THE CONQUEST OF SOUTH AMERICAN INFLATION

THOMAS SARGENT, NOAH WILLIAMS, AND TAO ZHA

ABSTRACT. We use the method of maximum likelihood to infer determinants of Latin American hyperinflations and stabilizations. Our model assigns potential 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 the fundamentals. Our maximum likelihood estimates allow us to interpret observed inflation rates in terms of shifts in the deficits, sequences of shocks that trigger temporary episodes of expectations driven hyperinflations, and occasional superficial reforms that correct expectations without reforming deficits. Our estimates also allow us to infer from inflation data the fiscal patterns that seem to have stabilized inflation to a low level on a permanent basis.

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

Date: July 10, 2006.

This paper has benefited from help and encouragement of many people. We thank Eduardo Ganapolsky, Karsten Jeske, Albert Marcet, and especially Christopher Sims, Dan Waggoner, and Michael Woodford for helpful discussions. Sagiri Kitao, Tomasz Piskorski, and Demian Pouzo provided their outstanding research assistance; Namgeun Jeong and Eric Wang provided indispensable assistance on clustering and parallel computing in the Linux operating system. We acknowledge the technical support of the Computing College of Georgia Institute of Technology on grid computation techniques. Finally, we thank Mike Chriszt, 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 herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System.

> 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, Brasil, 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 but 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 lack trustworthy monthly data on deficits and money supplies but trust our monthly series on inflation. 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 those densities with the monetary deficit data that we do have.

The reason that we posit an adaptive expectations scheme is not to turn the clock back to the days before the hall-mark 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

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

[2] Lucas (1986) and Marcet and Sargent (1989a,b) used adaptive expectations schemes to provide support for rational expectations as an equilibrium concept supported by a law of large numbers and to select among multiple rational expectations equilibria.

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 a 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. 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.

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 steadily 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 can push β_t above π_2^*, the shaded 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 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

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

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

increases in government deficits. 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 the deficit and the reform event. And 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 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: an exogenous jump 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 Imrohoroglu (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).

I.4. **Organization.** The remainder of this paper is organized as follows. Section II presents our model and a brief description of likelihood function for histories of inflation, consigning important details about the likelihood to appendix II.5. Section III gives a brief account of the concepts of self-confirming equilibria and conditional

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

self-confirming equilibria that will guide our empirical interpretations, consigning computational details to appendix B. Section V describes our estimation procedures and results. Section VI then assembles these results into a set of interpretable economic findings. To help us understand how far we have drifted from rational expectations, section **??** computes stationary points of unconditional and conditional self-confirming equilibria and compares them with stationary rational expectations inflation rates evaluated at our maximum likelihood parameter estimates. Section VIII concludes. Appendix **??** explores the fit of our model relative to some alternatives, while appendix C describes how we compute a subset of the rational expectations equilibria at our maximum likelihood parameter estimates.

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. 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 II.5.

The model consists of a demand function for money, a government budget constraint, 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]

$$\frac{M_t}{P_t} = \frac{1}{\gamma} - \frac{\lambda}{\gamma}\frac{P^e_{t+1}}{P_t}, \tag{1}$$

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

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

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

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

[6] Using 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.

[7] For 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.

where $0 < \lambda < 1$, $0 < \theta < 1$, $\gamma > 0$, $\bar{d}(s_t) > 0$, $s_t = \begin{bmatrix} m_t & v_t \end{bmatrix}$ is a Hamilton (1989)-Sclove (1983) Markov state that neither the agent inside the model nor we the econometricians observe; P_t is the price level at time t; M_t is nominal balances in percent of real output at time t; P_{t+1}^e is the public's expectation of the price level at time $t+1$; and $\eta_{d\,t}(m_t, v_t)$ 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 we assume equals the average amount of seigniorage financed by money creation in state m_t. The two Markov chains Q_m, Q_v induce a chain Q on the composite state $s_t = \begin{bmatrix} m_t & v_t \end{bmatrix}$ so that the transition matrix for this composite state is $Q = Q_m \otimes Q_v$.[8] Thus, a total number of states is $h = m_h \times v_h$.

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

$$\pi_{t+1}^e = \beta_t$$

where the superscript b stands for belief. The public updates belief β_t by using a constant-gain algorithm:

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

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

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

Model (1)-(5) makes inflation dynamics depend on $\gamma d_t(k)$ where $k \in \{1, \ldots, h\}$ 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 original system (1)-(5) remains the same. 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. □

The normalization is effectively a choice of units for the price level, which our model is silent about because we use it to deduce only the joint density over inflation sequences. 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

[8]We have also modelled 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.

identified separately, so that re-normalizing them in the manner of Proposition 1 gives the same equilibrium outcome.[9] For identification purposes, therefore, we normalize $\gamma = 1$ when searching for the maximum likelihood estimates and re-normalize it when comparing our computed d_t with the 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.

II.1. **Deterministic steady states.** For each state m_t, a deterministic version of model (1) - (5) can be obtained by fixing the state $m_t = m \in \{1, \ldots, m_h\}$ and setting $\eta_{dt}(v_t)$ to zero for all t .

Proposition 2. If

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

then there exist two steady state equilibria for π_t:

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

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

Proof. Sargent and Wallace (1987) show that

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

In stationary equilibrium, $\pi_t = \pi_{t-1}$. Substituting this equality into the above equation leads to (8) and (9). □

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

$$\pi_{\text{SS}}^{\max} \equiv \sqrt{\frac{\theta}{\lambda}}. \tag{10}$$

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

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 $\beta_0 < \pi_2^*(k)$. 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 dynamics property that stationary inflation *rises* when seigniorage *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.[10] 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 inflation

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

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:

$$1 - \lambda\beta_{t-1} > 0, \tag{12}$$

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

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

[10] Marcet and Nicolini (2003) began their study with the observation that Latin American data seem pock marked with recurrent big inflations of this type.

II.4. **Cosmetic reforms.** It is possible for a sequence of seigniorage shocks $\eta_{d\,t}$ to push β_t beyond a point when the inflation dynamics will continue 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 a random variable:

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

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

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

When $1 - \lambda\beta_t \leq \delta\theta(1 - \lambda\beta_{t-1})$, we also reset expectations so that $\beta_t = \pi_t^*(m_t)$.[11]

We use this device 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.** We denote the free parameters of the model as $\phi = \begin{bmatrix} \lambda & \bar{d}(m) & \xi_d(v) & \xi_\pi & \varepsilon & Q^m & Q^v \end{bmatrix}$ where $m = 1, \ldots, m_h$ and $v = 1, \ldots, v_h$. 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 $p(\pi^T, m^T, v^T, d^T, M^T, \beta^T|\phi)$, where we set $\beta_0 = \pi_0$ and the probabilities for the initial unobservable states m_0 and v_0 are set as 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 $\hat{\phi} = \text{argmax}_\phi p(\pi^T|\phi)$. We make inferences about the deficit from the conditional density $p(d^T|\pi^T, \hat{\phi})$. We consign a detailed description of how we constructed the likelihood function to appendix A.

III. Self-Confirming Equilibria

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

[11] We do this in order to guarantee that the model always implies positive price levels.

[12] There 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.

when $\varepsilon \downarrow 0$. We also describe functions corresponding to the $G(\beta)$ function in figure 1 that govern the dynamic behavior of β_t as $\varepsilon \downarrow 0$. Appendix C then defines rational expectations equilibria and links conditional and unconditional self-confirming equilibria to them.

An unconditional self-confirming equilibrium (SCE) is a β that satisfies $E\pi_t - \beta = 0$. For each Markov state k, a conditional self-confirming equilibrium is a $\beta(m)$ that satisfies $E[\pi_t|m_t = m] - \beta(m) = 0$ for $m \in \{1, \ldots, m_h\}$. A conditional SCE is just an unconditional SCE computed on the (false) assumption that the mean deficit state m_t will always be m. A conditional SCE is of interest of us because we estimate transition matrices Q_m that imply that the mean deficit state is 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.

Like the deterministic steady state REEs, if there exists an SCE, in general there exist two of them for each state m. We denote them as $\beta_1^*(m)$ and $\beta_2^*(m)$, where $\beta_1^*(m) \leq \beta_2^*(m)$.

For each country, we shall construct mean dynamics for the conditional SCE's for the estimated mean deficit in each deficit state m. We report them for each of our five countries in the top left panels of figures 3, 4, 5, 6, and 7. These curves are shaped just like $G(\beta)$ in figure 1. They will help us 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. And later, they'll help us evaluate how much our model deviates from rational expectations.

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.[13] 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.

[13]For analysis of escape events in other dynamic models, see Sargent (1999), Cho, Williams, and Sargent (2002), and Sargent and Williams (2003).

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:

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

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

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 is

$$\boldsymbol{\iota}\{\beta_{t-1} < 1/\lambda\} \sum_{k=1}^{h} \left[\Pr\left(s_t = k \mid \Pi_{t-1}, \phi\right) \int_{\underline{\omega}_t(k)}^{\overline{\omega}_t(k)} dF_{\eta_d}\left(\eta_{d\,t}(k)\right) \right], \tag{17}$$

where $\boldsymbol{\iota}(A)$ is an indicator function that returns 1 if the event A occurs and 0 otherwise and $F_{\eta_d}(x)$ is the cumulative density function (cdf) of $\eta_{d\,t}(k)$ evaluated at the value x. The probability of reform is

$$\boldsymbol{\iota}\{\beta_{t-1} \geq 1/\lambda\} + \boldsymbol{\iota}\{\beta_{t-1} < 1/\lambda\} \sum_{k=1}^{h} \left[\Pr\left(s_t = k \mid \Pi_{t-1}, \phi\right) \int_{\overline{\omega}_t(k)}^{\infty} dF_{\eta_d}\left(\eta_{d\,t}(k)\right) \right]. \tag{18}$$

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). The ingenious framework of Marcet and Nicolini (2003), however, enables us to estimate the structural parameters through the inflation likelihood derived in appendix II.5. On the other hand, we may ask too much of the model to pin down all the parameters. 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.[14] The value of δ implies that monthly inflation rates are bounded by $10,000\%$.[15]

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.[16] 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 the idea of Gibbs sampling and EM algorithm, breaks the set of model parameters into a few subsets and uses Sims's csminwel program to maximize the likelihood of one set of the model's parameters conditional on the other sets.[17] 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. This optimization process applied to only one starting point. We begin with a grid of 300 starting points; after convergence, we perturbs 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.[18] The other converged points typically have much lower likelihood values by at least a magnitude of hundreds in log value.

V.2. **Robustness analysis.** In addition to the specifications described in Section II, we have studied a number of alternative specifications. One could in principle let ϵ or β_t depend on regimes s_t. While this alternative creates no difficulty in analyzing

[14]One 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. For the other structural parameters, however, it is difficult or impossible to have a 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.

[15]Marcet and Nicolini (2003) set the bound at $5,000\%$.

[16]Such a problem is prevalent in the Bayesian estimation.

[17]The csminwel program can be found on http://sims.princeton.edu/yftp/optimize/.

[18]For each country, the whole optimization process is completed in 5-10 days on a cluster of 14 dual-processors, using the parallel and grid computing package called STAMPEDE provided to us by the Computing College of Georgia Institute of Technology.

the theoretical model, it is infeasible to compute the likelihood function because the unobservable variable β_t depends on a long 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. For most countries we estimated, the SCEs are close to the low steady state REEs.

For other alternative specifications, we have let π_t^* and d_t be serially correlated with their parameters to be freely estimated and allowed σ_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). None of these alternatives has improved the fit of our model.

VI. Findings

Tables 4-8 record maximum likelihood estimates of our models for Argentina, Bolivia, Brazil, Chile, and Peru, respectively, along with the estimated standard errors.[19] We use figures 3-7 to breath life into our maximum likelihood estimates. Each of these figures consists of five panels aligned to reveal features of our estimated models. 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, low, medium, and high, and two curves when it can take only two values. We regard these curves as 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 on the right compares bars that are seigniorage rates constructed from annual data with the $.16, .5, .84$ quantiles for d_t. See Appendix D for the details of how these numbers are computed. The dashed lines in the graph

[19]Following 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 $\log p(\Pi_T|\theta)$ evaluated at the MLEs. The estimated values of $\bar{d}$ are re-normalized to be consistent with what is reported in Figures 3-7.

contains the two-third probability of simulated annual deficits from our model; the solid line labelled "Model" represents the median of simulated annual deficits.[20] For all countries, it is encouraging that the deficits constructed from actual data seem to follow broad patterns about which the quantiles constructed from $p(d^T|\pi^T, \hat{\phi})$ are informative.

The third panel from the top on the right 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" when there are two mean states (or "L & M" when there are three mean states), is the probability that the deficit is in the low mean state (or either the low or the medium mean state in the three-mean case) as a function of time. The dashed line is the probability that an escape will have occurred next period.

The bottom right 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 autoregression constructed without imposing our economic model. In the graph for Brazil (see figure 5), the predicted value for the first hyperinflation is about 3 in log value. We keep the scale no greater than 0.7 in log value in order to make the reading of the actual and predicted inflation paths more discernible.

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

VI.1. **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 .2. Comparing β_t in the top right panel and with the probabilities in the third panel down on the right shows that the probability of escape becomes large when β approaches and finally exceeds that higher fixed point in 1989 and 1990. The low probability for the low deficit state in the third panel on the right shows that after 1975, Argentina lived with a chronically high mean deficit. The high deficit conditional dynamics indicate that if Argentina had been lucky enough to avoid sequences of adverse deficit shocks that drove β far enough above the stable rest point, it could avoid the kind of big inflation associated with an escape. The estimates say that it was thus lucky until the late 1980s, when the escape probability escalated and an escape occurred. The graphs attribute the 1991-1992 stabilization to Markov jump reductions in the mean and volatility of the deficit that shifted the conditional dynamics curve in a way that pushed expected

[20]We 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.

inflation rapidly downward (again see the top left panel), even though the estimated gain ε for Argentina is small.

VI.2. **Bolivia.** Our estimates indicate that the escape dynamics emphasized by Marcet and Nicolini (2003) played no role in Bolivia. The conditional dynamics in the top left panel of figure 4 assert that for both mean deficit states m, β_t never gets into the region where the unstable dynamics take over. This is confirmed by the escape probabilities in the third panel on the right. The estimates make an increase in the monetized deficit the culprit in causing the big inflation of the mid 1980s and give credit to a persistent decrease in the deficit in 1987 for reducing inflation.

VI.3. **Brazil.** Figure 5 shows estimates of a high mean deficit conditional dynamics curve that has no fixed points, which we interpret as asserting that when the deficit gets this high, an escape is likely to occur unless the country is lucky enough to have a sequence of negative shocks that push it far enough below that high conditional mean. The escape probabilities in the third panel down on the left tell a story of high and volatile deficits after 1985 and an escape probability that is volatile but high until 1994. The probabilities of the deficit levels in the third panel together with the two top panels confirms that 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 leave the mean deficit unaltered. Our model says that the 1994 stabilization is different, and was accompanied by a persistent reduction in the means and volatility of the monetized deficit. But before 1994, our estimates make Brazil a poster child for the mechanism featured in the Marcet and Nicolini (2003) model.

VI.4. **Chile.** Figure 6 tells the story that the big Chilean inflation of the mid 1970s was caused by a sustained run of high deficits (see the second and third panels on the right) that made escapes very likely (the conditional dynamics for the high mean deficit in the top left panel has essentially one rest point). The stabilization in the late 1970s is interpreted as a reduction of variance of shocks to deficits. The probabilities of high deficit, however, continue to put high weight on the high deficit until around 1981. The probabilities attached to lower mean monetized deficits rise after then, but remain volatile until the mid 1990s when they seem stay close to one. Fiscal reforms play some role in bringing down hyperinflation in the 1970s, but deficit shocks are the driving force in the conquest of Chilean inflation. Cosmetic reforms play little role for Chile, as their probabilities remain very low even during the runaway inflation period.

VI.5. **Peru.** Figure 7 interprets the big inflation after 1989 in terms of β, monetized deficit pairs that had traversed into the region that prompted a large and rapid escape and then a cosmetic reform of the Marcet and Nicolini (2003) type. This reform resets expectations and consequently the belief jumps down. Not until 1995 there is a fiscal reform as a jump in the probability assigned to the low or medium monetized deficit state.

VI.6. **Types of inflations and stabilizations.** Table 3 summarizes the key empirical findings displayed in Figures 3-7. The first column lists the three possible ways through which our model tells us hyperinflation can be arrested: a Marcet and Nicolini (2003) type of 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. The countries put in the appropriate boxes in the table are selected according to how 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 its monthly gross rate being 1.82. In the next two months, the inflation rate dropped to 1.15 and 1.07. The probability of cosmetic monetary reform is 67.5% for March, 75.6% for April, and 47.8% for May. Peru is another informative example. In August of 1990, 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 and jumps up to 100% in September. This reform resets the expectations and brings down the belief instantly.[21] 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).

[21]Without 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.

VI.7. **Parameter patterns.** Table 4-8 report the ML estimates of the structural parameters for all the five countries, along with the estimated standard errors.[22] The estimated values of $\bar{d}$ are re-normalized to be consistent with what is reported in Figures 3-7. As one can see, all the parameters 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. This phenomenon is also shown in Figure 9.

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*, though a low elasticity of the demand for money with respect to expected inflation, while it also discounted past rates of inflation the *most* though 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 to 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 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.

VI.8. **Comparison with Marcet and Nicolini (2003).** To illustrate how our model differs from theirs, for Argentina, we have formed the conditional density inflation conditioned by the history of inflation that is implied by the model of Marcet and Nicolini (2003) at their calibrated parameter values. Figure 8 reports the one-step forecasts and 90% probability distributional bands around them, together with actual inflation outcomes. This figure should be compared with the bottom panel of figure 3.

[22]Following 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 $\log p(\Pi_T|\theta)$ evaluated at the MLEs.

For Marcet and Nicolini's calibrated parameter values, we have computed 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 expected inflation. 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.

VI.9. **Ergodic distributions for deficits.** Figure 9 plots ergodic probabilities of the estimated average deficit level and the estimated standard deviation of deficit shocks. These probabilities are consistent with the computed SCEs. Clearly, Brazil and Chile have deficit processes that are conducive to persistent low inflation, while Argentina and Peru do not.

VI.10. **Alternative specifications.** 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, needless to say about comparing our model to a time-varying AR model. In the previous work, only certain moments or correlations were typically reported. 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 (see Appendix **??** for details).

For each country we have tried a large number of versions of our model, including 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 the 3×2 model. By the Schwarz criterion (SC) or Bayesian information criterion,[23] the 2×3 version of the model fits best for Argentina, Bolivia, and Chile and 3×2 version is the best for Brazil and Peru; all other versions including the constant-parameter case fit much 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.

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

VII. The Model's Fit

We have tried more than two dozen versions of our theoretical model and of the unrestricted atheoretical model. Within each model, the fit in most versions is substantially worse than our best-fit model and thus we do not report their results.

Table 9 reports the best-fit theoretical model for each country, compared with the constant-parameter theoretical model that has been commonly used in the literature and with the best-fit unrestricted regime-switching AR model.[24] For all the five countries, the best-fit 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-fit theoretical model is used as a baseline for comparison. The notation "df" stands for degrees of freedom in relation to the baseline model.

Figures 10-18 compare the log conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi})$ of our theoretical model with that of the best-fit 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.

VIII. Conclusion

Building on Sargent and Wallace (1987) and Marcet and Nicolini (2003), we develop a nonlinear general equilibrium model of hyperinflation. This model is fit to the data in Argentina, Bolivia, Brazil, Chile, and Peru. Unlike the previous literature, the time-series properties of this model are rigorously checked against the data in all these five countries. Our estimated model provides important insights that have not been explored in the existing literature. Our robust results show that a large amount of seigniorage is necessary for hyperinflation to occur repeatedly but the inflation dynamics depends crucially on a combination of many factors such as beliefs and fundamental shocks. On the other hand, fiscal reform in keeping the amount of seigniorage low is necessary for preventing the reoccurrence of high inflation. In other words, low inflation can be achieved by disciplinary fiscal policy and will be sustained if such a policy is to be maintained.

[24]Regime-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).

Appendix A. Deriving the Likelihood

We first derive a likelihood conditional on the hidden states s_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 $\eta_{d\,t}(k)$ is log-normal for $k = 1, \ldots, h$. Specifically, the probability density functions are

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

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

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

$$\begin{aligned} S_t &= \{s_0, s_1, \ldots, s_t\}, \\ \Pi_t &= \{\pi_{-1}, \pi_0, ..., \pi_t\}, \\ q &= \{q_{ij}\}\ \forall i, j = 1, \ldots, h, \\ \xi_d(s_t) &= 1/\sigma_d(s_t), \\ \xi_\pi &= 1/\sigma_\pi, \end{aligned}$$

and let ϕ be a collection of all structural parameters. We use the tilde above $\eta_{d\,t}(s_t)$ to indicate that $\tilde{\eta}_{d\,t}(s_t)$ is a *random* variable, whereas $\eta_{d\,t}(s_t)$ 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

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

where

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

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

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

where Pr[ ] 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

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

It follows that

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

□

After integrating out S_T, the overall likelihood is

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

where

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

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

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

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 section describes self-confirming equilibrium versions of our model, and section C 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 (**??**), we compute such equilibria for values of parameters that we estimated by maximizing the likelihood function of section II.

B.1. **Small gain convergence.** If the agent in our model were to implement a least squares estimator by replacing γ 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 agent under our model.

B.1.1. *Self-confirming equilibria.* Noah and Tom: I have rewritten this section and made all the notational corrections to be consistent with the previous sections. Questions: Does the following analysis take care of the cases where $\beta_{t-1} \leq 0$ or $\beta_t \leq 0$? What about the truncated log-normal distribution for π_t^*? Please check the notation, making sure it is consistent with the rest of the text.

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

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

where π_t is 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:

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

Hence, we can write (6) as:

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

[25] We have assumed that agents do not know the current regime s_t when forecasting inflation. Better informed agents would incorporate knowledge of s_t in forecasting inflation (see section **??**).

where

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

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

$$\tilde{g}(\beta, d_t, \pi_t^*) = g(\beta, \beta, d_t, \pi_t^*),$$
$$\xi(\beta) = (1-\beta\lambda),$$
$$\Psi_k(\beta, b) = \int_0^b \frac{1}{\xi(\beta)-d} d\, F_d(d(k)).$$

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

Since $\log d(k) \sim N(\log \bar{d}(k), \sigma_d^2(k))$ and $\log \pi_t^*(k) \sim N(\log \pi_1^*(k), \sigma_\pi^2)$, we have

$$\begin{aligned}
\bar{g}(\beta) &= E[\tilde{g}(\beta, d_t, \pi_t^*)] \\
&= \sum_{k=1}^h \left[\int_0^{(1-\delta\theta)\xi(\beta)} \frac{\theta(1-\lambda\beta)}{1-\lambda\beta-d} d\, F_d(d(k)) + \left[1 - F_d\left((1-\delta\theta)\xi(\beta)\right)\right] E(\pi^*(k))\right] \bar{q}_k - \beta \\
&= \sum_{k=1}^h \left\{\theta\xi(\beta)\Psi_k(\beta, (1-\delta\theta)\xi(\beta)) + \left[1 - \Phi\left(\frac{\log((1-\delta\theta)\xi(\beta)) - \log d(k)}{\sigma_d(k)}\right)\right] \pi^*(k) e^{\frac{\sigma_\pi^2}{2}}\right\} \bar{q}_k -
\end{aligned}$$

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

$$\dot{\beta} = \bar{g}(\beta) \tag{A9}$$

for $\delta > 0$ and a broad class of probability distributions of $\eta_{d\,t}(s_t)$ 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). □

The ODE (A9) governs the mean dynamics $\bar{g}$. For such an SCE to be a small-gain limit of our learning rule, β must be a stable equilibrium point of the ODE (A9). We don't have an explicit expression for $\bar{g}$, so we shall find a SCE numerically. Thus we look for stationary points $\bar{\beta}$ such that $\bar{g}(\bar{\beta}) = 0$. Since the system is scalar the stability condition is simply $\bar{g}'(\bar{\beta}) < 0$.

B.2. **Conditional SCE's.** For comparison, we are also interested in the SCE equilibria that would result if the economy were forever to remain in one regime. We refer to these as "conditional SCE equilibria." We need to insert Noah's new math here. Perhaps, the notion of stochastic ODE needs be introduced here.

They fixed points are stable points of the conditional mean dynamics $\bar{g}(\beta, j)$ that are defined implicitly above Noah: I think we should be more explicit here. and satisfy $\bar{g}\beta) = \sum_j \bar{g}(\beta, j)\bar{q}_j$. The stable conditional SCEs give the small-gain limit points if the regimes were fixed, while the stable SCE averages over the conditional equilibria.

Below in section (**??**), we shall 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.

Appendix C. Rational Expectations Equilibria

We now suspend (6) and consider a subset of the rational expectations equilibria of the model. We seek stationary Markov equilibria in which inflation and expected inflation are given by:

$$\begin{aligned} \pi_t &= \pi(s_t, s_{t-1}, d_t) \\ E_t \pi_{t+1} &= E_t \pi(s_{t+1}, s_t, d(s_{t+1}) + \eta_{d,t+1}(s_{t+1})) \\ &= \sum_{j=1}^{h} \int \pi(s_j, s_t, d(s_j) + \eta) dF(\eta|j) q_{s_t,j} \equiv \pi^e(s_t). \end{aligned}$$

Then going through calculations similar to those above we have:

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

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

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

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

$$\pi^e(i) = \sum_{j=1}^{h} \left\{ \theta \xi_i \Psi_j(\pi^e(j), \omega_{ij}) + \left[1 - \Phi\left(\frac{\log(\omega_{ij}) - \log d(j)}{\sigma_d(j)} \right) \right] \pi^*(j) e^{\frac{\sigma_\pi^2}{2}} \right\} q_{ij}, \tag{A10}$$

where $\xi_i = (1 - \pi^e(i)\lambda)/\gamma$ and Ψ_j is as above. Thus we have h coupled equations determining $\pi^e(s_t)$. 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 (A10) and computing $\pi^e(i)$ on the left side and iterating until convergence. Alternatively, any other numerical nonlinear equation solver can be used.

C.0.1. *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 true 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 see below, 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. In one country, Peru, we find that the shocks are so large that there is no REE. For the countries where a REE does exist, we focus on finding a stationary REE equilibrium that has the lowest inflation rates. This is the REE that is closest to the stable SCE.[26]

Appendix D. Seigniorage Rates: Actual Data and Model Implications

This is a new section. 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

$$d_{A,t}^{\text{Data}} = \frac{M_{A,t}^{\text{Agg}} - M_{A,t-1}^{\text{Agg}}}{Y_{A,t}^{\text{Agg}}} \tag{A11}$$

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 (A11) as possible, we compute the distribution of $d_{A,t}$ as follows. We first draw s_t from $\Pr(s_t|\hat{\phi}, \pi^T)$ and for a given s_t we then draw $d_t(s_t)$ and compute $d_{A,t}$ as an average of $d_t(s_t)$ over the twelve months 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.

[26]Thus, we do not exhaustively search for all stationary equilibria (let alone nonstationary equilibria, sunspots, and so forth).

Table 1. Parameters and their meanings

parameter(s)	feature
λ	demand for money
ε	expectations
$\bar{d}(m_i)$	log deficit mean
$\xi(v_i)$	log deficit inverse std
ξ_π	reform inverse std
Q^m	m- transition matrix
Q^v	v-transition matrix

Table 2. Money demand and adaptation parameters

Country	λ	ε
Argentina	.73	.023
Bolivia	.307	.232
Brazil	.613	.189
Chile	.875	.025
Peru	.74	.069

Table 3. 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)

Table 4. Argentina: MLEs for the 2×3 regime-switching model

λ : 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)]		
ξ_π : 16.78 (5.178)		
ϵ : 0.023 (0.001)		
Transition probability matrix for $\bar{d}(s_t)$:		
0.9789 (0.014)		0.0162
0.0211		0.9838 (0.007)
Transition probability matrix 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 5. Bolivia: MLEs for the 2×3 regime-switching model

λ : 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)]		
ξ_π : 26.52 (2.5114)		
ϵ : 0.232 (0.0375)		
Transition probability matrix for $\bar{d}(s_t)$:		
0.9629 (0.0237)		0.0041
0.0371		0.9959 (0.0028)
Transition probability matrix 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 6. Brazil: MLEs for the 3×2 regime-switching model

λ : 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)]		
ξ_π : 9.18 (10.8305)		
ϵ : 0.189 (0.0118)		
Transition probability matrix 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 for $\eta_{d\,t}(s_t)$:		
0.9344 (0.0292)		0.0969
0.0656		0.9031 (0.0338)

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

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

λ : 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)]		
ξ_π : 10.62 (3.8807)		
ϵ : 0.025 (0.0000)		
Transition probability matrix for $\bar{d}(s_t)$:		
0.9869 (0.0051)		0.0070
0.0131		0.9930 (0.0076)
Transition probability matrix for $\eta_{d\,t}(s_t)$:		
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 8. Peru: MLEs for the 3 × 2 regime-switching model

λ : 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)]		
ξ_π : 15.97 (2.4004)		
ϵ : 0.069 (0.0025)		
Transition probability matrix for $\bar{d}(s_t)$:		
0.9943 (0.0076)	0.0187	0.0000
0.5695	0.9626 (0.0142)	0.0350
0.0000	0.0187	0.9650 (0.0166)
Transition probability matrix for $\eta_{d\,t}(s_t)$:		
0.3016 (0.1310)		0.0453
0.6984		0.9547 (0.0127)

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

TABLE 9. Log likelihood adjusted by the Schwarz criterion

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

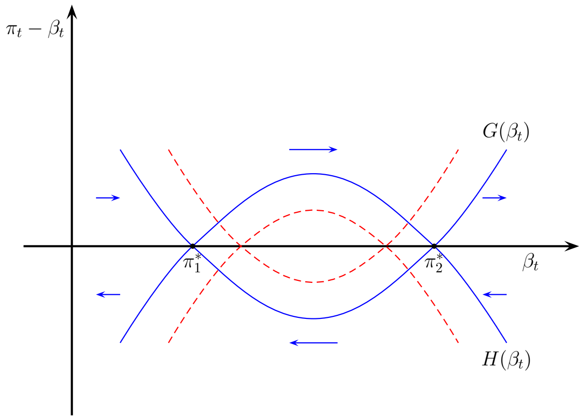

FIGURE 1. Mean adaptive dynamics and REE dynamics.

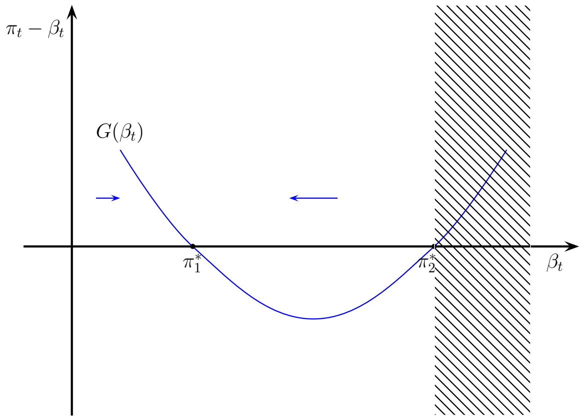

FIGURE 2. Adaptive dynamics and the 'escape event'.

— Argentina

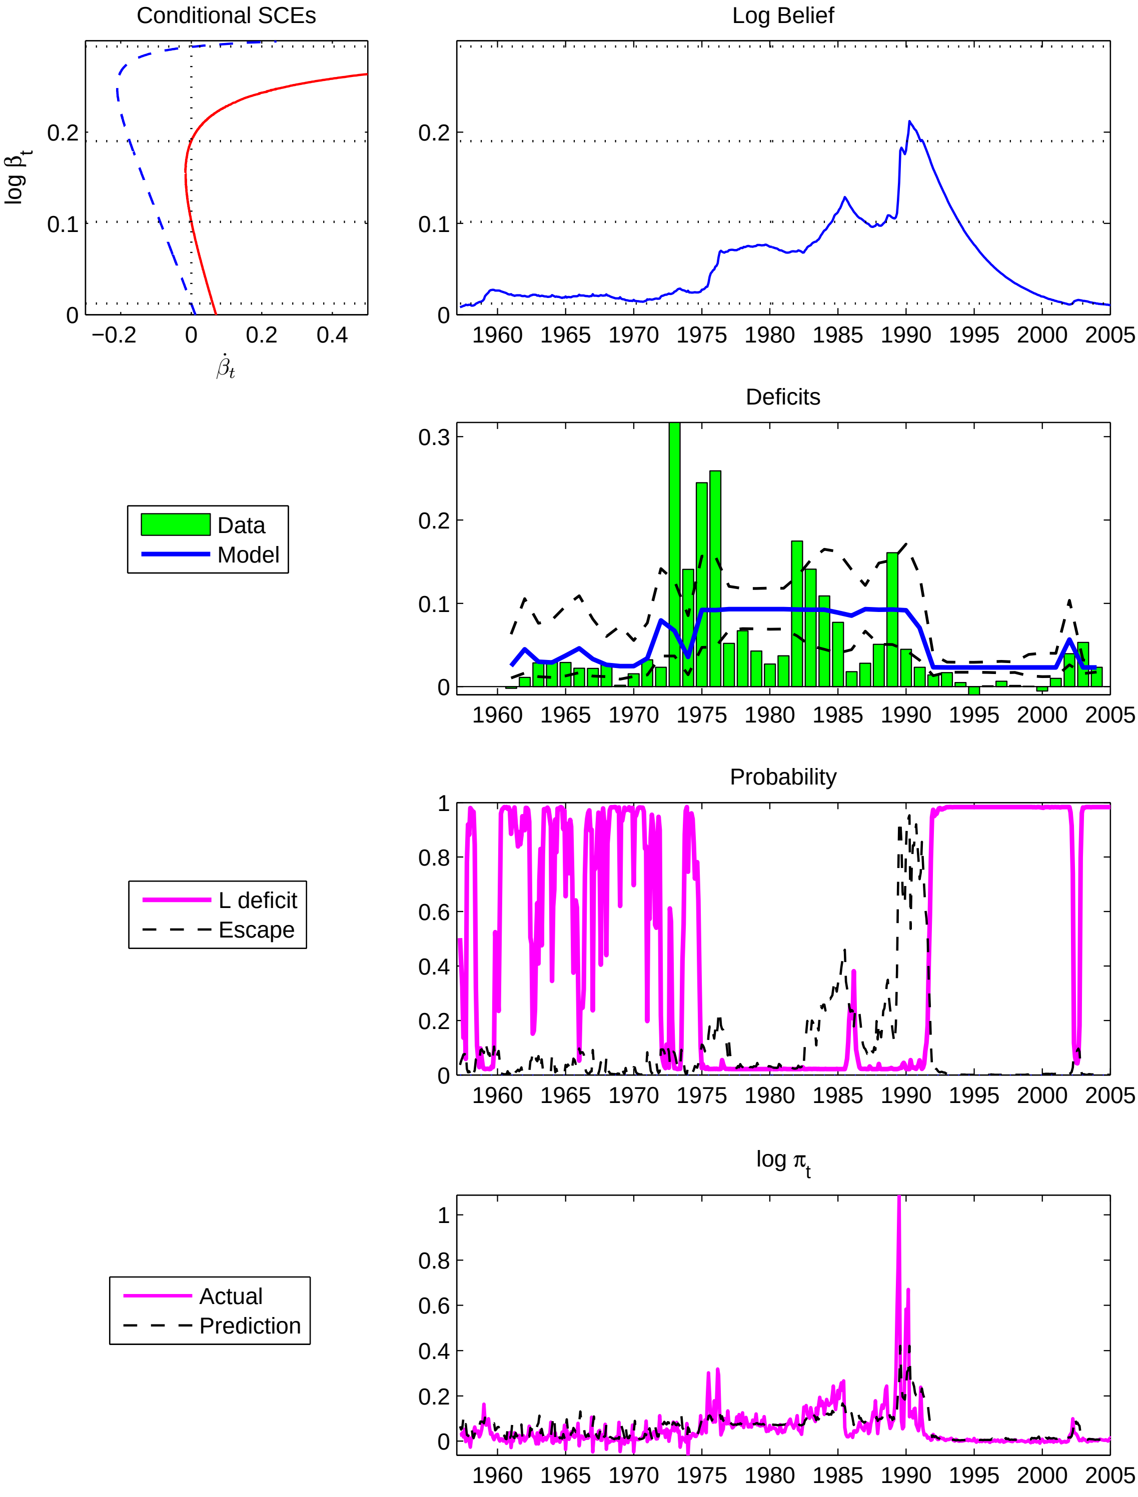

FIGURE 3. Argentina.

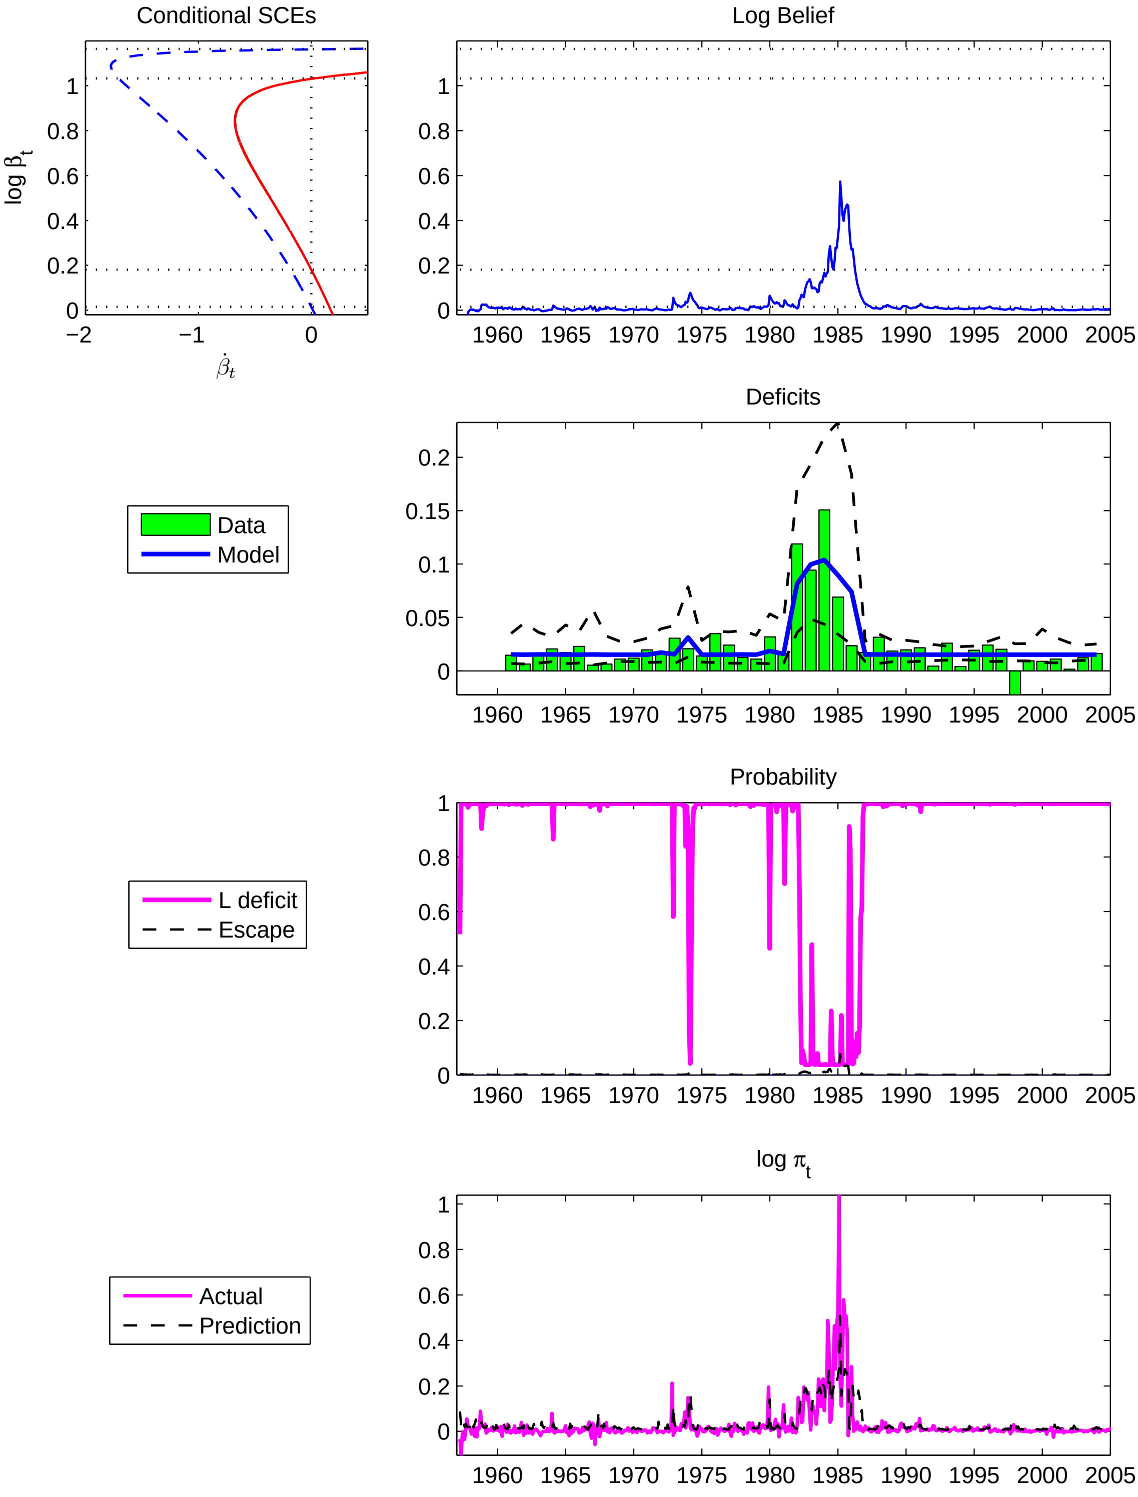

FIGURE 4. Bolivia.

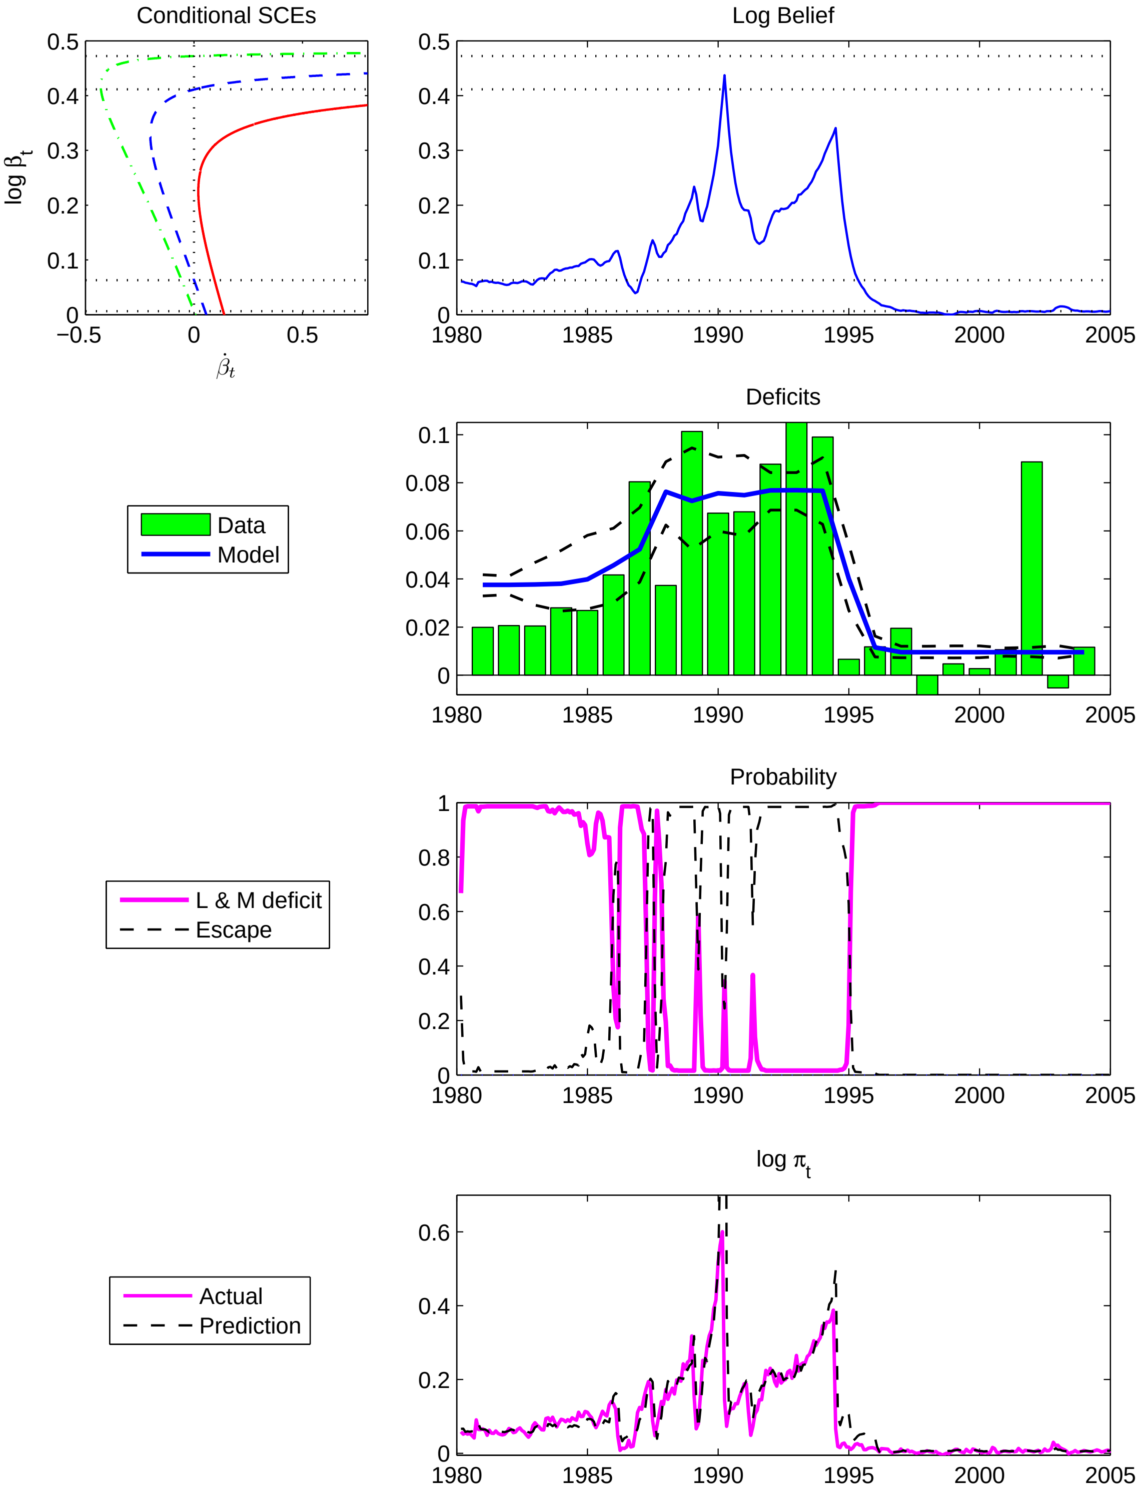

FIGURE 5. Brazil.

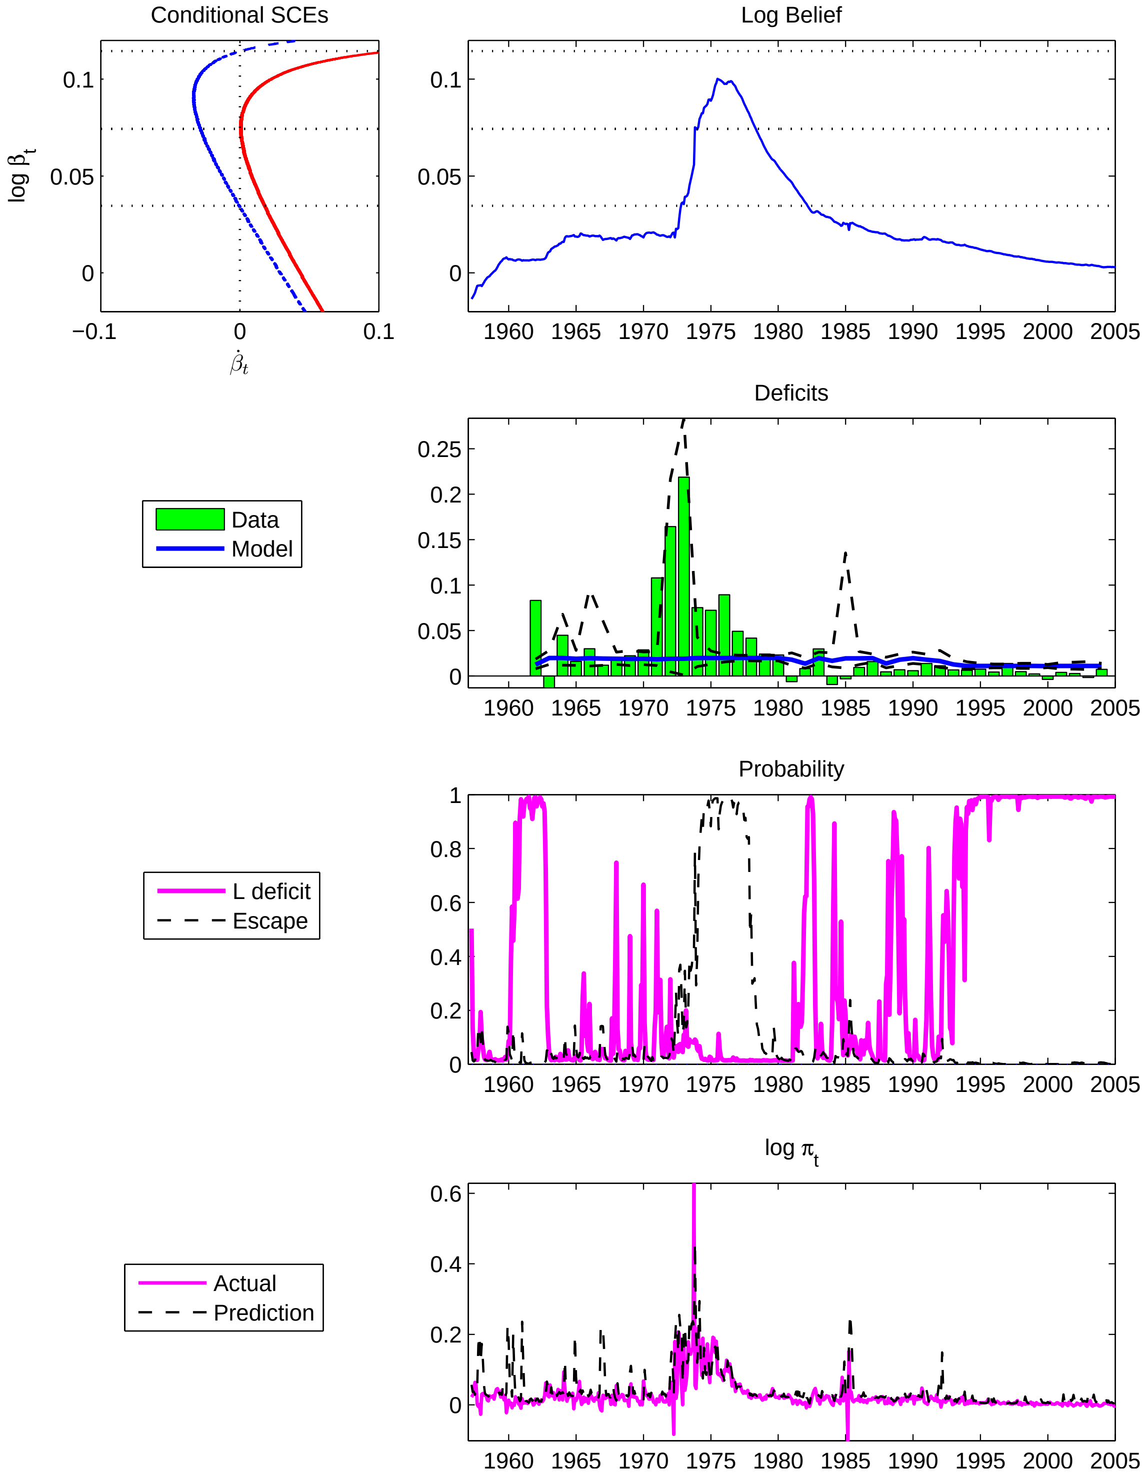

FIGURE 6. Chile.

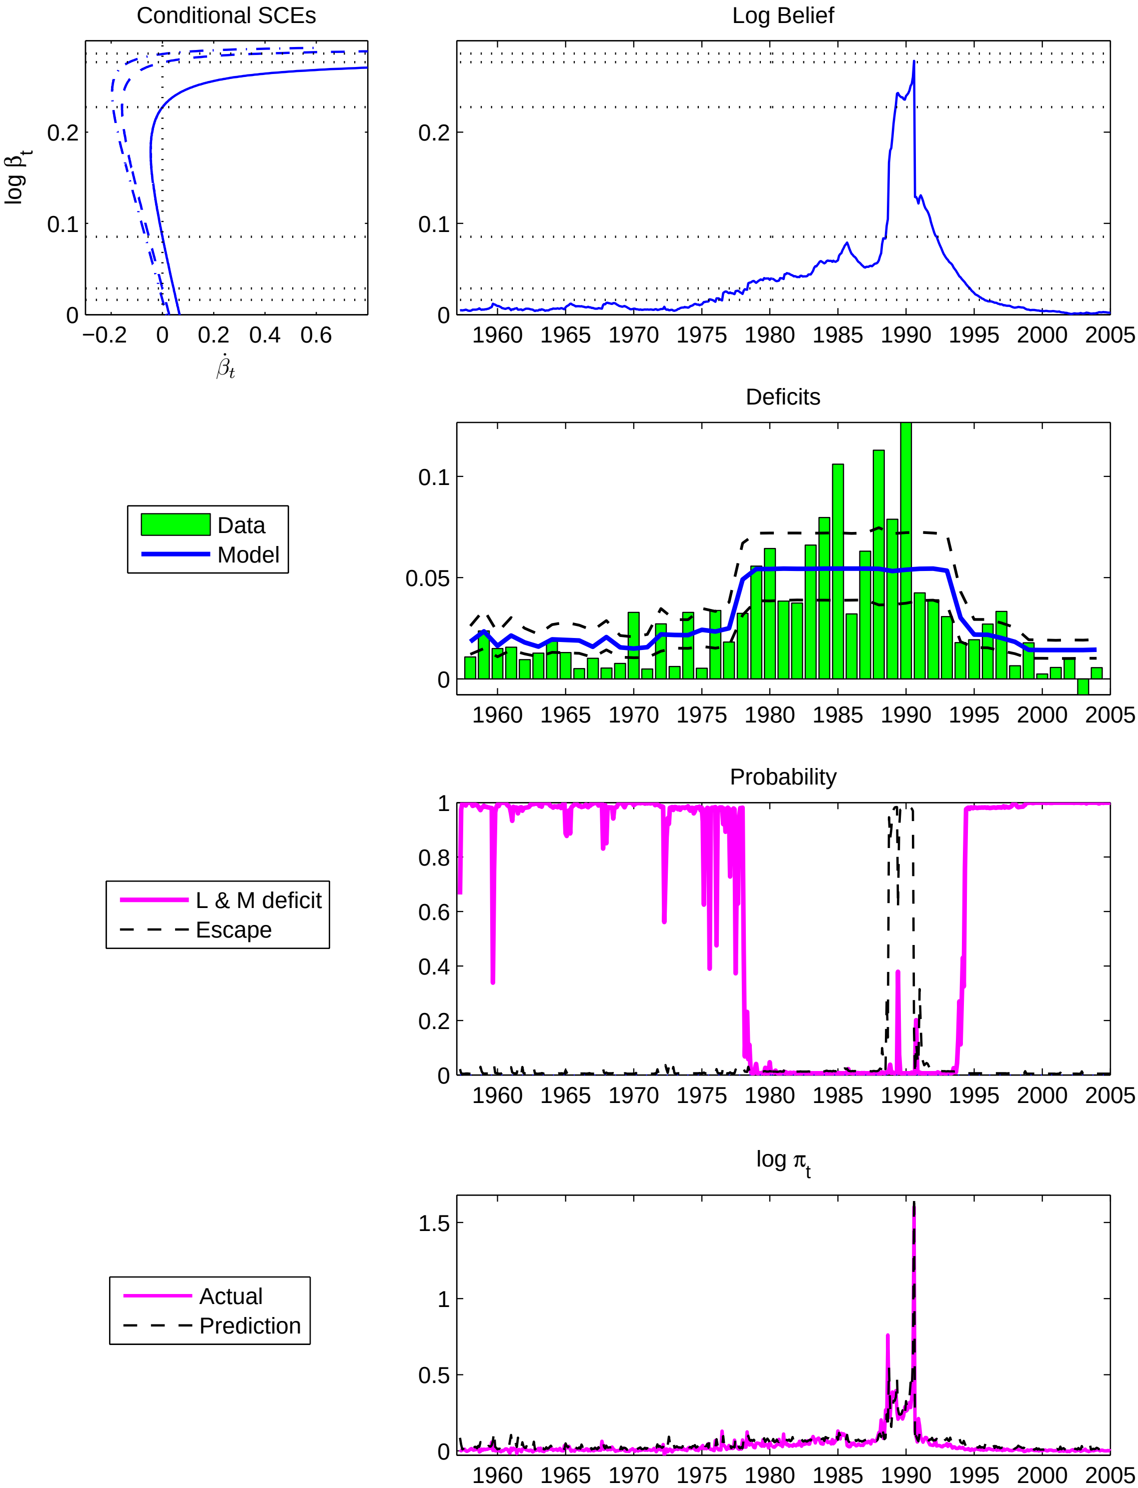

FIGURE 7. Peru.

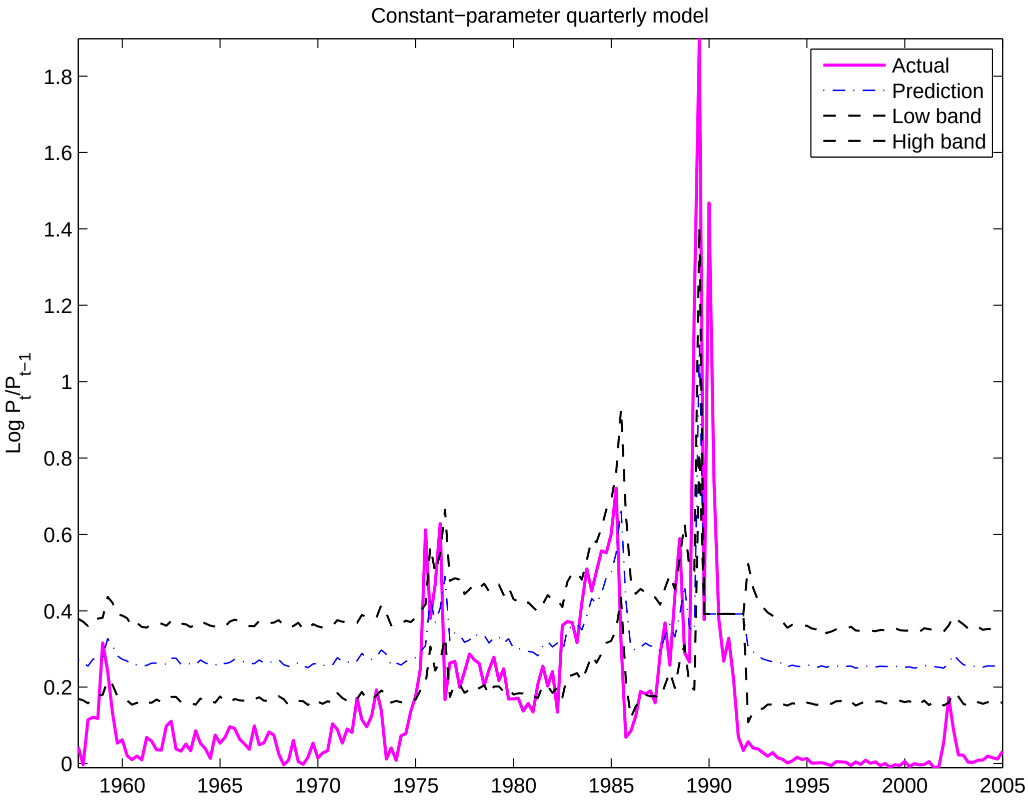

FIGURE 8. Marcet and Nicolini (2003) model's one-step prediction with 90% probability bands for Argentina.

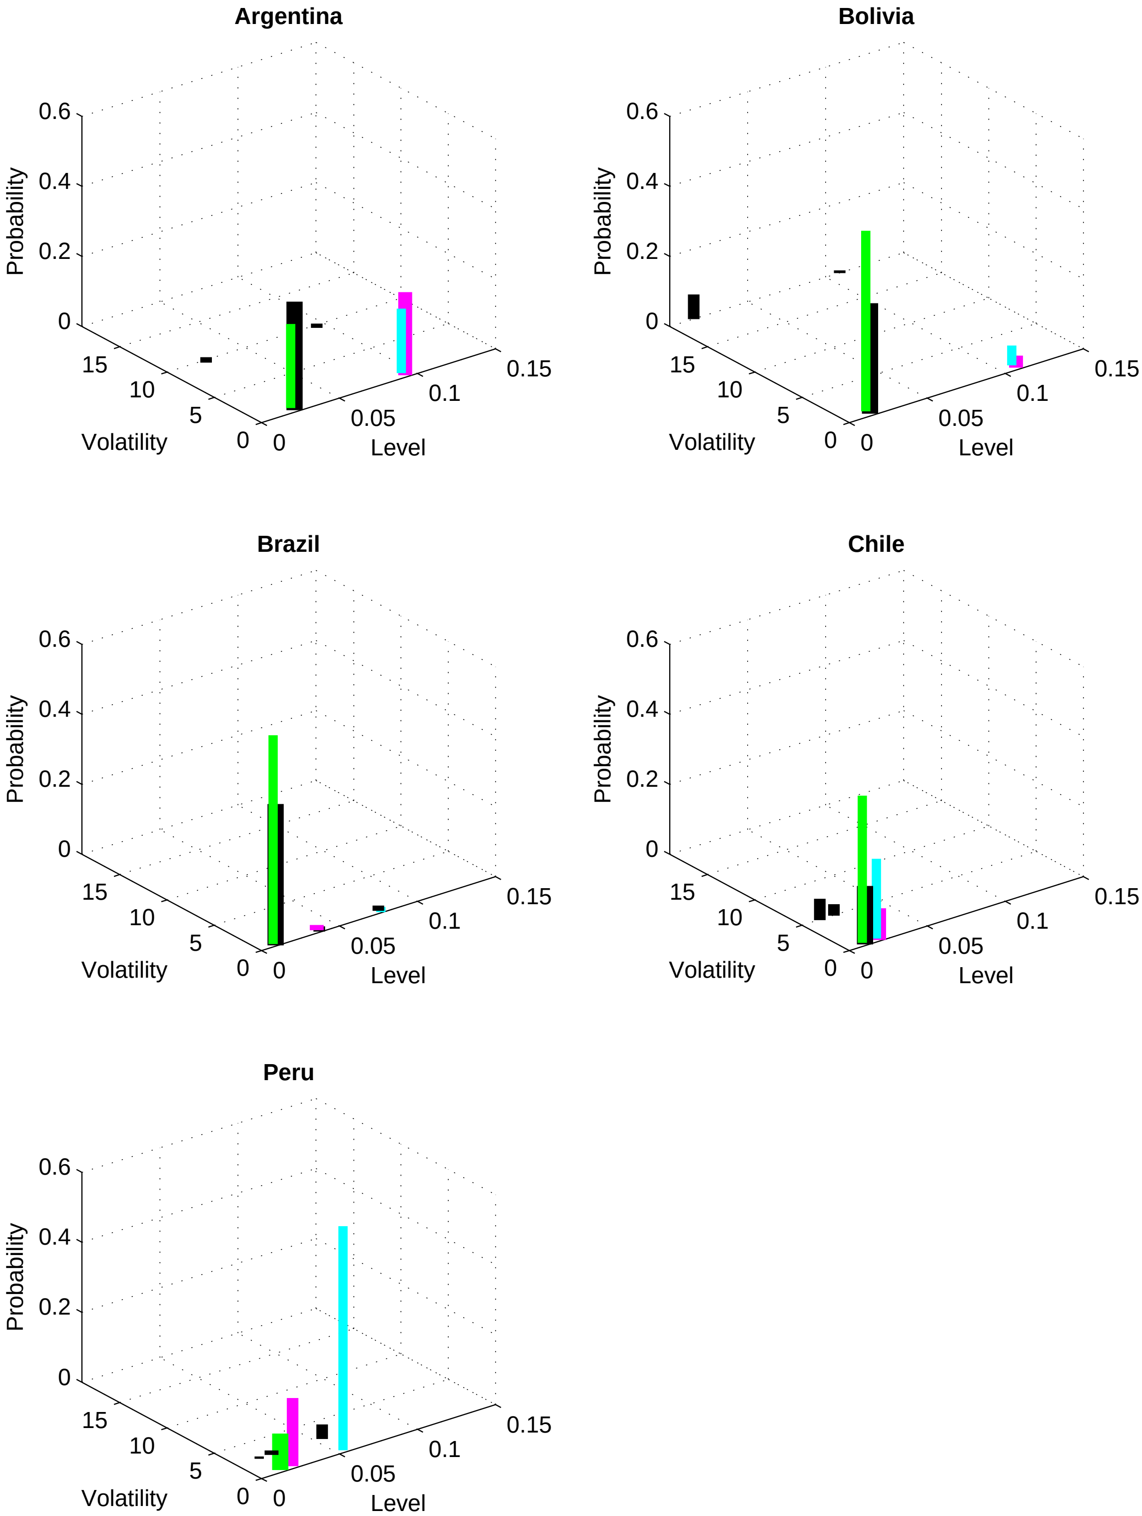

FIGURE 9. Ergodic probability given the estimated average deficit level (x-axis) and the estimated standard deviation of deficit shocks (y-axis).

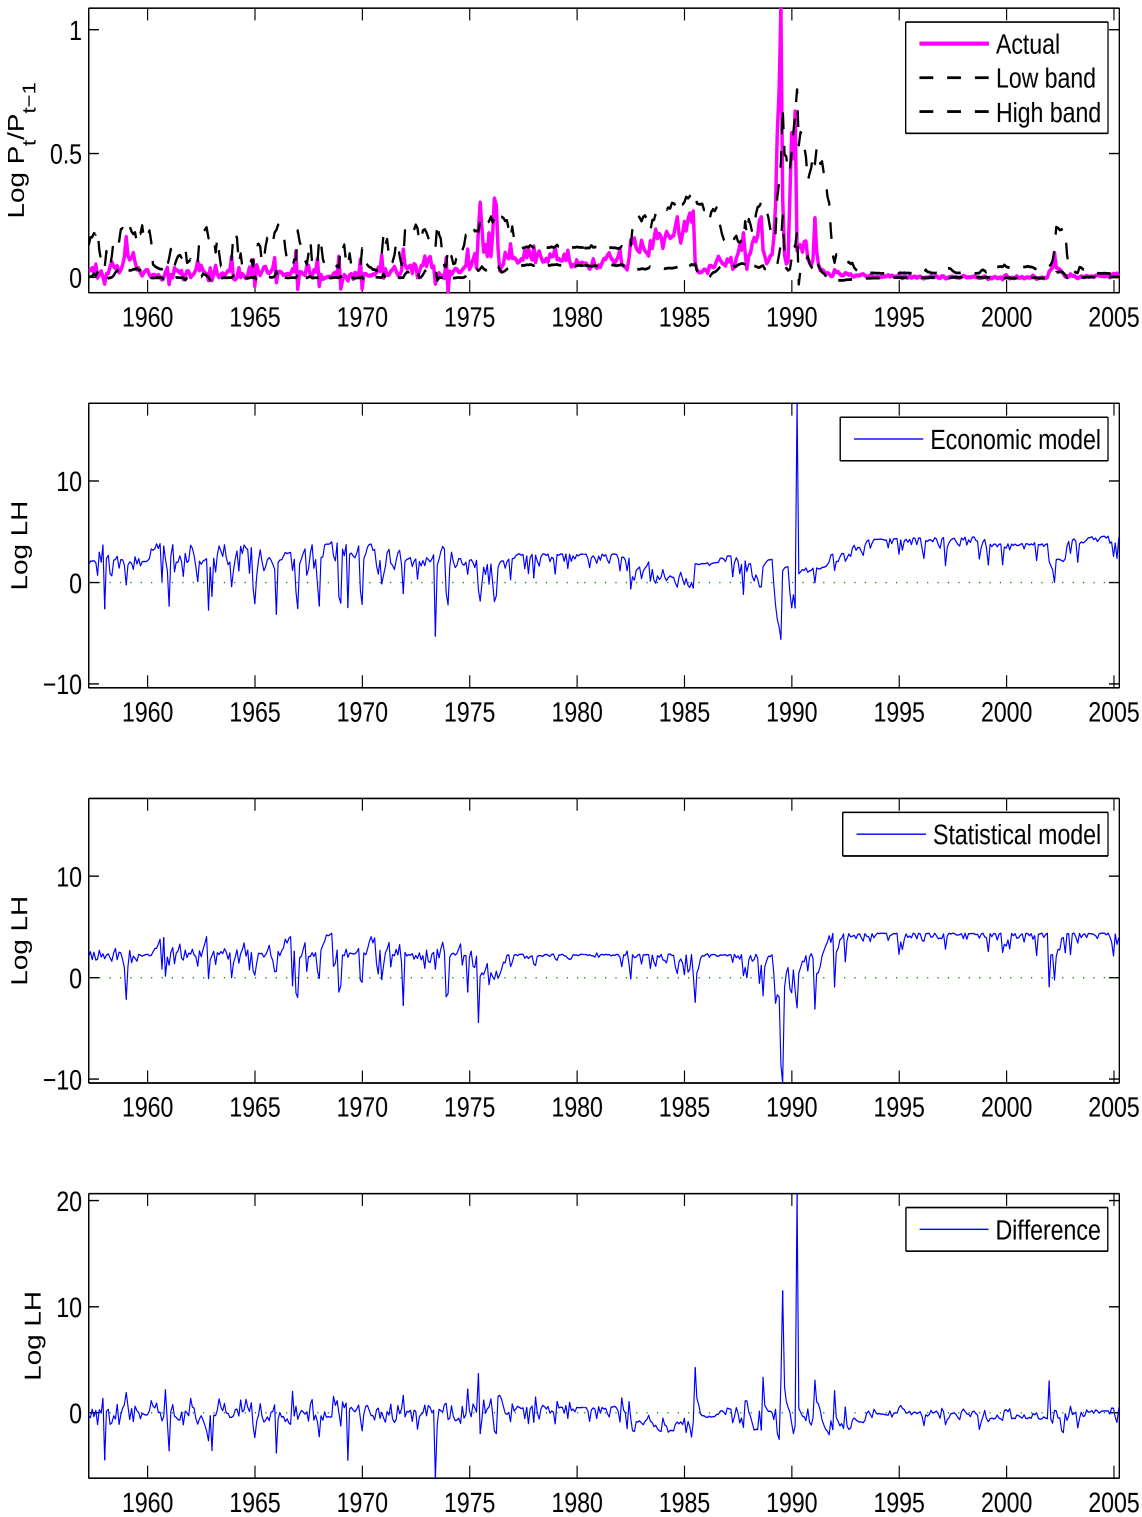

FIGURE 10. Argentina: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1},\hat{\phi}$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

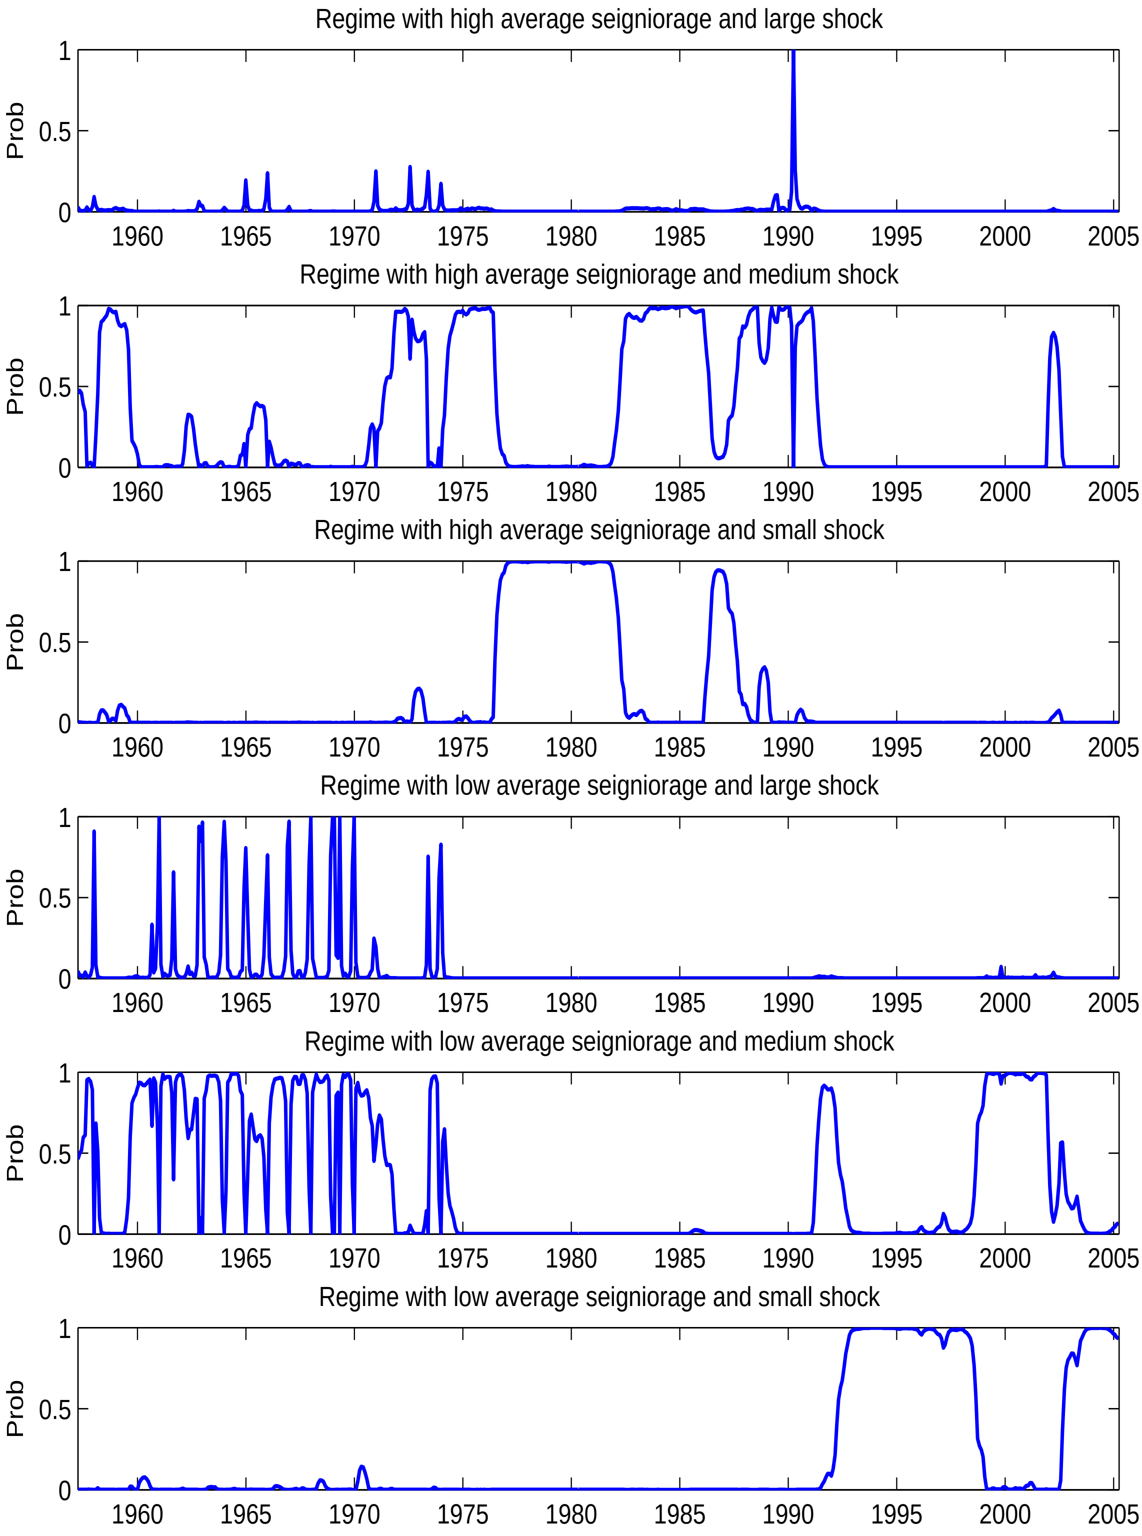

FIGURE 11. Argentina: smoothed probability of the regimes conditional on the MLEs and the data.

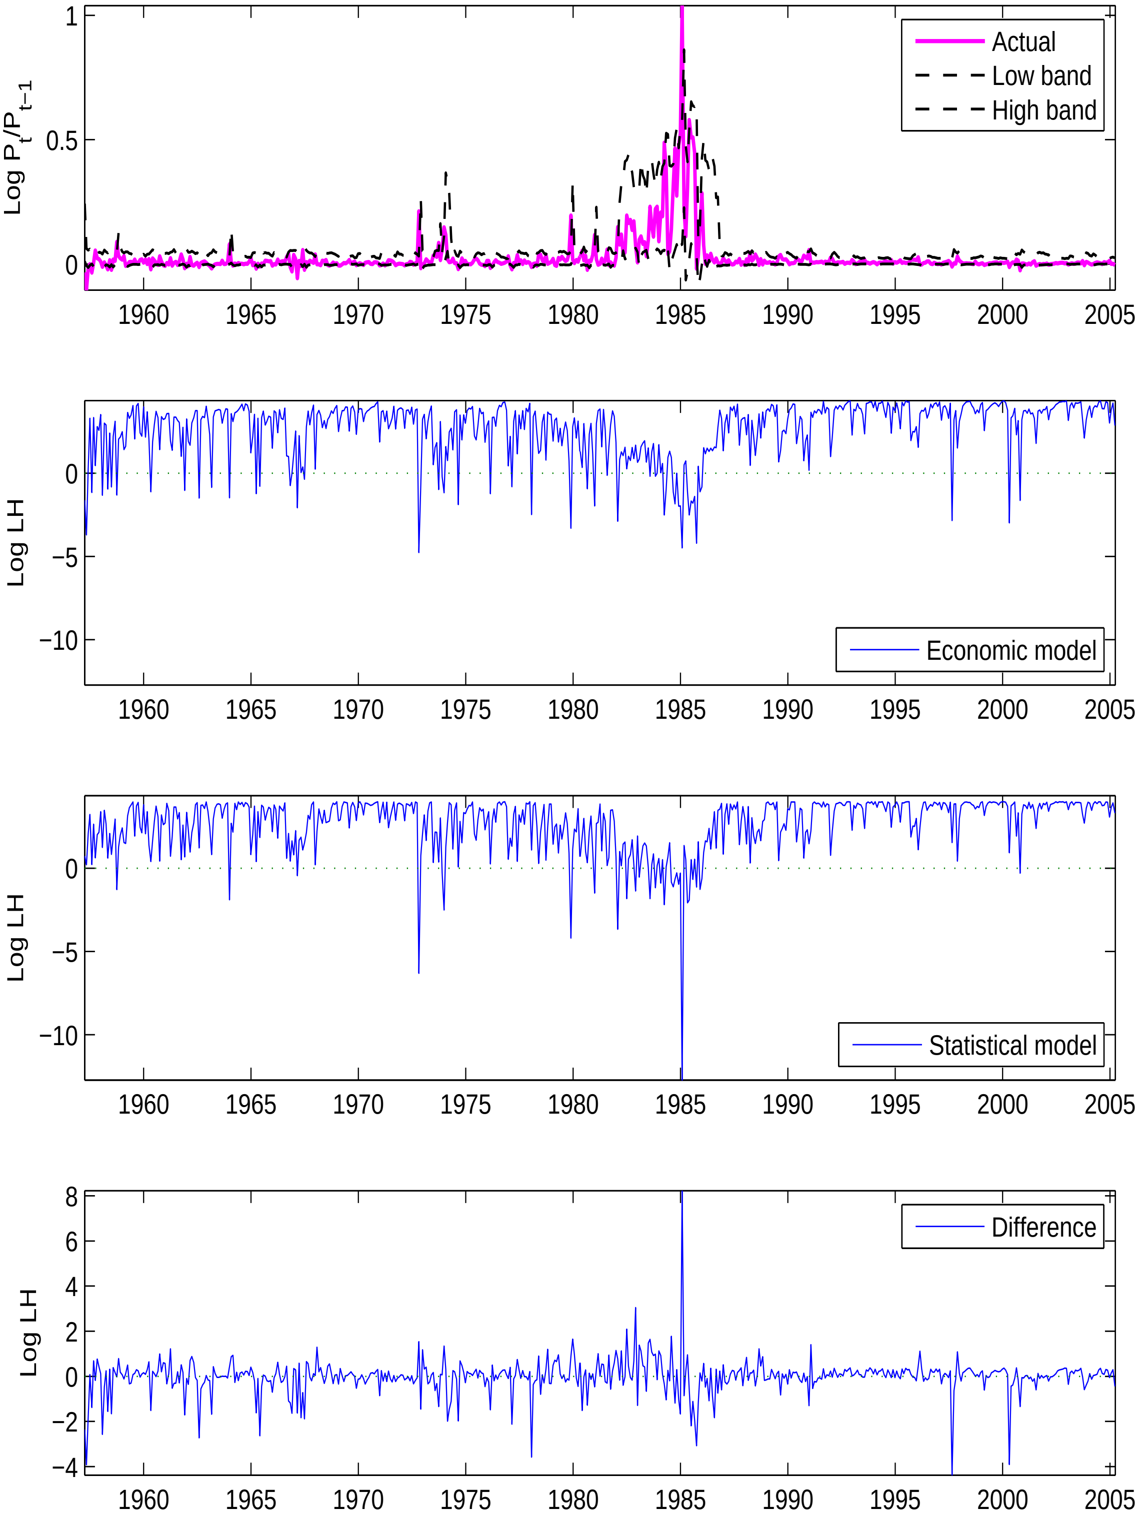

FIGURE 12. Bolivia: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1},\hat{\phi}$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

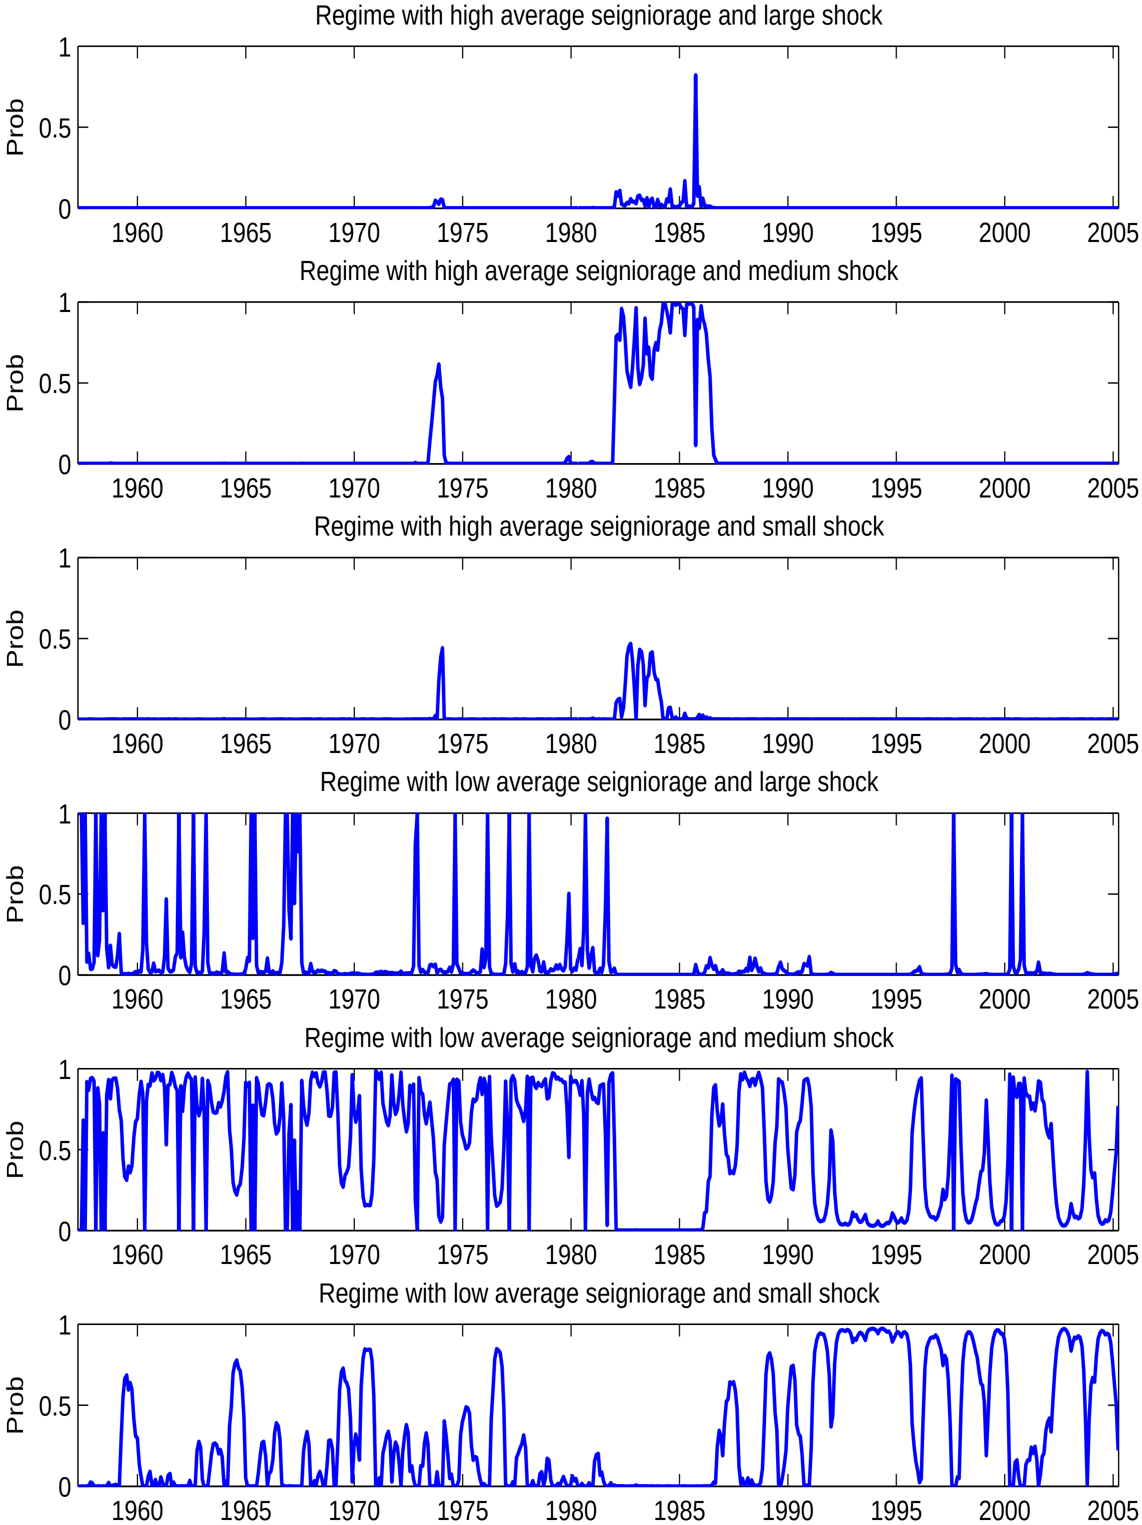

FIGURE 13. Bolivia: smoothed probabilities of the regimes conditional on the MLEs and the data.

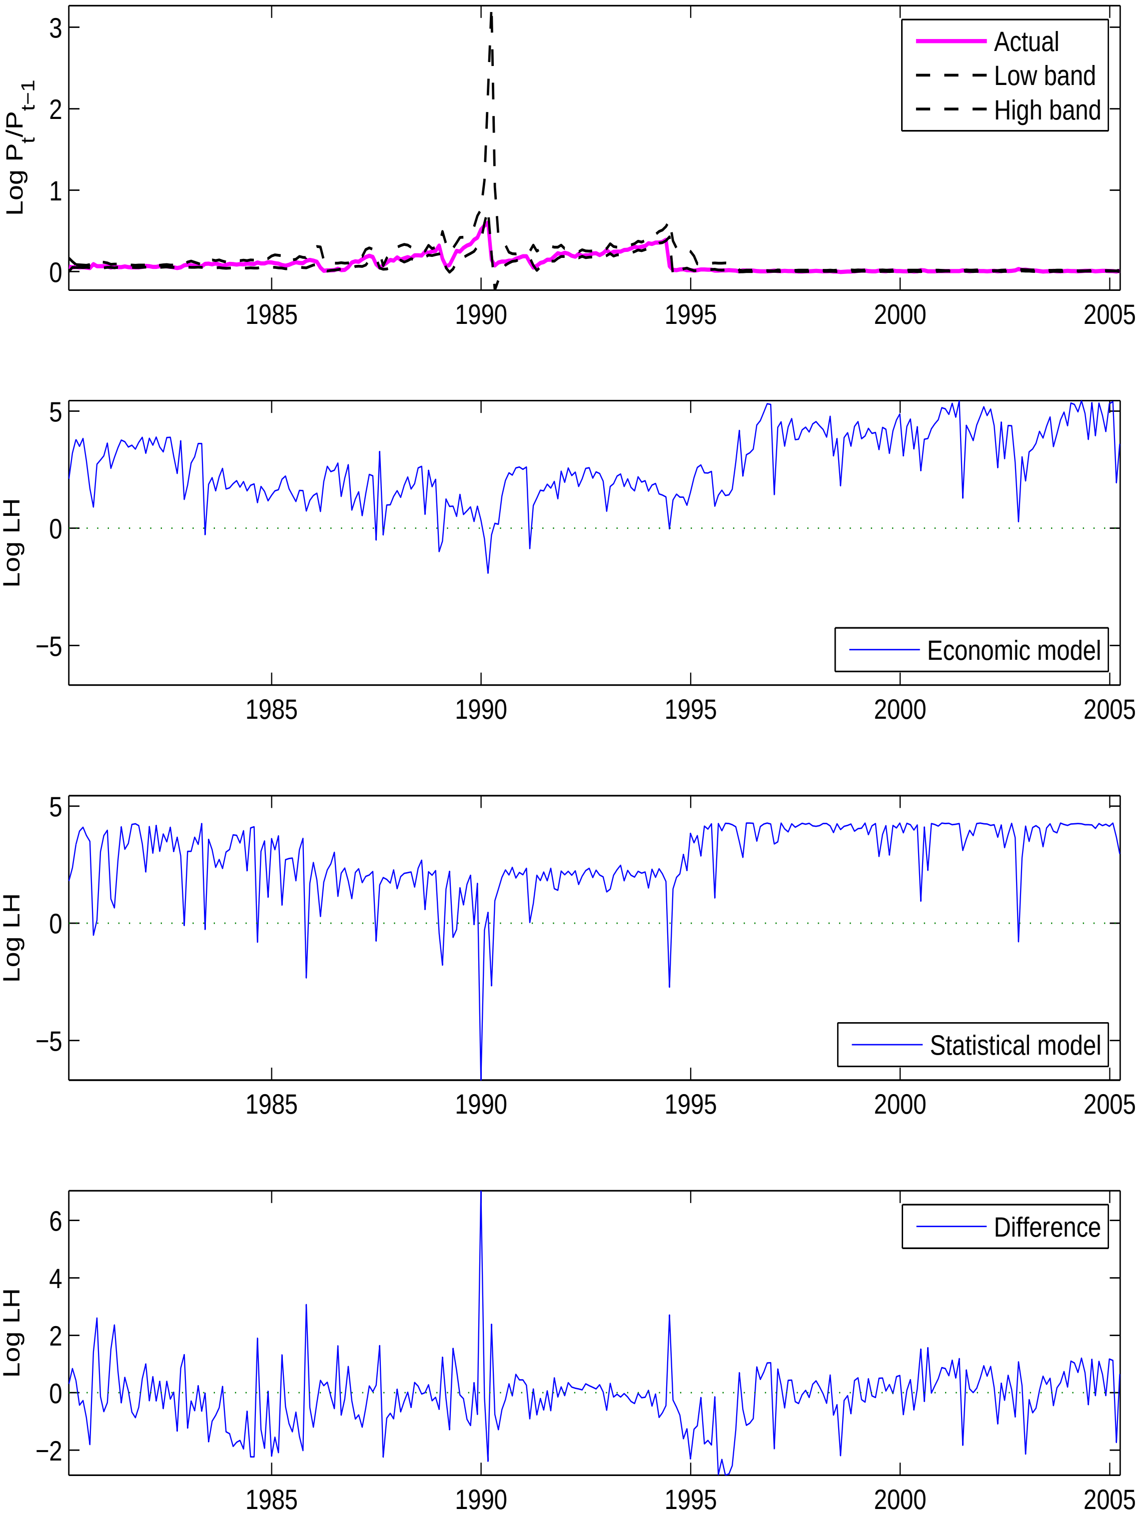

FIGURE 14. Brazil: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi}$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

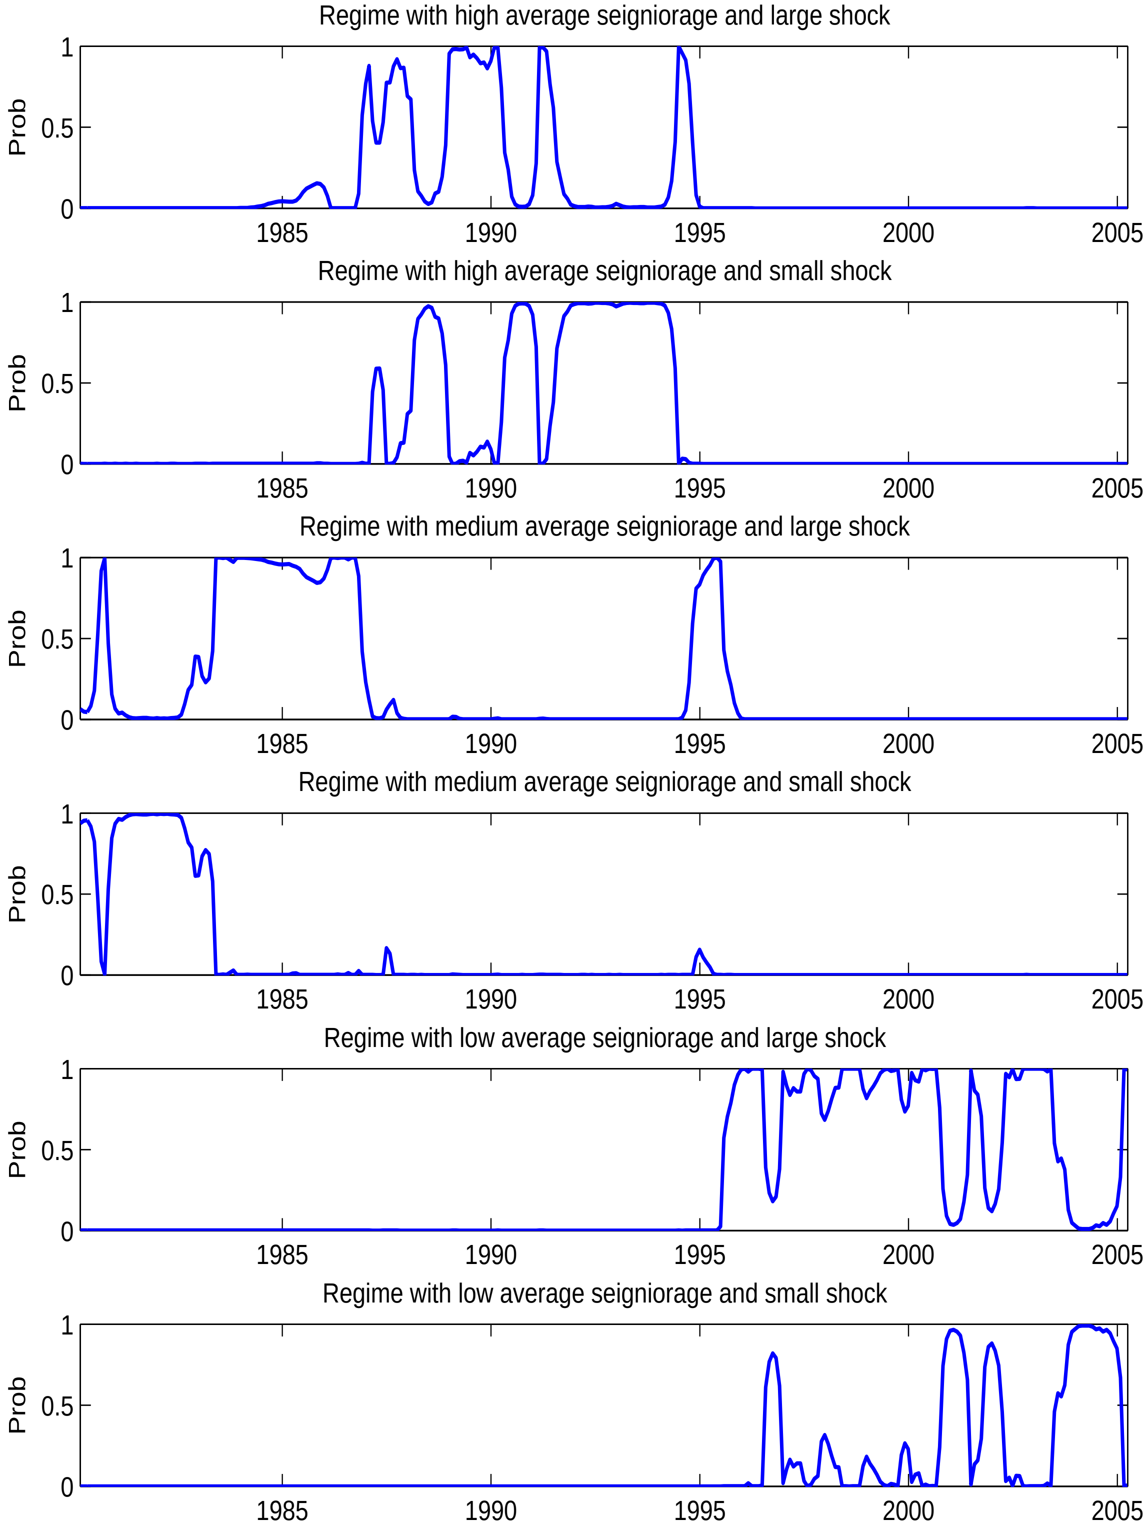

FIGURE 15. Brazil: smoothed probabilities of the regimes conditional on the MLEs and the data.

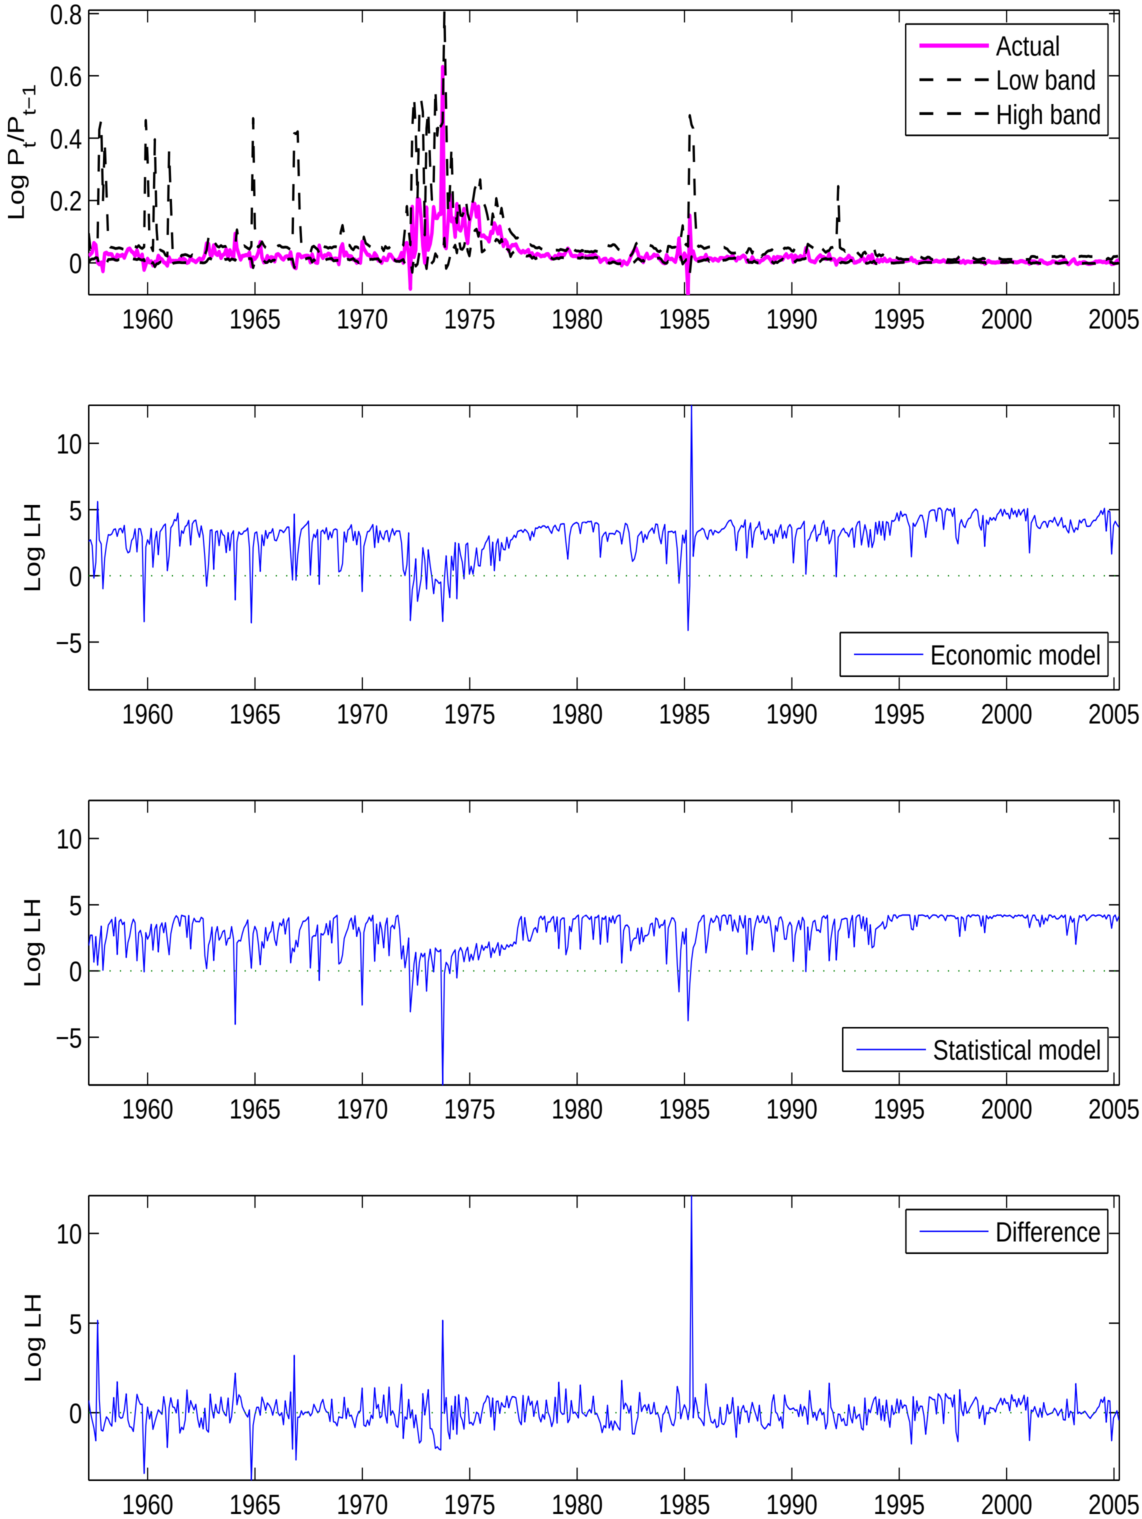

FIGURE 16. Chile: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi}$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

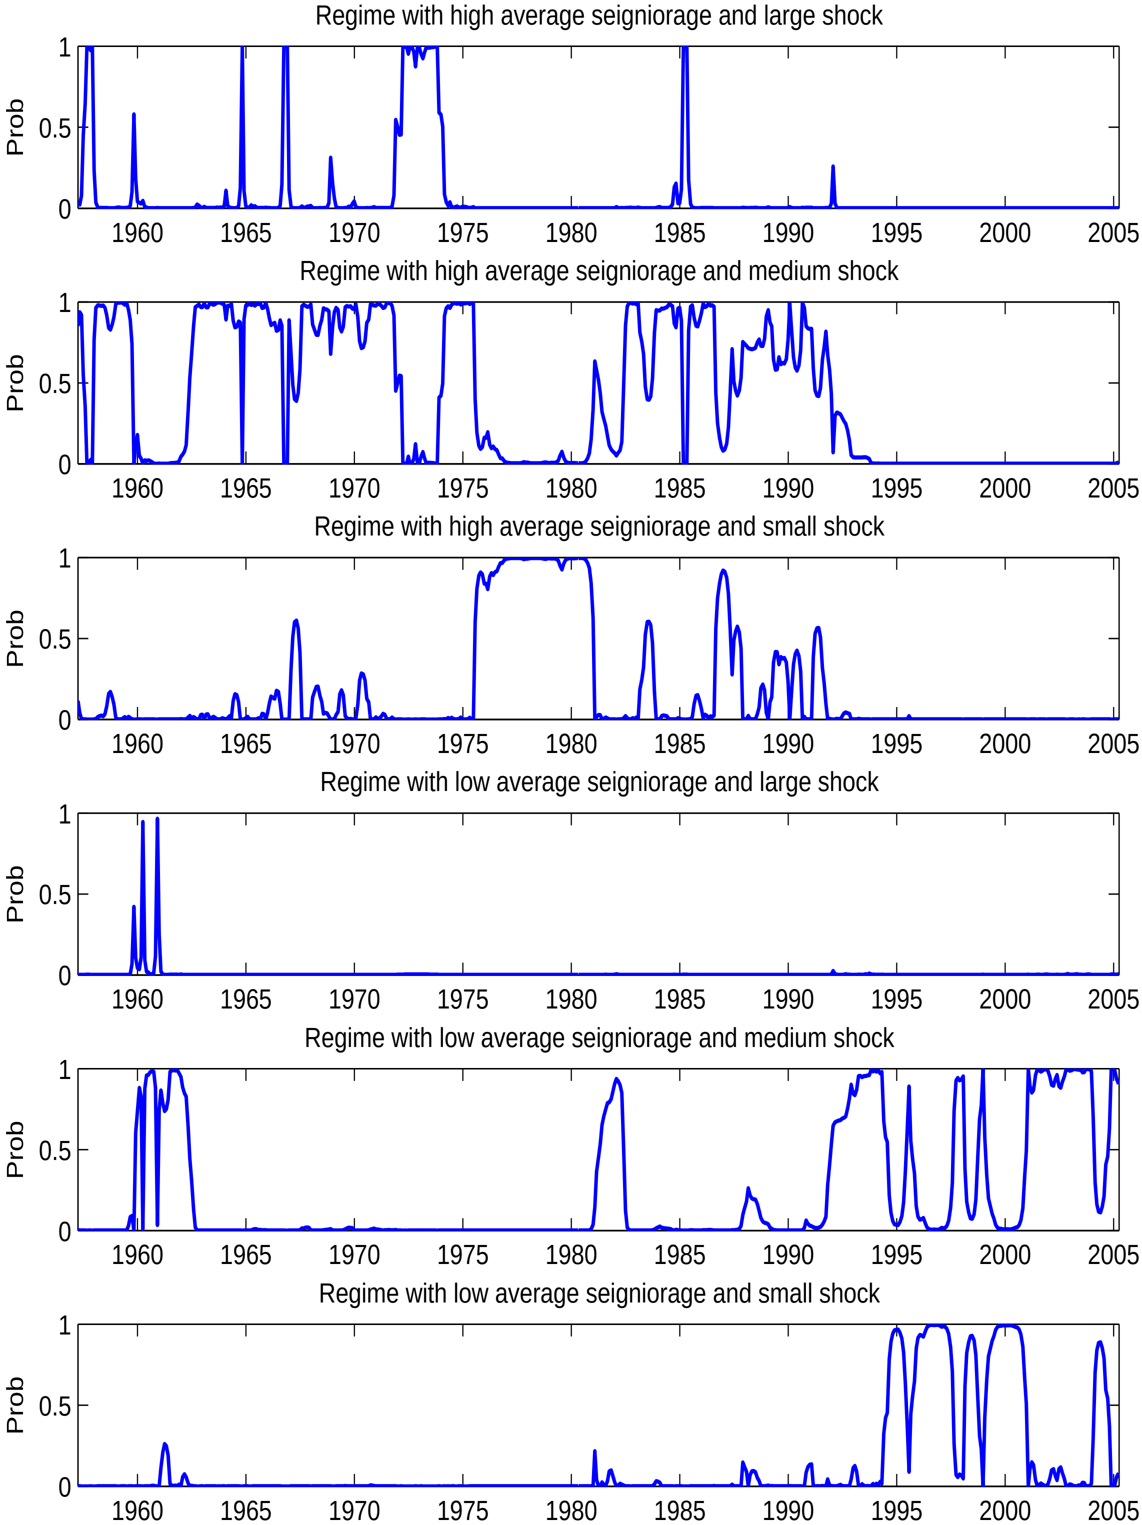

FIGURE 17. Chile: smoothed probabilities of the regimes conditional on the MLEs and the data.

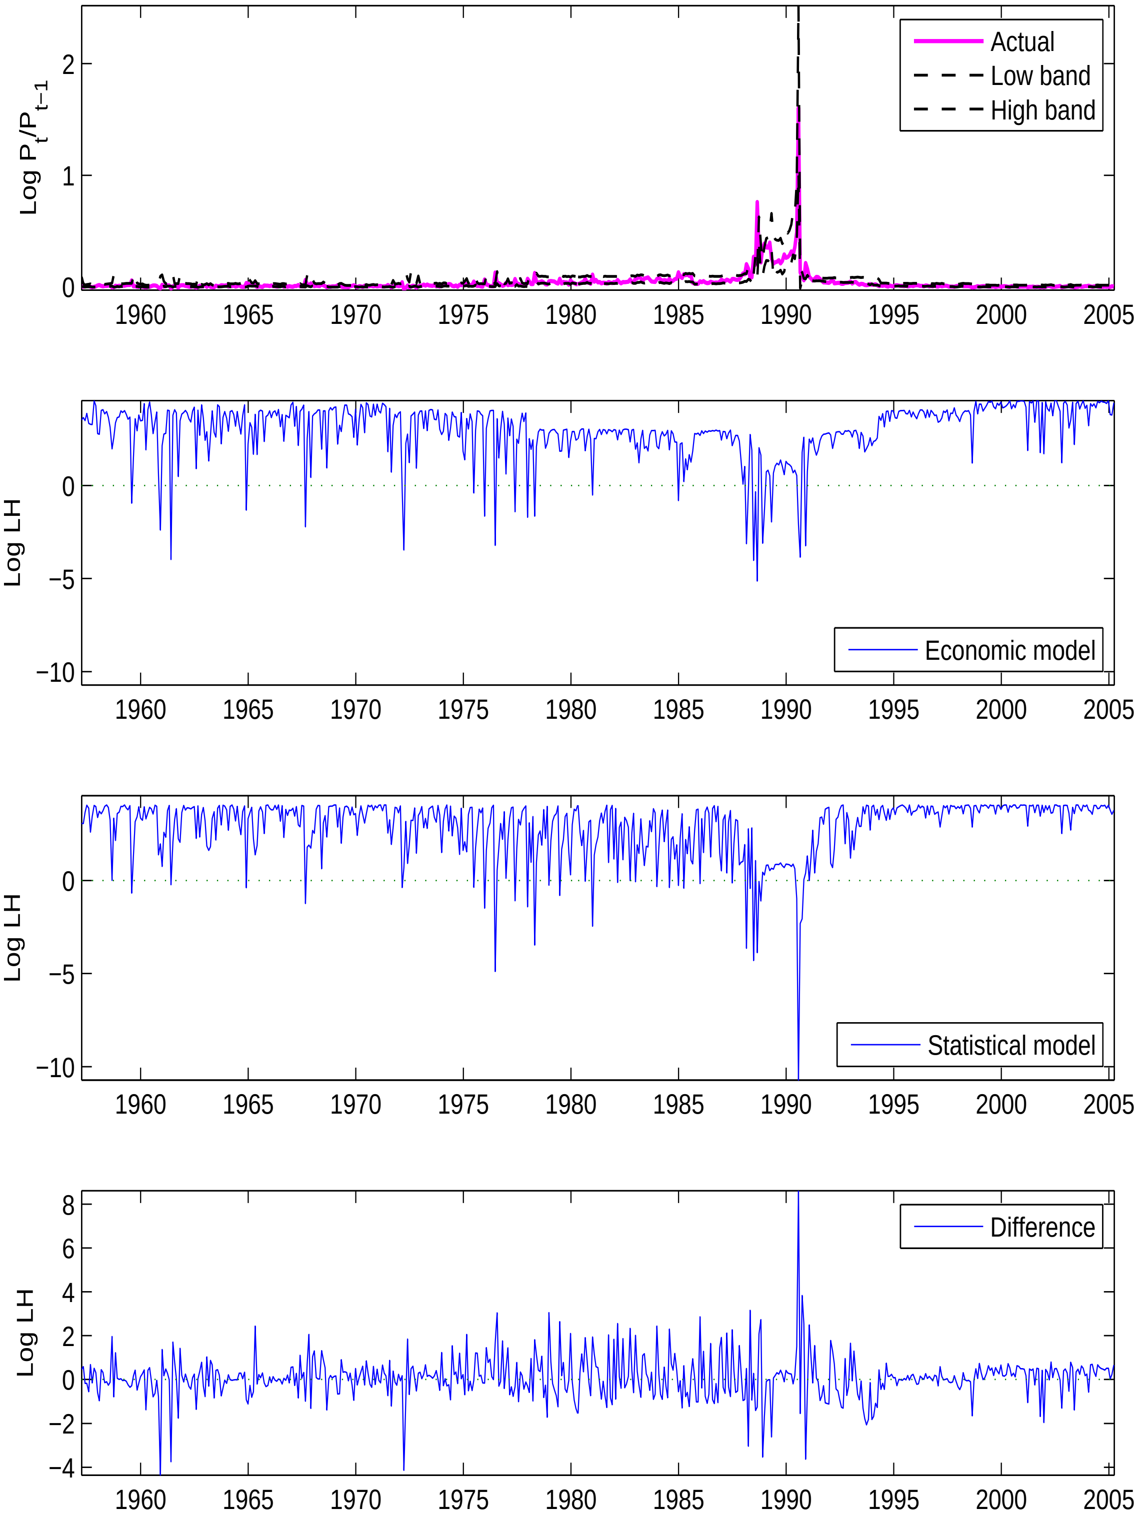

FIGURE 18. Peru: 90% probability bands of one-step predictions from our theoretical model, the log value of the conditional likelihood $p(\pi_t|\Pi_{t-1}, \hat{\phi}$ for both the theoretical and statistical models, and the difference in log conditional likelihood.

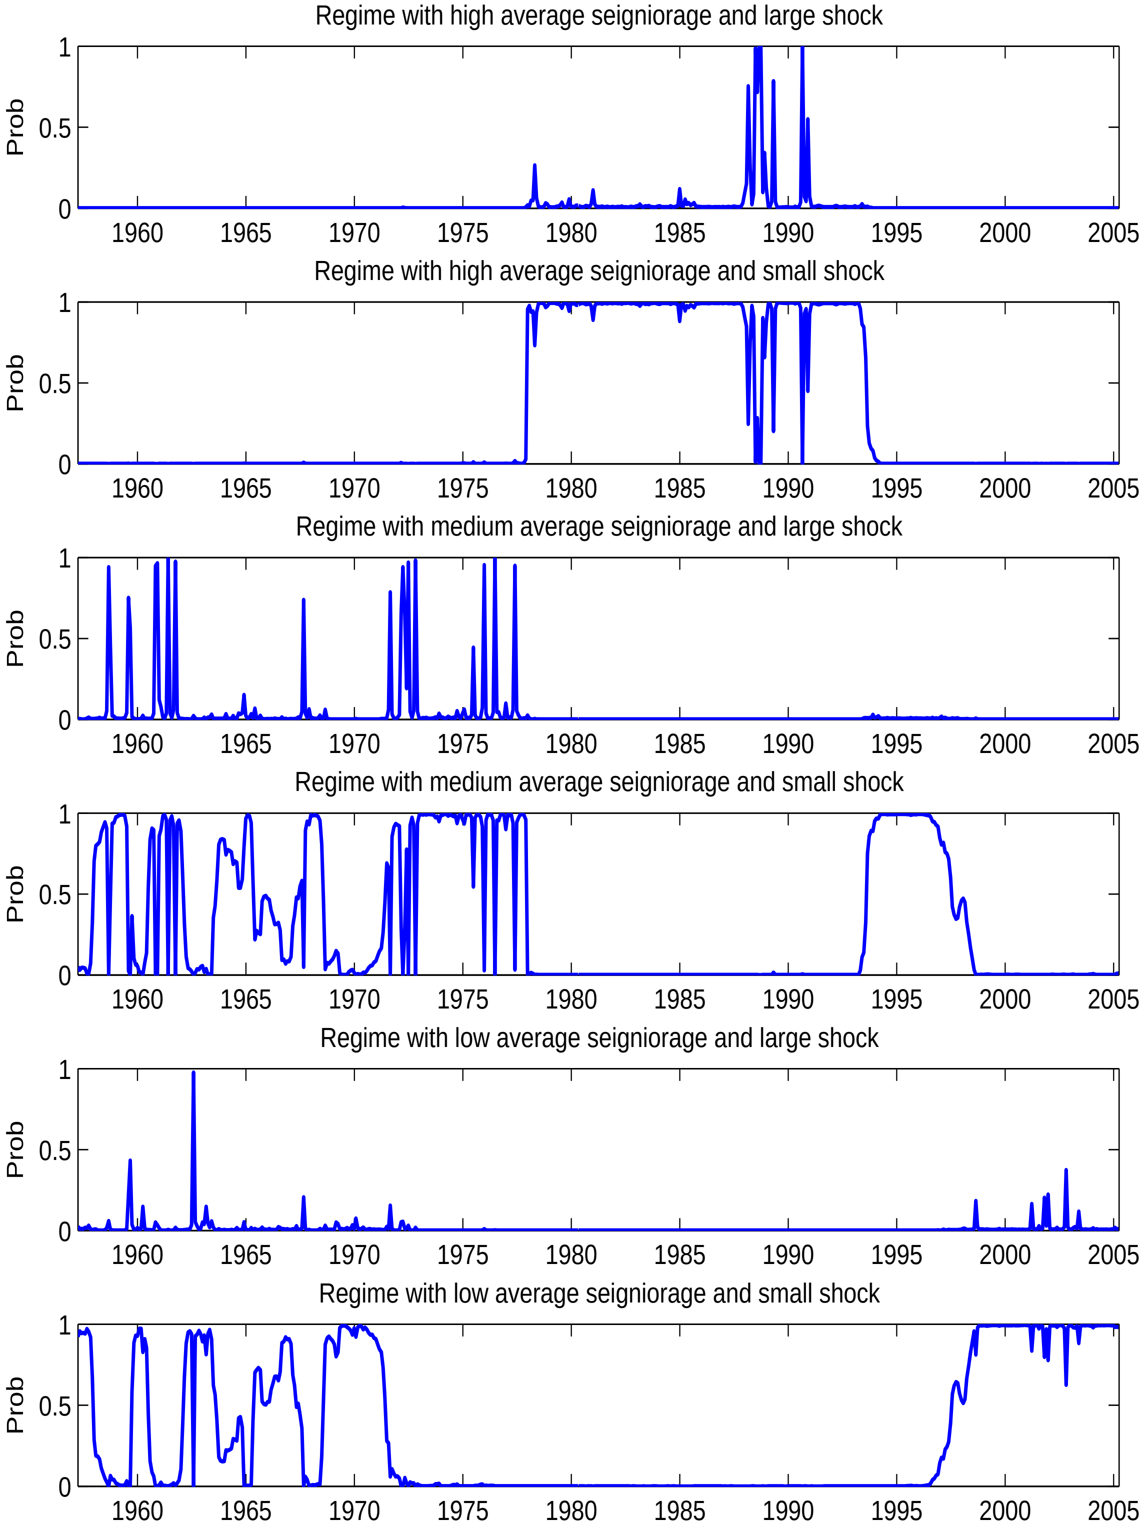

FIGURE 19. Peru: smoothed probabilities of the regimes conditional on the MLEs and the data.

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., 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.

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.

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.

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.

Sargent, T. J. (1999): *The Conquest of American Inflation.* Princeton University Press, Princeton, New Jersey.

Sargent, T. J., and N. Wallace (1987): "Inflation and the Government Budget Constraint," in *Economic Policy in Theory and Practice*, ed. by A. Razin, and E. Sadka, pp. 170–200. St. Martin's Press, New York, New York.

Sargent, T. J., 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 Segementation: 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 Moels 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.

New York University and Hoover Institution, Princeton University and NBER, Federal Reserve Bank of Atlanta