Learning Dynamics and Endogenous Currency Crises

In-Koo Cho
Department Of Economics
University of Illinois

Kenneth Kasa
Department of Economics
Simon Fraser University

November, 2002

Abstract. Currency crises are often followed by severe recessions. This is inconsistent with the predictions of second-generation currency crisis models. In these models, devaluations are supposed to restore competitiveness and stimulate the economy. Recent work on third-generation crisis models remedies this inconsistency by introducing foreign currency debt and adverse 'balance sheet effects'. However, like their earlier second-generation cousins, these newer currency crisis models rely on multiple equilibria and exogenous sunspots to explain the actual outbreak of a crisis. Hence, while recent third-generation models improve our *descriptions* of currency crises, they offer little improvement when it comes to *explaining* them.

In this paper, we address this shortcoming by introducing adaptive learning into the well known model of Aghion, Bacchetta, and Banerjee (2001). Even when equilibrium is unique in this model, we show that the 'escape dynamics' of the learning algorithm produce exactly the kind of Markov-Switching exchange rate behavior that is typically attributed to sunspots. An advantage of our approach is that currency crises become endogenous, in the sense that their stochastic properties can be related to assumptions about learning and other structural features of the economy.

JEL Classification #'s: F31, D83.

Learning Dynamics and Endogenous Currency Crises

In-Koo Cho and Kenneth Kasa

1 Introduction

Economists have made great strides during the past decade in understanding the dynamics of currency crises. Following the ERM Crisis (1992-93), Obstfeld (1994, 1997) developed a class of models based on an open-economy version of the Barro-Gordon model, which explained many of the puzzling features of this episode. Contrary to the predictions of the prevailing first-generation models, countries that left the EMS or widened their intervention bands did not do so because they "ran out of reserves". Instead, the decisions seemed to be motivated by the desire to avoid the unpleasant macroeconomic consequences of remaining in the system. Obstfeld's model formalized these trade-offs, and offered new insights into the nature of currency crises. He showed that when governments choose exchange rates sequentially in order to minimize a loss function, currency crises can become self-fulfilling prophecies, in the sense that expectations of a devaluation can elicit (ex post optimal) responses by the government that ratify those beliefs. Obstfeld's work triggered a flood of research during the 1990s on multiple equilibria in foreign exchange markets.[1]

Unfortunately, these so-called second-generation models encountered empirical problems almost immediately. One of the leading stylized facts of the Mexican and Asian Crises was the combination of devaluation and subsequent recession. Instead of devaluing in order to *avoid* a recession, the devaluations of Mexico and Asia seemed to be *causing* a recession. Clearly, something was missing from second-generation currency crisis models.

Although there are many reasons why a devaluation might prove to be contractionary, the recent literature on third-generation currency crisis models has focused on the role of foreign currency-denominated debt and its adverse "balance sheet effects".[2] This focus is empirically motivated, since foreign currency debt seemed to be at the heart of both the Mexican and Asian Crises.[3] In Mexico's case it was primarily the government that was exposed, whereas in Asia it was primarily the private sector. Either way, a sudden devaluation erodes net worth, and to the extent investment and borrowing capacity is constrained by net worth, due perhaps to information and incentive problems, expectations of a devaluation can turn out to be just as self-fulfilling as in the earlier second-generation class of models. The crucial difference is that now devaluations produce recessions.

Notwithstanding their contrasting predictions about the output effects of devaluations, second- and third-generation currency crisis models share one important feature – both models interpret a crisis as a sudden switch to a 'bad equilibrium'. As is now well known,

[1] See Jeanne (2000) and Flood and Marion (1999) for surveys. Krugman (1996) and Morris and Shin (1998) express skepticism about the relevance of multiple equilibria.

[2] Leading papers include Aghion, Bacchetta, and Banerjee (2000, 2001a,b), Cespedes, Chang, and Velasco (2001), Krugman (1999), Schneider and Tornell (2000), and Caballero and Krishnamurthy (2001).

[3] Burnside, Eichenbaum, and Rebelo (2001) and Aghion, Bacchetta, and Banerjee (2001a) provide a variety of exposure estimates for both Asia and Mexico.

when a model exhibits multiple (static) equilibria, it is usually possible to layer on an exogenous sunspot process that governs switches between (neighborhoods of) these equilibria. It is this exogenous sunspot process that is ultimately to blame for currency crises in these models. Hence, while much progress has been made in *describing* currency crises, remarkably little progress has been made in *explaining* them.[4]

This paper is motivated by our belief that sunspot models push the Rational Expectations Hypothesis too far, in the sense that they rely on an implausible degree of expectations coordination. Instead, we retreat from Rational Expectations and assume that agents must form their beliefs via an adaptive learning process, one that reflects doubts about both the structure of the economy and the credibility of the government's exchange rate policy. In this sense, we share some of the same misgivings as Morris and Shin (1998), who express doubts about the Common Knowledge assumptions of existing currency crisis models. Like them, our approach eliminates multiple equilibria. However, a key result of our paper is to show that despite the fact that equilibrium is unique in our model, this equilibrium is a *stochastic process*, and as a stochastic process it features exactly the sort of Markov-Switching dynamics that is usually attributed to sunspots. An advantage of our approach is that currency crises become endogenous, in the sense that their stochastic properties can be related to assumptions about learning and other structural features of the economy. In contrast, sunspot models place no testable restrictions on the dynamics of currency crises. They merely rationalize their occurrence ex post.[5]

As in the pioneering work of Sargent (1999), we discipline the use of adaptive expectations by requiring beliefs to be updated in a way that eliminates systematic forecast errors. We do this by imposing a least-squares orthogonality condition. However, again following Sargent (1999), we impute a subtle form of specification error to the government. Whereas in reality it is only unanticipated devaluations that matter, the government mistakenly believes that the exchange rate itself affects output. As a result, the evolving beliefs of the private sector inject 'parameter drift' into the government's approximating model, which the government responds to by placing more emphasis on recently observed data. This is accomplished by the use of a constant gain stochastic approximation algorithm.

Even though the underlying structure of the economy is stationary, the fact that the government discounts the past prevents beliefs from ever settling down to a fixed number. Instead, beliefs converge to a diffusion process. Understanding the dynamics of this diffusion process is the focus of our analysis. Using results from Williams (2001) and Cho, Williams, and Sargent (2002), we show that the dynamics consist of two distinct parts. The first part is called the 'mean dynamics'. The mean dynamics reflect the gov-

[4]Even those who advocate a fundamentals-based/first-generation account of recent currency crises often resort to sunspots when it comes to explaining their timing. See, e.g., Burnside, Eichenbaum, and Rebelo (2000).

[5]Of course, learning and multiple equilibria are not mutually exclusive. Woodford (1990) shows that sunspot equilibria can be learned. Kasa (2001) introduces adaptive learning into Obstfeld's (1997) escape clause model, and shows that learning dynamics, rather than sunspots, can generate switches between multiple steady states. More generally, Sargent (1993) and Evans and Honkapohja (2001) discuss the use of stability under learning as an equilibrium selection criterion. Still, in the spirit of Occam's Razor, we prefer an approach that delivers the dynamics we are after in the simplest possible model.

ernment's systematic efforts to eliminate forecast errors. This draws the system toward a 'self-confirming equilibrium', defined to be a situation where agents no longer have an incentive to revise their models. In our model, the self-confirming equilibrium features a currency that is 'overvalued'. This might seem counter-intuitive to those schooled in the Kydland-Prescott (1977) inflation bias tradition. However, remember that a key ingredient of our third-generation currency crisis model is the presence of adverse balance sheet effects. These balance sheet effects make surprise depreciations contractionary. In fact, given the usual distortion between natural and target output levels, the government in our model actually has an incentive to engage in surprise *appreciations*. Of course, the private sector is aware of this incentive and factors it into its expectations. As a result, the Nash equilibrium features an excessively strong currency.

The second part of the dynamics are more novel and more interesting. They are called the 'escape dynamics'. The escape dynamics feature sudden, recurrent movements away from the Nash equilibrium towards the Ramsey equilibrium. Since in the Ramsey outcome the exchange rate is *weaker*, a sudden shift toward Ramsey looks a lot like a currency crisis. In close analogy to Cho, Williams, and Sargent (2002), this switch is triggered by shocks that lead the government to discover a distorted version of the 'natural rate hypothesis', in the sense that it comes to believe (correctly) that there is no systematic relationship between the level of the nominal exchange rate and the level of output.

Putting the two pieces together, we obtain exchange rate paths that resemble the observed exchange rate histories of many developing countries. The mean dynamics produce gradual (real) appreciations of the currency. Then suddenly, when an escape occurs, the exchange rate depreciates and a crisis occurs. In contrast to the usual interpretation, the appreciation phase does not reflect a gradual loss of competitiveness which eventually forces a devaluation. Rather, it reflects the government's desire to avoid adverse balance sheet effects. Whether in fact there will be a recession when the government devalues depends crucially on the beliefs of the private sector. If the private sector has Rational Expectations, then no recession takes place. On the other hand, if the private sector must learn at the same time as the government is learning, then a sudden devaluation can indeed trigger a recession.

The rest of the paper is organized as follows. The next section outlines some stylized facts about currency crises. Our goal is to explain these facts. The third section develops our baseline third-generation crisis model. Although there are many possible models we could use as a platform, we employ a version of Aghion, Bacchetta, and Banerjee's (2001a,b) model (hereafter denoted ABB). We use this model because of its familiarity to many readers and its simplicity. As emphasized by ABB, this model potentially admits multiple Rational Expectations equilibria. In fact, ABB rely on multiple equilibria to generate currency crises. As noted above, we do not need multiple equilibria, although we cannot rule them out. Due to balance sheet effects and the government's desire to stimulate output, the Nash equilibrium features a stronger currency than the Ramsey equilibrium.

The fourth section incorporates adaptive learning into the model. To do this, we first posit a perceived law of motion for the government. This perceived law is assumed to be misspecified in the sense that it misinterprets the role of private sector expectations. Rather than doing Rational Expectations Econometrics, as described for example in Lucas and Sargent (1981), the government adopts the short-cut of allowing for parameter drift in its econometric model. This model takes the form of a vector autoregression of output and the (nominal) exchange rate on their own lagged values and on the current nominal interest rate. The private sector also has a perceived model, which it uses to forecast the exchange rate. In contrast to Sargent (1999) and Cho, Williams, and Sargent (2002), we do not impose a "Fed Watcher" assumption. In our model the public does not know the government's policy function. However, we do assume that the public's beliefs about the government are correctly specified, so that in a self-confirming equilibrium its beliefs about the exchange rate are aligned with the government's beliefs. We show in section four that a self-confirming equilibrium of this two-sided learning problem can be characterized by a nonlinear system of least-squares orthogonality conditions. Parameter values that satisfy this system have the property that neither the government nor the private sector has an incentive to revise its model of the economy.

The fifth section applies large deviations methods to characterize the escape dynamics of the learning algorithm. Large deviations methods can be thought of as a generalization of the Central Limit Theorem, which permit a rigorous treatment of rare events, like currency crises.[6] We show that crises depend solely on the beliefs of the government, although the beliefs of the private sector are important in determining the extent of any real effects from the crisis. Loosely speaking, this is due to the fact that the private sector does not take any direct action in response to its evolving beliefs. It merely revises its forecasts of the exchange rate.

Finally, section six contains simulations of the model. These simulations suggest that crises occur when the government "loses control" of an attempted modest depreciation. Given the feedback in the model, it turns out that a mild depreciation can quickly spiral into a vicious circle of currency depreciation and falling output. Each additional decline in output leads the government to respond by cutting interest rates which, given the right circumstance, can reinforce the currency depreciation. Although this seems to confirm some of the IMFs worse fears about currency crises, and supports its view that countries should respond to crises with tight monetary policy, the fact remains that crises have occurred in many countries *despite* a vigorous interest rate defense. The simulations suggest that in the model the early phase of a crisis is a period of relatively *low* interest rates, contrary to observed data. Although with a sufficiently high risk premium higher interest rates can be produced *in response* to a crisis, our adaptive learning model has a hard time getting interest rates to rise before a crisis. Our current work is attempting to reconcile this discrepancy.

[6]Freidlin and Wentzell (1998) contains a comprehensive treatment of continuous-time large deviations methods. Kifer (1990) and Dupuis and Kushner (1985, 1989) contain discrete-time results. Large deviations ideas were originally introduced into economics in the evolutionary game theory literature (see, e.g., Kandori, Mailath, and Rob (1993)). The use of large deviations methods to characterize escape dynamics and recurrent cycles was pioneered by Sargent (1999), and extended by Williams (2001) and Cho, Williams, and Sargent (2002).

2 Some Stylized Facts

Before developing the model, it is useful to first lay out the basic facts we are trying to explain. In a recent paper, Tornell and Westermann (2001) examine a wide range of macroeconomic variables before, during, and after a currency crisis. Their data consist of annual observations from a panel of 39 middle-income countries during the period 1980-1999. They argue that despite some obvious variation, both cross-sectionally and over time, there is a sense in which "currency crises are all alike".

Figure 1 reproduces plots of four key time series from Tornell and Westermann's paper. Each plot depicts the mean of the cross-section at a given point in time (after allowing for country-specific fixed effects). Each begins three years prior to a crisis and ends three years after a crisis. Figure 1a plots the growth rate of the real exchange rate (relative to tranquil periods), Figure 1b plots detrended real bank credit, Figure 1c plots detrended real output, and Figure 1d reports an interest rate spread, defined as the difference between lending rates and deposit rates.[7]

Figure 1

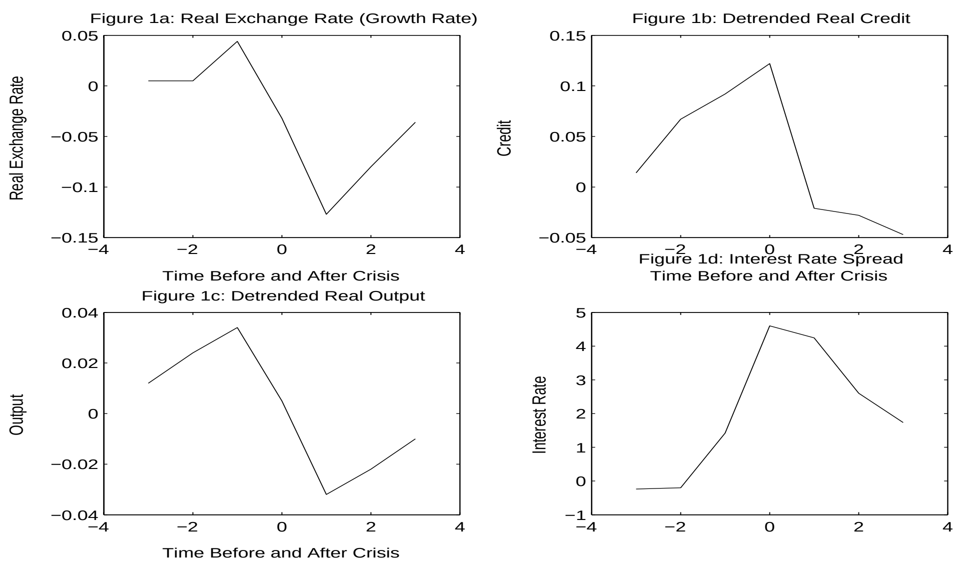

Inspection of these plots suggests the following stylized facts about currency crises:

Fact 1: *Relative to tranquil periods, the real exchange rate appreciates by about 5% prior to a crisis, and depreciates by about 16% during the crisis. Hence, on average the exchange rate depreciates by about 20%.*

[7]Detrending is accomplished via the HP filter. The interest rate spread is based on a subset of 11 countries: Argentina, Brazil, Chile, Indonesia, Finland, Korea, Malaysia, Mexico, Philippines, Sweden, and Thailand.

Fact 2: *Real GDP is above trend by about 3% the year before a crisis, and falls to about 3% below trend the year after a crisis. Hence, on average real GDP declines during a crisis by about 6% relative to trend.*

Fact 3: *Interest rate spreads rise by about one percentage point the year before a crisis, and then rise an additional 3.5 percentage points during the crisis. Hence, on average interest rates rise by about 4.5 percentage points.*

Fact 4: *Prior to a crisis there is a "lending boom", with real credit expanding by about 10% relative to trend. The crisis then precipitates a "credit crunch", with real credit falling by about 15% relative to trend.*

Tornell and Westermann (2001) proceed to develop a model that can explain these (and other) stylized facts. They argue persuasively that currency crises are part of a boom-bust cycle, which must be understood in its entirety; the boom and the bust are two sides of the same coin. Their model has the same two key ingredients as our model: (i) financial market imperfections that produce borrowing constraints, and (ii) distortions (due to government policy) that provide incentives for unhedged foreign currency debt. However, in our view their model shares two drawbacks with the rest of the third-generation currency crisis literature. First, they regard expectations as exogenous, and hence cannot explain the timing and frequency of crises. Second, and closely related to this, their model attributes the initial boom to an exogenous financial liberalization. Assuming this liberalization is a one-time event, their model is incapable of explaining a key feature of currency crises in many countries, namely, their *recurrence*. Our model attempts to overcome both these drawbacks.

3 A Third-Generation Currency Crisis Model

In this section we outline the model of ABB (2001a). This model will serve as a platform for our analysis of escape dynamics and currency crises. Since we are primarily interested in learning dynamics, the presentation here will be brief. The reader should consult ABB's paper for full details.

The defining characteristic of third-generation crisis models is the presence of unhedged foreign currency liabilities (i.e., balance sheet effects), which make (unanticipated) devaluatons contractionary. Particular models differ according to who incurs the liabilities and why. For example, in the models of Burnside, Eichenbaum, and Rebelo (2001) and Dooley (2000), it is the banking sector that is exposed, whereas in the models of Cespedes, Chang, Velasco (2001), Krugman (2000), Jeanne (2000), and ABB (2000, 2001a,b) it is firms that are exposed. Generally speaking, models that focus on the exposure of the banking sector tend to attribute the exposure to government deposit guarantees, whereas models that focus directly on firms tend to blame the exposure on asymmetric information problems. We follow ABB (2001a,b), and attribute balance sheet effects to moral hazard.

Most of the attention in this literature focuses on the combination of these balance sheet effects with financial market imperfections, which cause borrowing to be constrained by net worth. As the literature has demonstrated, this combination creates a potent propagation mechanism. With net foreign currency liabilities, devaluations erode net worth. Then, if borrowing is constrained by net worth, the decline in net worth produces a decline in investment and output, which then reinforces the original exchange rate decline. As in second-generation models, this circularity exposes the economy to multiple equilibria and sunspot fluctuations.

The ABB (2001a) model combines three essential ingredients. First, prices are assumed to be preset one period in advance. This produces real effects from nominal exchange rate changes. Second, financial market imperfections limit borrowing to be an endogenously determined multiple of net worth. This creates a 'financial accelerator'. Third, firms are assumed to be financed, at least partially, by foreign currency debt. As a result, exchange rate changes trigger the financial accelerator.

In full generality, these assumptions would produce a model that is quite complex. The contribution of ABB is to come up with a tractable formulation. We now proceed to outline this formulation.

3.1 Production and Price-Setting

Output of a single good is produced by competitive consumer/entrepreneurs according to a standard concave production function:

$$Y_t = f(k_t)\varepsilon_t \tag{1}$$

where Y_t is output, k_t is the capital stock, and ε_t is an i.i.d productivity shock. Capital is assumed to depreciate fully within the period.

Following ABB (2001b), we assume domestic entrepreneurs face a competitive fringe of foreign producers. Foreign firms have constant marginal costs. Both domestic and foreign firms must set prices at the beginning of each period, before the realization of the production shock and the exchange rate. Assuming foreign marginal costs are constant, and normalizing them to unity, implies that domestic firms must then set $P_t = E_{t-1}S_t$, where S_t denotes the nominal exchange rate, defined as the price of foreign currency. Hence, PPP holds ex ante, but not necessarily ex post.

3.2 The Credit Multiplier and the Currency Composition of Debt

Capital consists of the entrepreneur's own wealth, w_t, and any additional borrowed funds, d_t. That is, $k_t = w_t + d_t$. Firms can borrow either in terms of domestic currency, at interest rate i_{t-1}, or in terms of foreign currency, at (constant) interest rate i^*. As with price-setting, investment decisions must be made at the beginning of the period (so that the loan rate is the prevailing, prior period rate).

Debt contracts are only partially enforceable. In particular, borrowers can pay a cost, cP_tk_t, proportional to the amount borrowed, that allows them to abscond with the funds.

However, if a borrower does default, there is a probability, p, that the lender is able to track him down and collect anyway. Hence, assuming for now a domestic currency loan, a borrower will choose to repay if and only if:

$$P_tY_t - (1+i_{t-1})P_{t-1}d_t \geq P_tY_t - cP_tY_t - p(1+i_{t-1})P_{t-1}d_t$$

Collecting terms gives us the incentive compatability constraint, $d_t \leq \mu_t w_t$, where the 'credit multiplier', μ_t, is given by:

$$\mu_t = \frac{c}{(1-p)(1+r_{t-1}) - c} \tag{2}$$

where r_{t-1} is the real interest rate.

There are several things to note about this multiplier. First, notice that it increases with p. That is, firms can borrow more when the 'monitoring technology' improves. ABB interpret this as a proxy for financial market development. Second, because lending decisions are made before any shocks are realized, μ_t will be independent of the currency denomination of debt as long as Uncovered Interest Parity and (ex ante) PPP hold. Third, notice that μ_t is state dependent. That is, it varies with the real interest rate. Later, when we incorporate learning, we will approximate this dependence. Finally, notice that this model of debt is quite different from the influential model of Kehoe and Levine (1993). They assume debt contracts are enforced by future exclusion from the capital market. In contrast, borrowers in this model start each period with a clean slate.

As noted earlier, firms are free to borrow in either currency. One of the critical questions to emerge in the wake of recent financial crises is why a domestic economy would choose to become so exposed. That is, why is there so much unhedged foreign currency borrowing? By now, there are many (not necessarily mutually exclusive) theories. Perhaps the most common explanation relies on government bailout guarantees.[8] Alternatively, Jeanne (2000b) shows that foreign debt might play a signalling role, which lowers its interest rate. A third possibility is to appeal to moral hazard, as in ABB (2001b). They show that if the currency composition of a borrower's debt is not observable, it might be optimal to borrow abroad. All of these theories imply a distinction between privately optimal and socially optimal financing decisions, since individual firms ignore the effects of their borrowing decisions on the country's financial fragility.

Which of these theories is more important is irrelevant for us, because we follow ABB (2001a) and assume that debt is exogenous. In particular, we assume the foreign currency value of foreign debt is held constant at $\bar{d}^*$, while at the same time the real interest burden of domestic debt is also constant at $\bar{d}$. This is a tremendous simplification. It might not be such a bad assumption in models like ABB (2001b) and Burnside, Eichenbaum, and Rebelo (2001), where firms are led to a corner solution, and foreign debt is at its (credit constrained) maximum. Later we discuss how our results might be altered if the level and composition of debt were endogenous.

[8] See, e.g., Dooley (2000), Burnside, Eichenbaum, and Rebelo (2001), and Schneider and Tornell (2000).

In general, the dynamics of wealth and output depend on whether the borrowing constraint is binding. In fact, Krugman (1999) and ABB (2001a,b) emphasize that the distinction between a binding and a non-binding borrowing constraint lies at the heart of the model's nonlinearity and its capacity to generate multiple equilibria. However, since generating multiple equilibria is not our goal, we further assume that the borrowing constraint binds in every period. This implies restrictions on the production function and the support of the shocks. For example, it must be the case that $f'(k) > (1+i^*)(S^e/S)$ for all relevant values of k and S^e/S, where S^e/S is the expected (percentage) depreciation of the domestic currency. Otherwise, the firm would rather invest abroad.[9]

When the borrowing constraint is binding output can be written as a function of the entrepreneur's wealth:

$$Y_t = f((1+\mu_t)w_t)\varepsilon_t \tag{3}$$

To derive the law of motion for output we therefore need to derive the law of motion for wealth. Following ABB (2000, 2001a), we assume entrepreneurs consume (if they can) a fixed fraction, α, of their wealth. If wealth is zero they consume nothing. With full depreciation, wealth therefore evolves according to:

$$w_{t+1} = (1-\alpha)\frac{\Pi_t}{P_t} \tag{4}$$

where Π_t represents (nominal) profits net of debt repayments:

$$\Pi_t = P_t Y_t - (1+i_{t-1})P_{t-1}d_t - (1+i^*)\frac{S_t}{S_{t-1}}P_{t-1}d_t^* \tag{5}$$

where d_t is the (time-varying) real value of domestic debt and d_t^* is the (time-varying) real value of foreign debt. Combining eqs. (3), (4), and (5), and using the fact that nominal foreign currency debt is constant, i.e., $P_{t-1}d_t^*/S_{t-1} = \bar{d}^*$, and the fact that $(1+r_{t-1})d_t = \bar{d}$, delivers the following law of motion for output:

$$Y_t = f\left((1+\mu_t)(1-\alpha)\left[Y_{t-1} - \bar{d} - (1+i^*)\frac{S_{t-1}}{P_{t-1}}\bar{d}^*\right]\right)\varepsilon_t \tag{6}$$

Equation (6) is one of the two key equations of the model. ABB (2001a) call it the W-curve. Notice the role of balance sheet effects. Since prices are predetermined, a depreciation raises the foreign debt burden, which exerts a contractionary effect on (future) output. However, this is not quite the end of the story. As noted by ABB (2001a), since μ_t depends negatively on r_{t-1}, and since r_{t-1} depends negatively on S_{t-1}/P_{t-1} (due to Uncovered Interest Parity, $1+r_t = (1+i^*)(P_t/S_t)$), unanticipated depreciations also relax the borrowing constraint for a given level of wealth, since they reduce domestic real interest rates. This exerts an offsetting expansionary effect on output. So in general, as you might

[9] With full depreciation of the capital stock, and given the above assumptions on the relationship between the foreign interest rate and the domestic interest rate, the strategy of circumventing the borrowing constraint by accumulating a 'war chest' (Chamberlain and Wilson (2000)) would neither be possible nor in the entrepreneur's interest.

expect, the effect of a surprise depreciation on output is ambiguous. What is clear is that the negative effect is more likely to dominate when $\bar{d}^*$ is larger.

As it stands, equation (6) is still too general to be useful, despite the many assumptions that have already been made in deriving it. Therefore, we make one last assumption, and take a log-linear approximation. Since expectations are the focus of our analysis, when doing this we use the pricing rule to substitute $E_{t-2}S_{t-1}$ in place of P_{t-1}. This gives us:

$$y_t = \bar{y} + \rho y_{t-1} + \theta(s_{t-1} - E_{t-2}s_{t-1}) + v_{1t} \tag{7}$$

where lowercase letters are natural logs of the corresponding uppercase letters, and where the error term, v_{1t}, combines the i.i.d productivity shock, ε_t, and the approximation error, which is also presumed to be i.i.d.

Equation (7) is a standard open-economy 'expectations-augmented Phillips Curve', with three notable exceptions. First, lagged output enters the model, since there is a one-period lag between investment and output. Second, and more importantly, the 'slope' of the Phillips Curve is indeterminate. As noted earlier, the sign of θ depends on the relative importance of balance sheet effects. In what follows, we assume balance sheet effects are relatively strong, so that $\theta < 0$. Third, again because of the production lag, output in this model is predetermined. That is, exchange rate changes affect output with a one-period lag.

3.3 Financial Markets and Monetary Policy

ABB (2001a) close their model by combining the Uncovered Interest Parity condition with a standard money demand equation. This delivers a second equation relating the exchange rate to future output, called the IPLM curve. Changes in the money supply shift the IPLM curve. This allows ABB to make statements about how monetary policy should respond to the fait accompli of a currency crisis.

In this paper, we assume the government chooses the interest rate to minimize an explicit intertemporal loss function. This loss function reflects a trade-off between output stability, exchange rate stability, and interest rate stability.

$$\min_{\{i_t\}} E_t \frac{1}{2} \sum_{j=0}^{\infty} \delta^j [\omega_s(s_{t+j} - s^*)^2 + \omega_y(y_{t+j} - y^*)^2 + \omega_i(i_{t+j} - i^*)^2] \tag{8}$$

where the parameters s^*, y^* and i^* are arbitrary targets, and the ω parameters summarize the relative cost of the fluctuations. A fixed exchange rate regime, albeit an uninteresting one, would result if $\omega_y = \omega_i = 0$.

When setting the interest rate each period, the government confronts two constraints arising from the private sector. The first is the Phillips Curve in equation (7). The second is an Uncovered Interest Parity condition, modified to include a risk premium term that varies inversely with last period's output (relative to the steady state):[10]

$$i_t = i^* + E_t s_{t+1} - s_t - \phi(y_{t-1} - y^{ss}) + v_{2t} \tag{9}$$

where the i.i.d noise term v_{2t} reflects a combination of omitted risk factors and approximation error. ABB are able to ignore the risk premium since they only consider unanticipated one-time shocks. In general, the risk premium will be a nonlinear function of the net worth/capital ratio.[11] Our linearization is designed to capture the dominant time-varying element of this function. While in principle ϕ could be calibrated to data on bankruptcy rates and bankruptcy costs (or in our context, sovereign defaults), for simplicity we treat it as a free parameter.

3.4 Nash and Ramsey Equilibria

As ABB readily acknowledge, they do not discuss the potential importance of expectations and credibility in their model. Instead, they confine their attention to purely temporary, totally unanticipated shocks. From inspection of equations (7) and (9), however, it is clear that these issues are going to be central in any real world setting. Devaluations that are anticipated will be fully incorporated into prices and interest rates, and as a consequence have no real effects. But beyond the mere presence of expectations is the fact that the expectations concern the future actions of another agent, which raises issues of commitment and credibility.

Although not the focus of our analysis, it is useful for reference purposes to derive the Nash and Ramsey equilibria of the model with Rational Expectations. A Markov Perfect Nash equilibrium can be derived using a standard "guess and verify" strategy.

Proposition 3.1: *A Markov Perfect Nash equilibrium can be represented as the following vector ARMA(1,1) process:*

$$x_t = G_1 x_{t-1} + G_2 v_t + G_3 v_{t-1} \tag{10}$$

$$i_t = -F_1 x_{t-1} - F_2 v_{t-1} \tag{11}$$

where $x_t = (1, y_t, s_t)'$ *denotes the state vector, and* $v_t = (v_{1t}, v_{2t})'$ *is the shock vector.*

To prove this, first note that by appropriate choice of R we can write the government's problem as a standard discounted Linear-Quadratic Regulator (LQR):

$$\min_{i_t} E_t \sum_{j=0}^{\infty} \delta^j (x'_{t+j} R x_{t+j} + \omega_i i^2_{t+j}) \tag{12}$$

[10] Note, use of the same notation to denote the target interest rate and the foreign interest rate will not create confusion, since in what follows we always set them equal to each other.

[11] See, e.g., Cespedes, Chang and Velasco (2001) for a detailed derivation of the risk premium in a model that is quite similar to ours.

The only difficulty is that the state transition equation contains expectations:

$$A_0 x_t = A_1 x_{t-1} + A_2 E_t x_{t+1} + A_3(x_{t-1} - E_{t-2}x_{t-1}) + A_4 i_t + A_5 v_t \tag{13}$$

where the definitions of the A_i matrices follow directly from equations (7) and (9). Given this, conjecture a solution of the form given in equations (10) - (11), and use these conjectures to evaluate the expectations. Substituting back into (13) yields a conventional state transition equation:

$$\tilde{x}_t = J_1 \tilde{x}_{t-1} + J_2 i_t + J_3 v_t \tag{14}$$

where the J_i coefficient matrices are now functions of the A_i matrices and the undetermined coefficients in G_i and F_i, and the state vector has been augmented to include the shocks, $\tilde{x}_t = (x_t, v_t)'$.

We now have a standard LQR problem. The optimal decision rule can be computed by solving a Ricatti equation:

$$i_t = -\tilde{F}_1 \tilde{x}_{t-1} \tag{15}$$

where $\tilde{F}_1$ is function of all the parameters. Finally, plugging (15) into (14) delivers two sets of equilibrium fixed point conditions:

$$\begin{aligned} G_1 x_{t-1} + G_2 v_t + G_3 v_{t-1} &= [J_1(G_i, F_i) - J_2(G_i, F_i)\tilde{F}_1(G_i, F_i)]\tilde{x}_{t-1} + J_3(G_i, F_i) v_t \\ -F_1 x_{t-1} - F_2 v_{t-1} &= -\tilde{F}_1(G_i, F_i)\tilde{x}_{t-1} \end{aligned}$$

An equilibrium exists if we can find a stable G_1 matrix satisfying these conditions.

Although solution of these equations lies beyond the reach of pencil and paper methods, we have nonetheless verified the ARMA(1,1) structure of the equilibrium. It is important to remember, however, that even if a solution exists there is no guarantee that it is unique. For example, the representation in (10) - (11) might be overparameterized if the parameter values produce common factors in the AR and MA components. Alternatively, the parameters might permit arbitrary martingale difference sequences to be appended to the above solution. Since we do not focus on rational expectations equilibria, we do not worry about these possiblities.

The Nash equilbrium assumes the government is unable to commit to a policy, and consequently must take the expectations of the private sector as given. It is also of interest to consider a Ramsey equilibrium, where commitment is allowed. In a Ramsey equilibrium the government acts as a Stackelberg leader, and takes advantage of its ability to shape private sector expectations.

As is now well known, this problem is not recursive, since the government's current losses depend on its anticipated future actions. This gives the government an incentive to make promises that it may not want to keep ex post. Following the lead of Kydland and Prescott (1980), however, we can make the problem recursive by augmenting the state vector to include the current values of the Lagrange Multipliers on the two constraints arising from the private sector. This ensures the government keeps its promises.[12] Doing this gives us the following result:

[12]See Marcet and Marimon (1999) for extensions of this idea.

Proposition 3.2: *If the government can commit to an interest rate policy, there is a Ramsey Equilibrium characterized by the following system:*

$$\begin{bmatrix} x_t \\ i_t \\ \lambda_t \end{bmatrix} = \Lambda_1 \begin{bmatrix} x_{t-1} \\ i_{t-1} \\ \lambda_{t-1} \end{bmatrix} + \Lambda_2 v_t \tag{16}$$

where the state vector is, $x_t = (1, y_t, s_t)'$, *the shock vector is,* $v_t = (v_{1t}, v_{2t})'$, *and* λ_t *is a* 2×1 *vector of Lagrange Multipliers attached to the bottom two rows of the state transition equation in (13).*

To prove this we just need to form a Lagrangian by combining the objective function in (12) with the constraints in equation (13). The saddlepoint first-order conditions are:

$$\begin{aligned} \omega_i i_t - \lambda_t' e_{23} A_3 &= 0 && (17) \\ e_{23}[Rx_t + A_0' e_{23}' \lambda_t - \delta E_t A_1' e_{23}' \lambda_{t+1} - \delta^{-1} A_2' e_{23}' \lambda_{t-1}] &= 0 && (18) \\ e_{23}[A_0 x_t - A_1 x_{t-1} - A_2 E_t x_{t+1} + A_4 i_t + A_5 v_t] &= 0 && (19) \end{aligned}$$

where e_{23} is a 2×3 selection matrix that picks out the bottom two rows of a 3×3 matrix. It is important to note that these FOCs only apply for dates $t \geq 1$. With commitment, the initial date, $t = 0$, must be treated separately since we are allowing the government to exploit the private sector's initial conditions on this one date. Mathematically, this means we must solve this system with the initial condition $\lambda_{-1} = 0$, since by assumption previous promises are not binding at $t = 0$. It also means that the FOC with respect to x_t in (18) needs to be modified to include a $\lambda_1' A_4$ term, arising from the fact that at $t = 0$ prior private sector expectations are given. Economically, the way the government chooses to take advantage of the private sector depends on the relative cost of the fluctuations. If the initial value of output is sufficiently below the target, y^*, and if the government places sufficiently high weight on output fluctuations, then the it will choose to engage in a surprise *revaluation* at date $t = 0$. Doing this lowers the domestic currency cost of the foreign currency debt. This will increase net worth, relax borrowing constraints, and stimulate output.[13]

At this point, we can repeat the above guess and verify solution strategy on equations (17) - (19). In general, this produces a solution of the form given by (16), although in most applications there will be many zeros in the Λ_1 matrix.

4 Incorporating Adaptive Learning

In principle, we could at this point proceed to estimate and/or simulate the model given in the previous section by equations (10) - (11) or (16). Doing this might reveal some interesting effects of commitment on exchange rate dynamics. Following Woodford (1999),

[13] Of course, another way to exploit the initial condition would be to simply default on the initial foreign debt. We rule this out.

for example, we might expect that commitment produces 'smoother' interest rate and exchange rate processes. Also, it should be noted that even with Rational Expectations the model will generate *some* exchange rate dynamics, since the dynamics in y_t generate a time-varying incentive to depreciate, which then leads to expected and actual depreciations.[14] However, these dynamics will not look anything like currency crises. They will be smooth and episodic. Moreover, if the underlying shocks are symmetrically distributed, the (linearized) Rational Expectations Equilibria will have a hard time replicating the sharp nonlinearities that, almost by definition, characterize observed currency crises.

As noted in the Introduction, the conventional strategy for introducing nonlinearities is to exploit the nonlinearity of the budget constraint in models with borrowing constraints. This can produce multiple equilibria, and open the door to sunspot fluctuations. Our paper pursues an alternative strategy. We inject nonlinearity by introducing adaptive feedback between beliefs and outcomes. This puts us in a different 'space'. Rather than focus on switches between multiple steady states, our analysis focuses on the tail events of a single stochastic process.

We regard a learning approach to currency crises as more attractive and persuasive than a multiple equilibrium/sunspot approach, despite the fact that our learning algorithm is only 'boundedly rational' whereas sunspot equilibria are 'rational'. Although one could argue that the recurrent pattern of boom and bust, and exchange rate stability and collapse, has been operating in these countries long enough to be fully learned, we suspect that there is something about the sheer virulence of these episodes that makes them qualitatively different from learning the dynamics of, say, the inventory cycle or the Fed's reaction function.[15]

Of course, the downside to a boundedly rational learning approach is that it requires us to specify *two* models, one for the true structure of the economy and one for the beliefs or perceptions of the agents in the economy. The great virtue of Rational Expectations is that we only need to specify one. This means that the ultimate persuasiveness of our model will partly depend on a priori assessments of the plausiblity of agents' beliefs.

With this in mind, we try to be fairly unrestrictive and realistic in specifying the beliefs of the government and the private sector. In particular, we assume the government's perceived model takes the form of an unrestricted first-order vector autoregression (with intercepts) for the two state variables y_t and s_t:

$$x_t = \beta_1 x_{t-1} + \beta_2 i_t + u_{1t} \tag{20}$$

where $x_t \equiv (1, y_t, s_t)'$, β_1 is a 3×3 coefficient matrix, and β_2 is a 3×1 vector. By construction, the first row of β_1 is [1 0 0] and the first row of β_2 is 0.

Comparing equation (20) to the actual model in equations (7) and (9), one can see that the government's model contains a couple of misspecifications. First, the government

[14] In a similar model, Ireland (1999) argues that variation in 'time-consistency bias' can explain the rise and fall of U.S. inflation during the 1970s and 1980s.

[15] See Heymann and Leijonhufvud (1995) for an extended discussion of this point. On the other hand, see Baxter (1985) for a reasonably successful attempt to account for hyperinflations with Rational Expectations and Bayesian learning.

assumes interest rate changes have an immediate effect on output, whereas in reality they have a lagged effect. We adopt this specification for analytical convenience. If we are to maintain a low dimensional LQR structure we must either assume interest rates respond to lagged information and have contemporaneous effects on the state, or alternatively, assume that the interest rate responds to contemporaneous information and has a lagged effect on the state. In light of the model's simultaneity between the interest rate and the exchange rate, and the high frequency application we have in mind, we felt the first option was the more reasonable one. The second and more important misspecification is that the government's model fails to properly account for the expectations of the private sector. As a result, the evolving beliefs of the private sector manifest themselves as parameter drift in the government's model. Following Sargent (1999) and Cho, Williams, and Sargent (2002), we assume the government responds to this drift by placing more emphasis on recent data when updating the parameters of its model.

In our model, the only 'action' taken by the private sector in response to new information is to revise its expectations of next period's exchange rate.[16] Hence, the private sector just needs to formulate a model of the exchange rate. To be consistent with the government's beliefs, we assume it takes the following form:

$$s_t = \gamma_1 x_{t-1} + \gamma_2 i_t + u_{2t} \tag{21}$$

where γ_1 is a 1×3 vector.

If we now use these perceived laws of motion to evaluate the expectations in (13), it can readily be verified that the *actual* law of motion takes the following form:

$$x_t = \Gamma_1(\beta,\gamma)x_{t-1} + \Gamma_2(\beta,\gamma)x_{t-2} + C(\beta,\gamma)v_t \tag{22}$$

In deriving this equation, we assume the government chooses an interest rate policy to minimize (12), given its beliefs in (20). This is a standard LQR optimization problem. The solution yields an optimal feedback rule of the form, $i_t^p = -Fx_{t-1}$, where i_t^p denotes a 'planned' interest rate. The actual market interest rate is assumed to be equal to this planned rate plus an i.i.d shock, which is meant to capture random implementation errors or high frequency velocity or liquidity shocks in the interbank market. Thus, we have:

$$i_t = -F(\beta)x_{t-1} + v_{3t}$$

where the dependence of F on the beliefs of the government is made explicit. Note that because of the extra shock, v_t in (22) is now a 3×1 vector (i.e., $(v_{1t}, v_{2t}, v_{3t})'$).

4.1 Self-Confirming Equilibrium

The misspecification of the government's model prevents it from learning the Rational Expectations Equilibrium derived in section 2.4.[17] Despite this handicap, the government

[16]Note, this is likely a nonrobust feature of our model. For example, if foreign currency debt were *endogenous*, then the evolution of the private sector's beliefs could produce discrete changes in the system via changes in the stock of foreign currency liabilities.

[17]See Evans and Honkapohja (2001) for an extensive discussion of the circumstances under which it is and is not possible for agents to learn a model's Rational Expectations Equilibria.

and the private sector both act purposefully to eliminate systematic forecast errors. They do this by choosing the parameters of their perceived models to best fit the data. Since their models include intercepts, agents will be successful at avoiding systematic forecast errors. So at least in this sense our model is not vulnerable to the kind of criticism that was levelled at the original applications of adaptive expectations in the 1960s. Still, the misspecification does mean that agents can miss the data's higher order moments and, as a consequence, there will in general be 'patterns' in the forecast errors.

Although these patterns could be discovered if the agents were to explore alternative model specifications, we rule out this kind of experimentation. Following Sargent (1999), we adopt a weaker notion of equilibrium, which is well suited to models of boundedly rational learning. This equilibrium concept is called a 'self-confirming equilibrium'. A self-confirming equilibrium is weaker than a Rational Expectations Equilibrium in the sense that it merely requires beliefs to be confirmed 'along the equilibrium path'. Beliefs about 'off equilibrium path' events can be arbitrary. In the context of our model, off-equilibrium path play relates to the willingness of agents to entertain alternative model specifications. The beliefs of our agents will only be optimal within the parametric classes defined by (20) and (21). They will not be optimal in the Rational Expectations sense of fully conforming to the actual data generating process.[18]

If this restricted notion of optimality is defined in terms of minimizing the variance of one-step ahead forecast errors, then the following least squares normal equations characterize a self-confirming equilibrium:

$$E\left\{\begin{pmatrix} x_{t-1} \\ i_t \end{pmatrix}[x_t - \beta_1 x_{t-1} - \beta_2 i_t]\right\} = 0 \tag{23}$$

$$E\left\{\begin{pmatrix} x_{t-1} \\ i_t \end{pmatrix}[s_t - \gamma_1 x_{t-1} - \gamma_2 i_t]\right\} = 0 \tag{24}$$

where the expectations are evaluated using the distribution implied by the true data-generating process in (22). Parameter values that satisfy these equations have the property that agents do not have an incentive to revise the parameters of their models.

We can now state a more precise definition of our equilibrium concept.

Definition: *A* SELF-CONFIRMING EQUILIBRIUM *is a collection of regression coefficients* $(\hat{\beta}_1, \hat{\beta}_2, \hat{\gamma}_1, \hat{\gamma}_2)$ *and an interest rate policy,* $i_t^p = -Fx_{t-1}$ *such that when* x_t *is governed by (22), the regression coefficients satisfy the least-squares orthogonality conditions in (23)-(24).*

A self-confirming equilibrium is characterized in the following proposition.

[18]Evans and Honkapohja (2001) call this concept a Restricted Perceptions Equilibrium. Kirman (1983) is a precursor of this idea.

Proposition 4.1: *Given the perceived law of motion in equation (20), a self-confirming equilibrium is characterized by the following simultaneous nonlinear equations*:

$$\gamma_1 = e_3'\beta_1 \tag{25}$$
$$\gamma_2 = e_3'\beta_2 \tag{26}$$
$$\beta_1 = \Gamma_1(\beta,\gamma) + \beta_2 F(\beta) + \Gamma_2(\beta,\gamma)M_{xx}(1)M_{xx}(0)^{-1} \tag{27}$$
$$\beta_2 = C(\beta,\gamma)e_3 \tag{28}$$

where $M_{xx}(0)$ and $M_{xx}(1)$ are the 3×3 moment matrices Ex_tx_t' and Ex_tx_{t-1}', evaluated using the steady state distribution of x_t implied by (22).[19]

To derive these equations, simply substitute the actual process for x_t (given in (22)) into the least squares normal equations in (23) - (24), and collect terms.

Although these equations are highly nonlinear and must be solved numerically, their interpretation is straightforward. Equations (25) and (26) follow from the requirement that the beliefs of the private sector about the exchange rate be consistent with the government's beliefs. Equation (27) can be recognized as a version of the standard formula for 'omitted variable bias'. That is, the fitted regression coefficients on x_{t-1} equal their true values, $\Gamma_1 + \beta_2 F$, plus a bias term that equals the product of the coefficient of the 'excluded variable' (Γ_2) and the coefficients in a regression of the excluded on the included ($M_{xx}(1)M_{xx}(0)^{-1}$). Using (25) and (26) to substitute out the γ coefficients in (27) and (28), and exploiting the known structure of the matrices, it turns out that calculation of a self-confirming equilibrium can be reduced to a set of seven simultaneous nonlinear equations. Still, general proofs of existence, uniqueness, and stability are algebraically infeasible.

The above definition and characterization of a self-confirming equilibrium were stated in terms of population moments. In practice, agents do not know these moments. They must be estimated from the data. Following Sargent (1999) and Cho, Williams, and Sargent (2002), we assume agents do this via a recursive least squares procedure. Letting $\beta_t' = (\beta_{1t}, \beta_{2t})$ and $z_t' = (x_t', i_t)$ we can write the government's learning algorithm as follows:

$$\beta_t = \beta_{t-1} + a_g R_{gt}^{-1} z_{t-1}(x_t' - z_{t-1}'\beta_{t-1}) \tag{29}$$
$$R_{gt} = R_{g,t-1} + a_g(z_{t-1}z_{t-1}' - R_{g,t-1}) \tag{30}$$

The private sector uses an analogous learning algorithm:

$$\gamma_t = \gamma_{t-1} + a_p R_{pt}^{-1} z_{t-1}(s_t - z_{t-1}'\gamma_{t-1}) \tag{31}$$
$$R_{pt} = R_{p,t-1} + a_p(z_{t-1}z_{t-1}' - R_{p,t-1}) \tag{32}$$

where $\gamma_t' = (\gamma_{1t}, \gamma_{2t})$.

[19] These can be calculated by stacking (22) into a first-order system and then solving a matrix Lyapunov equation.

There are several points to notice about these algorithms. First, it is important to remember that the x_t and s_t processes that appear on the right-hand sides of (29) and (31) are the *true* data generating processes in (22). This makes the model 'self-referential', and complicates the analysis of the learning dynamics. Not only do outcomes affect beliefs via the learning algorithms in (29) - (32), but beliefs feedback to influence the observed data via equation (22). Second, the learning algorithms imply that beliefs are changing each period. However, when solving their optimization and forecasting problems, agents act as if their beliefs will never change. This is of course 'irrational', but we would argue that it is at least as descriptively accurate as assuming agents base current actions on forecasts of their future beliefs. Third, the crucial parameters in these algorithms are the two 'gain' parameters, a_g and a_p. These dictate how responsive beliefs are to new information. In a simple least squares procedure, the gains would decrease to zero at rate t^{-1}, reflecting the fact that each new piece of information adds less and less to the accumulated stock of prior experience. However, as noted in the Introduction, we assume agents pay more attention to recent data. They do this because they suspect the environment is nonstationary. This is accomplished by constraining the gain parameters to be (small) constants. This effectively discounts old data and allows agents to remain alert to potential 'regime changes'. Interestingly, we will see that constant gain learning algorithms can produce in an endogenous and self-confirming manner exactly the kind of instability that they are designed to guard against.

5 Escape Dynamics

Before studying the dynamics of the full model it is instructive to consider a stripped down version that highlights the essential elements. This example reveals an important feature of the escape dynamics; namely, they are determined by the beliefs of the government. While learning by the private sector is important in generating real effects during a crisis, the actual occurrence of crises and their frequency are determined solely by the evolving beliefs of the government.

5.1 An Example

To illustrate how the direction of escape from the self-confirming equilibrium is determined, let us examine a simplified version of the model. This example is virtually identical to the model investigated in Cho, Williams, and Sargent (2002), and is intended to reveal the links between our model and their model.

The example does not have a lagged output term on the right hand side of the open economy Phillips curve, and assumes unanticipated exchange rate changes have contemporaneous rather than lagged effects on output. Thus, the output equation takes the following form:

$$y_t = y_o + \theta(s_t - \beta_t) + v_{1t} \tag{33}$$

where β_t is the expected exchange rate at time t based on information available up to

$t-1$ and v_{1t} is i.i.d. white noise. Instead of perfect foresight, assume that β_t is formed according to the weighted average:

$$\beta_t = \beta_{t-1} + a_p\left(s_{t-1} - \beta_{t-1}\right) \tag{34}$$

where $a_p > 0$ is the gain function of the private sector. This rule produces correct forecasts on average, but misreads the dynamics. For simplicity, let us assume for a moment that the government minimizes the following objective function:

$$E[(y_t - y^*)^2 + (s_t - s^*)^2]$$

by controlling s_t, where the government's model of y_t is based on the misspecified Phillips curve:

$$y_t = \gamma_0 + \gamma_1 s_t. \tag{35}$$

We can parameterize the government's model by $\gamma = (\gamma_0, \gamma_1)$. Given γ, the government sets the exchange rate

$$H(\gamma) = \frac{-\gamma_1(\gamma_0 - y^*) + s^*}{1 + \gamma_1^2} \tag{36}$$

but the realized exchange rate at t is

$$s_t = H(\gamma) + v_{2t}$$

where v_{2t} is i.i.d white noise. We add v_{2t} to capture potential implementation errors.

The government updates the regression coefficients γ according to the recursive fixed gain algorithm:

$$\begin{bmatrix}\gamma_{0,t+1}\\ \gamma_{1,t+1}\end{bmatrix} = \begin{bmatrix}\gamma_{0,t}\\ \gamma_{1,t}\end{bmatrix} + a_g R_t^{-1}\begin{bmatrix}1\\ s_t\end{bmatrix}[y_t - \gamma_{0,t} - \gamma_{1,t}s_t] \tag{37}$$

$$R_{t+1} = R_t + a_g\left(\begin{bmatrix}1\\ s_t\end{bmatrix}[1 \quad s_t] - R_t\right) \tag{38}$$

where a_g is the gain function of the government's adaptive learning algorithm. Let

$$\lambda = \frac{a_p}{a_g}$$

which measures the relative speed of learning by the private sector compared to the government. In the "Fed Watcher" case of Cho, Williams, and Sargent (2002), where the private sector has rational expectations about the government's policy, $\lambda = \infty$.

Stacking up (34) and (37) along with (38) gives us,

$$\begin{bmatrix}\gamma_{0,t+1}\\ \gamma_{1,t+1}\\ \beta_{t+1}\end{bmatrix} = \begin{bmatrix}\gamma_{0,t}\\ \gamma_{1,t}\\ \beta_t\end{bmatrix} + a_g g(\gamma_{0,t}, \gamma_{1,t}, \beta_t, \lambda) + \mu_t \tag{39}$$

where μ_t is a martingale difference, and the associated ODE is

$$\begin{bmatrix} \dot{\gamma}_0 \\ \dot{\gamma}_1 \\ \dot{\beta} \end{bmatrix} = g(\gamma_0, \gamma_1, \beta, \lambda). \tag{40}$$

while the covariance matrix evolves according to (38). For later analysis, it will be convenient to have the explicit formulas for g and μ_t:

$$g(\gamma_{0,t}, \gamma_{1,t}, \beta_t, \lambda) = \begin{bmatrix} R_t^{-1} \begin{bmatrix} y_o + \theta(H(\gamma_t) - \beta_t) - \gamma_{0t} - \gamma_{1t} H(\gamma_t) \\ H(\gamma_t)(y_o + \theta(H(\gamma_t) - \beta_t) - \gamma_{0t} - \gamma_{1t} H(\gamma_t)) + (\theta - \gamma_{1t})\sigma_2^2 \end{bmatrix} \\ \lambda\,(H(\gamma_t) - \beta_t) \end{bmatrix} \tag{41}$$

and

$$\mu_t = \begin{bmatrix} R_t^{-1} \begin{bmatrix} -\theta v_{2t} + v_{1t} - \gamma_{1t} v_{2t} \\ -H(\gamma_t)(\gamma_{1t} - \theta) v_{2t} + H(\gamma_t) v_{1t} + (1 - \gamma_{1t}) v_{1t} v_{2t} + (v_{2t}^2 - \sigma_2^2)(\gamma_{1t} - \theta) \end{bmatrix} \\ \lambda v_{2t} \end{bmatrix}. \tag{42}$$

Under "normal" circumstances,

$$E\mu_t = 0.$$

For a small a_g, a_p, the stochastic process $\{\gamma_t, \beta_t\}_{t \geq 1}$ can be approximated by the trajectory induced by (40). The stationary solution of (40) is precisely the self-confirming equilibrium, and its stability can be easily verified by checking the eigenvalues of the Jacobian of the right hand side of (40).

Because we are investigating escapes from a stable point of the ODE, we need to consider "unusual" events. It turns out to be helpful to understand the escape dynamics if we assume that v_{it} has a binomial distribution with σ_i being realized with probability one half, and $-\sigma_i$ with probability one half. We shall focus on the case where σ_i is small.

The main difference from the model analyzed in Cho, Williams, and Sargent (2002) is that we *in principle* have to figure out the escape of $(\gamma_{0t}, \gamma_{1t}, \beta_t) \in \mathbb{R}^3$, while in Cho, Williams, and Sargent (2002), we only have to calculate the escape path of the first two components. We shall show that it suffices to focus on the escape dynamics of the first two components ($(\gamma_{0t}, \gamma_{1t})$) in our model, because the dynamics of β_t are completely "captured" by the dynamics of $(\gamma_{0t}, \gamma_{1t})$.

Let us consider, for a moment, a closely related, but highly artificial, model in which $(\gamma_{0t}, \gamma_{1t}, \beta_t)$ evolves according to (39) but μ_t is replaced by

$$\tilde{\mu}_t = \begin{bmatrix} R_t^{-1} \begin{bmatrix} \theta v_{2t} + v_{1t} - \gamma_{1t} v_{2t} \\ -H(\gamma_t)(\gamma_{1t} - \theta) v_{2t} + H(\gamma_t) v_{1t} + (1 - \gamma_{1t}) v_{1t} v_{2t} + (v_{2t}^2 - \sigma_2^2)(\gamma_{1t} - \theta) \end{bmatrix} \\ 0 \end{bmatrix}. \tag{43}$$

The only difference from the original recursive formula (39) is that the evolution of β_t is completely determined by the evolution of $\gamma_t = (\gamma_{0t}, \gamma_{1t})$. Conditioned on the history of γ_t, the probability distribution of β_t is degenerate, and therefore, there cannot be any escapes from the trajectory of β_t dictated by

$$\beta_{t+1} = \beta_t + a_g \lambda \left(H(\gamma_t) - \beta_t \right).$$

As a result, to understand the escape dynamics of $(\gamma_{0t}, \gamma_{1t}, \beta_t)$, it suffices to understand the escape dynamics of the first two components of the vector, because the dynamics of β_t are completely determined by the dynamics of $(\gamma_{0t}, \gamma_{1t})$.

We essentially "project" the path of $(\gamma_{0t}, \gamma_{1t}, \beta_t)$ onto the space of $(\gamma_{0t}, \gamma_{1t})$, and investigate the most likely direction of escape of $(\gamma_{0t}, \gamma_{1t})$. Because β_t must remain in the neighborhood of the surface determined by

$$\left\{ (\gamma_0, \gamma_1, \beta) \mid \dot{\beta} = H(\gamma_0, \gamma_1) - \beta \right\}$$

when $a_g, a_p > 0$ are sufficiently small, we can completely identify the escape path, once we identify the most likely escape path of $(\gamma_{0t}, \gamma_{1t})$.

We are now ready to examine the original recursive formula (39), (41) and (42) where the evolution of β_t is subject to the small perturbation v_{2t}. Recall that the key reason why we can focus on the dynamics of the first two components of $(\gamma_{0t}, \gamma_{1t}, \beta_t)$ is that the trajectory of β_t can be recovered from information about the trajectory of the first two components.

Fix t, and $(\gamma_{0t}, \gamma_{1t}, \beta_t)$. It is straightforward to verify that there is one-to-one correspondence between the realized values of (v_{1t}, v_{2t}) and $(\gamma_{0,t+1}, \gamma_{1,t+1})$. Thus, given the initial condition $(\gamma_{00}, \gamma_{10}, \beta_0)$, the sample path of $(\gamma_{0t}, \gamma_{1t})$ reveals the entire sample path of (v_{1t}, v_{2t}). Therefore, conditioned on the sample path of $(\gamma_{0t}, \gamma_{1t})$, and the initial condition $(\gamma_{00}, \gamma_{10}, \beta_0)$, we can recover the sample path of β_t. Formally, the sample path β_t is measurable with respect to the information generated by the sample path of $(\gamma_{0t}, \gamma_{1t})$ and the initial condition $(\gamma_{00}, \gamma_{10}, \beta_0)$. Thus, the distribution of the sample paths of β_t conditional on the sample path of $(\gamma_{0t}, \gamma_{1t})$ and initial condition $(\gamma_{00}, \gamma_{10}, \beta_0)$ is degenerate, and therefore, there is no possibility of escape from the path built from the information available from the evolution of $(\gamma_{0t}, \gamma_{1t})$. Therefore, it suffices to focus on the escape dynamics of the first two components of $(\gamma_{0t}, \gamma_{1t}, \beta_t)$.

Because we can focus on the first two components *without loss* of generality, we can invoke the analysis of Cho, Williams, and Sargent (2002) which corresponds to our model with

$$\lambda = \infty$$

so that the private sector can maintain

$$H(\gamma_t) = \beta_t \qquad \forall t \geq 1.$$

Fix a small neighborhood around the stable stationary solution in $\mathbb{R}^2$, and consider a probability distribution of the first exit points along the boundary. Because this probability distribution is hard to examine, we construct a new probability distribution of exit points, which is absolutely continuous with respect to the original probability distribution. Thus, if the original probability distribution is concentrated on a single point (as it should be according to standard results from large deviation theory), then so is the new probability measure.

The new probability distribution is constructed by focusing on the shortest escape path to a small neighborhood of the boundary. In Cho, Williams, and Sargent (2002), it suffices to show that we can focus on the escape path induced by a series of perturbations (v_{1t}, v_{2t}) which has exactly two values.[20] We found that among many possibilities, if $(v_{1t}, v_{2t}) \in \{(\sigma_1, \sigma_2), (-\sigma_1, -\sigma_2)\}$ along the escape path, then this path determines the most likely direction of escape from the stable solution of the ODE (or the self-confirming equilibrium). Moreover, as $a_g \to 0$, the set of escape paths collapses to the most likely path in a probabilistic sense.

A standard result from large deviation theory implies that as $a_p, a_g \to 0$, the set of escape paths collapses to this direction. A simple numerical analysis reveals that along the most likely escape path, the velocity vector of (γ_0, γ_1) can be very large, which explains the seemingly discontinuous jump from the self-confirming equilibrium to the Ramsey outcome in Cho, Williams, and Sargent (2002).

If $\lambda = \infty$, then the private sector's belief is instantly adjusted to the rapidly changing target set by the government, as its belief is changing rapidly along the dominant escape path. This instantaneous adjustment of the private sector's beliefs rules out real effects during the crisis.

However, for any $\lambda < \infty$, the velocity vector of β is uniformly bounded, because γ and β are uniformly bounded. Since γ is moving rapidly along the escape path, so is $H(\gamma)$. As a result, the gap between $H(\gamma)$ and β increases in the initial phase of the escape. Because the private sector initially underestimates the government's choice of the exchange rate, there is a real effect during the initial phase of escape: output falls during the exchange rate crisis. However, even along the escape path, the private sector's belief is adjusting to the target. Toward the end of the escape phase, the velocity of escape decreases, and the private sector can learn the target quite accurately. This increased accuracy at the end of the escape phase explains why the real effect tapers off toward the end of the exchange rate crisis.

On the other hand, during the convergent path to the self-confirming equilibrium, the adaptive learning scheme of the private sector can predict the government's policy quite well, because the government's belief $(\gamma_{0t}, \gamma_{1t})$ moves "slowly". Under "normal" circumstances, output stays around y_o. Systematic deviations can occur if and only if the economy is in the escape phase, which occurs rarely but recurrently, in our economy.

[20]Because each perturbation has a binomial distribution, (v_{1t}, v_{2t}) has four different realizations, each of which has probability one quarter. Thus, the sequence of perturbations that has precisely two different realizations of (v_{1t}, v_{2t}) is an unusual event.

6 Simulations

The relative complexity of our model has prevented us so far from obtaining an analytical characterization of the escape dynamics. For example, the presence of lags and additional state variables makes our model substantially more complicated than the models of Sargent (1999) and Cho, Williams, and Sargent (2002). Fortunately, the model is easy to simulate. We simply solve the nonlinear equations in (25)-(28) characterizing a self-confirming equilibrium to obtain starting values, then run the recursive constant gain stochastic approximation algorithms in (29)-(32) to generate time paths of beliefs and outcomes. This section presents results from these simulations.

6.1 Parameter Settings

At this stage we are not in a position to calibrate the model formally to observed currency crisis data. Instead, our goal is to replicate the qualitative comovements while staying in the ballpark on magnitudes.

It turns out that three parameters are especially important in governing the qualitative properties of the dynamics. The first is the gain parameter(s), (a_g, a_p), which governs the 'speed of learning'. The second is the risk premium, ϕ, which influences how interest rates react during a crisis. The third is the Phillips Curve slope, θ, which influences how output responds during a crisis. The remaining parameters primarily determine the scale of the data. Hence, we focus on a_g, a_g, ϕ, and θ.

We have some idea what θ should be given the fact that it's the coefficient on unanticipated depreciations in the log-linearized Phillips Curve in (7). Assuming a linear AK-style production function, a log-linearization of (6) reveals that $\rho = (1+\bar{\mu})(1-\alpha)\kappa$ and $\theta = -(1+\bar{\mu})(1-\alpha)\kappa(1+i^*)(\bar{d}^*/\bar{y})$, where κ is the output/capital ratio. This suggests that given a value for ρ, we can determine a reasonable value for θ from observed foreign currency debt/GDP ratios. For the simulations I simply set $\rho = 0.9$, which then suggests values for θ in the range .1 to .3. This follows from the exposure estimates of Burnside, Eichenbaum, and Rebelo (2001) and ABB (2001a), which indicate foreign debt/GDP ratios in the range of 10-35% in recent crisis countries.

The other two parameters, the gain and the risk premium, are harder to get a handle on. The risk premium parameter, ϕ, is disciplined by observations on how spreads respond to the state of the economy. A value of $\phi = 0.2$ implies that when GDP falls 10% below its steady state then domestic interest rates rise two percentage points above foreign interest rates. This seems in the ballpark, so we use this as a benchmark. However, we'll see that currency crisis data may call for a somewhat larger value.

As noted in Sargent (1999) and Cho, Williams, and Sargent (2002), the primary role of the gain is to determine the frequency of escapes. Escapes occur more frequently as the gain increases. However, the gain also determines the underlying unit of time and the rate at which information arrives and decisions are revised. Hence, it is not the case that the calendar date arrivals of currency crises can be freely altered by selecting the gain. We simply regard a_g and a_p as free parameters, and compare the results with alternative values.

Finally, the remaining parameters are selected with three considerations in mind. First, the steady state value of output and the target exchange rate are set arbitrarily to 6.0. This was done simply because the interest rate is around zero, and there are numerical advantages to having all the data be of the same order of magnitude. Second, the target output is set slightly above the steady state, which gives the government a persistent desire to stimulate the economy. This is a standard asssumption in these models. Third, the objective function weights were chosen so that output fluctuations are twice as costly as exchange rate and interest rate fluctuations. This seems reasonable, but absent a more rigorous derivation of the government's objectives, it is obviously somewhat arbitrary.

Table 1 contains a complete listing of the parameter settings for the various simulations.

6.2 Interpreting the Simulations

Since we are interested in sample paths rather than moments, we felt it was more informative to present plots of 'representative' time paths for a few key series. Of course, the drawback to this strategy is that it can be difficult to separate sampling variability from the model's underlying dynamics. However, the general character of the simulations, and in particular the magnitude and frequency of escapes are remarkably similar across simulation runs.[21]

Figures 2-4 report results from the benchmark parameterization, which has equal gain parameters, $a_p = a_g = .04$, and a relatively modest risk premium, $\phi = 0.20$. These simulations contain several interesting features. First, notice the mean dynamics feature a gradual (real) appreciation of the exchange rate. As noted in section 2, this is commonly observed in the data. It results here from the fact that the Nash exchange rate is below (i.e., stronger) than the Ramsey exchange rate. Second, notice that when a crisis erupts the exchange rate jumps from near its self-confirming value of around 1.0 to its target value of 6.0. This is a much larger depreciation than the 'average' crisis, as reported in section 2, but certainly not without precedent. We suspect that we could limit the severity of a crisis by altering the target exchange rate, s^*. Third, notice that a currency crisis triggers a severe recession, with output falling by about 10-15%. This is encouraging, since explaining this was one of our main goals. Finally, notice that these simulations reveal one major discrepancy with the data; namely, interest rates *fall* during a crisis and remain below their original level for some time.

Why does the model generate lower interest rates during a crisis? We conjecture that it reflects a situation in which the government "loses control" of an attempted mild depreciation. In the model, because the target output level exceeds the 'natural rate', there is a sense in which the government is looking for an opportunity to stimulate the

[21] The simulations were performed in MATLAB using the control and optimization routines **fsolve**, **dlyap**, and **dlqr**. The only delicate aspect of the simulations concerns instances where the parameters drift into an 'unstabilizable' region of the parameter space. This causes the LQR solver to bomb. We could circumvent this by incorporating a 'projection facility'. However, since the program only takes about 30 seconds to run on a Pentium IV, we just start over if this happens. The program is available upon request.

TABLE 1

PARAMETER SETTINGS FOR THE SIMULATIONS

	Figs 2-4	Figs 5a,b	Fig 6
δ	.95	.95	.95
ρ	.90	.90	.90
θ	$-.30$	$-.30$	$-.30$
y^{ss}	6.0	6.0	6.0
ϕ	.20	.40	.20
ω_s	0.5	0.5	0.5
ω_i	0.5	0.5	0.5
ω_y	1.0	1.0	1.0
y^*	8.0	8.0	8.0
s^*	6.0	6.0	6.0
σ_1^2	.005	.005	.005
σ_2^2	.02	.02	.02
σ_3^2	.0001	.0001	.0001
a_g	.04	.04	.08
a_p	.04	.04	.08
Self-Confirming Equilibrium			
y^{sce}	6.17	5.96	6.17
s^{sce}	0.74	2.51	0.74

Notes: (1) $\bar{y} = y^{ss}(1-\rho)$.

Figure 2

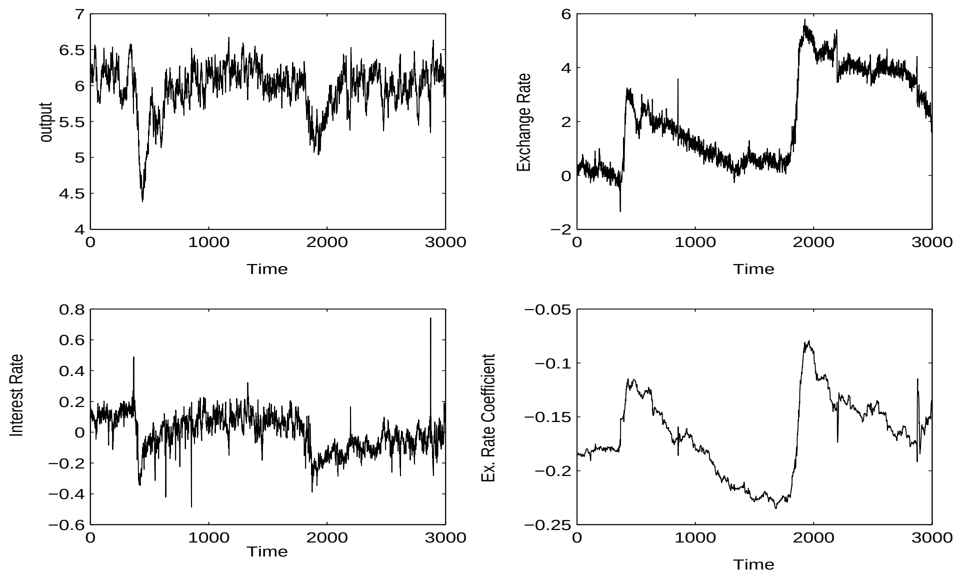

Figure 3

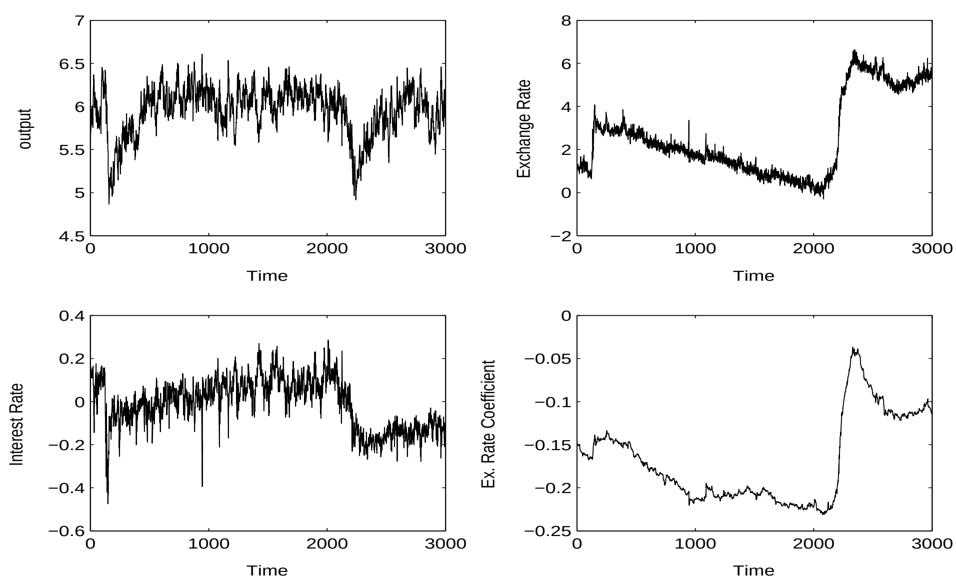

Figure 4a

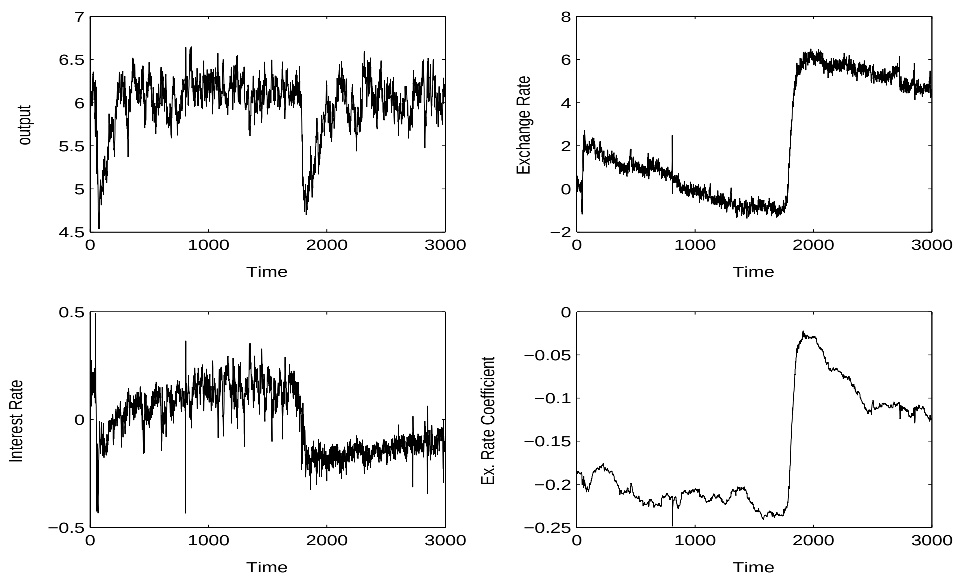

Figure 4b: Anatomy of a Currency Crisis

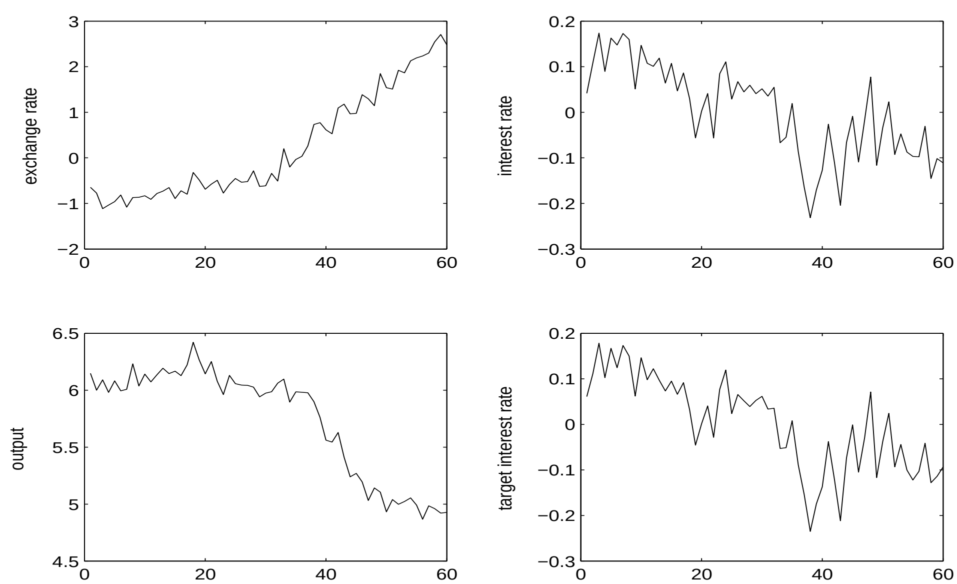

economy. Although it's aware of the potential for interest rate cuts to be contractionary, due to the presence of foreign debt, it doesn't know exactly what the economy's exposure is. It's evolving beliefs about this exposure are summarized by the reported plots of the exchange rate coefficients in the output equation. It is interesting to note that crises always occur when estimates of this coefficient drift up toward zero, away from the self-confirming value of around -0.25. This signals to the government that it is a good time to cut interest rates. From the Uncovered Interest Parity condition, along with the fact that expectations are adaptive and sluggish, the interest rate cut then triggers a currency depreciation. Unfortunately, the true value of the exchange rate coeffcient is more negative than the government believed, and output actually falls. This then leads the government to further cut the interest rate. At this point the government gets trapped into a vicious circle of currency depreciation and falling output, with each successive decline in output leading the government to try to offset it by further interest rate cuts.

The high frequency interactions between the interest rate, the exchange rate, and output are highlighted in Figures 4a and 4b, which provide a 'close-up' of an escape. In particular, Figure 4b plots observations 1751-1810, which include the early phase of the crisis revealed in Figure 4a. One can see immediately that crises are not literally jumps. The exchange rate moves rapidly but continuously from its initial value of around zero to a value around 3. If we take the underlying time unit to be a day, then the latter half of the figure indicates that the price of foreign currency is doubling every 10 days or so. Although there are instances where observed exchange rates literally jump overnight to a new higher level, this invariably happens in formalized pegged exchange rate regimes. Our model is not a model of a formal fixed exchange rate regime. The government cares about exchange rate fluctuations, but it also cares about other things.[22]

Although it's still a bit difficult to discern the timing, Figure 4b suggests that the interest rate starts to decline at about the same time, or slightly before, the exchange rate starts its rapid depreciation. Note also that the interest rate reduction seems to be a deliberate act of policy, not a mistake. This is revealed by the comparison of the actual interest rate with the target rate below it. Notice the two are virtually identical.

Perhaps the most troubling aspect of Figures 4a and 4b is the fact that the interest rate remains below its initial value for an extended period of time. There would appear to be few empirical counterparts of this feature, although Britain's departure from the EMS following the September 1992 attack is widely interpreted as allowing Britain to embark on a lower interest rate policy. However, Figures 5a and 5b reveal that this prediction of the model is sensitive to the value of the risk premium coefficient, ϕ. If we double ϕ to 0.4, so that a 10% reduction in output leads to a 4 percentage point increase in domestic interest rates, then we see that in response to the unfolding crisis the government lets the interest rate return more quickly to its pre-crisis level. Notice, however, that we still do not observe rates rising *above* their pre-crisis level. (Observe that the apparently slower evolution of the crisis in Figure 5 relative to Figure 4 is a figment of the different scalings in the two plots. In particular, Figure 5 only contains 1000 observations).

Finally, the last thing we examine in the simulations is the frequency of the escapes. Figures 2 - 5 suggest that crises occur approximately once every 1500 - 2000 periods. This suggests an interarrival time of about 6 to 8 years if the underlying time unit is a day, or about 30 to 40 years if the underlying time unit is a week (with 5 trading days per week). If we want to increase the frequency of escapes we just need to increase the gain

[22]Interestingly, the evidence in Tornell and Westermann (2001) suggests that many of the stylized facts of currency crises are independent of the formal 'regime' in place at the time.

Figure 5a: Bigger Risk Premium

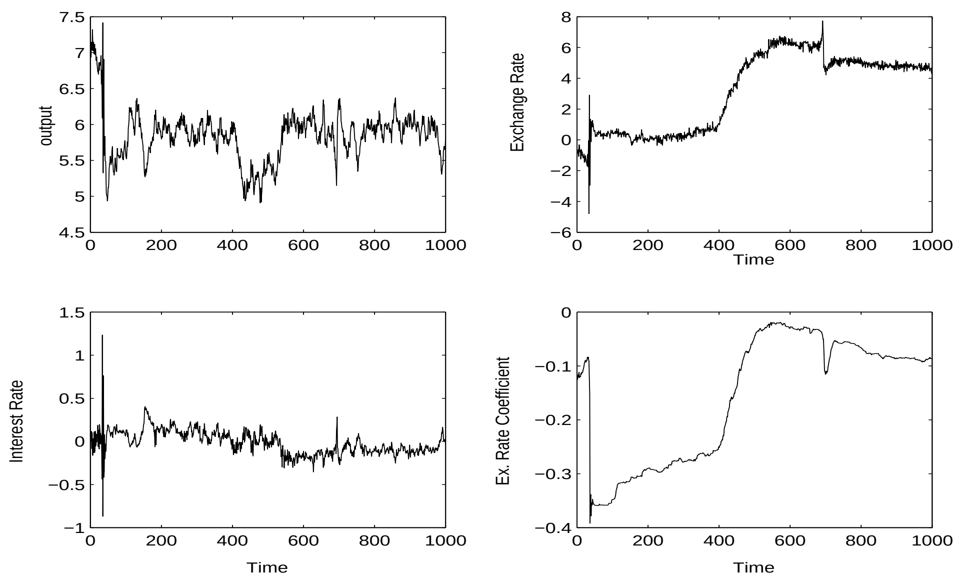

Figure 5b: A Closer Look at Observations 370-470

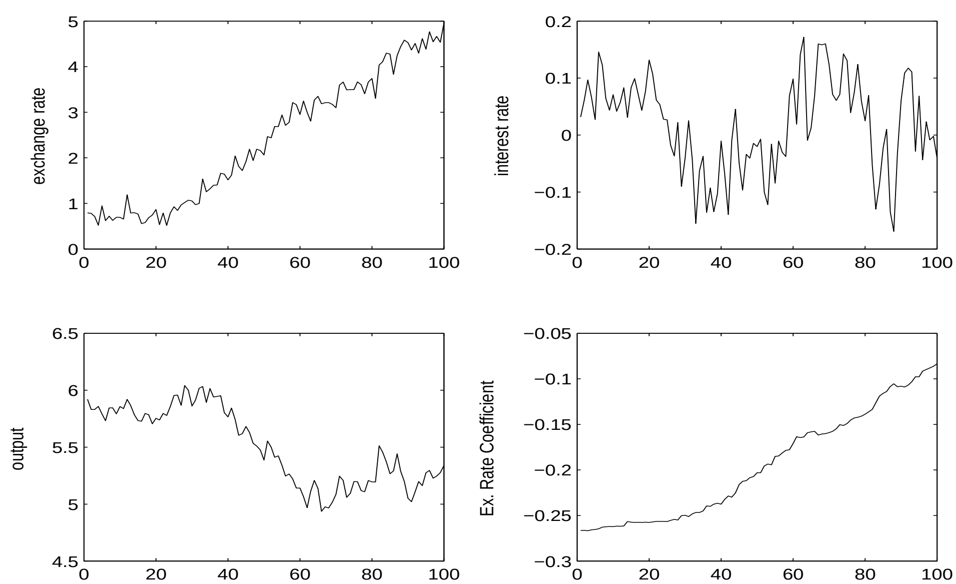

parameters of the learning algorithms.[23] Figure 6 shows what happens when we double the gain parameters to .08. Now a crisis erupts once every 800 periods or so, or roughly twice as often.[24] However, this is only happening because we assume agents are discounting data at a more rapid rate. With a gain parameter of 0.04, data that are 17 weeks old receive only half the weight of current data. Increasing the gain to 0.08 implies that this 'half-life' of data relevance falls to about 8 weeks.

Figure 6: Higher Gain Parameters

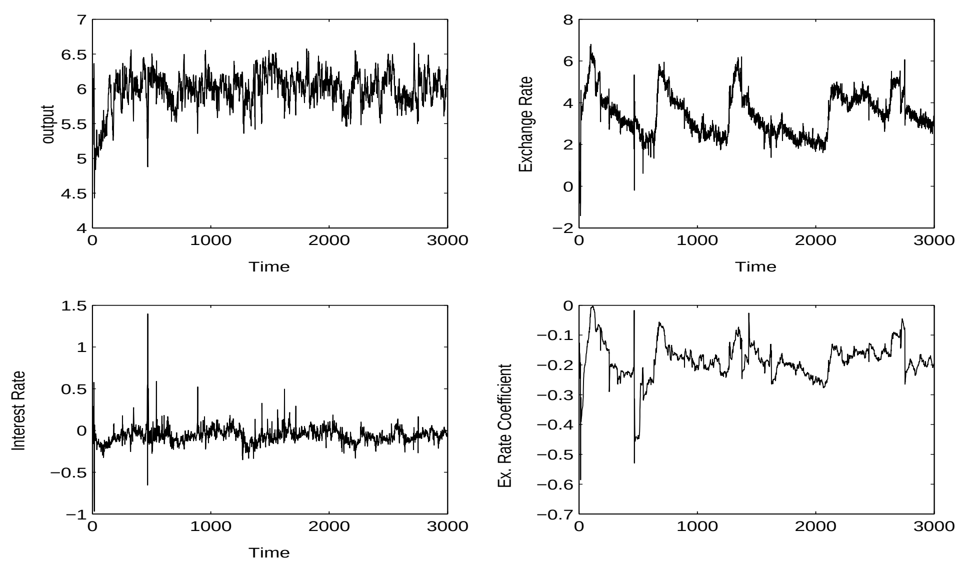

[23]We could also increase the variance of the shocks. This strategy is disciplined by the observed variances of the underlying series, however.

[24]Warning: One cannot extrapolate this to conclude that, for example, tripling the gain would cause crises to happen three times as fast. Large deviations theory predicts a nonlinear relationship between the gain and the mean escape times.

7 Conclusion

This paper has extended the third-generation currency crisis literature by modeling the high-frequency dynamics of currency crises. We do this by explicitly modeling the evolution of beliefs. We share the opinion with other contributors to this literature that beliefs lie at the heart of currency crises. However, rather than regarding these beliefs as responding in an implausibly coordinated way to exogenous sunspots, we regard beliefs as responding adaptively to recent experience. Despite the adaptive nature of beliefs, we show that the nonlinearity induced by self-referential feedback between beliefs and outcomes can produce what looks like 'switches' between multiple equilibria. However, in our model, currency crises are *not* switches between multiple equilibria. They reflect the 'escape dynamics' of a unique equilibrium stochastic process.[25]

Our model attributes currency crises to government miscalculation. Crises occur when the government underestimates the contractionary effects of currency depreciation. Unanticipated depreciations are contractionary due to the presence of unhedged foreign currency debt and its adverse balance sheet effects. We assume the government is unsure about the economy's exposure to these balance sheet effects, and must revise its beliefs about them recursively as it witnesses the economy's response to its interest rate policy. In close analogy to Sargent (1999) and Cho, Williams, and Sargent (2002), it occasionally but infrequently happens that a sequence of shocks occurs that leads the government to reduce its estimate of the output effects of nominal exchange rate changes. This happens when the shocks to the output equation and the Uncovered Interest Parity equation move in the same direction. The induced positive covariance between output and the exchange rate then produces an upward revision (i.e., toward zero) of the government's estimate of the output effects of depreciation. This signals to the government that it is safe to cut interest rates. Unfortunately, once the exchange rate starts to depreciate, the model's feedback causes the govenment to lose control of the situation.

Our model departs from Sargent and CWS in one important respect, which reverses the welfare effects of the escape dynamics. In Sargent and CWS, escapes are a *good thing*. Escapes occur when the government learns a version of the Natural Rate Hypothesis, which leads the government to reduce the inflation rate. This is unambiguously a good thing in their model since they assume the private sector has Rational Expectations, and hence, there are no output effects from inflation stabilization.

In contrast, we do not assume the private sector has Rational Expectations. Like the government, agents in the private sector must update their beliefs adaptively. This means that once the currency starts to depreciate, expectations of future depreciation are ignited. When this happens the government gets sucked into what Chari, Christiano, and Eichenbaum (1998) call an 'expectations trap'. If the government holds the line on interest rates then expectations of future depreciation produce depreciation now. In fact, with falling output, the presence of a risk premium in the Uncovered Interest Parity condition implies that the depreciation will *exceed* expectations, thus engendering a further decline in output. Our simulations suggest that initially the government actually adds fuel to the fire by cutting interest rates. They do this in a (misguided) effort to stabilize output. We also found that if the risk premium is large enough, the government will allow rates to rise shortly after the outbreak of the crisis. Without a strong risk premium, however, interest

[25]This distinction is similar to a point made by Glaeser and Scheinkman (2001) in their review of the social interactions literature. They point out that what may appear to be jumps between multiple equilibria may in fact be the result of a large 'social multiplier'.

rates stay (counterfactually) lower until the government's beliefs about the exchange rate return to the self-confirming equilibrium. At this point, the government correctly believes that a depreciation will be contractionary.

Although we feel our model successively and plausibly describes many of the stylized facts of currency crises, there are three directions in which we would like to extend our analysis. The first is to provide a more complete analytical characterization of the escape dynamics. Without this, our interpretations of the simulations will remain rather speculative. In principle, we could do this by applying the recent results of Williams (2001). This would produce a (high-dimensional) calculus of variations problem characterizing the escape path and the mean escape time. However, our model features two-sided learning, and this leads to some additional subtleties. For example, in our model there are more coefficients than underlying shocks. We suggested in section 5 that this means the escape dynamics are not as difficult to characterize as you might think. The stochastic singularity leads to a dimensionality reduction. Working this out in the full dynamic model is the subject of ongoing research. We conjecture that the singularity will manifest itself as a noninvertible weighting matrix in the quadratic form derived by Williams. Presumably this can be handled by taking some sort of generalized inverse.

A second avenue for future research would be to run more simulations with more parameter settings. Clearly, with more than a dozen parameters this could quickly degenerate into a mindless fishing expedition. However, it would be interesting to see whether the model can produce escape dynamics featuring rising interest rates before a crisis. One possibility would be to consider alternative objective function weights. In our opinion, however, the model's apparent failure to generate higher rates before a crisis should not be construed as a major flaw. After all, Obstfeld and Rogoff (1995, p. 86) pointed to the failure of interest rates to signal crises as a puzzle. Not only can we explain the failure of interest rates to rise before a crisis, we can actually generate crises accompanied by *lower* interest rates! Still, this property of the model suggests that our analysis is best suited to events like the Asian crisis, where fiscal deficits did not seem to be a factor (although see Burnside, Eichenbaum, and Rebelo (2001)).

Third, and perhaps most importantly, we would like to endogenize the foreign debt decision. Doing this would provide an interesting bridge to the analysis of Christiano, Gust, and Roldos (2002). They also study monetary policy during a crisis. However, their analysis is in some respects the mirror image of ours. They regard the crisis itself as an exogenous event, but allow foreign debt to be endogenous. They also find that interest rate reductions in response to a crisis can make matters worse.

REFERENCES

Aghion, Phillipe, Phillipe Bacchetta, and Abhijit Banerjee, 2000, A Simple Model of Monetary Policy and Currency Crises, *European Economic Review* 44, 728-38.

____________, 2001a, Currency Crises and Monetary Policy in an Economy with Credit Constraints, *European Economic Review* 45, 1121-50.

____________, 2001b, A Corporate Balance-Sheet Approach to Currency Crises, mimeo, University of Lausanne.

Baxter, Marianne, 1985, The Role of Expectations in Stabilization Policy, *Journal of Monetary Economics* 15, 343-62.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2000, On the Fundamentals of Self-Fulfilling Speculative Attacks, NBER Working Paper No. 7554.

____________, 2001, Hedging and Financial Fragility in Fixed Exchange Rate Regimes, *European Economic Review* 45, 1151-94.

Caballero, Ricardo J. and Arvind Krishnamurthy, 2001, International and Domestic Collateral Constraints in a Model of Emerging Market Crises, *Journal of Monetary Economics* 48, 513-48.

Cespedes, Luis Felipe, Roberto Chang and Andres Velasco, 2001, Balance Sheets and Exchange Rate Policy, mimeo, Harvard University.

Chamberlain, Gary, and Charles A. Wilson, 2000, Optimal Intertemporal Consumption under Uncertainty, *Review of Economic Dynamics* 3, 365-95.

Chari, V.V., Lawrence J. Christiano, and Martin Eichenbaum, 1998, Expectation Traps and Discretion, *Journal of Economic Theory* 81, 462-92.

Cho, In-Koo, Noah Williams, and Thomas J. Sargent, 2002, Escaping Nash Inflation, *Review of Economic Studies* 69, 1-40.

Christiano, Lawrence J., Christopher Gust, and Jorge Roldos, 2002, Monetary Policy in a Financial Crisis, mimeo, Northwestern Univ.

Dooley, Michael, 2000, A Model of Crises in Emerging Markets, *Economic Journal* 110, 256-72.

Dupuis, Paul and Harold J. Kushner, 1985, Stochastic Approximation via Large Deviations: Asymptotic Properties, *SIAM Journal of Control and Optimization* 23, 675-96.

____________, 1989, Stochastic Approximation and Large Deviations: Upper Bounds and w.p.1 Convergence, *SIAM Journal of Control and Optimization* 27, 1108-35.

Evans, George W. and Seppo Honkapohja, 2001, *Learning and Expectations in Macroeconomics*, Princetion Univ. Press.

Flood, Robert P. and Nancy Marion, 1999, Perspectives on the Recent Currency Crisis Literature, *International Journal of Finance and Economics* 4, 1-26.

Freidlin, M.I. and A.D. Wentzell, 1998, *Random Perturbations of Dynamical Systems*, Second Edition, Springer-Verlag.

Glaeser, Edward L., and Jose A. Scheinkman, 2001, Measuring Social Interactions, in *Social Dynamics*, (eds.) Steven N. Durlauf and H. Peyton Young, MIT Press.

Heymann, Daniel, and Axel Leijonhufvud, 1995, *High Inflation*, Oxford Univ. Press.

Ireland, Peter N., 1999, Does the Time-Consistency Problem Explain the Behavior of Inflation in the United States?, *Journal of Monetary Economics* 44, 279-91.

Jeanne, Olivier, 2000, Currency Crises: A Perspective on Recent Theoretical Developments, *Special Papers in International Economics No. 20*, International Finance Section, Princeton Univ.

_____________, 2000, Foreign Currency Debt and the Global Financial Architecture, *European Economic Review* 44, 719-27.

Kandori, Michihiro, George Mailath, and Rafael Rob, 1993, Learning, Mutation, and Long-Run Equilibria in Games, *Econometrica* 61, 29-56.

Kasa, Kenneth, 2001, Learning, Large Deviations, and Recurrent Currency Crises, forthcoming, *International Economic Review*.

Kehoe, Timothy J., and David K. Levine, 1993, Debt-Constrained Asset Markets, *Review of Economic Studies* 60, 865-88.

Kifer, Yuri, 1990, A Discrete-Time Version of the Wentzell-Freidlin Theory, *Annals of Probability* 18, 1676-92.

Kirman, Alan, 1983, On Mistaken Beliefs and Resultant Equilibria, in *Individual Forecasting and Aggregate Outcomes*, (eds.) Roman Frydman and Edmund S. Phelps, Cambridge Univ. Press.

Krugman, Paul, 1996, Are Currency Crises Self-Fulfilling?, *NBER Macroeconomics Annual*, MIT Press.

_____________, 1999, Balance Sheets, The Transfer Problem, and Financial Crises, in P. Isard, A. Razin, and A. Rose (eds.), *International Finance and Financial Crises: Essays in Honor of Robert P. Flood*, Kluwer, Dordrecht.

Kushner, Harold J. and G. George Yin, 1997, *Stochastic Approximation Algorithms and Applications*, Springer-Verlag.

Kydland, Finn E., and Edward C. Prescott, 1977, Rules Rather Than Discretion: The Inconsistency of Optimal Plans, *Journal of Political Economy* 85, 473-93.

_____________, 1980, Dynamic Optimal Taxation, Rational Expectations and Optimal Control, *Journal of Economic Dynamics and Control* 2, 79-91.

Lucas, Robert E., Jr., and Thomas J. Sargent (eds.), 1981, *Rational Expectations and Econometric Practice*, Univ. of Minnesota Press.

Marcet, Albert and Ramon Marimon, 1999, Recursive Contracts, mimeo, Univ. of Pompeu Fabra.

Marcet, Albert and Thomas J. Sargent, 1989, Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models, *Journal of Economic Theory* 48, 337-68.

Morris, Stephen and Hyun Song Shin, 1998, Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks, *American Economic Review* 88, 587-97.

Obstfeld, Maurice, 1994, The Logic of Currency Crises, *Cahiers Economiques et Monetaires* 43, 189-213.

______________, 1997, Destabilizing Effects of Exchange Rate Escape Clauses, *Journal of International Economics* 43, 61-77.

Obstfeld, Maurice, and Kenneth Rogoff, 1995, The Mirage of Fixed Exchange Rates, *Journal of Economic Perspectives* 9, 73-96.

Sargent, Thomas J., 1993, *Bounded Rationality in Macroeconomics*, Oxford University Press.

______________, 1999, *The Conquest of American Inflation*, Princetion University Press.

Schneider, Martin, and Aaron Tornell, 2000, Balance Sheet Effects, Bailout Guarantees and Financial Crises, NBER Working Paper No. 8060.

Tornell, Aaron, and Frank Westermann, 2001, Boom-Bust Cycles: Facts and Explanation, mimeo, UCLA.

Williams, Noah, 2001, Escape Dynamics in Learning Models, Ph.D. Dissertation, University of Chicago.

Woodford, Michael, 1990, Learning to Believe in Sunspots, *Econometrica* 58, 277-307.

______________, 1999, Optimal Monetary Policy Inertia, NBER Working Paper No. 7261.