



Resolving the Unbiasedness and Forward Premium Puzzles

Daniel L. Thornton

Working Paper 2007-014A
<http://research.stlouisfed.org/wp/2007/2007-014.pdf>

March 2007

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

Resolving the Unbiasedness and Forward Premium Puzzles

Daniel L. Thornton

Federal Reserve Bank of St. Louis

Phone (314) 444-8582

FAX (314) 444-8731

Email Address: thornton@stls.frb.org

November 2007

Abstract

There are two unresolved puzzles in the empirical foreign exchange literature. The first is the finding that tests of forward rate unbiasedness using the forward rate and forward premium equations yield markedly different conclusions. A companion puzzle—the *forward premium puzzle*—is the fact that the forward premium incorrectly predicts the direction of the subsequent change in the spot rate, which implies a massive rejection of uncovered interest parity. This paper resolves both puzzles.

JEL Classification: F31, G12, G15

Key words: uncovered interest parity, covered interest parity, forward premium, unbiasedness

The views expressed here are the author's and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of St. Louis. I would like to thank Menzie Chinn, Cletus Coughlin, David Dickey, Lucio Sarno, Giorgio Valente, and especially Ben McCallum, for useful comments and suggestions and John McAdams and Daniel McDonald for helpful research assistance.

1. Introduction

There are two unresolved empirical puzzles in the foreign exchange literature. The first puzzle arises when unbiasedness is tested using what McCallum (1994) calls the *forward rate* and *forward premium* equations. If the forward rate is an unbiased predictor of the future spot rate, the estimated slope coefficient should be insignificantly different from 1.0 for either specification. The puzzle arises because estimates of the slope coefficient from the forward rate specification are frequently close to, and sometimes insignificantly different from 1.0, while the estimates of the slope coefficient from the forward premium specification are frequently negative and nearly always significantly different from 1.0. The second puzzle, known as the *forward premium puzzle*, is the fact that the forward premium incorrectly predicts the direction of the subsequent change in the spot rate. Both puzzles have been investigated extensively.

This paper resolves both puzzles. The unbiasedness puzzle is resolved by showing that while the estimated slope coefficients from the forward rate and forward premium specifications are identical under the null hypothesis, they are not even comparable under the alternative. Hence, the fact that the estimates from the two specifications are very different is a consequence of the fact that, for a variety of reasons, unbiasedness does not hold perfectly. The noncomparability of the estimates from these alternative specifications also means that contrary to the suggestion in the literature, unbiasedness cannot be tested using the forward premium equation. This is consequence of the fact that comparable estimates of the slope coefficient do not exist under both the null and alternative hypotheses. The noncomparability estimates from these equations is invariant to whether exchange rates are $I(0)$ or $I(1)$ processes. Monte Carlo experiments

show that even tiny violations of unbiasedness can result in dramatically different estimates of the slope coefficient from the forward premium equation.

The resolution of the forward premium puzzle is a consequence of the near-random-walk behavior of exchange rates. Specifically, the near-random-walk behavior of the spot rate implies that there should be a weak and perhaps statistically insignificant relationship between changes in the spot rate and the forward premium. The negative relationship often reported in the literature is due to the strong positive correlation between the forward premium and the difference between the domestic and foreign interest rates, which is a consequence of covered interest parity (CIP), and the correlation between the change in the spot rate and the difference between the domestic and foreign interest rates. The latter correlation, which is sometimes positive and other times negative, is driven by economic fundamentals. Specifically, the correlation is negative when increases in the interest rate differential are due to expectations of higher domestic inflation relative to foreign inflation and positive when the behavior of the interest rate differential is due to changes in relative real interest rates. The estimates tend to be negative on average for many countries because, more often than not, the interest rate differential is driven by changes in expectations of relative inflation rates.

This resolution of the forward premium puzzle is shown to be consistent with the observed exchange rate data over different sample periods and a variety of exchange rates. Moreover, it accounts for a remarkably large proportion of the time variation in the estimates of the slope parameter from the forward premium equation.

The idea that the forward premium puzzle is due to expectations not being rational has its roots with Meese and Rogoff (1983), who showed that forecasts of the

exchange rate based on structural economic models were no better than random walk forecasts. Indeed, evidence that economic models can consistently outperform the simple random walk model remains elusive (e.g., Clarida, Sarno, Taylor, and Valente, 2003; Cheung, Chinn, and Pascual, 2005; Groen, 2005; Kilian and Taylor, 2003; and Rossi, 2006). Recent research has focused on the coexistence of economic fundamentals with the random walk behavior of exchange rates. For example, Engel and West (2005) show that “near” random walk behavior arises in a rational expectations present-value model if one of the economic fundamentals is nonstationary and the discount factor is “sufficiently” close to 1.0. Also, building on Frankel’s (1976) argument that economic fundamentals are detached from exchange rates because of swings in expectations about the future exchange rate, Bacchetta and van Wincoop (2004, 2006) show that the apparent disconnect between economic fundamentals and exchange rates can arise if market participants regularly change the weight that they give to various economic fundamentals—a behavior confirmed by Cheung and Chinn’s (2001) survey of U.S. foreign-exchange traders. Careful empirical work by Sarno and Valente (2006) suggests that such behavior can account for the observed behavior of some exchange rates. My resolution of the forward premium puzzle is in the spirit of this literature.

The remainder of the paper is divided into six sections. Section 2 presents the unbiasedness hypothesis and discusses the forward rate and forward premium tests for unbiasedness. Section 3 demonstrates the noncomparability of estimates of the slope coefficient from the forward rate and forward premium equations. The qualitative importance of the noncomparability of the estimates of the slope coefficient from the forward rate and forward premium equations for trivial violations of the unbiasedness

condition is examined in Section 4. Section 5 derives the forward premium equation used to test UIP. Section 6 presents a resolution of the forward premium puzzle based on near-random-walk behavior of exchange rates, CIP, and economic fundamentals. Section 7 concludes.

2. Tests of Unbiasedness

Unbiasedness of the forward exchange rate implies that

$$(1) \quad E_{t-1}s_t = f_{t-1},$$

where s_t denotes the log of the spot exchange rate expressed in terms of the home currency; f_t denotes the log of the 1-period-ahead forward exchange rate—the home-currency price of foreign exchange to be paid for and delivered in period t ; and E_{t-1} denotes the expectation conditional on all information available before s_t and f_t are determined.

Because $E_{t-1}s_t$ is unobservable, unbiasedness is most frequently tested under the assumption of rational expectations, i.e.,

$$(2) \quad s_t = E_{t-1}s_t + v_t,$$

where v_t is an *i.i.d.* random variable, $v_t \sim f(0, \sigma_v^2)$, that is independent of information available at period $t-1$. Substituting (1) into (2) yields

$$(3) \quad s_t = f_{t-1} + v_t.$$

Hodrick (1987) notes that (3) motivated researchers to test the unbiasedness proposition by estimating

$$(4) \quad s_t = \alpha + \beta f_{t-1} + v_t$$

and testing the hypothesis that $\alpha = 0$ and $\beta = 1$. In practice, however, only the hypothesis $\beta = 1$ is of concern (McCallum, 1994). Early investigations of forward rate unbiasedness (e.g., Frenkel, 1976, 1981; and Levich, 1978) relied on (4).

Following Meese and Singleton's (1982) evidence that foreign exchange rates are nonstationary, it has been common (indeed, nearly universal) to test the unbiasedness by estimating

$$(5) \quad \Delta s_t = \alpha + \beta(f_{t-1} - s_{t-1}) + \omega_t$$

and testing the hypothesis $\alpha = 0$ and $\beta = 1$.¹ Longworth (1981) was one of the first to test foreign exchange market efficiency using (5) and McCallum (1994) observes that tests of unbiasedness using (5) are "too numerous to mention." The preference for (5) over (4) stems from the apparent nonstationarity of exchange rates and the well-known fact that (4) and (5) are equivalent under the null hypothesis.

While it has been recognized for some time that estimates from the forward rate and forward premium equations are not identical when the maintained hypothesis is false (e.g., McCallum, 1994; Goodhart, McMahon, and Ngama, 1997; Sarno and Taylor, 2002; and Maynard, 2003), essentially unnoticed is the fact that the estimated slope coefficients from (4) and (5) are not comparable under the alternative hypothesis.² Indeed,

¹ Meese and Singleton (1982) suggested that testing the unbiasedness proposition using (4) "may be inappropriate, since the asymptotic distribution theory employed may not be valid" and suggested that (5) provided a better way to test the unbiasedness of the forward rate. Hodrick (1987, p. 30) provided support for this conclusion by demonstrating how (4) may generate estimates "supportive of the null hypothesis...even though it can be strongly rejected in specifications that use data in a form more likely to satisfy the assumption of time series stationarity." Specifically, if s_t and f_t are integrated order one, i.e., I(1), but cointegrated with a cointegrating vector (1, -1), estimates of β from (4) will be 1 regardless of the true value of β .

² The exception is McCallum (1994, p. 118), who notes that the "unconditional growth rate implications" from these two equations "are quite distinct when $\beta \neq 1.0$," i.e., when the null hypothesis does not hold.

considerable effort has been devoted to explaining why the estimates are so different (e.g., Hodrick, 1987; McCallum, 1994; Goodhart, McMahon, and Ngama, 1997; Barnhart, McNown, and Wallace, 1999; Maynard and Phillips, 2001; and Maynard, 2003).

3. The Noncomparability of Forward Rate and Forward Premium Tests of Unbiasedness

To see that estimates of β from (4) and (5) are not comparable when the null hypothesis does not hold, note that (4) is obtained after a normalization has been imposed. This is seen more clearly by considering the unnormalized version of (4), i.e.,

$$(6) \quad \lambda s_t = \alpha^* + \beta^* f_{t-1} + \varepsilon_t.$$

The normalization determines both the unit of measure and the direction of the minimization of the sum of squares.³ To see this, note that (6) can be rewritten as

$$(7) \quad s_t = \alpha^* / \lambda + (\beta^* / \lambda) f_{t-1} + \varepsilon_t / \lambda.$$

In the case of single-equation ordinary least squares (OLS) estimation, it is common practice to impose the normalization before (7) is estimated. The most convenient, and, hence, most commonly used, normalization is $\lambda = 1$. With this normalization, (7) is (4).

Moreover, with this normalization, $v_t = \varepsilon_t / \lambda$ is the distance from an orthogonal projection of s_t onto the space spanned by f_{t-1} .

Of course, the same analysis can be applied to (5). That is, before normalizing, (5) can be written as

$$(8) \quad \theta(s_t - s_{t-1}) = \alpha^{**} + \beta^{**} (f_{t-1} - s_{t-1}) + \mathcal{G}_t.$$

As before, (8) can be normalized and rewritten as

³ See Chow (1964) for a detailed analysis.

$$(9) \quad (s_t - s_{t-1}) = \alpha^* / \theta + (\beta^* / \theta)(f_{t-1} - s_{t-1}) + \mathcal{G}_t / \theta,$$

which, under the normalization that $\theta = 1$, can be re-parameterized as

$$(10) \quad (s_t - s_{t-1}) = \alpha'' + \beta''(f_{t-1} - s_{t-1}) + \omega_t.$$

This establishes that the unit of measure of ω_t is the distance from an orthogonal projection of $s_t - s_{t-1}$ onto $f_{t-1} - s_{t-1}$.

Note that (10) is identical to (5) except that β'' replaces β to emphasize the fact that the slope coefficients in (4) and (5) may be different. Indeed, they will be the same if and only if $\beta^* / \lambda = \beta^{**} / \theta$. This will occur if $\beta^* = \beta^{**}$ and $\lambda = \theta$. At any point in the parameter space other than $\beta = \beta'' = 1$, not only are the estimates of β from (4) and (5) different, they are not comparable; they have different units of measure, $v_t \neq \omega_t$.⁴

Despite the fact that (5) is routinely used to test the unbiasedness proposition, the noncomparability of the slope coefficient from (4) and (5) means that unbiasedness cannot be tested using (5), owing to the fact that estimates of β —equivalent to those from (4)—do not exist under the alternative hypothesis. Estimates of β from (5) are equivalent to (4) under the null hypothesis, but not even comparable under the alternative. Hence, in contrast to the assertion that (5) is a better way of testing unbiasedness than (4), the fact is unbiasedness cannot be tested using (5). This is true regardless of whether the spot and forward rates are I(0) or I(1) processes. The suggestion that (5) provides a more appropriate test of the unbiasedness hypothesis than (4) is simply wrong.

⁴ To see how this problem arose in the short-run demand for money literature, see Thornton (1985).

The noncomparability of (4) and (5) explains why the estimates from these specifications can be so different. Indeed, because of the general noncomparability of estimates of β from (4) and (5), there is no particular reason to be distressed by the fact that estimates from these equations are different. The unbiasedness puzzle is resolved by noting that there are good reasons to expect unbiasedness to be violated, at least to some extent. For example, unbiasedness will not hold if investors are risk averse because then a non-zero risk premium drives a wedge between f_{t-1} and $E_{t-1}s_t$. Alternatively, unbiasedness may be violated because of Jensen's inequality, or what is known in the foreign exchange literature as Siegel's paradox. While there is disagreement about the magnitude of the distortion caused by Jensen's inequality (e.g., McCulloch, 1975; Siegel, 1972; Goodhart, McMahon, and Ngama, 1997; and Dumas, Jennergren, and Naslund, 1995), there is no doubt that the existence of this mathematical fact prevents unbiasedness from holding exactly. As we are about to see, even a tiny violation of the unbiasedness hypothesis has significant consequences for the estimated slope coefficient from (5).

4. The Qualitative Importance of Non-comparability

The above analysis shows that estimates of the slope coefficients from (4) and (5) are not comparable. This section investigates how different these estimates can be for even a tiny violation of (1). To this end, a simple Monte Carlo experiment is conducted. The hypothetical spot rate, s_t^h , is generated by a simple AR(1) process, i.e.,

$s_t^h = \mu + \rho s_{t-1}^h + \xi_t$, $\rho < 1$ and $\xi_t \sim N(0, \sigma_\xi^2)$. The hypothetical forward rate is given by

$f_{t-1}^h = \lambda[\mu + (\rho + \theta_t)s_{t-1}^h]$, where $\theta_t \sim N(0, \sigma_\theta^2)$. Note that (1) holds if $\lambda = 1$ and $\theta_t = 0$

for all t . If $\lambda = 1$ and $\theta_t \neq 0$, (1) does not hold at every point in time, but holds on average, i.e., $E f_{t-1}^h = E s_t^h$. This is the type of violation that one might expect to be associated with Jensen's inequality. Finally, if $\lambda \neq 1$, the forward rate is biased even unconditionally, e.g., a risk premium or a peso problem.

Two parameterizations of the model are considered. The first assumes that exchange rates are stationary; the second assumes they are nonstationary. The parameters are chosen to match monthly data for the Swiss spot rate against the U.S. dollar for the period December 1978 – December 1998.⁵ Specifically, $\mu = 0.48$, $\rho = 0.98$, and $\sigma_\xi = 0.03$. The initial condition is the log of the Swiss exchange rate for December 1978, i.e., $s_0^h = 0.482$. For each case considered, 10,000 samples with a sample size of 300 are generated for both s_t^h and f_{t-1}^h after discarding the first 1,000 observations to lessen the effect of the initial condition. The sample size is similar to many empirical studies that use monthly data.

The distributions of the estimates of the slope coefficients from (4) and (5) under the null hypothesis, i.e., for $\lambda = 1$ and $\theta_t = 0$ for all t are presented in panels (a) and (b), respectively, of Figure 1. The estimates from both equations are centered on the true value of 1.0. The distribution of estimates from (4) are very tight, ranging from 0.9997 to 1.0003.

The distribution of the estimates from (5) is much wider and strongly skewed to the right. The positive skew reflects the finite sample bias caused by the extreme

⁵ The choice of the Swiss spot rate is arbitrary. The results are qualitatively the same regardless of the spot rate chosen.

persistence in the spot rate. This bias was first documented by Bekaert, Hodrick, and Marshall (1997) and disappears—albeit very slowly—as the sample size increases.⁶

To investigate what happens to the estimates from (4) and (5) when (1) does not hold, market participants are assumed to be aware that the spot rate is generated by an AR(1) process but are uncertain about the magnitude of ρ . Specifically, the true coefficient is $\rho + \theta_t$, but market participants believe it is ρ . The departure from the null hypothesis is very small. Specifically, $\sigma_\theta = 0.001$. The distributions of the estimates of the slope coefficients from (4) and (5) are presented in Figure 2, in panels (a) and (b), respectively. Given the modest violation of the null hypothesis, it is not surprising that the distribution of estimates from (4) in panel (a) is nearly identical to the distribution in panel (a) of Figure 1. The violation of (1) has an enormous impact on the estimates from (5), however. The estimates are relatively small. The mean estimate is 0.0155 and negative estimates are common. None of the negative estimates is as large as those frequently obtained in the literature, however.

The second parameterization assumes that exchange rates are nonstationary, i.e., $\rho = 1$. All other parameters are unchanged. Estimates of (4) and (5) under the null are presented in panels (a) and (b) of Figure 3. Estimates of β from (4) are very tightly clustered around 1.0 because of the well-known fact that the estimates are super consistent in this circumstance. The estimates from (5) are also tightly centered about 1.0 and are symmetric.

Now consider estimates of β from (4) and (5) under the assumption that $\lambda = 1$ and $\sigma_\theta = 0.001$, which are presented in Figure 4. Estimates of β from (4), shown in

⁶ The distribution remained slightly skewed for a sample size of 2,500 observations.

panel (a), are essentially the same as those reported in Figure 3. As is the case with the assumption of stationarity, estimates of β from (5) are very different. All of the estimates are positive, but considerably less than 1.0.⁷

Additional experiments were conducted letting λ equal 0.95 and 1.05, respectively, and using values of σ_θ ranging from 0.1 to 0.00001. The results are summarized here to conserve space. As might be expected, estimates of β from (4) are rather tightly distributed about the expected value, which is slightly greater than 1.0 when $\lambda < 1$ and slightly less than 1.0 when $\lambda > 1$, and the variance of the distributions declines as σ_θ is decreased from 0.1. Estimates from (5) continue to be close to zero, but are consistently negative when $\lambda < 1$ and consistently positive when $\lambda > 1$. The shape of the distribution is relatively insensitive to the choice of σ_θ .

These experiments demonstrate that, because of the non-comparability of the estimates from (4) and (5) when the null hypothesis does not hold, even relatively modest violations of the null hypothesis can generate very different parameter estimates from the two specifications. It is important to remember that the results in Figures 2 and 4 were for specific data-generating processes for the spot and forward rates. The results obtained for historical data depend on the true, but unknown, data-generating processes. Indeed, using a simple model economy Chakraborty (2007) has shown that small violations of the

⁷ An interesting question is how close σ_v^2 has to be to zero before non-comparability is not a significant issue (assuming $\lambda = 1$). Experiments suggest that when $\sigma_v = 0.0001$ the distribution of estimates from (5) looks more like the distribution in Figure 1, panel b, except that distribution is shifted further to the left, resulting in a much higher probability of a negative estimate. When $\sigma_v = 0.00001$, there is essentially no difference between the distribution and that reported in Figure 1, panel b.

unbiasedness condition can result in relatively large negative estimates of β , more typical of those found in the literature.

5. The Test for UIP

The forward premium equation is also used to test for UIP. To better understand the implication of the analyses in Sections 3 and 4 for tests of UIP, it is useful to see how the forward premium test for UIP is derived. UIP requires

$$(11) \quad i_{t-1} - i_{t-1}^* = E_{t-1} s_t - s_{t-1} + \zeta_{t-1},$$

where i_t and i_t^* are the domestic and foreign interest rates denominated in the domestic and foreign currencies, respectively, on assets that are identical and have the same maturity as the term of the forward contract. The error term, ζ_{t-1} , potentially reflects a variety of factors including time aggregation and a risk premium.

UIP states that expected gain from holding one currency rather than another must be equal to the opportunity cost of holding that currency rather than the other, which is just the nominal interest rate differential. UIP is essentially an arbitrage condition. When (11) does not hold investors have an expected profit opportunity.

Of course, there is also a riskless arbitrage opportunity, CIP, which states that

$$(12) \quad i_{t-1} - i_{t-1}^* = f_{t-1} - s_{t-1} + v_{t-1},$$

where v_{t-1} is an i.i.d. zero mean, constant variance random variable that reflects the fact that CIP will not hold identically because of transactions costs, measurement error, etc.⁸

Substituting (12) into (11) yields

$$(13) \quad f_{t-1} - s_{t-1} = E_{t-1} s_t - s_{t-1} + \zeta_{t-1} - v_{t-1}.$$

⁸ It is more common practice in the literature to present (12) without an error. However, there is no evidence that CIP holds without error.

The forward premium test for UIP is obtained by assuming (2), i.e., rational expectations.

With (2), (13) can be rewritten as

$$(14) \quad \Delta s_t = f_{t-1} - s_{t-1} + \eta_{t-1},$$

where $\eta_{t-1} = \zeta_{t-1} - v_{t-1}$, which is parameterized to yield the forward premium test for UIP,

i.e.,

$$(15) \quad \Delta s_t = \alpha' + \beta'(f_{t-1} - s_{t-1}) + \eta_{t-1}.$$

Note that the forward premium test of UIP does not assume unbiasedness: (1) is not used to derive (15); (11), (12), and (2) are required, but (1) is not. Hence, despite the fact that (15) and (5) are observationally equivalent, the forward premium tests for unbiasedness and UIP are conceptually separate and distinct.⁹

6. The Forward Premium Puzzle

Considerable resources have been devoted to “resolving” the forward premium puzzle—the fact that estimates of β from (15) are typically negative. Some of this literature focuses on the failure of UIP itself by arguing that the forward premium includes a time-varying risk premium. However, the majority of this work suggests that a time-varying risk premium cannot account for the massive rejection of UIP found in the literature. Hence, the violations of UIP found in the literature are generally attributed to something other than a failure of UIP per se.

Even though (1) is not used to derive (15), unbiasedness and UIP are closely linked. To see just how closely linked they are, note that substituting (1) into (11) yields

$$(16) \quad i_{t-1} - i_{t-1}^* = f_{t-1} - s_{t-1} + \zeta_{t-1},$$

⁹ McCallum (1994, p. 108-09) also notes that there is a distinction between these hypotheses.

which is observationally equivalent to (12). Consequently, if one assumes that unbiasedness holds, it would be extremely difficult to distinguish between CIP and UIP.¹⁰ In light of the overwhelming evidence that CIP holds, a rejection of UIP is tantamount to a rejection of unbiasedness. Moreover, we have already seen, small violations of the unbiasedness condition can result in negative estimates of β —a result common in the UIP literature.

Despite Chakraborty’s (2007) demonstration that small violations of the unbiasedness condition can result in estimates of β that are somewhat characteristic of those found in the literature, this section explores an alternative resolution of the forward premium puzzle. Specifically, it argues that the forward premium puzzle is due to the interaction of the near-random-walk behavior of exchange rates, the fact that CIP holds, and economic fundamentals that drive the behavior of exchange rates.

6.1 Random Walk Behavior and the Failure of Rational Expectations

It is useful to note that the “rational expectations” assumption used to derive (15) is very strong. One might think of (2) might be thought of as stochastic perfect foresight, with foresight being perfect if $v_t = 0$ for all t . The effect of (2) is to replace the unobservable variable, $E_{t-1}s_t$, with the observed variable, s_t . The conflict between the near-random-walk behavior of exchange rates and (2) can best be demonstrated by assuming that the spot rate is a pure random walk, i.e.,

$$(17) \quad s_t = s_{t-1} + v_t.$$

With this assumption, $E_{t-1}s_t = s_{t-1}$ and not $s_t + v_t$, as in (2). If in addition to the random walk assumption, it is assumed that unbiasedness holds, then $f_{t-1} = E_{t-1}s_t = s_{t-1}$ for all t

¹⁰ A distinction could only be made on the basis of the characteristics of the error term.

and $s_t - f_{t-1} = \Delta s_t$. The relationships among Δs_t , $f_{t-1} - s_{t-1}$, and $s_t - f_{t-1}$ are very different from the relationship if the rational expectations assumption (2) used to derive (15) holds. In this case, the $E_{t-1}s_t = s_t - v_t$. Under unbiasedness, $f_{t-1} = s_t + v_t$, so that $s_t - f_{t-1} = v_t$. Conceptually, if the rational expectations assumption (2) held, the white noise errors made in forecasting the spot rate, v_t , would be uncorrelated with shocks to the spot rate, v_t . However, if the spot rate was essentially unpredictable, as in the case of a near random walk stochastic process, $v_t \cong v_t$, and the correlation would be close to 1.0. Note too that $f_{t-1} - s_{t-1} = \Delta s_t + v_t$, and, hence, highly correlated with Δs_t .

Table 1 reports descriptive statistics of $f_{t-1} - s_{t-1}$, $s_t - f_{t-1}$, and Δs_t for nine currencies—the Canadian dollar, *CD*, UK pound, £, Swiss franc, *SF*, Japanese yen, ¥, Belgian franc, *BF*, Italian lira, *IL*, French franc, *FF*, Dutch guilder, *DG*, and German mark, *DM*. The data are end-of-month foreign currency/U.S. dollar spot and forward exchange rates. For the first four currencies listed, the data are for the period December 1978 – January 2002. For the last five currencies, the data are for the period December 1978 – December 1998.¹¹ Consistent with the evidence that exchange rates are near-random-walk processes, the descriptive statistics for $s_t - f_{t-1}$ and Δs_t are very similar. Indeed, the correlation between $s_t - f_{t-1}$ and Δs_t , presented in the second to last line of Table 1, is greater than 0.99 for every currency and the ratio of the variance of Δs_t to $s_t - f_{t-1}$ is very close to 1 for all of the currencies.

For all nine currencies, the maximum and minimum values of the forward premium are much smaller in absolute value than either $s_t - f_{t-1}$ or Δs_t . Moreover, the

¹¹ These data were used in Baillie and Kilic (2006) and were kindly provided by Rehim Kilic.

ratio of the variance of Δs_t to the forward premium is very large for all currencies, ranging from 30 for the CD to 225 for the ¥, with most being larger than 80. Moreover, the time-series properties of $f_{t-1} - s_{t-1}$ differ significantly from those of either $s_t - f_{t-1}$ or Δs_t . Hence, the data are very similar to what is expected if the spot rate is a random walk process and very different from what is expected if (2) held and the forward rate was an unbiased predictor of the spot rate.

The near random walk behavior of exchange rates is important for resolving the UIP puzzle. Specifically, if the spot rate were a pure random walk, estimates of β from (15) would be zero. Both f_{t-1} and s_{t-1} would reflect all of the information available at $t-1$ and Δs_t would reflect the new information, orthogonal to the information in either f_{t-1} or s_{t-1} . Estimates of β are typically not zero, however, and are sometimes statistically significant. This is shown in Table 2, which presents the estimates of (15) for each of the nine exchange rates over their respective sample periods. The table reports the OLS estimates α and β , the corresponding marginal significance level (s.l.), and the adjusted R-square (\bar{R}^2). The estimates are negative for all but two exchanges rates—*FF* and *IL*. However, estimates of \bar{R}^2 are very small, indicating a weak relationship between Δs_t and $(f_{t-1} - s_{t-1})$. Given the relative variances of the right- and left-hand sides of (15), reported in Table 1, it is hardly surprising that the estimated \bar{R}^2 are typically small or that estimates of β tend to be imprecisely estimated—the estimate of β is not significantly different from zero at conventional critical values for four of the nine currencies.

The weakness of the relationship between Δs_t and the forward premium is dramatically illustrated in Figure 5, which plots Δs_t (on the vertical axis) and the forward premium (on the horizontal axis) using the same scale. The 45-degree line indicates where all of the observations would lie if UIP held without error. Virtually no observations lie on or about the line.

Moreover, there is virtually no time variation in these figures. Hence, recent evidence that UIP fares better when asymmetries and nonlinearities between the spot and forward rate are accounted for, when the forward premium is at a premium rather than a discount, or during the first half of the 1980s (e.g., Wu and Zhang, 1996; Bansal, 1997; Bansal and Dahlquist, 2000; Zhou, 2002; Sarno, Valente, and Leon, 2006; and Baillie and Kilic, 2006) is likely to be fragile.

The fragility of the estimates of β is investigated further by estimating (15) using rolling regression with a 30-month window. These rolling regression estimates of β for each of the nine currencies are presented in Figures 6 and 7. The estimates are plotted on the last month of the corresponding sample. Estimates for the ¥ and the European currencies other than the pound are presented in Figure 6 because they behave similarly. Estimates for the *CD* and the £ are presented in Figure 7. Consistent with the results in Table 1, the estimates tend to be negative on average; however, they vary considerably over time, switching from negative to positive, sometimes rather quickly. Moreover, the estimates in Figure 6 behave similarly over time despite marked differences in the behavior of the forward premiums for some of these exchange rates.

What accounts for these results presented in Figures 6 and 7? The answer can be better understood by noting that the estimate of β is given by

$$(18) \quad \hat{\beta} = \frac{\sum_{t=1}^T (\bar{f}_{t-1} - \bar{s}_{t-1})(\bar{s}_t - \bar{s}_{t-1}) / T}{\sum_{t=1}^T (\bar{f}_{t-1} - \bar{s}_{t-1})^2 / T},$$

where the bar notes that the variables have been mean adjusted. Both the numerator and denominator of (18) tend to be small, but the denominator, which is the variance of the forward premium, is very small. Moreover, for all nine currencies, the denominator is much more stable over time than the numerator, which varies considerably over time, switching from negative to positive. Because the denominator is very small, relatively large changes in the numerator sometimes lead to very large positive or negative estimates of β .

The resolution of the forward premium puzzle comes from observing that while Δs_t and $f_{t-1} - s_{t-1}$ are nearly uncorrelated with each other, both are correlated with $i_{t-1} - i_{t-1}^*$. This is shown in Table 3, which presents the correlation between $i_{t-1} - i_{t-1}^*$ and $f_{t-1} - s_{t-1}$ and Δs_t for each of the nine currencies. One-month rates were not available for most of these countries over this time period. Hence, the foreign rates are 3-month rates or longer obtained from the IMF.¹² The U.S. rate is the secondary market rate on 3-month T-bills obtained from the Board of Governors of the Federal Reserve. The positive correlation between $i_{t-1} - i_{t-1}^*$ and $f_{t-1} - s_{t-1}$ for all currencies is a consequence of CIP. Indeed, if the maturity of interest rates matched that of the forward contracts and the observations were synchronous, the correlations would be very close to 1.0. The

¹² For Canada, France, Germany, Belgium, the U.K., and Japan, the rates are 3-month rates, on Treasury bills (Canada, France, the U.K., and Belgium), Interbank deposits (Germany), and time deposits (Japan). For Italy, the rate is the 3- to 12-month newly issued T-bills. For the Netherlands, it is the banker's call loan rate. For Switzerland, it is rate on federal debt register claims.

correlation between Δs_t and $i_{t-1} - i_{t-1}^*$ is much lower and is negative for every currency except the Italian lira, which is slightly positive.

Estimates of β are determined by the relationship of $i_{t-1} - i_{t-1}^*$ with both $f_{t-1} - s_{t-1}$ and Δs_t , rather than the relationship between $f_{t-1} - s_{t-1}$ and Δs_t , per se. Because of CIP, the correlation between $i_{t-1} - i_{t-1}^*$ and $f_{t-1} - s_{t-1}$ is always positive; therefore, the sign of β depends on the sign of the correlation of $i_{t-1} - i_{t-1}^*$ with Δs_t , being negative when the correlation is negative and positive when the correlation is positive. This is illustrated in Figure 8, which plots the 30-month rolling correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t , and the corresponding 30-month rolling regression estimates of β for the Italian lira. The estimates of β range from large negative values to large positive values, depending on whether the correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t is negative or positive. The correlation between these two time series is 0.788, suggesting that much of the time variation in the estimates of β is due to time variation in the correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t . Indeed, the signs of the correlation and estimates of β are the same 94 percent of time (200 of 212 estimates).

Table 4 presents the correlation between the 30-month rolling correlation of $i_{t-1} - i_{t-1}^*$ and Δs_t , and the corresponding estimates of β and percent of time when the correlation and the estimate of β had the same sign. Table 4 also reports the corresponding significance levels for these statistics based on simulations under the assumption that $i_{t-1} - i_{t-1}^*$, Δs_t , and $f_{t-1} - s_{t-1}$ are independent. The simulations use a 30-month window and sample sizes identical to those that yield the results in Tables 3 and 4

(see the appendix for details). The correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t accounts for much of the time variation in estimates of β for eight of the nine currencies. The exception is *CD*, where the correlation is only 0.212, which is not significantly different from what could be obtained if there were no relationship among these variables. Even for the *CD*, however, the correlations and the estimates of β were the same sign for 76 percent of the estimates and much larger than one would expect if $i_{t-1} - i_{t-1}^*$ and Δs_t were unrelated. For all of the other countries, the correlations are greater than 60 percent and the percent the same sign is 80 percent or higher. Both are statistically significant at very low significances levels.

The above analysis suggests that the forward premium puzzle is due to near-random-walk behavior of exchange rates and the fact that Δs_t and $f_{t-1} - s_{t-1}$ are both correlated with $i_{t-1} - i_{t-1}^*$. The strong positive correlation between $f_{t-1} - s_{t-1}$ and $i_{t-1} - i_{t-1}^*$ is a consequence of CIP. For most time periods and exchange rates the correlation between Δs_t and $i_{t-1} - i_{t-1}^*$ is negative; however, during some periods it is positive. A plausible explanation for the time variation in the relationship between Δs_t and $i_{t-1} - i_{t-1}^*$ is that the correlation is positive when the behavior of $i_{t-1} - i_{t-1}^*$ is due to changes in the relative real rates and negative when the behavior of $i_{t-1} - i_{t-1}^*$ is due to the relative behavior of domestic/foreign expected inflation rates. The latter dominates the former over longer sample periods for most countries. Consequently, over longer sample periods and for most countries, the estimate of β is negative. This explanation is consistent with evidence from the late 1970s and early 1980s that the response of interest rates to surprise

money announcements was due to changes in inflation rather than changes in real rates (e.g., Cornell, 1982, 1983; Engel and Frankel, 1984; Frenkel, 1981; and Thornton, 1989).

Furthermore, given the relative magnitudes of the variances of Δs_t and the forward premium, estimates of β from (15) can be very large—positive or negative. Because of the sensitivity of OLS to extreme observations, some of these estimates will be “statistically significant” in spite of the fact that the forward premium has little or no explanatory power for the change in the spot rate.

6.2 Robustness Check

The above evidence suggests that much of the failure of UIP is due to the near-random-walk behavior of exchange rate (which is at odds with the rational expectations assumption used to derive the standard test for UIP), CIP and economic fundamentals. If this explanation resolves the forward premium puzzle, it should be invariant to the sample period or to the particular pairs of exchange rates considered. To investigate the robustness of the result to exchange rates and sample periods, the same analysis was undertaken using monthly data for the *UKP* spot and 3-month forward exchange rates with the U.S. dollar (\$), the *FF*, the *DG*, and the *IL* for the period January 1921 through December 1936. These data are from Einzig (1937), a commonly used data source for testing UIP over this period. Data for the U.S. and U.K. 3-month Treasury bill rates for the same period was obtained from Global Financial Data.

Consistent with CIP, the correlation between $i_{t-3} - i_{t-3}^*$ and $f_{t-3} - s_{t-3}$ is positive for the four currencies, although somewhat lower and more variable than those for the U.S. dollar over the 1979-2000 sample period. Transactions and information costs are likely higher during this period, so that deviations from CIP are larger and more

persistent.¹³ Nevertheless, the results, summarized in Table 5, are similar to those obtained for the 1979-2000 period. The top panel presents the full sample estimates of (15). The estimates are typical of what is found in the literature. For three of the four currencies, the estimate of β is negative. Moreover, estimates are significantly different from zero for the \$ and the *DG*. The \bar{R}^2 is very small for the three European currencies, but is relatively large for the \$.

The bottom panel summarizes the results over a 30-month rolling sample. Estimates of β are unstable, ranging from large positive to large negative values. Moreover, as was the case for the U.S. dollar exchange rates over the 1979-2000 sample period, the correlation between $i_{t-3} - i_{t-3}^*$ and $s_t - s_{t-3}$ accounts for much of the time variation in the estimates of β . The correlation between these time series range from a low of about 0.50 for the *IL* to nearly 0.84 for the \$. The percent of the samples where the estimates of β and the correlation were the same sign is very high, ranging from 75.5 to 88.7 percent. Moreover, with one exception these estimates are statistically significant at the 5 percent level. The exception—the correlation for *IL*—is statistically significant at the 10 percent level.

As a further check the analysis is applied to the U.S. dollar spot and 1-month forward exchange rates for 17 countries for the period January 2000 through December 2006. The exchange rates and the corresponding 1-month U.S. and foreign interest rates were obtained from Bloomberg. These data are same-day exchange and interest rates for the last available observation in the corresponding month. The results are summarized in Table 6. While these data are not synchronous, the full sample correlations between

¹³ See Akram, Rime, and Sarno (2006) for data on the persistence of deviations from CIP in modern times.

$i_{t-1} - i_{t-1}^*$ and $f_{t-1} - s_{t-1}$ are above 90 percent for all but two currencies and are above 97 percent for twelve countries. Moreover, the maximum and minimum 30-month rolling correlations suggest relatively little time variation for most currencies with many estimates close to 1.0. These high correlations suggest that CIP is very close to holding for these data over much of the sample period. The exception is the Taiwanese dollar where the correlations range from a high of 0.740 to a low of -0.743.

In contrast, full-sample correlations between $i_{t-1} - i_{t-1}^*$ and Δs_t are negative for every country except Poland and maximum and minimum values of the 30-month rolling correlations indicate considerable time variation. For all countries, the minimum correlation is negative and typically relatively large, while for all but two countries—South Africa and Mexico—the maximum correlation is positive.

Table 7 presents full sample estimates of β for each currency along with the results from 30-month rolling correlations and regressions. The full sample estimates of β are negative except for the Taiwanese dollar and the Czech koruna. The latter is consistent with resolution of the forward premium puzzle presented here in that the Czech Republic is the only country whose full sample correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t is positive. Uncharacteristic of other sample periods, however, none of the estimates of β is statistically significant at the 5 percent significance level and only two are significant at the 10 percent level. Nevertheless, as before, estimates of β are unstable, ranging from large negative to large positive values. Moreover, the correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t accounts for much of the time variation in the estimates of β , and all of the correlations except the Taiwanese dollar are statistically significant. Furthermore, the

percent of estimates where the sign of β and the sign of the correlation between $i_{t-1} - i_{t-1}^*$ and Δs_t are the same is greater than 85 percent for 14 of the 17 currencies and are statistically significant for every currency except the Taiwanese dollar.

7. Conclusions

There are two related puzzles in the empirical foreign exchange literature. The first is the finding that there is a marked difference in the conclusion about the forward rate unbiasedness hypothesis depending on whether the hypothesis is tested using the forward rate or forward premium equation. The second—the *forward premium puzzle*—is the fact that the forward premium incorrectly predicts the direction of the subsequent change in the spot rate, which implies a massive rejection of UIP. This paper resolves both puzzles.

The first puzzle is resolved by showing that estimates of the slope parameter from the forward rate and forward premium equations are not comparable under the alternative hypothesis. Since there are several reasons why unbiasedness will not hold exactly, there is no reason to be concerned that the estimates from these equations are different—indeed, they are not comparable. Simple Monte Carlo experiments show that the estimates from these equations will be very different even for tiny violations of the null hypothesis.

The noncomparability of these parameters also means that, contrary to the suggestion in the literature that it is better to test unbiasedness using the forward premium rather than the forward rate specification, it is impossible to test unbiasedness using the forward premium specification. This conclusion stems from the fact that it is impossible

to obtain an estimate of the parameter under the alternative hypothesis that is comparable to the parameter under the null.

The resolution of the forward premium puzzle stems from the observation that the stochastic perfect foresight assumption used to derive the forward premium test of UIP is greatly at odds with the near random walk behavior of exchange rates. The near-random-walk behavior of exchange rates implies a weak relationship between changes in the spot rate and the forward premium, which is consistent with estimates of the forward rate equation found in the literature. CIP, however, implies that there is a strong positive correlation between the forward premium and the difference between domestic and foreign interest rates. Hence, the sign of the estimate of the slope coefficient from the forward premium equation depends on the sign of the correlation between the change in the spot rate and the difference between domestic and foreign interest rates, being positive when this correlation is positive and negative when the correlation is negative. This resolution is supported by the fact that time variation in the correlation between the change in the spot rate and the difference between domestic and foreign interest rates explains much of the time variation in the estimates of the slope coefficient. This is true for a variety of exchange rates and over several different sample periods. The sign of this correlation is determined by economic fundamentals. The correlation is positive when the behavior of the domestic/foreign interest rate differential is due to relative changes in real rates and negative when it is due to changes in expectations for inflation. The forward premium puzzle—the preponderance of negative estimates of the slope coefficient—is because, for most exchange rates and for most sample periods, the

difference between domestic and foreign interest rates reflects concerns about inflation rather than the behavior of real rates.

Of course, the above explanation need not account for all of the anomalous results obtained using (15). Other factors may also play a role. As noted in Sections 3 and 4, minor violations of the unbiasedness condition can generate negative estimates of the slope coefficient as well. Moreover, because exchange rates are not pure-random-walk processes, there may be some marginal predictability in the spot exchange rate beyond its current value (e.g., Clarida, Sarno, Taylor, and Valente, 2003). Nevertheless, the results presented here suggest that the near-random-walk behavior of the spot rate, CIP, and economic fundamentals that generate correlation between the domestic/foreign interest differential and the change in the spot rates are important—if not the most important—reasons for the empirical failure of UIP and, consequently, critical for resolving the forward premium puzzle.

References:

- Akram, Q.F., D. Rime, and L. Sarno, 2006, "Arbitrage in the Foreign Exchange Market: Turning on the Microscope," Norges Bank Working Paper No. 2005/12.
- Bacchetta, P. and E. van Wincoop., 2004, "A Scapegoat Model of Exchange Rate Fluctuations," *American Economic Review*, 94, 114-18.
- Bacchetta, P. and E. van Wincoop. 2006, "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *American Economic Review*, 96, 552-76.
- Baillie, R.T., and R. Kilic, 2006, "Do Asymmetric and Nonlinear Adjustments Explain the Forward Premium Anomaly?," *Journal of International Money and Finance*, 25, 22-47.
- Bakshi, G., and A. Naka, 1997, "Unbiasedness Of Forward Exchange Rates," *Financial Review*, 32, 145-162.
- Bansal, R., 1997, "An Exploration of the Forward Premium Puzzle in Currency Markets," *Review of Financial Studies*, 10, 369-403.
- Bansal, R., and M. Dahlquist. 2000, "The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies," *Journal of International Economics*, 51, 115-44.
- Barnhart, S. W., R. McNown, and M.S. Wallace, 1999, "Non-Informative Tests of the Unbiased Forward Exchange Rate," *Journal of Financial and Quantitative Analysis*, 34, 265-91.
- Bekaert, G., R.J. Hodrick, and D.A. Marshall, 1997, "On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Financial Economics*, 44, 309-48.
- Chakraborty, A. 2007. "Learning, the Forward Premium Puzzle and Market Efficiency," Economics Department, University of Tennessee, working paper.
- Cheung, Y-W., M.D. Chinn. 2001, "Currency Traders and Exchange Rate Dynamics: A Survey of the U.S. Market," *Journal of International Money and Finance*, 20, 439-71.
- Cheung, Y-W., M.D. Chinn, and A.G. Pascual, 2005, "Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?" *Journal of International Money and Finance*, 24, 1150-75.
- Chow, G.C., 1964, "A Comparison of Alternative Estimators for Simultaneous Equations," *Econometrica*, 32, 532-53.

- Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente, 2003, "The Out-of-sample Success of Term Structure Models as Exchange Rate Predictors: A Step Beyond." *Journal of International Economics* 60, 61–83.
- Cornell, B. 1982, "Money Supply Announcements, Interest Rates, and Foreign Exchange," *Journal of International Money and Finance*, 1, 201-08.
- Cornell, B. 1983, "The Money Supply Announcement Puzzle: Review and Interpretation," *American Economic Review*, 73, 644-57.
- Dumas, B., L.P. Jennergren, and B. Naslund, 1995, "Siegel's Paradox and the Pricing of Currency Options," *Journal of International Money and Finance*, 14, 213-23.
- Einzig, P. *The Theory of Forward Exchange*, MacMillian and Co., London, England, 1937.
- Engel, C.M. and J.A. Frankel, 1984, "Why Interest Rates React to Money Announcements: An Explanation from the Foreign Exchange Market," *Journal of Monetary Economics*, 13, 31-9.
- Engel, C.M. and K.D. West, 2005, "Exchange Rates and Fundamentals," *Journal of Political Economy*, 113, 485-517.
- Frenkel, J.A., 1976, "A Monetary Approach to the Exchange Rate: Doctrinal Aspects and Empirical Evidence," *Scandinavian Journal of Economics*, 78, 200-24.
- Frenkel, J.A., 1981, "Flexible Exchange Rates, Prices, and the Role of News: Lessons from the 1970s," *Journal of Political Economy*, 89, 665-705.
- Froot, K.A. and J.A. Frankel, 1989, "Forward Discount Bias: Is It an Exchange Rate Risk Premium?" *The Quarterly Journal of Economics*, 104, 139-61.
- Goodhart, C.A.E., P.C. McMahon, and Y.L. Ngama, 1997, "Why Does the Spot-Forward Discount Fail to Predict Changes in the Future Spot Rates?" *International Journal of Finance & Economics*, 2, 121-29.
- Groen, J.J.J., 2005, "Exchange Rate Predictability and Monetary Fundamentals in a Small Multi-Country Panel," *Journal of Money, Credit, and Banking*, 37, 495-516.
- Hodrick, R., *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, New York: Harwood Academic Publishers, New York, 1987.
- Kilian, L. and M.P. Taylor, 2003, "Why Is It So Difficult to Beat the Random Walk Forecast of Exchange Rates?" *Journal of International Economics*, 60, 85-107.

- Levich, R.M., 1978, "Further Results on the Efficiency of Markets for Foreign Exchange," in *Managed Exchange Rate Flexibility: The Recent Experience*, Federal Reserve Bank of Boston Conference Series no. 20. Boston: Federal Reserve Bank.
- Maynard, A., 2003, "Testing for Forward Rate Unbiasedness: On Regression in Levels and in Returns," *Review of Economics and Statistics*, 85, 313-27.
- Maynard, A., and P.C.B. Phillips, 2001, "Rethinking an Old Empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly," *Journal of Applied Econometrics*, 16, 671-708.
- McCallum, B.T., 1994, "A Reconsideration Of The Uncovered Interest Parity Relationship," *Journal of Monetary Economics*, 33, 105-32.
- McCulloch, J.H., 1975, "Operational Aspects of the Siegel Paradox," *Quarterly Journal of Economics*, 89, 170-72.
- Meese, R.A., and K.J. Singleton, 1982, "On Unit Roots and the Empirical Modeling of Exchange Rates," *Journal of Finance*, 37, 1029-35.
- Meese, R.A., and K. Rogoff. 1983, "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics*, 14, 3-24.
- Rossi, B. 2006, "Are Exchange Rates Really Random Walks? Some Evidence Robust to Parameter Instability," *Macroeconomic Dynamics*, 10, 20-38.
- Siegel, J., 1972, "Risk, Interest Rates and the Forward Exchange," *Quarterly Journal of Economics*, 86, 303-09.
- Sarno, L., G. Valente, and H. Leon, 2006, "Nonlinearity in Deviations from Uncovered Interest Parity: An Explanation of the Forward Bias Puzzle," *Review of Finance*, 10, 443-482.
- Sarno, L., G. Valente, 2006, "Exchange Rates and Fundamentals: Footloose or Evolving Relationship? Unpublished manuscript.
- Sarno, L. and M. Taylor, 2002, *The economics of exchange rates*, Cambridge University Press, Cambridge, U.K.
- Thornton, D.L., 1985, "Money Demand Dynamics: Some New Evidence," Federal Reserve Bank of St. Louis *Review*, 67, 14-23.
- Thornton, D.L., 1989, "The Effect of Unanticipated Money on the Money and Foreign Exchange Markets," *Journal of International Money and Finance*, 8, 573-87.

Wu, Y., and H. Zhang, 1996, "Asymmetry in Forward Exchange Rate Bias: A Puzzling Result," *Economics Letters*, 50, 407-11.

Zhou, S., 2002, "The Forward Premium Anomaly and the Trend Behavior of the Exchange Rates," *Economics Letters*, 76, 273-79.

Appendix

This appendix outlines the procedure for testing whether the correlations between the rolling correlations between the change in the spot rate and the interest rate differential and percentages of estimates where the signs of the estimated correlations and the estimates of β are the same are significantly different from what would have resulted had the series been independent. Specifically, I estimate an AR(1) model of Δs_t , $i_t - i_t^*$, and $f_t - s_t$ for each pair of countries and for each time period. I then generate 10,000 samples of the same size those used in the paper for each pair of currencies. The samples are obtained by bootstrapping the residuals from the estimated models. The residuals are sampled independently to guarantee independence. For each variable the initial conditions were the last observation for that variable in the corresponding historical sample and 1,000 observations were generated before the samples were taken to minimize the effect of the initial conditions.

An analysis identical to that in the text was then undertaken. For each of the 10,000 samples (15) was estimated using a 30-month rolling regression and the estimates of β retained.¹⁴ The 30-month rolling correlations of Δs_t and $i_{t-1} - i_{t-1}^*$ for identical samples were also estimated. The correlation between the 30-month rolling estimates of β and correlation of Δs_t and $i_{t-1} - i_{t-1}^*$ were calculated for each of the 10,000 replications. These data were used to obtain distributions of both the correlations under the null hypothesis and the percent of the estimates where the correlation and the estimate of β have the same sign.

¹⁴ There were a couple of estimates for 17 exchange rates over the January 2000 – December 2006 period where the sample estimates of β were slightly larger than 1.0. In these cases, β was set to be 0.99.

