

THE CONQUEST OF SOUTH AMERICAN INFLATION

THOMAS SARGENT, NOAH WILLIAMS, AND TAO ZHA

ABSTRACT. We infer determinants of Latin American hyperinflations and stabilizations by using the method of maximum likelihood to estimate a hidden Markov model that assigns roles both to fundamentals in the form of government deficits that are financed by money creation and to destabilizing expectations dynamics that can occasionally divorce inflation from fundamentals. Levels and conditional volatilities of monetized deficits drove most hyperinflations and stabilizations, with a notable exception in Peru where a cosmetic reform of the type emphasized by Marcet and Nicolini (2003) seems to have been at work.

Perhaps the simple rational expectations assumption is at fault here, for it is difficult to believe that economic agents in the hyperinflations understood the dynamic processes in which they were participating without undergoing some learning process that would be the equivalent of adaptive expectations.

Stanley Fischer, 1987

KEY WORDS: Fundamental and cosmetic policy reforms, likelihood, seigniorage, self-confirming equilibria, rational expectations, adaptation, escape dynamics.

Date: September 29, 2008.

This paper has benefited from help and encouragement of many people. We thank the referees and editors as well as Jeff Campbell, V.V. Chari, Chris Erceg, Eduardo Ganapolsky, Gary Hansen, Selahattin Imrohoroğlu, Karsten Jeske, Patrick Kehoe, Narayana Kocherlakota, David Levine, Federico Mandelman, Albert Marcet, Frederic Mishkin, Juan Pablo Nicolini, David Romer, Lee Ohanian, Monika Piazzesi, Christopher Sims, Martin Uribe, Dan Waggoner, and Michael Woodford for helpful discussions and comments. Sagiri Kitao, Christian Matthes, Tomasz Piskorski, and Demian Pouzo provided outstanding research assistance; Nameun Jeong and Eric Wang provided indispensable assistance on parallel computing in the Linux operating system. We acknowledge the technical support of the Computing College of Georgia Institute of Technology on grid computation techniques. Finally, we thank Mike Chriszt, Javier Donna, Pat Higgins, Jose Ricardo, Diego Vilan, and Elena Whisler for their help with both collecting the data and understanding institutional details. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. The research in this paper was partially supported by separate NSF grants to Thomas Sargent and Noah Williams.

I. INTRODUCTION

This paper formulates a nonlinear stochastic model of inflation, inflationary expectations, and money-creation-financed deficits and uses the implied density for the history of inflation to extract maximum likelihood parameter estimates for data from Peru, Argentina, Bolivia, Brazil, and Chile. The model extends earlier hyperinflation models of Sargent and Wallace (1987), Marcet and Sargent (1989), and Marcet and Nicolini (2003) in ways that allow us to distinguish such possible causes and cures of hyperinflations as persistent reforms and transitory shocks to seigniorage-financed fiscal deficits; explosions in inflationary expectations that are divorced from fiscal fundamentals; and monetary reforms that we call cosmetic because they do not alter fiscal fundamentals. Because each of those earlier papers assumed constant deficits, they were not designed to discriminate among these alternative causes and cures of big inflations.

We construct a hidden Markov model that features a demand function for money inspired by Cagan (1956), a budget constraint that determines the rate at which a government prints money, a stochastic money-financed deficit whose conditional mean and volatility are governed by a finite state Markov chain, and an adaptive scheme for the public's expected rate of inflation that allows occasional deviations from rational expectations that help to explain features of the data that a strict rational expectations version of the model cannot.¹ We trust our monthly series on inflation but lack trustworthy monthly or quarterly data on GDP and the money supply that would allow us to compute high-frequency seigniorage flows. Therefore, to estimate the model's free parameters, we form the density of a history of inflation, view it as a likelihood, and maximize it with respect to the parameters. For each country, we then form a joint density for the inflation and seigniorage histories at the maximum likelihood parameter estimates and use it to calculate a density for the seigniorage history conditional on the inflation history. As one of several validation exercises, we compare this seigniorage density with the annual seigniorage rate data available to us.

Our econometric model yields the following key empirical findings.

- (1) There is strong evidence that cosmetic reforms along the lines of Marcet and Nicolini (2003) brought down hyperinflation in Peru, although fundamental fiscal reforms eventually occurred to sustain low inflation later in the sample.
- (2) In Bolivia, movements in fiscal fundamentals both instigated and ended hyperinflation. This is an important finding because (i) destabilizing expectation effects play no role in generating hyperinflation and (ii) recurrent hyperinflations were caused only by regime changes in fundamentals. Thus, in Bolivia there seemed to be no divorce of inflationary expectations from fundamentals.
- (3) In Chile, changes in shock variances influenced the birth and death of hyperinflations, not cosmetic reforms.

¹Elliott, Aggoun, and Moore (1995) is a good reference about hidden Markov models.

- (4) Data for Brazil and Argentina tell mixed stories. In both countries, cosmetic reforms played roles in temporarily lowering inflation rates, but they were all eventually followed by fundamental reforms.

Thus, our maximum-likelihood estimates enable us to disentangle forces that started and ended different episodes of hyperinflation.

The rest of the paper is organized as follows. Section II describes in detail the stochastic specification of monetized deficits that we use to extend the models of Sargent and Wallace (1987), Marcet and Sargent (1989), and Marcet and Nicolini (2003) in ways that have a better chance of explaining time series that include episodes of both low and high inflation. Section III generalizes the notion of self-confirming equilibria in ways that we use to define escapes and cosmetic reforms in section IV. Section V then defines the likelihood function. Section VI indicates how higher moments of the inflation data that are partly driven by the hidden Markov states allow us to identify parameters from data on inflation alone. Then section VII describes technical details underlying our maximum-likelihood estimation procedure. Section VIII reports maximum likelihood estimates and uses them to interpret inflation histories for our five countries. Section IX compares conditional self-confirming equilibria to rational expectations equilibria computed at our maximum likelihood parameter estimates and argues that they are close. Section X reports a variety of robustness checks while section XI makes concluding observations. Five appendixes contain technical details and descriptions of our data.

II. MODEL

II.1. Money demand, government budget constraint, deficit. The money demand equation and government budget constraint are:²

$$\frac{M_t}{P_t} = \frac{1}{\gamma} - \frac{\lambda}{\gamma} \frac{P_{t+1}^e}{P_t}, \quad (1)$$

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

where P_t is the price level at time t ; M_t is nominal balances as a percent of output at time t ; P_{t+1}^e is the time t public's expectation of the price level at time $t + 1$, and $d_t(m_t, v_t)$ is the part of the government's real deficit that must be covered by printing money. Here $0 < \lambda < 1$, $0 < \theta < 1$, $\gamma > 0$. Equation (1) asserts that the demand for real balances varies inversely with the public's expected rate of inflation $\frac{P_{t+1}^e}{P_t}$. Equation (2) asserts that the growth of nominal balances per unit of output equals d_t , the part of the government deficit that is monetized. Here the parameter θ adjusts both for growth in real output and possibly any direct taxes on cash balances. We use the term seigniorage to denote the money-financed government deficit. We

²For an interpretation of this money-demand equation as a saving decision in a general equilibrium model, see Marimon and Sunder (1993), Marcet and Nicolini (2003), and Ljungqvist and Sargent (2004). The government budget constraint equation (2) was used by Friedman (1948) and Fischer (1982), among many others.

assume that it is exogenous and governed by a hidden Markov process

$$d_t(m_t, v_t) = \bar{d}(m_t) + \eta_{dt}(s_t), \quad (3)$$

$$\Pr(m_{t+1} = i | m_t = j) = q_{m,ij}, \quad i, j = 1, \dots, m_h, \quad (4)$$

$$\Pr(v_{t+1} = i | v_t = j) = q_{v,ij}, \quad i, j = 1, \dots, v_h. \quad (5)$$

Here $s_t \equiv (m_t \ v_t)$ is a Markov state, as in Hamilton (1989) and Sclove (1983), that we the econometricians do not observe; and $\eta_{dt}(s_t)$ is an i.i.d. random shock. We assume that $\log d_t(m_t, v_t)$ is normally distributed with mean $\log \bar{d}(m_t)$ and variance $\sigma_d(v_t)^2$. This distribution implies that $\bar{d}(m_t)$ is both the geometric mean and the median of $d_t(m_t, v_t)$ and that $\eta_{dt}(s_t)$ has the following probability density function:

$$p_d(\eta | s_t) = \begin{cases} \frac{\exp\left[-\frac{[\log(\bar{d}(m_t) + \eta) - \log \bar{d}(m_t)]^2}{2\sigma_d^2(v_t)}\right]}{\sqrt{2\pi}\sigma_d(v_t)(\bar{d}(m_t) + \eta)} & \text{if } \eta > -\bar{d}(m_t) \\ 0 & \text{if } \eta \leq -\bar{d}(m_t) \end{cases}. \quad (6)$$

The log-normal distribution of d_t makes seigniorage positive and captures the skewness of the inflation distribution observed in the data. In our empirical work, we experimented with other distributions that allow for negative seigniorage (e.g., the normal distribution), but these did not improve the model's fit.

The Markov component m_t governs the mean seigniorage while the component v_t governs the volatility of the random shock $\eta_{dt}(s_t)$. We use the conventions that the m_t index runs from high seigniorage to low seigniorage and that the v_t index runs from high volatility to low volatility. We index hidden states in this way to be consistent with our emphasis on the impact of fiscal reforms as policy switches from a first regime (high seigniorage) to a second regime (low seigniorage), although we also discuss cases where fiscal policy switches from the second regime to the first regime as in Section III below.

Each column of each transition probability matrix $Q_\ell = [q_{\ell,ij}]$ for $\ell = m, v$ sums to 1. The Markov chains (Q_m, Q_v) induce a chain on the composite state $s_t = (m_t \ v_t)$ with transition matrix $Q = Q_m \otimes Q_v$.³ The total number of states is $h = m_h \times v_h$. The Markov-switching structure contributes flexibility that allows our model to fit both high and low inflation episodes. When we discuss the theoretical mechanism displayed in figure 1 and our empirical findings in section VIII, we shall indicate some extensive interactions among the Markov states m_t , the seigniorage shocks η_{dt} , and agents' expectations β_t .

II.2. Expectations.

II.2.1. *Rational expectations equilibrium.* The pieces introduced so far, namely equations (1) (2), (3), (4), and (5) are sufficient to define rational expectations equilibria that pin down P_{t+1}^e as the mathematical expectation $E_t P_{t+1}$ taken with respect to the joint probability distribution over sequence of outcomes induced by these equations.

³We have also considered cases where m_t and v_t are not independent, but the fit of these versions of the model is much worse. See Section X for a detailed discussion.

As we discuss in section IX, these equilibria are useful benchmarks, but our paper focuses on an alternative expectation formation mechanism.

II.2.2. *Learning with constant gain.* Let

$$\pi_{t+1}^e \equiv P_{t+1}^e/P_t = \beta_t.$$

As our baseline specification, we follow Marcet and Sargent (1989) and Marcet and Nicolini (2003) and assume that the public updates belief β_t by using a constant-gain algorithm:

$$\beta_t = \beta_{t-1} + \varepsilon(\pi_{t-1} - \beta_{t-1}), \quad (7)$$

where $0 < \varepsilon \ll 1$ and π_t is the gross inflation rate at time t , defined as

$$\pi_t = P_t/P_{t-1}.$$

Specification (7) is consistent with a growing literature in macroeconomics that uses calibration or econometric techniques to compare time series data with models in which some agents use constant-gain learning algorithms to form their beliefs. Other examples are the calibration study of big inflations by Marcet and Nicolini (2003) and the econometric studies of models of Phillips curves and monetary policies by Chung (1990), Sargent (1999), Sargent, Williams, and Zha (2006b), and Carboni and Ellison (2007),⁴ and of the term structure of interest rates by Piazzesi and Schneider (2006).

II.2.3. *Learning with state-dependent gains.* While for most of this paper we retain the constant gain specification, in section X we report some results with two alternative learning rules. Keeping the same notation, we suppose agents have a $m_h \times 1$ vector of beliefs β_t whose m^{th} component evolves according to one of the following two learning mechanisms:

$$\beta_{t,m} = \beta_{t-1,m} + \Pr(m|\pi^{t-1}, \phi) \varepsilon_m (\pi_{t-1} - \beta_{t-1,m}), \quad (8)$$

or

$$\beta_{t,m} = \beta_{t-1,m} + \varepsilon_m (\pi_{t-1} - \beta_{t-1,m}), \quad (9)$$

where $m = 1, \dots, m_h$. The learning rule (8) implies that the current gain depends on the probability of being in state m , given the previous data, while the learning rule (9) does not have this dependence. It is important to note that the rule (8) does not nest the constant-gain learning rule (7), but the rule (9) includes the constant-gain rule as a special case by simply restricting $\varepsilon_m = \varepsilon$ for all m . At time t , β_t equals β_{t,m_t} , where m_t is a realized mean-seigniorage state. These specifications can allow agents to discount past data more rapidly during high inflation episodes and are computationally feasible because they allow us to write the likelihood recursively. Other state-dependent gain specifications, such as a version of (7) with a scalar belief β_t but with a switching ε_{m_t} , are computationally infeasible because they do not allow us to represent the likelihood function recursively.

⁴These papers all use models descended from one proposed and simulated by Sims (1988).

II.3. Further restrictions on expectations. By using (1),(2), and (7), we obtain the following formula for equilibrium inflation:

$$\pi_t = \frac{\theta(1 - \lambda\beta_{t-1})}{1 - \lambda\beta_t - d_t(s_t)}, \quad (10)$$

provided that both the numerator and denominator are positive. Thus, beliefs β_{t-1} and β_t must evidently satisfy the following inequalities:

$$1 - \lambda\beta_{t-1} > 0, \quad (11)$$

$$1 - \lambda\beta_t - d_t(s_t) > \delta\theta(1 - \lambda\beta_{t-1}). \quad (12)$$

Condition (11) ensures that both the price level and money stock in the last period are positive. Given this restriction on β_{t-1} , condition (12) ensures that both the price level and money stock in the current period are positive as well. Furthermore, it imposes a small value $\delta > 0$ to bound the denominator of (10) away from zero so that inflation is bounded above by $1/\delta$. This bound is necessary for the existence of a self-confirming equilibrium, as will be explained in Appendix C.

We have thus far included nothing in our specification to guarantee that beliefs β_t will respect restrictions (11) and (12). Next we turn to explaining why we need to impose some additional restrictions to insure that they are satisfied, and to motivate the restrictions that we choose.

III. MEAN DYNAMICS TOWARD SELF-CONFIRMING EQUILIBRIA

To complete our specification, we use objects that are associated with different notions of a self-confirming equilibrium (SCE), following Sargent (1999) and Cho, Williams, and Sargent (2002). These are useful reference points that help us make precise the senses in which agents' beliefs are consistent with the inflation outcomes that they observe.

Definition III.1. An unconditional SCE is a probability distribution over inflation histories π^T and a β that satisfy $E\pi_t - \beta = 0$.

Definition III.2. For each $m \in \{1, \dots, m_h\}$, a fixed- m SCE is a probability distribution over inflation histories π^T and a $\beta(m)$ that satisfy

$$E[\pi_t | m_t = m \forall t] - \beta(m) = 0.$$

Definition III.3. For each $m \in \{1, \dots, m_h\}$ and $v \in \{1, \dots, v_h\}$ a fixed- m - v SCE is a probability distribution over inflation histories π^T and a $\beta(m, v)$ that satisfy

$$E[\pi_t | m_t = m, v_t = v \forall t] - \beta(m, v) = 0.$$

Self-confirming equilibria and conditional self-confirming equilibria pertain are defined in terms of orthogonality conditions that govern β_t in large samples as the gain $\varepsilon \rightarrow 0$. An unconditional SCE states that agents' beliefs β are correct unconditionally: β is the unconditional expectation of π (which itself is a function of β). In this paper, we focus more on fixed- m SCEs, which better characterize time paths of

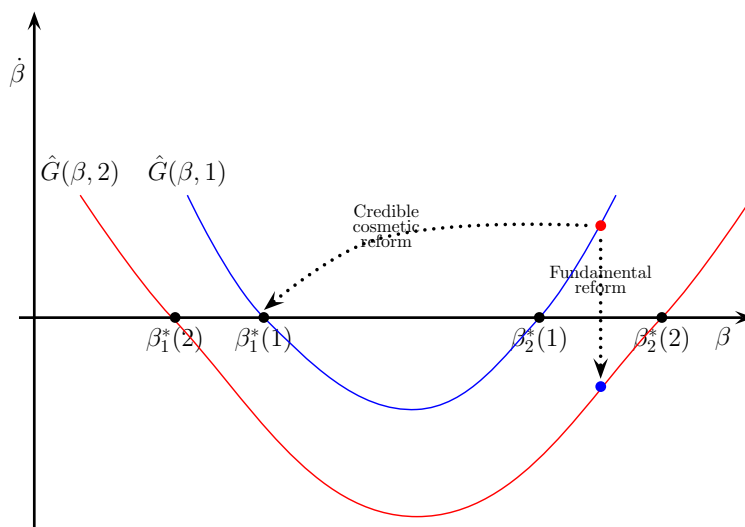


FIGURE 1. Conditional mean dynamics, fixed- m SCEs, and cosmetic and fundamental reforms.

beliefs. A fixed- m SCE is just an unconditional SCE that is computed on the assumption that the seignorage state m_t will always be m . This assumption is of course false, but when m is persistent, it serves as a good, useful approximation. In Section IX, we also discuss a particular set of fixed- m - v SCEs and relate them to deterministic steady-state equilibria. There we also show that the beliefs in the fixed- m SCEs approximate rational expectations equilibrium beliefs.

Our definitions of SCEs hold for fixed β , while the learning rule (7) describes how agents' beliefs β_t adjust so that the orthogonality conditions are approximately satisfied. By averaging appropriately, we can describe the typical behavior of the stochastic process for beliefs by a deterministic differential equation. In particular, the orthogonality condition in definition III.1 is associated with *mean dynamics* $\dot{\beta} = G(\beta)$ that describe the average dynamic behavior of β_t as $\varepsilon \rightarrow 0$. Similarly, the orthogonality condition in definition III.1 is associated with *conditional mean dynamics* $\dot{\beta} = \hat{G}(\beta, m)$ that describe the average dynamic behavior of β_t as $\varepsilon \rightarrow 0$ and as the seignorage state m becomes increasingly persistent. SCEs and conditional SCEs are the rest points of G and $\hat{G}(m)$, respectively. The fact that our estimates of the transition matrix Q_m make the hidden states very persistent renders the conditional mean dynamics especially interesting to us. Appendix C defines mean dynamics, describes how to compute the functions G and $\hat{G}(m)$ and their rest points, and makes precise the sense in which mean dynamics govern the behavior of β_t in our stochastic system.

Figure 1 depicts typical $\hat{G}(\beta, m)$ functions. In section VIII, we shall plot the corresponding \hat{G} functions associated with our maximum likelihood estimates for five countries. For each m , there are two conditional SCEs, a low-inflation $\beta_1^*(m)$ and a high inflation $\beta_2^*(m)$. Evidently, (i) $\beta_1^*(m)$ is a stable fixed point of the conditional mean dynamics, (ii) $\beta_2^*(m)$ is an unstable fixed point, and (iii) $\beta_2^*(m)$ marks the edge of the domain of attraction of the stable SCE. When β_t exceeds $\beta_2^*(m)$, the

conditional mean dynamics propel both beliefs and inflation upward, making β_t on average grow without bound and eventually imperilling conditions (11) and (12). Furthermore, figure 1 shows that when m switches from the low (second) to the high (first) seignorage state (recall our convention that $\bar{d}(2) < \bar{d}(1)$), the stable SCE shifts up: $\beta_1^*(2) < \beta_1^*(1)$. Increases in seignorage are associated with increases in expected inflation. With this increase, the domain of attraction of the stable SCE also shrinks, as $\beta_2^*(2) > \beta_2^*(1)$. Thus, not only does greater seignorage lead to higher average inflation, but there is also a greater chance of rapid increases in inflation after beliefs enter the unstable region. Later, we shall use the conditional mean dynamics to interpret the causes and cures of inflation in five countries. Figure 1 also illustrates two of the types of reforms that can arrest explosive inflation paths, which we now discuss.

IV. ESCAPES AND REFORMS

In this section, we add features to our model that become active when $\beta_t > \beta_2^*(m)$, a situation that the section III as identified as one in which on average both inflation and expectations of inflation threaten to increase beyond points where the model breaks down because either (11) or (12) is violated. We adopt the following language:

Definition IV.1. An *escape* (from the domain of attraction of a stable conditional SCE) is said to occur when $\beta_t > \beta_2^*(m)$.

Definition IV.2. A *reform* is called for whenever without the reform, condition (11) or (12) would be violated so long as the hidden state m were to remain unchanged.

Definition IV.3. A *fundamental reform* occurs when the exogenous seignorage state m jumps to make (11) and (12) satisfied at the initial β_{t-1} .

Definition IV.4. A *cosmetic reform* occurs when (1) a reform is called for, (2) m remains unchanged, and (3) π_t and perhaps β_t is reset according to rules defined in Definition IV.6 or IV.7 below.

Figure 1 illustrates how two types of reforms can lead to lower inflation. A fundamental reform comes from a switch in the seignorage state. After the switch, the belief β is now in the domain of attraction of the low SCE $\beta_1^*(2)$. A credible cosmetic reform that resets inflation and beliefs is shown by a jump in β from a high level down to $\beta_1^*(1)$. Because the economy remains in a high seignorage state, a repetition of a high inflation episode is more likely to reoccur after a cosmetic reform.

For given levels of expectations β_t, β_{t-1} , values of seignorage shocks η_{dt} that contribute to escapes can be defined in terms of the following two critical levels:⁵

$$\underline{\omega}_t(m_t) = 1 - \lambda\beta_t - \frac{\theta(1 - \lambda\beta_{t-1})}{\beta_2^*(m_t)} - \bar{d}(m_t), \quad (13)$$

$$\bar{\omega}_t(m_t) = 1 - \lambda\beta_t - \delta\theta(1 - \lambda\beta_{t-1}) - \bar{d}(m_t), \quad (14)$$

⁵If $\beta_2^*(m_t)$ does not exist, we replace this term in (13) by π_{SS}^{\max} defined in (A4).

By using (10) it can be shown that $\underline{\omega}_t(m_t)$ is the value of the shock η_{dt} that leaves $\pi_t = \beta_2^*(m_t)$.⁶ Thus, because they imply that $\pi_t > \beta_2^*(m)$ via the adaptive mechanism (7), values of $\eta_{dt} > \underline{\omega}_t(m_t)$ contribute to pushing β_t outside the domain of attraction of the stable fixed- m SCE $\beta_1^*(m_t)$. By using (12), we can deduce that $\bar{\omega}_t(m_t)$ is a value of the shock $\eta_{dt}(s_t)$ that puts π_t equal to its upper bound δ^{-1} . This discussion leads to the following.

Definition IV.5. A realization of η_{dt} is said to be escape-provoking if $\eta_{dt} \in [\underline{\omega}_t(m_t), \bar{\omega}_t(m_t)]$.

Conditional on the current-period state being $s_0 = (m_0, v_0)$, the probability that an escape-provoking event occurs or continues is $F_d(\bar{\omega}_t(m_0)|s_0) - F_d(\underline{\omega}_t(m_0)|s_0)$, where $F_d(x|s_0)$ is the cumulative density function of $\eta_{dt}(s_0)$ evaluated at the value x , constructed from the pdf in (6). In other words, this is the probability that the inflation rate is greater than $\beta_2^*(m_0)$ but remains less than the level δ^{-1} that would prompt the reform. The probability that an escape-provoking event occurs at time t , given the $t - 1$ observable data set, is

$$\iota\{\beta_{t-1} < 1/\lambda\} \sum_{s_0=1}^h \left[\Pr(s_t = s_0 | \Pi_{t-1}, \phi) (F_d(\bar{\omega}_t(m_0)|s_0) - F_d(\underline{\omega}_t(m_0)|s_0)) \right], \quad (15)$$

where $\iota(A)$ is an indicator function returning 1 if the event A occurs and 0 otherwise.

Given β_{t-1} , if $\beta_{t-1} \geq 1/\lambda$, a reform is called for with probability one; otherwise, the price level or money stock would continue to be negative even at time t . If $\beta_{t-1} < 1/\lambda$, values of $\eta_{dt} > \bar{\omega}_t(m_t)$ cause the model to break down if both β_t and π_t are left alone. In this case, we impose a cosmetic reform. Conditional on being in state s_0 , the probability that a cosmetic reform at time t occurs is $1 - F_d(\bar{\omega}_t(m_0)|s_0)$, which is evidently the probability that the seignorage shock exceeds the level that sets the inflation rate to its upper bound of δ^{-1} and prompts the reform. Specifically, the probability of a cosmetic reform at time t , given the $t - 1$ data set, is:

$$\iota\{\beta_{t-1} \geq 1/\lambda\} + \iota\{\beta_{t-1} < 1/\lambda\} \sum_{s_0=1}^h \left[\Pr(s_t = s_0 | \Pi_{t-1}, \phi) [1 - F_d(\bar{\omega}_t(m_0)|s_0)] \right]. \quad (16)$$

We now give precise definitions of two types of cosmetic reforms. Both types are used in our estimation and each proves important in different cases.

Definition IV.6. (*A cosmetic reform that resets inflation but leaves beliefs untouched.*) Whenever $\beta_{t-1} \geq 1/\lambda$ so that (11) is violated or whenever $\eta_{dt}(s_t) > \bar{\omega}_t(m_t)$ so that and (12) is violated, we simply reset inflation to the low deterministic rational-expectations-equilibrium inflation rate $\pi_1^*(m_t)$ (computed in appendix A) plus some noise $\eta_{\pi_t}(m_t)$:

$$\pi_t = \pi_t^*(m_t) \equiv \pi_1^*(m_t) + \eta_{\pi_t}(m_t). \quad (17)$$

⁶When we use the state-dependent learning rules then the realized β_{t-1} depends on m_{t-1} and thus we must consider $\bar{\omega}_t(m_t, m_{t-1})$ and $\underline{\omega}_t(m_t, m_{t-1})$ explicitly.

Ideally, we would have liked to reset inflation to the low SCE inflation rate. The reason that we choose the low REE inflation rate instead is entirely practical: the SCEs are difficult to compute, while the REEs are easy and the lower REE approximates the lower SCE well. As shown in Section IX, the low steady state REE inflation rate is close to the low SCE rate with the low value of m (average-seigniorage state) and the low value of v (volatility state). It is computationally inexpensive to calculate $\pi_1^*(m_t)$ according to Proposition 1 in Appendix A. The noise term $\eta_{\pi t}(m_t)$ is modelled as an i.i.d. random shock such that

$$0 < \pi_t^*(m_t) < 1/\delta.$$

The probability density of the noise $\eta_{\pi t}(m_t)$ takes the following particular form:

$$p_{\pi}(\eta_{\pi}|m_t) = \begin{cases} \frac{\exp\left[-\frac{[\log(\pi_1^*(m_t) + \eta_{\pi}) - \log \pi_1^*(m_t)]^2}{2\sigma_{\pi}^2}\right]}{\sqrt{2\pi}\sigma_{\pi}(\pi_1^*(m_t) + \eta_{\pi})\Phi\left(\frac{-\log(\delta) - \log(\pi_1^*(m_t))}{\sigma_{\pi}}\right)} & \text{if } -\pi_1^*(m_t) < \eta_{\pi} < 1/\delta - \pi_1^*(m_t) \\ 0 & \text{otherwise} \end{cases}, \quad (18)$$

where $\Phi(x)$ is again the standard normal cdf of x with the convention that $\log(0) = -\infty$ and $\Phi(-\infty) = 0$. This truncated distribution ensures that inflation is below the upper bound $1/\delta$ and maintains the skewness property of the inflation distribution. In our empirical work, we again experimented with other distributional forms, but they did not improve the model's fit.

Definition IV.7. (*A cosmetic reform with inflation and beliefs both reinitialized.*) If condition (12) would otherwise be violated, then we reset $\beta_t = \pi_t$ as well as resetting $\pi_t = \pi_t^*(m_t) \equiv \pi_1^*(m_t) + \eta_{\pi t}(m_t)$.

Remark IV.8. The cosmetic reforms of Marcet and Nicolini (2003) did not reset beliefs β_t . Resetting beliefs helps fit the data in some cases. A cosmetic reform of the definition IV.6 type can be said to be a *less credible* reform because while inflation is brought down, high beliefs still linger and will be brought down only if π stays low long enough to lower them through adaption. A cosmetic reform of the definition IV.7 type can be interpreted as a *more credible* reform because beliefs adjust immediately. In our empirical work, we find cases, like Peru, where definition IV.7 cosmetic reforms improve the model's fit much relative to definition IV.6 reforms. For other cases, the data seem to favor definition IV.6 reforms.

In summary, our cosmetic reforms are crude devices designed to mimic various things that Latin American governments did in the 1980s to arrest inflation “on the cheap” without tackling fiscal deficits.⁷

TABLE 1. Parameters and their meanings

Parameter(s)	Feature
λ	demand for money
ε	adaptation rate
$\bar{d}(m)$	log seignorage mean
$\sigma(v)$	log seignorage std
σ_π	reform std
Q_m	m - transition matrix
Q_v	v -transition matrix

V. THE LIKELIHOOD FUNCTION

We fix the parameters (θ, δ) in estimation. Denote the remaining free parameters of the model as $\phi = [\lambda \ \bar{d}(m) \ \sigma_d(v) \ \sigma_\pi \ \varepsilon \ \text{vec}(Q_m) \ \text{vec}(Q_v)]$, where $m = 1, \dots, m_h$ and $v = 1, \dots, v_h$. For convenience, table 1 contains a reminder of the meanings of these parameters. Let z^t denote the history $[z_1, \dots, z_t]$. Given a parameter vector, the model induces a joint density $p(\pi^T, m^T, v^T, d^T, M^T, \beta^T | \phi)$, where we set $\beta_0 = \pi_0$ and we set the probability for the initial unobservable state s_0 as described in appendix B. We take the initial observable π_0 as given. The initial value M_0 is a function of β_0 and d_0 has no effect on the likelihood so long as π_0 is given. We take the marginal density $p(\pi^T | \phi)$ as our likelihood function and compute the estimator $\hat{\phi} = \text{argmax}_\phi p(\pi^T | \phi)$. We make inferences about seignorage from the conditional density $p(d^T | \pi^T, \hat{\phi})$. Appendix B describes details.

VI. IDENTIFICATION

The density $p(\pi^T | \phi)$ allows us to infer all of the parameters in ϕ from a record of inflation rates π^T . Here we indicate how the following three important features of the inflation times series contain important identifying information:

- the level of (log) inflation in a high-seignorage state;
- the changing patterns of conditional volatilities of inflation rates;
- the skewness of inflation rates.

We illustrate what drives identification by drawing from $p(\pi^T | \phi)$, i.e., by simulating the model, for two artificial settings of the parameter vector ϕ . In Economy 1, the discounting parameter λ is 0.30 and the median seignorage rate, $\bar{d}(m)$, is 0.10 in the first state (high-seignorage) and 0.01 in the second state (low-seignorage). In Economy 2, $\lambda = 0.89$, $\bar{d}(1) = 0.003$ and $\bar{d}(2) = 0.002$. All other parameters are held fixed across the two economies at the values $\sigma_d(1) = \sigma_d(2) = 0.67$, $\sigma_\pi = 0.1$, $\varepsilon = 0.025$, $q_{m,11} = 0.99$, and $q_{m,22} = 0.99$. Figure 2 reports time series of inflation

⁷Examples of cosmetic monetary reforms are exchange rate pegs, direct price controls, and new currency introductions. See Dornbusch (1985) for a contemporary discussion and Marcet and Nicolini (2003) for a discussion, as well as our discussions in section VIII.

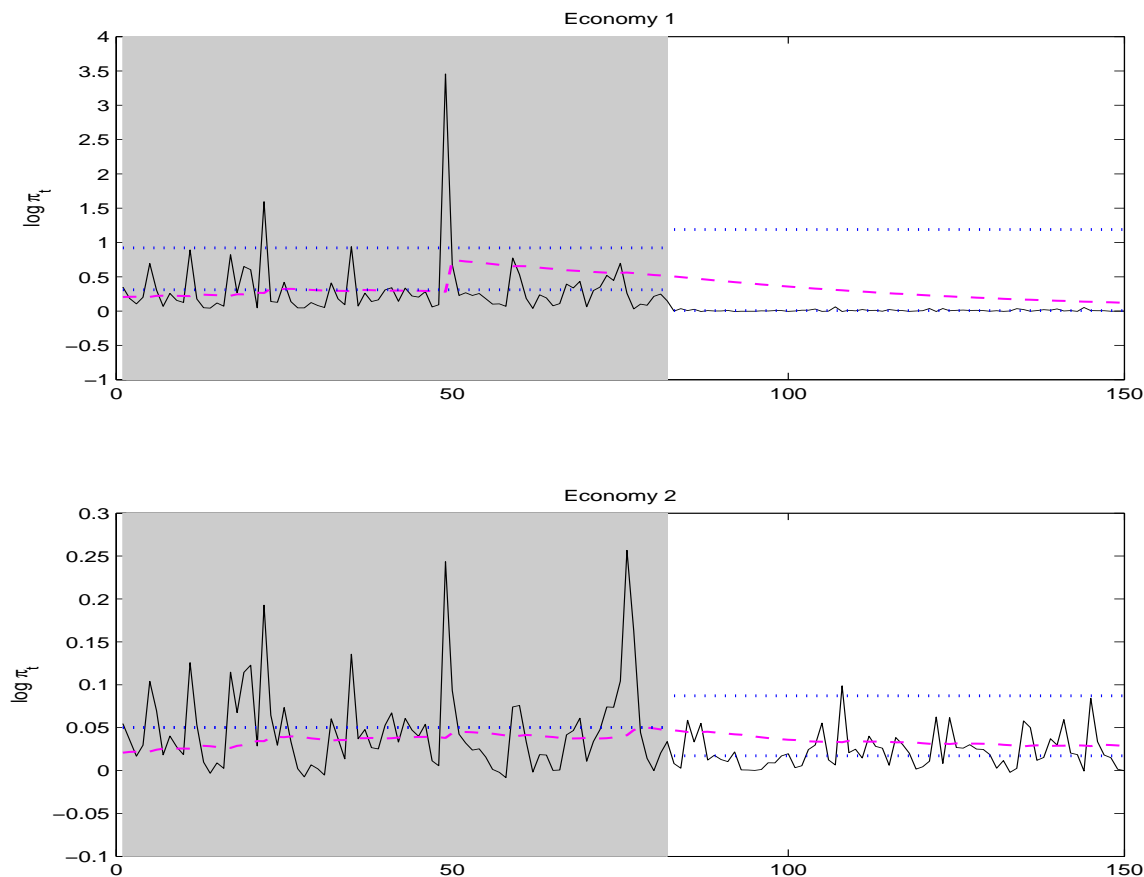


FIGURE 2. Simulated inflation data with $\sigma_d(1) = \sigma_d(2) = 0.67$, $\sigma_\pi = 0.1$, and $\varepsilon = 0.025$ for both economies. Top panel (Economy 1): $\lambda = .3$, $\bar{d}(1) = 0.1$ and $\bar{d}(2) = 0.01$. Bottom panel (Economy 2): $\lambda = 0.89$, $\bar{d}(1) = 0.003$ and $\bar{d}(2) = 0.002$. The dotted lines represent the fixed- m SCEs and the dashed lines represent the evolution of $\log \beta_t$. In the shaded areas, the economies are in the high seignorage state ($m = 1$).

generated with the same starting value for inflation (but with the initial fifty observations thrown away) and with identical draws of seignorage level m_t and identical draws of standardized normal random variables being multiplied by $\sigma_d(v_t)$ to generate seignorage sequences for the two economies (see equation (6)). In addition to inflation, the figure plots the fixed- m SCEs using dotted lines and beliefs β_t using dashed lines. Notice that for Economy 2, the high and low SCEs are virtually identical in the high-seignorage state. That makes escape-provoking events very likely when beliefs approach the higher SCE $\beta_2^*(m)$, as they do in the figure. But for Economy 1, a substantial gap between SCEs in the high-seignorage state makes escape-provoking events less likely, even though inflation rates are much higher. The sample paths are such that escapes occur for *neither* economy, since $\beta_t < \beta_2^*(m)$ always for both economies. Thus, along the sample paths shown, there are no cosmetic reforms for either economy. Instead, when inflation falls, it is either due to seignorage shocks

(in the form of η_{dt}) or shifts in m . As we shall see from the parameter values and interpretations reported in Section VII, Economy 1 resembles Bolivia and Economy 2 resembles Chile.

We now focus on the high-seigniorage state marked by the shaded area in Figure 2. Although the patterns of the inflation series are similar for both economies, inflation in Economy 1 is much higher than in Economy 2 and the volatility pattern differs. The magnitude of inflation and the volatility pattern are both determined by the values of λ and $\bar{d}(1)$. The low discount parameter $\lambda = 0.30$ means that the steady state maximum seigniorage rate is 0.21. Since the maximum seigniorage rate is twice as large as $\bar{d}(1) = 0.10$, the domain of attraction to the low SCE inflation rate is large, lowering the probability of escape-provoking shocks η_{dt} . This is indicated by the wide gap between the two SCEs in the shaded area in the top panel of Figure 2. Consequently, Economy 1 is one in which cosmetic reforms are likely to play little role in generating recurrent hyperinflations. Instead, shifts in the hidden Markov process governing seigniorage are the main force contributing to big reductions in inflation.

In Economy 2, the high discount parameter $\lambda = 0.89$ leads to a much lower average level of inflation than in Economy 1 (note that the scale in the bottom panel of Figure 2 is much smaller than that in the top panel). This high parameter value implies that the steady state maximum seigniorage rate is 0.0038. Because this maximum is only 27% higher than the median seigniorage rate $\bar{d}(1) = 0.003$, the probability of escape-provoking events is much higher than in Economy 1. This enlarges the scope for expectations β to have a role that is divorced from fluctuations in “fundamentals” (i.e., seigniorage). Escapes from the domain of attraction of $\beta_1^*(1)$ (here approximately equal to $\beta_2^*(1)$) are more likely to occur in Economy 2, though none actually occur in the sample path depicted in the bottom panel of Figure 2. Instead, as in Economy 1, hyperinflations are driven solely by shocks η_{dt} to seigniorage. Thus, the sample path displayed for Economy 2 depicts another case where cosmetic reforms like those of Marcet and Nicolini (2003) play no role in generating recurrent hyperinflations: the equilibrium conditions (11) and (12) are satisfied along the particular high-inflation sample path displayed.

The gain parameter, ε , has important effects on the volatility and skewness of inflation and the likelihood of cosmetic reforms. In Economy 2, for example, if we increase ε from 0.025 to 0.04, the beliefs will quickly exceed the higher SCE $\beta_2^*(1)$ and cosmetic reforms will take place. Consequently, log inflation can jump up to as high as 3.5 and jump down to as low as -1.5 . Unless we observe such a skewed and volatile inflation series, the estimated gain is likely to be much smaller than 0.04 and thus cosmetic reforms are unlikely to happen. In some countries, however, cosmetic reforms play a crucial role as we shall see in section VIII.

In summary, the overall magnitude, the volatility pattern, and the skewness of inflation enable one to infer the underlying parameters such as λ , $\bar{d}(m)$, and ε . For the purpose of illustration, in this section we have held all other parameters are fixed. These other parameters have also have important effects, especially on higher moments of inflation. We turn to these in the next section.

VII. ESTIMATION

VII.1. Estimation procedure. In estimation we use the monthly CPI inflation for each country published in the International Financial Statistics. These data sets are relatively reliable and have samples long enough to span episodes of both hyperinflation and low inflation. The long sample makes it reasonable to use the Schwarz criterion to measure the fit of our parsimonious model. The sample period is 1957:02–2005:04 for Argentina, Bolivia, Chile, and Peru and 1980:01–2005:04 for Brazil.

There are no reliable or even available data on GDP, money, and the government deficit in many hyperinflation countries even on a quarterly basis because of “poorly developed statistical systems” (Bruno and Fischer, 1990). However, as discussed in Section VI, we are able to estimate the structural parameters through the inflation likelihood derived in appendix B. As discussed in section V, we fix some parameters that here take the values $\beta_0 = \pi_0$, $\theta = 0.99$, and $\delta = 0.01$. The value of θ is consistent with economic growth and some cash taxes. The value of δ implies that monthly inflation rates are bounded by 10,000%, while Marcet and Nicolini (2003) set the bound at 5,000%. Although we do not use them in estimation, we do have annual data on seignorage that are described in appendix E. As discussed below, we compare them with the distribution of seignorage levels predicted by the model. In Section X, we discuss some exercises with the limited quarterly seignorage data we were able to obtain.

The long samples make the likelihoods of inflation well shaped around their global peaks. There are local peaks but often the likelihood values there are very small relative to the maximum likelihood (ML) value. Nonetheless, if one chooses a bad starting point to search for the ML estimate, the numerical search algorithm is likely to stall at a local peak. Thus, obtaining the maximum likelihood estimates (MLEs) proves to be an unusually challenging task. The optimization method we use combines the block-wise BFGS algorithm developed by Sims, Waggoner, and Zha (forthcoming) and various constrained optimization routines contained in the commercial IMSL package. The block-wise BFGS algorithm, following the idea of Gibbs sampling and EM algorithm, breaks the set of model parameters into subsets and uses Christopher A. Sims’s `csminwel` program to maximize the likelihood of one set of the model’s parameters conditional on the other sets. Maximization is iterated at each subset until it converges. Then the optimization iterates between the block-wise BFGS algorithm and the IMSL routines until it converges. The convergence criterion is the square root of machine epsilon.

Thus far we have described the optimization process for only one starting point. We begin with a grid of 300 starting points; after convergence, we perturb each maximum point in both small and large steps to generate additional 200 new starting points and restart the optimization process again; the MLEs attain the highest likelihood

TABLE 2. Log likelihood adjusted by the Schwarz criterion

	Constant (1 × 1)	Best-fit	Best-fit AR(2)	log posterior odds
Argentina	980.9 (df=8)	1275.4	1346.1 (df=0)	-70.7
Bolivia	1248.1 (df=8)	1540.0	1547.2 (df=0)	-7.1
Brazil	510.7 (df=9)	814.6	853.6 (df=-1)	-39.0
Chile	1422.3 (df=8)	1745.9	1721.7 (df=0)	24.14
Peru	1378.1 (df=8)	1711.7	1658.7 (df=0)	52.8

value.⁸ The other converged points typically have much lower likelihood values by at least a magnitude of hundreds of log values.

VIII. FINDINGS

In this section, we present and interpret our main empirical results with the constant-gain learning rule. Results for state-dependent gains are discussed later in Section X. Before going through the analysis country by country, we look at how best-fitting models are determined and how some key parameters vary across countries.

VIII.1. Fits. Since our theoretical model is highly restricted, one would not expect its fit to approach that of a standard autoregressive (AR) model, let alone a time-varying AR model. In previous work with models related to ours, such as Marcet and Nicolini (2003), only particular moments or correlations implied by the model were typically reported and compared to the data. By contrast, we evaluate the fits of our models for all five countries by comparing various versions of our theoretical model with one another and also with fits attained with flexible, unrestricted statistical models.

For each country we have tried more than two dozen versions of our theoretical model and of the unrestricted atheoretical models. Among other things, we varied the number of Markov states and their interdependence and considered alternative specifications for distributions of the shocks η_{dt} . If the number of states is 3 for m_t and 2 for v_{2t} , we call it a 3×2 model. By the Schwarz criterion (SC) or Bayesian information criterion, the 2×3 version of the model fits best for Argentina, Bolivia, and Chile and the 3×2 version is the best for Brazil and Peru; all other versions including the 1×1 constant-parameter model fit worse. With the 3-state case, we follow Sims, Waggoner, and Zha (forthcoming) and restrict the probability transition matrix to be of the following form:

$$\begin{bmatrix} \chi_1 & (1 - \chi_2)/2 & 0 \\ 1 - \chi_1 & \chi_2 & 1 - \chi_3 \\ 0 & (1 - \chi_2)/2 & \chi_3 \end{bmatrix},$$

where χ_j 's are free parameters to be estimated.

⁸The `csminwel` program can be found on <http://sims.princeton.edu/yftp/optimize/>. For each country, the whole optimization process is completed in 5-10 days on a cluster of 14 dual-processors, using the parallel and grid computing package called STAMPEDE provided to us by the Computing College of Georgia Institute of Technology.

Table 2 reports the log likelihood (adjusted by the Schwarz criterion) of the best-fitting theoretical model for each country, the constant-parameter theoretical model of Marcet and Nicolini (2003) where there are no state-dependent parameters, and the best-fitting unrestricted regime-switching AR model. The constant-parameter model fits poorly for every country. For all five countries, the best-fitting atheoretical model is the 2×2 hidden Markov AR(2) model that allows the two states driving coefficients to be independent of the two states driving shock variances. We use our best-fitting theoretical model as a baseline for comparison. The notation “df” stands for degrees of freedom in relation to the baseline model. The last column reports the posterior odds of the best-fitting economic model relative to the best-fitting statistical model. The table shows that our model fits worse than the best fitting atheoretical model for Argentina, Bolivia, and Brazil, but better for Chile and Peru.

However, even in countries such as Argentina and Brazil, where our model fits worse overall, the discrepancy is largely driven by the superior performance of the statistical models in the non-hyperinflation episodes. For the periods of hyperinflation, our model does much better than the atheoretical statistical models (see Sargent, Williams, and Zha (2006a) for detailed evidence). We prefer to evaluate our model not by fit alone but in terms of the economic interpretations of hyperinflations and stabilizations that it enables. We now discuss these.

VIII.2. Parameter patterns. Table 3 reports the maximum likelihood estimates of our model for Peru, Argentina, Bolivia, Brazil, and Chile.⁹ There are interesting cross-country differences in the important discounting or elasticity parameter λ in equation (1) and the gain parameter ε that controls the rate at which past observations are discounted in the expectations scheme (7). Bolivia has the lowest λ and the highest ε , indicating that it discounted future money creation rates the *most* through a low elasticity of the demand for money with respect to expected inflation, while it also discounted past rates of inflation the *most* through a high gain in the expectations scheme. Comparing Bolivia’s (λ, ε) with Chile’s shows expected inflation to be more important in the demand for money and expectations to discount past observations much less in Chile.

In general, as discussed in Section VI, the smaller λ is, the less likely it is that an escape will take place because the domain of attraction of the low SCE inflation rate is larger. Once β is in the escape region, a large value of ε tends to accelerate increases in both inflation and beliefs β . An informative example is Brazil where both λ and ε are large. For Argentina, Chile, and Peru, the value of λ is even larger and consequently the escape-provoking probabilities are quite high during the hyperinflation period. For Bolivia, the value of λ is quite low. Thus, even though its estimated gain is higher than those in the other countries, the domain of attraction of the low SCE is large enough to prevent the escape event from occurring during

⁹The estimated standard errors of these estimates are reported in Sargent, Williams, and Zha (2006a). They are small. We do not report them in this version to save space.

TABLE 3. Estimates of ϕ from best-fit models for all five countries

Parameters	Peru	Argentina	Bolivia	Brazil	Chile
λ	0.740	0.730	0.307	0.613	0.875
$\bar{d}(1)$	0.0129	0.0165	0.0752	0.0481	0.00317
$\bar{d}(2)$	0.0052	0.0040	0.0105	0.0234	0.00175
$\bar{d}(3)$	0.0033	*	*	0.0060	*
$\sigma_d(1)$	2.538	9.615	18.87	0.355	4.926
$\sigma_d(2)$	0.312	0.675	0.756	0.092	2.298
$\sigma_d(3)$	*	0.264	3.252	*	0.435
σ_π	0.063	0.060	0.308	0.101	0.094
ε	0.069	0.023	0.232	0.189	0.025
$q_{m,11}$	0.9943	0.9789	0.9629	0.9845	0.9869
$q_{m,22}$	0.9626	0.9838	0.9959	0.9732	0.9930
$q_{m,33}$	0.9650	*	*	1.0000	*
$q_{v,11}$	0.3016	0.4395	0.3344	0.9344	0.7627
$q_{v,22}$	0.9547	0.9260	0.8180	0.9031	0.9310
$q_{v,33}$	*	0.9713	0.8513	*	0.9131

Note: The symbol * means “not applicable.”

the hyperinflation period. We return to these observations in our country-by-country analysis below.

VIII.3. Peru (cosmetic reforms). Figure 3 uses our maximum likelihood estimate $\hat{\phi}$, as reported in Table 3, and the implied joint distribution $p(\pi^T, m^T, v^T, d^T, M^T, \beta^T | \hat{\phi})$ to interpret inflation episodes in Peru. First, for $\hat{\phi}$, we use procedures described in appendix C to construct the conditional mean dynamics $\hat{G}(\beta, m)$ for the hidden $\bar{d}(m)$ states, $m = 1, \dots, m_h$. The top left panel depicts the functions $\dot{\beta} = \hat{G}(\beta, m)$ discussed above, with the zeros being the fixed- m SCEs. We projected the fixed- m SCEs as horizontal dotted lines into the top right panel, which plots our estimates of the public’s inflation beliefs β_t that are implied by the initial β_0 , our estimate of ε , and π^T .

The second panel from the top on the right shows bars that are seigniorage rates constructed from annual data as well as 0.16, 0.5, and 0.84 probability quantiles for d_t for the estimated density $p(d^T | \pi^T, \hat{\phi})$ (see Appendix E for details about how these numbers are computed in the data and the model). The dashed lines in the graph contain two-thirds of the probability distribution of simulated annual seigniorage from our model; the solid line labelled “Model” represents the median of simulated annual seigniorage.

The third panel from the top records probabilities of two events that we have computed from the joint density $p(\pi^T, d^T | \hat{\phi})$. The thick solid line, denoted “L & M Seigniorage” is the probability that the median-seigniorage state m is in either the low

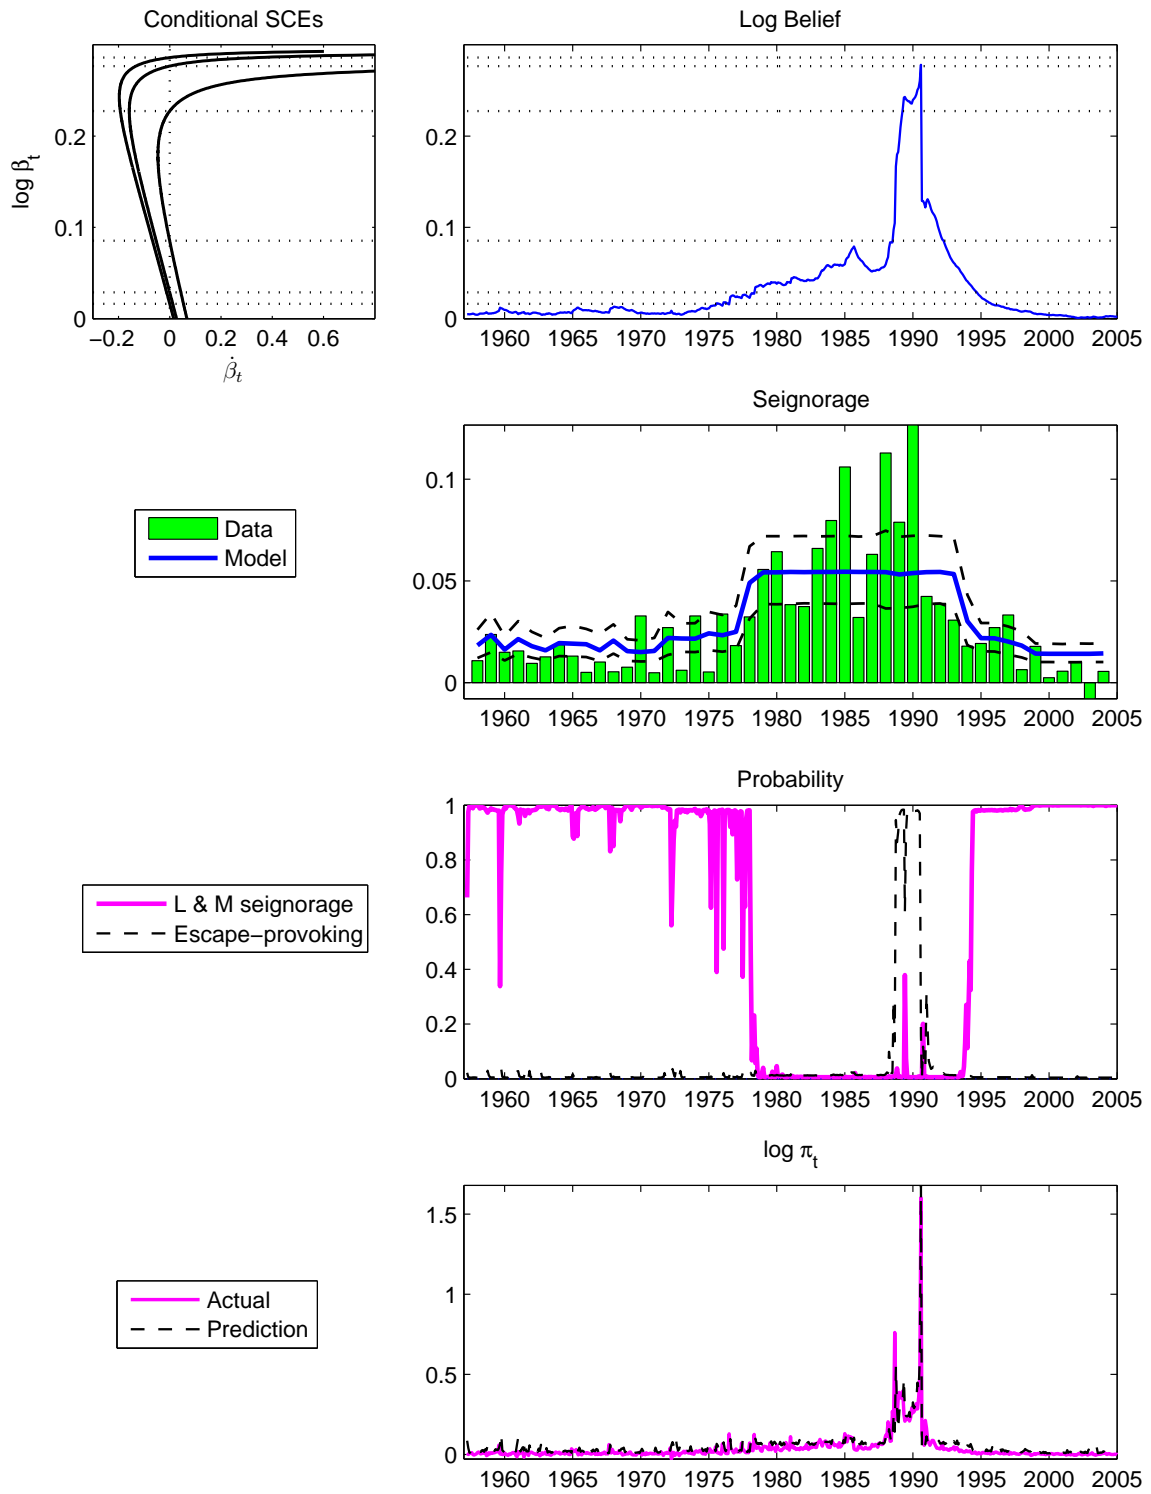


FIGURE 3. Peru.

or the medium state as a function of time. The dashed line is the probability that an escape-provoking event will occur next period, computed as described in section IV above. Finally, the bottom panel shows the actual inflation history π^T and the history of one-step ahead estimates produced by our model evaluated at $\hat{\phi}$, conditioning on earlier inflation rates.

Now turning to what these plots say about Peru, we see in the third panel that throughout the 1960s and most of the 1970s, Peru had low or middle $\bar{d}(m)$ states. The average seignorage does not vary substantially across these states, as the second panel shows that the median prediction is relatively flat, which roughly matches the annual seignorage data. Expected inflation was relatively low throughout this period (top panel), remaining near the conditional SCEs that are at 0.016 and 0.029 for the low and medium states, and the actual inflation rate was relatively low and stable as well (bottom panel). However, around 1978 Peru entered into a high-seignorage state (third panel) that persisted until 1993. This is confirmed by the persistently high seignorage revenues shown in the second panel. But inflation did not accelerate immediately. Rather, beliefs drifted upward throughout the 1980s to the stable conditional SCE in the high state (top panel) and inflation climbed slowly along with it (bottom panel).

But being in the high seignorage state made Peru vulnerable to a sequence of positive shocks η_{dt} that would threaten to send expectations β_t above the escape threshold β_2^* . In the late 1980s expected inflation increased rapidly, traversing into the region that prompted a large and rapid escape, as indicated in the belief dynamics in the top panel and the sharp increase in the escape-provoking probability in the third panel. Inflation itself increased even more dramatically (bottom panel), triggering a credible cosmetic reform as in definition IV.7 in which both inflation *and* beliefs are reset. For our constant-gain learning models, Peru is the only country in our data set for which our model asserts that such a double-barrelled cosmetic reform certainly occurred. (For learning with state-dependent gains, we find another case where the data favor such a double-barrelled cosmetic reforms.) We interpret the resetting of β as indicating that to the public the cosmetic reform is credible in the sense that the public believed it to be effective in cutting future inflation rates. Consequently, expected and actual inflation jumped down dramatically in 1992. Moreover, this cosmetic reform seemed to have been successful. Consistent with the high-seignorage SCE around 0.08, inflation remained relatively low and stable throughout the rest of the sample (bottom panel), even though the economy remained in the high average seignorage state for a considerable time. Interestingly, the run-up in inflation followed the problems with price controls and the nationalization of banks in 1987. The stabilization occurred when President Fujimori took office in 1991. Evidence of a fiscal reform is absent until around 1994, when the probability assigned to the low or medium seignorage state increased nearly to one (third panel). Thus, Peru seems to be a case where unorthodox cosmetic reforms along the lines of Marcet and Nicolini (2003) were successful in vanquishing hyperinflation.

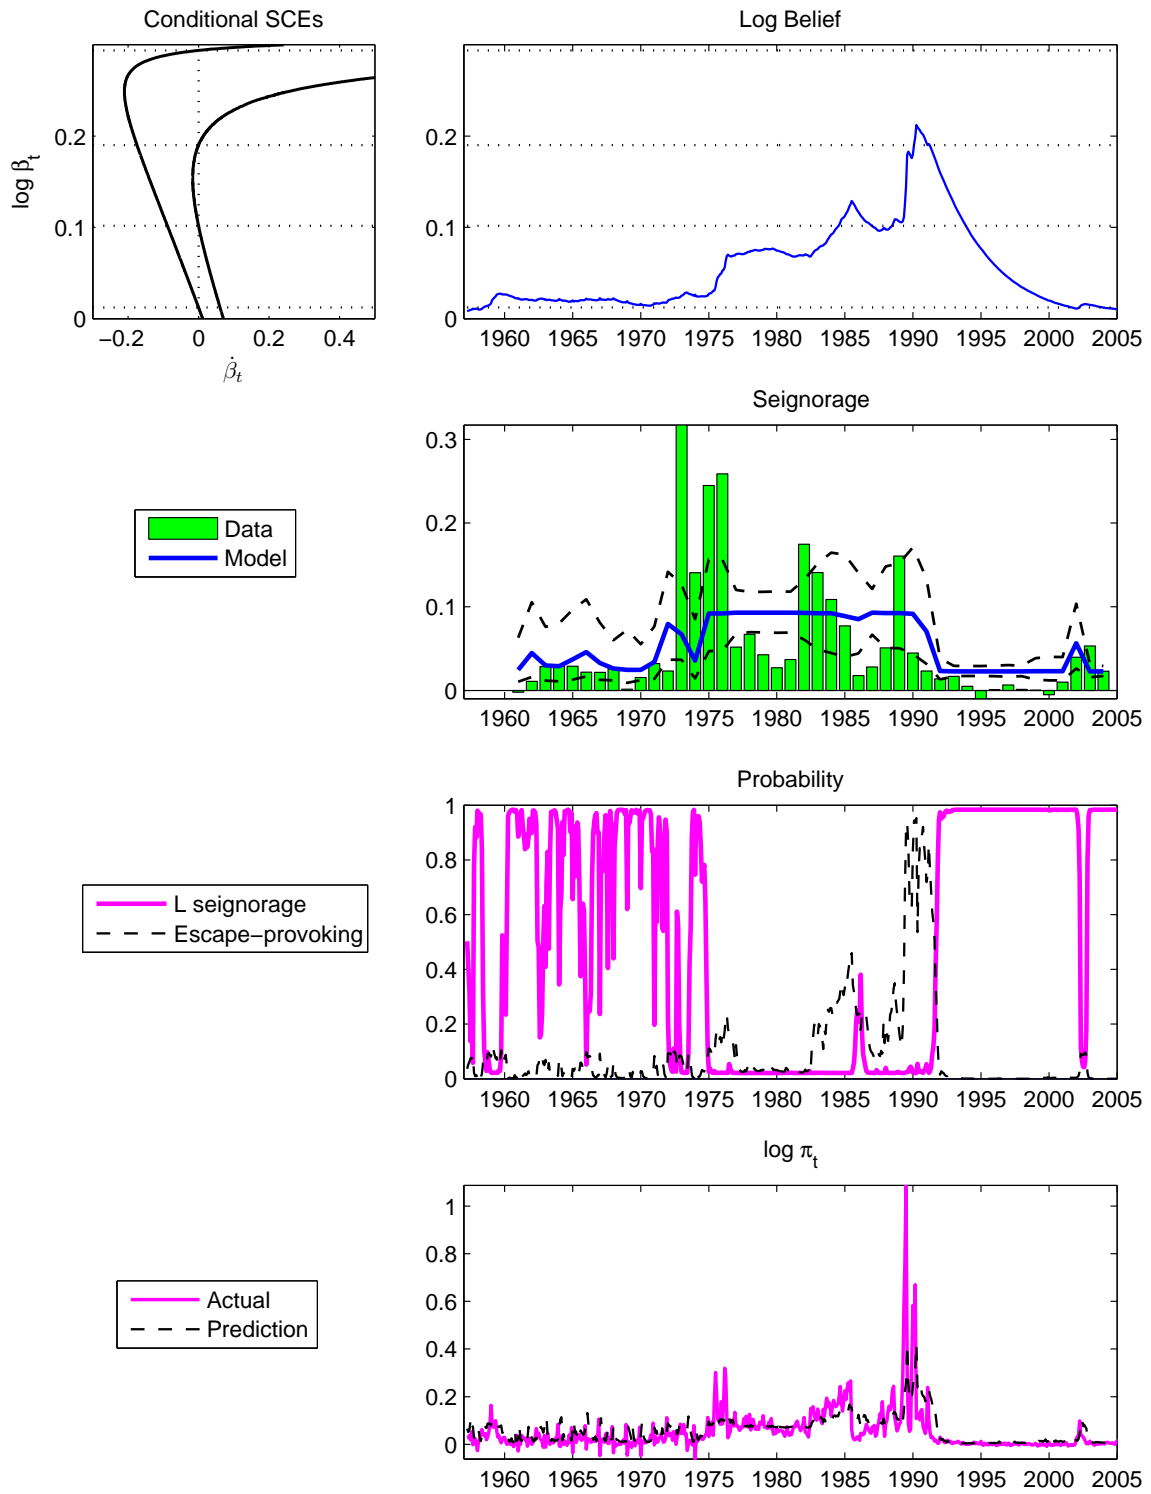


FIGURE 4. Argentina.

VIII.4. Argentina (fundamental reforms). In figure 4 we plot the same five panels for Argentina. Comparing β_t in the top right panel with the probabilities in the third panel down shows that throughout the 1960s up until 1975 the economy was repeatedly in the low seignorage state, the probability of an escape-provoking event was very low, and expected inflation hovered around the low conditional SCE. The low probability for the low seignorage state in the third panel on the right shows that after 1975 up until 1991, Argentina lived with a chronically high $\bar{d}(m)$. The second panel from the top shows that our model predicts higher and more volatile seignorage throughout this period, and this is largely confirmed in the annual seignorage data. Throughout this period, expected inflation drifted higher and higher (as shown in the first panel on the right), first tending toward the stable conditional SCE around 0.1 associated with the high m state, then going even higher. The bottom panel shows that actual inflation drifted upward during this period, with spikes of very high inflation in 1976 and 1984, driven largely by the shocks to seignorage. Again, looking at the third panel, the probability of an escape-provoking event becomes large when β approaches and finally exceeds that higher fixed point near 0.2 in 1989 and 1990. As expected inflation increased rapidly in 1990 actual inflation (as shown in the bottom panel) increased even more rapidly, leading to a dramatic hyperinflationary episode.

The SCE dynamics conditional on high $\bar{d}(m)$ indicate that if Argentina had been lucky enough to avoid sequences of adverse shocks that drove β substantially above the stable rest point near 0.1, it could have avoided the kind of big inflation associated with an escape. Our estimates say that it was actually lucky in this way until the late 1980s, when the escape-provoking event probability escalated and an escape occurred. From 1991-1992, the inflation fell rapidly as shown in the bottom panel. Our model attributes this stabilization to switches in the Markov states governing the mean and volatility of seignorage, which remained in a lower and less volatile state for most of the rest of the sample. Again, this is confirmed by the second panel, which shows lower seignorage throughout the later 1990s, apart from a period of volatility associated with the crisis in 2002. This change in the state m shifted the conditional dynamics curve $\hat{G}(\beta, m)$ (in the top left panel) in a way that pushed expected inflation rapidly downward, with beliefs (in the top right panel) heading toward the lower conditional SCE once again. Although the change in the seignorage state m and the decline in the actual inflation rate occurred relatively quickly, the small estimated gain ε made expected inflation fall gradually throughout the last years of the sample. Thus, our results suggest that the convertibility plan in 1991 was successful not because it pegged the exchange rate, but because it was backed up by a fiscal reform that persistently lowered seignorage. In section X, we show how some of our results for Argentina are altered when we consider a state-dependent learning rule.

VIII.5. Bolivia (fiscal determination). Our results for Bolivia tell a substantially different story. Our estimates suggest that the escape dynamics played no role in Bolivia. The most striking thing about the conditional dynamics in the top left panel of figure 5 is how spread out the conditional SCEs are in each seignorage state m .

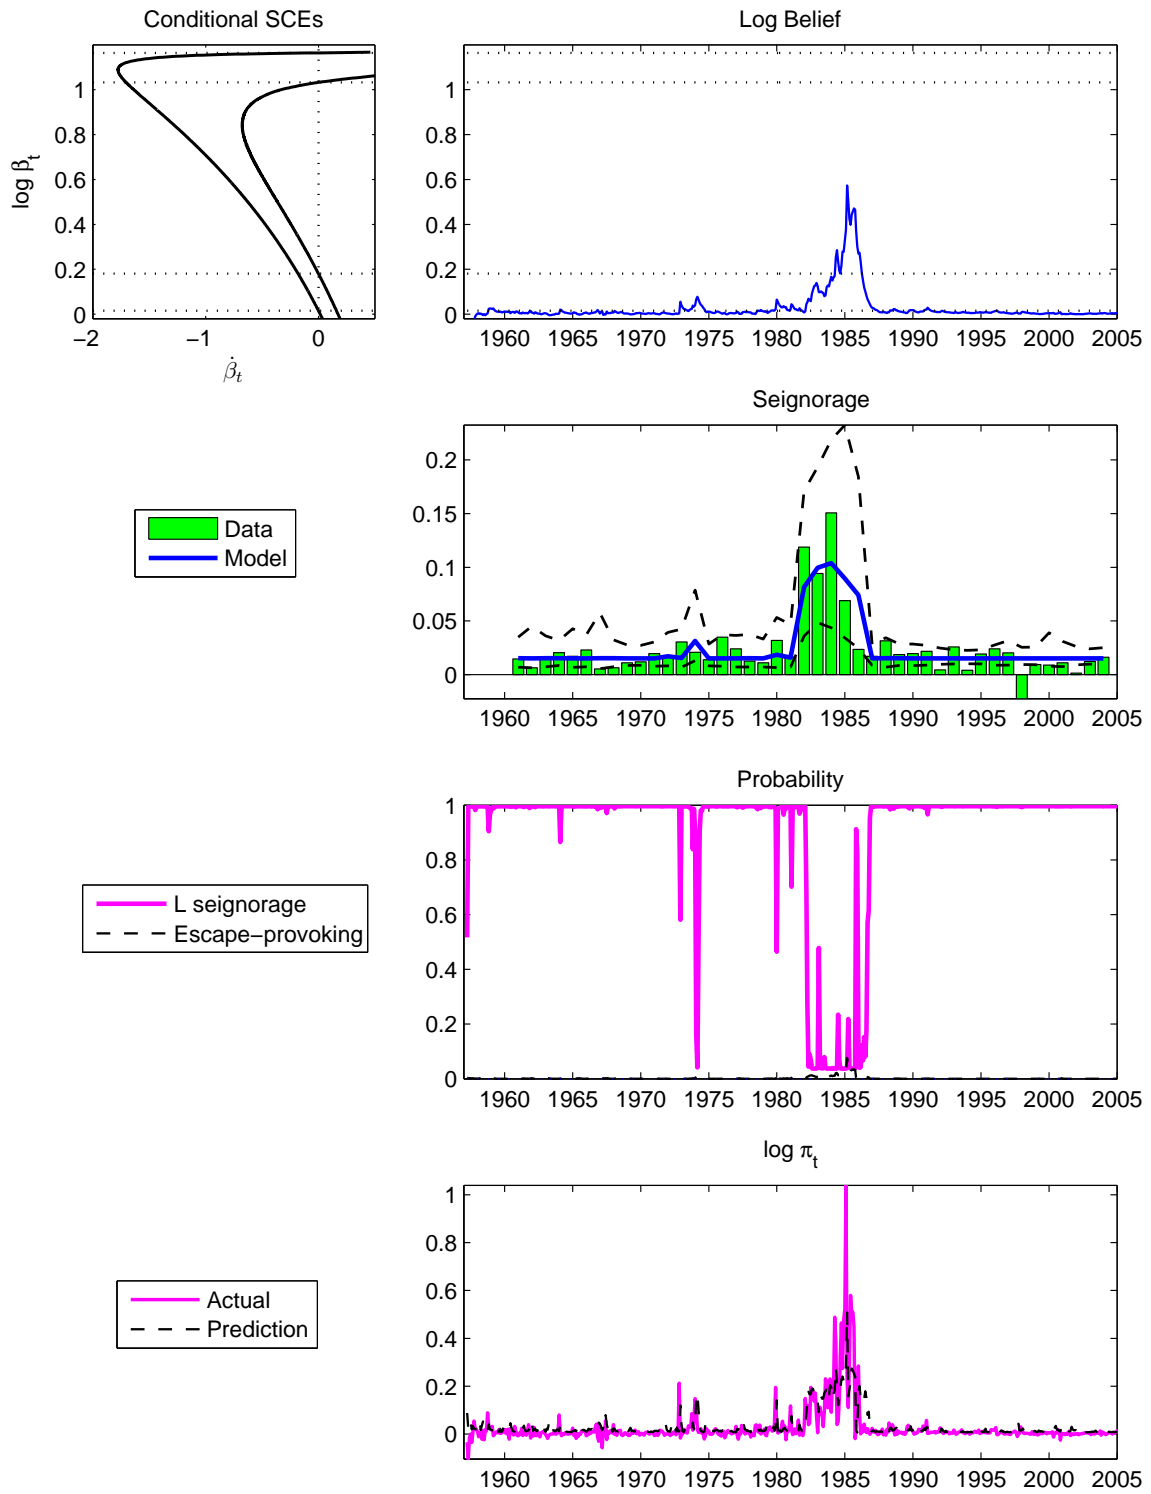


FIGURE 5. Bolivia.

As noted above, the spread between these SCEs is heavily influenced by the money demand elasticity parameter λ , whose estimated value in Bolivia is quite low. Thus, the stable conditional SCEs are near zero and 0.2 in the low and high m states, respectively, but the high SCEs that mark the edge of the domain of attraction are both over 1.0. As the top right panel shows, the beliefs β_t never get into the region where the unstable dynamics take over. This is confirmed by the third panel of the figure, which shows that the escape-provoking event probabilities are very small throughout the entire sample, so small that it is hard to detect by eye.

Since the learning dynamics play very little part in Bolivia, our model suggests that the dynamics of inflation in this country are almost entirely driven by the dynamics of seignorage revenue. As shown in the third panel, our estimates indicate that a switch to the high $\bar{d}(m)$ state m took place around 1982, a view that is confirmed by the seignorage data shown in the second panel. Throughout most of the sample our model predicts a low and relatively stable level of seignorage, with a large and volatile period in the mid-1980s, and this is essentially what the data show. With this switch to a higher m state, expected inflation increases (top panel) and the country experiences a hyperinflation (bottom panel) driven both by the higher mean inflation and large shocks (notice the relatively large discrepancy between the predicted and actual inflation in this period). However, after the 1985 “shock therapy” reforms were implemented, the economy switched back to the low and more stable seignorage states (third panel), actual seignorage is lower and more stable (second panel), expected inflation falls back down toward the lower stable conditional SCE (top panel), and actual inflation is stabilized at a low level (bottom panel). This country thus illustrates the importance of allowing the data to determine the causes of hyperinflation, whether due to learning dynamics or largely driven by fundamentals. Bolivia is a prime example of the importance of the fiscal determination of hyperinflation.

VIII.6. Brazil (cosmetic reforms followed by fundamental reforms). Brazil, as shown in figure 6, presents an interesting case study with two main episodes of hyperinflation that appear to have been ended by different means. First note that in the top left panel the low seignorage state has a conditional SCE near zero, the medium state’s SCE is near 0.06, but the high $\bar{d}(m)$ conditional dynamics curve has no fixed points. We interpret this as asserting that when the economy is in this state, expected inflation is likely to increase steadily and an escape will occur unless the country is lucky enough to have a sequence of negative shocks that push it far enough below that high conditional mean.

Our estimates suggest that from 1980-1985 the economy was in the medium seignorage state, as evidenced by the state probabilities in the third panel down and the predictions and actual levels of seignorage in the second panel. Throughout this period, expected inflation was near the medium $\bar{d}(m)$ SCE (top right panel) and actual inflation was moderately high but relatively stable (bottom panel). However, between 1985 and 1987 the economy shifted to the high seignorage state, remaining there until 1994. Again, this is clear from the state probabilities in the third panel and the

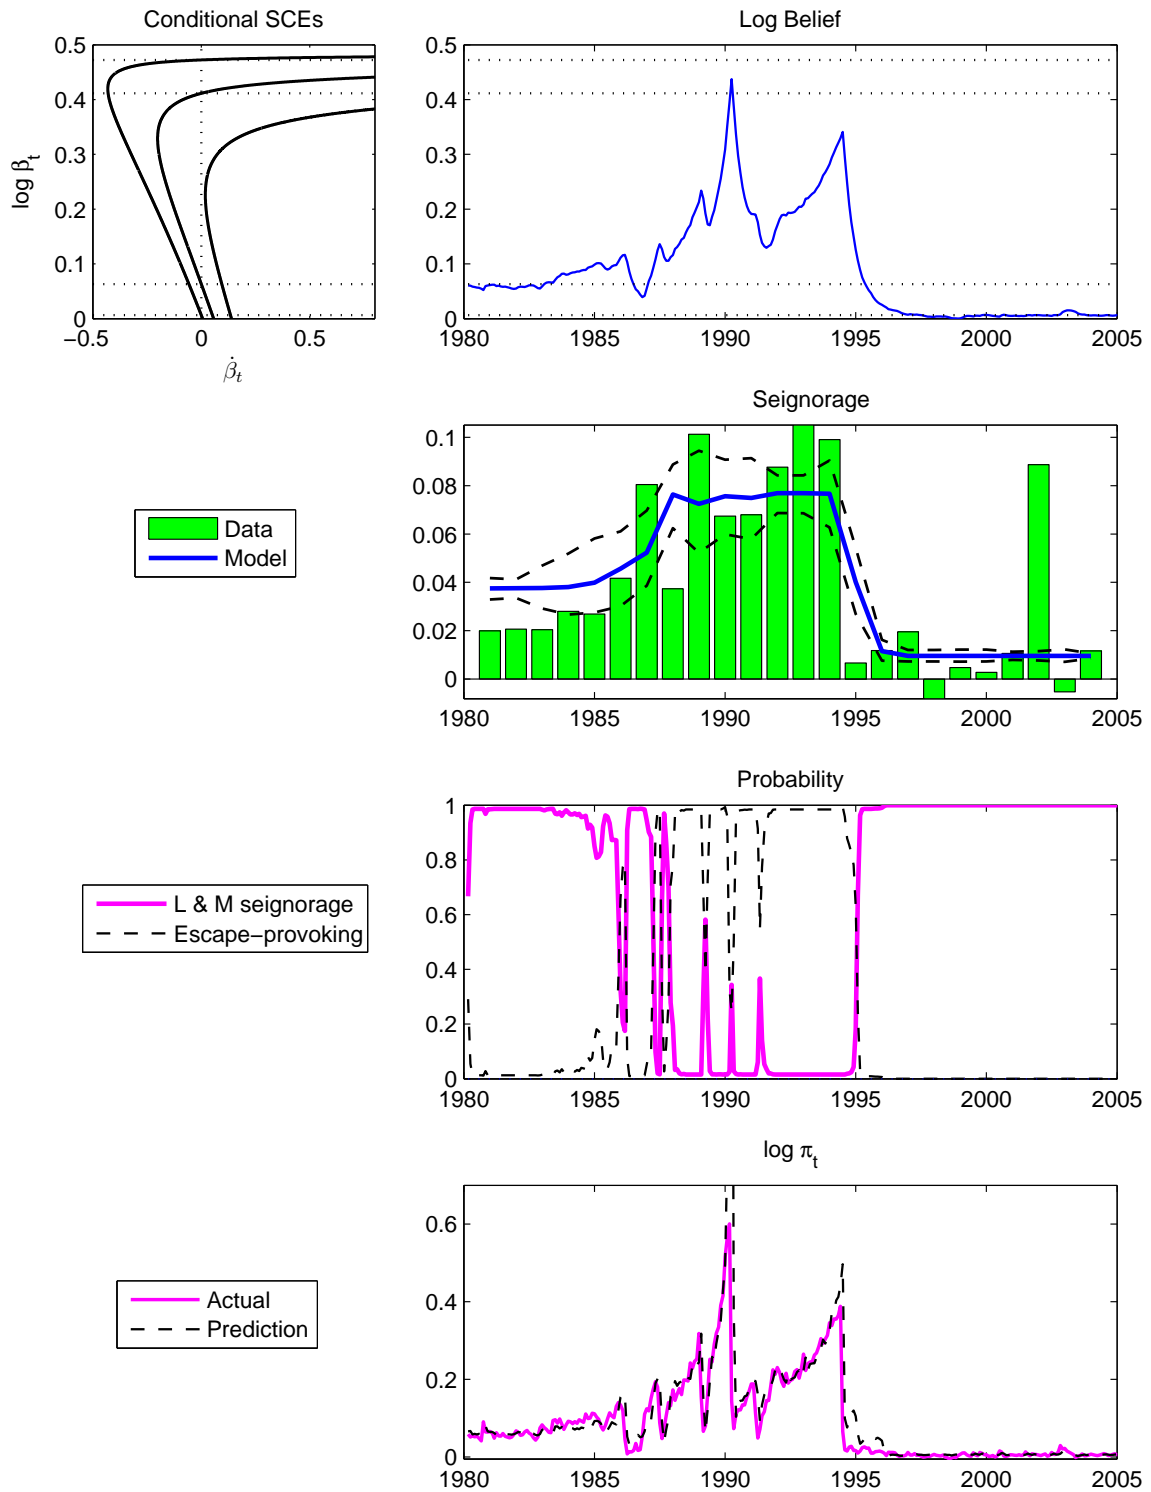


FIGURE 6. Brazil.

seignorage predictions and data in the second panel. Once the economy entered this state, the unstable learning dynamics kicked in. The escape-provoking event probabilities in the third panel rose repeatedly after 1985, remaining mostly near one after 1987 up until 1994, and there was high and volatile seignorage during this period. Expected inflation increased rapidly from 1987 through 1991 (top panel) and actual inflation skyrocketed (bottom panel). In the graph, the predicted value for this first hyperinflation is about 3 in log points. We truncate the figure at 0.7 log points in order to make the actual and predicted inflation paths more discernible.

Actual inflation fell from its peak in 1991 (bottom panel), while the economy continued to run large deficits that necessitated money creation (second panel). Thus, our model interprets the recurrent inflations and stabilizations before 1994 in the manner of Marcet and Nicolini (2003), namely, as recurrent escapes followed by cosmetic reforms. Unlike Peru, these reforms are instances of the less comprehensive cosmetic reform from definition IV.6. In them, inflation is reset but not beliefs, reflecting the incomplete credibility that the public attaches to the reform. Nonetheless, these reforms did succeed in lowering expected inflation sharply due to the large learning gain (top panel). This reduction was only temporary, as expected inflation rose rapidly again until 1994, with actual inflation rising again to another peak (bottom panel). Our model says that the 1994 stabilization is different from the earlier cosmetic reforms; this time the stabilization was accompanied by a persistent reduction in the mean and volatility of seignorage. This is evident in the lower predicted seignorage in the second panel, which largely accords with the lower and more stable actual levels (apart from 2002). After 1994, beliefs fell rapidly down to the low seignorage SCE (top panel), and actual inflation remained stable at a low level (bottom panel). Moreover, as shown in table 3, our estimates of the transition probabilities Q_m suggest that the high and medium $\bar{d}(m)$ states are transitory, and thus our model predicts the sustained stable inflation that accompanies the low seignorage state. These findings are consistent with the policy experience of Brazil, which, before adopting the stable *Real* plan in 1994 that systematically reduced seignorage, had during 1986-1990 introduced a succession of plans with wage and price freezes and new currencies. Thus, Brazil provides an interesting example of some futile cosmetic reforms ultimately being followed by a successful sustained fiscal reform.

VIII.7. Chile (variances of fiscal shocks). In figure 7 we consider the case of Chile, whose experience again is rather distinct from the other countries. First note that the scale in the top panels is significantly smaller than the other countries, with the low seignorage state having conditional SCEs near 0.03 and 0.12, and the high state having essentially one rest point near 0.07. Thus, even the escapes in Chile are consistent with much lower inflation rates than in Brazil. Moreover, the seignorage levels themselves do not vary sizeably across states, as the median model prediction in the second panel is essentially flat over the entire sample. The probabilities of the low $\bar{d}(m)$ state in the third panel are relatively volatile before 1994 (as reflected by relatively volatile inflation rates), but these do not translate into volatile predictions.

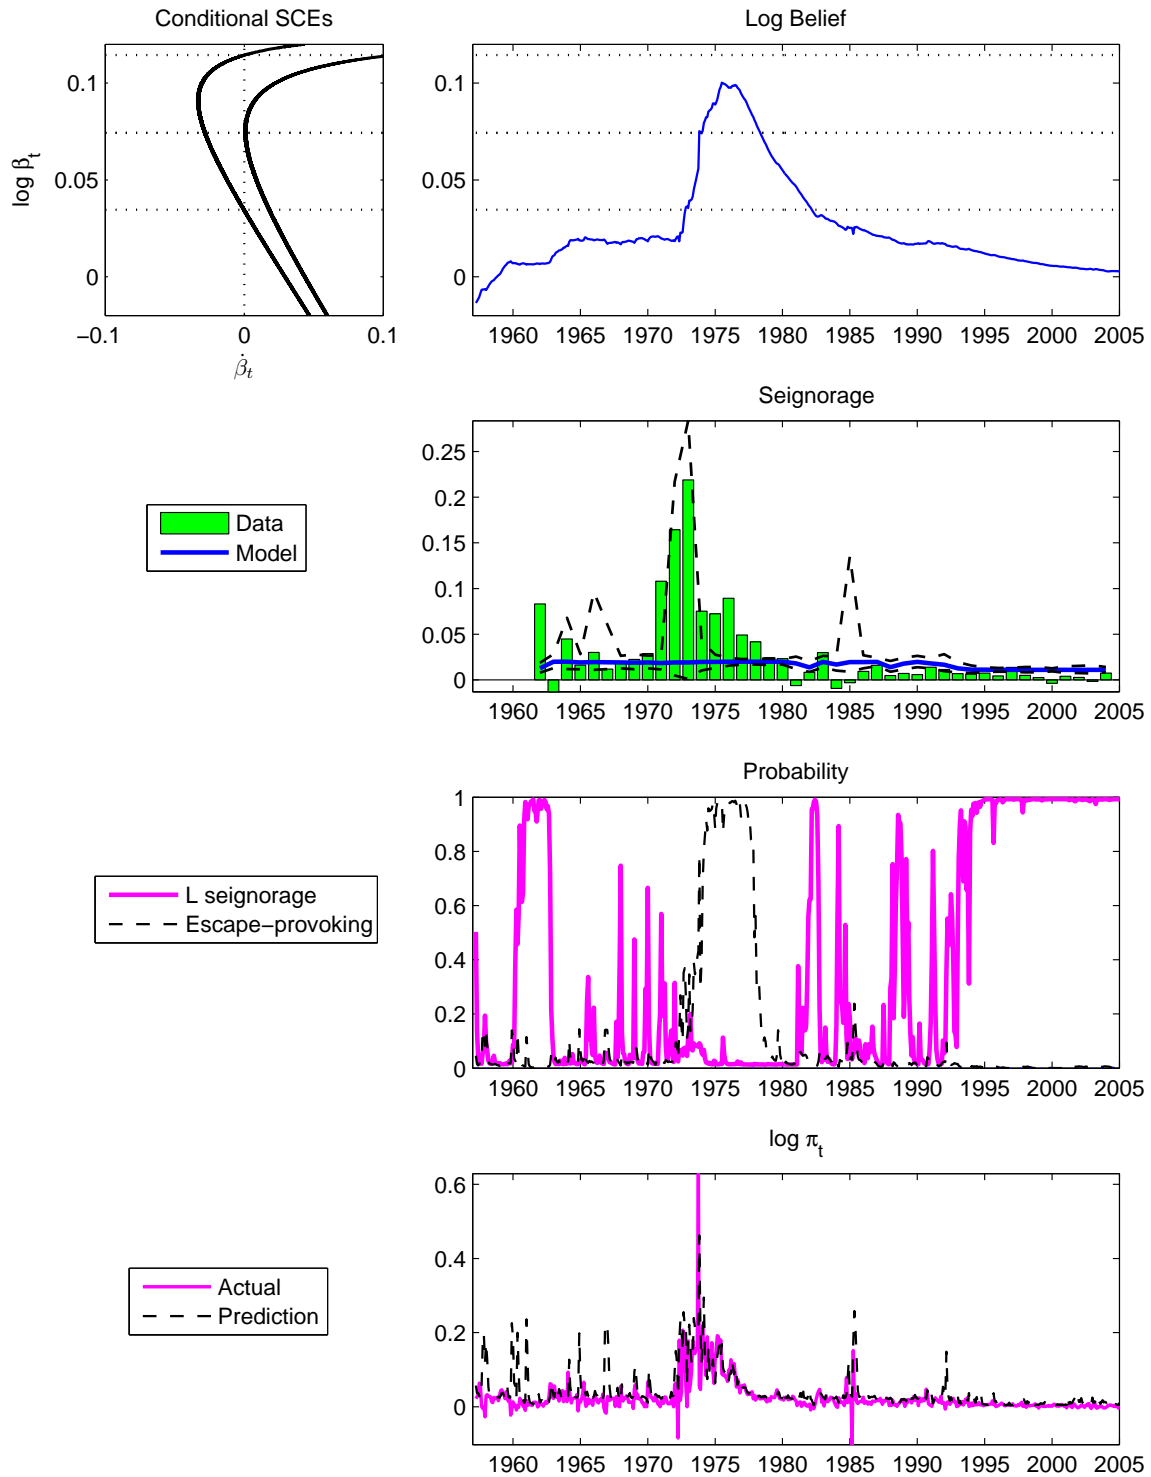


FIGURE 7. Chile.

Thus, our estimates suggest that the buildup and spike in inflation in the mid 1970s (bottom panel) was caused by a sustained run of high seignorage, largely driven by economy entering the high shock variance state. This is evident in the second

panel, where although the median prediction remains flat in the 1970s, there is a large tail evident in the distribution of seignorage, so that the predictive distribution covers the increase that we observe in the data. These shocks caused beliefs to drift upward (top panel), increasing the probability of an escape-provoking event (third panel), and leading to the hyperinflation observed (bottom panel). The importance of shocks in 1970s Chile is consistent with the instability that accompanied the Allende government and the subsequent coup.

Because the buildup in inflation was largely driven by shocks to the seignorage, the stabilization in the late 1970s is interpreted as a reduction of variance of these shocks. Beliefs drift continually downward after 1978 (top panel), and inflation remains relatively low throughout the rest of the sample (bottom panel), apart from a short-lived spike around 1985. Cosmetic reforms play little role for Chile, as their probabilities remain very low even during the runaway inflation period. Thus, Chile is again an example of the importance of fiscal policy for inflation. But although fiscal reforms play some role in bringing down hyperinflation in the 1970s, seignorage shocks are the driving force in the conquest of Chilean inflation. After the tumult of the 1970s the economy engaged in a more stable fiscal policy, resulting in relatively stable inflation.

IX. COMPARING STEADY STATE, SCEs, AND REEs

We have examined what our estimates imply for the self-confirming and rational expectations equilibria discussed above and defined formally in appendices C and D. We have emphasized the parts of our story for the dynamics of hyperinflation that require retreating from rational expectations. However, in many cases the retreat is relatively minor. In particular, as shown in Sargent, Williams, and Zha (2006a), the fixed- m SCEs discussed in the previous sections are close to the REE beliefs. In most cases, the differences are well within half of a percentage point, with the largest differences being slightly more than one percentage point. Similarly, unconditional SCEs and REEs are very close as well. Appendix D shows that the fixed- $m - v$ SCEs with the low volatility states are also very close to the deterministic steady states $\pi_1^*(m)$. That justifies our use of $\pi_1^*(m)$ in defining cosmetic reforms.

X. ROBUSTNESS

As discussed above, we have tried many different specifications of our model that do not fit as well as our baseline specification. Here we report one change that does matter. We also discuss the implications of our results for quarterly seignorage data.

X.1. State-dependent gains. We re-estimate the model with both of the state-dependent learning rules described in Section II.2.3. For countries other than Argentina, neither rule gives the model a better fit to the data than the constant-gain rule. This is not surprising because inflation can decline drastically even when beliefs adjust slowly, as shown in Section VI. For Argentina, however, both state-dependent rules improve the fit significantly. The best-fit model is a 2×3 Markov process with

the rule (9); the log likelihood increases to 1338.7 from 1275.4 under the constant-gain rule, which, according to Table 3, makes the economic model's fit competitive with the best-fit statistical model. The elasticity parameter λ is estimated to be 0.83, somewhat higher than the estimate in the constant-gain case. The gain for the high-seigniorage state is estimated to 0.052, much higher than the estimated gain (0.023) in the constant-gain case, implying that agents discount the past inflation data faster during the high-seigniorage periods. The estimate for the gain parameter during the low-seigniorage state is 0.013, lower than the estimate of the constant gain.

Other estimated results are plotted in Figure 8. In the left panel on the first row of graphs, the dotted lines represent SCEs conditional on the prevailing seigniorage state m and the solid line represents the evolution of beliefs β_t . In the high-seigniorage state, the two SCEs (high and low) are the same (marked by the middle dotted line). But in the low-seigniorage state, the high and low SCEs are far apart.

During the high-seigniorage periods, as beliefs approach the high SCE represented by the middle dotted line, escape-provoking events become very likely. Beliefs increase drastically after 1988, thanks to the large gain. After the largest increase of inflation in July 1989, cosmetic reforms take place and both inflation and beliefs are reset in August 1989. Shortly after, the second largest inflation rate occurs in March 1990 because seigniorage remains high. The high inflation rate did not last due to what the model identifies to be favorable fiscal conditions consisting of an initially lower volatility of fiscal shocks, followed by fundamental fiscal reforms. The left panel in the second row of graphs shows this pattern of seigniorage rates backed out from the estimated model. Again, these model predictions are remarkably consistent with the data. The right panel in the first row of Figure 8 displays one-step predictions of inflation rates, which track both hyperinflation and low inflation well. The low-inflation data, like the high-inflation data, are crucial in helping to identify factors that determine the rise and fall of hyperinflation.

X.2. Quarterly data on seigniorage. Because changes in prices are rapid during a hyperinflation period, it would be informative to compare our model's predictions with high-frequency data on seigniorage. For many countries studied in this paper, however, it is difficult or even impossible to find reliable quarterly data on seigniorage. While monthly or quarterly data on nominal money is generally available, there is little quarterly data on GDP. The International Finance Statistics (IFS), the arguably most reliable data source compiled by International Monetary Fund (IMF), has quarterly GDP data for Argentina from 1993Q1 on, Bolivia from only 1995Q1 on, and for Brazil from 1991Q1 on. We were able to obtain some more historical data for Argentina going back to 1970Q1 from the Ministry of Economy and Production in Argentina.¹⁰ We use this data set on GDP to compute quarterly seigniorage rates. As shown in the right panel on the second row of Figure 8, the model's predictions

¹⁰We are grateful to Juan Pablo Nicolini for his help on collecting an additional data set on the quarterly Argentinean GDP series.

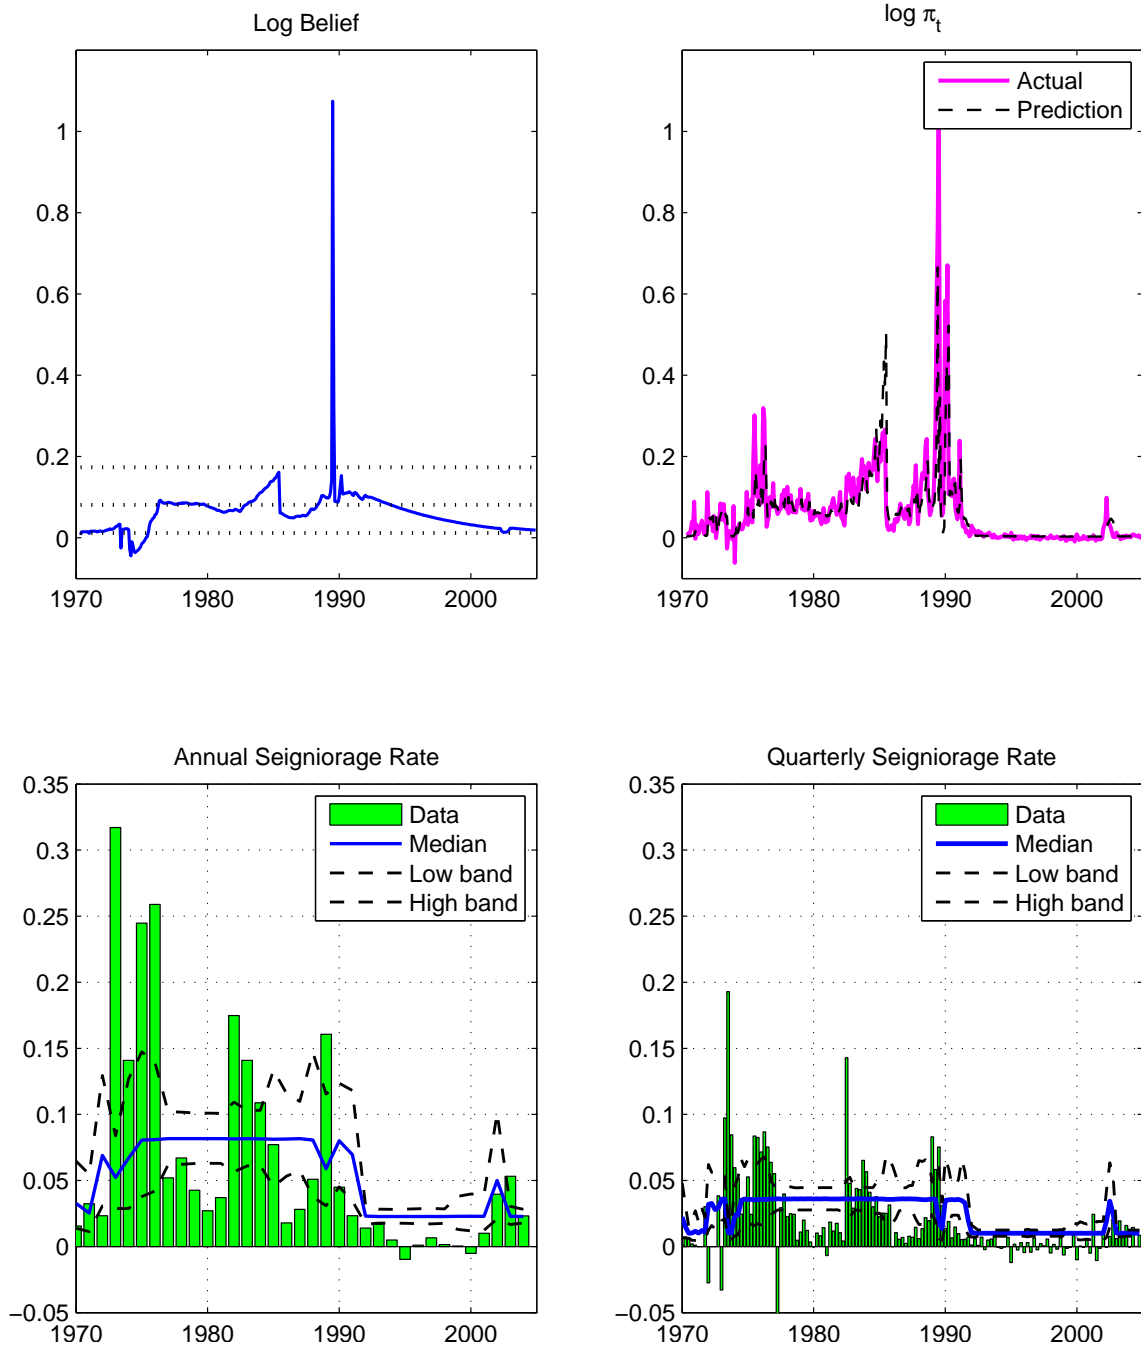


FIGURE 8. Argentina: the estimated results from the best (2×3) model with the state-dependent learning rule (9). The dotted lines in the left panel on the first row of graphs represent SCEs conditional on the prevailing seigniorage m state.

(normalized around the historical mean implied by the data), along with 68% error bands, are consistent with the high-frequency data.

Comparing annual seigniorage rates, used by Fischer (1982), with quarterly seigniorage rates computed from our additional data set, one can see from Figure 8 that the high-low patterns are quite similar, although the overall scale of quarterly rates tends to be lower than annual rates. The pattern is what matters for our model because the extra parameter γ is free to adjust to control the overall scale, as shown in Proposition 2 in Appendix B. Moreover, it is not always true that the magnitude of annual seigniorage rates is significantly higher than that of quarterly rates, even in high-money-creation periods when prices change drastically. For example, we obtained some limited quarterly data from Brazil. In 1993, the annual seigniorage rate was 0.1051 while the quarterly rates ranged from 0.0763 in Q1 to 0.1274 in Q4.

Another problem relates to considerable uncertainty surrounding the quality of quarterly data on both money and GDP. For example, using the data reported by the Ministry of Economy and Production in Argentina, the seigniorage rate for 1989Q2 is only 5.8%. But according to Ahumada, Canavese, Sanguinetti, and Sosa (1993) and Ahumada, Canavese, Alvaredo, and Di Tella (2000) who use different data definitions and sources, the quarterly rate for 1989Q2 is computed to be as high as 10.5%, and within this quarter the maximum seigniorage rate is estimated to reach 15.7%. Despite the dispersion in magnitude, however, the overall pattern of high and low seigniorage rates tends to be consistent across different data frequencies and data sources.

XI. CONCLUDING REMARKS

Our empirical results identify episodes in which different causes sparked big rises and falls in inflation. Table 4 briefly summarizes some of the key empirical patterns contained in Figures 3- 7 from section VIII. Our model tells us that inflation can rise because of changes in the fundamentals, whether through high seigniorage levels or high shock variances or both, or through potentially explosive expectation dynamics caused by “escape provoking” events. The two columns in the table sort episodes into inflations that were driven by such escape provoking events or solely by alterations in the fiscal fundamentals. Escapes are said to occur when inflationary expectations escape from a state-dependent domain that attracts inflationary expectations to a state-dependent self-confirming equilibrium pinned down by the fiscal fundamentals. The rows of table 4 indicate three possible ways that our model tells us hyperinflation can be stopped: a superficial monetary reform that mechanically resets inflation without altering the seigniorage state, a fundamental fiscal reform activated by a change in the Markov state of seigniorage level, and no reform in the seigniorage level but a change in the conditional shock variance state.

We have used joint densities evaluated at maximum likelihood parameter values for each country to assign episodes to appropriate boxes in the table. As a rule for making these assignments, we categorize an episode according to whether our model assigns high probabilities (i.e., over 60%) of an escape-provoking event or of a cosmetic reform. In March of 1990, for example, Brazilian inflation reached its peak with a

TABLE 4. Causes for the rise and fall of *hyperinflation* across countries

	Escape-provoking	Fundamentals
Cosmetic reform	Brazil (87-91) Peru (87-92)	
Fundamental fiscal reform	Argentina (87, 90-91) Brazil (92-95)	Bolivia (82-86)
Reform only in conditional variance, not level, of signiorage	Chile (71-78)	Argentina (75-76, 76-77, 83-86)

monthly gross rate of 1.82. In the next two months, the inflation rate dropped to 1.15 then 1.07. We estimated the probability of a cosmetic monetary reform to be 67.5% for March, 75.6% for April, and 47.8% for May. Peru is another example. In August of 1990, the Peruvian monthly inflation rate reached 4.97, was brought down to 1.14 in September, and stayed at a relatively low level around 1.1 for a number of months thereafter. This unusually volatile fluctuation enables us to estimate the probability of a cosmetic reform that is only 10.8% in August but jumps to 100% in September. Expected inflation had become so high that it rendered a cosmetic reform inevitable. Thus, a cosmetic reform along the lines introduced by Marcet and Nicolini (2003) generated the Peruvian stabilization.

Economic stories vary across other episodes in our sample. Our estimates indicate that the high and volatile inflation episode in Brazil finally ended in 1994 only with a sustained fiscal reform. For Argentina (from 1987 to 1991) and Bolivia, fiscal reforms also played a dominant role in conquering hyperinflation. For Argentina (from 1976 to 1986) and Chile, reductions in the variance of shocks to seigniorage made essential contributions to ending the hyperinflations.

In conclusion, our econometrics suggest that adjustments in levels and conditional volatilities of monetized deficits seem to have stabilized inflation processes in most of the hyperinflations, with a notable exception in Peru where a cosmetic reform of the type emphasized by Marcet and Nicolini (2003) seems to have been at work.

APPENDIX A. STEADY STATES OF DETERMINISTIC MODEL

We now report equilibria from a perfect foresight version of the model where agents observe and condition on the seigniorage state m . We work with a deterministic version of model (1) - (5) obtained by fixing the state $m_t = m \in \{1, \dots, m_h\}$ and setting η_{dt} to zero for all t . Such equilibria are useful reference points in the analysis of our stochastic adaptive model. There are two steady states associated with each m .

Proposition 1. If

$$\bar{d}(m) < 1 + \theta\lambda - 2\sqrt{\theta\lambda}, \quad (\text{A1})$$

then there exist two steady state equilibria for π_t :

$$\pi_1^*(m) = \frac{(1 + \theta\lambda - \bar{d}(m)) - \sqrt{(1 + \theta\lambda - \bar{d}(m))^2 - 4\theta\lambda}}{2\lambda}, \quad (\text{A2})$$

$$\pi_2^*(m) = \frac{(1 + \theta\lambda - \bar{d}(m)) + \sqrt{(1 + \theta\lambda - \bar{d}(m))^2 - 4\theta\lambda}}{2\lambda}. \quad (\text{A3})$$

Proof. Sargent and Wallace (1987) show that

$$\pi_t = (\lambda^{-1} + \theta - \bar{d}(m)\lambda^{-1}) - \frac{\theta}{\lambda\pi_{t-1}}.$$

In stationary equilibrium, $\pi_t = \pi_{t-1}$. Substituting this equality into the above equation leads to (A2) and (A3). \square

We shall impose (A1) in our empirical work. Note that the maximum value that $\bar{d}(m)$ can take and still have a steady state (SS) inflation rate exist is $1 + \theta\lambda - 2\sqrt{\theta\lambda}$. When $\bar{d}(m)$ attains this maximum value, the two SS inflation rates both equal

$$\pi_{\text{SS}}^{\text{max}} \equiv \sqrt{\frac{\theta}{\lambda}}. \quad (\text{A4})$$

Proposition 1 tells us that a steady state rational expectations equilibrium (REE) inflation rate is bounded above by $1/\lambda$ and this bound is attained when $\bar{d} = 0$. Eckstein (1987) noted the peculiar property that among stationary rational expectations equilibria, the biggest inflations occur when the budget deficit is zero. This perverse property does not characterize our learning model with adaptive expectations specified in Section II.3. In our model, it is only the belief that is bounded by $1/\lambda$; inflation itself can escalate beyond $1/\lambda$ up to $1/\delta$, because we are free to set δ to be arbitrarily small such that $\lambda \geq \delta$.

APPENDIX B. THE LIKELIHOOD

B.1. Normalization. The model (1)-(5) makes inflation dynamics depend on $\gamma d_t(s)$, where $s \in \{1, \dots, h\}$, and not on the individual parameters γ and $d_t(s)$ separately. Therefore, we have

Proposition 2 (Normalization). The dynamics of π_t are unchanged if both $d_t(s)$ and $1/\gamma$ are normalized by the same scale.

Proof. Let $d_t(s)$ and $1/\gamma$ be multiplied by any real scalar κ . If we redefine P_t to be P_t/κ , the system (1)-(5) remains unaffected. The redefinition of the price level simply means that the price index is re-based, which affects the dynamics of neither M_t nor π_t . \square

The normalization is effectively a choice of units for the price level, about which our model is silent because we deduce a joint density over inflation sequences only. Proposition 2 explains why we deviate from the procedure of Marcat and Nicolini (2003), who treated γ and $\bar{d}(m)$ as separate parameters, and who interpreted the

calibrated value of d_t as measuring monetized deficits as a share of GDP. With only inflation data these parameters cannot be identified separately, so that re-normalizing them in the manner of Proposition 2 gives the same equilibrium outcome.¹¹ Therefore we normalize $\gamma = 1$ when maximizing the likelihood function. After we have estimated the free parameters, we re-normalize γ to match the pertinent country's price *level* for the purpose of computing estimates of d_t to compare with some annual seignorage data in section VIII. It is important to note that such normalization affects only the mean of $\log d_t$ or the median $\bar{d}(m_t)$, but not the standard deviation of $\log d_t$.

We first derive a likelihood conditional on the hidden composite states $s_t = [m_t v_t]$ and then integrate over states to find the appropriate unconditional likelihood. Recall that we have specified the distributions for η_{dt} and $\eta_{\pi t}$ in (6) and (18). Denote

$$\begin{aligned} s^t &= \{s_1, \dots, s_t\}, \\ \xi_d(s_t) &= 1/\sigma_d(s_t), \\ \xi_\pi &= 1/\sigma_\pi, \end{aligned}$$

and again let ϕ be a collection of all structural parameters. We use the tilde above $\eta_{dt}(s_t)$ to indicate that $\tilde{\eta}_{dt}(s_t)$ is a *random* variable, whereas $\eta_{dt}(s_t)$ is the realized value associated with π_t . The following proposition provides the key component of the overall likelihood function.

Proposition 3. Given the pdfs (18) and (6), the conditional likelihood is

$$\begin{aligned} p(\pi_t | \pi^{t-1}, s^T, \phi) &= p(\pi_t | \pi^{t-1}, s_t, \phi) \\ &= C_{1t} \frac{|\xi_\pi| \exp \left[-\frac{\xi_\pi^2}{2} (\log \pi_t - \log \pi_1^*(s_t))^2 \right]}{\sqrt{2\pi} \Phi(|\xi_\pi|(-\log(\delta) - \log(\pi_1^*(s_t))) \pi_t)} \\ &\quad + C_{2t} \left(\frac{\theta |\xi_d(s_t)| (1 - \lambda \beta_{t-1})}{\sqrt{2\pi} [(1 - \lambda \beta_t) \pi_t - \theta(1 - \lambda \beta_{t-1})] \pi_t} \right. \\ &\quad \left. \exp \left[-\frac{\xi_d^2(s_t)}{2} \left[\log[(1 - \lambda \beta_t) \pi_t - \theta(1 - \lambda \beta_{t-1})] - \log \pi_t - \log d(s_t) \right]^2 \right] \right), \end{aligned} \tag{A5}$$

where

$$\begin{aligned} C_{1t} &= \iota \{ \beta_{t-1} \geq 1/\lambda \} + \iota \{ \beta_{t-1} < 1/\lambda \} \\ &\quad \left(1 - \Phi \left[|\xi_d(s_t)| (\log(\max[(1 - \lambda \beta_t) - \delta \theta(1 - \lambda \beta_{t-1}), 0]) - \log d(s_t)) \right] \right), \\ C_{2t} &= \iota \{ \beta_{t-1} < 1/\lambda \} \iota \left\{ \frac{\theta(1 - \lambda \beta_{t-1})}{\max(1 - \lambda \beta_t, \delta \theta(1 - \lambda \beta_{t-1}))} < \pi_t < \frac{1}{\delta} \right\}. \end{aligned}$$

Proof. We need to prove that

$$\int_0^{1/\delta} p(\pi_t | \pi^{t-1}, s_t, \phi) d\pi_t = 1.$$

¹¹For a general discussion of normalization in econometrics, see Hamilton, Waggoner, and Zha (2007).

One can show from (18) and (6) that the right side of equation (A5) is equivalent to

$$\begin{aligned} & \iota \{ \beta_{t-1} \geq 1/\lambda \} p_\pi(\pi_t - \pi_1^*(s_t) | s_t) + \iota \{ \beta_{t-1} < 1/\lambda \} \\ & \left[\iota \left\{ \frac{\theta(1 - \lambda\beta_{t-1})}{\max(1 - \lambda\beta_t, \delta\theta(1 - \lambda\beta_{t-1}))} < \pi_t < \frac{1}{\delta} \right\} p_d(\eta_{dt}(s_t) | s_t) \frac{d\eta_{dt}(s_t)}{d\pi_t} \right. \\ & \quad \left. + \Pr[\tilde{\eta}_{dt}(s_t) \geq \bar{\omega}_t(s_t)] p_\pi(\pi_t - \pi_1^*(s_t) | s_t) \right], \end{aligned}$$

where $\Pr[\cdot]$ is the probability that the event in the brackets occurs.

Consider the case where $\beta_{t-1} < 1/\lambda$ (the other case is trivial). Denote

$$L_t = \frac{\theta(1 - \lambda\beta_{t-1})}{\max(1 - \lambda\beta_t, \delta\theta(1 - \lambda\beta_{t-1}))}.$$

It follows that

$$\begin{aligned} & \int_0^{1/\delta} p(\pi_t | \pi^{t-1}, s_t, \phi) d\pi_t \\ &= \int_{L_t}^{1/\delta} p_d(\eta_{dt}(s_t) | s_t) \frac{d\eta_{dt}(s_t)}{d\pi_t} d\pi_t + \Pr[\tilde{\eta}_{dt}(s_t) > \bar{\omega}_t(s_t)] \int_0^{1/\delta} p_\pi(\pi_t - \pi_1^*(s_t) | s_t) d\pi_t \\ &= \int_{-\bar{d}(s_t)}^{\bar{\omega}_t(s_t)} p_d(\eta_{dt}(s_t) | s_t) d\eta_{dt}(s_t) + \Pr[\tilde{\eta}_{dt}(s_t) \geq \bar{\omega}_t(s_t)] \\ &= \Pr[\tilde{\eta}_{dt}(s_t) < \bar{\omega}_t(s_t)] + \Pr[\tilde{\eta}_{dt}(s_t) \geq \bar{\omega}_t(s_t)] = 1. \end{aligned}$$

□

After integrating out s^T , the overall likelihood is

$$p(\pi^T | \phi) = \prod_{t=1}^T p(\pi_t | \pi^{t-1}, \phi) = \prod_{t=1}^T \left\{ \sum_{s_t=1}^h [p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)] \right\}, \quad (\text{A6})$$

where

$$\Pr(s_t | \pi^{t-1}, \phi) = \sum_{s_{t-1}=1}^h [\Pr(s_t | s_{t-1}, q) \Pr(s_{t-1} | \pi^{t-1}, \phi)]. \quad (\text{A7})$$

The probability $\Pr(s_{t-1} | \pi^{t-1}, \phi)$ can be updated recursively. We follow Sims, Waggoner, and Zha (forthcoming) and set

$$\Pr(s_0 | \pi^0, \phi) = 1/h.$$

For $t = 1, \dots, T$, the updating procedure involves the following computation:

$$\Pr(s_t | \pi^t, \phi) = \frac{p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)}{\sum_{s_t=1}^h [p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)]}. \quad (\text{A8})$$

As shown in Sims, Waggoner, and Zha (forthcoming), one can also use the above recursive structure to compute the smoothed probability of s_t , $\Pr(s_t | \pi^T, \phi)$.

APPENDIX C. SELF-CONFIRMING EQUILIBRIA AND MEAN DYNAMICS

This appendix describes how to compute the self-confirming equilibria (SCE) defined in section III, as well as formally defining the mean dynamics which lead to them.

C.1. Unconditional SCEs. Definition III.1 provides the moment condition which an SCE must satisfy. We now formally compute it and the corresponding mean dynamics G discussed in section III. Let:

$$\omega(\beta_t, \beta_{t-1}) = 1 - \lambda\beta_t - \delta\theta(1 - \lambda\beta_{t-1}).$$

Given the model and the specifications of cosmetic reforms in equation (17), we have:

$$\pi_t = \mathbf{1}(d_t(s_t) < \omega(\beta_t, \beta_{t-1})) \frac{\theta(1 - \lambda\beta_{t-1})}{1 - \lambda\beta_t - d_t(s_t)} + \mathbf{1}(d_t(s_t) \geq \omega(\beta_t, \beta_{t-1})) \pi_t^*(s_t).$$

Hence, we can write the learning rule (7) as:

$$\beta_{t+1} = \beta_t + \varepsilon g(\beta_t, \beta_{t-1}, d_t, \pi_t^*) \quad (\text{A9})$$

where

$$g(\beta_t, \beta_{t-1}, d_t, \pi_t^*) = \mathbf{1}(d_t < \omega(\beta_t, \beta_{t-1})) \frac{\theta(1 - \lambda\beta_{t-1})}{1 - \lambda\beta_t - d_t} + \mathbf{1}(d_t \geq \omega(\beta_t, \beta_{t-1})) \pi_t^*(s_t) - \beta_t.$$

We then define the following terms:

$$\begin{aligned} \tilde{g}(\beta, d_t, \pi_t^*) &= g(\beta, \beta, d_t, \pi_t^*), \\ \Psi_s(\beta, b) &= \int_0^{b - \bar{d}(s)} \frac{1}{1 - \lambda\beta - \bar{d}(s) - \eta} dF_d(\eta|s). \end{aligned}$$

Here we abuse notation and let \bar{d} depend on the composite state s rather than m alone. It follows that $\tilde{\omega}(\beta) \equiv \omega(\beta, \beta) = (1 - \delta\theta)(1 - \lambda\beta)$ and that $\Psi_s(\beta, b)$ is finite as $b \rightarrow 1 - \lambda\beta$ because δ in equation (12) is bounded away from zero.

Let \bar{q}_s denote the unconditional probability of the event $\{s_t = s\}$, as implied by the ergodic distribution of Q . Then we can write the unconditional expectation:

$$\begin{aligned} G(\beta) &\equiv E[\tilde{g}(\beta, d_t, \pi_t^*)] \\ &= \sum_{k=1}^h \left[\int_0^{\tilde{\omega}(\beta) - \bar{d}(k)} \frac{\theta(1 - \lambda\beta)}{1 - \lambda\beta - \bar{d}(k) - \eta} dF_d(\eta|k) + [1 - F_d(\tilde{\omega}(\beta)|k)] \bar{\pi}^*(k) \right] \bar{q}_k - \beta \\ &= \sum_{k=1}^h \left\{ \theta(1 - \lambda\beta) \Psi_k(\beta, \tilde{\omega}(\beta)) + \left[1 - \Phi \left(\frac{\log \tilde{\omega}(\beta) - \log \bar{d}(k)}{\sigma_d(k)} \right) \right] \bar{\pi}^*(k) \right\} \bar{q}_k - \beta \end{aligned}$$

Note then from definition III.1 that an unconditional SCE is a β satisfying $G(\beta) = 0$. Such an unconditional SCE by characterizes the limit of beliefs of the adaptive agents under our model in a way made precise in the following result.

Proposition 4. As $\varepsilon \rightarrow 0$ the beliefs $\{\beta_t\}$ from (A9) converge weakly to the solution of the ordinary differential equation (ODE):

$$\dot{\beta} = G(\beta) \quad (\text{A10})$$

for $\delta > 0$ and a broad class of probability distributions of $\eta_{dt}(s_t)$ and $\eta_{\pi t}(m_t)$ (including those specified in (18) and (6)).

Proof. Under our assumptions about distributions and the truncation rule, this follows from Kushner and Yin (1997). \square

The ODE (A10) describes the *mean dynamics* G , meaning that for small gains ε , the belief trajectories tend to track those of the ODE. Moreover as $t \rightarrow \infty$ the ODE (A10) converges to an SCE which is stable, i.e. which satisfies $G'(\beta) < 0$.

C.2. Conditional SCEs. As discussed in section III we are also interested in those SCEs which assume that the seignorage state m (and possibly the volatility state v) remain fixed for all time. Thus, we now formally define the *conditional mean dynamics* $\hat{G}(\beta, m)$. Let $\bar{q}_{v,k}$ denote the unconditional probability of the event $\{v_t = k\}$, which is an element of the ergodic distribution of Q_v . Then define

$$\begin{aligned} \hat{G}(\beta, m) &\equiv E[\tilde{g}(\beta, d_t(m_t, v_t), \pi_t^*) | m_t = m \forall t] \\ &= \sum_{k=1}^{v_h} \left[\int_0^{\tilde{\omega}(\beta) - \bar{d}(m)} \frac{\theta(1 - \lambda\beta)}{1 - \lambda\beta - \bar{d}(m) - \eta} dF_d(\eta | [m, k]) \right] \bar{q}_{v,k} \\ &\quad + \sum_{k=1}^{v_h} [1 - F_d(\tilde{\omega}(\beta) - \bar{d}(m) | [m, k])] \bar{\pi}^*(m) \bar{q}_{v,k} - \beta \\ &= \sum_{k=1}^{v_h} \left\{ \theta(1 - \lambda\beta) \Psi_{[m,k]}(\beta, \tilde{\omega}(\beta)) + \left[1 - \Phi \left(\frac{\log \tilde{\omega}(\beta) - \log \bar{d}(m)}{\sigma_d(k)} \right) \right] \bar{\pi}^*(m) \right\} \bar{q}_{v,k} - \beta. \end{aligned}$$

Then the fixed- m SCEs from definition III.2 satisfy $G(\beta(m), m) = 0$. The fixed- m - v SCEs from definition III.3 can be characterized in a similar manner. In section IX we show that the conditional SCE beliefs give good approximations to rational expectations beliefs under our estimated parameters.

We show how such conditional SCEs characterize beliefs by taking a two time-scale limit in which the gain ε goes to zero but the probabilities of switching m states go to zero at a faster rate. For simplicity, we suppose that case when $m_t \in \{1, 2\}$. Then note that we can write:

$$E_t m_{t+1} = m_t + Q_m(-m, m)(3 - 2m_t)$$

where $Q_m(-m, m)$ is the off-diagonal element of column m of Q_m . Therefore we can write the evolution of m_t as:

$$m_{t+1} = m_t + Q_m(-m, m)(3 - 2m_t) + v_{t+1}$$

where $E_t v_{t+1} = 0$. Now we consider a slow variation limit where $Q_m \rightarrow I$, and thus we scale $Q_m(m, -m)$ by a small parameter α , which also implies that the martingale difference term v_{t+1} inherits the scaling. Thus, we extend the system that we analyze from (A9) to:

$$\begin{aligned} \beta_{t+1} &= \beta_t + \varepsilon g(\beta_t, \beta_{t-1}, d_t(m_t, v_t), \pi_t^*) \\ m_{t+1} &= m_t + \alpha [Q_m(-m, m)(1 - 2m_t) + v_{t+1}] \end{aligned} \tag{A11}$$

Following Tadić and Meyn (2003) we consider the limit where $\varepsilon \rightarrow 0$ and $\alpha \rightarrow 0$ but where $\alpha \ll \varepsilon$. In this limit, m_t varies more slowly than the beliefs β_t , and thus we can effectively treat the state m as fixed.

Proposition C.1. *Let $\varepsilon \rightarrow 0$ and $\alpha \rightarrow 0$ such that $\alpha/\varepsilon \rightarrow 0$ and $\varepsilon^{3/2}/\alpha \rightarrow 0$. Then for $m_0 = m$ the beliefs $\{\beta_t\}$ from (A11) converge weakly to the solution of the ordinary differential equation (ODE):*

$$\dot{\beta} = \hat{G}(\beta, m) \quad (\text{A12})$$

for $\delta > 0$ and a broad class of probability distributions of $\eta_{dt}(s_t)$ and $\eta_{\pi t}(s_t)$ that include those specified in (18) and (6).

Proof. (Sketch.) This follows from Corollary 2 to Theorem 2 in Tadić and Meyn (2003). Since they focus on convergence of the “slow” process m_t as well they require stronger stability conditions on the ODE than are necessary for our result. \square

The ODE (A10) describes the *conditional mean dynamics* G , meaning that for small gains ε and persistent Markov chains Q_m , the belief trajectories tend to track those of the ODE. Moreover as $t \rightarrow \infty$ the ODE (A12) converges to a stable fixed- m SCE, i.e. which satisfies $\hat{G}_\beta(\beta(m), m) < 0$.

APPENDIX D. RATIONAL EXPECTATIONS EQUILIBRIA

We now suspend the adaptive learning rule (7) and consider a subset of the rational expectations equilibria of the model. Further, while previously we’ve assumed that agents do not observe the seignorage state m_t , we now assume that rational agents do condition on it.

D.1. Computing equilibria. We seek stationary Markov equilibria in which inflation and expected inflation are given by:

$$\begin{aligned} \pi_t &= \pi(s_t, m_{t-1}, d_t) \\ E[\pi_{t+1} | m_t] &= E[\pi(s_{t+1}, m_t, \bar{d}(m_{t+1}) + \eta_{d,t+1}(s_{t+1})) | m_t] \\ &= \sum_{j=1}^{m_h} \sum_{k=1}^{v_h} \int \pi([j, k], m_t, \bar{d}(j) + \eta) dF_d(\eta | [j, k]) \bar{q}_{v,k} Q_m(m_t, j) \\ &\equiv \pi^e(m_t). \end{aligned}$$

Note that we assume that the volatility state v_t is unobserved and that agents’ subjective distribution over this state is given by the ergodic distribution \bar{q}_v . Then going through calculations similar to those above we have:

$$\pi(s_t, m_{t-1}, d_t) = \frac{\theta(1 - \lambda\pi^e(m_{t-1}))}{1 - \lambda\pi^e(m_t) - d_t(s_t)}.$$

This only holds when the denominator is positive, so we truncate as in section IV, giving:

$$\begin{aligned} \pi(s_t, m_{t-1}, d_t) = & \mathbf{1}(d_t(s_t) < \omega(\pi^e(m_t), \pi^e(m_{t-1}))) \frac{\theta(1 - \lambda\pi^e(m_{t-1}))}{1 - \lambda\pi^e(m_t) - d_t(s_t)} \\ & + \mathbf{1}(d_t(s_t) \geq \omega(\pi^e(m_t), \pi^e(m_{t-1}))) \pi_t^*(m_t) \end{aligned}$$

Letting $\omega_{ij} = \omega(\pi^e(j), \pi^e(i))$ and taking expectations of both sides conditional on information at $t - 1$ and setting $m_{t-1} = i$ yields:

$$\pi^e(i) = \sum_{j=1}^{m_h} \sum_{k=1}^{v_h} \left\{ \theta(1 - \lambda\pi^e(i)) \Psi_{[j,k]}(\pi^e(j), \omega_{ij}) + \left[1 - \Phi \left(\frac{\log(\omega_{ij}) - \log \bar{d}(j)}{\sigma_d(k)} \right) \right] \bar{\pi}^*(j) \right\} \bar{q}_{v,k} Q_m(i, j). \quad (\text{A13})$$

Thus, we have m_h coupled equations determining $\pi^e(m_t)$. Substituting this solution into the expression for $\pi(\cdot)$ then gives the evolution of inflation under rational expectations. The equations are sufficiently complicated that an analytic solution is not available, and hence we must look for equilibria numerically. A simple iterative solution method for the equations consists of initializing the $\pi^e(j)$ on the right side of (A13) and computing $\pi^e(i)$ on the left side and iterating until convergence.

We typically find that there are two conditional SCEs in each state. Loosely speaking, REEs average across the conditional SCEs, taking into account the probability of state switches. So, there are typically four REEs that switch between values close to the conditional SCEs in each state. However, when shocks to seignorage become large enough there may be only one conditional SCE in a state, or a conditional may fail SCE to exist altogether. Depending on the weight that these high-shock states have in the invariant distribution, the unconditional SCE may also fail to exist. Similarly, there may be fewer rational expectations equilibria or none at all.

D.2. Comparing steady states, SCEs, and REEs. For each country, Table 5 reports the deterministic SS from appendix A, and fixed- m - v SCEs of the definition III.3 type, and the low inflation REEs closest to these SCEs. We focus on the SCEs conditional on each state m but always with the low volatility state v . These SCEs are of particular interest because they are likely to be close to the low SS inflation rates that we in use to defining our cosmetic reforms. The table shows that this set of conditional SCEs are very close to the SS inflation rates, justifying the way we define the cosmetic reforms based on the low SS inflation rates. The reported SCEs are close to the REEs but do differ somewhat, particularly in countries other than Brazil. This is at least partly because the REE beliefs have more mean reversion, in that they recognize that the economy will eventually switch to another state, while the SCEs do not.

APPENDIX E. SEIGNORAGE RATES: ACTUAL DATA AND MODEL IMPLICATIONS

Because we have no reliable data on real output and money on a monthly basis, we construct a time series of annual deficits financed by money creation. Following

TABLE 5. Steady State (SS) of nonstochastic model, Conditional SCEs, and REEs: Log values of beliefs and low deterministic SS inflation rates in each state m coupled with the low volatility state v . The symbol * means “not applicable.”

Country	SS/SCE/REE	Low m and low v	Medium m and low v	High m and low v
Peru	SS	0.0028	0.0108	0.0498
Peru	SCE	0.0035	0.0120	0.0543
Peru	REE	0.0183	0.0330	0.0833
Argentina	SS	0.0052	*	0.0687
Argentina	SCE	0.0057	*	0.0737
Argentina	REE	0.0126	*	0.0768
Bolivia	SS	0.0052	*	0.1115
Bolivia	SCE	0.0059	*	0.1203
Bolivia	REE	0.0083	*	0.1356
Brazil	SS	0.0057	0.0592	0.2063
Brazil	SCE	0.0057	0.0595	0.2126
Brazil	REE	0.0060	0.0690	0.2123
Chile	SS	0.0044	*	0.0198
Chile	SCE	0.0046	*	0.0202
Chile	REE	0.0393	*	0.0586

Fischer (1982), we calculate annual seigniorage rates from actual data as

$$d_{A,t}^{\text{Data}} = \frac{M_{A,t}^{\text{Agg}} - M_{A,t-1}^{\text{Agg}}}{Y_{A,t}^{\text{Agg}}} \quad (\text{A14})$$

where the subscript “A” stands for annual and the superscript “Agg” stands for aggregate. $M_{A,t}^{\text{Agg}}$ is aggregate reserve money for the year containing the month indexed by t and $Y_{A,t}^{\text{Agg}}$ is aggregate nominal GDP in that year. For this calculation, there is no parameter θ involved because we work directly on the aggregate data on money.

To make the simulated data from our model as close to (A14) as possible, we compute the distribution of $d_{A,t}$ as follows. We first draw s_t from $\Pr(s_t|\hat{\phi}, \pi^T)$ and for a given s_t we then draw $d_t(s_t)$ and compute $d_{A,t}$ as an average of $d_t(s_t)$ over the twelve months of the year. The simulated data $d_{A,t}$ only approximate the actual data $d_{A,t}^{\text{Data}}$ because of the following differences. The price index data P_t used for our model is CPI, not the GDP deflator. For the actual data, $d_{A,t}^{\text{Data}}$ is calculated as a ratio of two sums or aggregates. For the simulated data, $d_{A,t}^{\text{Data}}$ is computed as a sum of monthly money creations in percent of real output.

In our estimation, d_t is arbitrarily normalized. When comparing to actual data, we need to re-normalize it. We do so by matching the average of medians of simulated annual seignorage to the average of actual seignorage rates over the sample for Argentina, Bolivia, Brazil, and Peru. For Chile, we use the average over the sample excluding the hyperinflation period 1971-1975 during which large simulated

seignorage levels are caused by a large shock variance. The effect of this relatively large variance is shown by the skewed distribution marked by the dashed bands in the second-row graph of Figure 7. Note that changes in shock variances have no effect on the median of simulated seignorage.

REFERENCES

- AHUMADA, H., A. CANAVESE, F. G. ALVAREDO, AND I. T. DI TELLA (2000): "Un Análisis Comparativo del Impacto Distributivo del Impuesto Inflacionario y de un Impuesto Sobre el Consumo," *Económica*, XLVI(2), 3–36.
- AHUMADA, H., A. CANAVESE, P. SANGUINETTI, AND W. SOSA (1993): "Efectos Distributivos del Impuesto Inflacionario: Una Estimación del Caso Argentino Are Hyperinflation Paths Learnable?," *Economía Mexicana*, II(2), 329–85.
- BRUNO, M., AND S. FISCHER (1990): "Seigniorage, Operating Rules, and the High Inflation Trap," *Quarterly Journal of Economics*, May, 353–374.
- CAGAN, P. (1956): "The Monetary Dynamics of Hyperinflation," in *Studies in the Quantity Theory of Money*, ed. by M. Friedman. University of Chicago Press, Chicago, IL.
- CARBONI, G., AND M. ELLISON (2007): "Learning and the Great Inflation," CEPR Discussion Papers 6250, C.E.P.R. Discussion Papers.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.
- CHUNG, H. (1990): *Did Policy Makers Really Believe in the Phillips Curve? An Econometric Test*. University of Minnesota, Minneapolis, USA, ph.d. dissertation edn.
- DORNBUSCH, R. (1985): "Comment," in *Inflation and indexation: Argentina, Brazil, and Israel*, ed. by J. Williamson. MIT Press, Cambridge, MA.
- ECKSTEIN, Z. (1987): "Comments on Thomas Sargent and Neil Wallace, 'Inflation and the Government Budget Constraint'," in *Economic Policy in Theory and Practice*, ed. by A. Razin, and E. Sadka, pp. 201–202. St. Martin's Press, New York, New York.
- ELLIOTT, R., L. AGGOUN, AND J. MOORE (1995): *Hidden Markov Models: Estimation and Control*. Springer-Verlag, New York and Berlin.
- FISCHER, S. (1982): "Seigniorage and the Case for a National Money," *Journal of Political Economy*, 90(2), 295–313.
- (1987): "Comments on Thomas Sargent and Neil Wallace, 'Inflation and the Government Budget Constraint'," in *Economic Policy in Theory and Practice*, ed. by A. Razin, and E. Sadka, pp. 203–207. St. Martin's Press, New York, New York.
- FRIEDMAN, M. (1948): "A Monetary and Fiscal Framework for Economic Stability," *American Economic Review*, 38, 245–264.
- HAMILTON, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57(2), 357–384.

- HAMILTON, J. D., D. F. WAGGONER, AND T. ZHA (2007): "Normalization in Econometrics," *Econometric Reviews*, 26(2-4), 221–252.
- KUSHNER, H. J., AND G. G. YIN (1997): *Stochastic Approximation Algorithms and Applications*. Springer-Verlag.
- LJUNGQVIST, L., AND T. J. SARGENT (2004): *Recursive Macroeconomic Theory*. The MIT Press, Cambridge, Massachusetts, second edn.
- MARCET, A., AND J. P. NICOLINI (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93(5), 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989): "Least-Squares Learning and the Dynamics of Hyperinflation," in *International Symposia in Economic Theory and Econometrics*, ed. by W. Barnett, J. Geweke, and K. Shell, pp. 119–137. Cambridge University Press, Cambridge, England.
- MARIMON, R., AND S. SUNDER (1993): "Indeterminacy of Equilibria in a Hyperinflation World: Experimental Evidence," *Econometrica*, 61(5), 1073–1107.
- PIAZZESI, M., AND M. SCHNEIDER (2006): "Equilibrium Yield Curves," in *NBER Macroeconomics Annual 2006*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, pp. 389–442. MIT Press.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton, New Jersey.
- SARGENT, T. J., AND N. WALLACE (1987): "Inflation and the Government Budget Constraint," in *Economic Policy in Theory and Practice*, ed. by A. Razin, and E. Sadka, pp. 170–200. St. Martin's Press, New York, New York.
- SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006a): "The Conquest of South American Inflation," NBER Working Paper No. 12606.
- (2006b): "Shocks and Government Beliefs: The Rise and Fall of American Inflation," *American Economic Review*, 96(4), 1193–1224.
- SCLOVE, S. L. (1983): "Time-Series Segmentation: A Model and a Method," *Information Sciences*, 29(1), 7–25.
- SIMS, C. A. (1988): "Projecting Policy Effects with Statistical Models," *Revista de Analisis Economico*, 3, 3–20.
- SIMS, C. A., D. F. WAGGONER, AND T. ZHA (forthcoming): "Methods for Inference in Large Multiple-Equation Markov-Switching Models," *Journal of Econometrics*.
- TADIĆ, V. B., AND S. P. MEYN (2003): "Asymptotic Properties of Two Time-Scale Stochastic Approximation Algorithms with Constant Step Sizes," Proceedings of the 2003 American Control Conference.