Moreover, since agents use realized inflation to update their beliefs, there is a feedback between realized and expected inflation. To provide intuition about the behavior of inflation in this case, let us write (14) as

$$H(\beta_t, \beta_{t-1}, \pi_t, \pi_{t-1}, \pi_{t-2}, d_t) = 0$$

[7] For simplicity we set the variance of the prior to be equal to the steady state Kalman filter uncertainty about β_t, which is given by

$$\sigma_0^2 = \frac{-\sigma_\eta^2 + \sqrt{\left(\sigma_\eta^2\right)^2 + 4\sigma_\eta^2\sigma_u^2}}{2}.$$

Therefore, the value of α_t is constant and given by

$$\frac{1}{\alpha} = \frac{\sigma_\eta^2 + \sigma_0^2}{\sigma_u^2 + \sigma_\eta^2 + \sigma_0^2}.$$

and define $h(\beta, \pi, d) \equiv H(\beta, \beta, \pi, \pi, \pi, d)$. The function h is useful to provide an approximation of the behavior of inflation in a situation in which $d_t = \delta$, $\beta_t \approx \beta_{t-1}$, and $\pi_t \approx \pi_{t-1} \approx \pi_{t-2} \approx \beta_{t-1}$. In such a case, (14) boils down to the quadratic equation (6), which implies that the rational expectations stationary inflation rates are also stationary inflation rates under learning. Moreover, one can show the low inflation rational expectations equilibrium is stable under learning but the high inflation one is not.[8]

These dynamics resemble those reported in Marcet and Sargent (1989b) and Marcet and Nicolini (2003) and are attractive in explaining periods of relatively large but stable inflation followed by a rapid burst in inflation. The reason is that as long as expected inflation is around the stable low inflation equilibrium $\boldsymbol{\pi}_1(\delta)$, realized inflation will remain in that region. However, if a sequence of positive shocks to the primary deficit drive inflation expectations beyond the value of $\boldsymbol{\pi}_2(\delta)$, the unstable dynamics on the wrong side of the Laffer curve take control and inflation can grow very quickly.

2.6 The Exchange Rate Regime

We introduce policy in the form of a stabilization plan based on *exchange rate rules* (ERR) when inflation gets out of hand. The role of the ERR is to stop the inflation dynamics under learning and bring inflation back down to the stable region. In particular, we assume that when inflation goes beyond a certain threshold β^U, the government establishes a crawling peg, and asks for special short-run financial assistance, setting $s_t = 1$ in (2).[9] While this policy is enforced, inflation is determined by the crawling peg and equation (2) determines the financing needs F_t to keep inflation under control. The government enforces this regime so as to attain an inflation target $\overline{\beta}$ in exactly T periods, and it lifts interventions once inflation falls below $\hat{\beta} < \beta^U$.

[8] However, notice that the stationary inflation rate $\boldsymbol{\pi_i}(\delta)$ is stable under learning only if

$$\frac{\partial \pi}{\partial \beta}\Big|_{\beta=\boldsymbol{\pi_i}(\delta)} = -\frac{\partial h/\partial \beta}{\partial h/\partial \pi}\Big|_{\beta=\boldsymbol{\pi_i}(\delta)} = \frac{\phi\gamma\boldsymbol{\pi_i}(\delta) - \phi\gamma}{\phi - \phi\gamma\boldsymbol{\pi_i}(\delta) - \delta} < 1, \tag{15}$$

as long as the denominator in (15) is positive. Whenever there is a positive price that clears the money market, the denominator in (15) is positive. We can always extend the model to include the case in which there are reserves that can be depleted to ensure the existence of such a price, in the spirit of Marcet and Nicolini (2003). We can verify that this is indeed the case if and only if

$$\boldsymbol{\pi_i}(\delta) < \frac{\phi + \phi\gamma - \delta}{2\phi\gamma}. \tag{16}$$

Using (6), it is easy to show that this condition is satisfied by the smallest root $\boldsymbol{\pi}_1(\delta)$, but not by the largest $\boldsymbol{\pi}_2(\delta)$.

[9] ERR could also be triggered to restore equilibrium. The two cases in which this could happen are if the money demand becomes negative or if, given the realization of the deficit process, the money demand is too low for an equilibrium to exist.

3 Quantitative Performance

In this section we calibrate and simulate the model. We show how likely hyperinflations are in equilibrium, and emphasize the role of deficit persistence. These examples quantify the amplification effect that expectations can have on the equilibrium inflation rate.

3.1 Calibration

In the model, one period represents a quarter. The baseline parameterization is summarized in Table 1.

Deficit process and money demand We estimate an AR(1) process for Argentina's primary deficit as a percentage of GDP and calibrate the values for ρ and σ_ϵ using the results of this estimation. We also assume that austerity will eventually prevail, so we set the long-run value of the deficit, δ, to be zero.

The money demand parameters ϕ and γ determine the shape of the Laffer curve that relates the inflation rate to the amount of revenues it raises. We use those reported in Marcet and Nicolini (2003), which target the inflation rate that maximizes the stationary Laffer curve and the maximum seigniorage as a percentage of GDP for the case of Argentina during the eighties. These are key parameters, since the distance between the two stationary inflation rates determines the size of the stable set and therefore the likelihood of a hyperinflation to unravel.

Belief system We use the belief system specified in (10) and assume that the learning mechanism is the same whether or not the ERR is enforced in a given period. This amounts to assume no credibility in the stabilization process. In all cases below, we set the initial prior to the low inflation steady state ($\beta_0 = \boldsymbol{\pi}_1(\delta)$), and the initial uncertainty to guarantee that the Kalman gain $1/\alpha$ is stationary. These choices offer two advantages. First, the dynamics of inflation is not be affected by the initial conditions. Second, as discussed above, in the case in which $1/\alpha = 0$, the dynamics of inflation under learning coincide with the case of rational expectations.

The belief system also requires to set the value of the the inverse of the Kalman gain $1/\alpha$, which summarizes the uncertainty that private agents hold about long run inflation. Below, we consider values of $1/\alpha = \{0.01, 0.05, 0.10, 0.20\}$.

Exchange Rate Regime Policy must specify when and how fast to intervene, the inflation target while the regime is enforced, and when to lift interventions. In all cases below we use

Table 1: Parameters for Baseline Economy

Parameter	Symbol	Value
Long-run deficit	δ	0
Persistence of deficit	ρ	.93
SD of shocks to deficit	σ_ϵ	.01
Money demand parameter	ϕ	.36
Money demand parameter	γ	.39

the same target $\overline{\beta} = \boldsymbol{\pi}_1(\delta)$, and the same threshold to lift interventions $\hat{\beta} = (1 + \boldsymbol{\pi}_2(\delta))/2$. This choice implies that the ERR is in place until expectations fall back down to the stable set, so the dynamics imply that — absent a new series of negative shocks — the economy converges to the low inflation equilibrium. We then consider a class of policies parameterized by the the other two policy instruments, β^U and T, that indicate when and how fast to intervene.[10]

3.2 Simulation

The previous description indicates that model economies differ in the value of $1/\alpha$ and the ERR policy in place, indexed by β^U and T. For each economy, we simulate $N = 10,000$ paths of length $\tau = 60$ for the deficit process and compute realized and expected inflation, beliefs, demand for real balances, reserve accumulation, and an indicator variable for whether an ERR was triggered, assuming the data generating process corresponds to the learning equilibrium.

3.3 Probability of Hyperinflations

To emphasize that the proper modeling of the deficit process makes a big difference for the results, we first compare economies with different degrees of persistence in the deficit. In each case, the probability of experiencing n hyperinflations is estimated using the fraction of simulated paths for which there were exactly n hyperinflations.

Table 2 displays the results of this exercise, and each panel considers a different value for $1/\alpha$. Low values imply small departures from rational expectations since in the limit, when $1/\alpha \to 0$, beliefs remain constant at the low inflation steady state. In the first column we consider different degrees of deficit persistence including, in the last row of each panel, the estimated one of $\rho = 0.93$. The table indicates that when there is very little uncertainty about long run inflation – $1/\alpha$ is low–, there is virtually no hyperinflationary episodes; they only

[10]The choice of how fast to intervene as captured by the parameter T allows us to consider the trade-off between "shock" and "gradualism," using the language in the literature. It is also motivated by the experiences of many countries that chose a crawling peg with declining rates to smoothly lower inflation, as opposed to other experiences in which the exchange rate was fixed, setting a devaluation rate of zero, at the moment of switching to the ERR.

Table 2: Probability of n Hyperinflations with $(\beta^U, T) = (150\%, 1)$. The numbers reported correspond to model economies with the same ERR policy in place but that differ in their degree of uncertainty about long run inflation, as captured by $1/\alpha$, and in the persistence of the deficit given by ρ. The probability of experiencing n hyperinflations is estimated using the fraction of simulated paths for which there were exactly n hyperinflationary episodes.

$1/\alpha = 0.01\%$				
ρ	0	1	2	≥ 3
0.50	100.0	0.0	0.0	0.0
0.70	100.0	0.0	0.0	0.0
0.90	84.4	8.4	4.2	3.0
0.94	70.5	10.9	7.5	11.1
$1/\alpha = 0.05\%$				
ρ	0	1	2	≥ 3
0.50	100.0	0.0	0.0	0.0
0.70	99.8	0.1	0.0	0.0
0.90	65.9	16.9	8.7	8.5
0.94	54.8	15.6	10.9	18.6
$1/\alpha = 0.10\%$				
ρ	0	1	2	≥ 3
0.50	100.0	0.0	0.0	0.0
0.70	98.5	1.5	0.1	0.0
0.90	54.3	21.8	12.2	11.7
0.94	45.8	19.1	12.8	22.3
$1/\alpha = 0.20\%$				
ρ	0	1	2	≥ 3
0.50	100.0	0.0	0.0	0.0
0.70	91.8	7.6	0.6	0.0
0.90	41.4	27.4	15.8	15.5
0.94	36.3	22.7	15.8	25.2

occur when persistence is very large but are largely infrequent. However, as we introduce more uncertainty, it doesn't take much for hyperinflations to become recurrent events, especially when persistence is high. This suggests that while hyperinflations are more likely the more we depart from rational expectations, and these departures can be arbitrarily small when the deficit process is highly persistent. In fact, the table shows that with a persistence as low as $\rho = 0.5$, these episodes virtually disappear.

These results emphasize that positive persistence involves an important extension relative to Marcet and Nicolini (2003), which only considered i.i.d. shocks to the primary deficit. By allowing $\rho > 0$, we not only gain in realism, since shocks to the deficit are indeed very persistent, but also verify that properly calibrating its value turns out to be crucial for quantifying the risks of hyperinflation when the access to debt markets is limited.

4 Policy analysis

Hyperinflations are stopped by a policy regime switch that adopts a crawling peg. To do so, the government requires external financing to satisfy the budget constraint. We now show how large and persistent these funding requirements are according to the model. We first present the computations for the policy $\{\beta^U = 150\%, T = 1\}$. We then evaluate alternative policies that imply either an earlier intervention (lower value for β^U) or a more gradual intervention (larger values for T), or both.

4.1 Financing requirements to stop the hyperinflations

The cost of implementing ERRs consists of the resources required to satisfy the government budget constraint while the policy is in place. Once the regime is triggered, its enforcement requires depleting reserves or having access to external funds, as monetization of the deficit is not available. However, to the extent that the regime brings expected inflation down and the demand for real balances recovers, it is possible to collect seigniorage and use it to accumulate reserves or repay the loan.[11] We now quantify these financing needs for several model economies that differ in the parameters of their ERR policy.

Table 3 displays the results of this exercise for the case in which $\alpha = 0.20$. The first column contains the values we consider for the policy parameters β^U and T. Each row reports the cumulative balance of an account that is set to zero at the moment the ERR is triggered and pays a real interest rate equal to zero.[12] In all cases, we report the median value across all simulations. For example, the number -2.9 at the top of the second column means that in the quarter in which the ERR is adopted, the median external funds required to satisfy the government budget constraint is 2.9% of yearly GDP. The number 2.3% at the top of the third column implies that by the second quarter, the government has accumulated assets. The rest of the table can be read in a similar fashion. We report the cumulative balance in the account for one period more than the policy variable T.

A remarkable feature of Table 3 is that need for external financing is typically a temporary phenomenon. Typically, the government would be able to pay the debt and, when this occurs,

[11] Note that it is possible that a sequence of good shocks brings the deficit d_t into negative territory. In this case, (2) implies that the money stock should go down, which may generate a deflation. In our simulations, we assume that when $d_t < 0$, then $M_t = M_{t-1}$ so that net inflation is zero and the extra resources are accumulated in reserves at the Treasury.

[12] The period is a quarter, so a risk-free interest rate would be very close to zero. One could easily impute a non-zero interest rate, but given the magnitude of the numbers, the table would barely change.

Table 3: Cumulative Median Change in Reserves ($1/\alpha = 0.20$, in % of yearly GDP). The change in reserves corresponds to the resources needed to satisfy the government budget constraint (2). Negative changes occur when the ERR is in place and positive ones when the extra seigniorage coming from the recovery of the money demand is larger than the realized deficit. In all cases, we report the median value across all simulations.

ERR T	**Periods After Intervention**				
	1	2	3	4	5
$\beta^U = 150\%$					
1	−2.9	2.3	–	–	–
2	−1.7	−2.4	2.6	–	–
3	−1.4	−2.4	−2.4	−2.1	–
4	−1.4	−2.4	−3.0	−3.2	−3.0
$\beta^U = 100\%$					
1	−3.3	1.3	–	–	–
2	−2.1	−1.1	1.5	–	–
3	−1.9	−1.5	−0.7	0.0	–
4	−1.8	−1.7	−1.2	−0.9	−0.3

it actually ends up with additional resources. This may seem surprising, but hyperinflations are actually bad for tax purposes because they are the result of unstable dynamics that appear on the wrong side of the Laffer curve. Along these dynamics, the inflation tax is decreasing with inflation. At the same time, real money balances are shrinking. Once the ERR is put into place, both processes revert. In particular, real money balances grow substantially, which means that nominal money is growing at a rate higher than the inflation rate.[13]

An important message arises from Table 3: early interventions and shock policies are more demanding in terms of external financing. In the case of early interventions, this can be noted by comparing the negative balances in the first period after intervention when $\beta^U = 150\%$ with those corresponding to $\beta^U = 100\%$. In the case of shock policies, the comparison must be done across rows as we increase T and keep β^U fixed. These results seem to imply a trade-off. A policy that acts early and brings down inflation quickly is better, but it may call for larger external support. Note, however, that for higher values of T, the negative balances are similar, suggesting that if external financing is an issue, one may have to adopt gradual policies, but ideally ones that bring inflation down in two or three quarters.

In all the cases considered in the table, hyperinflations occur in equilibrium. To make a choice among these alternative policies we must be able to measure the welfare cost of each alternative. To this issue we turn next.

[13]Interestingly, this process of reserve accumulation has been a common feature in the stabilization plans in Latin America in the 1980s.

4.2 Welfare costs of hyperinflations

To measure the welfare cost of hyperinflations, we rely on a standard overlapping generations monetary model of a small open economy. This model provides a microfoundation for the money demand equation proposed in Section 2. We present here a brief description of the environment of such model and refer the reader to Marcet and Nicolini (2003) for details.

Consider an economy in which agents live for two periods. Their lifetime endowments are $(1, e)$, their utility function is logarithmic and their discount factor is b. The only instrument available to move resources over time is money. We assume that $e < b$, so that a monetary equilibrium exists. In this environment, a monetary policy that maintains a constant quantity of money so that equilibrium inflation is zero implements a Pareto-optimal allocation in which consumption is given by:

$$c = \frac{1+e}{1+b} \text{ and } x = b\frac{1+e}{1+b}.$$

On the other hand, the non-monetary equilibrium corresponds to autarky, in which consumption is given by:

$$c = 1 \text{ and } x = e.$$

We can measure the value of the monetary system as the minimum fraction of consumption that the agent would demand to be compensated for moving from the monetary economy to autarky. This corresponds to the value of Δ that satisfies:

$$\ln\frac{1+e}{1+b} + b\ln b\frac{1+e}{1+b} = \ln(1+\Delta) + b\ln(1+\Delta)e. \tag{17}$$

This value is bounded as long as $e > 0$. One can also show that it is decreasing on e, and it approaches zero as $e \to b$. Any equilibrium in which inflation is positive and bounded in some or all periods, will imply a utility for the different generations that is higher than in autarky. Therefore, the welfare cost of the hyperinflationary equilibria is bounded above by Δ. Thus, the choice of the parameters e and b determines the value to society of having a monetary system, which in itself puts a bound on the welfare costs of a monetary system that does not work perfectly. The parameters of the money demand in Table 1 map one-to-one to values for e and b. Using these values in (17) delivers a value of Δ of roughly 0.10, or 10% of total consumption.

We measure the welfare cost of hyperinflations by computing the compensating variation of eliminating all inflationary dynamics that arise from learning. Specifically, for any given realization of the deficit for 100 periods, we compute the utility attained in equilibrium when we set $1/\alpha = 0$, which corresponds to the RE equilibrium. Then, for the same realization, we compute the utility attained in equilibrium when we set $1/\alpha$ equal to 0.01, 0.05, 0.10, and 0.20. Then, for each value of $1/\alpha$, we compute the percentage of consumption that agents would be

Table 4: Consumption Equivalent Welfare Change (in %). ERR policy is $(\beta^U, T) = (150\%, 1)$. The two calibrations correspond to economies in which the value of the monetary system is 10% and 30%. In each case, the first column calculates the welfare change using the full sample path, and the second does it restricting attention to the 10 quarters that precede the first hyperinflation.

$1/\alpha$	Gains from Money: 10%		Gains from Money: 30%	
	Full	10q	Full	10q
0.20	0.8	1.6	1.1	2.7
0.10	0.5	1.5	0.9	2.5
0.05	0.4	1.4	0.8	2.3
0.01	0.2	1.3	0.5	2.2

willing to forgo under the RE equilibrium to avoid the inflation rates that arise for positive values of $1/\alpha$. We repeat this exercise for 10,000 different simulations and compute the average consumption change.[14] In each case, we perform two different computations. We calculate the welfare change using the full sample path and also restricting attention to the 10 periods that precede the first hyperinflation and subsequent adoption of the ERR.

The results appear in Table 4. The first column lists the values for $1/\alpha$. The second and third columns report the consumption equivalent welfare change for the 100 periods, and for the 10 periods leading to the first hyperinflation in each simulation, respectively. The first measure is the one standard in the literature, and as can be seen, the numbers are sizable. For example, when $1/\alpha = 0.20$, the cost of the hyperinflations is around 0.8% of consumption in each of the 100 periods. This corresponds to 8% of the total gain of having a monetary system. The second measure is higher by construction. Again, when $1/\alpha = 0.20$, the cost of the hyperinflations is around 1.6% of each quarter consumption, which amounts to 16% of the total value of a monetary system during those periods.

As mentioned above, the value of Δ that solves equation (17) for the calibration presented in Table 1 is around 10% of total consumption. This number appears rather low to us for modern economies. Thus, we repeat the computations for the same parameter values, except that we increase the endowment in the first period and reduce it in the second period while maintaining a constant output, such that the value of Δ that solves equation (17) now becomes 30% of total output. The resulting numbers are reported in columns 4 and 5 of Table 4. The numbers are substantially larger in this case.

In Table 4, the mandate of the ERR policy was to intervene when inflation was above the threshold $\beta^U = 150\%$ and to bring inflation down fast, in one period ($T = 1$). We now compute

[14] Notice that this procedure implies that we compute welfare using the true distribution of prices, rather than the distribution that agents believe is the true distribution.

Table 5: Consumption Equivalent Welfare Change (in %). Full sample comparison when the welfare gain from monetary system is 30% for model economies that differ in terms of the parameters of the ERR policy. Early interventions are associated to a lower β^U and Gradualism to higher T. The welfare gain is computed as the value of Δ that satisfies (17).

$1/\alpha$	$T=1$	$T=2$	$T=3$	$T=4$
$\underline{\beta^U = 150\%}$				
0.20	1.1	1.3	1.7	2.2
0.10	0.9	1.1	1.4	1.7
0.05	0.8	0.9	1.2	1.4
0.01	0.5	0.6	0.7	0.9
$\underline{\beta^U = 100\%}$				
0.20	0.8	1.0	1.0	1.2
0.10	0.6	0.7	0.7	0.9
0.05	0.5	0.6	0.6	0.7
0.01	0.2	0.3	0.3	0.3

the welfare cost for alternative policy parameters. In particular, we set $\beta^U = 100\%$ and allow for values of T all the way up to 4. This allows us to assess the convenience of an early intervention and to compare shock policies versus gradual ones. The results are reported in Table 5. The first row in the table considers alternative values for the policy parameter T. The first column in the table reports the values considered for β^U and $1/\alpha$. The first numbers in the second column correspond to numbers reported in the fourth column of Table 4. These results suggest that the costs of a gradual policy that takes four periods to bring down inflation are sizable. For $\beta^U = 150\%$ and $1/\alpha = 0.20$, the cost increases by at least 75% across the board. In addition, an early intervention translate into considerable welfare benefits. For instance, when $T = 1$ and $1/\alpha = 0.20$, adopting an ERR earlier decreases the welfare cost by roughly 30%, from 1.1% of consumption when $\beta^U = 150\%$, to 0.8% when $\beta^U = 100\%$.

We would like to emphasize that none of these results are qualitatively surprising: earlier interventions imply lower inflation rates, so welfare ought to be higher. Similarly, gradual policies imply higher equilibrium inflation rates, so welfare ought to be smaller. The value of these computations is that it provides magnitudes that can be compared with the costs of these ERRs: the external financing required to satisfy the government budget constraint while the ERR is in place. The welfare costs in Table 4 when computed as a fraction of total consumption in the 10 quarters leading to the hyperinflation and the ERR ranges from 1.6% of GDP to 2.7% depending on what we assume about the value of the monetary system. These numbers, accumulated over 10 quarters (2.5 years) represent between 4% and 6.8% of the GDP in one year. These welfare gains represent money on the table. To grab a good share of that, a very short loan of a smaller magnitude is required — a loan that can be paid in full very quickly and that in all cases leaves the government with some extra revenue. These ERRs appear to be very close to a free lunch.

5 Justifying the System of Beliefs

In the learning equilibrium inflation is a function of past deficit, as (14) indicates, but agents think that it evolves according to (10), which is obviously not the same. The proposed belief system satisfies, however, various desirable principles. First, there is a clear sense in which the deviation from RE is can be made arbitrarily small. Second, beliefs are close to the data behavior, and to the extent that the equilibrium outcome of our model reproduces the behavior of the data, we can expect that agents in the real world can hold this system of beliefs and that this will in fact render the system close to the model behavior.[15] Third, the system of beliefs is close to observed surveys of inflation expectations, which are conducted continuously in many countries.[16]

In addition to these criteria, we want to ensure that, given the system of beliefs, the equilibrium is such that the agents see their beliefs as a reasonable description of the inflation that they observe. More precisely, we want to check if agents that observe the model outcome are unable to statistically reject their belief system within a few periods. In this way, we can think of the considered system of beliefs as being consistent with the model of inflation that we, as economists, consider. To this task we turn next.

5.1 Testable Restrictions

We now focus on the conditions under which agents would question their belief system in a learning equilibrium. To do this, we consider the implications of the equilibrium conditions and the belief system for the vector $x_t = (e_t, d_t - \rho d_{t-1})$, where $e_t \equiv (\pi_t - \rho_\pi \pi_{t-1}) - (\pi_{t-1} - \rho_\pi \pi_{t-2})$, and we evaluate these implications using simulated data. The following proposition adapts the results in Adam et al. (2016). It lists a set of necessary and sufficient second order conditions for the statement that inflation and deficit data are indeed generated by the model.

Proposition 2 *Let d_t be AR(1) with innovation ϵ_t as in (3). There exists a belief system as the one described in Section 2.4 consistent with the autocovariance function of $\{x_t\}$ if and only*

[15] Various authors (see for example, Stock and Watson (2007)) have chosen a similar model to explain actual inflation in various countries. Although they often use a more involved model including time-varying volatility, it often has the main features of our system of beliefs, namely, serial correlation and a permanent shock to average inflation.

[16] Many authors (e.g., Roberts (1997), among others) have fit the model in (10) to inflation surveys.

if the following restrictions hold:

1. $\mathbb{E}[x_{t-i}e_t] = 0$ *for all* $i \geq 2$,

2. $\mathbb{E}[(\epsilon_t + \epsilon_{t-1})e_t] = 0$,

3. $\Sigma b^2 + \mathbb{E}[e_t e_{t-1}] < 0$,

4. $\mathbb{E}[e_t] = 0$,

where $\Sigma = \sigma_\epsilon^2$ *and* $b = \mathbb{E}[\epsilon_t e_t]$ *correspond to the coefficient of a regression of* e_t *on* ϵ_t *in population.*

The proof is presented in the Appendix D. We test these moment restrictions using the parameterization displayed in Table 1. If we find that these restrictions cannot be rejected in the samples we consider, we conclude that agents could be holding the system of beliefs for inflation as stated above in the model at hand. Now we provide tests for these restrictions.

5.2 Statistics

Restrictions 1, 2, and 4 represent first moment restrictions of the form $\mathbf{E}[y_t] = 0$ for $y_t = e_t q_t$, for various $q_t \in \mathbf{R}^n$. In order to test these restrictions, we estimate $\mathbf{E}[y_t]$ through its corresponding sample mean:

$$\frac{1}{T}\sum_{t=1}^{T} y_t.$$

Using standard arguments, the statistic

$$\hat{Q}_T = T\left(\frac{1}{T}\sum_{t=1}^{T} y_t\right)' \hat{S}^{-1}\left(\frac{1}{T}\sum_{t=1}^{T} y_t\right) \xrightarrow{d} \chi_n^2,$$

as $T \to \infty$ for some $\hat{S}$ that is a consistent estimator of the asymptotic variance of $\frac{1}{T}\sum_{t=1}^{T} y_t$. We use[17]

$$\hat{S} = T \cdot \mathbf{E}\left[(\bar{y}_t - \mathbf{y})(\bar{y}_t - \mathbf{y}))'\right] = \sum_{\nu=-\infty}^{\infty} \Gamma_\nu = \Gamma_{-1} + \Gamma_0 + \Gamma_1.$$

In order to test restriction 3, we use a one-sided test of the form $H_0 : \alpha < 0$, where α is set to satisfy

$$\mathbf{E}\left[(\epsilon_t b + e_{t-1})e_t - \alpha\right] = 0.$$

[17] Here we exploit the MA(1) property of e_t and use the fact that beyond the first lead and lag, these autocovariances matrices must be equal to zero.

GMM sets the estimate of b to the OLS coefficient of a regression of e_t on ϵ_t and the estimate of α precisely to $b'\Sigma b + E(e_{t-1}e_t)$.

5.3 Rejection Frequencies

The belief system can be evaluated by checking whether the rejection frequencies exceed a predetermined significance level. In what follows, we calculate rejection frequencies using the theoretical asymptotic distribution of $\hat{Q}_T$.[18] In the case of Restrictions 1, 2, and 4, asymptotic theory implies that $\hat{Q}_T \to \chi^2_n$ as the sample size increases. In testing restriction 1, we use three lags of each element of x_t, and we always include a constant term in the instrument vector q_t. Notice that by including a constant, restriction 4 is embedded in the joint hypothesis testing performed for restriction 1. In the case of restriction 3, the asymptotic properties of the GMM estimator of α imply that, under the null hypothesis, it will be normally distributed and centered at 0.

Rejection Frequencies using the Asymptotic Distribution. Table 6 displays the results of testing the restrictions of Proposition 1 using the asymptotic theoretical distribution of the statistics described in the previous paragraphs. Since restriction 1 required some discretionary choice regarding the set of instruments, we single out its results.

The results indicate that agents will find it difficult to reject their beliefs based on the observation of realized inflation in a span of 10 years (40 periods) when the signal-to-noise ratio of their beliefs is higher, since this implies a higher stationary $1/\alpha$, in this case, 0.05. On the contrary, the results show that for values closer to the RE equilibrium, the rejection frequencies are higher than 10%, particularly for restrictions 2 and 3. This emphasizes the notion that hyperinflations add a persistent component that has, because of the formation of expectations, a life on its own. Thus, for values of $1/\alpha$ that generate several hyperinflations, agents are less likely to reject the belief system, which makes hyperinflations themselves more likely.

5.4 Alternative Belief Systems

We end this section with a brief discussion of alternative ways to model the belief system. The discussion is motivated by the observation that embracing the internal rationality approach as we do in the paper implies abandoning the hope of equilibrium uniqueness. To do so, one needs to take a strong stand on a unique specification and parameterization of the belief system. Being able to use theory and data to uniquely pin down a value for $1/\alpha$, for instance, seems to us to be beyond the ability of current macroeconomics. We see the value of this approach as

[18] The results are very similar if one uses empirical small sample distributions.

Table 6: Rejection Frequencies at the 5% level for $(\beta^U, T) = (150\%, 1)$
This table reports rejection frequencies obtained from testing restrictions of Proposition 2 using simulated data. The set of instruments for restriction 1 includes a constant and three lags of $d_t - \rho d_{t-1}$ (1a), three lags of e_t (1b), or three lags of x_t (1c).

	T		
	40	60	100
$1/\alpha = 0.05$			
Restriction 1a	14.4 %	15.2 %	19.0 %
Restriction 1b	1.6 %	1.5 %	1.2 %
Restriction 1c	1.7 %	1.6 %	1.4 %
Restriction 2	17.7 %	19.9 %	21.7 %
Restriction 3	13.5 %	7.3 %	2.8 %
$1/\alpha = 0.20$			
Restriction 1a	6.8 %	6.6 %	6.0 %
Restriction 1b	0.7 %	0.7 %	0.6 %
Restriction 1c	1.0 %	0.9 %	0.7 %
Restriction 2	9.8 %	9.7 %	8.4 %
Restriction 3	9.9 %	5.3 %	2.4 %

finding classes of belief systems that deliver relatively similar outcomes and which cannot be rejected by the agents in the model. For example, the probability distribution over the number of hyperinflations changes when one changes $1/\alpha$ from 0.10 to 0.20, while the test results are good for both values. How can one use the model to look into the data and distinguish the two values? Our approach does not allow one to make that distinction, and the specific welfare computations do depend on the value of that parameter.

The belief system described in (10) can be extended along several dimensions. For instance, agents in the model never anticipate that in order to stop a hyperinflation, the government may impose an ERR policy. We can easily relax this assumption by assuming that agents place a positive probability to the event that an ERR is actually imposed. To entertain this possibility, we replace (10b) by the following:

$$\pi_t^\star = \begin{cases} \pi_{t-1}^\star + \eta_t & \text{if } z_t = 0 \\ \overline{\beta} + \upsilon_t & \text{if } z_t = 1 \end{cases} \tag{18}$$

where z_t indicates whether an ERR has been imposed and $\overline{\beta}$ is the ERR target. In addition, we assume that agents believe that, once inflation goes beyond a certain threshold, the probability

Table 7: Probability of n Hyperinflations with $(\beta^U, T) = (150\%, 1)$. The numbers reported correspond to model economies with the same ERR policy in place but that differ in their degree of uncertainty about long run inflation, as captured by $1/\alpha$. The probability of experiencing n hyperinflations is estimated using the fraction of simulated paths for which there were exactly n hyperinflationary episodes.

Belief system given by (10a) and (10b)				
$1/\alpha$	0	1	2	≥ 3
0.01	70.5	10.9	7.5	11.1
0.05	54.8	15.6	10.9	18.6
0.10	45.8	19.1	12.8	22.3
0.20	36.3	22.7	15.8	25.2
Belief system given by (10a), (18), and (19)				
$1/\alpha$	0	1	2	≥ 3
0.01	70.1	11.3	7.7	10.9
0.05	55.6	15.4	10.9	18.0
0.10	47.6	17.3	12.2	22.9
0.20	38.1	19.3	13.2	29.4

of an exchange rate regime becomes positive:

$$\Pr[z_{t+1} = 1] = \begin{cases} 0 & \text{if } \pi_t \leq \overline{\pi} \\ p^{ERR} & \text{if } \pi_t > \overline{\pi} \end{cases}. \tag{19}$$

Given a threshold $\overline{\pi}$, one can impose the additional restriction that the probability that agents use p^{ERR} is the actual probability in the model, given the parameters of the ERR policy. One way to assess whether this formulation delivers similar results as the original belief system is by comparing the frequency of hyperinflations under two alternatives; we do this in Table 7. While these numbers indicate this extended belief system makes hyperinflations less likely, the effect is rather small. We take this as an indication that this extension delivers similar results to the baseline belief system.

A natural further exploration is to allow agents to use data on the deficit in the signal extraction problem. As long as agents also use inflation data, the main results will not be affected, since the key driver of the model is the feedback between actual and perceived inflation. On the other hand, it is likely to improve the performance of the tests, since agents would now optimally use the deficit data, which in our current formulation they are ignoring. Overall, we conjecture that these or other extensions are unlikely to change the main message of the paper.

6 Conclusions

In this paper we perform policy analysis in a model in which agents do not have rational expectations. Since deviations from rationality are more relevant in changing environments, we use a model of hyperinflations and study the effect of policies that fix the exchange rate to bring inflation down and stabilize the economy. Our results show that intervening early and fast result in higher welfare, but might require a larger need for external funds.

References

Adam, K. and A. Marcet (2011): "Internal rationality, imperfect market knowledge and asset prices," *Journal of Economic Theory*, 146, 1224–1252.

Adam, K., A. Marcet, and J. Beutel (2017): "Stock Price Booms and Expected Capital Gains," *American Economic Review*, 107, 2352–2408.

Adam, K., A. Marcet, and J. P. Nicolini (2016): "Stock market volatility and learning," *Journal of Finance*, 71, 33–82.

Eusepi, S., C. G. Gibbs, and B. Preston (2022): "Monetary policy trade-offs at the zero lower bound," CAMA Working Papers 2022-26, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

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

Marcet, A. and J. P. Nicolini (2003): "Recurrent hyperinflations and learning," *American Economic Review*, 93, 1476–1498.

Marcet, A. and T. J. Sargent (1989a): "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48, 337–368.

——— (1989b): *Least-squares learning and the dynamics of hyperinflation*, Cambridge: Cambridge University Press, 119–137.

Nicolini, J. P. (1996): "Ruling out speculative hyperinflations: The role of the government," *Journal of Economic Dynamics and Control*, 20, 791–809.

Obstfeld, M. and K. Rogoff (1983): "Speculative hyperinflations in maximizing models: Can we rule them out?" *Journal of Political Economy*, 91, 675–687.

Roberts, J. M. (1997): "Is inflation sticky?" *Journal of Monetary Economics*, 39, 173–196.

Sargent, T. and N. Wallace (1987a): "Inflation and the Government Budget Constraint," in *Economic Policy in Theory and Practice*, ed. by A. Razin and E. Sadka, London: Palgrave Macmillan, 170–207.

Sargent, T., N. Williams, and T. Zha (2009): "The conquest of South American inflation," *Journal of Political Economy*, 117, 211–256.

Sargent, T. J. (1982): "The Ends of Four Big Inflations," in *Inflation: Causes and Effects*, National Bureau of Economic Research, Inc, NBER Chapters, 41–98.

Sargent, T. J. and N. Wallace (1987b): "Inflation and the government budget constraint," *Economic Policy in Theory and Practice*, 170–200.

Stock, J. H. and M. W. Watson (2007): "Why has U.S. inflation become harder to forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

Vimercati, R. B., M. S. Eichenbaum, and J. Guerreiro (2021): "Fiscal Stimulus with Imperfect Expectations: Spending vs. Tax Policy," Working Paper 29134, National Bureau of Economic Research.

A Proof of Proposition 1

If $d_0 = \delta$, it is straightforward that $\underline{\pi}(d_0) = \boldsymbol{\pi_1}(\delta)$. Hence, the goal is to show that $\underline{\pi}(d_0)$ exists when $d_0 \neq \delta$. The logic consists of showing that, given d_0, one can find a value for π_0 to be exactly at $\boldsymbol{\pi_1}(d_t)$ in period t, where $\boldsymbol{\pi_1}(d_t)$ denotes the solution to the quadratic equation (6) given the level of primary deficit d_t. Notice that such value is well defined for all d_t as long as $d_0 \in \mathbf{D}$.

Lemma 1 *If $d_0 \neq \delta$, there exists a monotone and bounded sequence of initial conditions $\{\beta_t\}$ such that for all t, $\pi_0 = \beta_t$ implies $\pi_t = \boldsymbol{\pi_1}(d_t)$.*

This result indicates that $\lim_{t\to\infty} \beta_t$ is well defined. We denote this limit by $\underline{\pi}(d_0)$, making explicit its dependence on the initial value of the primary deficit d_0. Furthermore, it also indicates that if $\pi_0 = \underline{\pi}(d_0)$, then $\lim_{t\to\infty} \pi_t = \lim_{t\to\infty} \boldsymbol{\pi_1}(d_t) = \boldsymbol{\pi_1}(\delta)$. Hence, the second statement of Proposition 1 can be viewed as a corollary of Lemma 1. Although the two cases $d_0 < \delta$ and $d_0 > \delta$ need to be considered separately, the proof is analogous, so we present only the one corresponding to $d_0 < \delta$.

Proof of Lemma 1. The proof is by induction. For the initial step, observe that

$$\pi_0 = \boldsymbol{\pi_1}(\delta) \implies \pi_1 = (1-\rho)\boldsymbol{\pi_1}(\delta) + \rho\mathbf{F}(\boldsymbol{\pi_1}(\delta), d_0) > \boldsymbol{\pi_1}(\delta) > \boldsymbol{\pi_1}(d_1),$$

where the first inequality follows from the following property about F:

$$\mathbf{F}(\pi, d) > \pi \iff \pi \in (\boldsymbol{\pi_1}(d), \boldsymbol{\pi_2}(d)),$$

and the second follows from the fact that $d < \delta$ implies $\boldsymbol{\pi_1}(d) < \boldsymbol{\pi_1}(\delta)$. On the other hand, we also have that

$$\pi_0 = \boldsymbol{\pi_1}(d_0) \implies \pi_1 = (1-\rho)\mathbf{F}(\boldsymbol{\pi_1}(d_0), \delta) + \rho\boldsymbol{\pi_1}(d_0) < \boldsymbol{\pi_1}(d_0) < \boldsymbol{\pi_1}(d_1).$$

Since $\mathbf{F}$ is continuous and monotone, it follows that there exists a unique $\beta_1 \in (\boldsymbol{\pi_1}(d_0), \boldsymbol{\pi_1}(\delta))$ such that if $\pi_0 = \beta_1$, then $\boldsymbol{\pi_1} = \boldsymbol{\pi_1}(d_1)$.

For the inductive step, suppose there exists β_t such that $\pi_0 = \beta_t$ implies $\pi_t = \boldsymbol{\pi_1}(d_t)$. Then it follows that

$$\pi_t = \boldsymbol{\pi_1}(d_t) \implies \pi_{t+1} = (1-\rho)\mathbf{F}(\boldsymbol{\pi_1}(d_t), d_t) + \rho\boldsymbol{\pi_1}(d_t) < \boldsymbol{\pi_1}(d_t) < \boldsymbol{\pi_1}(d_{t+1}),$$

and we again have that

$$\pi_t = \boldsymbol{\pi_1}(\delta) \implies \pi_{t+1} = (1-\rho)\boldsymbol{\pi_1}(\delta) + \rho\mathbf{F}(\boldsymbol{\pi_1}(\delta), d_t) > \boldsymbol{\pi_1}(\delta) > \boldsymbol{\pi_1}(d_{t+1}).$$

Hence, there exists $\beta_{t+1} \in (\beta_t, \boldsymbol{\pi_1}(\delta))$ such that $\pi_0 = \beta_{t+1}$ implies $\pi_{t+1} = \boldsymbol{\pi_1}(d_{t+1})$. This also implies that $\beta_{t+1} > \beta_t$ for all t, so the sequence is monotone and bounded. ■

To complete the proof of the proposition, take an arbitrary sequence $\{\pi_t\}$ that evolves according to (5) and suppose $\pi_0 > \underline{\pi}(d_0)$. Two cases need to be considered. First, if $\pi_0 \geq \max\{\boldsymbol{\pi_1}(\delta), \boldsymbol{\pi_1}(d_0)\}$, then (5) and the properties of F imply that statement 3 is satisfied. Second, if $\pi_0 \in (\underline{\pi}(d_0), \max\{\boldsymbol{\pi_1}(\delta), \boldsymbol{\pi_1}(d_0)\})$, then the proof of Lemma 1 indicates there exists a period t in which $\pi_t > \max\{\boldsymbol{\pi_1}(\delta), \boldsymbol{\pi_1}(d_t)\}$, and therefore (5) and the properties of F again tell us that statement 3 holds. The proof of the first statement, when $\pi_0 < \underline{\pi}(d_0)$, follows exactly the same steps.

B Linearization of Equation 4

To linearize equation (4), we treat π_t, π_t^e, and π_{t+1}^e as different variables. Thus, we obtain

$$\hat{\pi}_t = \frac{\delta}{\phi - \phi\gamma\boldsymbol{\pi}_1(\delta) - \delta}\hat{d}_t - \frac{\phi\gamma\boldsymbol{\pi}_1(\delta)}{\phi - \phi\gamma\boldsymbol{\pi}_1(\delta)}\hat{\pi}_t^e + \frac{\phi\gamma\boldsymbol{\pi}_1(\delta)}{\phi - \phi\gamma\boldsymbol{\pi}_1(\delta) - \delta}\hat{\pi}_{t+1}^e.$$

Rational expectations implies that, around the low inflation steady state, $\hat{\pi}_t^e = 0$ for all t. Therefore, the last two terms on the right-hand side cancel out, and we obtain the equation in the main text.

C Estimating Inflation Persistence

To calculate inflation paths, we use the following system:

$$\begin{aligned}
d_t &= (1-\rho)\delta + \rho d_{t-1} + \epsilon_t, && (20)\\
\pi_t &= \frac{\phi - \phi\gamma(\rho_\pi \pi_{t-2} + (1-\rho_\pi)\,\beta_{t-1})}{\phi - \phi\gamma(\rho_\pi \pi_{t-1} + (1-\rho_\pi)\,\beta_t) - d_t}, && (21)\\
\alpha_t &= 1 + \alpha_{t-1}, && (22)\\
\beta_t &= \beta_{t-1} + \frac{1}{\alpha}\left(\frac{\pi_{t-1} - \rho_\pi \pi_{t-2}}{1-\rho_\pi} - \beta_{t-1}\right), && (23)
\end{aligned}$$

with the initial condition $\{\pi_{-1}, d_0, \alpha_0, \beta_0\} = \{\boldsymbol{\pi}_1, \delta, \overline{\alpha}, \boldsymbol{\pi}_1\}$, where $\boldsymbol{\pi}_1$ is the low inflation steady state in the RE equilibrium, and $\overline{\alpha}$ is just a positive integer. We assume that the variance of

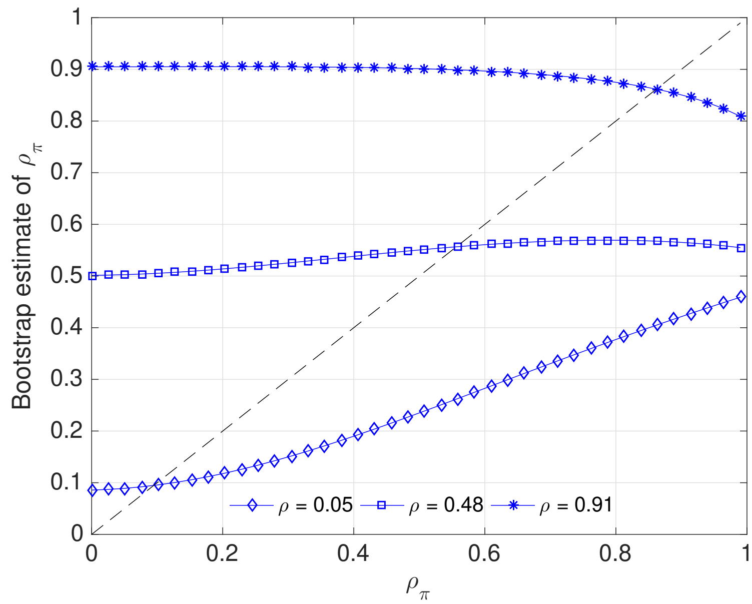

Figure 1: Inflation persistence: bootstrap estimate vs. true parameter: The points at which the solid lines cross the dashed line represent the fixed points $\rho_\pi = \hat{\rho}_\pi(N, T)$. For this simulation we set $N = 2,500$ and $T = 2,500$.

the i.i.d. shock ϵ_t is small enough so that the system remains stable.

We generate N samples of length T for each variable defined in (20)-(23). Let $\hat{\rho}_\pi(N, T)$ be the bootstrap estimate of the persistence of inflation, which is a function of all primitives of the model, including ρ_π and ρ. We allow $\rho_\pi \neq \rho_\delta$ and restrict attention to positive autocorrelation coefficients. Figure 1 displays $\hat{\rho}_\pi(N, T)$ as a function of ρ_π. The figure shows that for any ρ, there is a unique ρ_π such that $\rho_\pi = \hat{\rho}_\pi(N, T)$. We are interested in those fixed points because they correspond to the case in which agents are able to estimate ρ_π using past data. Since we proved that in an RE equilibrium, $\rho = \rho_\pi$, considering $\rho_\pi \neq \rho_\delta$ implies that we are allowing the persistence of inflation to be biased in a learning equilibrium. The nature of that bias is portrayed in Figure 2, which plots the fixed point $\rho_\pi^\star = \rho_\pi = \hat{\rho}_\pi(N, T)$ as a function of ρ. The figure indicates that the bias tends to be negative for low persistence of d_t and positive for high persistence of d_t.

D Proof of Proposition 2

The proof considers the general formulation in which agents use information about the innovations to the deficit to adjust their expectations regarding future inflation.

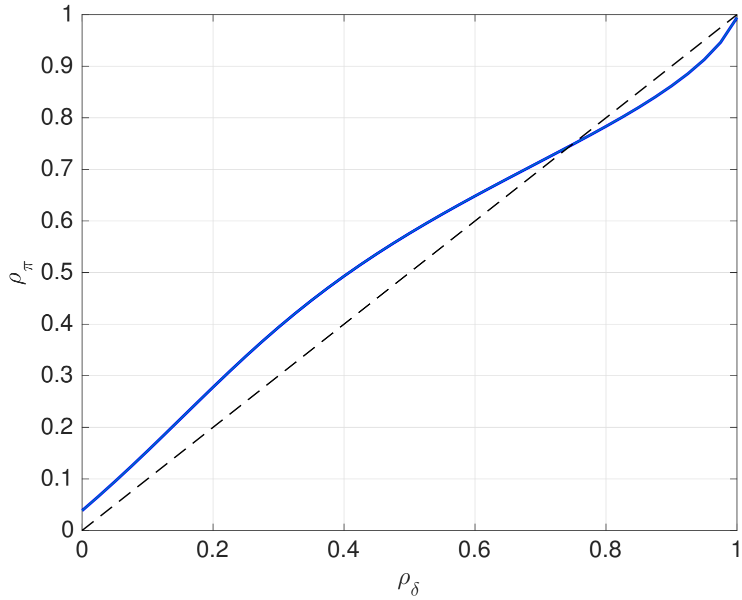

Figure 2: Inflation persistence vs. Deficit persistence: The solid line represents the fixed points portrayed in Figure 1 for different values of ρ. For this simulation we set $N = 2,500$ and $T = 2,500$.

Restriction 1. Note that, according to the belief system

$$\begin{aligned} e_t &= \pi_t^* + \psi\epsilon_t + u_t - \left(\pi_{t-1}^* + \psi\epsilon_{t-1} + u_{t-1}\right) \\ &= \eta_t + \psi\epsilon_t - \psi\epsilon_{t-1} + u_t - u_{t-1}, \end{aligned}$$

whereas

$$d_t - \rho d_{t-1} = (1-\rho)\delta + \epsilon_t,$$

so that Restriction 1 holds for $i \geq 2$.

Restriction 2. To prove Restriction 2, observe that

$$\mathbb{E}\left[(d_t - \rho d_{t-1})e_t\right] = \mathbb{E}\left[((1-\rho)\delta + \epsilon_t)(\eta_t + \psi\epsilon_t - \psi\epsilon_{t-1} + u_t - u_{t-1})\right] = \psi\sigma_\epsilon^2,$$

while

$$\mathbb{E}\left[(d_{t-1} - \rho d_{t-2})e_t\right] = \mathbb{E}\left[((1-\rho)\delta + \epsilon_{t-1})(\eta_t + \psi\epsilon_t - \psi\epsilon_{t-1} + u_t - u_{t-1})\right] = -\psi\sigma_\epsilon^2,$$

so that Restriction 2 also holds.

Restriction 3. Note that

$$\begin{aligned}\mathbb{E}\left[e_t e_{t-1}\right] &= \mathbb{E}\left[(\eta_t + \psi\epsilon_t - \psi\epsilon_{t-1} + u_t - u_{t-1})(\eta_{t-1} + \psi\epsilon_{t-1} - \psi\epsilon_{t-2} + u_{t-1} - u_{t-2})\right] \\ &= -\psi^2\sigma_\epsilon^2 - \sigma_u^2.\end{aligned}$$

Now consider the projection of $d_t - \rho d_{t-1}$ on e_t. The projection is given by

$$\frac{\mathrm{Cov}\left[e_t, d_t - \rho d_{t-1}\right]}{\mathrm{Var}\left[d_t - \rho d_{t-1}\right]}(d_t - \rho d_{t-1}),$$

where

$$\begin{aligned}\mathrm{Cov}\left[e_t, d_t - \rho d_{t-1}\right] &= \mathrm{Cov}\left[\eta_t + \psi\epsilon_t - \psi\epsilon_{t-1} + u_t - u_{t-1}, \epsilon_t\right] = \psi\sigma_\epsilon^2, \\ \mathrm{Var}\left[d_t - \rho d_{t-1}\right] &= \sigma_\epsilon^2,\end{aligned}$$

so that the projection is given by $\psi\epsilon_t$, and the variance of the projection is $\psi^2\sigma_\epsilon^2$. Plugging these results into the right-hand side of Restriction 3 delivers

$$\mathbb{E}\left[e_t, e_{t-1}\right] + b'^2\sigma_\epsilon^2 - \sigma_u^2 + \psi^2\sigma_\epsilon^2 = -\sigma_u^2 < 0.$$

Restriction 4. Since all the perturbation terms have zero expectation, Restriction 4 follows immediately.

E Inflation Persistence

The equilibria characterized in Proposition 1 for the case $d_0 < \delta$ are depicted in Figure 3. Figure 4 displays sample paths of both inflation and primary deficit according to the linearized system of equation 9.

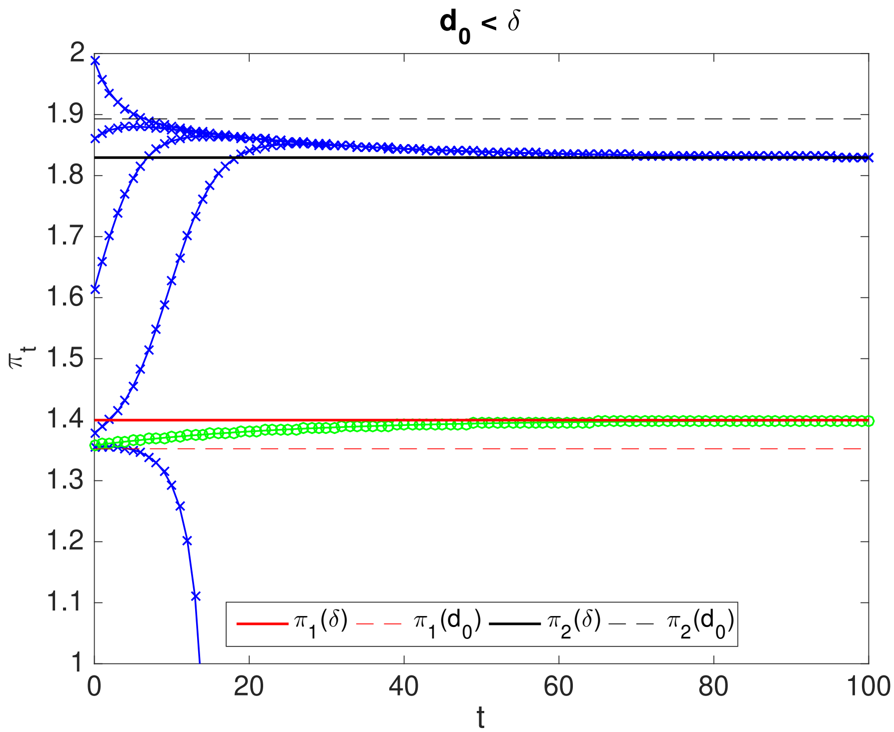

Figure 3: Inflation paths under rational expectations that evolve according to (5). The horizontal lines correspond to the solutions to the quadratic equation when $d = \delta$ (solid line) and $d = d_0$ (dashed line). The line with circles that converges to the low inflation steady state starts at $\pi_0 = \underline{\pi}(d_0)$, which is characterized in Proposition 1.

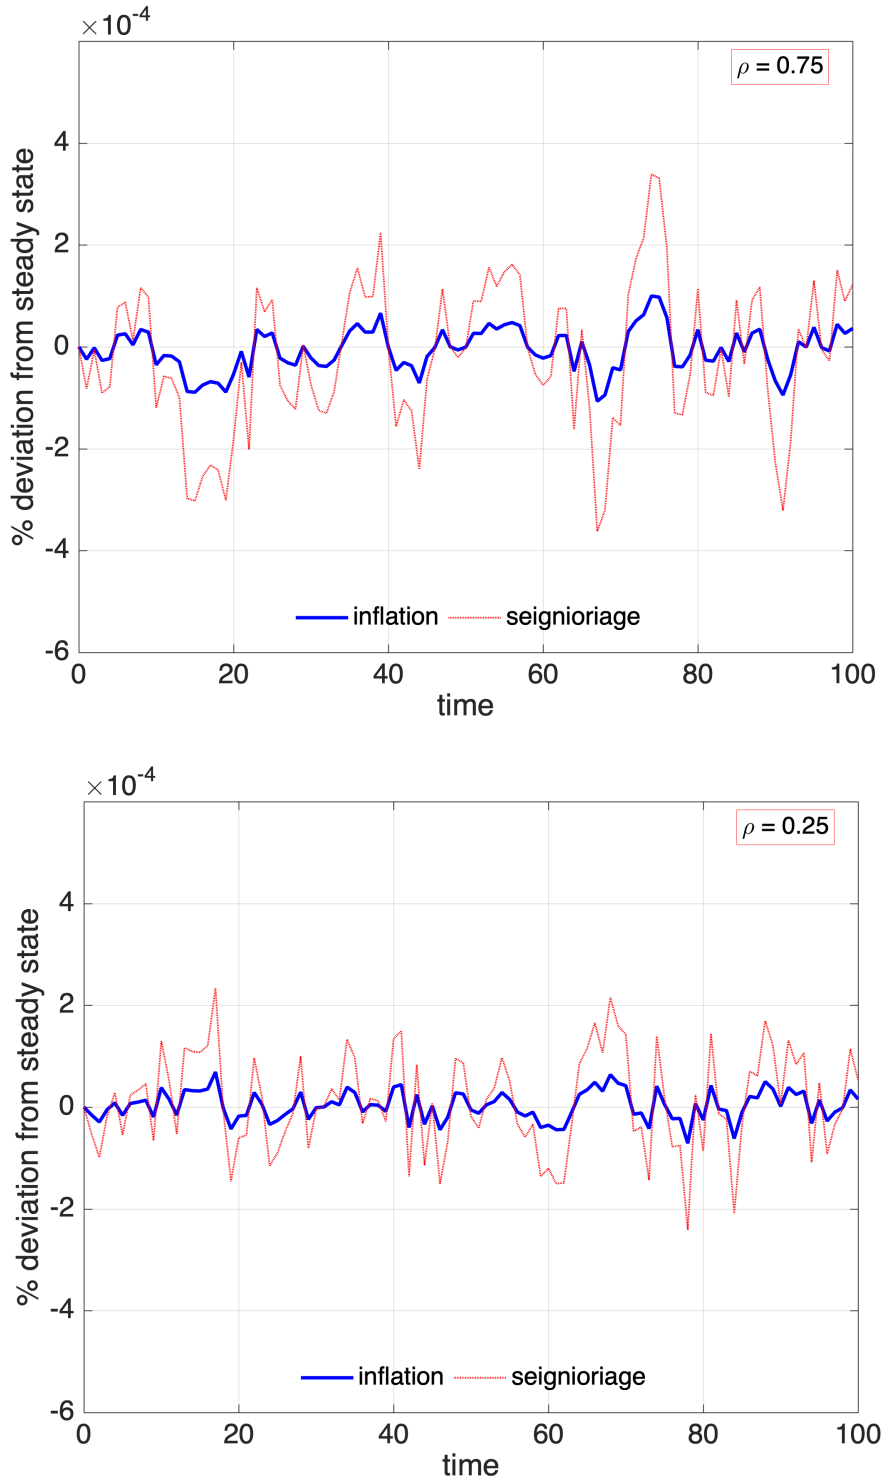

Figure 4: Sample paths for Inflation and Deficit around the low inflation steady state in the linearized rational expectations equilibrium.