

Can Perpetual Learning Explain the Forward Premium Puzzle?*

Avik Chakraborty George W. Evans
University of Tennessee University of Oregon

June 30, 2006. Revised August 28, 2006.

Abstract

Under rational expectations and risk neutrality the linear projection of exchange rate change on the forward premium has a unit coefficient. However, empirical estimates of this coefficient are significantly less than one and often negative. We investigate whether replacing rational expectations by discounted least squares (or “perpetual”) learning can explain the result. We calculate the asymptotic bias under perpetual learning and show that there is a negative bias that becomes strongest when the fundamentals are strongly persistent, i.e. close to a random walk. Simulations confirm that perpetual learning is potentially able to explain the forward premium puzzle.

JEL classifications: D83, D84, F31, G12, G15

Keywords: Learning, exchange rates, forward premium.

1 Introduction

The ‘Forward Premium Puzzle’ is a long-standing empirical paradox in international finance. The puzzle refers to the finding that the forward exchange rate consistently predicts the expected depreciation in the spot exchange

*Support from National Science Foundation Grant No. SES-0617859 is gratefully acknowledged. We are indebted to Stephen Haynes and Joe Stone, for many helpful discussions and comments, and for comments received at the “Learning Week” workshop held at the Federal Reserve Bank of St. Louis in July, 2006.

rate but with a smaller magnitude and often the opposite sign than specified by rational expectations. A large literature documents and attempts to explain the puzzle, but mostly with very mixed success. This paper proposes a resolution from a new perspective.

According to theory, if the future rate of depreciation in the exchange rate is regressed on the forward premium (the forward rate less the current spot rate in logarithms), then the slope coefficient on the forward premium should be unity provided the agents are risk-neutral and do not make systematic errors in their forecast. More formally, if s_t is the natural log of the current spot exchange rate (defined as the domestic price of foreign exchange), Δs_{t+1} is the depreciation of the natural log of the spot exchange rate from period t to $t + 1$, i.e., $\Delta s_{t+1} = s_{t+1} - s_t$, and F_t is the natural log of the one-period forward rate at period t , then in the true regression equation

$$\Delta s_{t+1} = \alpha + \beta(F_t - s_t) + u_{t+1}, \quad (1)$$

β is unity and u_{t+1} is uncorrelated with the forward premium $F_t - s_t$. It follows that $E\hat{\beta} = 1$, where $\hat{\beta}$ is the least squares estimate of the slope coefficient on the forward premium.

This theoretical result is based on assumptions of risk-neutrality and rational expectations. If agents are risk neutral then they must set today's forward rate equal to their expectation about the future spot rate, i.e. $F_t = \hat{E}_t s_{t+1}$, where $\hat{E}_t s_{t+1}$ denotes their expectation of s_{t+1} formed at time t . If, moreover, their expectations are rational then $\hat{E}_t s_{t+1} = E_t s_{t+1}$, where $E_t s_{t+1}$ denotes the true mathematical expectation of s_{t+1} conditioned on information available at time t , assumed to include F_t and s_t . With rational expectations, agents' forecast errors $u_{t+1} = s_{t+1} - E_t s_{t+1}$ satisfy $E_t u_{t+1} = 0$, i.e. agents do not make systematic forecasting errors. Combining risk neutrality and rational expectations we obtain

$$s_{t+1} = F_t + u_{t+1},$$

and thus the depreciation of exchange rate from t to $t + 1$ is given by

$$\Delta s_{t+1} = (F_t - s_t) + u_{t+1}$$

where $E_t u_{t+1} = 0$, which gives the theoretical prediction $E\hat{\beta} = 1$.

A large volume of research has empirically tested the hypothesis $\beta = 1$, and concluded that the least squares estimate $\hat{\beta}$ is significantly less than

1. In fact, in the majority of cases, $\hat{\beta}$ is less than zero.¹ We reproduce part of Table 1 from Mark and Wu (1998) documenting the existence of the puzzle. In the table they used quarterly data ranging from 1976.I to 1994.I on USD (dollar) rates of GBP (pound), DEM (deutsche-mark) and JAY (yen) as well as three cross rates². The evidence thus strongly refutes the theoretical prediction that $\beta = 1$, and hence apparently contradicts the efficient market hypothesis. This is the much renowned “forward premium puzzle” (or “forward premium anomaly”).

Table 1
Regressions of Quarterly Depreciation on 3-Month Forward Premium

$$\Delta s_{t+1} = \alpha + \beta(F_t - s_t) + \varepsilon_{t+1}$$

	USD/GBP	USD/DEM	USD/JAY	GBP/DEM	GBP/JAY	DEM/JAY
	1976:I-1994:I					
$\hat{\alpha}_{OLS}$	-1.340 (0.895)	0.638 (0.886)	3.294 (0.964)	1.622 (1.116)	7.702 (1.687)	1.041 (0.648)
$\hat{\beta}_{OLS}$	-1.552 (0.863)	-0.136 (0.839)	-2.526 (0.903)	-0.602 (0.782)	-4.261 (1.133)	-0.755 (1.042)

The key to the resolution of the puzzle seems to be hidden in the ordinary least squares formula for $\hat{\beta}$. Assuming $\beta = 1$ we have

$$\hat{\beta} = \frac{\widehat{\text{cov}}(\Delta s_{t+1}, F_t - s_t)}{\widehat{\text{var}}(F_t - s_t)} = 1 + \frac{\widehat{\text{cov}}[(F_t - s_t), u_{t+1}]}{\widehat{\text{var}}(F_t - s_t)},$$

where $\widehat{\text{cov}}$ and $\widehat{\text{var}}$ denote sample covariance and sample variance. Therefore, $\widehat{\text{cov}}[(F_t - s_t), u_{t+1}] < 0$ is needed to explain the downward bias in $\hat{\beta}$.

Existing research follows two major approaches. One of them assumes that investors in the foreign exchange market are risk-averse. Consequently, the forward rate not only incorporates their expectation about the future depreciation but also includes a risk-premium as a hedge against the risk from investing in a more volatile asset characterized by a higher rate of return. As a result, expected depreciation is not a conditionally unbiased forecast of actual depreciation. Despite its intuitive appeal, empirical studies have shown the difficulty of the risk premium approach in providing a satisfactory

¹Froot and Thaler (1990) and Engel (1996) provide comprehensive reviews of this puzzling observation.

²For more details about the data see Mark and Wu (1998).

explanation of the puzzle.³ This has led to a general skepticism of the risk-premium explanation.

The other approach centers around the potential ability of non-rational expectations to explain the results. This potential is apparent from some of the other findings related to exchange rate behavior.⁴ Our paper is motivated by this research, which suggests the importance of deviations from rational expectations in foreign exchange markets. If traders do not have perfectly rational expectations, their forecast errors may be correlated with previous period's information and this would introduce an observed bias in the forward premium regression results.⁵ The question we want to examine is whether a natural form of bounded rationality would yield $\text{cov}[(F_t - s_t), u_{t+1}] < 0$ and hence explain the systematic under-prediction of future depreciation.⁶

In fact, we require only a small and quite natural deviation from rational expectations, based on the econometric learning approach increasingly utilized in macroeconomics. Recent applications include the design of monetary policy (Bullard and Mitra (2002), Evans and Honkapohja (2003), and Orphanides and Williams (2005a)), recurrent hyperinflations in Latin America (Marcet and Nicolini (2003)), US inflation and disinflation (Sargent (1999), Orphanides and Williams (2005b), Bullard and Eusepi (2005)), asset prices (Timmermann (1993), Brock and Hommes (1998), Bullard and Duffy (2001), Adam et. al. (2006)), and currency crises and exchange rates (Kasa (2004),

³Fama (1984) demonstrates that, for this to happen, the variance of the risk premium must be greater than the variance of expected depreciation, and their covariance must be negative. These requirements do not appear to be supported empirically.

⁴De Long et. al. (1990) demonstrated that the presence of both rational and non-rational traders in the market tends to distort asset prices significantly away from the fundamental values and therefore has the potential to explain many financial market anomalies. Mark and Wu (1998) demonstrated that the behavior of the variance and covariance of the risk premium as required by Fama (1984) does not have empirical support, while the existence of noise traders in the market under certain numerical assumptions yields results compatible with the data.

⁵Chakraborty and Haynes (2005) demonstrate, in the context of deviations from rationality, that nonstationarity in the relevant variables can explain the related puzzle of little or no bias in "level" specification between the future spot and current forward rate, yet significant negative bias with frequent sign reversals in the standard forward premium specification.

⁶In connection with the closely related issue of uncovered interest parity, McCallum (1994) argues that monetary policy response to exchange rate changes may account for the econometric findings. As he notes, this and the view that expectations are less than fully rational are potentially complementary explanations.

Kim and Mark (2005)).

In the present paper we show that when the fundamentals driving the exchange rate are strongly persistent, a downward bias in $\hat{\beta}$ *necessarily* arises for arbitrarily small deviations from rational expectations due to learning, and that this bias is potentially strong enough to reverse the sign of the relationship.

Our key assumption is that while agents do know the true form of the relationship between the fundamentals and the exchange rate that would hold under rational expectations, they do not know the parameter values and must estimate them from observed data.⁷ In the model we analyze, the exchange rate s_t , under rational expectations, satisfies

$$s_{t+1} = bv_t + u_{t+1},$$

where v_t is the observed value of the fundamentals, assumed exogenous, and u_{t+1} is unforecastable white noise. Under rational expectations b takes a particular value \bar{b} that depends on the model parameters and on the parameters of the stochastic process v_t . The rational one-step ahead forecast is then given by $E_t s_{t+1} = \bar{b}v_t$. However, we instead make the assumption that the agents do not know the true value of b and must estimate it from the data by running a regression of s_{t+1} on v_t .

More specifically, agents estimate b by “constant gain” or “discounted” least-squares learning of the type studied by Sargent (1999), Bischi and Marimon (2001), Cho et. al. (2002), Kasa (2004), Williams (2004) and Orphanides and Williams (2005a).⁸ Orphanides and Williams refer to this as “perpetual” learning, since agents remain perpetually alert to possible structural change. We show that under this form of learning the agents’ estimates b_t are centered at the RE value \bar{b} , but gradually and randomly move around this value as the estimates respond to recent data. Because b_t is not exactly equal to \bar{b} in every period, we have a deviation from full rational expectations. However, agents are in many ways very rational and quite sophisticated in their learning: they know the form of the relationship and estimate the true parameter value, adjusting their estimates, in response to recent forecast errors, in accordance with the least squares principle.

⁷Lewis (1989) used Bayesian learning to provide an explanation for the forward premium puzzle. However, the model could not explain the persistence of prediction errors, since the magnitude of the error shrinks, over time, to zero.

⁸For a general discussion of constant gain learning, see Chapter 14 of Evans and Honkapohja (2001).

Is this form of least-squares learning sufficient to explain the forward premium puzzle? We argue that indeed it may. Using theoretical results from the macroeconomics learning literature, we can derive the stochastic process followed by b_t under learning and derive an approximation for the asymptotic bias of the least-squares estimate $\hat{\beta}_t$ of the forward premium slope coefficient. This bias turns out to depend on all the structural parameters in the model, including the autoregressive coefficient ρ of the fundamentals process, which we model as a simple AR(1) process. We are interested in results for the case of large $0 < \rho < 1$ and especially as $\rho \rightarrow 1$, since in this limiting case the exchange rate under rational expectations would follow a random walk, in accordance with the well-known empirical results of Meese and Rogoff (1983). Our principal finding is that precisely in this case the downward bias is substantial. Perpetual learning therefore appears capable of entirely explaining the forward premium puzzle.

2 Framework

2.1 A simple exchange rate model

To illustrate our central point we use a very simple monetary exchange rate model based on purchasing power parity, risk-neutrality and covered interest parity.⁹ The equations are as follows:

$$F_t = \hat{E}_t s_{t+1} \quad (2)$$

$$i_t = i_t^* + F_t - s_t \quad (3)$$

$$m_t - p_t = d_0 + d_1 y_t - d_2 i_t \quad (4)$$

$$p_t = p_t^* + s_t. \quad (5)$$

Here s_t is the log of the price of foreign currency, F_t is the log of the forward rate at t for foreign current at $t+1$, and $\hat{E}_t s_{t+1}$ denotes the market expectation of s_{t+1} held at time t . Equation (2) assumes risk neutrality and equation (3) is the closed parity condition, with i_t and i_t^* the domestic and foreign interest rate, respectively. Equation (4) represents money market equilibrium, where m_t is the log money supply, p_t is the log price level and y_t is log real GDP. Finally the purchasing power parity condition is given by (5), where p_t^* is the log foreign price level. The parameters d_1, d_2 are assumed to be positive.

⁹See, for example, Frenkel (1976), Mussa (1976) and Engel and West (2005).

These equations can be solved to yield the reduced form

$$s_t = \theta \hat{E}_t s_{t+1} + v_t, \quad (6)$$

where $\theta = d_2/(1 + d_2)$, so that $0 < \theta < 1$.

$$v_t = (1 + d_2)^{-1}(m_t - p_t^* - d_0 - d_1 y_t + d_2 i_t^*)$$

represents the “fundamentals.” We will treat v_t as an exogenous stochastic process, which implicitly assumes the “small country” case with exogenous output.¹⁰ We will focus on the case in which v_t is an observable stationary AR(1) process¹¹

$$v_t = \delta + \rho v_{t-1} + \varepsilon_t$$

with serial correlation parameter $0 < \rho < 1$. For application of the theoretical learning results we need to make the technical assumption that v_t has compact support.¹² Our results would also apply to the case in which v_t is trend-stationary, with compact support around a known deterministic trend (and could be extended to the case in which the trend is unknown).

In modeling expectation formation by the agents we make the assumption that their forecasts $\hat{E}_t s_{t+1}$ are based on a reduced form econometric model of the exchange rate, specifically $s_t = a + b v_{t-1} + \eta_t$, where η_t is treated as exogenous white noise, using coefficients that are estimated from the data using discounted least-squares. Specifically, we assume that at the beginning of time t , agents have estimates a_{t-1}, b_{t-1} of the coefficients a, b , based on data through time $t - 1$. These, together with the observed current value of the fundamentals v_t , are used to forecast the next period’s exchange rate $\hat{E}_t s_{t+1} = a_{t-1} + b_{t-1} v_t$. The fundamentals, together with the forecasts, determine the exchange rate according to (6), and then at the end of period t the parameter estimates are updated to a_t, b_t , for use in the following period. We now turn to a detailed discussion of the learning rule and the theoretical results for the system under learning.

¹⁰For the large country case see Chakraborty (2005, 2006).

¹¹It would be straightforward to allow for an additional unobserved white noise shock.

¹²This rules out the normal distribution, but is compatible with a truncated normal distribution in which the distribution is restricted to an (arbitrarily large) closed interval. Our assumption of compact support ensures that v_t has finite moments of all orders.

3 Formal Results under Learning

3.1 Stochastic approximation results

For theoretical convenience we examine the system

$$\begin{aligned} s_t &= \theta \hat{E}_t s_{t+1} + v_t \\ v_t &= \rho v_{t-1} + \varepsilon_t, \end{aligned}$$

where $\varepsilon_t \sim iid(0, \sigma_\varepsilon^2)$ and $0 \leq \rho < 1$. Here we have normalized the intercept to zero, which is equivalent to assuming that agents know its true value and that we are looking at the system in deviation from the mean form. In the RE (rational expectations) solution

$$s_t = \bar{b}v_{t-1} + \bar{c}\varepsilon_t, \text{ where } \bar{b} = (1 - \rho\theta)^{-1}\rho \text{ and } \bar{c} = (1 - \rho\theta)^{-1}.$$

Instead, market participants estimate the coefficient b by constant gain least squares.¹³ This is most conveniently expressed in recursive form.¹⁴ The estimate based on data through time t is given by the algorithm

$$\begin{aligned} b_t &= \gamma R_{t-1}^{-1} v_{t-1} (s_t - b_{t-1} v_{t-1}) \\ R_t &= \gamma (v_t^2 - R_{t-1}), \end{aligned} \tag{7}$$

where $\gamma > 0$ is a small positive constant. R_t can be viewed as an estimate of the second moment of the fundamentals. Since forecasts are formed as

$$\hat{E}_t s_{t+1} = b_{t-1} v_t, \tag{8}$$

the exchange rate under learning is given by

$$s_t = (\theta b_{t-1} + 1) v_t. \tag{9}$$

Using stochastic approximation results it can be shown that the mean path of b_t and R_t can be approximated by the differential equations¹⁵

$$\begin{aligned} db/d\tau &= R^{-1} \sigma_v^2 ((\theta\rho - 1)b + \rho) \\ dR/d\tau &= \sigma_v^2 - R, \end{aligned} \tag{10}$$

¹³If $\delta \neq 0$ then the REE is $s_t = \bar{a} + \bar{b}v_{t-1} + \bar{c}\varepsilon_t$, where \bar{b}, \bar{c} are unchanged and $\bar{a} = (1 - \theta)^{-1}(1 - \rho\theta)^{-1}\delta$. Under learning agents would estimate (a, b) using constant gain recursive least squares. The numerical results of Section 5 allow for $\delta \neq 0$.

¹⁴See, e.g., Marcet and Sargent (1989), Sargent (1999) or Evans and Honkapohja (2001).

¹⁵See the Appendix for technical details.

where $\tau = \gamma t$. This differential equation system has a unique equilibrium $(\bar{b}, \bar{R}) = ((1 - \rho\theta)^{-1}\rho, \sigma_v^2)$ that is globally stable, so that, whatever the initial values for the learning algorithm, we have $E b_t \rightarrow \bar{b}$ as $t \rightarrow \infty$.

Under ordinary (“decreasing gain”) least-squares learning γ is replaced by $1/t$ and it can be shown that in the limit we obtain fully rational expectations, i.e. $b_t \rightarrow \bar{b}$ with probability one as $t \rightarrow \infty$. We instead focus on the natural modification in which ordinary least-squares is replaced by constant gain least squares, as above, so that γ is a small fixed positive number, e.g. $\gamma = 0.02$ or $\gamma = 0.05$. This assumption – that agents weight recent data more heavily than past data – is being increasingly studied in the macroeconomic literature, as noted in the introduction.

Why would constant gain learning be natural to employ? As emphasized by Sargent (1999), applied econometricians and forecasters recognize that their model is subject to misspecification and structural change. Constant gain least-squares is a natural way to allow for potential structural change taking an unknown form, because it weights recent data more heavily than older data. This procedure is well known in the statistics and engineering literature, see for example, Chapters 1 and 4, Part I, of Benvensite et. al. (1990). As noted by Orphanides and Williams (2005a), an additional theoretical advantage is that it converts the model under learning to a stationary environment, so that results can be stated in a way that does not depend on the stage of the learning transition. In effect, under constant gain least squares, agents are engaged in perpetual learning, always alert for possible changes in structure.

Of course the appropriate choice the of gain parameter γ will be an issue of some importance. In principle this parameter might be chosen by agents in an optimal way, reflecting the trade-off between tracking and filtering. This is discussed in Benvensite et. al. (1990) and analyzed in a simple economic set-up in Evans and Ramey (2006). In the current paper, in line with most of the literature, we do not directly confront this issue, but instead investigate how our results depend on the value of the gain. Empirical macroeconomic evidence on forecaster expectations and forecast performance for GDP growth and inflation,¹⁶ suggest values of the gain for quarterly data in the range $\gamma = 0.02$ to $\gamma = 0.05$. Reasonable values for γ in our setting will depend on the amount of perceived structural change in the link between the exchange rate and fundamentals and may, therefore, be different.

¹⁶See Orphanides and Williams (2005b) and Branch and Evans (2006).

Under constant gain learning, a natural result is obtained that goes beyond the decreasing gain asymptotic convergence result. Rational expectations can still be viewed as a limiting case, but constant gain learning turns out to yield surprising results for small deviations from this limit. Our central starting point is the unsurprising result that with a small constant gain $\gamma > 0$, the parameter b_t remains random as $t \rightarrow \infty$, with a mean equal to the RE value \bar{b} , and with a small variance around \bar{b} . We have the following:

Proposition 1 *Consider the model under constant gain learning. For $\gamma > 0$ sufficiently small, and γt sufficiently large, b_t is approximately normal with mean \bar{b} and variance γC , where*

$$C = \frac{1 - \rho^2}{2(1 - \rho\theta)^3},$$

and the autocorrelation function between b_t and b_{t-k} is approximately $e^{-(1-\theta\rho)\gamma k}$.

The proof is given in the Appendix. Thus, provided the process has been running for sufficiently long so that the influence of initial conditions is small, the distribution of b_t at each time t can be approximated by

$$b_t \sim N(\bar{b}, \gamma C),$$

for $\gamma > 0$ small. Note that rational expectations arises as the limit in which $\gamma \rightarrow 0$, since in this case at each time t the parameter estimate b_t has mean \bar{b} and zero variance. Thus for small $\gamma > 0$ we are indeed making small deviations from rationality.

Up to this point the results may appear straightforward and fairly uncontroversial: under perpetual gain learning with small constant gain $\gamma > 0$, the agents' estimate of the key parameter used to forecast exchange rates has a mean value equal to its RE value, but is stochastic with a standard deviation depending on the structural parameters and proportional to $\sqrt{\gamma}$. However, the implications for the forward premium puzzle are dramatic, as we will now see.

Using Proposition 1 we can obtain the implications for the bias of the least squares estimate $\hat{\beta}$, in the forward premium regression (1), under the null hypothesis $H_0 : \alpha = 0, \beta = 1$, when private agents forecast exchange rates using constant gain least squares updating with a small gain γ . For

convenience we assume that $\alpha = 0$ is imposed so that the econometrician estimates a simple regression without intercept.¹⁷

The Appendix establishes the following result:

Proposition 2 *Under the null hypothesis H_0 the asymptotic bias $\text{plim } \hat{\beta} - 1$, for $\gamma > 0$ sufficiently small, is approximately equal to*

$$B(\gamma, \theta, \rho) = -\frac{\gamma(1-\theta)(1+\rho)(1-\theta\rho)}{\gamma(1-\theta)^2(1+\rho) + 2(1-\rho)(1-\theta\rho)}.$$

Thus for all parameter values $0 \leq \theta < 1$ and $0 \leq \rho < 1$, we have a negative bias, which is particularly strong for ρ near 1. More specifically we have:

Corollary 3 *$B(\gamma, \theta, \rho) < 0$ for all $0 \leq \theta < 1$, $0 \leq \rho < 1$ and $0 < \gamma < 1$, and the size of the approximate bias $|B(\gamma, \theta, \rho)|$ is increasing in γ and in ρ and decreasing in θ . For $\gamma > 0$ sufficiently small, we obtain the limiting approximations*

$$\lim_{\rho \rightarrow 1}(\text{plim } \hat{\beta} - 1) = -1 \text{ and } \text{plim } \hat{\beta} - 1 = -\frac{\gamma(1-\theta)}{\gamma(1-\theta)^2 + 2} \text{ if } \rho = 0.$$

Corollary 3 implies that, for small γ , the value of $\text{plim } \hat{\beta}$ approaches 0 as $\rho \rightarrow 1$. Below, in Section 3.2, we investigate the situation numerically and find that small samples can further magnify the bias: for typical sample sizes and plausible values of γ , median values of $\hat{\beta}$ are *negative* as $\rho \rightarrow 1$.

Finally we can also examine the t -statistic for the test of $H_0 : \beta = 1$, given by $t_{\hat{\beta}} = (\hat{\beta} - 1)/RE(\hat{\beta})$. Since for all $0 \leq \rho < 1$ we have $\text{plim } \hat{\beta} - 1 < 0$ it follows that:

Corollary 4 *For $\gamma > 0$ sufficiently small, $t_{\hat{\beta}} \rightarrow -\infty$ as the sample size $T \rightarrow \infty$.*

Our results are stated for sufficiently small γ because this is needed to invoke the stochastic approximation results. Below we look at the quality of the approximation for plausible values of $\gamma > 0$. The theoretical results are illustrated in Figure 1, which shows the approximation $\text{plim } \hat{\beta} = 1 + B(\gamma, \theta, \rho)$, as a function of ρ over $0 \leq \rho \leq 1$, for fixed $\theta = 0.6$, and for three values $\gamma = 0.01, 0.05$ and $\gamma = 0.10$.

¹⁷This makes no difference asymptotically. Below we numerically investigate how inclusion of the intercept in the test regression affects the small sample results.

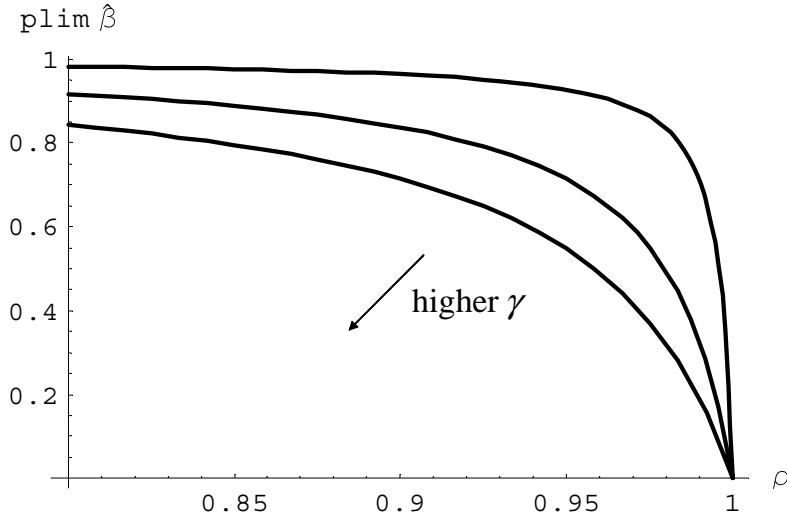


Figure 1: Theoretical $\text{plim}(\hat{\beta})$ for $\theta = 0.6$ and $\gamma = 0.01, 0.05$ and 0.10 .

As expected, the asymptotic bias depends upon γ , and for sufficiently small $\gamma > 0$ the size of the bias, given ρ , is proportional to γ . For any given $0 \leq \rho < 1$, as $\gamma \rightarrow 0$ we approach the rational expectations limit and in this limit the bias of $\hat{\beta}$ is zero. However, a striking and surprising feature of our results is the behavior of $\text{plim } \hat{\beta}$ as $\rho \rightarrow 1$ for fixed γ : given γ , the asymptotic bias of $\hat{\beta}$ approaches -1 as $\rho \rightarrow 1$, regardless of the size of γ . The intuition for this result is given below, in Section 4. Here we emphasize the powerful implications for the forward premium test, which we state as follows:

Corollary 5 *For any $\varepsilon > 0$ there exists $\gamma > 0$ and $\hat{\rho} < 1$ such that for all $\hat{\rho} \leq \rho < 1$ we have both $E(b_t - \bar{b})^2 < \varepsilon$ for all t and $\text{plim } \hat{\beta} < \varepsilon$.*

Thus, for learning gain parameters sufficiently small, provided the autocorrelation parameter of the fundamentals process is sufficiently high, the deviation from rational expectations will be arbitrarily small, at every point in time, as measured by mean square error, and yet the downward bias in the forward premium regression can be made arbitrarily close to -1 .

3.2 Numerical and Small Sample Results

To determine the quality of the approximation we have simulated paths of b_t , using equations (7) and (9), and computed numerical estimates of $E(b_t)$ and $\sqrt{\text{var}(b_t)}$. The results are given in Tables 2 and 3.¹⁸ The approximation appears generally satisfactory although the quality deteriorates as $\rho \rightarrow 1$, especially for larger values of γ . These findings are to be expected: our theoretical results give the limiting results for small γ and rely on stationarity of the fundamental process v_t . Although the theoretical results hold for all $|\rho| < 1$, provided $\gamma > 0$ is sufficiently small, it is not surprising to see deviations for fixed γ as v_t approaches a nonstationary process via $\rho \rightarrow 1$.¹⁹ Table 2 clearly shows that the sample mean is quite close to the predicted RE value \bar{b} for most parameter values, with sample means slightly less than \bar{b} for larger values of γ . Table 3 shows a fairly good match with theory for most parameter values, with the main deviation being underprediction of the standard deviation for large ρ and γ .

We next look at numerical results concerning the forward premium regressions. We first examine the quality of the theoretical approximation results given in Proposition 2 and then determine the small sample properties. Our focus now is on the predicted bias that arises in the forward premium regression (1) when agents forecast using perpetual least squares learning with a small constant gain. Table 4 reports the simulation results for a large sample $T = 20,000$. Tables 5a and 5b give the small sample results, for $T = 120$ and $T = 360$, realistic samples sizes with quarterly and monthly data, respectively, both for $\hat{\beta}$ and for the t -statistic of the test of $H_0 : \beta = 1$.

Table 4 presents the comparison between $\hat{\beta}$ values predicted by Proposition 2 and the mean values generated by simulations under learning with different combinations of parameter values. Although our theoretical results have not been formally demonstrated for the limit case $\rho = 1$, a pure random walk, we include simulations for this value as well. Again, it appears that the theoretical prediction is fairly accurate for small values of γ . As noted earlier, deviations for larger γ are understandable since the theory developed here is valid for small γ . The key qualitative predictions of Proposition 2, and Corollary 3, hold in the numerical results of Table 4. In particular, an

¹⁸In all of our numerical results we have chosen ε_t to be *iid* with a standard normal distribution. We have also set $\delta = 0$ unless otherwise specified.

¹⁹Very small values of $\gamma > 0$ present numerical difficulties since extremely large sample sizes would then be needed to reliably estimate the mean and standard deviation.

increase in γ or ρ (and the smaller value of θ) leads to a smaller value of $\text{plim } \hat{\beta}$. For $\gamma = 0.05$ or $\gamma = 0.10$ the simulation results in Table 4 show an even stronger downward bias in $\hat{\beta}$ than is predicted by our theoretical results.

We now consider the small sample results given in Table 5. The sample size employed in Table 5a of $T = 120$ corresponds to thirty years of non-overlapping quarterly data and in Table 5b $T = 360$ corresponds to thirty years of non-overlapping monthly data. Although the results again show a substantial downward bias in $\hat{\beta}$ for an important range of parameter values, there are significant differences in the small sample results and the pattern is more erratic. On the one hand, there are cases of positive bias that arise with small γ , lower ρ and higher θ . On the other hand, especially for ρ close to or equal to one, the downward bias is even more extreme. Inspection of the detailed results show a substantial number of extreme values for $\hat{\beta}$ and the t -statistic (which is why we report their median values).

One of the reasons for the complex small sample results can be seen from the following argument. If we have both a small gain γ and a small sample size T the value of b_t will vary little within the sample. Useful insights can thus be obtained by considering the limiting case of $b_t = b$ fixed over the sample period at some value possibly different from \bar{b} . If agents believe that $s_t = bv_{t-1} + c\varepsilon_t$, we have $F_t = \hat{E}_t s_{t+1} = bv_t$ and $s_t = (1 + \theta b)v_t$ so that the forward premium is

$$F_t - s_t = ((1 - \theta)b - 1)v_t. \quad (11)$$

and the forecast error $u_{t+1} = s_{t+1} - bv_t$ is given by

$$u_{t+1} = (1 + \theta b)(\rho v_t + \varepsilon_{t+1}) - bv_t.$$

Although we cannot calculate $E(\hat{\beta})$ for a finite T it is revealing to compute

$$a(b) = \frac{\text{cov}[(F_t - s_t), u_{t+1}]}{\text{var}(F_t - s_t)} = -\frac{(1 - \theta\rho)(b - \bar{b})}{(1 - \theta)(b - (1 - \theta)^{-1})}, \quad (12)$$

which is the asymptotic bias that would result as $T \rightarrow \infty$ if b were kept fixed.

The asymptotic bias is negative for $b < \bar{b}$ and less than -1 for $b > 1/(1 - \theta)$. However for $\bar{b} < b < 1/(1 - \theta)$ the asymptotic bias is positive and there is a singularity at $b = 1/(1 - \theta)$, with both arbitrarily large negative and positive values in a neighborhood of the singularity.²⁰ Calculating $a(b)$ is artificial since it holds $b_t = b$ fixed as $T \rightarrow \infty$, whereas under perpetual

²⁰This phenomenon disappears in the limiting case $\rho = 1$ since then $\bar{b} = 1/(1 - \theta)$.

learning b_t is a stochastic process centered at \bar{b} . However, it clearly indicates the complexities that can be expected in small sample simulations.

In Table 5c we show, for selected parameter values of interest, how the differences between the asymptotic results of Table 4 and the small sample results of Table 5a and 5b depend on the sample size. In Table 5c we also investigate the small sample effect of including an intercept in the test regression. It can be seen that in small samples the inclusion of an intercept in the test regression further magnifies the deviation from the asymptotic results. This effect is particularly striking for the smaller gain parameter value $\gamma = 0.01$, in line with the argument of the preceding paragraph. For $\rho = 1$ we obtain negative values of $\hat{\beta}$ for all values of $\hat{\beta}$ considered and when an intercept is included in the test regression the effect can be pronounced.²¹ Whether or not an intercept is included, as the sample size T becomes large there is convergence to the theoretical and large sample results given earlier.

On balance the findings of Table 5 reinforce the theoretical results of Section 3.1 and the central thrust of this paper. For ρ near or equal to 1, and for empirically plausible values of γ , the median value of $\hat{\beta}$ is not only biased downwards from 1, but *negative* values for $\hat{\beta}$ would be entirely unsurprising. Thus for fundamentals processes that are close to a random walk, perpetual learning clearly has the potential to explain the forward premium puzzle.

4 Discussion

What is the source of the downward bias to $\hat{\beta}$ that we have established theoretically and numerically? In this section we provide the intuition for the case in which the fundamentals follow a random walk, i.e. $\rho \rightarrow 1$. Our starting point is the result that $b_t \sim N(\bar{b}, \gamma C)$. Since, for small $\gamma > 0$, the parameter b_t is near \bar{b} and moves very gradually over time, it is useful again to consider the impact on $\hat{\beta}$ of an arbitrary value for b held fixed at a value close to but not equal to \bar{b} . As $\rho \rightarrow 1$ the fixed b asymptotic bias function (12) satisfies $a(b) \rightarrow -1$ at every point *other* than the singularity, which for $\rho = 1$ coincides with the RE solution \bar{b} . This is fully consistent with the theoretical findings of Section 3.1. What is the underlying reason for this result?

²¹Chakraborty (2005) shows that similar qualitative results are obtained for ARIMA(p,1,q) estimates of the fundamental processes.

When $\rho = 1$, the fundamentals v_t follow a pure random walk, the RE solution is $s_t = (1 - \theta)^{-1}v_t$, or equivalently $s_t = (1 - \theta)^{-1}v_{t-1} + (1 - \theta)^{-1}\varepsilon_t$, and $F_t = E_t s_{t+1} = (1 - \theta)^{-1}v_t$. Thus under RE

$$\begin{aligned} s_t &= \bar{b}v_t = F_t \text{ where } \bar{b} = (1 - \theta)^{-1}, \text{ and} \\ F_t - s_t &\equiv 0 \text{ and } u_{t+1} = s_{t+1} - F_t = \bar{b}\varepsilon_{t+1}. \end{aligned}$$

Consider now the situation for $b \neq \bar{b}$. As discussed in the Introduction, $\hat{\beta}$ is biased downward from one if $\text{cov}_t(F_t - s_t, u_{t+1}) < 0$.²² If agents believe that $s_t = bv_{t-1} + c\varepsilon_t$, we have from (11) that

$$F_t - s_t = (1 - \theta)(b - \bar{b})v_t$$

when $\rho = 1$. The intuition is clearest if we split u_{t+1} into

$$u_{t+1} = \Delta s_{t+1} - (F_t - s_t),$$

i.e. the difference between Δs_{t+1} and the forward premium. Then

$$\begin{aligned} \text{cov}_t(F_t - s_t, u_{t+1}) &= \text{cov}_t(\Delta s_{t+1}, F_t - s_t) - \text{var}_t(F_t - s_t) \\ &= -\text{var}_t(F_t - s_t) < 0 \text{ if } b \neq \bar{b}, \end{aligned}$$

since in the random walk case $\Delta s_{t+1} = b\varepsilon_{t+1}$, *whatever the value of b* , and since $\text{cov}_t(\varepsilon_{t+1}, F_t - s_t) = 0$.

To summarize, under the true regression model $H_0 : \alpha = 0, \beta = 1$, but with (arbitrarily) small deviations from RE, the error term u_{t+1} in the forward premium regression is negatively correlated with the forward premium because u_{t+1} is simply the difference between the (unforecastable) exchange rate change and the forward premium itself. This negative correlation is present unless $b = \bar{b}$ i.e. RE holds exactly, in which case $\text{var}_t(F_t - s_t) = 0$. Furthermore, for $b \neq \bar{b}$ we have $\text{cov}_t(F_t - s_t, u_{t+1}) / \text{var}_t(F_t - s_t) = -1$, for all t . Since this holds for all $b \neq \bar{b}$, since under learning b_t will be close to but (with probability one) not equal to \bar{b} , and since with a small gain $\gamma > 0$ the agents' estimates b_t will be almost constant over time, it is not surprising that Proposition 2 was able to establish a downward bias of $\text{plim}(\hat{\beta} - 1) = -1$ for the limiting case $\rho \rightarrow 1$.

²²Here we use conditional covariances and variances because for $b \neq \bar{b}$ the unconditional moments are not well-defined when $\rho = 1$. However, as seen below, the conditional moments are independent of t . Furthermore, the unconditional moments are well-defined for all $0 < \rho < 1$ and $\lim_{\rho \rightarrow 1}(\text{cov}_t(u_{t+1}, F_t - s_t) / \text{var}_t(F_t - s_t)) = -1$.

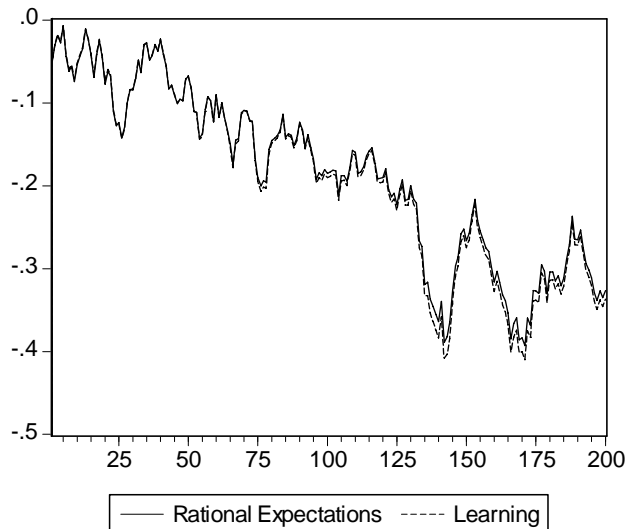


Figure 2: Model generated time paths of log exchange rate, for $\theta = 0.6$, $\rho = 0.99$, under rational expectations and under learning ($\gamma = 0.05$).

What *is*, perhaps, unexpected and surprising is that *arbitrarily small* deviations from RE yield a downward bias near -1 for ρ near 1. The reason for this is that the asymptotic bias depends on the *ratio* $\text{cov}(F_t - s_t, u_{t+1}) / \text{var}(F_t - s_t)$. Under RE $\text{cov}(F_t - s_t, u_{t+1}) = 0$ for all $0 \leq \rho \leq 1$ but $\text{var}(F_t - s_t) \rightarrow 0$ as $\rho \rightarrow 1$. Thus under RE the ratio is always zero except at $\rho = 1$, when the ratio is undefined since $F_t - s_t \equiv 0$. Under learning we also have $\text{plim}(\widehat{\text{cov}}((F_t - s_t), u_{t+1})) \rightarrow 0$ and $\text{plim}(\widehat{\text{var}}(F_t - s_t)) \rightarrow 0$ as $\gamma \rightarrow 0$ but the ratio is close to -1 for $\rho < 1$ near 1.

Although under our approach there are persistent deviations from RE, the parameter b_t is centered at and stays close to its RE value. Consequently, although expectations are not fully rational, the mistakes are both small and not consistently wrong in a way that is easily detectable. Indeed the time path of the exchange rate will typically be close to the RE path, even when the forward premium regression gives a value for $\hat{\beta}$ with the wrong sign.

Figures 2–4 give the results of a typical simulation of our model over $T = 200$ periods, with parameters set at $\theta = 0.6$ and $\rho = 0.99$.²³ Figure 2 gives

²³The standard deviation of the innovation to the fundamentals has been chosen so that

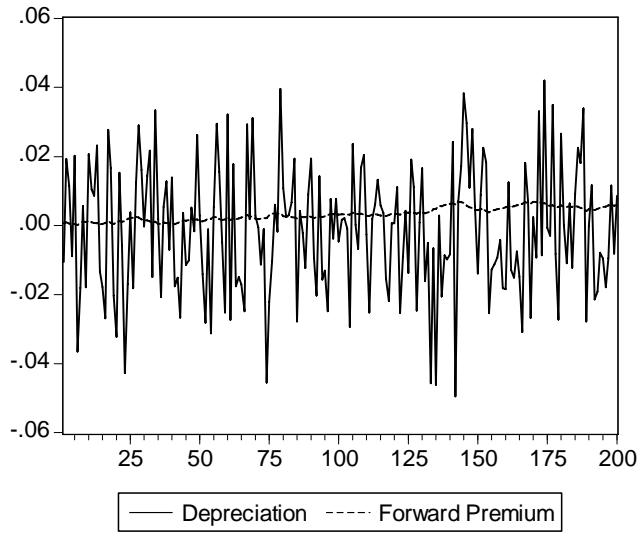


Figure 3: Model generated data with rational expectations, $\theta = 0.6$ and $\rho = 0.99$. Test statistic $\hat{\beta} = 1.10$.

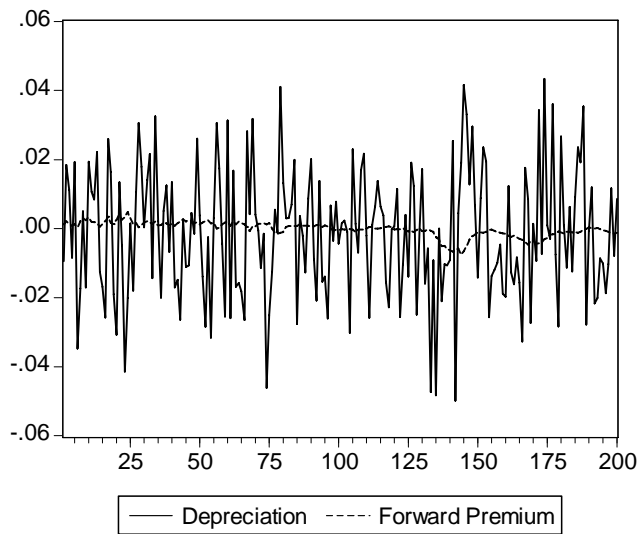


Figure 4: Model generated data with learning, $\theta = 0.6$, $\rho = 0.99$, and $\gamma = 0.05$. Test statistic $\hat{\beta} = -0.49$.

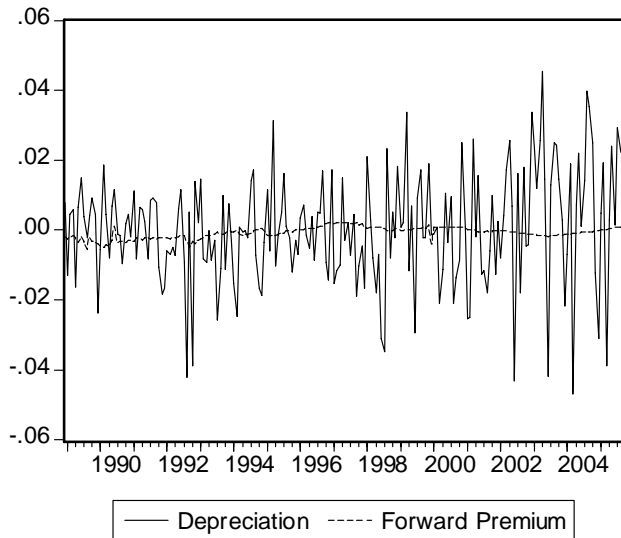


Figure 5: Time path of USD/CAD (Canadian Dollar) log exchange rate (monthly data December 1988 - September 2005). Test statistic $\hat{\beta} = -0.60$.

the time paths for the log of the exchange rate under rational expectations and under least-squares learning with constant gain $\gamma = 0.05$. The two time paths are generated by the same sequence of exogenous random shocks. Some mild “overshooting” under learning is evident, which is another immediate implication of learning for ρ near 1.²⁴ Figure 3 gives the corresponding simulation results under RE for depreciation and the forward premium, while Figure 4 gives the same variables under learning.²⁵ Although the qualitative features of the simulated data are the same under RE and under learning, the estimate $\hat{\beta}$ from the RE data is 1.10, and is insignificantly different from 0 (with $t_{\hat{\beta}} = 0.14$). In contrast, the corresponding value under learning is $\hat{\beta} = -0.49$ (with $t_{\hat{\beta}} = -2.42$), a typical illustration of the forward premium

the scale for depreciation is similar to that seen in the Canadian-US data.

²⁴The somewhat greater variation of s_t under learning is consistent with the excess volatility results of Kim and Mark (2005). The extent of overshooting and excess volatility seen in our simulations depends on the parameter settings.

²⁵Note that the explanatory power of the forward premium regressions is low, matching another standard finding in the data. This phenomenon was stressed by McCallum (1994), and is apparent, for example, in the Canadian-US data shown in Figure 5.

puzzle.²⁶ For comparison Figure 5 presents the quarterly depreciation and forward premium data for the Canadian dollar price of the US dollar over 1988.Q4 - 2005.Q3. Qualitatively this data resembles the simulated results in Figures 3 and 4, but the forward premium regression results of $\hat{\beta} = -0.60$ clearly are more in accordance with our model under learning.²⁷

5 Extensions and Further Discussion

In this section we briefly take up several alternative formulations and extensions in order to illustrate the robustness of our results.

5.1 Present value formulation

Applying the law of iterated expectations to the reduced form model (6) implies that $\hat{E}_t s_{t+j} = \theta \hat{E}_t s_{t+j+1} + \hat{E}_t v_{t+j}$. By recursive substitution and assuming that $\lim_{j \rightarrow \infty} \left| \theta^j \hat{E}_t s_{t+j} \right| = 0$ we obtain

$$s_t = \sum_{j=0}^{\infty} \theta^j \hat{E}_t v_{t+j}, \quad (13)$$

providing the sum converges. Equation (13) is sometimes called the “forward” or “present value” solution to (6), and it is the unique nonexplosive rational expectations solution for $|\theta| < 1$. When v_t follows an AR(1) process $v_t = \delta + \rho v_{t-1} + \varepsilon_t$, it is easily shown that under RE (13) yields $s_t = \delta \theta (1 - \theta)^{-1} (1 - \theta \rho)^{-1} + (1 - \theta \rho)^{-1} v_t$, for $-1 < \rho \leq 1$, which of course agrees with the solution given in Section 2.1.

In the model (6) with learning there are two natural approaches, depending on whether we treat s_t as determined directly from (6) by v_t and $\hat{E}_t s_{t+1}$, or whether we think of s_t as determined by the “discounted” sum of expected fundamentals (13). Both approaches have been used in the literature on learning and asset prices, e.g. both are used in Timmermann (1996). In this paper we have used the “self-referential approach,” based directly on (6), both because it squarely rests on the open-parity condition stressed in

²⁶The results shown are typical, but we remark that there is a wide variation across simulations. For example, for this parameter setting, the 25% and 75% quartiles for $\hat{\beta}$ are approximately -1.3 and 0.17 .

²⁷For a comprehensive empirical analysis see Chakraborty (2006).

the exchange rate literature and because it emphasizes that exchange rates are determined by short-run expected exchange rate movements as well as by fundamentals.

However, it is of interest to know if our results are also obtained if the present value formulation (13) is used, where the role of learning is confined to estimation of the fundamentals process. This is the approach emphasized in Kim and Mark (2005) in their analysis of the potential for learning to explain exchange rate volatility and the observed links between exchange rates and fundamentals. We again examine the results under constant gain learning. Thus we assume that agents estimate $v_t = \delta + \rho v_{t-1} + \varepsilon_t$ by constant gain RLS,²⁸ using data through time $t - 1$. This yields estimates $(\hat{\delta}_{t-1}, \hat{\rho}_{t-1})$, which agents use to compute $\hat{E}_t v_{t+j} = \hat{\delta}_{t-1} \sum_{i=0}^{j-1} \hat{\rho}_{t-1}^i + \hat{\rho}_{t-1}^j v_t$ at time t . Computing (13) the exchange rate s_t is given by²⁹

$$s_t = \hat{\delta}_{t-1} \theta (1 - \theta)^{-1} (1 - \theta \hat{\rho}_{t-1})^{-1} + (1 - \theta \hat{\rho}_{t-1})^{-1} v_t,$$

with the forward rate given $F_t = \hat{E}_t s_{t+1} = \theta^{-1} (s_t - v_t)$.

Table 6 gives the finite sample results. It can be seen that the results are consistent with the key results of main part of the paper. For $\rho < 1$ close to one and $\rho = 1$, and for gain parameters consistent with the learning literature, we find $\hat{\beta}$ strongly biased downward and often negative. Thus perpetual learning also leads to results in line with the forward premium puzzles in this alternative formulation.

5.2 Random structural change and endogenous gain

The motivation for constant gain least squares is that it allows agents to better track any structural change that occurs. Throughout the paper we have analyzed the impact of the use of constant gain learning in a model in which, in fact, there is no structural change. In effect, we have studied the implications solely of the use by agents of a learning rule with greater robustness to structural change than ordinary (decreasing gain) least-squares learning. This is in keeping with most of the now substantial literature on constant-gain or perpetual learning, reviewed earlier. However, a natural extension

²⁸The estimates are given by $\phi_t = \gamma R_{t-1}^{-1} X_{t-1} (v_t - \phi_{t-1} X_{t-1})$, $R_t = \gamma (X_t X_t' - R_{t-1})$, where $\phi_t' = (\delta_t, \rho_t)$ and $X_t' = (1, v_t)$.

²⁹To ensure that the sum converges we impose a ‘‘projection facility’’ that maintains estimates ϕ_t at their previous value if ρ_t would otherwise exceed 1.05. For discussions of projection facilities see Marcet and Sargent (1989) and Evans and Honkapohja (2001).

would examine the results for a model incorporating unknown structural change and agents using constant gain least squares learning.³⁰

We now briefly consider such an extension, adapting the structural change model used in Evans and Ramey (2006). The fundamentals process is now assumed to be

$$v_t = \delta + \mu_t + \rho v_{t-1} + \varepsilon_t,$$

where μ_t is a regime switching process taking the form

$$\mu_t = \begin{cases} \mu_{t-1} & \text{with probability } 1 - q \\ \zeta_t & \text{with probability } q, \end{cases}$$

where $0 \leq q \leq 1$ and ζ_t is an *iid* process. In our numerical results we assume that ζ_t is has the uniform distribution over the closed interval $[-L, L]$, where $L > 0$. The form of the process μ_t is assumed to be unknown to the agents, who continue to forecast using $\tilde{E}_t s_{t+1} = a_{t-1} + b_{t-1} v_t$, with the parameters (a_{t-1}, b_{t-1}) estimated using constant gain least-squares.

Following Evans and Honkapohja (1993), Marcet and Nicolini (2003) and Evans and Ramey (2006), we now also impose that the gain parameter γ is set at a (Nash equilibrium) value that minimizes the one-step ahead mean-square forecast error for individual agents, given that other agents use this value. Thus agents are choosing the gain parameter γ optimally in the MSE sense, trading off the benefits of lower γ , which increases filtering and thus reduces random fluctuations in estimated parameters, against the benefits of larger γ , which improves tracking of structural change.

Table 7 presents results for an illustrative numerical exercise with $\delta = 1$, $L = 0.4$ and $\theta = 0.6$. The probability of structural shift q is set at 2%, 5% or 10% per period. Equilibrium γ are approximate values computed numerically. The results confirm that with perpetual learning the results of the standard test regressions are entirely in accordance with the forward premium puzzle when the fundamentals follow an AR(1) process with ρ close to one. Estimated $\hat{\beta}$ are negative or close to zero, again reinforcing the central finding of this paper. The magnitudes of the t-statistics for the test of $H_0 : \beta = 1$ are now larger than in Tables 5a,b, as a result of the random structural shifts. These magnitudes, of course, would depend on the average size of the shifts, which is governed by L .

³⁰The evidence for structural change is considered in Chakraborty (2006).

5.3 Infrequent structural breaks

In the previous subsection we considered a model in which there are continuing occasional structural shifts, often small, where the dates of any structural changes are unknown to the agents. This is one plausible view of how structural change affects the economy. An alternative, e.g. Timmermann (1993), is that structural shifts are infrequent events and that the time of the shifts, though not the size of their impact, may be known to agents as soon as the shift occurs. In this set-up, agents can be expected to use a decreasing gain least-squares estimator as long as the structure is unchanged. However, when a known structural break occurs, the gain is increased to a larger value, with decreasing gain then employed until the next structural break.

The updating recursive algorithm for $\phi'_t = (a_t, b_t)$ is

$$\begin{aligned} \phi_t &= m_t^{-1} R_{t-1}^{-1} X_{t-1} (s_t - \phi_{t-1}' X_{t-1}), \\ R_t &= m_t^{-1} (X_t X_t' - R_{t-1}), \text{ where } X_t' = (1, v_t), \\ m_t &= \begin{cases} \bar{\gamma}^{-1} & \text{if structural change in } t-1 \\ m_{t-1} + 1 & \text{otherwise} \end{cases} \end{aligned}$$

For a structural break at $t = 1$ the gain sequence of $m_t = 1/t$ starting at $t = 2$ corresponds to ordinary least squares (with starting value $t = 2$ because there are two parameters to estimate). This is implemented with $\bar{\gamma} = 0.5$. Choosing a smaller value of $\bar{\gamma}$ would smooth initial estimates by placing additional weight on the prior estimate.

Again, we perform a small numerical experiment to study the robustness of our results. To do so we suppose that $v_t = \delta + \rho v_{t-1} + \varepsilon_t$ and start the system in the RE equilibrium. We then consider a 25% increase in δ with $\bar{\gamma} = 0.5$ or $\bar{\gamma} = 0.2$. The structural change occurs at $t = 1$, and that it has occurred becomes known to agents at the end of the period. Table 8 gives the results.³¹ The results are broadly in line with our main results. For the values of ρ tabled, there is a strong downward bias in $\hat{\beta}$ in every case except with $\theta = 0.9$ and the smaller gain increase to $\bar{\gamma} = 0.2$. In fact the downward bias emerges also in this case for ρ even closer to one. For example, with $\theta = 0.9$, $\bar{\gamma} = 0.2$ and $T = 360$ we get median $\hat{\beta} = 0.71$ for $\rho = 0.998$ and $\hat{\beta} = -2.53$ for $\rho = 0.999$.

³¹We here use values ρ very close to 1 in place of $\rho = 1$ so that the mean of v_t is well-defined. Qualitatively similar results showing a downward bias to $\hat{\beta}$ are obtained for a 25% decrease in the mean of the fundamentals.

Of course, as $T \rightarrow \infty$ we will find $\hat{\beta} \rightarrow 1$ since decreasing gain least-squares learning converges asymptotically to the RE. However, recurrent infrequent structural breaks can be expected to lead to a substantial downward bias in $\hat{\beta}$ for fundamentals processes that are close to a random walk.³² Our main conclusions thus appear robust also to this alternative formulation with infrequent structural breaks. Provided ρ is near to or equal to one, least-squares learning by market agents is consistent with the forward-premium puzzles results found in the literature.

6 Conclusions

The forward premium anomaly is a long outstanding puzzle that has proved difficult to explain based on risk premia and other orthodox approaches. While it has long been recognized that the anomalous empirical results might be due to irrationality in the exchange markets, the present paper shows that an adaptive learning approach increasingly employed in the macroeconomics literature appears able to reproduce the key empirical results. Modeling expectations by constant gain least-squares learning ensures that deviations from rational expectations are both small and persistent in realistic ways. Agents continue to update their parameter estimates because of concern for structural change, in a way similar to the use of rolling data windows. The result is perpetual learning by agents that keeps expectations close to RE, but with small random deviations due to revisions to the forecast rule driven by recent forecast errors.

We have shown theoretically that as the fundamentals process approaches a random walk, an empirically realistic case, even arbitrarily small deviations from RE, in accordance with perpetual learning, induce a large downward asymptotic bias in the estimated forward premium regression coefficient. Simulations for small sample results reinforce this result and indicate that negative values for this coefficient are fully consistent with our theory. Alternative formulations of learning and explicit incorporation of different types of structural shifts lead to qualitatively similar findings. The results of this paper thus suggest that the learning theory approach to expectation formation in the foreign exchange markets should be considered a serious contender in future empirical work on the forward premium puzzle.

³²In work in progress we examine the exchange-rate results under learning in greater detail for both frequent and infrequent structural change.

Appendix: Technical Details and Proofs

Proof of Proposition 1: We are considering the system (7). Combining these equations with (9) we obtain

$$\begin{aligned} b_t &= \gamma R_{t-1}^{-1} v_{t-1} ((\theta b_{t-1} + 1)v_t - b_{t-1}v_{t-1}) \\ R_t &= \gamma(v_t^2 - R_{t-1}). \end{aligned}$$

This takes the form

$$\Lambda_t = \Lambda_{t-1} + \gamma \mathcal{H}(\Lambda_{t-1}, X_t), \quad (\text{A.1})$$

where $\Lambda'_t = (b_t, R_t)$ and $X'_t = (v_t, v_{t-1})$ and where the components of \mathcal{H} are

$$\begin{aligned} \mathcal{H}_b(\Lambda_{t-1}, X_t) &= R_{t-1}^{-1} v_{t-1} ((\theta b_{t-1} + 1)v_t - b_{t-1}v_{t-1}) \\ \mathcal{H}_R(\Lambda_{t-1}, X_t) &= v_t^2 - R_{t-1}. \end{aligned}$$

Systems of the form (A.1) are known as stochastic recursive algorithms (SRAs), and have been widely studied in the learning literature.

The algorithm is initialized with some starting point $\Lambda_0 = a = (b_0, R_0)'$. We apply Proposition 7.8 and Theorem 7.9 of Evans and Honkapohja (2001), which are based on the stochastic approximation results of Benveniste, Metivier and Priouret (1990). That the required assumptions hold for the system at hand can be established using arguments analogous to those given on pp. 334-335 of Evans and Honkapohja (2001) for the cobweb model.

The stochastic approximation results for constant gain algorithms of this form are stated in terms of a continuous time process $\Lambda^\gamma(\tau)$. Let $\tau_t^\gamma = \gamma t$ and define

$$\Lambda^\gamma(\tau) = \Lambda_t \text{ if } \tau_t^\gamma \leq \tau < \tau_{t+1}^\gamma.$$

Thus $\Lambda^\gamma(\tau)$ is the continuous time interpolation of the discrete time process Λ_t under study. Here we make explicit the dependence on γ in $\Lambda^\gamma(\tau)$, which is implicit in Λ_t . Next, consider the differential equation

$$d\Lambda/d\tau = h(\Lambda(\tau)), \text{ where } h(\Lambda) \equiv E\mathcal{H}(\Lambda, X_t).$$

For $\Lambda(\tau)' = (b(\tau), R(\tau))$ we compute $h(b, R)' = (h_b(b, R), h_R(b, R))$ where

$$\begin{aligned} h_b(b, R) &= R^{-1} \sigma_v^2 ((\theta \rho - 1)b + \rho) \\ h_R(b, R) &= (\sigma_v^2 - R). \end{aligned}$$

This is the differential equation system (10) introduced in Section 3.1.

The differential equation $d\Lambda/d\tau = h(\Lambda(\tau))$ is well defined everywhere except at $R = 0$ and the RE solution $\bar{b} = \rho/(1 - \theta\rho)$, $\bar{R} = \sigma_v^2$ is globally stable. Let $\tilde{\Lambda}(\tau, a)$ denote the solution to this differential equation with initial condition $\Lambda(0) = a$. Finally, define

$$U^\gamma(\tau) = \gamma^{-1/2}(\Lambda^\gamma(\tau) - \Lambda(\tau, a)).$$

$U^\gamma(\tau)$ is the continuous time stochastic process which is used to approximate Λ_t for small γ . Proposition 7.8 of Evans and Honkapohja (2001) yields the following. For any fixed time $T > 0$, as $\gamma \rightarrow 0$, the stochastic process $U^\gamma(\tau)$, $0 \leq t \leq T$ converges weakly to the solution of the stochastic differential equation

$$dU(\tau) = D_\Lambda h(\tilde{\Lambda}(\tau, a))U(t)dt + \mathcal{R}^{1/2}(\tilde{\Lambda}(\tau, a))dW(\tau),$$

with initial condition $U(0) = 0$, where $W(\tau)$ is a standard vector Wiener process. Here \mathcal{R} is the 2×2 matrix with (i, j) element

$$\mathcal{R}^{ij}(\Lambda) = \sum_{k=-\infty}^{\infty} \text{cov}[\mathcal{H}_i(\Lambda, X_k), \mathcal{H}_j(\Lambda, X_0)], \text{ for } i, j = 1, 2.$$

Since $\tilde{\Lambda}(\tau, a)$ remains close to $\bar{\Lambda} = (\bar{b}, \bar{R})'$ for all $\tau \geq 0$ (and converges asymptotically to $\bar{\Lambda}$ as $\tau \rightarrow \infty$), for starting points near $\bar{\Lambda}$ (or for τ sufficiently large) $U^\gamma(\tau)$ can be approximated, for small γ by

$$dU(\tau) = D_\Lambda h(\bar{\Lambda})U(t)dt + \bar{\mathcal{R}}^{1/2}dW(\tau),$$

where $\bar{\mathcal{R}} \equiv \mathcal{R}(\bar{\Lambda})$. The stationary solution to this equation (e.g. see pp. 114-5 of Evans and Honkapohja (2001)) is a Gaussian process with mean zero and autocovariance function

$$EU(\tau)U(\tau - \hat{\tau})' = \rho(\hat{\tau}) = e^{sB}\mathcal{C} \text{ for } \hat{\tau} \geq 0, \text{ where}$$

$$B = D_\Lambda h(\bar{\Lambda}) \text{ and } \mathcal{C} = \int_0^\infty e^{uB}\bar{\mathcal{R}}e^{uB'} du.$$

From Theorem 7.9 of Evans and Honkapohja (2001) we also have the asymptotic result that for any sequences $\gamma_k \rightarrow 0$ and $\tau_k \rightarrow \infty$ the sequence $U^{\gamma_k}(\tau_k)$ converges in distribution to a normal random variable with mean 0 and variance \mathcal{C} . Computing the relevant quantities we have

$$B = D_\Lambda h(\bar{\Lambda}) = \begin{pmatrix} \theta\rho - 1 & 0 \\ 0 & -1 \end{pmatrix}, \text{ so that } e^{uB} = \begin{pmatrix} e^{u(\theta\rho-1)} & 0 \\ 0 & e^{-u} \end{pmatrix}.$$

Our focus is on the behavior of b_t , the first component of Λ_t . We have

$$\bar{\mathcal{R}}^{11} = \bar{R}^{-2} \sum_{k=-\infty}^{\infty} \Gamma(k), \text{ where}$$

$$\Gamma(k) = \text{cov}(Y_t, Y_{t-k}) \text{ and } Y_t = (1 + \theta\bar{b})v_t v_{t-1} - \bar{b}v_{t-1}^2.$$

The (1, 1) element of $e^{uB}\bar{\mathcal{R}}e^{uB'}$ is $\bar{\mathcal{R}}^{11}e^{2(\theta\rho-1)u}$ so that

$$C \equiv \mathcal{C}^{11} = \bar{\mathcal{R}}^{11} \int_0^\infty e^{2(\theta\rho-1)u} du = \frac{\bar{\mathcal{R}}^{11}}{2(1-\theta\rho)}.$$

Next, note that

$$\begin{aligned} Y_t &= (1 + \theta\bar{b})(\rho v_{t-1} + \varepsilon_t)v_{t-1} - \bar{b}v_{t-1}^2 \\ &= (\rho + \bar{b}(\theta\rho - 1))v_{t-1}^2 + (1 + \theta\bar{b})v_{t-1}\varepsilon_t \\ &= (1 - \rho\theta)^{-1}v_{t-1}\varepsilon_t. \end{aligned}$$

Thus

$$\Gamma(0) = \frac{\sigma_v^2 \sigma_\varepsilon^2}{(1 - \rho\theta)^2} \text{ and } \Gamma(k) = 0 \text{ for } k \neq 0,$$

so that

$$\bar{\mathcal{R}}^{11} = R^{-2}\Gamma(0) = \frac{\sigma_\varepsilon^2}{\sigma_v^2(1 - \rho\theta)^2} = \frac{1 - \rho^2}{(1 - \rho\theta)^2},$$

and

$$C = \frac{1 - \rho^2}{2(1 - \rho\theta)^3}.$$

The above implies that for small $\gamma > 0$ and large τ the stochastic process $U_b^\gamma(\tau) = \gamma^{-1/2}(b^\gamma(\tau) - \bar{b})$ is approximately Gaussian with mean zero, variance C and autocovariance $\rho(\hat{\tau}) = e^{(\theta\rho-1)\hat{\tau}}$. Since for $\gamma > 0$ small $\tau \approx \gamma t$ it follows that for small γ and large t the distribution of b_t is approximately normal with mean \bar{b} and variance γC and that the autocorrelation function $E((b_t - \bar{b})(b_{t-k} - \bar{b}))/E(b_t - \bar{b})^2$ is approximately equal to $e^{(\theta\rho-1)\gamma k}$. This establishes Proposition 1.

Proof of Proposition 2: The asymptotic bias is given by

$$\text{plim}_{T \rightarrow \infty} \hat{\beta}_T - 1 = \frac{\text{plim}_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T (F_t - s_t)u_{t+1}}{\text{plim}_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T (F_t - s_t)^2}.$$

From $s_t = (\theta b_{t-1} + 1)v_t$ and $F_t = \hat{E}_t s_{t+1} = b_{t-1}v_t$ we have $F_t - s_t = v_t(b_{t-1}(1 - \theta) - 1)$ and

$$\begin{aligned} T^{-1} \sum_{t=1}^T (F_t - s_t)^2 &= T^{-1} \sum_{t=1}^T (b_{t-1}(1 - \theta) - 1)^2 v_t^2 \\ &= (NP)^{-1} \sum_{k=1}^P \sum_{j=1}^N (b_{N(k-1)+j-1}(1 - \theta) - 1)^2 v_{N(k-1)+j}^2, \end{aligned}$$

where for convenience we look at T such that $T = PN$. Provided N and P are large and $\gamma > 0$ is sufficiently small relative to N we have $b_{N(k-1)+j-1}^2 \approx b_{N(k-1)}^2$ for $j = 1, \dots, N$ and

$$\begin{aligned} T^{-1} \sum_{t=1}^T (F_t - s_t)^2 &\approx P^{-1} \sum_{k=1}^P (b_{N(k-1)}(1 - \theta) - 1)^2 N^{-1} \sum_{j=1}^N v_{N(k-1)+j}^2 \\ &\approx \sigma_v^2 P^{-1} \sum_{k=1}^P (b_{N(k-1)}(1 - \theta) - 1)^2 \\ &\approx \sigma_v^2 (1 - \theta)^2 E (b_t - (1 - \theta)^{-1})^2 \\ &= \sigma_v^2 (1 - \theta)^2 \left(\gamma C + \left(\frac{\rho}{1 - \theta \rho} - \frac{1}{1 - \theta} \right)^2 \right). \end{aligned}$$

where we have used the weak law of law numbers first for $N^{-1} \sum_{j=1}^N v_{N(k-1)+j}^2 \xrightarrow{p} E v_t^2 = \sigma_v^2$ and then for $P^{-1} \sum_{k=1}^P (b_{N(k-1)}(1 - \theta) - 1)^2 \xrightarrow{p} E (b_t - (1 - \theta)^{-1})^2$.

From $u_{t+1} = s_{t+1} - F_t = ((\theta b_t + 1)\rho - b_{t-1})v_t + (\theta b_t + 1)\varepsilon_{t+1}$ we have

$$\begin{aligned} (F_t - s_t)u_{t+1} &= -(1 - \theta)(1 - \theta\rho) (b_{t-1} - (1 - \theta)^{-1}) \\ &\quad \times (b_{t-1} - \bar{b} - \theta\rho(1 - \theta\rho)^{-1}(b_t - b_{t-1})) v_t^2 \\ &\quad + (b_{t-1}(1 - \theta) - 1)(\theta b_t + 1)v_t \varepsilon_{t+1}. \end{aligned}$$

But $\text{plim } T^{-1} \sum_{t=1}^T (b_{t-1}(1 - \theta) - 1)(\theta b_t + 1)v_t \varepsilon_{t+1} = 0$ from the law of large numbers since $E_t(b_{t-1}(1 - \theta) - 1)(\theta b_t + 1)v_t \varepsilon_{t+1} = 0$. Thus for N and P large and $\gamma > 0$ sufficiently small relative to N we have

$$\begin{aligned} &T^{-1} \sum_{t=1}^T (F_t - s_t)u_{t+1} \\ &\approx -(1 - \theta)(1 - \theta\rho) \times \\ &\quad (NP)^{-1} \sum_{k=1}^P \sum_{j=1}^N (b_{N(k-1)+j-1} - (1 - \theta)^{-1}) \times \\ &\quad (b_{N(k-1)+j-1} - \bar{b} - \theta\rho(1 - \theta\rho)^{-1}(b_{N(k-1)+j} - b_{N(k-1)+j-1})) v_{N(k-1)+j}^2 \end{aligned}$$

$$\begin{aligned}
&\approx -(1-\theta)(1-\theta\rho)P^{-1} \times \\
&\quad \sum_{k=1}^P (b_{N(k-1)} - (1-\theta)^{-1}) (b_{N(k-1)} - \bar{b}) N^{-1} \sum_{j=1}^N v_{N(k-1)+j}^2 \\
&\approx -\sigma_v^2(1-\theta)(1-\theta\rho)^{-1}E(b_t - (1-\theta)^{-1})(b_t - \bar{b}) \\
&= -\sigma_v^2(1-\theta)(1-\theta\rho)\gamma C.
\end{aligned}$$

Taking the ratio $T^{-1} \sum_{t=1}^T (F_t - s_t)u_{t+1} / (T^{-1} \sum_{t=1}^T (F_t - s_t)^2)$ we obtain

$$\text{plim } \hat{\beta} - 1 = -\frac{(1-\theta\rho)\gamma C}{(1-\theta)\gamma C + [(1-\rho)^2 / ((1-\theta)(1-\theta\rho)^2)]}.$$

Substituting for C the expression obtained in Proposition 1, and simplifying, we get the result claimed.

Proof of Corollary 3: $B(\gamma, \theta, \rho) < 0$ follows immediately from 1, as do the limiting values at $\rho = 0$ and as $\rho \rightarrow 1$. The remaining properties follow by differentiation of $B(\gamma, \theta, \rho)$ with respect to each argument and using the inequalities $0 < \gamma \leq 1$, $0 < \theta < 1$ and $0 \leq \rho < 1$.

Proof of Corollary 4: The t-statistic is $t_{\hat{\beta}} = (\hat{\beta} - 1) / SE(\hat{\beta})$ where $SE(\hat{\beta}) = T^{-1/2}\hat{\sigma} / (T^{-1} \sum_{t=1}^T (F_t - s_t)^2)$ and $\hat{\sigma}^2 = T^{-1} \sum_{t=1}^T \hat{u}_t^2$. Here $\hat{u}_t = s_{t+1} - s_t - \hat{\beta}(F_t - s_t) = (\theta b_t + 1)v_{t+1} - (\theta b_{t-1} + 1)v_t - \hat{\beta}(b_{t-1}(1-\theta) - 1)$. Since \hat{u}_t converges in distribution, as $t \rightarrow \infty$, to a stationary random variable with finite second moments, it follows that $\hat{\sigma}^2$ converges in probability to a finite number. Similarly, at least for small γ , $(F_t - s_t)^2$ is asymptotically stationary with finite moments, and so $T^{-1} \sum_{t=1}^T (F_t - s_t)^2$ converges in probability to a finite positive number. Thus $SE(\hat{\beta}) \rightarrow 0$ as $T \rightarrow \infty$. Since for all $0 \leq \rho < 1$ we have that $\text{plim } \hat{\beta} - 1 < 0$ and the result follows.

Proof of Corollary 5: By Proposition 1, for $\gamma > 0$ sufficiently small, $E(b_t - \bar{b})^2 \approx \gamma C(\rho)$ where $C(\rho) = (1-\rho^2) / 2(1-\rho\theta)^3$. For any given $0 < \theta < 1$, C is continuous in ρ for all $0 \leq \rho \leq 1$. Therefore $C(\rho)$ is bounded uniformly over $0 \leq \rho \leq 1$ and thus over $0 \leq \rho < 1$. Thus, for any $\varepsilon > 0$ we can choose $\gamma > 0$ sufficiently small such that $E(b_t - \bar{b})^2 < \varepsilon$. Given this γ , Proposition 2 and Corollary 3 imply that by choosing $\rho < 1$ sufficiently large we can simultaneously ensure that $\text{plim } \hat{\beta} < \varepsilon$.

Table 2: Mean ratios for different parameter combinations.

		ρ				
		0.8	0.9	0.95	0.97	0.99
		Mean	Mean	Mean	Mean	Mean
θ	γ	ratio	ratio	ratio	ratio	ratio
0.6	0.005	0.997	0.998	0.997	0.998	0.998
	0.01	0.996	0.994	0.994	0.994	0.995
	0.02	0.989	0.989	0.989	0.991	0.991
	0.05	0.972	0.974	0.977	0.979	0.982
	0.1	0.951	0.958	0.962	0.964	0.971
0.9	0.005	0.997	1.001	0.998	1.001	0.999
	0.01	0.999	0.998	0.995	0.996	0.997
	0.02	0.993	0.994	0.994	0.998	0.995
	0.05	0.981	0.983	0.987	0.993	0.991
	0.1	0.967	0.977	0.981	0.981	0.983

Note: Results from 100 simulations with sample size of 20,000 after discarding first 4000 data points. The ratios given are $Mean(\hat{b}_{simulation})/E(\hat{b}_{theory})$.

Table 3: Standard deviation ratios for different parameter combinations.

		ρ				
		0.8	0.9	0.95	0.97	0.99
		S.D.	S.D.	S.D.	S.D.	S.D.
θ	γ	ratio	ratio	ratio	ratio	ratio
0.6	0.005	0.957	0.981	1.010	1.009	1.125
	0.01	0.996	1.004	1.029	1.099	1.235
	0.02	0.993	1.026	1.095	1.142	1.406
	0.05	1.014	1.082	1.197	1.298	1.694
	0.1	1.036	1.134	1.289	1.457	1.962
0.9	0.005	0.933	0.940	0.960	0.929	1.023
	0.01	0.992	0.987	0.975	1.065	1.164
	0.02	0.990	1.009	1.068	1.099	1.311
	0.05	1.008	1.058	1.152	1.237	1.546
	0.1	1.025	1.099	1.215	1.333	1.710

Note: Results from 100 simulations with sample size of 20,000 after discarding first 4000 data points. The ratios given are $\widehat{SD}(\hat{b}_{simulation})/SD(\hat{b}_{theory})$.

Table 4: Theoretical and simulated $\hat{\beta}$ for large samples.

		ρ							
		0.98		0.99		0.995		1.0	
θ	γ	$\hat{\beta}_{theory}$	$\hat{\beta}_{sim}$	$\hat{\beta}_{theory}$	$\hat{\beta}_{sim}$	$\hat{\beta}_{theory}$	$\hat{\beta}_{sim}$	$\hat{\beta}_{theory}$	$\hat{\beta}_{sim}$
0.6	0.01	0.83	0.94	0.71	0.78	0.55	0.39	0	-0.01
	0.02	0.71	0.75	0.55	0.35	0.38	-0.09	0	-0.02
	0.03	0.62	0.48	0.45	0.01	0.29	-0.22	0	-0.02
	0.05	0.50	0.09	0.33	-0.22	0.20	-0.24	0	-0.04
	0.1	0.32	-0.31	0.19	-0.31	0.11	-0.23	0	-0.06
0.9	0.01	0.95	1.00	0.91	1.00	0.83	0.94	0	-0.01
	0.02	0.91	0.96	0.83	0.87	0.71	0.53	0	-0.01
	0.03	0.87	0.90	0.77	0.67	0.62	0.09	0	-0.03
	0.05	0.80	0.69	0.66	0.21	0.49	-0.31	0	-0.08
	0.1	0.65	0.03	0.48	-0.50	0.31	-0.55	0	-0.11

Note: Results from 100 simulations with sample size of $T = 20,000$ after discarding first 20,000 data points. $\hat{\beta}_{sim}$ is the mean value across simulations. No intercept in test regression.

Table 5a: Simulated $\hat{\beta}$ and $t_{\hat{\beta}}$ for sample size $T = 120$.

		ρ							
		0.98		0.99		0.995		1.0	
θ	γ	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$
0.6	0.01	1.10	0.10	0.96	-0.02	0.74	-0.15	-0.82	-0.77
	0.02	0.87	-0.11	0.29	-0.51	-0.30	-0.76	-0.50	-0.96
	0.03	0.54	-0.47	-0.18	-0.90	-0.44	-1.11	-0.55	-1.15
	0.05	-0.04	-1.03	-0.46	-1.39	-0.45	-1.47	-0.36	-1.35
	0.1	-0.44	-2.10	-0.44	-2.14	-0.40	-2.12	-0.28	-1.85
0.9	0.01	1.22	0.20	1.39	0.25	1.70	0.26	-1.33	-0.52
	0.02	1.22	0.21	1.26	0.15	0.76	-0.10	-1.21	-0.66
	0.03	1.15	0.12	0.83	-0.10	0.08	-0.37	-1.01	-0.78
	0.05	0.86	-0.11	0.24	-0.47	-0.73	-0.95	-1.03	-0.96
	0.1	-0.07	-0.98	-0.82	-1.36	-1.23	-1.55	-0.76	-1.23

Note: Results from 1000 simulations with sample size of $T = 120$ after discarding the first 20000 data points. Table gives medians of $\hat{\beta}_{sim}$ and of $t_{\hat{\beta}}$ for testing $H_0 : \beta = 1$, without intercept in the test regression.

Table 5b: Simulated $\hat{\beta}$ and $t_{\hat{\beta}}$ for sample size $T = 360$.

		ρ							
		0.98		0.99		0.995		1.0	
θ	γ	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$
0.6	0.01	0.96	-0.06	0.77	-0.26	0.39	-0.51	-0.47	-1.14
	0.02	0.74	-0.47	0.19	-0.89	-0.31	-1.31	-0.29	-1.47
	0.03	0.48	-0.88	-0.14	-1.46	-0.43	-1.79	-0.20	-1.70
	0.05	0.02	-1.77	-0.33	-2.31	-0.37	-2.48	-0.15	-2.15
	0.1	-0.36	-3.51	-0.38	-3.63	-0.35	-3.65	-0.12	-2.99
0.9	0.01	1.09	0.15	1.24	0.25	1.19	0.14	-1.11	-0.85
	0.02	1.04	0.07	1.07	0.06	0.57	-0.34	-0.72	-1.01
	0.03	0.97	-0.06	0.69	-0.32	-0.14	-0.87	-0.53	-1.13
	0.05	0.74	-0.42	0.12	-1.02	-0.78	-1.55	-0.50	-1.35
	0.1	0.06	-1.61	-0.68	-2.29	-0.77	-2.50	-0.40	-1.80

Note: Results from 1000 simulations with sample size of $T = 120$ after discarding the first 20000 data points. Table gives medians of $\hat{\beta}_{sim}$ and of $t_{\hat{\beta}}$ for testing $H_0 : \beta = 1$, without intercept in test regression.

Table 5c: Effect of sample size on estimated $\hat{\beta}$.

Sample size	$\theta = 0.6$ and $\gamma = 0.01$				$\theta = 0.6$ and $\gamma = 0.02$			
	$\rho = 0.99$		$\rho = 1.0$		$\rho = 0.99$		$\rho = 1.0$	
	Intercept:		Intercept:		Intercept:		Intercept:	
	without	with	without	with	without	with	without	with
100	1.24	2.96	-0.85	-11.3	0.57	1.08	-1.04	-4.02
200	0.85	1.86	-0.37	-4.62	0.39	0.66	-0.35	-2.37
500	0.82	0.92	-0.31	-2.04	0.28	0.35	-0.26	-0.95
1000	0.79	0.89	-0.26	-0.75	0.32	0.33	-0.11	-0.46
2000	0.79	0.85	-0.13	-0.37	0.35	0.36	-0.09	-0.20
5000	0.79	0.81	-0.10	-0.17	0.30	0.30	-0.06	-0.11
10000	0.79	0.79	-0.05	-0.08	0.33	0.34	-0.05	-0.07
20000	0.79	0.79	-0.04	-0.07	0.31	0.31	-0.02	-0.03

Note: Results from 100 simulations after discarding the first 20000 data points. Table gives medians of $\hat{\beta}_{sim}$ for test regression without and with intercept.

Table 6: Constant-gain learning of fundamentals process for $T = 120$.

		$\theta = 0.6$			$\theta = 0.9$		
		γ			γ		
		0.01	0.05	0.1	0.01	0.05	0.1
ρ							
0.98	$\hat{\beta}$	0.74	-0.78	-0.61	-0.11	-3.78	-3.09
	$t_{\hat{\beta}}$	-0.23	-2.76	-3.60	-1.00	-6.21	-5.99
0.99	$\hat{\beta}$	-0.10	-0.79	-0.63	-1.00	-3.66	-3.15
	$t_{\hat{\beta}}$	-0.71	-2.82	-3.66	-1.32	-6.17	-6.02
1.0	$\hat{\beta}$	-1.07	-0.75	-0.62	-2.11	-3.39	-3.15
	$t_{\hat{\beta}}$	-1.22	-2.92	-3.63	-1.95	-5.95	-6.07

Note: Results from 1000 simulations with sample size of $T = 120$ after discarding the first 20000 data points. Table gives medians of $\hat{\beta}_{sim}$ and of the t-statistics $t_{\hat{\beta}}$ for testing $H_0 : \beta = 1$. Test regression includes intercept.

Table 7: Constant-gain learning with random structural shifts.

		Equilibrium gain γ	$T = 120$		$T = 360$	
			$\hat{\beta}$	$t_{\hat{\beta}}$	$\hat{\beta}$	$t_{\hat{\beta}}$
$\rho = 0.985$	$q = 0.02$	0.060	-0.58	-2.49	-0.06	-4.06
$\rho = 0.99$	$q = 0.05$	0.087	-0.23	-3.17	-0.11	-5.32
	$q = 0.10$	0.041	-0.81	-2.39	-0.24	-3.60
$\rho = 0.995$	$q = 0.05$	0.095	-0.36	-3.69	0.01	-5.43
	$q = 0.10$	0.061	-0.63	-2.71	-0.12	-4.47

Note: q = probability of structural shift. $\theta = 0.6$ and $\delta = 1$. Regime switching process with $L = 0.4$. γ is approximate Nash equilibrium gain. Results from 100 simulations after discarding first 20000 data points. Test regression include intercept. Table gives medians of $\hat{\beta}$ and $t_{\hat{\beta}}$.

Table 8: $\hat{\beta}$ for decreasing-gain learning with single structural break.

		$\theta = 0.6$		$\theta = 0.9$	
		$\bar{\gamma} = 0.5$	$\bar{\gamma} = 0.2$	$\bar{\gamma} = 0.5$	$\bar{\gamma} = 0.2$
$T = 120$	$\rho = 0.98$	-0.16	0.43	0.49	1.65
	$\rho = 0.99$	-0.30	0.00	0.38	1.65
	$\rho = 0.995$	-0.66	-0.97	-0.67	0.97
	$\rho = 0.997$	-1.17	-2.02	-1.73	-0.03
$T = 360$	$\rho = 0.98$	0.07	0.61	0.37	1.14
	$\rho = 0.99$	0.01	0.43	0.26	1.19
	$\rho = 0.995$	-0.01	0.27	0.22	1.08
	$\rho = 0.997$	-0.17	-0.11	0.11	0.84

Note: Results from 2000 simulations. Table gives medians of $\hat{\beta}$. Test regressions include intercept. 25% increase in mean of fundamentals.

References

- [1] Adam, Klaus, Marcet, Albert and Nicolini, Juan Pablo, “Stock Market Volatility and Learning,” mimeo., April 2006.
- [2] Benveniste, J., M. Metivier and P. Priouret, *Adaptive algorithms and stochastic approximations*. Springer-Verlag: Berlin, 1990.
- [3] Bischi, Gian-Italo and Marimon, Ramon, “Global Stability of Inflation Target Policies with Adaptive Agents.” *Macroeconomic Dynamics*, 2001, 5, 148-179.
- [4] Branch, William and Evans, George W. “A Simple Recursive Forecasting Model.” *Economics Letters*, 2006, 91, 158-166.
- [5] Brock, William A. and Hommes, Cars. “Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model.” *Journal of Economic Dynamics and Control*, 1998, 22, pp. 1236-1274.
- [6] Bullard, James and Duffy, John. “Learning and Excess Volatility.” *Macroeconomic Dynamics*, 2001, 5, 272-302.
- [7] Bullard, James and Eusepi, Stefano. “Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?” *Review of Economic Dynamics*, 2005, 8, pp. 324-359.
- [8] Bullard, James and Mitra, Kaushik. “Learning about Monetary Policy Rules.” *Journal of Monetary Economics*, 2002, 49, pp. 1105-1129.
- [9] Chakraborty, Avik. *Learning and the Forward Premium Puzzle*, May 2005, Ph.D. dissertation, University of Oregon.
- [10] Chakraborty, Avik. “Learning, the Forward Premium Puzzle and Market Efficiency,” 2006, mimeo., University of Tennessee, March 2006.
- [11] Chakraborty, Avik and Haynes, Stephen. “Econometrics of the Forward Premium Puzzle,” 2005, Univ. of Oregon Working Paper No. 2005-18.
- [12] Cho, In-Koo, Williams, Noah and Sargent, Thomas J. “Escaping Nash Equilibrium.” *Review of Economic Studies*, 2002, 69, pp. 1-40.

- [13] De Long, J. Bradford, Shleifer, Andrei, Summers, Lawrence H. and Waldmann Robert J. "Noise Trader Risk in Financial Markets." *Journal of Political Economy*, August 1990, 98(4), pp. 703-738.
- [14] Engel, Charles. "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence." *Journal of Empirical Finance*, 1996, 3, pp. 123-192.
- [15] Engel, Charles and West, Kenneth D. "Exchange Rates and Fundamentals." *Journal of Political Economy*, 2005, 113(3), pp. 485-517.
- [16] Evans, George W. and Honkapohja, Seppo. "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations." Federal Reserve Bank of San Francisco *Economic Review*, 1993 (1), pp. 3-13.
- [17] Evans, George W. and Honkapohja, Seppo. *Learning and expectations in macroeconomics*. Princeton, NJ: Princeton University Press, 2001.
- [18] Evans, George W. and Honkapohja, Seppo. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies*, 2003, 70, pp. 807-824.
- [19] Evans, George W. and Ramey, Garey. "Adaptive Expectations, Underparameterization and the Lucas Critique." *Journal of Monetary Economics*, 2006, 53, pp. 249-264.
- [20] Fama, Eugene F. "Forward and Spot Exchange Rates." *Journal of Monetary Economics*, November 1984, 14(3), pp. 319-338.
- [21] Frenkel, Jacob A. "A Monetary Approach to the Exchange Rate: Doctrinal Aspects and Empirical Evidence." *Scandinavian Journal of Economics* 1976, 78, pp. 200-224.
- [22] Froot, Kenneth and Thaler, Richard. "Anomalies: Foreign Exchange." *Journal of Economic Perspectives*, Summer 1990, 4(3), pp. 179-192.
- [23] Kasa, Kenneth. "Learning, Large Deviations and Recurrent Currency Crises." *International Economic Review*, 2004, 45, pp. 141-173.
- [24] Kim, Young Se and Mark, Nelson C. "Exchange Rates and Fundamentals under Adaptive Learning." November 27, 2005, mimeo.

- [25] Lewis, Karen K. “Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange.” *American Economic Review*, September 1989, 79(4), pp. 621-636.
- [26] Marcet, Albert and Nicolini, Juan P. “Recurrent Hyperinflations and Learning.” *American Economic Review*, December 2003, 93(5), pp. 1476-1498.
- [27] Marcet, Albert and Sargent, Thomas J. “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models.” *Journal of Economic Theory*, 1989, 48, pp. 337-68.
- [28] Mark, Nelson C. and Wu, Yangru. “Rethinking Deviations from Uncovered Interest parity: The Role of Covariance Risk and Noise.” *Economic Journal*, November 1998, 108, pp. 1686-1706.
- [29] McCallum, Bennett T. “A Reconsideration of the Uncovered Interest Parity Relationship.” *Journal of Monetary Economics*, 1994, 33, pp. 105-132.
- [30] McGough, Bruce. “Shocking Escapes.” *Economic Journal*, 2006, 116, pp. 507-528.
- [31] Meese, Richard A., and Rogoff, Kenneth. “Empirical Exchange Rate Models of the Seventies: Do They Fit Out-Of-Sample?” *Journal of International Economics*, February 1983, 14(1-2), pp. 3-24.
- [32] Mussa, Michael L. “The Exchange Rate, the Balance of Payments and Monetary and Fiscal Policy under a Regime of Controlled Floating.” *Scandinavian Journal of Economics* 1976, 78, pp. 229-248.
- [33] Orphanides, Athanasios and Williams, John C. “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in Ben S. Bernanke and Michael Woodford, eds., *The inflation targeting debate*. Chicago: NBER and the University of Chicago Press, 2005a, pp. 201-234.
- [34] Orphanides, Athanasios and Williams, John C. “The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations.” *Journal of Economic Dynamics and Control*, November 2005b, 29(11), pp. 1927-1950

- [35] Sargent, Thomas J. *The conquest of American inflation*. Princeton, NJ: Princeton University Press, 1999.
- [36] Timmermann, Allan. “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices.” *Quarterly Journal of Economics*, November 1993, 108(4), pp. 1135-1145.
- [37] Timmermann, Allan. “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning.” *Review of Economic Studies*, 1996, 63(4), pp. 523-57.
- [38] Williams, Noah. “Escape Dynamics in Learning Models,” Working Paper, Department of Economics, Princeton University, 2004.