

The Conquest of South American Inflation

Thomas Sargent, Noah Williams, and Tao Zha

Working Paper 2006-20

November 2006

The Conquest of South American Inflation

Thomas Sargent, Noah Williams, and Tao Zha

Working Paper 2006-20

November 2006

Abstract: We infer determinants of Latin American hyperinflations and stabilizations by using the method of maximum likelihood to estimate a hidden Markov model that potentially 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. Our maximum likelihood estimates allow us to interpret observed inflation rates in terms of variations in the deficits, sequences of shocks that trigger temporary episodes of expectations driven hyperinflations, and occasional superficial reforms that cut inflation without reforming deficits. Our estimates also allow us to infer the deficit adjustments that seem to have permanently stabilized inflation processes. Our results show how the available inflation, deficit, and other macroeconomic data had left informed economists like Rudiger Dornbusch and Stanley Fischer undecided about the ultimate sources of inflation dynamics.

JEL classification: C5, E5, H3

Key words: self-confirming equilibria, rational expectations, adaptation, inflation, seigniorage, deficits, escape dynamics

The authors thank Eduardo Ganapolsky, Karsten Jeske, Albert Marcet, David Romer, and especially Christopher Sims, Dan Waggoner, and Michael Woodford for helpful discussions. Sagiri Kitao, Tomasz Piskorski, and Demian Pouzo provided outstanding research assistance. Namgeun Jeong and Eric Wang provided indispensable assistance on clustering and parallel computing in the Linux operating system. The authors acknowledge the technical support of the Computing College of the Georgia Institute of Technology on grid computation techniques. They also thank Mike Chriszt, Jose Ricardo, Diego Vilan, and Elena Whisler for their help with both collecting the data and understanding institutional details. Sargent thanks the National Science Foundation for research support. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Thomas Sargent, Professor of Economics, Department of Economics, New York University, 269 Mercer Street, New York, NY 10003, 212-998-3548, thomas.sargent@nyu.edu; Noah Williams, Assistant Professor, Department of Economics, Princeton University, 001 Fisher Hall, Princeton, NJ 08544-1021, 609-258-4019, noahw@princeton.edu; or Tao Zha, Research Economist and Senior Policy Adviser, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, 404-498-8353, tzha@earthlink.net.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's Web site at www.frbatlanta.org. Click "Publications" and then "Working Papers." Use the WebScriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

THE CONQUEST OF SOUTH AMERICAN INFLATION

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

The Larida proposal starts from the premise that indexation, not monetized deficits, is the cause of inflation, and I share that view completely. Money creation, and more importantly, rising velocity because of monetary deregulation, are at best the air in the tires; indexation is decidedly the engine of inflation.

Rudiger Dornbusch, 1985

I. INTRODUCTION

I.1. A hidden Markov model. This paper estimates a hidden Markov model for inflation in five South American countries, Argentina, Bolivia, Brazil, Chile, and Peru.¹ Ours is a “back-to-basics” model. It 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 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 hard-to-detect deviations from rational expectations that help to explain features of the data that a strict rational expectations version of the model cannot. We trust our monthly series on inflation but lack trustworthy monthly data on deficits and money supplies. 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 deficit histories at the maximum likelihood parameter estimates and use it to calculate a density for the deficit history conditional on the inflation history. As one of several validation exercises, we compare this deficit density with the annual monetary deficit data that we do have. Unlike purely statistical models, we can use our model to infer the causes of hyperinflations

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

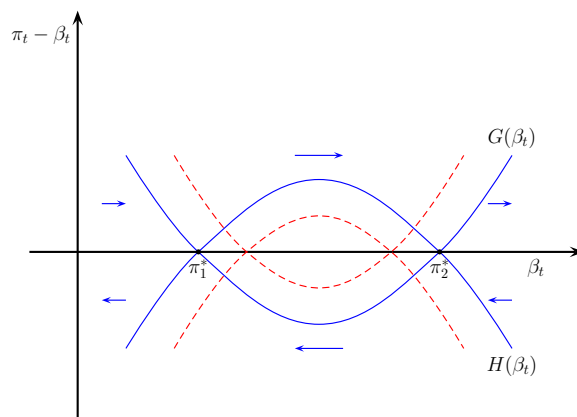


FIGURE 1. Mean adaptive dynamics and REE dynamics.

and stabilizations at different times and places. Our model offers explanations of particular inflations and stabilizations that differ substantially across episodes.

Our purpose in positing an adaptive expectations scheme is not to turn the clock back to the days before the hallmark cross-equation restrictions of the rational expectations revolution caused expectations to disappear as free variables in dynamic models.² On the contrary, we shall exploit rational expectations restrictions and self-confirming equilibria when we analyze salient features of our model's dynamics that allow it to fit the inflation data. But like Marcet and Nicolini (2003), our model retreats from rational expectations by adding an adaptation parameter that gives people's expectations dynamics that help our model explain the data, partly by eliminating some perverse out-of-steady state rational expectations dynamics and partly by allowing occasional expectations-driven big inflations. Though we use different procedures to highlight this, we shall argue along with Marcet and Nicolini (2003) that the departures of our model from rational expectations are not large.

I.2. Basic idea. We start with the insight of Marcet and Sargent (1989b) and Marcet and Nicolini (2003) that an adaptive expectations version of a hyperinflation model shares steady states with its rational expectations version, but has more plausible out-of-steady-state dynamics. Figures 1 and 2 summarize the dynamics of the model of Sargent and Wallace (1987) and Marcet and Sargent (1989b) and show the basic ingredients of our model.

²Sargent (1981) and most of the articles in Lucas and Sargent (1981) argue for making parameters measuring expectations disappear from econometric models. Lucas (1986) and Marcet and Sargent (1989a,b) used adaptive expectations schemes to justify rational expectations as an equilibrium concept supported by a law of large numbers and to select among multiple rational expectations equilibria.

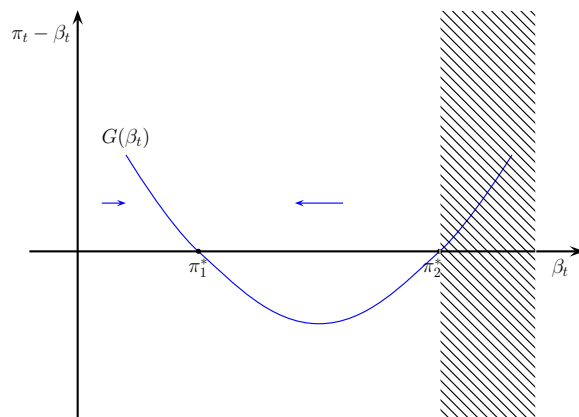


FIGURE 2. Adaptive dynamics and the “escape event”.

Here β_t denotes the public’s expected gross rate of inflation at date t ; $H(\beta)$ and $G(\beta)$ describe the actual inflation π_t determined by rational expectations (or perfect foresight) dynamics and some least squares learning (or adaptive expectations) dynamics, respectively. The dashed curves correspond to a higher deficit level than do the solid curves. Figure 1 indicates that while the rational expectations dynamics and the learning dynamics share fixed points (the zeros of H and G), they identify different fixed points as stable ones: the high expected inflation fixed point π_2^* is stable under the rational expectations dynamics, while the lower expected inflation fixed point π_1^* is stable under the learning dynamics.³ The two fixed points are on different sides of the peak of the Laffer curve, so increases in the deficit raise π_1^* , but lower π_2^* .⁴ Around the lower fixed point, increases in the government deficit increase inflation, while they lower inflation at the higher fixed point. An attractive feature of the learning dynamics is that they dispose of the implausible higher fixed point, with its perverse comparative statics, as a focal point for analysis.

Now think of stochastic versions of a model under the learning dynamics in which the G curve shifts in a stochastically stationary way as shocks impinge on the deficit or in which shocks impinge directly on the inflation rate without shifting the G curve. In such a stochastic version of the model with learning, the learning dynamics will tend to push expected inflation toward a stochastic counterpart of the lower fixed point so long as expectations remain within its domain of attraction, i.e., so long as they remain beneath π_2^* . But occasionally shocks might push β_t above π_2^* , the shaded

³Lucas (1986), Marcet and Sargent (1989a), and Evans and Honkapohja (2001) recommend selecting rational expectations equilibria that are stable under least squares learning.

⁴See Marcet and Sargent (1989b) for details.

region in figure 2 in which the G dynamics cause actual and expected inflation to increase without limit. When β exceeds π_2^* , we say that an *escape* (from the domain of attraction of π_1^*) has occurred. Marcet and Nicolini (2003) exploit the insight that occasional escapes from the domain of attraction of π_1^* could capture the recurrent bursts of inflation in Latin America that seemed not to coincide with any marked increases in government deficits. However, note that when deficits are higher on average, as in the dashed lines in figure 1, not only is inflation higher on average due to the higher π_1^* , but the higher fixed point π_2^* is lower and thus it will be easier for beliefs to escape. To make this explanation fit together, Marcet and Nicolini (2003) supplemented the basic model of Marcet and Sargent (1989b) with a story about mechanical reforms that end an escape episode by exogenously interrupting the G dynamics and resetting actual inflation (and therefore π_t) well within the domain of attraction of π_1^* under the G dynamics.

We adopt the idea of Marcet and Nicolini (2003) that there are recurrent escapes, but differ from them in our stochastic specification of the deficit and the reform event. Instead of calibrating the model as they do, we form a likelihood function, maximize it, then use the equilibrium probability distribution that we estimate to extract interpretations of the observed hyperinflation in terms of “normal” dynamics driven by deficits and “extraordinary” dynamics driven by escape dynamics. One of our objectives is to use the likelihood function for inflation to spot when escape events and reform events occurred. In addition to Marcet and Nicolini’s mechanical reforms that eventually arrest escaping inflation, the richer dynamics that we attribute to the deficit allow another type of reform not allowed by Marcet and Nicolini (2003): an exogenous shift in the deficit regime. This type of reform allows us to fit our model over longer periods than would be appropriate for the Marcet and Nicolini (2003) specification, in particular, the periods that include both recurrent hyperinflations as well as enduring stabilizations.

I.3. Related literature. Sargent and Wallace (1987) formed the likelihood function for a rational expectations version of a model closely related to the one that we shall study here. Their model has a continuum of rational expectations equilibria but is nevertheless overidentified. A single parameter in the likelihood function indexes the continuum of equilibria. Imrohoroglu (1993) estimated the Sargent and Wallace (1987) model using data from German hyperinflation of the early 1920s and made inferences about the prevailing rational expectations equilibrium. Because it assumed a constant mean deficit, the econometric setup in Sargent and Wallace (1987) and

Imrohroglu (1993) was not designed to explain data series spanning periods of hyperinflation and their subsequent stabilizations.⁵ To explain such data, we modify the model of Marcet and Nicolini (2003) while adhering to the maximum likelihood philosophy of Sargent and Wallace (1987). Another strand of the literature uses game theoretic ideas to develop models in which hyperinflations may be an important part of an equilibrium path. In Zarazaga (1993) and Mondino, Sturzenegger, and Tommasi (1996), conflict between different groups in determining policy results in an equilibrium path where inflation may switch between phases of high and low rates. While our model shares the features of switches in the rates of inflation, the mechanism underlying it is quite different.

Models of escapes have been used to model inflation rates and exchange rates at moderate levels of inflation by Sargent (1999), Cho, Williams, and Sargent (2002), Sargent and Williams (2003), Cho and Kasa (2003), Kasa (2004), Tetlow and von zur Muehlen (2004), and Ellison, Graham, and Vilmunen (2006).

I.4. Organization. The remainder of this paper is organized as follows. Section II presents our model and a brief description of the likelihood function for histories of inflation. We consign important details about the likelihood to appendix A. Section III gives a brief account of the concepts of self-confirming equilibria and conditional self-confirming equilibria that will guide our empirical interpretations. We relegate computational details to appendix B. In section IV, we describe how to compute the probabilities of two important events, namely, when beliefs escape and when monetary reforms take place. Section V describes our estimation procedures and results. Section VI then assembles these results into a set of interpretable economic findings. Section VII then explores the fit of our model relative to some alternatives. To help us measure how far we have drifted from rational expectations, section VIII computes stationary points of conditional self-confirming equilibria and compares them with stationary rational expectations inflation rates, defined and computed in appendix C, evaluated at our maximum likelihood parameter estimates. Section IX concludes. Appendix D describes the deficit data and discusses how we compute statistics from our model to compare to this data. Finally, appendix E collects some additional results.

II. THE MODEL

Given a vector of parameters, the model induces a probability distribution over sequences of inflation rates, money creation rates, deficits, and a hidden Markov state.

⁵A recent paper by Adam, Evans, and Honkapohja (in press) is another theoretical model that deals with a single episode of hyperinflation.

We use this joint distribution to deduce a marginal distribution for a sequence of inflation rates as a function of the model's parameters: this is our likelihood function. We maximize it to get parameter estimates. In this section, we describe the economic forces at play on the way to constructing the likelihood function to be presented in appendix A.

The model consists of a demand function for money, a government budget constraint, an exogenous process for fiscal deficits, and a formulation that by slightly retreating from rational expectations occasionally gives expectations a life of their own in shaping the evolution of inflation.⁶ The money demand equation, the government budget constraint, and the law of motion for deficits are:⁷

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

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

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

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

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

where $0 < \lambda < 1$, $0 < \theta < 1$, $\gamma > 0$, $\bar{d}(s_t) > 0$. Here $s_t = [m_t \ v_t]$ is a Markov state, as in Hamilton (1989) and Sclove (1983), that neither the agents inside the model nor we the econometricians observe; 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 public's expectation of the price level at time $t + 1$; and $\eta_{dt}(v_t)$ is an i.i.d. random shock. Each column of each transition probability matrix $Q_\ell = [q_{\ell,ij}]$ for $\ell = m, v$ sums to 1. The coefficient $\bar{d}(m_t)$ measures the average deficit, which via the government budget constraint equals the average amount of seigniorage financed by money creation in state m_t . Thus, in what follows we speak interchangeably of the fiscal or monetary deficit. The two Markov chains (Q_m, Q_v) induce a chain Q on the composite state $s_t = [m_t \ v_t]$ with transition matrix $Q = Q_m \otimes Q_v$.⁸ The total number of states is $h = m_h \times v_h$.

⁶Using adaptive rather than rational expectations also strengthens the role of the deficit as a fundamental that determines inflation. See Marimon and Sunder (1993) and the remarks in section II.2.

⁷For an interpretation of this equation as a saving decision in a general equilibrium model, see Marimon and Sunder (1993), Marcet and Nicolini 2003) and Ljungqvist and Sargent (2004). Equation (2) was used by Friedman (1948) and Fischer (1982), among many others.

⁸We 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 V.2 for a detailed discussion.

Rather than imposing rational expectations, we follow Marcet and Sargent (1989b) and Marcet and Nicolini (2003) and assume that:

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

The public updates the belief β_t by using a constant-gain algorithm:

$$\beta_t = \beta_{t-1} + \varepsilon(\pi_{t-1} - \beta_{t-1}), \quad (6)$$

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

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

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

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

Proof. Let $d_t(k)$ 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 does not affect the dynamics of either M_t or 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 1 explains why we have chosen to deviate from the procedure of Marcet and Nicolini (2003), who treated γ and $\bar{d}(m)$ as separate parameters, and who interpreted the calibrated value of d_t to measure fiscal deficits as a share of GDP. We think that procedure is misleading because these parameters cannot be identified separately, so that re-normalizing them in the manner of Proposition 1 gives the same equilibrium outcome.⁹ For identification purposes, therefore, when searching for the maximum likelihood estimates, we normalize $\gamma = 1$. We shall eventually re-normalize it to match the pertinent price level when we compute estimates of d_t to compare with some monetary deficit data. 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 $\log d_t$.

⁹For a general discussion of normalization in econometrics, see Hamilton, Waggoner, and Zha (2004).

II.1. Deterministic steady states. A deterministic version of model (1) - (5) can be obtained by fixing the state $m_t = m \in \{1, \dots, m_h\}$ and setting $\eta_{dt}(v_t)$ to zero for all t . We now report equilibria from a perfect foresight version of the model where agents observe and condition on the mean deficit state m . Such equilibria are useful reference points in the analysis of our stochastic adaptive model. As is well-known by now, there are two deterministic steady states associated with each m .

Proposition 2. If

$$\bar{d}(m) < 1 + \theta\lambda - 2\sqrt{\theta\lambda}, \quad (7)$$

then there exist two steady state equilibria for π_t :

$$\pi_1^*(m) = \frac{(1 + \theta\lambda - \bar{d}(m)) - \sqrt{(1 + \theta\lambda - \bar{d}(m))^2 - 4\theta\lambda}}{2\lambda}, \quad (8)$$

$$\pi_2^*(m) = \frac{(1 + \theta\lambda - \bar{d}(m)) + \sqrt{(1 + \theta\lambda - \bar{d}(m))^2 - 4\theta\lambda}}{2\lambda}. \quad (9)$$

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 (8) and (9). \square

We shall impose (7) 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

$$\pi_{SS}^{\max} \equiv \sqrt{\frac{\theta}{\lambda}}. \quad (10)$$

II.2. Limit points of near deterministic dynamics. Marcet and Sargent (1989b) and Marcet and Nicolini (2003) show that when the gain ε is sufficiently small, π_t converges to π_1^* when the initial belief satisfies $\beta_0 < \pi_2^*(k)$.¹⁰ Marcet and Sargent (1989b) describe how this outcome “reverses the dynamics” under rational expectations studied by Sargent and Wallace (1987), according to which π_t converges to the *high* steady state inflation rate $\pi_2^*(k)$. The $\pi_2^*(k)$ stationary point exhibits the perverse comparative statics property that stationary inflation *rises* when seigniorage

¹⁰They actually establish their convergence results for a learning scheme where the constant ε is replaced by $1/t$. Our convergence results in appendix B use the constant gain specification, which changes the nature of the convergence.

falls. We impute the constant gain (or adaptive expectations) learning scheme to agents for two reasons. First, we want to arrest the perverse comparative dynamics associated with rational expectations because we believe that normally higher deficits actually cause higher inflation and that imposing this feature on the model will help to explain the data. Second, as noted earlier by Marcet and Sargent (1989b) and Marcet and Nicolini (2003), in the presence of sufficiently large shocks, the adaptive expectations scheme creates the possibility that some big inflations are driven by dynamics of inflation expectations that are divorced from the fundamental force that normally causes inflation, namely, the deficit. We shall soon discuss such dynamics under the moniker “escape dynamics”. But first we state some restrictions on parameters and outcomes that are necessary for our equilibrium to be well defined.

II.3. Restrictions on parameters and outcomes. We return to the stochastic version of the model. By using (1)-(2) and (6), we obtain the following formula for equilibrium inflation:

$$\pi_t = \frac{\theta(1 - \lambda\beta_{t-1})}{1 - \lambda\beta_t - d_t(s_t)}, \quad (11)$$

provided that both the numerator and denominator are positive. As shown in the next section, the denominator must be bounded away from zero to ensure that the moments of inflation exist and that the inflation dynamics converge. Therefore, to guarantee existence of an equilibrium with positive prices and positive real balances, we impose the following restrictions:

$$1 - \lambda\beta_{t-1} > 0, \quad (12)$$

$$1 - \lambda\beta_t - d_t(s_t) > \delta\theta(1 - \lambda\beta_{t-1}). \quad (13)$$

Condition (13) uses a small value $\delta > 0$ to bound the denominator of (11) away from zero. It follows that inflation is bounded by $1/\delta$. Because the steady state rational expectations equilibrium (REE) inflation rate is bounded by $1/\lambda$ according to Proposition 2, it follows that $\lambda \geq \delta$.

II.4. Cosmetic reforms. It is possible for a sequence of seigniorage shocks η_{dt} to push β_t beyond a point when the inflation dynamics will tend to push β_t upward without limit, in turn leading to an explosion of inflation driven by adverse expectations dynamics. Unless we do something to arrest these dynamics, the model breaks down in the sense that conditions (12) and (13) will ultimately be violated. We do this by mechanically imposing a “reform” event: whenever these conditions are violated, we simply reset inflation to the low steady state $\pi_1^*(m_t)$ plus some noise:

$$\pi_t = \pi_t^*(m_t) \equiv \pi_1^*(m_t) + \eta_{\pi t}(m_t), \quad (14)$$

TABLE 1. Parameters and their meanings

Parameter(s)	Feature
λ	demand for money
ε	expectations
$\bar{d}(m)$	log deficit mean
$\xi(v)$	log deficit inverse std
ξ_π	reform inverse std
Q_m	m - transition matrix
Q_v	v -transition matrix

where $\eta_{\pi t}(m_t)$ is an i.i.d. random shock such that

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

When in addition $1 - \lambda\beta_t \leq \delta\theta(1 - \lambda\beta_{t-1})$, then the model breaks down even with $d_t = 0$. In this case we also reset expectations so that $\beta_t = \pi_t$.¹¹

We use this inelegant device of cutting inflation and expected inflation to represent the cosmetic reforms that Latin American governments occasionally used in the 1980s to arrest inflation without really altering the stochastic process for money-financed deficits.¹²

II.5. Likelihood functions. As discussed in section V.1 below, we fix the parameters (θ, δ) in estimation. Denote the remaining free parameters of the model as $\phi = [\lambda \ \bar{d}(m) \ \xi_d(v) \ \xi_\pi \ \varepsilon \ \text{vec}(Q_m) \ \text{vec}(Q_v)]$, where $m = 1, \dots, m_h$ and $v = 1, \dots, v_h$. The new parameters $\xi_d(v)$ and ξ_π are the inverses of the variances for structural shocks to deficits and inflation and will be discussed in detail in appendix A. For convenience, table 1 contains a reminder of the interpretations of these parameters. Let π^t be a history of inflation from 1 to t , and similarly for the other variables. Given a parameter vector, the model induces a joint density $p(\pi^T, m^T, v^T, d^T, M^T, \beta^T | \phi)$,

¹¹Marcet and Nicolini (2003) did not have this second clause in their reforms. We use it because it helps to fit the data. It has the interpretation that when agents see that the first clause of the cosmetic reform takes place *with certainty*, they are likely to adjust their belief quickly. If the cosmetic reform is less certain (i.e., its probability is less than one), the public will continue to update its belief according to the constant-gain algorithm. In principle, one could make reform of inflation expectations depend on states by assuming that the agents can observe states or infer states from a history of data. As discussed in section V.2, this specification would make estimating the model an insurmountable task.

¹²There were many cosmetic reforms in Latin America in the 1980s that sought to stabilize inflation “on the cheap” without tackling fiscal deficits. See Dornbusch (1985) for a contemporary discussion and Marcet and Nicolini (2003) for a discussion of superficial monetary reforms.

where we set $\beta_0 = \pi_0$ and the probability for the initial unobservable composite state $s_0 = [m_0 \ v_0]$ is set as described in appendix A. We follow the convention that the initial observable π_0 is always taken as given. The initial value M_0 is a function of β_0 and d_0 has no effect on the likelihood as long as π_0 is given. We take the marginal density $p(\pi^T|\phi)$, viewed as a function of ϕ , as our likelihood function and compute the estimator $\hat{\phi} = \operatorname{argmax}_{\phi} p(\pi^T|\phi)$. We make inferences about the deficit from the conditional density $p(d^T|\pi^T, \hat{\phi})$. Appendix A contains a detailed description of how we construct the likelihood function.

III. SELF-CONFIRMING EQUILIBRIA

Although we are mostly interested in the dynamics of the model that arise under learning, it is useful to have a benchmark equilibrium notion that characterizes these dynamics. Following Sargent (1999) and Cho, Williams, and Sargent (2002), we adopt the notion of a self-confirming equilibrium (SCE) as a natural limiting concept and reference point for learning models. In an SCE, agents' beliefs are consistent with their observations, though their beliefs may be incorrect about off-equilibrium events. In appendix B, we define self-confirming equilibria and conditional self-confirming equilibria in terms of orthogonality conditions that will govern β in large samples when $\varepsilon \rightarrow 0$. We also describe functions corresponding to the $G(\beta)$ function in figure 1 that govern the dynamic behavior of β_t as $\varepsilon \rightarrow 0$. Appendix C then defines rational expectations equilibria and links them to conditional and unconditional self-confirming equilibria. Here we briefly describe the main concepts.

An unconditional SCE is a β that satisfies $E\pi_t - \beta = 0$. For each mean deficit state m , a conditional SCE is a $\beta(m)$ that satisfies

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

for $m \in \{1, \dots, m_h\}$. A conditional SCE is thus just an unconditional SCE computed on the (false) assumption that the mean deficit state m_t will always be m . Like the deterministic steady state REEs, in general there are two conditional SCEs for each state m . We denote them as $\beta_1^*(m)$ and $\beta_2^*(m)$, where $\beta_1^*(m) \leq \beta_2^*(m)$. As in our discussion above, $\beta_1^*(m)$ is a stable attractor of the beliefs, while $\beta_2^*(m)$ marks the edge of the domain of attraction of this stable SCE.

As we describe in appendix B, the conditional SCEs serve as good approximations to REEs for very persistent average deficit states. They interest us because our estimate of $p(m_t = m | \pi^T, \hat{\phi})$ implies that the mean deficit states are very persistent. As we shall see, this makes a conditional SCE $\beta(m)$ a good approximation to the expected inflation rate in state m in a rational expectations equilibrium. It also

promises to make the conditional mean dynamics a good guide to the motion of β in our adaptive model in mean deficit state m . More formal details justifying and qualifying these assertions are given in appendix B.

For each country, we construct mean dynamics for the conditional SCEs in each deficit state m . We report them for each of our five countries in our key figures 3-7 below. In section VII, they will help us to interpret the inflation histories in our five countries in terms of convergence to a lower fixed point, and escapes above the higher fixed point associated with the mean deficit in state m . Later in section VIII, they will help us evaluate how much our model deviates from a rational expectations benchmark.

IV. PROBABILITIES OF ESCAPE AND COSMETIC REFORM

When a sequence of seigniorage shocks η_{dt} pushes β_t above the *unstable* SCE $\beta_2^*(m)$, we say that the inflation dynamics have *escaped* from the domain of attraction of the low SCE inflation rate. When an escape has proceeded so far that a breakdown threatens in the sense that (12)-(13) are violated, we impose the reform discussed in Section II.4.

Escapes and reforms contribute important features to the likelihood function. To give a formal definition of probabilities of escape and reform, we introduce the following notation:

$$\underline{\omega}_t(m_t) = 1 - \lambda\beta_t - \frac{\theta(1 - \lambda\beta_{t-1})}{\beta_2^*(m_t)} - \bar{d}(m_t), \quad (15)$$

$$\bar{\omega}_t(m_t) = 1 - \lambda\beta_t - \delta\theta(1 - \lambda\beta_{t-1}) - \bar{d}(m_t). \quad (16)$$

If $\beta_2^*(m_t)$ does not exist, we replace this term in (15) by π_{SS}^{\max} defined in (10). In the escape region, because actual inflation π_t is higher than β_t , both perceived inflation and inflation itself tend to escalate and thus hyperinflation is likely to occur. The probability of this escape event at time t given the $t - 1$ information set is

$$\mathbf{1}\{\beta_{t-1} < 1/\lambda\} \sum_{k=1}^h [\Pr(s_t = k \mid \Pi_{t-1}, \phi) (F_{\eta_d(k)}(\bar{\omega}_t(k)) - F_{\eta_d(k)}(\underline{\omega}_t(k)))] , \quad (17)$$

where $\mathbf{1}(A)$ is an indicator function that returns 1 if the event A occurs and 0 otherwise and $F_{\eta_d(k)}(x)$ is the cumulative density function (cdf) of $\eta_{dt}(k)$ evaluated at the value x . The probability of reform at time t given the $t - 1$ information set is:

$$\mathbf{1}\{\beta_{t-1} \geq 1/\lambda\} + \mathbf{1}\{\beta_{t-1} < 1/\lambda\} \sum_{k=1}^h [\Pr(s_t = k \mid \Pi_{t-1}, \phi) (1 - F_{\eta_d(k)}(\bar{\omega}_t(k)))] . \quad (18)$$

V. ESTIMATION

V.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 cover the 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 above, we are able to estimate the structural parameters through the inflation likelihood derived in appendix A. On the other hand, we may ask too much of the model to pin down all the parameters from inflation data alone. Therefore we fix the values of the following three parameters as $\beta_0 = \pi_0$, $\theta = 0.99$, and $\delta = 0.01$. The value of θ is consistent with economic growth and some cash taxes.¹³ The value of δ implies that monthly inflation rates are bounded by 10,000%.¹⁴ Although we do not use it in estimation, we do have annual data on seignorage which is described in appendix D. As discussed below, we use this data to validate our model by comparing the distribution of deficits predicted by the model with the actual data on seignorage.

Because of the long sample, the likelihood of inflation is well shaped around its peak. There are local peaks but often the likelihood values at these locations are essentially zero relative to the maximum likelihood (ML) value. Nonetheless, if one chooses a poor starting point to search for the ML estimate, the numerical algorithm is likely to lead to an estimate at a local peak.¹⁵ 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 (2006) and various constrained optimization routines contained in the commercial IMSL package. The block-wise BFGS algorithm, following

¹³One could impose a prior distribution of θ with values ranging from 0.96 to 1.0. This is one of few parameters we have a strong prior on. It is difficult or impossible, however, to have a common prior distribution on the other structural parameters because the likelihood shape differs considerably across countries. If we center a tight prior around the location as odds with the likelihood peak, the model would be unduly penalized. It would be more informative to study the likelihood itself and let the data determine what the model estimates are for each country. One could interpret our likelihood approach as having a diffuse prior on the other structural parameters.

¹⁴Marcet and Nicolini (2003) set the bound at 5,000%.

¹⁵Such a problem is prevalent in Bayesian estimation.

the idea of Gibbs sampling and EM algorithm, breaks the set of model parameters into a few subsets and uses Sims’s `csmnwel` program to maximize the likelihood of one set of the model’s parameters conditional on the other sets.¹⁶ This 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 are obtained at the highest likelihood value.¹⁷ The other converged points typically have much lower likelihood values by at least a magnitude of hundreds of log values.

V.2. Robustness analysis. In addition to the specifications described in section II, we have studied a number of alternative specifications described in section VII below. None of these alternatives has improved the fit of our model. Further, one could in principle let ε or β_t depend on regimes s_t . While this alternative creates no difficulty in analyzing the theoretical model, it is infeasible to compute the likelihood function because the unobservable variable β_t depends on the entire history of regimes s_t (even though s_t itself follows a Markov chain). For the model with bounded rationality, moreover, it may make sense to assume that the agents do not know the regimes while it is natural to assume that they know the regimes in the rational-expectations model.

VI. FINDINGS

In this section, we present and interpret our main empirical results. Before going through the analysis country by country, we look at how some key parameters vary across countries.

VI.1. Parameter patterns. Tables 6-10 in appendix E report the maximum likelihood estimates of our model for Argentina, Bolivia, Brazil, Chile, and Peru, respectively, along with the estimated standard errors.¹⁸ As one can see, all the parameters

¹⁶The `csmnwel` program can be found on <http://sims.princeton.edu/yftp/optimize/>.

¹⁷For 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.

¹⁸Following Sims (2001) and Hamilton, Waggoner, and Zha (2004), the standard errors are derived from the covariance matrix that is computed as the inverse of the Hessian of $\log p(\Pi_T|\theta)$

TABLE 2. Money demand and adaptation parameters

Country	λ	ε
Argentina	0.73	0.023
Bolivia	0.307	0.232
Brazil	0.613	0.189
Chile	0.875	0.025
Peru	0.74	0.069

are tightly estimated except ξ_π in the case of Brazil. The standard error for the element in the second row and second column of Q_m for Chile implies a high likelihood that the low deficit regime would last forever. As for Peru, one can see that the high deficit regime is more persistent than the low and medium deficit regimes.

For our five countries, table 2 reports our maximum likelihood estimates of 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 (6). There are interesting cross country differences in these parameters. 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, 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 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 probabilities of escape 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 the hyperinflation period. We return to these observations in our country-by-country analysis below.

VI.2. The key figures. We use figures 3-7 below to breathe life into our maximum likelihood estimates for each country. Each of these figures consists of five panels

evaluated at the MLEs. The estimated values of \bar{d} are re-normalized to be consistent with the data, as we describe in appendix D.

aligned to reveal features of our estimated model. The top left panel contains two or three curves that depict the conditional mean dynamics for β and whose zeros depict conditional self-confirming equilibria evaluated at the maximum likelihood parameters $\hat{\phi}$ for the country under study. The SCEs are conditional on the different estimated average deficits measured by $\bar{d}(m)$. There are three curves when the Markov state for the mean deficit can take three values and two curves when we allow it to take only two values. These curves are empirical renditions of the $G(\beta)$ functions in figure 1 at the different levels of the mean deficits we have estimated. We have projected the zeros from these figures as horizontal dotted lines into the top right panel, which plots our estimates of the public’s inflation beliefs β_t over time. These dotted lines tell us the stable (the lower values) and unstable (the higher values) of beliefs for each deficit level and help us to identify the range of β ’s that qualify as escapes and reforms.

The panel that is the second from the top compares bars that are seigniorage rates constructed from annual data along with the 0.16, 0.5, 0.84 probability quantiles for d_t predicted by our model (which recall uses only inflation data for estimation). See Appendix D for the details of how these numbers are computed in the data and the model. The dashed lines in the graph contains two-thirds of the probability distribution of simulated annual deficits from our model; the solid line labelled “Model” represents the median of simulated annual deficits.¹⁹ For all countries, it is striking that the deficits constructed from actual data seem to follow relatively closely the patterns that the model predicts and about which the quantiles constructed from $p(d^T|\pi^T, \hat{\phi})$ are informative.

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 Deficit” when there are two mean states (or “L & M Deficit” when there are three mean states), is the probability that the mean deficit regime is in the low deficit state (or either the low or the medium state in the three-mean deficit case) as a function of time. The dashed line is the probability that an escape will occur next period, computed as described in section IV above. Figures 10-14 in appendix E plot the full suite of probabilities of the mean deficit and deficit shock variance regimes for each country.

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. Later, we shall compare the fit of this model with a good-fitting

¹⁹We report quantiles because our model makes the distribution of d_t a fat-tailed mixture of log normal distributions. When the deficit shock variance is large with a high average amount of seigniorage, the deficit distribution is quite skewed with a very fat tail.

autoregression constructed without imposing our economic model. Here we note simply that the model tends to track the data relatively well, particularly in the hyperinflation episodes.

We now use these figures to tell what our estimates say about the histories of inflation in these five countries.

VI.3. Argentina. The top left panel of figure 3 shows curves for two conditional self-confirming equilibrium dynamics. The one associated with the high m state has its higher fixed point at a value for $\log \beta$ of about 0.2, while the low m state has a low stable point near zero. Comparing β_t in the top right panel with the probabilities in the third panel down on the right shows that throughout the 1960s up until 1975 the economy was repeatedly in the low deficit regime, the probability of escape was very low, and expected inflation hovered around the low conditional SCE. The low probability for the low deficit state in the third panel on the right shows that after 1975 up until 1991, Argentina lived with a chronically high mean deficit. The second panel from the top shows that our model predicts higher and more volatile deficits throughout this period, and this is largely confirmed in the annual deficit 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 the deficits. Again, looking at the third panel, the probability of escape 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 high deficit conditional dynamics indicate that if Argentina had been lucky enough to avoid sequences of adverse deficit shocks that drove β 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 thus lucky until the late 1980s, when the escape 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 the deficit, which remained in a lower and less volatile regime for most of the rest of the sample. Again, this is confirmed by the second panel, which shows smaller deficits throughout the later 1990s, apart from a period of volatility associated with the crisis in 2002. This change in the mean deficit state shifted the conditional dynamics curve (in the top left panel) in a way that pushed expected inflation rapidly downward, with beliefs (in

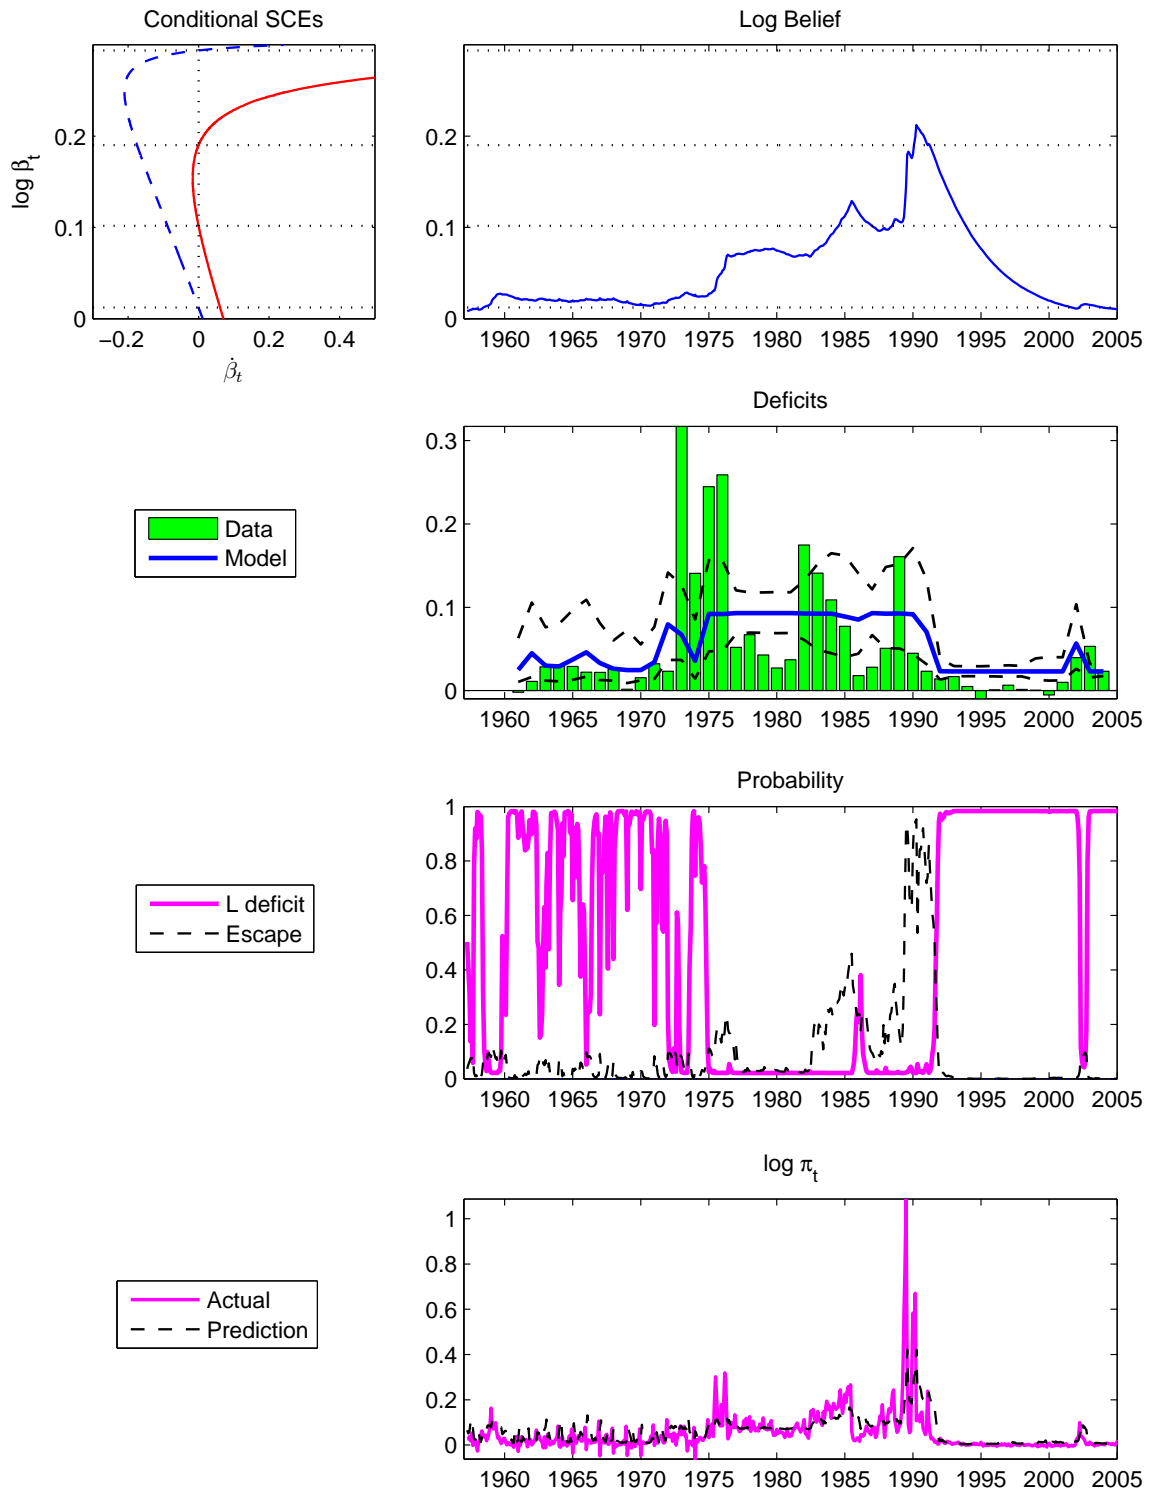


FIGURE 3. Argentina.

the top right panel) eventually converging down toward the lower conditional SCE once again. Although the change in the deficit state and the decline in the actual inflation rate occurred relatively quickly, expected inflation fell throughout the last years of the sample because the estimated gain ε is small.

VI.4. Bolivia. Our results for Bolivia tell a substantially different story. Our estimates suggest that the escape dynamics emphasized by Marcet and Nicolini (2003) played no role in Bolivia. The most striking thing about the conditional dynamics in the top left panel of figure 4 is how spread out the conditional SCEs are in each regime. As noted above, this is governed in our model 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 probabilities are very small throughout the entire sample, so small that it is hard to detect by eye from the third panel on the right.

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 mean deficit state took place around 1982, a view that is confirmed by the deficit data shown in the second panel. Throughout most of the sample our model predicts a low and relatively stable level of monetized deficits, 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, around 1987 the economy switches back to the low and more stable deficit regimes (third panel), actual deficits are 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.

VI.5. Brazil. Brazil, as shown in figure 5, presents an interesting case study with two main episodes of hyperinflation that appear to have been ended by different

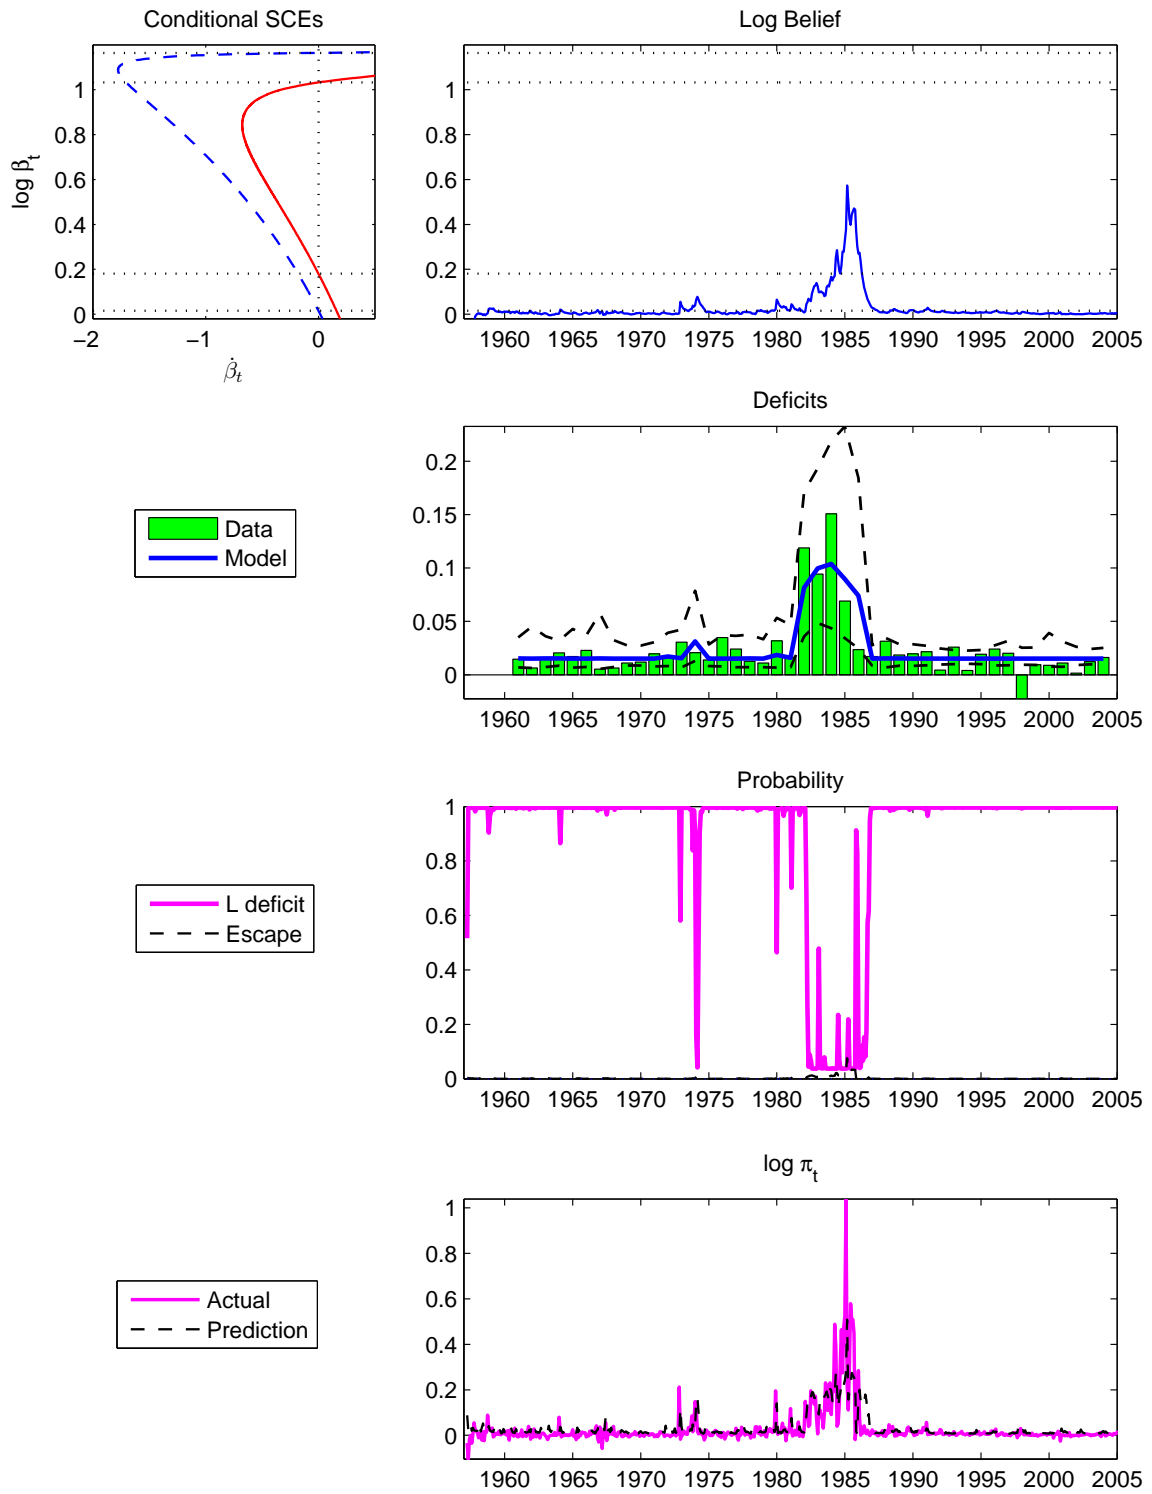


FIGURE 4. Bolivia.

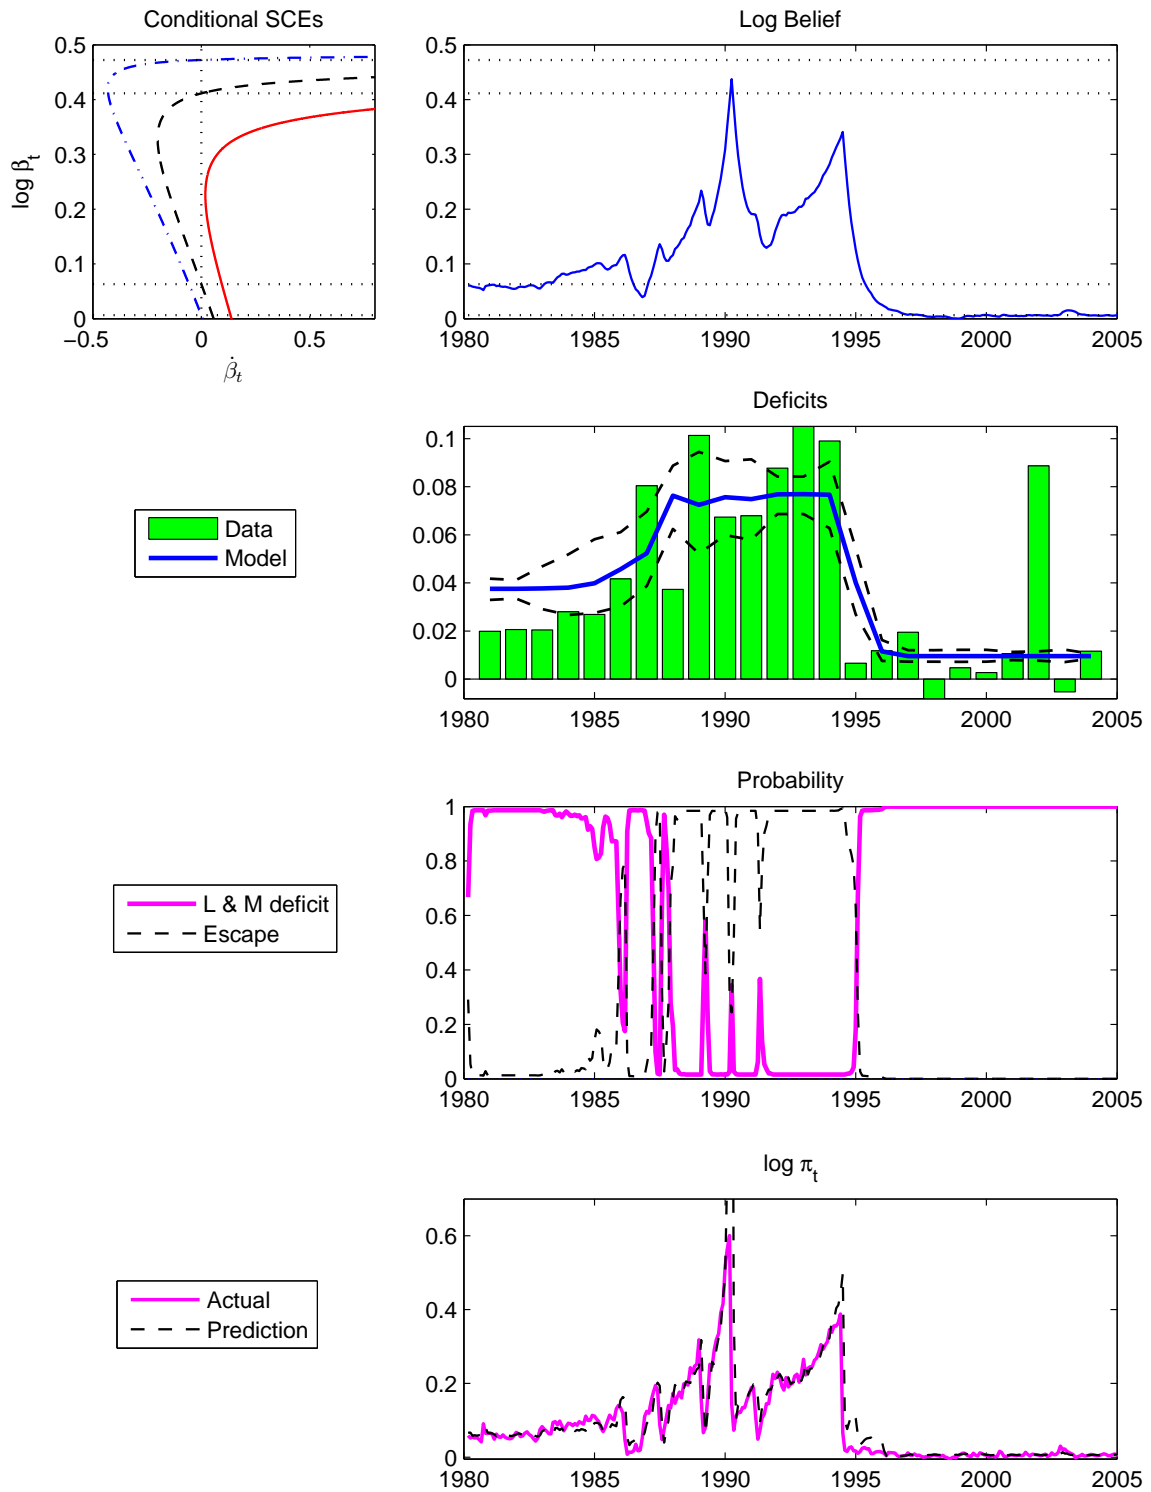


FIGURE 5. Brazil.

means. First note that in the top panel on the left the low average deficit state has a conditional SCE near zero, the medium deficit's SCE is near 0.06, but the high mean deficit conditional dynamics curve has no fixed points. We interpret this as asserting that when the economy is in this deficit regime, 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 deficit state, as evidenced by the regime probabilities in the third panel down and the predictions and actual levels of deficits in the second panel. Throughout this period, expected inflation was near the medium deficit 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 average deficit state, remaining there until 1994. Again, this is clear from the regime probabilities in the third panel and the deficit predictions and data in the second panel.²⁰ Once the economy entered this regime, the unstable learning dynamics kicked in. The escape probabilities in the third panel rose repeatedly after 1985, remaining mostly near one after 1987 up until 1994, and there were volatile but high deficits 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 superficial mechanical reforms that cut inflation but leave the mean deficit unaltered. These reforms succeeded in lowering expected inflation temporarily (top panel), but it rose rapidly again until 1994, with actual inflation rising again to another peak (bottom panel). However, unlike the earlier cosmetic reforms, our model says that the 1994 stabilization is different, and was accompanied by a persistent reduction in the mean and volatility of the monetized deficit. This is evident in the lower predicted deficits in the second panel, which largely accords with the lower and more stable actual deficits (apart from 2002). After 1994, beliefs fell rapidly down to the low deficit SCE (top panel), and actual inflation remained stable at a low level (bottom

²⁰Note that regime probabilities are computed on a monthly basis while the annual deficit is an average of monthly deficits as discussed in appendix D.

panel). Moreover, as shown in table 8 in appendix E, our estimates of the transition probabilities Q_m suggest that the high and medium deficit regimes are transitory, and thus our model predicts the sustained stable inflation that accompanies the low deficit regime. Thus, Brazil provides an interesting example of the ultimate futility of the cosmetic reforms, followed by a successful sustained fiscal reform.

VI.6. Chile. In figure 6 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 deficit state having conditional SCEs near 0.03 and 0.12, and the high deficit 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 deficit levels themselves do not vary sizeably across regimes, as the median model prediction in the second panel is essentially flat over the entire sample. The probabilities of the low deficit in the third panel are relatively volatile before 1994 (as reflected by relatively volatile inflation rates), but these do not translate into volatile predictions. 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 deficits, 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 deficits, 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 escape (third panel), and leading to the hyperinflation observed (bottom panel).

Because the buildup in inflation was largely driven by shocks, the stabilization in the late 1970s is interpreted as a reduction of variance of shocks to deficits. 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, deficit 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.

VI.7. Peru. Figure 7 considers the case of Peru. The third panel down makes clear that throughout the 1960s and most of the 1970s, the economy was in the low and middle average deficit states. The average deficits do not vary substantially across

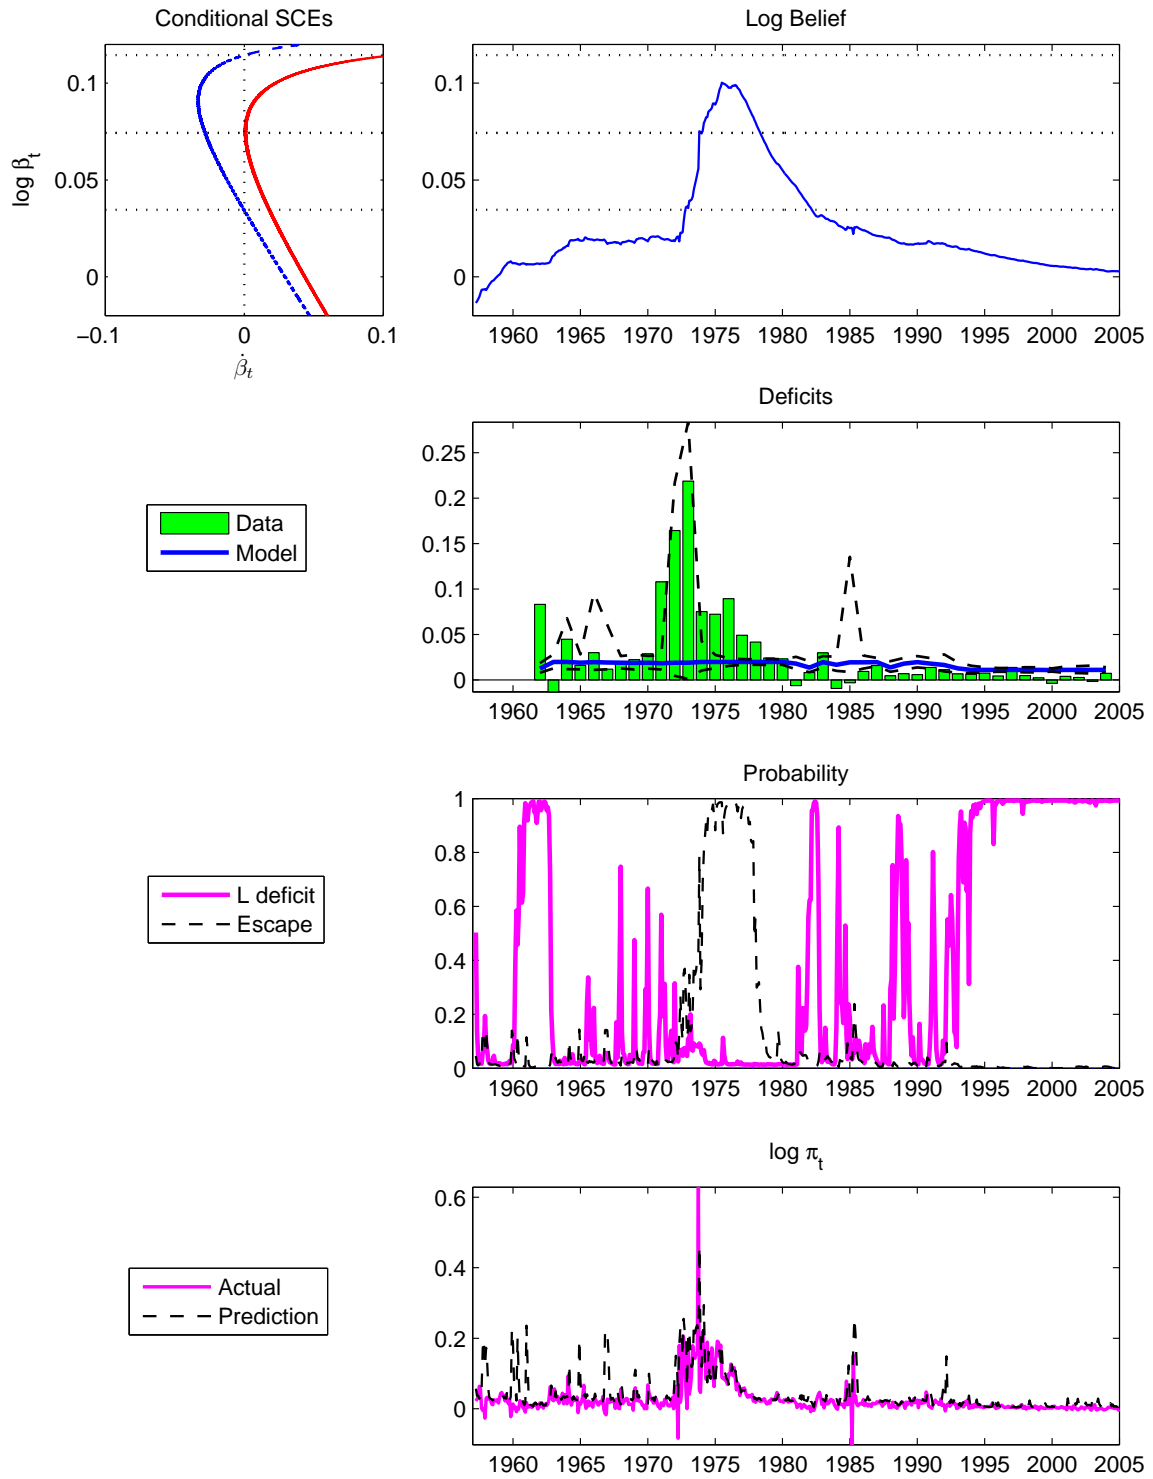


FIGURE 6. Chile.

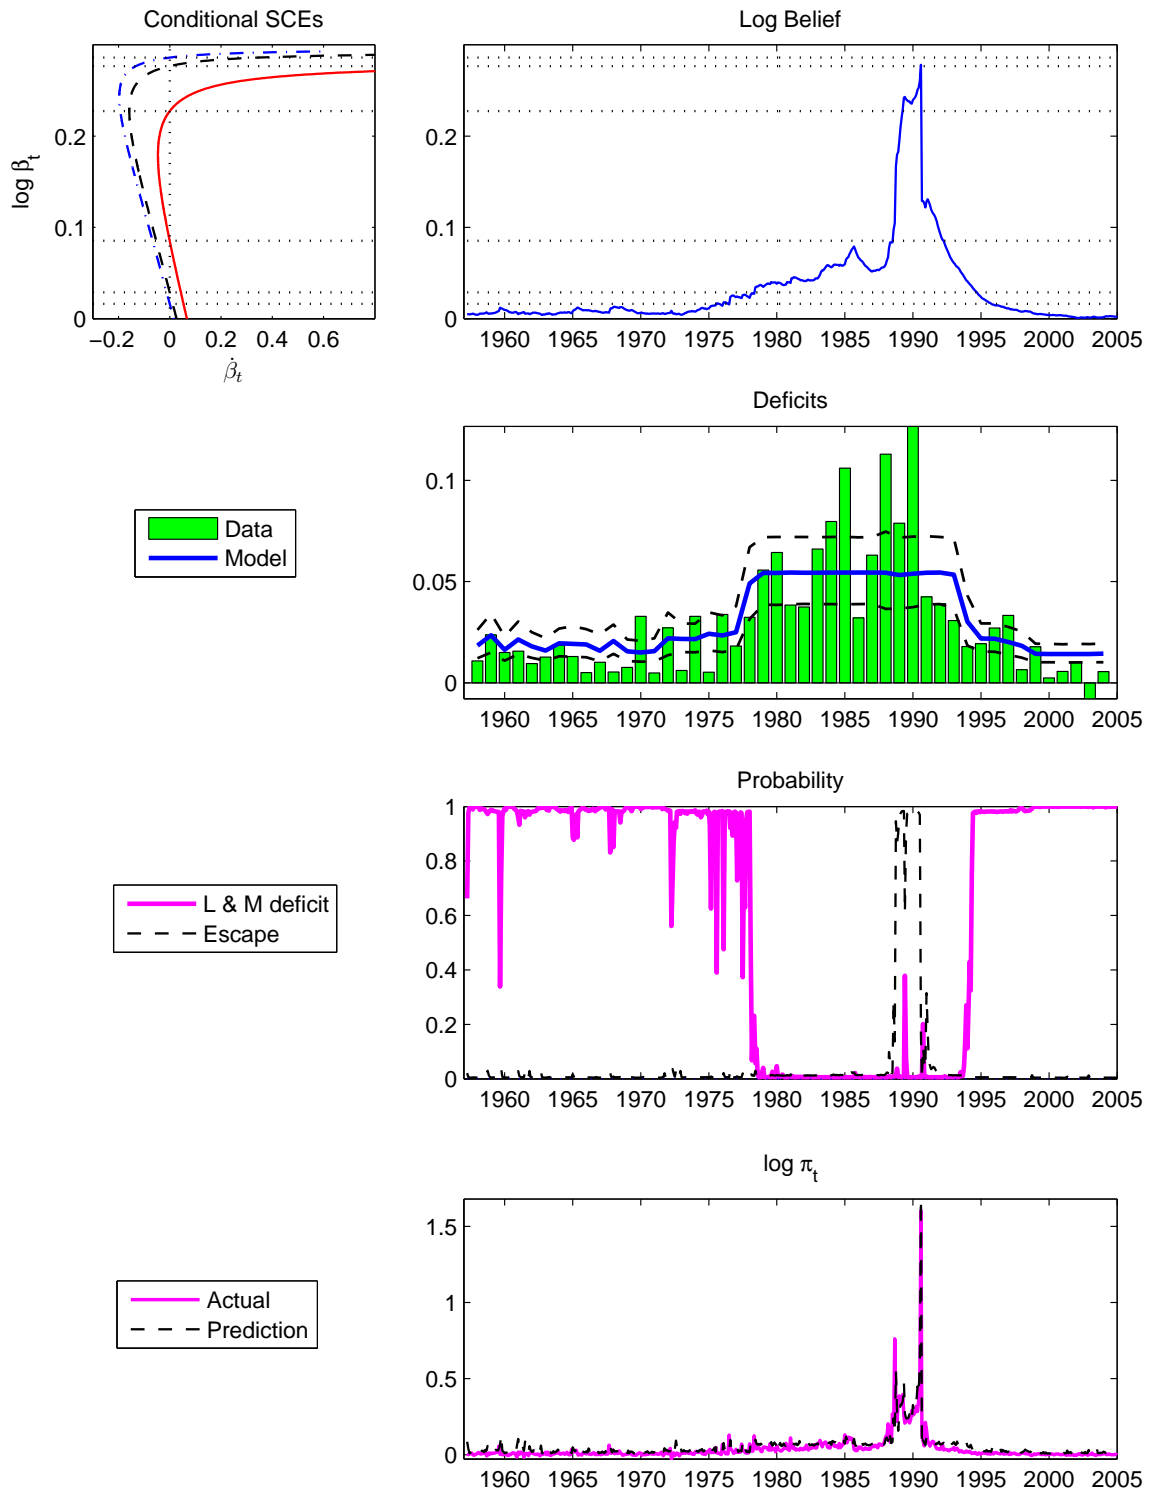


FIGURE 7. Peru.

these regimes, as the second panel shows that the median prediction is relatively flat, which roughly matches the annual deficit 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 our model suggests that the economy entered into a high deficit regime (third panel) that persisted all the way 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 with it (bottom panel).

But as we have seen in the other countries, the high deficit state is precarious and a sequence of deficit shocks was enough to upset the balance. In 1989 expected inflation increased rapidly, traversing into the region that prompted a large and rapid escape, as is evident in the belief dynamics in the top panel and the sharp increase in the escape probability in the third panel. Inflation itself increased even more dramatically (bottom panel), triggering a definite cosmetic reform. In other countries, cosmetic reforms occur with probability less than one and thus only inflation itself is reset. Peru is the only country that experiences a cosmetic reform with certainty, which evokes expectations to be reset as well. We interpret this double-barreled intervention as a reform that, although it does not alter the deficit dynamics, is credible in the sense that the public believed it to be effective in cutting inflation. Consequently, expected and actual inflation jump down dramatically in 1992. Moreover, unlike the failure of the initial cosmetic reforms in Brazil where there is no SCE in the high-deficit regime, these reforms seemed to have been successful in Peru, where there is a large gap between the low and high SCEs in the high-deficit state. Consistent with the high-deficit 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 deficit state for a considerable time. Evidence of a fiscal reform is absent until 1995, when the probability assigned to the low or medium monetized deficit state increased nearly to one (third panel). Thus, Peru seems to be the case where the unorthodox cosmetic reforms discussed by Marcet and Nicolini (2003) were the most successful in vanquishing hyperinflation.

VI.8. Comparison with Marcet and Nicolini. We now illustrate how our model differs from the model of Marcet and Nicolini (2003), which is geared toward explaining the experience of Argentina. We have formed the distribution conditioned on the history of inflation that is implied by their quarterly model at their calibrated parameter values. Figure 8 reports the one-step forecasts and 90% probability distributional

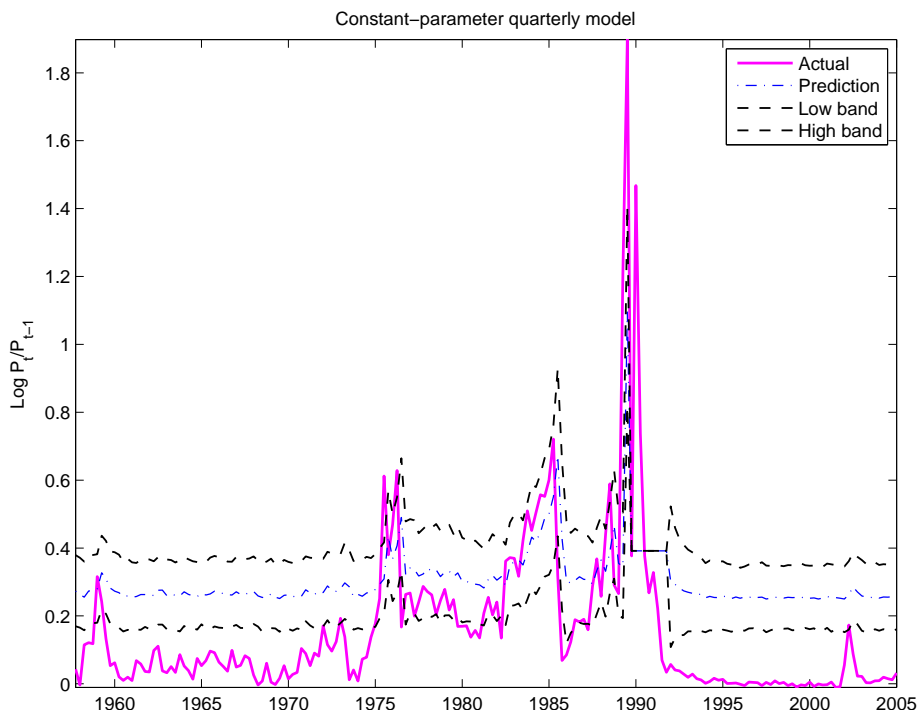


FIGURE 8. Marcet and Nicolini (2003) model's one-step prediction with 90% probability bands for Argentina.

bands around them, together with actual inflation outcomes. This figure should be compared with the bottom panel of figure 3. For Marcet and Nicolini's calibrated parameter values, we have found that there is no SCE, so that the ordinate of their $G(\beta)$ curve (see figure 1) always exceeds zero. This means that inflation expectations are perpetually along an escape path that must terminate with a mechanical monetary reform that will reset inflation itself. This seems to be the reason that in figure 8 the Marcet-Nicolini constant-parameter economic model over-predicts actual inflation in the relatively low inflation periods preceding 1975 and following 1991. By considering a more general deficit process, our model is thus able to capture both the hyperinflations and the stabilizations.

VII. THE MODEL'S FIT

Since our theoretical model is highly restricted, one would not expect its fit to come even close to be as good as a standard autoregressive (AR) model, let alone a time-varying AR model. In previous work with models like ours, such as Marcet and Nicolini (2003), only certain moments or correlations of the model were typically reported and

TABLE 3. Log likelihood adjusted by the Schwarz criterion

	Constant	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

compared to the data. By contrast, in this paper we take the fit of our model seriously and report it against the flexible, unrestricted statistical models. We compare not only various versions within our model but also our model with different types of AR models.

For each country we have tried more than two dozen versions of our theoretical model and of the unrestricted atheoretical models. Some of the different variations include the models with constant parameters, with 2-5 states for $\bar{d}(s_t)$ and $\eta_{dt}(s_t)$ jointly, for $\bar{d}(s_t)$ only, for $\eta_{dt}(s_t)$ only, and for $\bar{d}(s_{1t})$ and $\eta_{dt}(s_{2t})$ where s_{1t} and s_{2t} are independent state variables. 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,²¹ 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 constant-parameter case fit worse. With the 3-state case, we follow Sims, Waggoner, and Zha (2006) 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.

In addition, we have tried a number of different specifications of the model. For example, we have let π_t^* and d_t be serially correlated with their parameters freely estimated, and allowed ξ_π to be time varying. We have also allowed $\bar{d}(m_t)$ to be negative, used a number of different distributions for η_π and η_d , including the truncated normal distribution used by Marcet and Nicolini (2003), and introduced more lagged inflation variables in the learning rule (6). Again, within each model, the fit in most versions is substantially worse than our best-fitting model and thus we do not report their results.

²¹See Sims (2001) for detailed discussions of how to use the SC for model comparison.

Table 3 reports the log likelihood (adjusted by the Schwarz criterion) of the best-fitting theoretical model for each country, the constant-parameter theoretical model as in Marcet and Nicolini (2003), and the best-fitting unrestricted regime-switching AR model.²² The constant-parameter model for every country fits poorly; this result is consistent with the analysis in section VI.8. For all the five countries, the best-fitting atheoretical model is the 2×2 AR(2), which allows the two states in coefficients to be independent of the two states in shock variances. Our best-fitting theoretical model is used 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, while it fits better for in the cases of 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 hyperinflations, our model does as well or better than the atheoretical models. These points are illustrated by figures 15-19 in appendix E, where we compare the log conditional likelihood $p(\pi_t | \pi^{t-1}, \hat{\phi})$ of our theoretical model with that of the best-fitting statistical model. Clearly, the fit is much better for our theoretical model than the statistical model during the period of hyperinflation. Take Argentina as an example. The log likelihood for the non-hyperinflation periods 1957:04-1974:12 and 1993:01-2004:04 is 1014.4 for the statistical model and 948.1 for the theoretical model. This difference is 66.2, which captures most of the difference between the fits of the two models. Similarly, the log likelihood for the hyperinflation period 1979:01-1987:12 in Bolivia is 133.0 for our theoretical model and 117.6 for the statistical model, so the fit is much better for our model during this period. Thus, while our model is not an unqualified empirical success, it does a reasonably good job of capturing much of the inflation dynamics in these countries. Moreover the validation of our model provided by the monetary deficit data, as well as the insight that it provides into the causes and consequences of inflation make us regard it as a clear economic success.

VIII. COMPARING SCEs AND REEs

We now examine what our estimates imply for the self-confirming and rational expectations equilibria discussed above and defined formally in appendices B and

²²Regime-switching AR models considered here are simply a special case of regime-switching VAR models developed by Sims and Zha (2006) and by Sims, Waggoner, and Zha (2006).

TABLE 4. Conditional SCEs and REEs. Log value of beliefs in each regime.

Country	SCE/REE	Low m	Medium m	High m
Argentina	SCE	0.0123	-	0.1018
Argentina	REE	0.0188	-	0.0902
Bolivia	SCE	0.0153	-	0.1804
Bolivia	REE	0.0165	-	0.1711
Brazil	SCE	0.0064	0.0632	0.2256
Brazil	REE	0.0063	0.0726	0.2299
Chile	SCE	0.0346	-	0.0743
Chile	REE	0.0378	-	0.0612
Peru	SCE	0.0163	0.0288	0.0853
Peru	REE	0.0186	0.0332	0.0833

C. As we have emphasized, key parts of our story for the dynamics of hyperinflations require retreating from rational expectations. We argue here that the retreat is relatively minor.

In particular, Table 4 lists the conditional SCEs for each country, along with the beliefs from a rational expectations equilibrium (REE) consistent with low inflation. (As we discuss in appendix C, there are many other REEs as well.) The table makes quite clear that the conditional SCEs that govern the dynamics of our adaptive agents' beliefs are very 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. The values in the low m (average deficit) state in particular are very close for all countries.²³ Moreover the REEs are computed under the assumption that agents observe and condition on the mean deficit state m_t . Since our adaptive agents' beliefs converge in distribution to these conditional SCEs, this suggests that on average agents in our model do not make systematic forecast errors. Even though they do not observe and do not condition on the Markov state governing the average deficit, they are able to adapt their beliefs over time in such a way that this omission is not very costly to them. Thus, our departure from full rationality is rather small. But by making this departure, we change in important ways the dynamics of the model and allow the data to inform us as to the underlying causes and consequences of the South American inflations.

²³Brazil was a special case in that there is no conditional SCE in the high average deficit regime. Similarly, we were not able to find a low inflation REE. The table reports the "near equilibria" that have the lowest forecast errors, and correspond to the minimal value of the mean dynamics in the high seignorage state.

TABLE 5. Causes for the rise and fall of *hyperinflation* across countries

	Escape	No Escape
Cosmetic Reform	Brazil (87-91) Peru (87-92)	
Fiscal Deficit Reform	Argentina (87-91) Brazil (92-95)	Bolivia (82-86)
No Reform (Driven by Deficit Shocks)	Chile (71-78)	Argentina (76-86)

IX. CONCLUDING REMARKS

Building on Sargent and Wallace (1987) and Marcet and Nicolini (2003), we have formulated a nonlinear stochastic model of inflation, expectations, and money-supply financed deficits and then fit the implied density over histories of inflation to data for Argentina, Bolivia, Brazil, Chile, and Peru. Figures 3-7 summarize the stories that our maximum likelihood estimates tell about the inflation histories of these countries.

IX.1. Types of inflations and stabilizations. Table 5 briefly summarizes some of the key empirical patterns revealed by Figures 3-7. The first column lists three possible ways that our model tells us hyperinflation can be stopped: a superficial monetary reform that mechanically resets inflation without altering the deficit regime, a fiscal deficit reform activated by a change in the mean monetary deficit, and no reform in the mean monetary deficit but a change in the conditional deficit shock variance. The top row lists the two possible causes of hyperinflation: a high probability of self-perpetuating escalation of inflation governed by escape dynamics, and a large deficit shock variance coupled with a small probability of escape. We chose to put countries into appropriate boxes in the table according to whether our model assigns high probabilities (i.e., over 60%) of escape or of a cosmetic reform.

In March of 1990, for example, Brazilian inflation reached its peak with a monthly gross rate of 1.82. In the next two months, the inflation rate dropped to 1.15 then 1.07. The probability of cosmetic monetary reform is 67.5% for March, 75.6% for April, and 47.8% for May. By contrast, the high and volatile inflation episode in Brazil finally ended in 1994 with a sustained fiscal reform. Peru is another informative 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. The probability of cosmetic reform is only 10.8% in August but

jumps to 100% in September. Expected inflation is so high that the reform resets inflation expectations as well as inflation, which brings down the belief instantly.²⁴ Thus, the cosmetic reform is crucial for interpreting the fall of hyperinflation in Peru.

For Argentina (from 1987 to 1991), Bolivia, and Brazil, fiscal reforms play a dominant role in conquering hyperinflation. A reduction of the variance of shocks to deficits can be also important, as seen in Chile and Argentina (from 1976 to 1986).

IX.2. Ergodic distributions for deficits. Figure 9 plots the ergodic probabilities of the estimated average deficit level and the estimated standard deviation of deficit shocks. As we have seen, in the high average deficit states, inflation tends to be high on average and there is typically a sizeable probability of escaping to hyperinflation. This chance is even greater when the variance of the deficit shocks is larger. These ergodic probabilities are therefore helpful in predicting the long run inflation experiences of the different countries in our sample. Countries whose ergodic distribution has significant mass in the high average deficit and high shock variance states will thus tend to experience repeated episodes of high inflation and hyperinflation. In contrast, countries where most of the mass is in low mean deficits and small shocks will tend to have low and stable inflation. Clearly, Brazil and Chile have deficit processes that are conducive to persistent low inflation, while Argentina and Peru do not. Bolivia is an intermediate case, suggestive of mostly low inflation perhaps interrupted by occasional periods of high inflations.

IX.3. The epigraphs. We conclude by returning to the thoughtful epigraphs of Dornbusch (1985) and Fischer (1987) with which we began. The articles in which those quotes appear testify to how the available inflation, deficit, and other macroeconomic data had left informed observers like Dornbusch and Fischer undecided about the ultimate sources of inflation dynamics. While not including the indexation arrangements that Dornbusch suspected contributed so much to inflation dynamics, our model makes important concessions in his direction by treating the mean deficit as a hidden Markov state whose effects on inflation are confounded by noise and obscured by shocks to deficits and perturbations to expectations and therefore have to be gently coaxed from the data. And we have taken Fischer's doubts about the pure rational expectations model of Sargent and Wallace (1987) seriously by backing away from rational expectations just enough to make room for the escapes and reforms that help our model to explain the data and interpret the histories.

²⁴Without this resetting, the cosmetic reform of Marcet and Nicolini (2003) would have stayed with probability one for the next eleven months because of the unusually high values of expected hyperinflation.

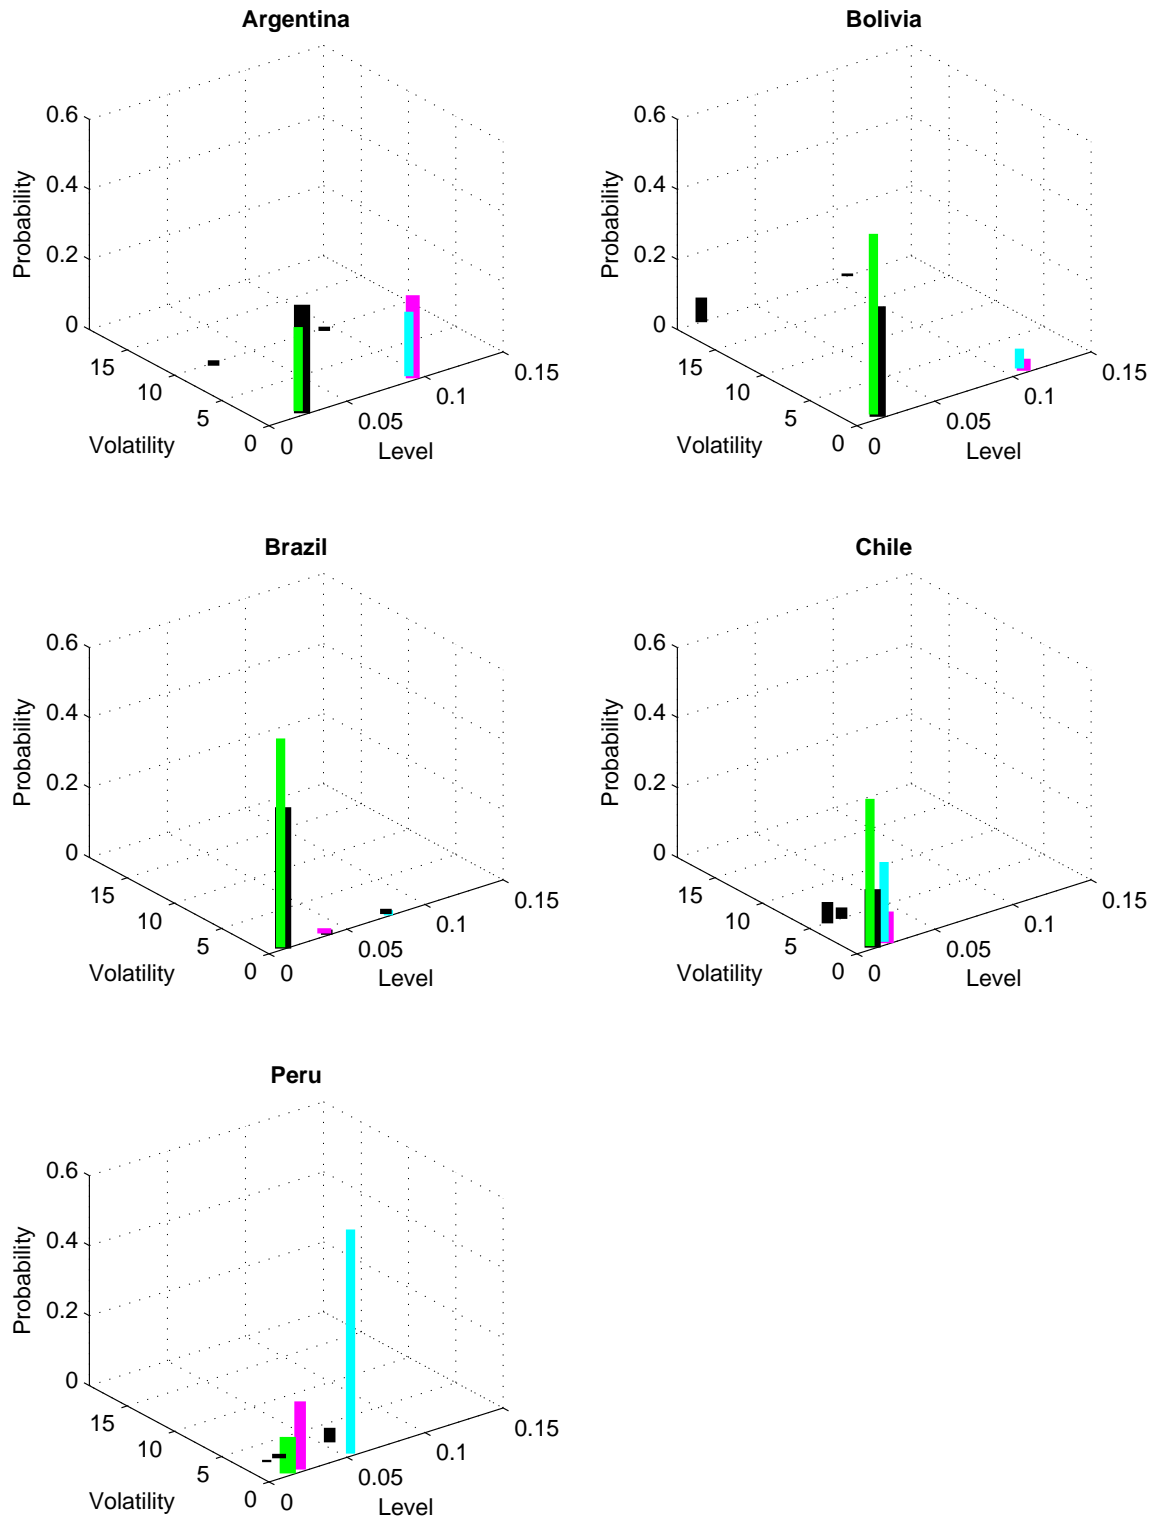


FIGURE 9. Ergodic probability given the estimated average deficit level (x-axis) and the estimated standard deviation of deficit shocks (y-axis).

APPENDIX A. DERIVING THE LIKELIHOOD

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. We assume that the probability distribution of $\eta_{\pi t}(k)$ is truncated log-normal and that the distribution of $\eta_{dt}(k)$ is log-normal for $k = 1, \dots, h$ where $h = m_h \times v_h$. This general setup includes both the case where m_t and v_t are independent and the case where $m_t = v_t$. Specifically, the probability density functions are

$$p_{\pi}(\eta_{\pi t}(k)) = \begin{cases} \frac{\exp\left[-\frac{[\log(\pi_1^*(k) + \eta_{\pi t}(k)) - \log \pi_1^*(k)]^2}{2\sigma_{\pi}^2}\right]}{\sqrt{2\pi}\sigma_{\pi}(\pi_1^*(k) + \eta_{\pi t}(k))\Phi\left(\frac{-\log(\delta) - \log(\pi_1^*(k))}{\sigma_{\pi}}\right)} & \text{if } -\pi_1^*(k) < \eta_{\pi t}(k) < 1/\delta - \pi_1^*(k), \\ 0 & \text{otherwise} \end{cases}, \quad (\text{A1})$$

$$p_d(\eta_{dt}(k)) = \begin{cases} \frac{\exp\left[-\frac{[\log(\bar{d}(k) + \eta_{dt}(k)) - \log \bar{d}(k)]^2}{2\sigma_d^2(k)}\right]}{\sqrt{2\pi}\sigma_d(k)(\bar{d}(k) + \eta_{dt}(k))} & \text{if } \eta_{dt}(k) > -\bar{d}(k), \\ 0 & \text{if } \eta_{dt}(k) \leq -\bar{d}(k) \end{cases}, \quad (\text{A2})$$

where $\Phi(x)$ is the standard normal cdf of x . We use the convention that $\log(0) = -\infty$ and $\Phi(-\infty) = 0$. Equation (A2) implies that the geometric mean of $d_t(s_t)$ is $\bar{d}(s_t)$. 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 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 (A1) and (A2), the conditional likelihood is

$$\begin{aligned}
p(\pi_t | \pi^{t-1}, s^T, \phi) &= p(\pi_t | \pi^{t-1}, s_t, \phi) \\
&= C_{1t} \frac{|\xi_\pi| \exp \left[-\frac{\xi_\pi^2}{2} (\log \pi_t - \log \pi_1^*(s_t))^2 \right]}{\sqrt{2\pi} \Phi(|\xi_\pi|(-\log(\delta) - \log(\pi_1^*(s_t))) \pi_t)} \\
&\quad + C_{2t} \left(\frac{\theta |\xi_d(s_t)| (1 - \lambda \beta_{t-1})}{\sqrt{2\pi} [(1 - \lambda \beta_t) \pi_t - \theta(1 - \lambda \beta_{t-1})] \pi_t} \right. \\
&\quad \left. \exp \left[-\frac{\xi_d^2(s_t)}{2} \left[\log[(1 - \lambda \beta_t) \pi_t - \theta(1 - \lambda \beta_{t-1})] - \log \pi_t - \log d(s_t) \right]^2 \right] \right), \tag{A3}
\end{aligned}$$

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

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

With some algebraic work, one can show from (A1) and (A2) that Equation (A3) is equivalent to the following expression

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

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

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

□

After integrating out s^T , the overall likelihood is

$$\begin{aligned}
p(\pi^T | \phi) &= \prod_{t=1}^T p(\pi_t | \pi^{t-1}, \phi) \\
&= \prod_{t=1}^T \left\{ \sum_{s_t=1}^h [p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)] \right\},
\end{aligned} \tag{A4}$$

where

$$\Pr(s_t | \pi^{t-1}, \phi) = \sum_{s_{t-1}=1}^h [\Pr(s_t | s_{t-1}, q) \Pr(s_{t-1} | \pi^{t-1}, \phi)]. \tag{A5}$$

The probability $\Pr(s_{t-1} | \pi^{t-1}, \phi)$ can be updated recursively. We follow Sims, Waggoner, and Zha (2006) and set

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

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

$$\Pr(s_t | \pi^t, \phi) = \frac{p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)}{\sum_{s_t=1}^h [p(\pi_t | \pi^{t-1}, s_t, \phi) \Pr(s_t | \pi^{t-1}, \phi)]}. \tag{A6}$$

As shown in Sims, Waggoner, and Zha (2006), one can also use the above recursive structure to compute the smoothed probability of s_t , $\Pr(s_t | \pi^T, \phi)$.

APPENDIX B. SELF-CONFIRMING EQUILIBRIA

This appendix describes self-confirming equilibrium versions of our model, while appendix C below describes rational expectations equilibria. We do not estimate either of these types of equilibria. However, by estimating the adaptive model of section II, we recover all the parameters that are required to compute such equilibria. In section VIII, we compute such equilibria for our estimated parameter values.

B.1. Small gain convergence. If the agents in our model were to implement a least squares estimator by replacing ε in the updating rule (6) by t^{-1} , we would expect β_t to converge to a constant level of expected inflation that equals the actual unconditional mean rate of inflation. Such a constant average level of gross inflation is a special case of a *self-confirming equilibrium* (SCE) as described by Sargent (1999).²⁵ We find such an unconditional SCE by computing a small gain limit for the beliefs of the adaptive agents under our model. We also consider a small variation limit in which the transition probabilities of the average deficit state m_t become degenerate, leading us to a notion of a *conditional* self-confirming equilibrium. As we've seen, these equilibria are important reference points for characterizing the belief dynamics.

B.1.1. Self-confirming equilibria. A self-confirming equilibrium (SCE) is a fixed point of beliefs, β , that is consistent with what the agents observe and solves the following population orthogonality condition:

$$E[\pi_t - \beta] = 0, \tag{A7}$$

where π_t is itself a function of β .

Let:

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

As we implement a “reform” by setting π_t randomly in the way described in equation (14), 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 (6) as:

$$\beta_{t+1} = \beta_t + \varepsilon g(\beta_t, \beta_{t-1}, d_t, \pi_t^*) \tag{A8}$$

²⁵We have assumed that agents do not know the current regime $s_t = (m_t, v_t)$ when forecasting inflation. Better informed agents would incorporate knowledge of s_t in forecasting inflation (see section C).

where

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

Let $F_{d(k)}(x)$ be the cdf of $d(k)$ at the value x . We define the following terms:

$$\begin{aligned} \tilde{g}(\beta, d_t, \pi_t^*) &= g(\beta, \beta, d_t, \pi_t^*), \\ \xi(\beta) &= (1 - \beta\lambda), \\ \Psi_k(\beta, b) &= \int_0^b \frac{1}{\xi(\beta) - x} dF_{d(k)}(x). \end{aligned}$$

It follows that $\bar{\omega}(\beta) \equiv \omega(\beta, \beta) = (1 - \delta\theta)\xi(\beta)$ and that $\Psi_k(\beta, b)$ is finite as $b \rightarrow \xi(\beta)$ because δ in equation (13) is bounded away from zero.

Recall that when a reform event takes place, $\pi_t^*(m_t)$ has a truncated log-normal distribution, and denote its mean as $\bar{\pi}^*(m_t)$. Also denote \bar{q}_k as the unconditional probability of the event $\{s_t = k\}$, which is an element of the ergodic distribution of Q . Then since $\log d(k) \sim N(\log \bar{d}(k), \sigma_d^2(k))$ we can write the unconditional expectation:

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

Proposition 4. As $\varepsilon \rightarrow 0$ the beliefs $\{\beta_t\}$ from (A8) converge weakly to the solution of the ordinary differential equation (ODE):

$$\dot{\beta} = G(\beta) \tag{A9}$$

for $\delta > 0$ and a broad class of probability distributions of $\eta_{dt}(s_t)$ and $\eta_{\pi t}(s_t)$ (including those specified in (A1) and (A2)).

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

The ODE (A9) governs the mean dynamics G , in that for small gains the belief trajectories tend to track those of the ODE. Thus, for beliefs to converge in a weak sense to an SCE, that SCE must be a stable equilibrium point of the ODE (A9). We don't have an explicit expression for G , so we shall find the SCE numerically. Thus, we look for a stationary point β^* such that $G(\beta^*) = 0$. Since the system is scalar the stability condition is simply $G'(\beta^*) < 0$. This gives formal content to the diagram in Figure 1 and the related discussion.

B.2. Conditional SCEs. For comparison, we are also interested in the self-confirming equilibria that would result if the economy were forever to remain in one average deficit regime. Thus, instead of (A7) the orthogonality condition is now:

$$E[\pi_t - \beta | m_t = m \forall t] = 0, \quad (\text{A10})$$

where again π_t is itself a function of β . We refer to these as conditional self-confirming equilibria. Since our estimated average deficit regimes are very persistent, we expect beliefs to adapt to these slowly varying regimes. We justify the consideration of conditional SCEs by refining our small-gain limit above. We now consider a two time-scale limit, in which the gain in the belief updating goes to zero but the probabilities of switching mean deficit regimes go to zero at a faster rate.

For simplicity, consider the case where the mean deficit regime state takes on two values: $m_t \in \{0, 1\}$. Then note that we can write the evolution of the mean deficit state m_t as:

$$m_{t+1} = m_t + \eta_{t+1}(m_t)$$

where $\eta_{t+1}(0) = 0$ with probability $Q_m(1, 1)$, $\eta_{t+1}(0) = 1$ with probability $Q_m(2, 1)$, $\eta_{t+1}(1) = 0$ with probability $Q_m(2, 2)$, and $\eta_{t+1}(1) = -1$ with probability $Q_m(1, 2)$. Thus, we have:

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

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

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

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 the martingale difference term v_{t+1} inherits the scaling. Thus, we extend the system that we analyze from (A8) to:

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

Now, 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 for the belief evolution we can effectively treat the mean deficit state as fixed. Thus, we extend our previous mean dynamics $G(\beta)$ above to the *conditional mean dynamics* $\hat{G}(\beta, m)$ that we now define. To do so, let $F_{\eta_d(v)}(x)$ be the cdf of $\eta_{dt}(v)$ at x , denote $\bar{q}_{v,k}$ as the unconditional probability of the event $\{v_t = k\}$, which is an

element of the ergodic distribution of Q_v , and define:

$$\Psi_v(\beta, b, m) = \int_0^{b-\bar{d}(m)} \frac{1}{\xi(\beta) - \bar{d}(m) - x} dF_{\eta_d(v)}(x).$$

Then we have:

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

Then we have the following result.

Claim 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):

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

for $\delta > 0$ and a broad class of probability distributions of $\eta_{dt}(s_t)$ and $\eta_{\pi t}(s_t)$ (including those specified in (A1) and (A2)).

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 conditional self-confirming equilibria are the collections of fixed points of the conditional dynamics, and thus are the β^* that satisfy $\hat{G}(\beta^*, m) = 0$. Again, for a conditional SCE to be a limit point of the learning it must be stable and thus satisfy $\hat{G}_\beta(\beta^*, m) < 0$. In section (VIII), we study how well the conditional SCE beliefs approximate rational expectations beliefs under our estimated parameters.

B.3. Qualifications. It is important to note that we have convergence in a weak sense of convergence in distribution. For any constant positive gain, when regimes change and as the deficit is hit by shocks, beliefs will continue to fluctuate. These fluctuations become proportionately smaller when the gain ε is smaller, but for any positive gain the beliefs will have a non-degenerate distribution. As the gain shrinks, this distribution collapses to a point mass on the solution of the ODE. Proposition 4 describes only the average behavior of beliefs for small gains. There may be extended

periods in which beliefs are away from the SCE, particularly when some regimes may be experienced for extended periods. This is why the conditional SCEs in Claim 1 are useful reference points for the analysis. Over the relatively long spells in which the average deficit state is constant the beliefs would tend to be centered on the conditional SCE. However, the escape dynamics play an important role as well, and they persist with positive gains.

APPENDIX C. RATIONAL EXPECTATIONS EQUILIBRIA

We now suspend the adaptive learning rule (6) and consider a subset of the rational expectations equilibria of the model. Further, while previously we've assumed that agents within the model do not observe the Markov state governing seignorage, we now assume that rational agents do condition on the average deficit state m_t .

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

$$\begin{aligned}\pi_t &= \pi(s_t, m_{t-1}, d_t) \\ E[\pi_{t+1}|m_t] &= E[\pi(s_{t+1}, m_t, \bar{d}(m_{t+1}) + \eta_{d,t+1}(v_{t+1}))|m_t] \\ &= \sum_{j=1}^{h_m} \sum_{k=1}^{h_v} \int \pi([j, k], m_t, \bar{d}(j) + x) dF_{\eta_d(k)}(x) \bar{q}_{v,k} Q_m(m_t, j) \\ &\equiv \pi^e(m_t),\end{aligned}$$

where we use the notation $s_t = [j, k] \equiv [m_t = j, v_t = k]$. Note that we assume that the state v_t governing seignorage shocks 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:

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

Again this only holds when the denominator is positive (which is the more stringent condition), so we truncate as above, giving:

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

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:

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

where $\xi_i = 1 - \pi^e(i)\lambda$ and Ψ_k is as above. Thus, we have h_m 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. Alternatively, any other numerical nonlinear equation solver can be used.

C.2. Multiplicity and nonexistence. Though there are multiple rational expectations equilibria of the model, there is typically a unique SCE that is stable under learning. As we've seen, in the deterministic counterpart of the model there are two REEs. With small enough shocks, we also find that there are two conditional SCEs in each regime. As discussed above, the SCEs average across these conditional SCEs. Thus, for example with two possible regimes and two conditional SCEs in each regime, there would typically be two SCEs, with one of them stable. REEs also average across the conditional SCEs, taking into account the probability of regime switches. So, for example, with two conditional SCEs in each regime, there are typically four REEs that switch between values close to the conditional SCEs in each regime. However, when shocks to seignorage become large enough there may be only one conditional SCE in a regime, or it could even occur that a conditional SCE fails to exist altogether. Depending on the weight that these high-shock regimes have in the invariant distribution, the SCE may also fail to exist. Similarly, there may be fewer rational expectations equilibria or none at all.

As we've seen, these observations are empirically relevant, as in some countries our estimates imply very large seignorage shocks in some regimes. Nevertheless, in all cases we find that a stable SCE exists, even though there may not be a conditional SCE in the high shock regimes. This suggests that beliefs may tend to diverge in the regimes with high shocks, with agents expecting ever-growing inflation (up to the truncation point). But the regimes usually will not last long enough for this to actually happen, and the lower shock regimes tend to bring beliefs back down.

APPENDIX D. SEIGNIORAGE 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 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 $\text{Pr}(s_t | \hat{\phi}, \pi^T)$ 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 for the year containing all these months indexed by t . The simulated data $d_{A,t}$ is only an approximation to the actual data $d_{A,t}^{\text{Data}}$ because of these 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 deficits to the average of actual annual deficits 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 deficits are caused by a large variance of deficit shocks. 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 6. Note that changes in shock variances has no effect on the median of simulated deficits.

APPENDIX E. ADDITIONAL TABLES AND FIGURES

Tables 6-10 present the full maximum likelihood estimates of our model for the five countries in our sample.

TABLE 6. Argentina: MLEs for the 2×3 regime-switching model

$\lambda : 0.730 (0.0104)$		
$[\bar{d}(1); \bar{d}(2)] : [0.0937 (0.0009); 0.0228 (0.0002)]$		
$[\xi_d(1); \xi_d(2); \xi_d(3)] : [0.104 (0.050); 1.482 (0.074); 3.784 (0.226)]$		
$\xi_\pi : 16.78 (5.178)$		
$\epsilon : 0.023 (0.001)$		
Transition probability matrix Q_m for $\bar{d}(s_t)$:		
0.9789 (0.014)		0.0162
0.0211		0.9838 (0.007)
Transition probability matrix Q_v for $\eta_{dt}(s_t)$:		
0.4395 (0.139)	0.0370	0.0000
0.5605	0.9260 (0.021)	0.0287
0.0000	0.0370	0.9713 (0.018)

Note: the numbers in the parentheses are estimated standard errors.

TABLE 7. Bolivia: MLEs for the 2×3 regime-switching model

$\lambda : 0.307 (0.038)$		
$[\bar{d}(1); \bar{d}(2)] : [0.1088 (0.0078); 0.0151 (0.0006)]$		
$[\xi_d(1); \xi_d(2); \xi_d(3)] : [0.053 (0.0396); 1.322 (0.0732); 3.252 (0.3157)]$		
$\xi_\pi : 26.52 (2.5114)$		
$\epsilon : 0.232 (0.0375)$		
Transition probability matrix Q_m for $\bar{d}(s_t)$:		
0.9629 (0.0237)		0.0041
0.0371		0.9959 (0.0028)
Transition probability matrix Q_v for $\eta_{dt}(s_t)$:		
0.3344 (0.1067)	0.0910	0.0000
0.6656	0.8180 (0.0426)	0.1487
0.0000	0.0910	0.8513 (0.0405)

Note: the numbers in the parentheses are estimated standard errors.

TABLE 8. Brazil: MLEs for the 3×2 regime-switching model

$\lambda : 0.613 (0.0073)$		
$[\bar{d}(1) \bar{d}(2) \bar{d}(3)] : [0.0771 (0.0020); 0.0375 (0.0006); 0.0096 (0.0001)]$		
$[\xi_d(1) \xi_d(2)] : [2.818 (0.1672); 10.929 (1.0010)]$		
$\xi_\pi : 9.18 (10.8305)$		
$\epsilon : 0.189 (0.0118)$		
Transition probability matrix Q_m for $\bar{d}(s_t)$:		
0.9845 (0.0127)	0.0134	0.0000
0.0155	0.9732 (0.0224)	0.0000
0.0000	0.0134	1.0000
Transition probability matrix Q_v for $\eta_{dt}(s_t)$:		
0.9344 (0.0292)		0.0969
0.0656		0.9031 (0.0338)

Note: the numbers in the parentheses are estimated standard errors.

TABLE 9. Chile: MLEs for the 2×3 regime-switching model

$\lambda : 0.875 (0.0000)$		
$[\bar{d}(1) \bar{d}(2)] : [0.0200 (0.0000); 0.0110 (0.0000)]$		
$[\xi_d(1) \xi_d(2) \xi_d(3)] : [0.203 (0.0619); 2.298 (0.1036); 6.985 (0.5367)]$		
$\xi_\pi : 10.62 (3.8807)$		
$\epsilon : 0.025 (0.0000)$		
Transition probability matrix Q_m for $\bar{d}(s_t)$:		
0.9869 (0.0051)		0.0070
0.0131		0.9930 (0.0076)
Transition probability matrix Q_v for $\eta_{dt}(s_t)$:		
0.7627 (0.0740)	0.0345	0.0000
0.2373	0.9310 (0.0193)	0.0869
0.0000	0.0345	0.9131 (0.0289)

Note: the numbers in the parentheses are estimated standard errors.

TABLE 10. Peru: MLEs for the 3×2 regime-switching model

$\lambda : 0.740 (0.0001)$		
$[\bar{d}(1) \bar{d}(2) \bar{d}(3)] : [0.0542 (0.0003); 0.0219 (0.0001); 0.0139 (0.0001)]$		
$[\xi_d(1) \xi_d(2)] : [0.394 (0.0714); 3.208 (0.1107)]$		
$\xi_\pi : 15.97 (2.4004)$		
$\epsilon : 0.069 (0.0025)$		
Transition probability matrix Q_m for $\bar{d}(s_t)$:		
0.9943 (0.0076)	0.0187	0.0000
0.0057	0.9626 (0.0142)	0.0350
0.0000	0.0187	0.9650 (0.0166)
Transition probability matrix Q_v for $\eta_{dt}(s_t)$:		
0.3016 (0.1310)		0.0453
0.6984		0.9547 (0.0127)

Note: the numbers in the parentheses are estimated standard errors.

Figures 10-14 plot the smoothed (two-sided) probabilities of the regimes conditional on our estimates and the full data sample. Each panel of the figures plots the probability of a particular combination (m, v) of the mean deficit and deficit shock variance states.

Figures 15 - 19 summarize the relative fit of our model for each country. In each figure, the top panel plots the inflation data along with the 90% probability bands for the one-step ahead forecasts from our model. The second panel plots the time series of the log of the conditional likelihood for our model, while the third panel plots the corresponding time series for the best-fitting regime-switching autoregressive statistical model. The bottom panel plots the difference between the two previous log conditional likelihoods. Together these bottom three panels summarize the relative fit of our model versus the statistical model in different time periods.

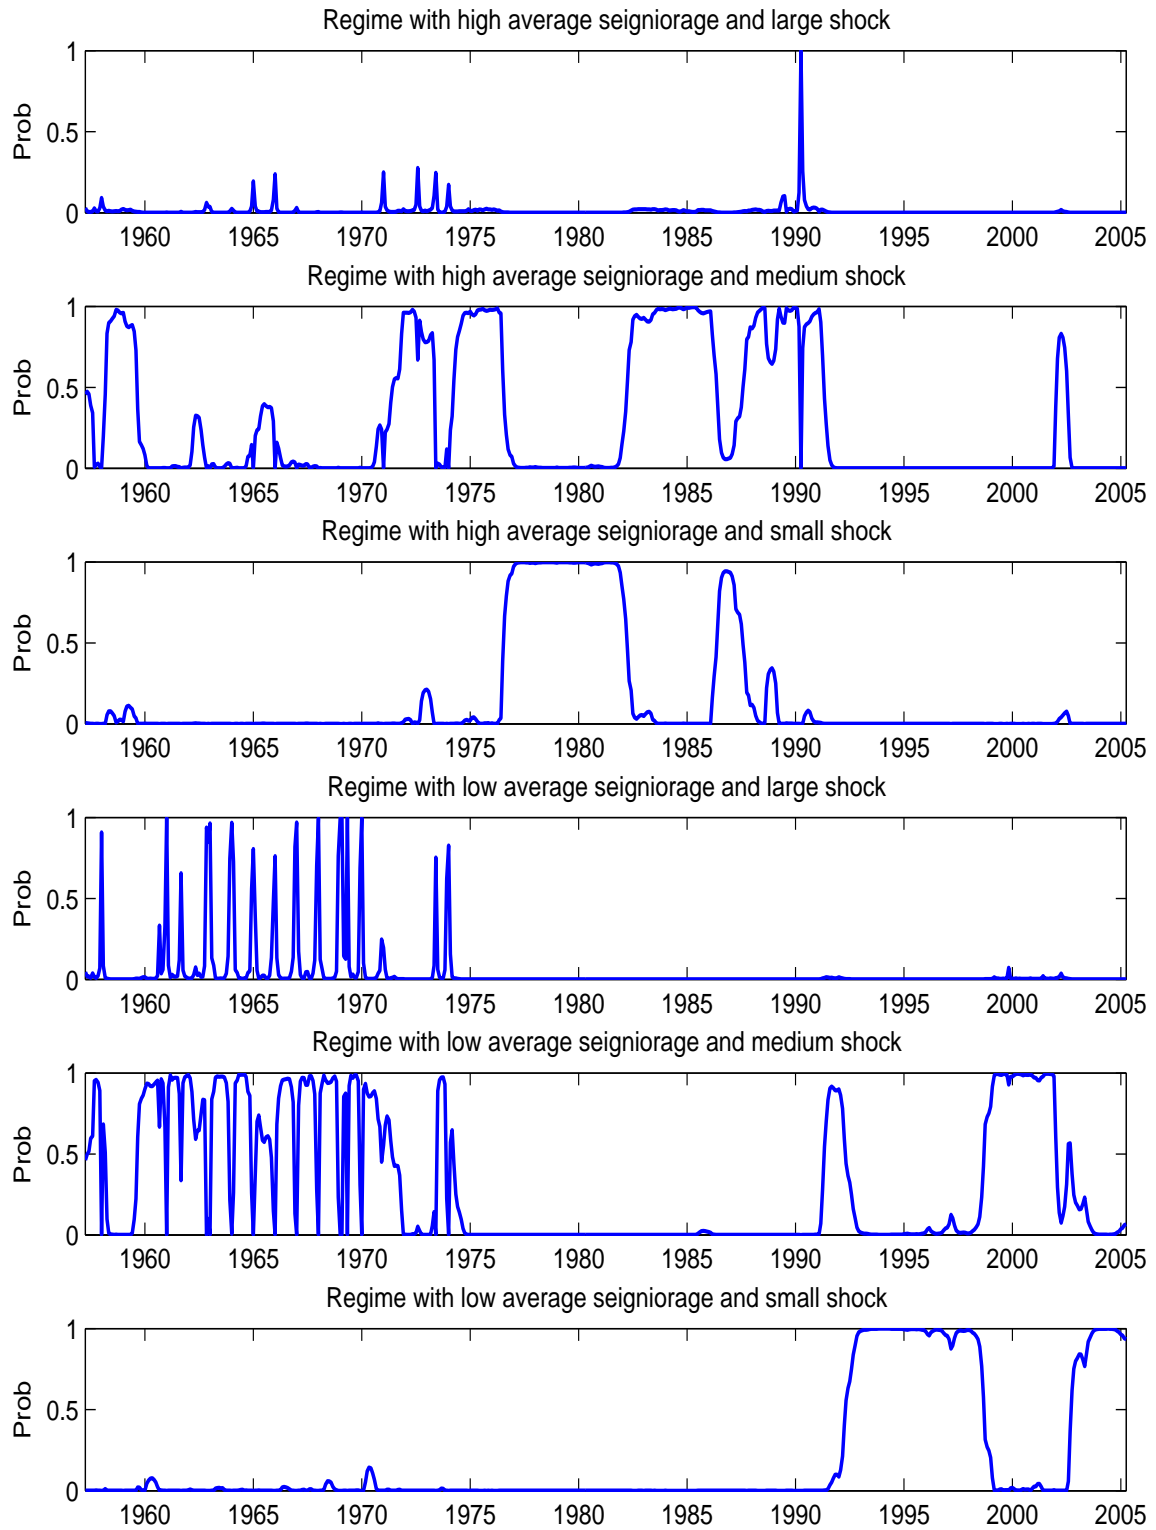


FIGURE 10. Argentina: smoothed probability of the regimes conditional on the MLEs and the data.

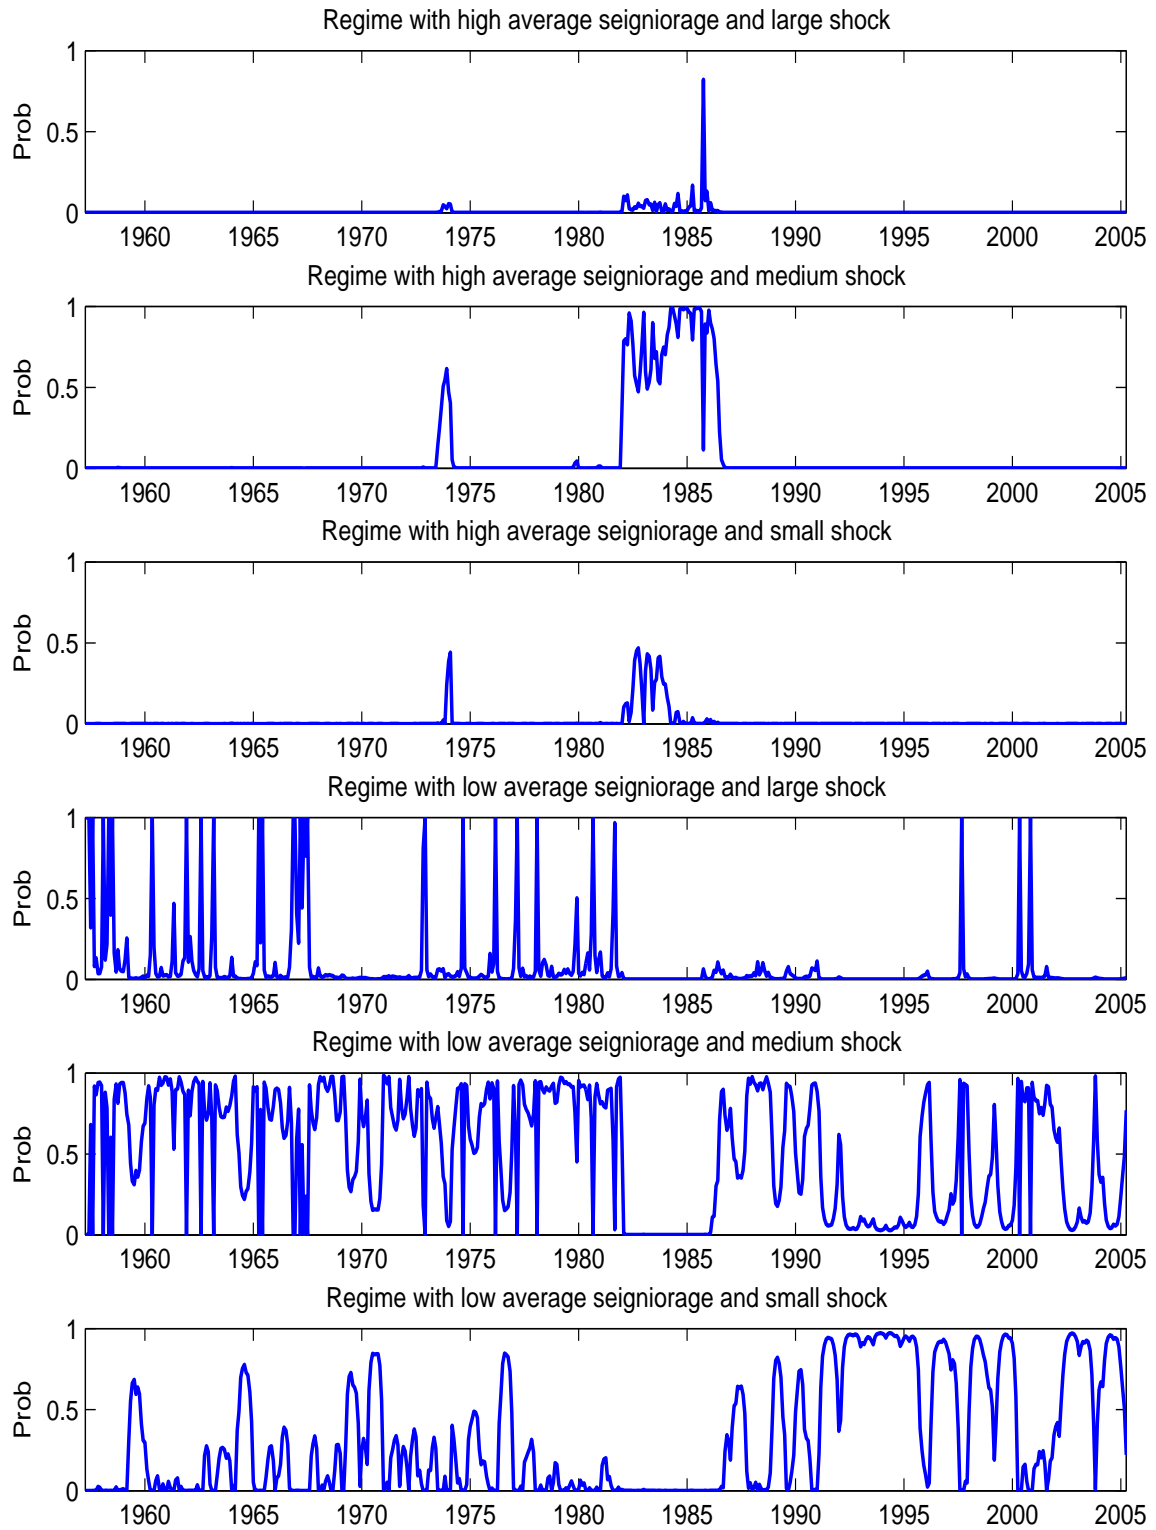


FIGURE 11. Bolivia: smoothed probabilities of the regimes conditional on the MLEs and the data.

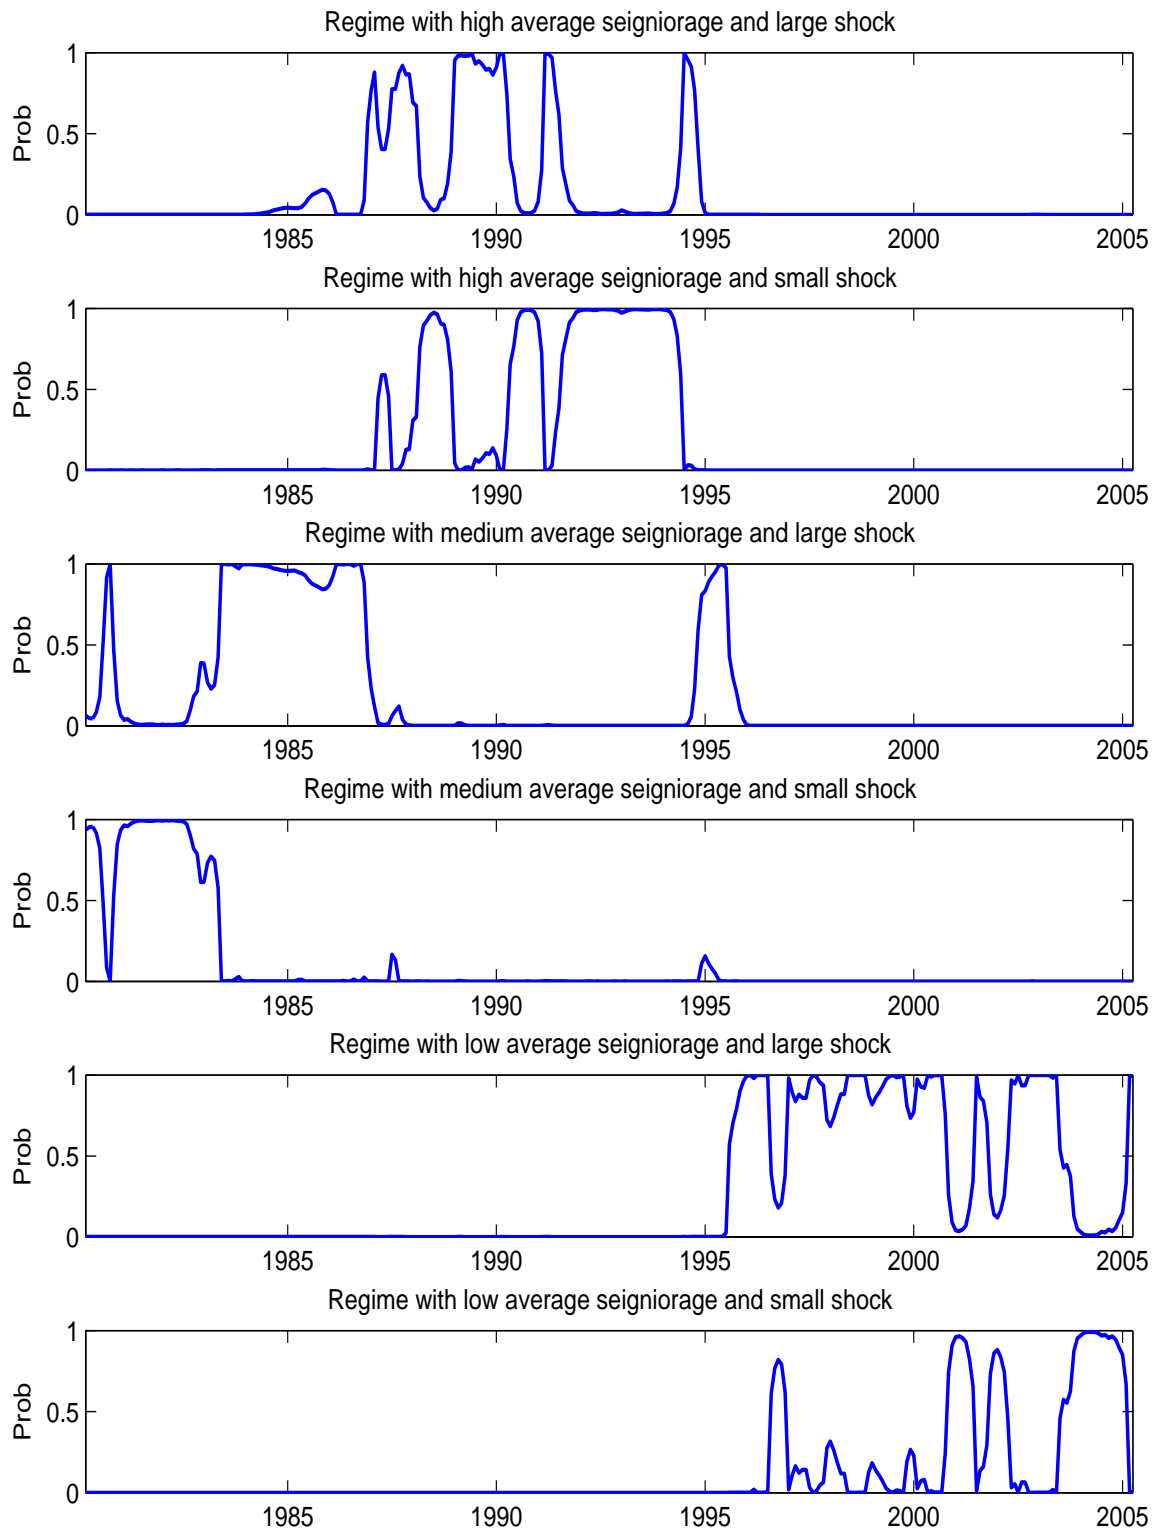


FIGURE 12. Brazil: smoothed probabilities of the regimes conditional on the MLEs and the data.

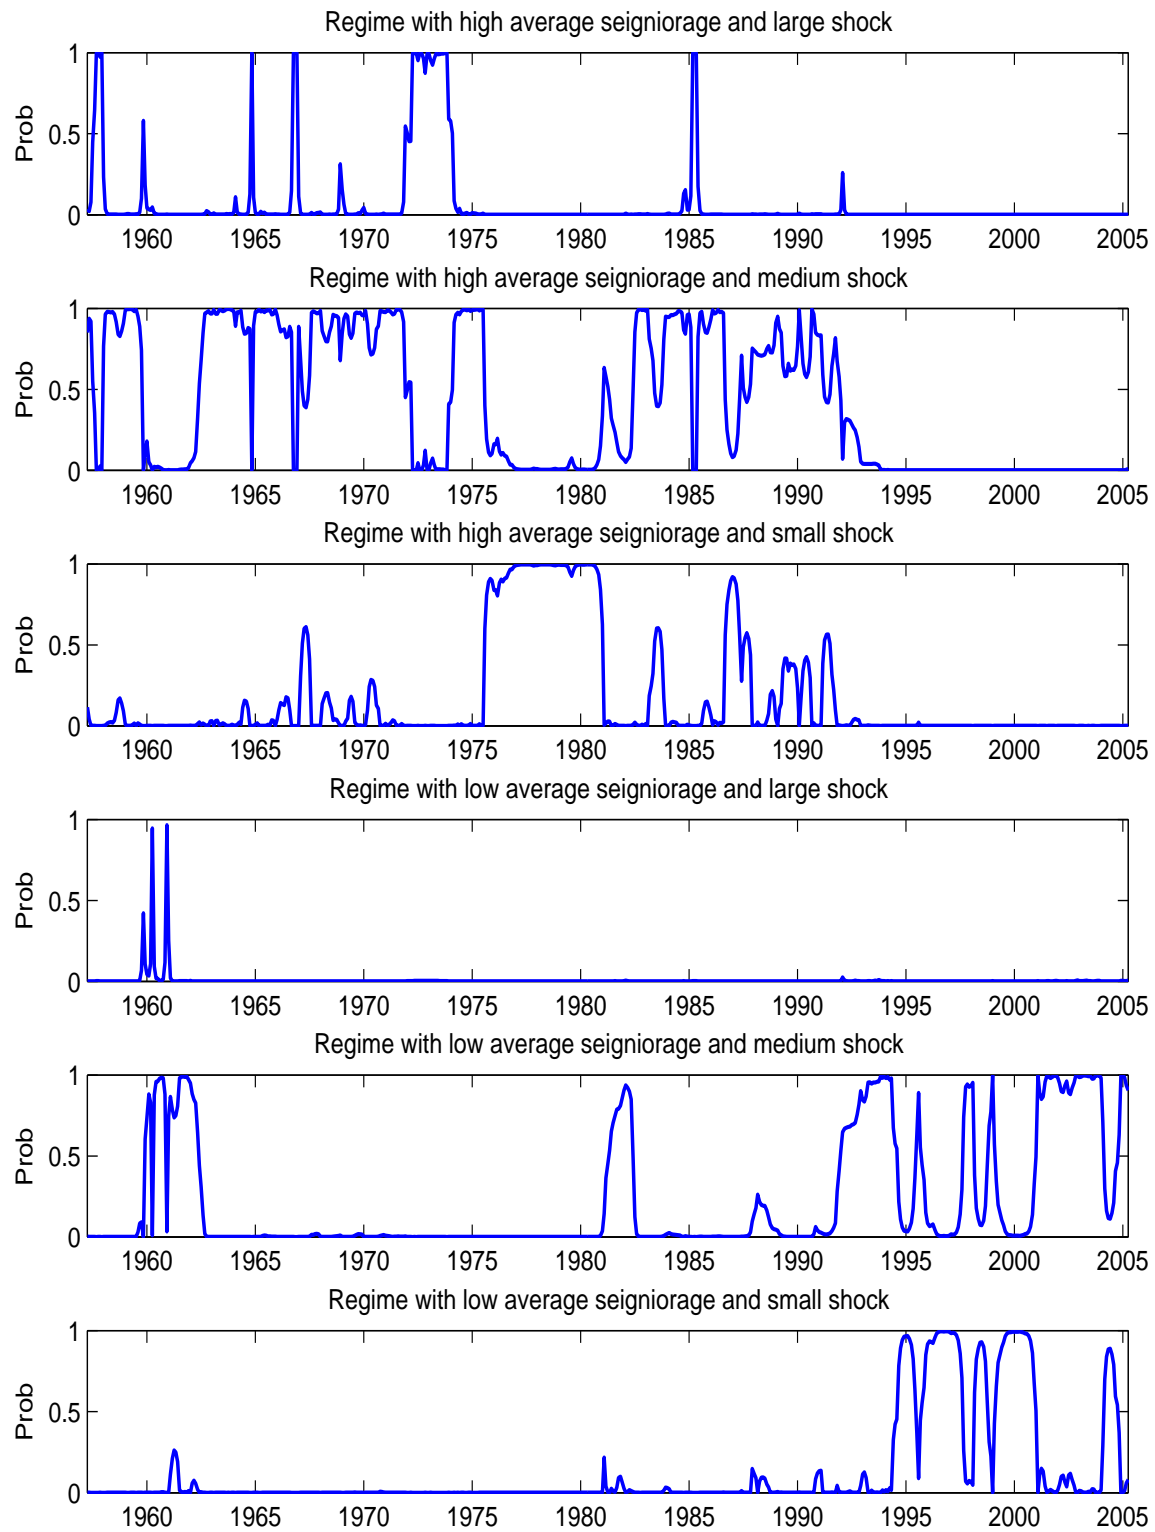


FIGURE 13. Chile: smoothed probabilities of the regimes conditional on the MLEs and the data.

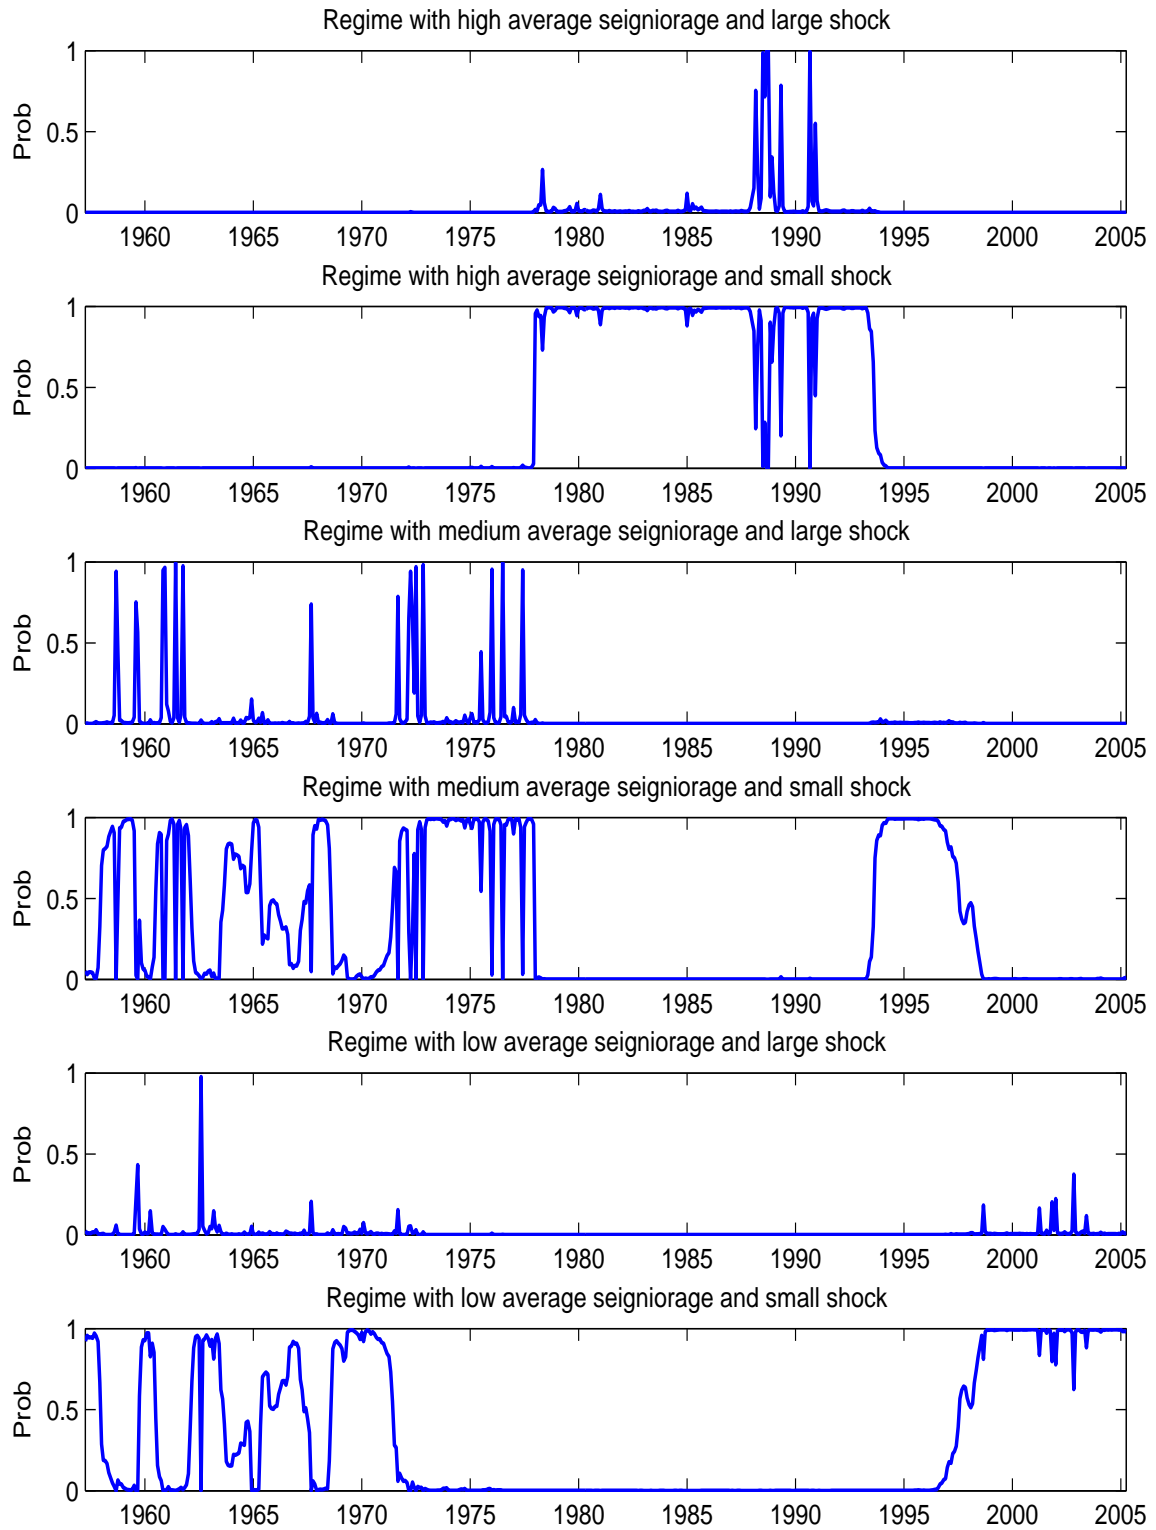


FIGURE 14. Peru: smoothed probabilities of the regimes conditional on the MLEs and the data.

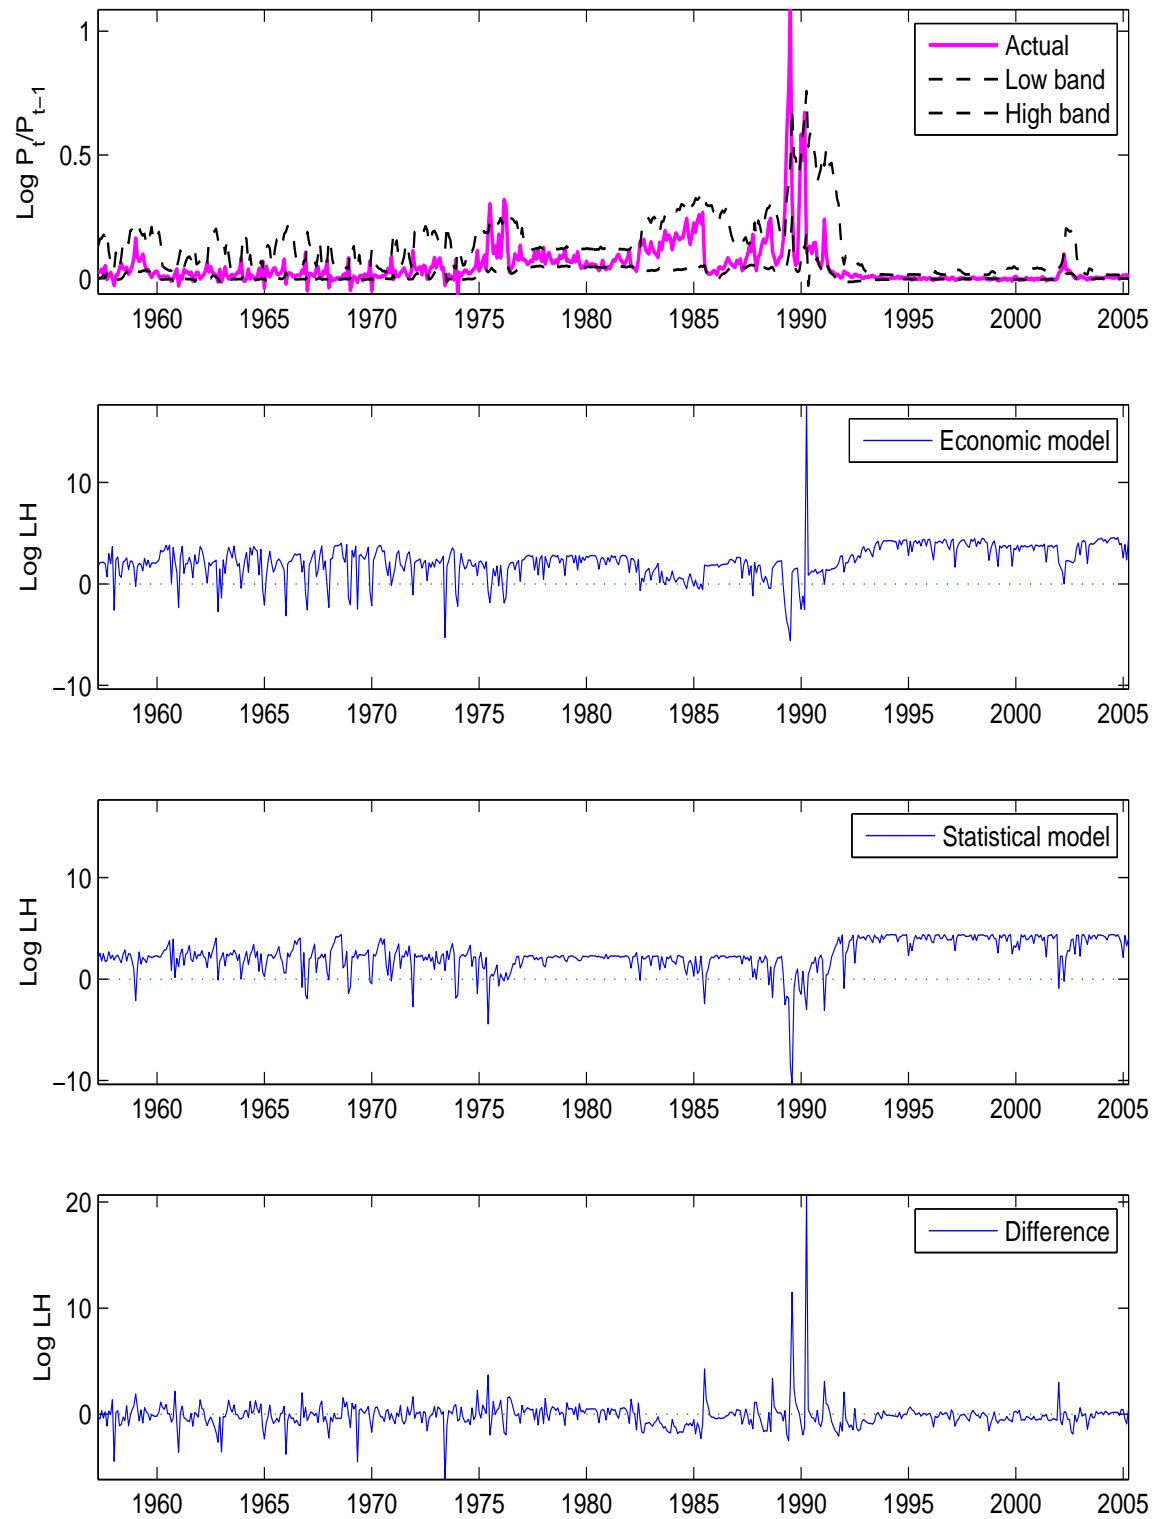


FIGURE 15. Argentina: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

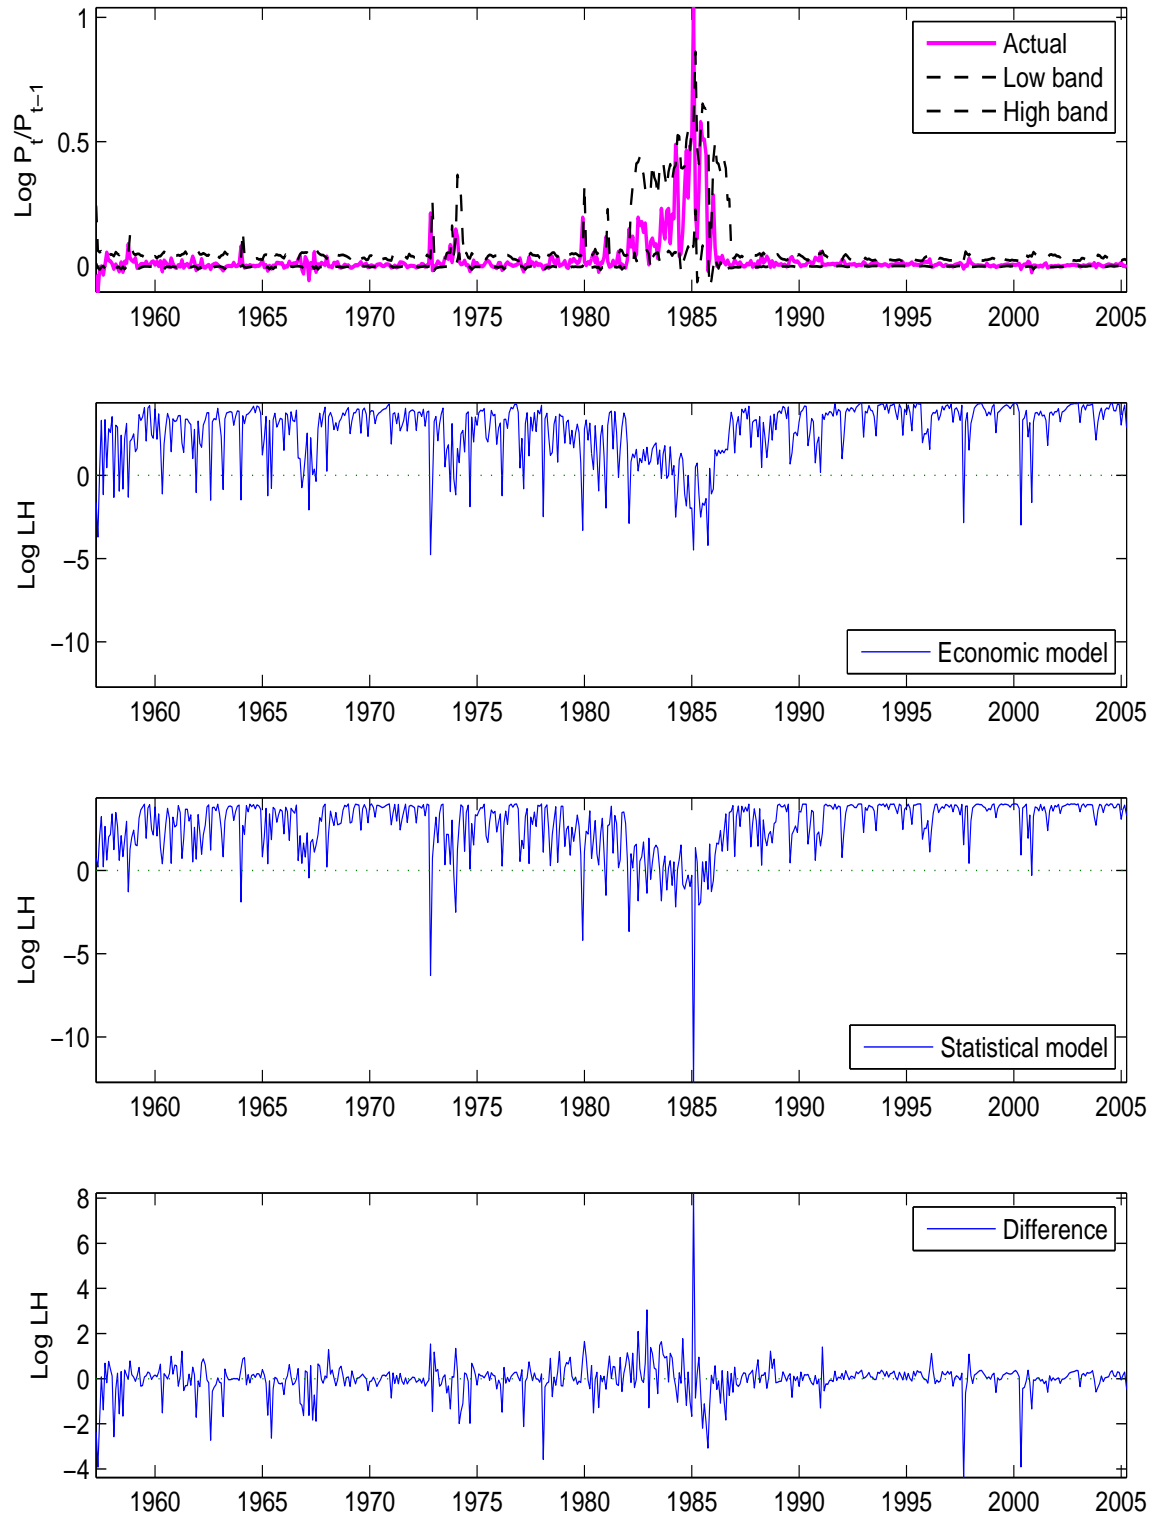


FIGURE 16. Bolivia: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

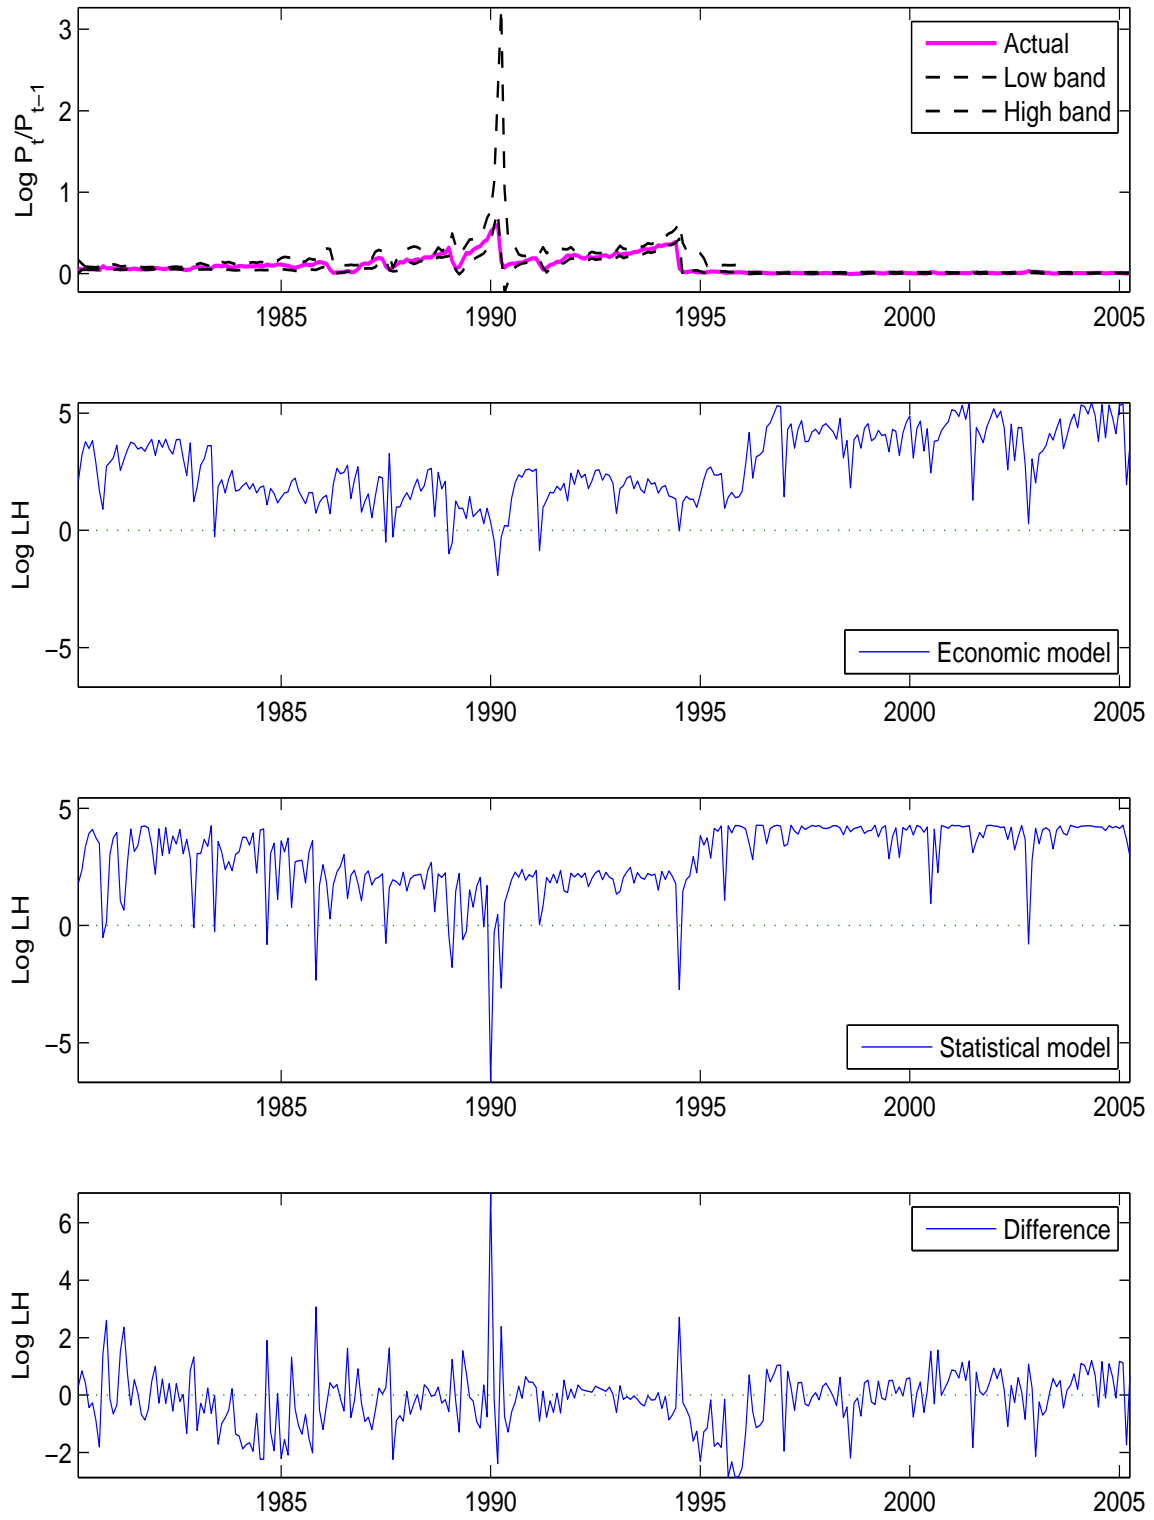


FIGURE 17. Brazil: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

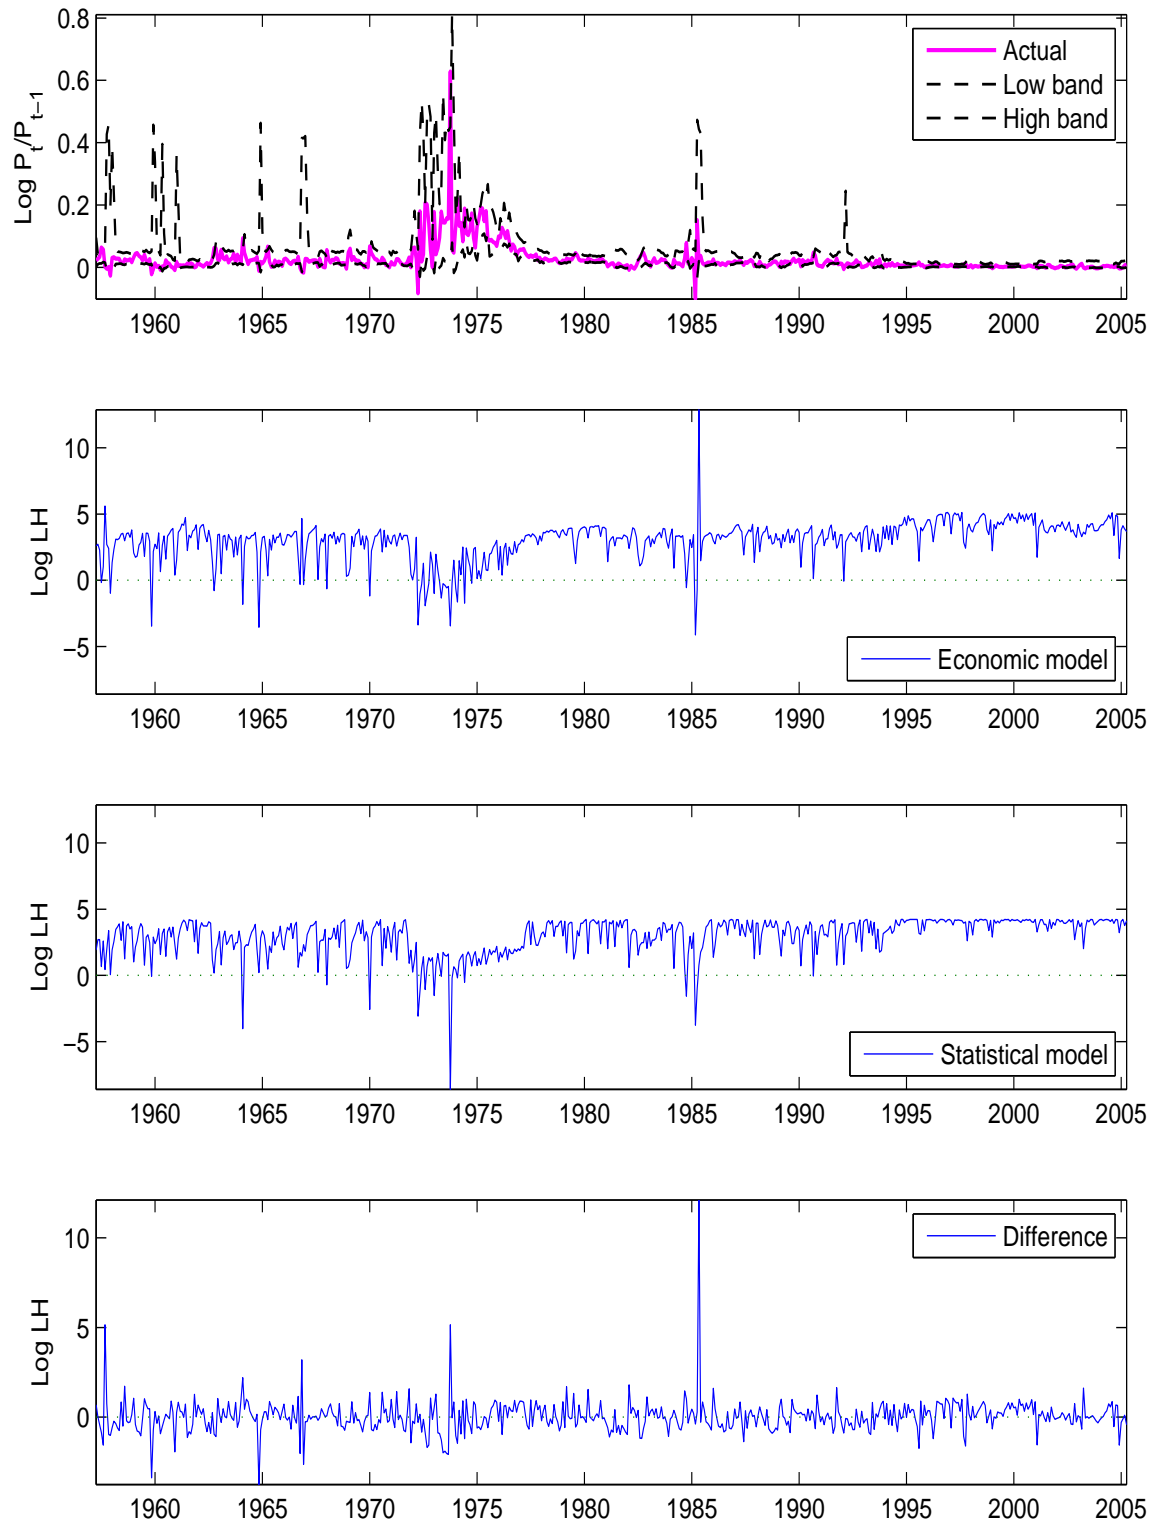


FIGURE 18. Chile: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

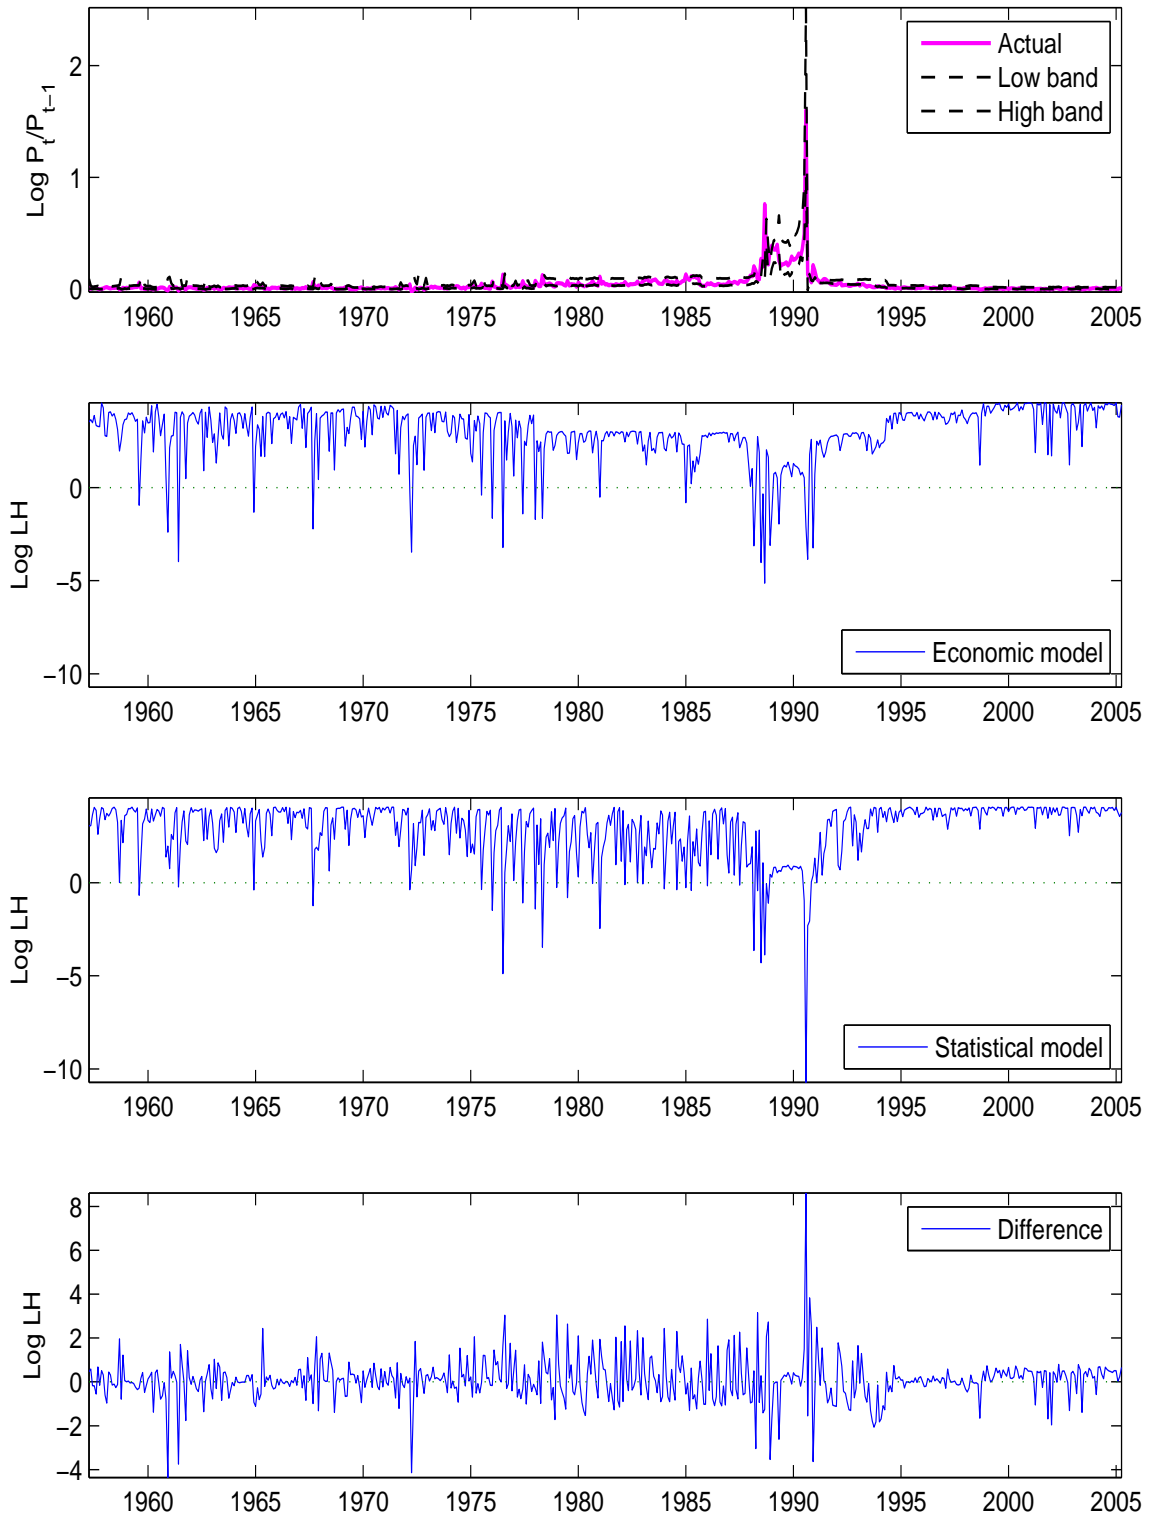


FIGURE 19. Peru: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

REFERENCES

- ADAM, K., G. W. EVANS, AND S. HONKAPOHJA (in press): "Are Hyperinflation Paths Learnable?," *Journal of Economic Dynamics and Control*.
- 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.
- CHO, I.-K., AND K. KASA (2003): "Learning Dynamics and Endogenous Currency Crises," *Computing in Economics and Finance 2003*, Society for Computational Economics.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.
- DORNBUSCH, R. (1985): "Comment," in *Inflation and indexation: Argentina, Brazil, and Israel*, ed. by J. Williamson. MIT Press, Cambridge, MA.
- ELLIOTT, R., L. AGGOUN, AND J. MOORE (1995): *Hidden Markov Models: Estimation and Control*. Springer-Verlag, New York and Berlin.
- ELLISON, M., L. GRAHAM, AND J. VILMUNEN (2006): "Strong contagion with weak spillovers," *Review of Economic Dynamics*, 9(2), 263–283.
- EVANS, G., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- 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," *The 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 (2004): "Normalization in Econometrics," Federal Reserve Bank of Atlanta Working Paper 2004-13.
- IMROHOROGLU, S. (1993): "Testing for Sunspot Equilibria in the German Hyperinflation," *Journal of Economic Dynamics and Control*, 17(3), 289–317.
- KASA, K. (2004): "Learning, Large Deviations, And Recurrent Currency Crises," *International Economic Review*, 45, 141–173.

- 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.
- LUCAS, JR., R. E. (1986): "Adaptive Behavior and Economic Theory," *Journal of Business*, 59, S401–S426.
- LUCAS, JR., R. E., AND T. J. SARGENT (1981): *Rational Expectations and Econometrics Practice*. University of Minnesota Press, Minneapolis, Minnesota.
- MARCET, A., AND J. P. NICOLINI (2003): "Recurrent Hyperinflations and Learning," *The American Economic Review*, 93(5), 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989a): "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.
- (1989b): "Least-Squares Learning and the Dynamics of Hyperinflation," 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.
- MONDINO, G., F. STURZENEGGER, AND M. TOMMASI (1996): "Recurrent High Inflation and Stabilization: A Dynamic Game," *International Economic Review*, 37, 981–996.
- SARGENT, T. J. (1981): "Interpreting Economic Time Series," *Journal of Political Economy*, 89(2), 213–48.
- (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., AND N. WILLIAMS (2003): "Impacts of Priors on Convergence and Escapes from Nash Inflation," Manuscript, New York University and Princeton University.
- SCLOVE, S. L. (1983): "Time-Series Segmentation: A Model and a Method," *Information Sciences*, 29(1), 7–25.
- SIMS, C. A. (2001): "Stability and Instability in US Monetary Policy Behavior," Manuscript, Princeton University.

- SIMS, C. A., D. F. WAGGONER, AND T. ZHA (2006): “Generalized Methods for Restricted Markov-Switching Models with Independent State Variables,” Unfinished Manuscript.
- SIMS, C. A., AND T. ZHA (2006): “Were There Regime Switches in US Monetary Policy?,” *The American Economic Review*, 96, 54–81.
- 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.
- TETLOW, R., AND P. VON ZUR MUEHLEN (2004): “Avoiding Nash Inflation: Bayesian and robust responses to model uncertainty,” *Review of Economic Dynamics*, 7, 869–899.
- ZARAZAGA, C. (1993): *Hyperinflations, Institutions, and Moral Hazard in the Appropriation of Seignorage*. University of Minnesota, Minneapolis, USA, ph.d. dissertation edn.