Journal of Monetary Economics 55 (2008) 477–490

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Can perpetual learning explain the forward-premium puzzle? ☆

Avik Chakraborty [a], George W. Evans [b,*]

[a] University of Tennessee, USA
[b] Department of Economics, University of Oregon, Eugene, OR 97403-1285, USA

ARTICLE INFO

Article history:
Received 19 June 2007
Received in revised form
4 March 2008
Accepted 7 March 2008
Available online 12 March 2008

JEL classification:
D83
D84
F31
G12
G15

Keywords:
Learning
Exchange rates
Forward premium
Expectations

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 show that replacing rational expectations by discounted least-squares (or "perpetual") learning generates a negative bias that becomes strongest when the fundamentals are strongly persistent, i.e. close to a random walk. Perpetual learning can explain the forward-premium puzzle while simultaneously replicating other features of the data, including positive serial correlation of the forward premium and disappearance of the anomaly in other forms of the test.

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 rate with a smaller magnitude and often the opposite sign than specified by rational expectations (RE). 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.

As McCallum (1994) emphasized, the forward-premium puzzle (or "anomaly") coexists with a number of other stylized facts. The R^2 of the forward-premium regression is typically very low, and the forward premium itself is positively correlated. When alternative forms of the test are performed, the anomaly disappears. Our proposed explanation—that private agents are engaged in econometric learning—is capable not only of generating the forward-premium puzzle, but also of matching these other stylized features of the data.

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

☆ Support from National Science Foundation Grant no. SES-0617859 is gratefully acknowledged. We are greatly indebted to the referee for comments, to Stephen Haynes and Joe Stone for many helpful discussions and comments, and for comments received at the July 2006 "Learning Week" workshop, FRB of St. Louis, the November 2006 ECB Conference on Monetary Policy, Asset Markets and Learning, Frankfurt, and in numerous seminars.
* Corresponding author. Tel.: +1 541 346 4662; fax: +1 541 346 1243.
E-mail address: gevans@uoregon.edu (G.W. Evans).

0304-3932/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jmoneco.2008.03.002

Table 1
Regressions of quarterly depreciation on 3-month forward premium

	USD/GBP	USD/DEM	USD/JAY	GBP/DEM	GBP/JAY	DEM/JAY
$\Delta s_{t+1} = \alpha + \beta(F_t - s_t) + u_{t+1}$						
1976:I– 1994:I						
$\hat{\alpha}_{\text{OLS}}$	−1.340	0.638	3.294	1.622	7.702	1.041
	(0.895)	(0.886)	(0.964)	(1.116)	(1.687)	(0.648)
$\hat{\beta}_{\text{OLS}}$	−1.552	−0.136	−2.526	−0.602	−4.261	−0.755
	(0.863)	(0.839)	(0.903)	(0.782)	(1.133)	(1.042)

Note: Reproduced from Table 1 of Mark and Wu (1998).

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

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

β is unity and u_{t+1} is uncorrelated with the forward premium $F_t - s_t$. It follows that Plim $\hat{\beta} = 1$, 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 RE. 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 RE, 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 RE 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 Plim $\hat{\beta} = 1$.

A large volume of research has empirically tested the hypothesis $\beta = 1$, and concluded that the least-squares estimate $\hat{\beta}$ is often significantly less than 1. In fact, in the majority of cases, $\hat{\beta}$ is less than zero. Froot and Thaler (1990), McCallum (1994) and Engel (1996) provide comprehensive reviews of the empirical results. 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 rates.[1] The evidence thus strongly refutes the theoretical prediction that $\beta = 1$, and hence apparently contradicts the efficient market hypothesis. This is the renowned "forward-premium puzzle" (or "forward-premium anomaly").

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 expectations 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 explanation of the puzzle. Fama (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. This has led to a general skepticism of the risk-premium explanation.

The other main approach centers around the potential ability of nonrational expectations to explain the results. This potential is apparent from some of the other findings related to exchange-rate behavior. De Long et al. (1990) demonstrates that the presence of both rational and nonrational 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

[1] For more details about the data see Mark and Wu (1998). In Table 1, standard errors of estimates are in parentheses.

(1998) show that while the behavior of the variance and covariance of the risk premium as required by Fama (1984) does not have empirical support, the existence of noise traders in the market under certain numerical assumptions yields results compatible with the data. Furthermore, Lewis (1989) used Bayesian learning to provide an explanation for the forward-premium puzzle, though the model could not explain its persistence since the magnitude of the prediction errors shrinks, over time, to zero. Finally, for 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. However, as he notes, this and the view that expectations are less than fully rational are potentially complementary explanations.

Our paper is motivated by this research, which suggests the potential importance of deviations from RE in foreign exchange markets. If traders do not have perfectly RE, their forecast errors may be correlated with previous period's information and this would introduce an observed bias in the forward-premium regression results.[2] 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.

In fact, we require only a small and quite natural deviation from RE, 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; 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; McGough, 2006), asset prices (Timmermann, 1993; Brock and Hommes, 1998; Bullard and Duffy, 2001; Adam et al., 2006; Branch and Evans, 2008), and currency crises and exchange rates (Kasa, 2004; Kim, 2006).

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 RE due to learning. In particular our theoretical results imply the large sample limiting value of $\text{plim}(\hat{\beta}) = 0$ in the empirically realistic case in which the fundamentals follow (or approximate) a random walk. In addition, we show numerically that this downward bias is magnified in small sample sizes, yielding mean negative $\hat{\beta}$ in line with those observed empirically.

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 RE, they do not know the parameter values and must estimate them from observed data. In the model we analyze, the exchange rate s_t, under RE, 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 RE 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).[3] 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 RE. 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. Furthermore, for small gains, the value of b_t and the agents' forecasts will be quite close to RE.

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 for $\rho \to 1$, since in this limiting case the exchange rate under RE would follow a random walk, in accordance with the well-known empirical results of Meese and Rogoff (1983). A large econometric literature, initiated by the findings of Nelson and Plosser (1982), has established that most macroeconomic time-series either contain a unit root or a near-unit root. Empirically the case of interest thus concerns fundamentals processes with ρ close to or equal to one. Under learning this will lead to exchange rates that are close to a random walk. Our principal finding is that precisely in this case the downward bias in the forward-premium regression is substantial. Perpetual learning therefore appears capable of entirely explaining the forward-premium puzzle.

Some other features of the empirical evidence appear also to be consistent with perpetual learning. McCallum (1994) documents other regularities of the foreign exchange market that are related to the forward-premium puzzle. These include very low R^2 for the forward-premium regressions, estimates of β close to one in level form tests and in n-period

[2] Chakraborty 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.

[3] For a general discussion of constant-gain learning, see Evans and Honkapohja (2001, Chapter 14).

difference versions of the test with $n \geqslant 2$, and strong positive serial correlation of the forward premium itself. We will verify numerically that in addition to providing an explanation of the standard forward-premium anomaly, perpetual learning is able to reproduce these other key aspects of the data.

2. Framework

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[4]:

$$F_t = \hat{E}_t s_{t+1}, \tag{2}$$
$$i_t = i_t^* + F_t - s_t, \tag{3}$$
$$m_t - p_t = d_0 + d_1 y_t - d_2 i_t, \tag{4}$$
$$p_t = p_t^* + s_t. \tag{5}$$

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 currency at $t+1$, and $\hat{E}_t s_{t+1}$ denotes the market expectation of s_{t+1} held at time t. Eq. (2) assumes risk neutrality and Eq. (3) is the closed parity condition, with i_t and i_t^* the domestic and foreign interest rate, respectively. Eq. (4) represents money market equilibrium, where m_t is log money supply, p_t is log price level and y_t is log real GDP. 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.

These equations can be solved to yield the reduced form

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

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.[5] We will focus on the case in which v_t is an observable stationary AR(1) process[6]

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

with $0<\rho<1$. For application of the theoretical learning results we need to make the technical assumption that v_t has compact support.[7] 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). As discussed above, we are particularly interested in the case ρ close to one, but the theory we develop will be valid for any $0<\rho<1$.

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.

3. Formal results under learning

We now develop results for the monetary exchange-rate model under learning. We begin with theoretical results based on stochastic approximation tools, which is the main technique for analyzing least-squares learning in macroeconomic models. Since these formal results focus on approximations that are valid for small gains and large sample sizes, we also present numerical simulation results for a range of gain parameters and plausible sample sizes.

3.1. Stochastic approximation results

For theoretical convenience we examine the system

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

[4] See, for example, Frenkel (1976), Mussa (1976) and Engel and West (2005). This model is the simplest version of the "asset market approach" to exchange rates. Engel and West (2005) describe the various ways in which the model can be generalized.

[5] For the large country case see Chakraborty (2005, 2008).

[6] It would be straightforward to allow for an additional unobserved white noise shock.

[7] This 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.

where $\varepsilon_t \sim iid(0, \sigma_\varepsilon^2)$ and $0 \leqslant \rho < 1$. 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 solution

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

and market participants forecast using the known coefficient $\bar{b}$. Under learning, they instead estimate this coefficient by constant-gain least squares.[8,9] This is most conveniently expressed in recursive form.[10] The estimate based on data through time t is given by the algorithm

$$b_t = b_{t-1} + \gamma R_{t-1}^{-1} v_{t-1}(s_t - b_{t-1}v_{t-1}),$$

$$R_t = R_{t-1} + \gamma(v_t^2 - R_{t-1}), \tag{7}$$

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 equations[11]

$$\mathrm{d}b/\mathrm{d}\tau = R^{-1}\sigma_v^2((\theta\rho - 1)b + \rho),$$

$$\mathrm{d}R/\mathrm{d}\tau = \sigma_v^2 - R, \tag{10}$$

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 $\mathrm{E}b_t \to \bar{b}$ as $t \to \infty$.

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 RE, i.e. $b_t \to \bar{b}$ with probability one as $t \to \infty$. We instead focus on the natural modification in which ordinary least squares are replaced by constant-gain least squares, as above, so that γ is a small fixed positive number, e.g. $\gamma = 0.02$ or 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 are 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, Benveniste et al. (1990, Chapters 1 and 4, Part I). 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 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 Benveniste 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,[12] suggest values of the gain for quarterly data in the range $\gamma = 0.02$–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.

Under constant-gain learning, a natural result is obtained that goes beyond the decreasing gain asymptotic convergence result. RE 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 $t \to \infty$, with a mean equal to the RE value $\bar{b}$, and with a small variance around $\bar{b}$. We have the following[13]:

---

[8] If $\delta \neq 0$ 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.

[9] Under RE the regression obeys standard assumptions for $0 < \rho < 1$. When $\rho = 1$ agents would be estimating a cointegrating relationship.

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

[11] Technical details are given in the "Supplementary materials: technical details and proofs" for this paper, available on the Science Direct website.

[12] See Orphanides and Williams (2005b) and Branch and Evans (2006).

[13] For proofs of propositions, see the "Supplementary materials: technical details and proofs" for this paper, available at the Science Direct website. The proof of Proposition 1 applies results given in Evans and Honkapohja (2001, Chapter 7), which in turn are based on the stochastic approximation results of Benveniste et al. (1990).

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 $e^{-(1-\theta\rho)\gamma k}$.

It follows that, 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 RE arises as the limit in which $\gamma \to 0$, 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.[14] We have the following result:

Proposition 2. *Under the null hypothesis H_0 the asymptotic bias* 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 $0 \leqslant \theta < 1$ and $0 \leqslant \rho < 1$, 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 $0 \leqslant \theta < 1$, $0 \leqslant \rho < 1$ 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*

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

Corollary 3 implies that, for small γ, the value of plim $\hat{\beta}$ approaches 0 as $\rho \to 1$. 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 $\rho \to 1$.

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

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

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 Fig. 1, which shows the approximation plim $\hat{\beta} = 1 + B(\gamma, \theta, \rho)$, as a function of ρ over $0 \leqslant \rho \leqslant 1$, for fixed $\theta = 0.6$, and for three values $\gamma = 0.01$, 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 $0 \leqslant \rho < 1$, as $\gamma \to 0$ we approach the RE 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 plim $\hat{\beta}$ as $\rho \to 1$ for fixed γ: given γ, the asymptotic bias of $\hat{\beta}$ approaches -1 as $\rho \to 1$, regardless of the size of γ. The intuition for this result is given 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 $\hat{\rho} \leqslant \rho < 1$ we have both* $E(b_t - \bar{b})^2 < \varepsilon$ *for all t and* 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 RE 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.

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

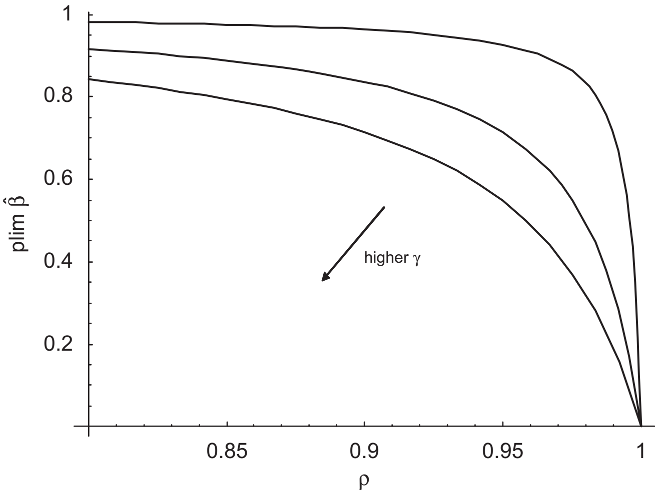

Fig. 1. Theoretical plim($\hat{\beta}$) for $\theta = 0.6$ and $\gamma = 0.01$, 0.05 and 0.10.

Table 2
Theoretical and simulated $\hat{\beta}$ for large samples

θ	γ	ρ							
		0.98		0.99		0.995		1.0	
		$\hat{\beta}_{\text{theory}}$	$\hat{\beta}_{\text{sim}}$	$\hat{\beta}_{\text{theory}}$	$\hat{\beta}_{\text{sim}}$	$\hat{\beta}_{\text{theory}}$	$\hat{\beta}_{\text{sim}}$	$\hat{\beta}_{\text{theory}}$	$\hat{\beta}_{\text{sim}}$
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. $\hat{\beta}_{\text{sim}}$ is the mean value across simulations. No intercept in test regression.

3.2. Numerical and small sample results

Our theoretical results are based on the asymptotic limit for large samples and small gains. Using Eqs. (7) and (9) we now simulate paths for b_t and s_t and investigate numerically, for realistic sample sizes and plausible gain parameters, the bias that arises in the forward-premium regression (1). Table 2 reports the simulation results for a large sample $T = 20{,}000$ and a range of gains $\gamma > 0$. Table 3a and b gives the small sample results, for $T = 120$ and 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 studies the impact of sample size in more detail. We focus on the empirically plausible cases of $\rho < 1$ close to one and the limit case $\rho = 1$, i.e. a pure random walk.

Table 2 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.[15] The key qualitative predictions of Proposition 2, and Corollary 3, hold in the numerical results of Table 2. In particular, an increase in γ or ρ (and the smaller value of θ) leads to a smaller value of plim $\hat{\beta}$. For $\gamma = 0.05$ or 0.10 (and in most cases for $\gamma = 0.03$) the simulation results in Table 2 show an even stronger downward bias in $\hat{\beta}$ than is predicted by our theoretical results for small $\gamma > 0$.

We next consider the small sample results given in Table 3. The sample size employed in Table 3a of $T = 120$ corresponds to 30 years of nonoverlapping quarterly data and in Table 3b of $T = 360$ corresponds to 30 years of

[15] In 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.

Table 3
Simulated $\hat{\beta}$ and $t_{\hat{\beta}}$ for sample size (a) $T=120$ and (b) $T=360$

θ	γ	ρ							
		0.98		0.99		0.995		1.0	
		$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$
(a)									
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
(b)									
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 sizes of $T=120$ and $T=360$ after discarding the first 20,000 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 4
Effect of sample size on estimated $\hat{\beta}$

Sample size	$\theta=0.6$ and $\gamma=0.02$			
	$\rho=0.99$		$\rho=1.0$	
	Intercept		Intercept	
	Without	With	Without	With
100	0.57	1.08	−1.04	−4.02
200	0.39	0.66	−0.35	−2.37
500	0.28	0.35	−0.26	−0.95
1000	0.32	0.33	−0.11	−0.46
2000	0.35	0.36	−0.09	−0.20
5000	0.30	0.30	−0.06	−0.11
10, 000	0.33	0.34	−0.05	−0.07
20, 000	0.31	0.31	−0.02	−0.03

Note: Results from 100 simulations after discarding the first 20,000 data points. Table gives medians of $\hat{\beta}_{sim}$ for test regression without and with intercept.

nonoverlapping 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

$$F_t - s_t = ((1-\theta)b - 1)v_t \tag{11}$$

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

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

which is the asymptotic bias that would result as $T \to \infty$ 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.[16] Calculating $a(b)$ is artificial since it holds $b_t = b$ fixed as $T \to \infty$, whereas under perpetual 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 4 we show, for selected parameter values of interest, how the differences between the asymptotic results of Table 2 and the small sample results of Table 3a and b depend on the sample size. In Table 4 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. For $\rho = 1$ we obtain negative values of $\hat{\beta}$ for all sample sizes, and if an intercept is included in the test regression the effect can be pronounced.[17] 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.[18] The theoretical prediction $\text{plim}(\hat{\beta}) = 0$ for $\rho = 1$ is in principle testable in data sets with large sample sizes.

One other small sample result, not shown in the tables, is nonetheless worth emphasizing: the variation in $\hat{\beta}$ across simulations is substantial for small samples. Consider, for example, $\theta = 0.6$, $\gamma = 0.05$ and a sample size of $T = 120$. For $\rho = 1$ the first and third quartiles for $\hat{\beta}$ are approximately $(Q1, Q3) = (-1.26, 0.26)$ if no intercept is included and $(Q1, Q3) = (-3.13, -0.77)$ with an intercept in the test regression. Similarly for $\rho = 0.99$ the quartiles are approximately $(Q1, Q3) = (-0.96, 0.29)$ without intercept and $(Q1, Q3) = (-1.42, 0.55)$ with intercept. For larger samples this range shrinks and it becomes small in very large samples. However, for typical sample sizes like $T = 120$ or 360 one would expect to see considerable variation in $\hat{\beta}$ across data sets. A substantial cross-country variation is indeed evident in Table 1.

On balance these numerical findings 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.

Although we have focussed attention on the ability of perpetual learning to explain the forward-premium puzzle results when least squares are used to estimate (1), it is important to know whether perpetual learning is consistent with the empirical exchange-rate evidence in other dimensions. McCallum (1994) pointed out two other features of the forward-premium regression that stand out empirically. First, when the forward model is tested in levels form, i.e. by estimating

$$s_{t+1} = \alpha + \beta F_t + u_{t+1}.$$

McCallum found that estimated values of β were close to one in the data in line with theory. To see if this feature is matched by the learning model we constructed analogs of Table 3a and b for the levels form regression. Our numerical findings are that the median values of $\hat{\beta}_{\text{sim}}$ lie in the range 0.975–1.000 for every entry in the tables. Thus the perpetual learning model successfully reproduces this feature of the data.[19]

Second, there are other forms of the test of bias in the forward exchange rate that are identical under the null but differ under the alternative. In particular, a multiperiod version of (1) could be examined, taking the form

$$s_{t+1} - s_{t+1-n} = \alpha + \beta(F_t - s_{t+1-n}) + u_{t+1} \quad \text{for } n = 1, 2, 3 \ldots .$$

The standard test, based on (1), sets $n = 1$. McCallum (1994) estimates the multiperiod version for $n = 1, 2$ and 3 and finds that the forward-premium anomaly is present *only* for $n = 1$. For $n = 2$ and 3 his estimates of β are close to one, in line with theory. Again we investigate this issue numerically using simulated data from our learning model. Table 5 presents the results for $n = 2$ and $T = 120$. These results provide a striking contrast to the corresponding results for $n = 1$ given in

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

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

[18] The reason for the greater downward bias of the estimate of β when an intercept is included, and the sample size is small, would be worth investigating in future research.

[19] The table for $T = 120$ can be found in the "Supplementary materials: extensions and further discussion," available on the Science Direct website. The relationship between the levels and the forward-premium regressions is discussed in Chakraborty and Haynes (2005).

Table 5
Simulated $\hat{\beta}$ and $t_{\hat{\beta}}$ for regression of $(s_{t+1} - s_{t-1})$ on $(F_t - s_{t-1})$ and sample size $T = 120$

θ	γ	ρ							
		0.98		0.99		0.995		1.0	
		$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$	$\hat{\beta}_{sim}$	$t_{\hat{\beta}}$
0.6	0.01	1.01	0.10	1.01	0.11	1.01	0.11	1.01	0.07
	0.02	1.02	0.16	1.02	0.19	1.02	0.17	1.01	0.10
	0.03	1.02	0.23	1.02	0.24	1.02	0.26	1.02	0.21
	0.05	1.04	0.42	1.04	0.39	1.02	0.26	1.02	0.22
	0.1	1.05	0.56	1.05	0.57	1.05	0.52	1.04	0.40
0.9	0.01	1.01	0.06	1.01	0.07	1.01	0.13	1.01	0.07
	0.02	1.02	0.21	1.02	0.20	1.02	0.23	1.02	0.22
	0.03	1.03	0.34	1.03	0.34	1.03	0.31	1.02	0.27
	0.05	1.04	0.41	1.04	0.42	1.05	0.52	1.05	0.50
	0.1	1.09	0.95	1.08	0.87	1.09	0.95	1.08	0.91

Note: Results from 1000 simulations with sample size of $T = 120$ after discarding the first 20,000 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 3a: for $n = 2$ there is no forward-premium puzzle, fully in line with the empirical results of McCallum (1994). Similar results to Table 5 are obtained for $n = 3$ and $T = 120$[20] and for $n = 2, 3$ and $T = 360$. Thus the perpetual learning model correctly predicts both the cases in which the forward-premium puzzle is found empirically and the cases in which it is not observed empirically.[21]

4. Discussion

What is the source of the downward bias to $\hat{\beta}$ that we have established theoretically and numerically? We now provide the intuition for the limiting case in which the fundamentals follow a random walk. 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 $\rho \to 1$ the fixed b asymptotic bias function (12) satisfies $a(b) \to -1$ 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?

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

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

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 $s_t = b v_{t-1} + c\varepsilon_t$, 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

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

since in the random walk case $\Delta s_{t+1} = (1 + \theta b)\varepsilon_{t+1}$, *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

[20] This table is given in the "Supplementary materials: extensions and further discussion" for this paper, available on the Science Direct website.

[21] There is also a sizable literature on the predictability of the exchange rate over multiple periods. We reserve for future research the theoretical implications and empirical results for this aspect of the perpetual learning model.

[22] Here 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 $\lim_{\rho \to 1}(\text{cov}_t(u_{t+1}, F_t - s_t)/\text{var}_t(F_t - s_t)) = -1$.

Fig. 2. Model generated time paths of log exchange rate, for $\theta = 0.6$, $\rho = 1$, under rational expectations and under learning with $\gamma = 0.04$.

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 $\rho \to 1$.

What *is*, perhaps, unexpected and surprising is that *arbitrarily small* deviations from RE yield a downward bias approaching -1 as ρ approaches 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 $0 \leqslant \rho \leqslant 1$ but $\text{var}(F_t - s_t) \to 0$ as $\rho \to 1$. 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 $\text{plim}(\widehat{\text{cov}}((F_t - s_t), u_{t+1})) \to 0$ and $\text{plim}(\widehat{\text{var}}(F_t - s_t)) \to 0$ as $\gamma \to 0$ but the ratio is close to -1 for $\rho < 1$ near 1. Furthermore, as our numerical computations have shown, for $\rho = 1$ itself, the system under learning is well-behaved, with an asymptotic bias for $\hat{\beta}$ of -1 and an even stronger downward bias for $\hat{\beta}$ in small samples.

Figs. 2 and 3 give the results of a typical simulation of our model over $T = 200$ periods, with parameters set at $\theta = 0.6$ and $\rho = 1$.[23] Fig. 2 gives the time paths for the log of the exchange rate under RE, and under least-squares learning with constant gain $\gamma = 0.04$. The two time paths, which are generated by the same sequence of exogenous random shocks, are almost indistinguishable. Some mild "overshooting" under learning can be seen under close inspection, which is another immediate implication of learning for $\rho \to 1$.[24] Fig. 3 gives the corresponding simulation results for depreciation and the forward premium under learning. The estimated value from the forward-premium regression in this simulation is $\hat{\beta} = -0.66$, with a t-statistic of $t = -2.23$, a typical illustration of the forward-premium puzzle.[25] Similar results are obtained for values of ρ that are close to but less than one.

For comparison Fig. 4 presents the monthly depreciation and forward-premium data (from Bloomberg) for the log Canadian dollar price of the US dollar, December 1988–September 2005. The data and the regression estimate $\hat{\beta} = -0.60$, with a t-statistic of $t = -2.04$, are quite similar to the simulated results under learning given in Fig. 3.[26] Note that the explanatory power of the forward-premium regressions is low in each, which is another standard finding in the data. This phenomenon was stressed by McCallum (1994).[27]

A feature of our analysis worth emphasizing is that exchange-rate paths under learning and under RE can be almost identical, and yet lead to very different results for the test of the forward-premium bias. This finding, illustrated in Figs. 2 and 3, results from two factors. First, with small gains the model under learning has estimates b_t that are close to the RE value. As a result forecasts and hence exchange rates are close to what would be observed under RE. Second, for ρ close to or equal to one, variation over time of the forward premium $F_t - s_t$ is small, leading to the low explanatory power of $F_t - s_t$ in predicting exchange-rate changes that was noted by McCallum (1994). Because the bias of $\hat{\beta}$ is determined by the *ratio* of $\widehat{\text{cov}}(\Delta s_{t+1}, F_t - s_t)$ to $\widehat{\text{var}}(F_t - s_t)$, the downward bias resulting from $\widehat{\text{cov}}(\Delta s_{t+1}, F_t - s_t) < 0$ can still be large, despite the

[23] The standard deviation of the innovation to the fundamentals has been chosen so that the scale for depreciation is similar to that seen in the Canadian–US data.

[24] The somewhat greater variation of s_t under learning is consistent with the excess volatility results of Kim (2006). The extent of overshooting and excess volatility seen in our simulations depends on the parameters θ, ρ and γ.

[25] The results shown are typical, but we remark that there is a wide variation across simulations. For example, for this parameter setting, 80% of the values of $\hat{\beta}$ lie in the interval $(-3.67, 0.13)$.

[26] For a comprehensive empirical analysis see Chakraborty (2008).

[27] See also Engel and West (2005).

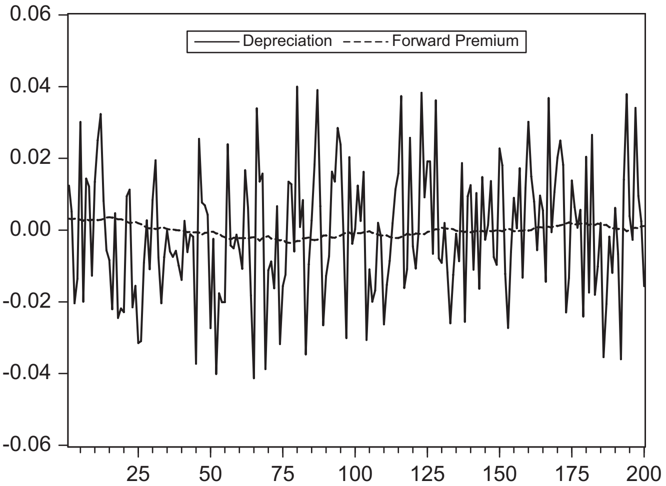

Fig. 3. Model generated data with learning, for $\theta = 0.6$, $\rho = 1$ and $\gamma = 0.04$. Test statistic $\hat{\beta} = -0.66$.

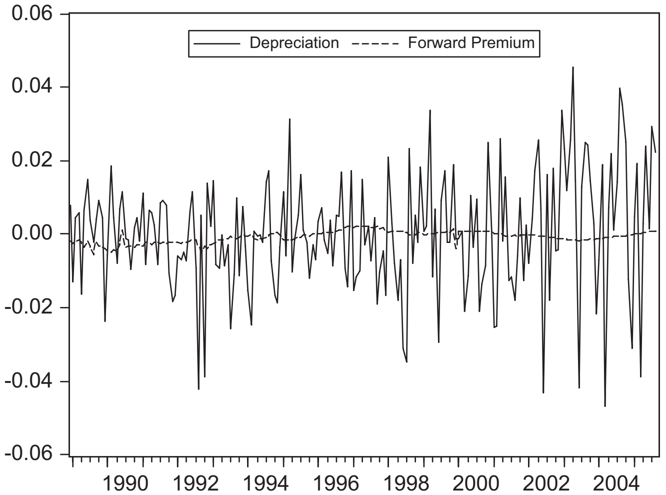

Fig. 4. Time path of USD/CAD (Canadian Dollar) log exchange rate (monthly data December 1988–September 2005). Test statistic $\hat{\beta} = -0.60$.

minimal impact on the s_t path. An implication of these results is that profiting from the deviation from RE would require taking large and risky positions in the foreign exchange market.

In the Introduction we argued that nonrational expectations were a potential explanation of the forward-premium puzzle. A natural question to raise is whether other types of deviation from RE would yield results similar to perpetual learning. To investigate this, consider modeling expectations as equal to the RE value plus a white noise expectation shock, i.e. $F_t = \hat{\mathrm{E}}_t s_{t+1} = \bar{b} v_t + \eta_t$, where $\eta_t \sim iid(0, \sigma_\eta^2)$ is an exogenous process independent of the v_t process. It can be shown that this leads to

$$F_t - s_t = (1 - \theta\rho)^{-1}(\rho - 1)v_t + (1 - \theta)\eta_t,$$
$$\Delta s_{t+1} = (1 - \theta\rho)^{-1}(\rho - 1)v_t - \theta\eta_t + (1 - \theta\rho)^{-1}\varepsilon_{t+1} + \theta\eta_{t+1}.$$

For $|\rho| < 1$ the asymptotic bias $\operatorname{plim} \hat{\beta} - 1 = \operatorname{cov}(\Delta s_{t+1}, F_t - s_t)/\operatorname{var}(F_t - s_t) - 1$ can be computed to be

$$\operatorname{plim} \hat{\beta} - 1 = -\frac{(1 - \theta)(\sigma_\eta^2/\sigma_v^2)}{(1 - \rho)^2/(1 - \rho\theta)^2 + (1 - \theta)^2(\sigma_\eta^2/\sigma_v^2)}.$$

Hence there is indeed a downward bias. For given ρ and σ_v^2 the asymptotic bias tends to zero as $\sigma_\eta^2 \to 0$. However, for given σ_η^2 and σ_v^2 (or for given σ_η^2 and σ_ε^2) the asymptotic bias tends to $-(1 - \theta)^{-1}$ as $\rho \to 1$. Thus as the fundamentals tend to a random walk, there is again a strong downward bias that could potentially explain the forward-premium puzzle. However, there are other empirical implications that in our view make the perpetual learning model clearly preferable.

In the case of white noise expectation shocks, the autocorrelation of the forward premium $F_t - s_t$ is low for ρ near one. In fact, it can be seen that $F_t - s_t \to (1-\theta)\eta_t$ as $\rho \to 1$, so that the first-order autocorrelation $\text{cor}(F_t - s_t, F_{t-1} - s_{t-1}) \to 0$ as $\rho \to 1$. This is counterfactual since the forward premium is known to be heavily positively serially correlated. For example, McCallum (1994) reports estimated first-order autocorrelations of $F_t - s_t$, for monthly data over 1978.01–1990.07 of 0.829, 0.897 and 0.906 for the \$/DM, \$/£ and \$/yen, respectively (and first-order autocorrelations are known to be downward biased in small samples). Similarly, for the Canadian data illustrated above, the estimated first-order autocorrelation is 0.902. In contrast to the case of white noise expectation shocks, strong positive autocorrelation of the forward premium will be present in the perpetual learning model for ρ near one. Under learning $F_t - s_t = ((1-\theta)b_{t-1} - 1)v_t$, and Proposition 1 implies that for small gains γ the agents' parameter estimates b_t are strongly positively serially correlated. It follows that as $\rho \to 1$ the forward premium itself will be positively serially correlated. Intuitively, under constant-gain least-squares learning b_t is revised each period by small amounts, leading to positive correlation in b_t around the RE value. This in turn induces positive serial correlation of the forward premium $F_t - s_t$, and when ρ is close to one the degree of positive serial correlation of $F_t - s_t$ will be high, as found in the data.

We conclude our discussion with brief comments on several extensions.[28] Our formulation of learning based on (6) uses a "self-referential" approach in which s_t depends on $\hat{E}_t s_{t+1}$ as well as on the current fundamentals v_t. An alternative is to solve (6) forward for s_t as a present-value type of sum of current and expected future v_{t+j}, $j = 0, 1, 2, \ldots$. Under learning agents estimate the v_t process using constant-gain least squares. This formulation leads to similar results to the self-referential version. Other variations incorporate actual structural change into the model: either unobserved structural change of the type described in Evans and Ramey (2006), or infrequent but observed structural change. In both cases we find that for ρ near one the forward-premium puzzle arises under learning. Thus our central findings appear robust to the detailed specification of the model under learning.

We have developed our results under learning using the basic canonical monetary exchange-rate model. We have isolated a simple and powerful mechanism for explaining the forward-premium anomaly, based on the combination of econometric learning by private agents and random walk or near random-walk behavior of the fundamentals driving the exchange rate. More elaborate models would allow for risk aversion, incomplete price adjustment, incomplete information processing, or heterogenous expectations. We anticipate that the impact of learning will remain prominent in more complex exchange-rate models.

5. 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 RE 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. Specifically, we obtain for this case the limiting value of $\text{plim}(\hat{\beta}) = 0$ in sharp contrast to the RE coefficient value of $\beta = 1$. In our numerical results we find that this downward bias is magnified in small samples, and that the large negative values of this coefficient observed empirically are fully consistent with our theory. 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.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at 10.1016/j.jmoneco.2008.03.002.

References

Adam, K., Marcet, A., Nicolini, J.-P., 2006. Stock market volatility and learning. Mimeo.
Benveniste, J., Metivier, M., Priouret, P., 1990. Adaptive Algorithms and Stochastic Approximations. Springer, Berlin.
Bischi, G.-I., Marimon, R., 2001. Global stability of inflation target policies with adaptive agents. Macroeconomic Dynamics 5, 148–179.
Branch, W.A., Evans, G.W., 2006. A simple recursive forecasting model. Economics Letters 91, 158–166.
Branch, W.A., Evans, G.W., 2008. Learning about risk and return: a simple model of bubbles and crashes. Mimeo.
Brock, W.A., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics and Control 22, 1236–1274.

[28] For details see the "Supplementary materials: extensions and further discussion," available on the Science Direct website.

Bullard, J., Duffy, J., 2001. Learning and excess volatility. Macroeconomic Dynamics 5, 272–302.
Bullard, J., Eusepi, S., 2005. Did the great inflation occur despite policymaker commitment to a Taylor rule? Review of Economic Dynamics 8, 324–359.
Bullard, J., Mitra, K., 2002. Learning about monetary policy rules. Journal of Monetary Economics 49, 1105–1129.
Chakraborty, A., 2005. Learning and the forward premium puzzle. Ph.D. Dissertation, University of Oregon.
Chakraborty, A., 2008. Learning, the forward premium puzzle and market efficiency. Macroeconomic Dynamics, forthcoming.
Chakraborty, A., Haynes, S., 2005. Econometrics of the forward premium puzzle. Working Paper No. 2005-18, University of Oregon (revised 2008).
Cho, I.-K., Williams, N., Sargent, T.J., 2002. Escaping Nash equilibrium. Review of Economic Studies 69, 1–40.
De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703–738.
Engel, C., 1996. The forward discount anomaly and the risk premium: a survey of recent evidence. Journal of Empirical Finance 3, 123–192.
Engel, C., West, K.D., 2005. Exchange rates and fundamentals. Journal of Political Economy 113, 485–517.
Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.
Evans, G.W., Honkapohja, S., 2003. Expectations and the stability problem for optimal monetary policies. Review of Economic Studies 70, 807–824.
Evans, G.W., Ramey, G., 2006. Adaptive expectations, underparameterization and the Lucas critique. Journal of Monetary Economics 53, 249–264.
Fama, E.F., 1984. Forward and spot exchange rates. Journal of Monetary Economics 14, 319–338.
Frenkel, J.A., 1976. A monetary approach to the exchange rate: doctrinal aspects and empirical evidence. Scandinavian Journal of Economics 78, 200–224.
Froot, K., Thaler, R., 1990. Anomalies: foreign exchange. Journal of Economic Perspectives 4 (3), 179–192.
Kasa, K., 2004. Learning, large deviations and recurrent currency crises. International Economic Review 45, 141–173.
Kim, Y.S., 2006. Exchange rates and fundamentals under adaptive learning. Mimeo.
Lewis, K.K., 1989. Changing beliefs and systematic rational forecast errors with evidence from foreign exchange. American Economic Review 79, 621–636.
Marcet, A., Nicolini, J.-P., 2003. Recurrent hyperinflations and learning. American Economic Review 93, 1476–1498.
Marcet, A., Sargent, T.J., 1989. Convergence of least-squares learning mechanisms in self-referential linear stochastic models. Journal of Economic Theory 48, 337–368.
Mark, N.C., Wu, Y., 1998. Rethinking deviations from uncovered interest parity: the role of covariance risk and noise. Economic Journal 108, 1686–1706.
McCallum, B.T., 1994. A reconsideration of the uncovered interest parity relationship. Journal of Monetary Economics 33, 105–132.
McGough, B., 2006. Shocking escapes. Economic Journal 116, 507–528.
Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: Do they fit out-of-sample? Journal of International Economics 14, 3–24.
Mussa, M.L., 1976. The exchange rate, the balance of payments and monetary and fiscal policy under a regime of controlled floating. Scandinavian Journal of Economics 78, 229–248.
Nelson, C.R., Plosser, C.I., 1982. Trends and random walks in macroeconomic time-series: some evidence and implications. Journal of Monetary Economics 10, 139–162.
Orphanides, A., Williams, J.C., 2005a. Imperfect knowledge, inflation expectations, and monetary policy. In: Bernanke, B.S., Woodford, M. (Eds.), The Inflation Targeting Debate. NBER, University of Chicago Press, Chicago, pp. 201–234.
Orphanides, A., Williams, J.C., 2005b. The decline of activist stabilization policy: natural rate misperceptions, learning and expectations. Journal of Economic Dynamics and Control 29, 1927–1950.
Sargent, T.J., 1999. The Conquest of American Inflation. Princeton University Press, Princeton, NJ.
Timmermann, A., 1993. How learning in financial markets generates excess volatility and predictability in stock prices. Quarterly Journal of Economics 108, 1135–1145.
Williams, N., 2004. Escape dynamics in learning models. Working Paper, Department of Economics, Princeton University.