European Financial Management, Vol. 6, No. 2, 2000, 121–148

An explanation of the forward premium 'puzzle'

Richard Roll* and Shu Yan
The Anderson School at UCLA, Los Angeles, CA 90095-1481, USA,
e-mail: rroll@anderson.ucla.edu; syan@anderson.ucla.edu

Abstract

Existing literature reports a puzzle about the forward rate premium over the spot foreign exchange rate. The premium is often negatively correlated with subsequent changes in the spot rate. This defies economic intuition and possibly violates market efficiency. Rational explanations include non-stationary risk premia and econometric mis-specifications, but some embrace the puzzle as a guide to profitable trading.

We suggest there is really no puzzle. A simple model fits the data: forward exchange rates are unbiased predictors of subsequent spot rates. The puzzle arises because the forward rate, the spot rate, and the forward premium follow nearly non-stationary time series processes. We document these properties with an extended sample and show why they give the delusion of a puzzle.

Keywords: *foreign exchange; anomalies; non-stationary time series.*

JEL classification: *F*31, *G*15.

1. The 'puzzle' in the forward exchange premium

The spread between the forward foreign exchange rate and the concurrent spot exchange rate has exhibited some puzzling behavior. During particular sample periods, it is negatively correlated with subsequent changes in the spot exchange rate.[1] This has attracted a lot of attention, for it defies even the most basic common sense.

Interest rate parity offers a striking panorama on the puzzle. The forward exchange premium on a given date must equal the difference between default-free nominal interest rates in the two countries,

$$R_{D,t} - R_{F,t} = f_t - s_t = \text{forward premium},$$

*Keynote speaker at the European Financial Management Association Conference, Paris, France, June 1999.

[1] As near as we can tell, the empirical phenomenon was first documented by Bilson (1981) and Longworth (1981), who cite each others' working papers. Similar results are in Cumby and Obstfeld (1984), Fama (1984), Hodrick and Srivastava (1984), and Huang (1984). These authors use data beginning in the mid-1970s, some extending into the early 1980s. An early review was provided by Boothe and Longworth (1986).

where f_t and s_t are the natural logarithms of, respectively, the forward and spot exchange rates observed on date t, (domestic currency units per foreign unit), and R denotes the continuously-compounded nominal zero-coupon interest rate over the same term as the forward contract, with additional subscripts D and F denoting domestic and foreign. Within the bounds of transaction costs, interest rate parity is a pure arbitrage condition which holds at every t for freely convertible currencies.

But if the forward premium is *negatively* related to subsequent change in the spot rate, the domestic currency will appreciate (depreciate) when its nominal interest rate is higher (lower.) This suggests a simple and profitable trading scheme: Buy the bonds of the country whose nominal interest rate is higher. That trading rule produces a greater local currency return for sure; and it will likely be supplemented by an exchange rate appreciation. To the extent that a trader can borrow in the low interest rate country and use the proceeds to buy bonds in the high interest rate country, the forward premium puzzle represents an alluring opportunity.

The negative correlation of the puzzle cannot, however, be very close to perfect because the trading scheme above would then work most of the time. Indeed, the correlation found in past research is small in magnitude, though allegedly negative and significant. The smallish correlation opens the door for other explanations such as a time varying risk premium. If the forward rate contains a risk premium proportional to the nominal interest differential, trading profits could just be compensation for risk.

The trouble is a lack of theory about why there should be a risk premium in the first place. Various models have been concocted, none very convincing in our view. True, a risk averse international exchange trader might eschew large positions because of the inherent volatility. But volatility besets either trading direction; it alone cannot explain why the higher interest rate is persistently associated with larger trading profits.

We began this project with our own half-baked explanation based on decomposing the nominal interest rate into its real interest and expected inflation components. We reasoned that a real interest differential should induce a capital flow that might cause a subsequent currency appreciation in the higher (real) interest country. The expected inflation component, however, should have the opposite effect, a higher expected inflation being associated with a subsequent currency depreciation. Consequently, the 'puzzle' might arise in periods when the real interest rate differential exceeds the expected inflation differential. Our imagined process depended also on real interest incentives eliciting sluggish and perhaps inefficient capital flows. It did *not* envision a risk premium explanation of the difference in real interest rates, a possibility investigated by Korajczyk (1985).

As we examined the data, however, we began to realize that such an elaborate explanation might be incorrect and not even necessary. Repeating the empirical tests of past researchers on an updated data set, we discovered that the 'puzzle' was an ephemeral phenomenon. It was not present during some sub-periods, though of course this alone is not an adequate reason for dismissing its existence entirely. Nonetheless, we asked whether there might be a simpler theory consistent with the data *in all periods*. The surprising answer is yes. The theory is the oldest and perhaps the simplest in international finance: viz., the forward exchange rate is an unbiased predictor of the subsequent spot rate.

Section 2 describes our data sample. In section 3, we present a battery of diagnostic tests for stationarity, documenting the possibility of a unit root in both forward and spot exchange rates and, most important for understanding the puzzle, in the forward premium. We also replicate and confirm past empirical tests and update them with an extended sample. Given the empirical finding of possible non-stationarity in the forward premium, section 4 discusses whether such a result is theoretically even plausible. Section 5 presents some standard econometric corrections for non-stationarity in regression models, documents their shortcomings, and then presents an entirely new non-regression method for estimating the relation between forward and spot exchange rates. In section 6, we reveal the final result of our empirical sleuthing: why it is possible and even likely to uncover the appearance of a puzzle. Section 7 summarizes and concludes.

2. Data

Our foreign exchange data consist of monthly prices of the US dollar in eight major currencies: the Canadian dollar, French franc, German mark, Italian lira, Japanese yen, Dutch guilder, Swiss franc, and British pound. For July 1974 through December 1994, the data were from Harris Bank; earlier versions of this data set were examined in such well-known papers as Fama (1984) and Hodrick (1987). We also obtained from Datastream a second sample, originally generated by the National Westminster Bank and covering the period March 1985 through July 1997.

All calculations were performed with the log forward and spot exchange rates, $f_t = \log_e(F_t)$ and $s_t = \log_e(S_t)$, respectively, where F_t is the *1-month* forward rate of exchange, (foreign currency per US\$), and S_t is the spot exchange rate, both observed on trading day t. For the National Westminster Bank data set, t is close to the 15th of each month. If the 15th is not a trading day, the previously-available observation is used. The Harris Bank data are Friday closing observations spaced at 4-week intervals. There is a slight misalignment because the forward rates are typically 30 days to maturity, not 28.

By using monthly data only, we finesse the tricky problem of over-lapping observations that engaged the ingenuity of many past researchers, cf. Hansen and Hodrick (1980). There is, of course, a cost in the sense that our results may not be generally applicable to longer-term forward rates and their implicit forecast horizons.

3. Non-stationarity in exchange rate time series

3.1. *Naïve regressions*

A simple exploratory model of the forward rate's forecasting ability would be a regression in the levels,

$$s_{t+1} = a + bf_t + \varepsilon_{t+1}, \tag{1}$$

where a and b are coefficients to be estimated and ε_t is a disturbance. In fact, as shown in the first panel of Table 1, this model makes a good superficial impression; the intercept, a, is close to zero and the slope, b, is near 1.0 when estimated with the longest sample, the Harris Bank data over approximately 20 years. F tests fail to reject the null hypothesis, $\mathbf{H}_0$: $\{a = 0,\ b = 1\}$.

A closer examination, however, reveals that something is seriously amiss. Within sub-periods,[2] results are obtained such as those reported in the middle panels of

[2] The Harris data were divided to align the first date of the second sub-period with NatWest's beginning date.

Table 1

Forward exchange rate forecasts of subsequent spot exchange rates regression in levels: $s_{t+1} = a + bf_t + \varepsilon_{t+1}$

s_t and f_t are the logs of, respectively, the spot and forward exchange rates on date t in currency units per US$. a, b, and ε_t are, respectively intercept, slope, and disturbance estimates. Standard errors are given in parenthesis below estimated coefficients. The F-statistic is for a Wald test of $\mathbf{H_0}$: $\{a=0, b=1\}$; its p-value is in parenthesis. The Newey-West heteroskedasticity correction is employed. Monthly observations from two data sources, Harris Bank and NatWest Bank.

	a	*b*	*F*	*a*	*b*	*F*	*a*	*b*	*F*	*a*	*b*	*F*
	July 74–December 94			July 74 February 85			March 85–December 94			March 85–July 97		
	Harris									National Westminster		
Canada	0.002	0.99	0.67	0.002	1.00	1.43	−0.001	1.00	0.93	−0.001	1.00	1.17
	(0.002)	(0.009)	(0.51)	(0.002)	(0.01)	(0.24)	(0.004)	(0.02)	(0.40)	(0.004)	(0.01)	(0.31)
France	0.021	0.99	0.86	−0.022	1.01	2.56	0.129	0.92	21.79	0.131	0.92	16.58
	(0.020)	(0.01)	(0.42)	(0.017)	(0.01)	(0.08)	(0.032)	(0.02)	(0.00)	(0.027)	(0.02)	(0.00)
Germany	0.004	0.99	0.14	0.008	1.00	1.55	0.029	0.94	18.26	0.031	0.94	13.94
	(0.009)	(0.01)	(0.87)	(0.017)	(0.02)	(0.22)	(0.010)	(0.01)	(0.00)	(0.08)	(0.01)	(0.00)
Italy	0.054	0.99	0.83	−0.075	1.01	1.52	0.428	0.94	8.11	0.323	0.95	4.32
	(0.054)	(0.01)	(0.44)	(0.058)	(0.01)	(0.22)	(0.159)	(0.02)	(0.00)	(0.159)	(0.02)	(0.02)
Japan	−0.000	1.00	0.42	0.279	0.95	3.04	0.179	0.96	7.52	0.216	0.95	6.49
	(0.029)	(0.01)	(0.66)	(0.116)	(0.02)	(0.05)	(0.061)	(0.01)	(0.00)	(0.072)	(0.01)	(0.00)
Netherlands	0.005	0.99	0.14	−0.005	1.01	1.52	0.037	0.94	18.49	0.039	0.94	15.88
	(0.011)	(0.01)	(0.87)	(0.018)	(0.02)	(0.22)	(0.012)	(0.01)	(0.00)	(0.09)	(0.01)	(0.00)
Switzerland	0.009	0.99	0.87	0.033	0.96	2.59	0.026	0.93	18.18	0.025	0.93	14.26
	(0.007)	(0.01)	(0.42)	(0.015)	(0.02)	(0.08)	(0.009)	(0.02)	(0.00)	(0.007)	(0.01)	(0.00)
UK	−0.016	0.97	1.37	0.015	1.02	1.95	−0.062	0.88	10.44	−0.060	0.89	6.71
	(0.010)	(0.02)	(0.26)	(0.010)	(0.02)	(0.15)	(0.016)	(0.04)	(0.00)	(0.019)	(0.04)	(0.00)

Table 1. Notice that the slope coefficients for certain countries (e.g. Japan and Switzerland) within sub-periods are *both* further from unity than the coefficient estimated with the entire data sample. Moreover, the F tests frequently reject $\mathbf{H}_0$: $\{a=0, b=1\}$. How can this happen?

We think it reveals long-term swings in the levels of both spot and forward rates. Making a *reductio ad absurdum* illustrative example, imagine two distinct sub-periods with widely-divergent levels, say, as in the case of the French franc, a level of 5 francs per dollar in one period and 10 francs in the other period. *Within* each sub-period, imagine also that $b \leqslant 0$. Putting all the data into one naïve regression, however, fits a line with a slope considerably closer to unity, simply by dint of the large difference in means between sub-periods. This is the essence of the 'spurious' regression problem that besets the analysis of *non*-stationary time series.

Past empirical literature abounds with failures to reject non-stationarity in exchange rate levels, both spot and forward. For example, see Meese and Singleton (1982), Doukas and Rahman (1987), Baillie and Bollerslev (1989), and Corbae, Lim and Ouliaris (1992), all of whom used data terminating in the mid-1980s at the latest.

Table 2 adds support to the suspicion of non-stationarity with formal stationarity tests of exchange rate levels through 1997. Both the augmented Dickey/Fuller test (1979, 1981) and the Phillips/Perron (1987, 1988) test are reported for the Harris Bank data and the National Westminster data. The critical values are approximately −4.0 and −3.5 at the 1% and 5% levels, respectively, (MacKinnon, 1991); test statistics less than these values would reject the hypothesis of a unit root and accept stationarity

For the entire sample period of the Harris Bank data, a unit root hypothesis cannot be rejected for either the forward or the spot exchange rate of any country. This suggests that they could be non-stationary and, consequently, that an OLS fit of model (1) might produce spurious results. With non-stationary dependent and explanatory variables, the estimated OLS coefficients of (1) have no asymptotic distribution and cannot be relied upon to reveal the true state of nature, even in the largest samples.

The Harris Bank data yield smaller (in absolute value) Dickey/Fuller and Phillips/ Perron test statistics in the first half of the sample, (July 1974 through February 1985) than in the second half, (March 1985 through December 1994). Indeed, in the second half, the unit root hypothesis is close to the 5% rejection level for many of the series. These results are augmented by the National Westminster data, which cover the second half of the Harris Bank data plus an additional 31 months, (through July 1997.) In every country but Japan, for both spot and forward rates, the Phillips/ Perron test is larger in absolute magnitude. Unlike the Dickey/Fuller test, it exceeds the 5% critical value in six instances, and the 1% critical value for the UK. The additional 31 months had a mixed impact on the unit root test statistics; for most series, the Dickey/Fuller test decreased in significance while the Phillips/Perron test increased.

It should be recognized that divergent results across sub-periods and methods are emblematic of time series that are close to non-stationary. So, perhaps the variations within Table 2 are to be expected and represent nothing more than sampling error in the test statistics. Of course, one can never be sure that a time series really does have a unit root, but caution suggests that these series are close enough that they should not be used in simple regressions without a healthy dose of skepticism about the estimated regression coefficients.

Table 2

Unit root tests of non-stationarity in exchange rate levels.

Augmented Dickey/Fuller and Phillips/Perron statistics with intercept and time trend.

	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94
	Dickey/Fuller			Phillips/Perron		
Spot rates, Harris Data						
Canada	−1.544	−2.075	−0.145	−1.595	−2.347	0.256
France	−1.380	−0.763	−3.351	−1.340	−0.698	−2.931
Germany	−1.801	−0.265	−3.349	−1.845	−0.230	−2.873
Italy	−1.675	−0.902	−2.402	−1.693	−0.926	−2.023
Japan	−2.350	−1.645	−3.200	−2.323	−1.650	−2.664
Netherlands	−1.764	−0.335	−3.302	−1.787	−0.298	−2.840
Switzerland	−2.301	−0.369	−3.209	−2.446	−0.448	−2.993
UK	−2.258	−0.129	−2.311	−2.292	−0.178	−2.609
Forward rates, Harris data						
Canada	−1.564	−2.096	0.108	−1.624	−2.385	0.217
France	−1.393	−0.823	−3.379	−1.347	−0.744	−2.937
Germany	−1.814	−0.262	−3.365	−1.855	−0.221	−2.887
Italy	−1.675	−0.937	−2.416	−1.703	−0.960	−2.033
Japan	−2.366	−1.659	−3.177	−2.338	−1.634	−2.645
Netherlands	−1.774	−0.303	−3.317	−1.794	−0.265	−2.850
Switzerland	−2.313	−0.358	−3.212	−2.454	−0.439	−2.992
UK	−2.268	−0.132	−2.329	−2.306	−0.187	−2.628
MacKinnon critical values: a smaller (more negative) test statistic rejects non-stationarity						
1%	−4.000	−4.035	−4.041	−4.000	−4.033	−4.039
5%	−3.430	−3.447	−3.450	−3.429	−3.446	−3.448

NatWest data, March 1985–July 1997

	Spot Rates		Forward Rates	
	Dickey/Fuller	Phillips/Perron	Dickey/Fuller	Phillips/Perron
Canada	−1.171	−1.331	−1.181	−1.352
France	−2.902	−3.483	−2.942	−3.510
Germany	−2.832	−3.248	−2.870	−3.281
Italy	−2.747	−3.104	−2.759	−3.119
Japan	−2.817	−2.676	−2.816	−2.674
Netherlands	−2.796	−3.238	−2.829	−3.266
Switzerland	−3.066	−3.553	−3.079	−3.562
UK	−2.742	−4.200	−2.758	−4.232
MacKinnon critical values: a smaller (more negative) test statistic rejects non-stationarity				
1%	−4.024	−4.022	−4.024	−4.022
5%	−3.442	−3.441	−3.442	−3.441

3.2. The puzzle in detail

The spectre of non-stationary forward and spot exchange levels was no doubt part of the reason previous researchers adopted alternatives to regression (1). Perhaps the most studied model subtracted the lagged spot rate from both sides of (1) to give

$$s_{t+1} - s_t = a + b\ (f_t - s_t) + \varepsilon_{t+1}, \tag{2}$$

The dependent variable is the first difference in the (log) spot exchange rate, which is more likely to be stationary than the level, though a formal test of stationarity was not conducted by most past researchers. The explanatory variable is the forward premium. Regression (2) brought out the 'puzzle' because the estimated coefficient b was negative; Bilson (1981), Cumby and Obstfeld (1984) and Fama (1984). These studies and others are summarized in the survey monograph by Hodrick (1987, chs 3–4).

Table 3 presents the fit of regression model (2) to our extended data set. In the earlier sub-period, July 1974 through February 1985, which coincides roughly with the periods used in the above-cited studies, the estimated slope coefficient b in (2) is indeed negative for every country. The Wald F test of $\mathbf{H_0}$: $\{a=0,\ b=1\}$ rejects at a high level of significance in all cases, but the slope coefficient is twice its standard error for only four countries (though it is close for others.)

In the second sub-period, March 1985 through December 1994 for the Harris data or through July 1997 for the NatWest data, nine of the 16 coefficients are negative, but only two, Canada for the Harris data and Japan for the NatWest data, can be regarded as less than zero at a standard significance level. The Wald test also produces more ambiguous results; the p values are less than 0.10 for only six of the 16 regressions.

Goodhart, McMahon and Ngama (1992) report similar instability across sub-periods in coefficients estimated from (2). Using samples of both weekly and monthly data for different calendar periods (some extending through 1987), they conclude that negative estimates of b in (2) '... are due to the presence of structural breaks and/or outliers in the data. After allowing for these all the beta estimates become insignificantly different from zero' (p. 138).

Our regressions for the entire Harris sample period (July 1974–December 1994) have uniformly negative slope coefficients, (Table 3), but they are significantly below zero in only four of eight countries. The Wald test rejects $\mathbf{H_0}$: $\{a=0,\ b=1\}$ in six of eight countries.

In all of the regressions with model (2), explanatory power is very low, with only one R-square exceeding 5%, (Netherlands in the first sub-period) and most are even smaller.

3.2. Another unit root problem

Is (2) a well-specified OLS model? Tests of the first difference in log spot exchange rates, the dependent variable in (2), soundly reject the unit root hypothesis for all currencies and time periods.[3] First differences in exchange rates thus *appear* to be stationary.

[3] To save space, the results are not reported here, but they will be furnished on request to interested readers.

Table 3

Forward premium forecasts of subsequent changes in spot exchange rates regression with monthly observations: $s_{t+1} - s_t = a + b(f_t - s_t) + \varepsilon_{t+1}$.

s_t and f_t are the logs of, respectively, the spot and forward exchange rates on date t in currency units per US\$. a, b, and ε_t are, respectively intercept, slope, and disturbance estimates. Standard errors are given in parenthesis below estimated coefficients. The F-statistic is for a Wald test of $\mathbf{H}_0$: $\{a=0, b=1\}$; its p-value is in parenthesis. The R-square is unadjusted. The Newey-West heteroskedasticity correction is employed.

	a	b	F	R^2	a	b	F	R^2	a	b	F	R^2	a	b	F	R^2
	July 74–Dec 94				July 74–Feb 85				Mar 85–Dec 94				Mar 85–Dec 97			
	Harris												NatWest data			
Canada	0.004	−1.56	16.30	0.034	0.004	−1.14	11.97	0.023	0.004	−2.04	8.44	0.033	0.001	−0.64	6.44	0.006
	(0.001)	(0.46)	(0.00)		(0.003)	(0.66)	(0.00)		(0.003)	(0.98)	(0.00)		(0.001)	(0.52)	(0.00)	
France	0.002	−0.70	2.25	0.000	0.009	−1.09	3.74	0.019	−0.006	0.27	2.85	0.003	−0.003	−0.55	2.07	0.001
	(0.003)	(0.84)	(0.11)		(0.003)	(0.94)	(0.03)		(0.004)	(1.33)	(0.06)		(0.004)	(1.46)	(0.13)	
Germany	−0.003	−0.74	2.18	0.004	−0.006	−2.50	4.10	0.024	−0.007	1.52	1.95	0.015	−0.004	0.86	0.76	0.004
	(0.003)	(0.84)	(0.11)		(0.005)	(1.43)	(0.02)		(0.003)	(1.26)	(0.15)		(0.003)	(1.18)	(0.47)	
Italy	0.004	−0.10	3.40	0.000	0.017	−1.12	9.95	0.045	−0.011	2.12	2.44	0.021	−0.002	0.09	2.01	0.000
	(0.004)	(0.46)	(0.03)		(0.005)	(0.48)	(0.00)		(0.008)	(2.27)	(0.09)		(0.007)	(1.89)	(0.14)	
Japan	−0.008	−1.68	9.83	0.022	−0.004	−1.39	7.82	0.023	−0.010	−1.68	3.65	0.007	−0.012	3.41	4.50	0.031
	(0.002)	(0.66)	(0.00)		(0.003)	(0.62)	(0.00)		(0.004)	(2.00)	(0.03)		(0.004)	(1.64)	(0.01)	
Netherlands	−0.003	−1.54	5.71	0.019	−0.003	−2.63	10.26	0.064	−0.007	1.41	2.13	0.009	−0.001	−0.15	0.83	0.000
	(0.002)	(0.76)	(0.00)		(0.003)	(0.83)	(0.00)		(0.004)	(1.51)	(0.12)		(0.003)	(1.11)	(0.44)	
Switzerland	−0.007	−1.28	4.34	0.013	−0.013	−2.49	5.72	0.032	−0.006	0.47	1.28	0.001	−0.0051	−0.42	0.97	0.001
	(0.004)	(0.77)	(0.01)		(0.007)	(1.12)	(0.00)		(0.004)	(1.70)	(0.28)		(0.0038)	(1.57)	(0.38)	
UK	0.006	−1.84	5.78	0.021	0.009	−1.48	6.11	0.026	0.001	−1.22	1.47	0.002	0.001	−0.91	0.82	0.002
	(0.003)	(0.84)	(0.00)		(0.003)	(0.75)	(0.00)		(0.006)	(2.44)	(0.23)		(0.004)	(2.33)	(0.44)	

However, as reported in Table 4, the forward premium, (the explanatory variable in (2)), *is* close to being non-stationary. For the Harris Bank overall data period, the Phillips/Perron test rejects the unit root hypothesis at the 1% level for five of the eight countries while Dickey/Fuller rejects for only France and Italy. Surprisingly, the test statistics are closer to significant in the first sub-period, in contrast to the spot and forward rate levels (Table 2). During the second sub-period, no test statistic is even close to rejecting the unit root hypothesis. After the data are purged of probable errors,[4] the extended sample of the NatWest data finds neither test statistic significant over the 1985–97 period. The unit root statistics are uniformly small in absolute magnitude, intimating non-stationarity.

Baillie and Bollerslev (1989) state that the forward premium is stationary for their data, which cover 1980–85. We could not, however, find any test in their paper to support such a conclusion. The cointegration test in their section 2 seems to have been conducted only with the forward rate forecast error, $s_{t+1}-f_t$, not with the forward premium, f_t-s_t. As we shall see, there are ample theoretical reasons why the forecast error *should* be stationary and persuasive empirical results that it is. Neither the theory nor the empirical support applies to the forward premium.

Admittedly, a truly non-stationary forward premium would be something of a paradox. If (2) is a well-specified model, with disturbances unrelated to the right-side variable *and* with $b \neq 0$, then *both* variables must share a unit root property; both must be either stationary or non-stationary. Yet the statistical tests suggest that $s_{t-1}-s_t$ is always stationary and f_t-s_t is sometimes not. Which test should we believe?

One possibility is that f_t-s_t truly has a unit root. In this case, $s_{t+1}-s_t$ also has a unit root but the noise contributed by ε_t is so overwhelming that only an extremely large sample would detect the unit root's presence. The forward premium, f_t-s_t, however, is substantially less noisy, so its unit root *is* revealed in our modest-sized samples.

Table 5 presents a small-scale investigation of this possibility with a simulation of a unit root variable masked by varying degrees of noise. Since the simulated variable y is composed of a random walk plus additive noise, it is truly non-stationary. As an illustrative example, we applied the Phillips/Perron test to y and find, as suspected, a dramatic effect of noise on the test's efficacity. The standard deviation of the test statistic increases steadily with the level of noise relative to the volatility of x, the non-stationary component of x. Shockingly, for noise levels 4 and 5 times greater than the variability (standard deviation) of y, even the mean of the simulated unit root statistic is more negative than critical rejection values. At the highest level of noise in Table 5, many more than half of the simulated statistics *incorrectly* reject the presence of a unit root at the 1% level.

We interpret this simulation as providing some credence to the idea that *both* variables in (2), f_t-s_t and $s_{t-}-s_t$ really might have unit roots but the test statistics are powerless to detect its presence in $s_{t+1}-s_t$ because of the noise and modest sample size.

An alternative explanation of the results in Table 4 is that neither variable really has a unit root but the forward premium is *nearly* non-stationary. Unit root tests have notoriously weak power in such circumstances; i.e. they fail to reject the false hypothesis of a unit root unless the sample size is very large.

[4] These results are given in the columns headed 'corrected data'.

Table 4
Unit root tests of forward premium stationarity.

	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94	Dickey/Fuller	Phillips/Perron	Dickey/Fuller	Phillips/Perron
	Dickey/Fuller			Phillips/Perron			Original data		Corrected data*	
	Harris						NatWest, Mar 85–July 97			
Canada	−3.405	−3.093	−1.673	−4.148	−3.655	−2.000	−2.068	−4.079	−0.827	−1.303
France	−4.242	−3.436	−1.192	−7.594	−6.249	−1.911	−2.026	−9.042	−1.138	−1.808
Germany	−2.195	−3.185	−0.708	−2.129	−3.510	−0.572	−0.708	−2.945	−0.850	−0.793
Italy	−4.418	−3.525	−1.954	−6.910	−5.480	−2.867	−2.277	−3.155	−2.154	−2.546
Japan	−3.856	−3.176	−0.697	−4.096	−4.108	−0.811	−2.108	−8.541	−1.386	−1.276
Netherlands	−3.085	−4.131	−0.922	−4.116	−5.011	−0.760	−1.471	−5.746	−0.925	−0.911
Switzerland	−2.754	−2.111	−0.805	−2.511	−2.075	−0.813	−0.893	−1.498	−1.365	−1.276
UK	−3.252	−3.321	−0.698	−3.602	−3.662	−0.755	−1.647	−4.390	−1.613	−2.044
	MacKinnon critical values									
1%	−4.00	−4.04	−4.04	−4.00	−4.03	−4.04	−4.02	−4.02	−4.05	−4.03
5%	−3.43	−3.45	−3.45	−3.43	−3.45	−3.45	−3.44	−3.44	−3.45	−3.44

* Data were purged of errors identified by applying interest rate parity; i.e. by comparing the forward premium with the nominal interest differential. If the deviation was greater than 0.001, that date was deleted. Since the number of data errors differed across countries, the MacKinnon critical values are slightly different also, but they are very close to the average numbers reported in the table.

Table 5

Simulation of Phillips/Perron statistic for a truly non-stationary variable with a substantial component of stationary noise.

x follows the non-stationary process, $x_t = x_{t-1} + \nu_t$, and y contains non-stationary and stationary components: $y_t = x_t + \varepsilon_t$, where ν and ε are normally and independently-distributed with zero means and standard deviations σ_ν and σ_ε respectively. There are 10,000 replications of a simulated time series sample of size $T = 500$.

The rejection levels of the Phillips/Perron test are: −2.57 (1%), −1.94 (5%), and −1.62 (10%). Note: $\sigma_\nu = 1$ is fixed, and only σ_ε varies.

σ_ε	Mean	Median	Standard Deviation	Maximum	Minimum
			Phillips/Perron Statistic		
1	5.41	5.33	1.67	12.44	−7.19
2	4.50	5.16	2.48	8.23	−20.37
3	0.35	2.22	4.39	5.83	−25.00
4	−3.11	−1.44	5.16	3.51	−25.31
5	−5.71	−4.42	5.39	1.80	−24.22

A third possibility is that $b = 0$ in (2), implying that forward exchange rates contain no useful information whatsoever about future spot exchange rates. Goodhart *et al.* (1992) point out that if the spot exchange rate actually follows a random walk, (and is thus non-stationary), then regression (2) *should* produce $b = 0$. This would resolve the empirical paradox, allowing a stationary $s_{t+1} - s_t$ to co-exist with a non-stationary $f_t - s_t$. It would also explain why the estimated slope coefficient in (2) is negative in some periods. If the true value of b is zero, its estimate should be negative in about half of all samples. Moreover, if the regression's independent variable, $f_t - s_t$, is actually non-stationary, inferences are problematic; the estimated b could be negative in sizable samples.

One feature shared by all three possibilities is that the forward premium is non-stationary or nearly so. The economics of this condition are worthy of a short digression.

4. The economics of non-stationarity in the forward premium

When first seeing the empirical results in Table 4 indicating that the forward premium is possibly non-stationary, we wondered whether a unit root process was even plausible for this particular variable. It seemed more likely that inter-market arbitrage, international capital flows, and perhaps government action would keep the forward premium from wandering aimlessly. Given interest rate parity, this appeared all the more probable because non-stationarity in the forward premium implies non-stationarity in the default-free nominal interest rate difference between the two countries. Could $R_{D,t} - R_{F,t}$ actually have a unit root? Could it conceivably drift unfettered; or do international economic forces drive it toward some long-term level?

Some insight about this issue can be garnered from the Fisher 'Open' Relation, i.e. by expressing the nominal interest differential as the sum of a real rate differential plus the difference in expected inflation rates,

$$R_{D,t} - R_{F,t} = r_{D,t} - r_{F,t} + E_t(I_{D,t+1} - I_{F,t+1}),$$

where r is the real rate and I is the actual inflation rate, both with appropriate country and time subscripts.[5]

A truly risk-free real interest differential should *not* have a unit root. A difference in genuinely riskless real rates should elicit capital flows from the low to the high (real) interest country and ultimately eradicate any economically meaningful gap. Consequently, if the forward premium has a unit root, it must be caused by (1) a unit root in the expected inflation differential or (2) a non-stationary time-varying risk premium differential in the real rate, or a combination of both.

Unlike the real interest differential, there are no private international forces acting to equilibrate inflationary expectations. True, if the monetary policies of the two countries are linked, even casually, one would anticipate some mean reverting tendency in inflation and presumably in inflationary expectations. But countries can and history has proven that they often do pursue widely divergent monetary policies, bringing broad departures of inflation from any perceptible long-term level. Hyperinflations are examples of the explosions attendant upon non-stationary processes.

But even hyperinflations eventually come to an end. The difference in inflationary expectations does not wander in perpetuity so it cannot really be non-stationary over an infinite horizon. Over periods of months and even years, however, it has often demonstrated the propensity to mimic a unit root process. If a model such as (2) is fit with data from those periods, the results could be misleading.

Turning now to the conjecture that a risk premium is the source of non-stationarity, we first note that the interest rate parity arbitrage condition equates the forward exchange premium to the *default-free* nominal interest difference. Any risk premium embedded in the underlying real rates must be attributable to inflation risk because the *nominal* payoffs are certain. This implies that higher moments of the inflation process, including but not limited to volatility or covariation with other risk factors, are the underlying driving forces behind any non-stationarity in the forward premium.

Consequently, if the forward premium has a unit root, then so too must some moment of the inflation process. It could be the first moment or some higher moment, or several moments, of either country's inflation.

Table 6 presents unit root tests for the nominal interest differential and, for the NatWest data only, the actual inflation differential (not inflationary expectations).[6] In the second sub-period, the unit root hypothesis cannot be rejected for the nominal rate differential for any country. However, for some countries the unit root statistic exceeds (in absolute amount) the 5% critical value during the first sub-period and in the overall period since 1974. There are conspicuous differences between the Dickey/Fuller and the Phillips/Perron tests, but we believe the latter is susceptible to contamination by data error outliers. Also, a few cases of rejection are to be expected

---

[5] 'D' denotes domestic and 'F' denotes foreign.

[6] The Harris data do not include actual inflation.

Table 6

Unit root tests of nominal interest differentials and inflation.

	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94	July 74–Dec 94	July 74–Feb 86	Mar 85–Dec 94	Dickey/Fuller	Phillips/Perron	Dickey/Fuller	Phillips/Perron
	Dickey/Fuller			Phillips/Perron			Interest differential		Monthly inflation	
	Nominal interest rate differential, Harris, monthly data						NatWest, March 85–July 97, Monthly			
Canada	−2.975	−2.578	−1.522	−3.664	−3.230	−1.776	−1.220	−1.206	−4.515	−13.295
France	−3.808	−3.156	−1.242	−7.108*	−6.011	−1.700	−1.012	−1.295	−6.502	−10.199
Germany	−2.127	−3.183	−0.553	−1.996	−3.554	−0.495	−0.978	−0.752	−4.635	−9.358
Italy†	−3.861	−3.004	−1.774	−7.881*	−5.969	−2.195	−2.049	−2.491	−4.349	−9.292
Japan†	−3.421	−2.682	−0.727	−3.579	−2.832	−0.771	−1.263	−1.314	−5.956	−19.858
Netherlands	−2.981	−4.163	−0.707	−3.849	−4.824	−0.732	−1.160	−0.968	−6.224	N/A
Switzerland	−2.661	−2.050	−0.871	−2.521	−2.176	−0.887	−0.906	−0.950	−4.162	−11.219
UK	−3.129	−3.224	−0.704	−3.701	−3.833	−0.797	−1.568	−1.686	−4.469	−9.567

* These numbers are large because of some apparently bad data in the first Harris sub-period. In general, all the interest differential series and forward premium series have more volatility in the first sub-period, and therefore have smaller (more negative) test statistics.

† There are missing observations in both Italian and Japanese data. The MacKinnon critical values for these two countries are slightly different from but close to those of the other countries.

just by chance. Given interest rate parity, these results reflect similar tests using forward premium data.[7]

The actual inflation difference appears to be stationary since all test statistics exceed their critical values, most by a wide margin. There are two possibilities to explain these results: (1) near non-stationarity in the nominal rate difference is caused by a near non-stationary risk premium in the underlying real rate difference, and (2) inflationary expectations are near non-stationary but this is masked by a large amount of *unexpected* inflation (which is itself stationary or nearly so), cf. the supporting simulation results in Table 5.

A time-varying risk premium was posited by Fama (1984), who argued that when exchange markets are efficient, a negative coefficient in regression (2) implies a risk premium negatively correlated with the expected change in spot rates. Provided that (2) is well-specified, Fama's argument is impeccable. If, however, the explanatory variable in (2) is non-stationary or nearly so, as indeed seems likely, the estimated coefficient can range widely and be negative for extended subperiods, even without the negative correlation suggested by Fama, a correlation in abject want of theoretical support.[8]

A risk premium could be the source of the puzzle simply because it is non-stationary; it need not be negatively correlated with expected changes in spot rates. As we just pointed out, however, the source could also be non-stationarity in inflationary expectations. One thing is certain: regression (2) is suspiciously close to being ill-specified.

5. Further tests of forward rate predictability in the presence of non-stationarity

5.1. Standard methods

When dealing with possibly non-stationary time series, the standard econometric cure is to work with first differences. Applying this procedure to both sides of (2), the resulting regression equation is

$$s_{t+1} - 2s_t + s_{t-1} = a + b(f_t - f_{t-1} - s_t + s_{t-1}) + \varepsilon_{t+1}. \tag{3}$$

A fit of (3) for the Harris data is reported in the first panel of Table 7.[9] While the unbiased forward rate hypothesis $\mathbf{H}_0$:$\{a=0, b=1\}$ cannot be rejected by an F test for most series and time periods, there is an obvious lack of power. With so much noise, virtually any conceivable hypothesis would fail to be rejected. The simple first-differencing method does not shed much light on the 'puzzle'.

A more powerful version of the standard test can be developed by conditionally assuming the forward rate is an unbiased predictor, i.e. that

$$E_t(s_{t+1}) = f_t$$

[7] Cf. Table 4.

[8] Based on our argument that a risk premium in the forward exchange premium must be related to higher moments of inflation, Fama's negative correlation implies a link between, say, a high inflation volatility in one country and an expected appreciation of that country's currency. This seems counter-intuitive, to say the least.

[9] To conserve space, we omit similar results for the NatWest data.

and markets are rational in the sense that

$$s_{t+1} = E_t(s_{t+1}) + \varepsilon_{t+1},$$

where the prediction error, ε_{t+1}, is completely random and stationary. Combining these assumptions gives a model analogous to (2):

$$s_{t+1} - s_t = E_t(s_{t+1}) - s_t + \varepsilon_{t+1},$$

First differencing to mitigate any unit root problems,

$$s_{t+1} - 2s_t + s_{t-1} = [E_t(s_{t+1}) - E_{t-1}(s_t)] - [s_t - s_{t-1}] + \varepsilon_{t+1} - \varepsilon_t,$$

or, equivalently, (again using the unbiasedness of the forward rate),

$$s_{t+1} - 2s_t + s_{t-1} = f_t - 2s_t + s_{t-1} + \varepsilon_{t+1}. \tag{4}$$

The second panel of Table 7 reports a fit of (4) to the Harris data. The unbiased forward rate forecast hypothesis, $\mathbf{H_0}$:$\{a=0,\ b=1\}$, appears to be fairly well supported. The F test rejects this hypothesis for no country in any period at a 1% level. It is close to rejection at the 5% level for Japan during the entire period and Canada during the first sub-period. The estimated slope, b, is more than two standard errors from unity in a few other instances, but this could just be happenstance; a few spurious 'significant' coefficients are to be expected when many are calculated.

Again, however, there is a problem; it is intimated by the surprisingly high level of explanatory power. The same highly volatile component, $s_t - s_{t-1}$, appears on both sides of (4). Even if $s_{t+1} - s_t$ and $f_t - s_t$ were unrelated, or even negatively related, adding $s_t - s_{t-1}$ to both sides of the regression virtually guarantees a significant positive estimate of b.

First differencing in either form (3) or (4) is thus problematic. Both regressions unfairly favor the null hypothesis that the forward rate is an unbiased predictor, (3) because of weak power and (4) because of a spurious positive correlation between the dependent and explanatory variables.

5.2. Other literature

Dissatisfaction with OLS regressions in any of the forms (1) through (4) has motivated the development of sophisticated correctives. Baillie and Bollerslev (1989) summarize the issue well:

> ...it is generally not appropriate simply to first difference the data in order to achieve stationarity. This fairly standard procedure imposes too many unit roots. Estimating a model in levels involving only nonstationary variables, however, imposes too few unit roots, and standard asymptotically based inference procedures will not apply (p. 176.)

Their solution is to fit a cointegrating vector, i.e. to test whether OLS estimates of the coefficients a and b in (1) produce stationary residuals. They do. In addition, Baillie and Bollerslev impose $a=0$, $b=1$ in (1) and report that the resulting residuals also are stationary.

A similar finding is reported by Hakkio and Rush (1989) who utilize the 'error-correction' equivalent characterization of cointegration. Although they find that spot and forward rates are cointegrated (for the pound sterling and deutschemark, respectively) they also are unable to reject the joint hypothesis that (1) there is no risk

Table 7

First differencing the forward premium and the subsequent change in spot rates.

s_t and f_t are the logs of, respectively, the spot and forward exchange rates on date t in currency units per US$. a, b, and ε_t are, respectively intercept, slope, and disturbance estimates. Standard errors are given in parenthesis below estimated coefficients. The F-statistic is for a Wald test of $\mathbf{H_0}$:$\{a=0, b=1\}$; its p-value is in parenthesis. The R^2 is unadjusted. The Newey-West heteroskedasticity correction is employed. Monthly Harris data.

Standard first difference regression: $s_{t+1} - 2s_t + s_{t-1} = a + b(f_t - f_{t-1} - s_t + s_{t-1}) + \varepsilon_{t+1}$

	a	b	F	R^2	a	b	F	R^2	a	b	F	R^2
	July 74–December 94				July 74–February 85				March 85–December 94			
Canada	0.000	−2.48	2.87	0.014	0.000	−0.79	0.77	0.002	−0.000	−7.88	9.14	0.066
	(0.000)	(1.47)	(0.06)		(0.001)	(1.52)	(0.46)		(0.001)	(2.08)	(0.00)	
France	0.000	0.51	0.05	0.001	0.001	0.89	0.07	0.005	−0.001	−3.82	0.98	0.007
	(0.001)	(1.59)	(0.95)		(0.002)	(1.65)	(0.93)		(0.002)	(3.49)	(0.38)	
Germany	−0.000	−2.91	0.67	0.003	0.000	−3.24	0.76	0.008	−0.001	1.56	0.04	0.000
	(0.001)	(3.37)	(0.51)		(0.001)	(3.53)	(0.47)		(0.002)	(12.09)	(0.96)	
Italy	−0.000	−1.66	4.38	0.016	0.001	−2.01	5.44	0.053	−0.001	1.76	0.06	0.003
	(0.001)	(0.90)	(0.01)		(0.002)	(0.94)	(0.01)		(0.002)	(3.42)	(0.94)	
Japan	0.000	0.09	0.09	0.000	0.000	−0.57	0.26	0.001	0.000	10.87	0.91	0.013
	(0.001)	(2.13)	(0.91)		(0.002)	(2.24)	(0.77)		(0.002)	(7.32)	(0.41)	
Netherlands	0.000	−0.48	0.58	0.000	0.001	−0.33	0.52	0.000	−0.001	−6.28	0.18	0.002
	(0.001)	(1.38)	(0.56)		(0.002)	(1.36)	(0.60)		(0.002)	(12.46)	(0.84)	
Switzerland	−0.000	−5.47	2.70	0.010	0.001	−3.61	1.57	0.008	−0.001	−19.36	3.18	0.029
	(0.001)	(2.79)	(0.07)		(0.002)	(2.67)	(0.21)		(0.002)	(8.21)	(0.05)	
UK	−0.000	−2.88	0.98	0.005	0.000	−1.64	0.46	0.004	−0.001	−18.06	1.83	0.026
	(0.001)	(2.77)	(0.38)		(0.001)	(2.88)	(0.63)		(0.002)	(10.05)	(0.17)	

Regression equation based on unbiased forward rate: $s_{t+1} - 2s_t + s_{t-1} = a + b(f_t - 2s_t + s_{t-1}) + \varepsilon_{t+1}$

Canada	0.000	0.89	2.39	0.443	0.002	0.91	3.08	0.441	−0.002	0.91	2.31	0.463
	(0.001)	(0.05)	(0.09)		(0.001)	(0.07)	(0.05)		(0.001)	(0.08)	(0.10)	
France	−0.002	0.97	0.46	0.489	0.003	0.98	0.65	0.490	−0.008	1.01	2.91	0.515
	(0.002)	(0.07)	(0.63)		(0.003)	(0.11)	(0.52)		(0.004)	(0.11)	(0.06)	
Germany	−0.000	0.99	0.02	0.498	0.005	0.99	1.59	0.487	−0.007	1.05	1.80	0.532
	(0.002)	(0.08)	(0.98)		(0.003)	(0.11)	(0.21)		(0.004)	(0.11)	(0.17)	
Italy	−0.002	0.93	1.18	0.465	0.002	0.91	1.00	0.433	−0.006	0.97	2.09	0.499
	(0.002)	(0.06)	(0.31)		(0.003)	(0.07)	(0.37)		(0.004)	(0.10)	(0.13)	
Japan	−0.002	0.89	2.97	0.448	0.002	0.85	2.76	0.423	−0.007	0.98	2.59	0.492
	(0.002)	(0.05)	(0.05)		(0.003)	(0.007)	(0.07)		(0.003)	(0.07)	(0.08)	
Netherlands	−0.001	0.96	0.19	0.483	0.005	0.92	1.54	0.453	−0.007	1.05	2.03	0.535
	(0.002)	(0.08)	(0.83)		(0.003)	(0.12)	(0.22)		(0.004)	(0.11)	(0.14)	
Switzerland	−0.000	0.93	0.42	0.467	0.004	0.90	1.69	0.442	−0.006	1.00	1.44	0.509
	(0.003)	(0.07)	(0.66)		(0.004)	(0.13)	(0.19)		(0.004)	(0.11)	(0.24)	
United Kingdom	−0.000	0.89	1.49	0.445	0.004	0.83	2.57	0.416	−0.006	0.96	2.43	0.484
	(0.002)	(0.07)	(0.23)		(0.003)	(0.10)	(0.08)		(0.004)	(0.11)	(0.09)	

premium *and* (2) the exchange market is informationally efficient. This result is, however, based on an F test whose reliability could be questioned given the non-Gaussian nature of exchange rate data.

Recent papers by Phillips, McFarland and McMahon (1996) and by Phillips and McFarland (1997) emphasize these distributional peculiarities and apply tests suitable for non-Gaussian variates. The first paper examines unique data from the 1920s (French and Belgian francs, Italian lira, and US dollar against the British pound) with a least absolute deviation regression method. The evidence favors cointegration between the forward rate and subsequent spot rate. Moreover, the authors report the interesting result that the slope coefficient b in (1) is much closer to unity when estimated with robust methods than with OLS. However, $a=0$, $b=1$ is rejected for the two francs and the lira (though accepted for the dollar). The second paper uses very similar methods applied to different data, the Australian/US dollar, 1984–91. Some conclusions are different, e.g. robust estimates of b are *farther* from unity than OLS estimates. Again, however, the hypothesis $a=0$, $b=1$ is rejected.

Data errors can explain thick tails and induce many other problems. For example, Cornell (1989) demonstrated that data errors in forward rates can make particular tests uncover non-existent risk premiums. We noted earlier that data errors have the potential to induce gross alterations in unit root statistics. For example, the Phillips/Perron test of forward premium stationarity produced a 'highly significant' statistic, −8.541, for Japan using all of the NatWest data observations, (Table 4). But when the forward rates were corrected using interest rate parity, the test statistic was only −1.276, far below the critical value and thus not rejecting non-stationarity after all. Using the interest differential instead of the forward premium (the two are identical except for trading costs, which are minimal in exchange markets), Table 6 supports the second result; the Phillips/Perron test statistic for Japan is a very close −1.314. As this case demonstrates clearly, data errors can actually reverse inferences.

Even without data errors, changes over time in exchange rate volatility can induce thick tails, a well-known result of mixing distributions even when the elements of the mixture are Gaussian. Goodhart *et al.* (1992) emphasize this effect when they point to 'regime shifts'. Along with data errors, it is a plausible explanation of the thick-tailed distributions documented by Phillips *et al.* (1996) and by Phillips and McFarland (1997). In both papers, the authors mention particular episodes that might be identified as shifts in regime: the 'Poincaré bear squeeze' of 1924 in the first paper and a 1988 change in Australian banking regulations in the second paper.

Bossaerts (1995) provides an ingenious and different rationale for thick tails and other exchange rate phenomena. He hypothesizes that exchange market traders have been learning on the job since 1973, when the major trading countries adopted floating exchange regimes. If overlapping generations of traders learn in a 'frequentist' (i.e. non-Bayesian) way, the sample paths of particular test statistics can be predicted to have a certain appearance. Bossaerts shows that actual sample paths adhere closely to the predicted paths.

All of these papers are very well-executed, but none contains a complete explanation of the basic puzzle. Most are, in fact, strangely silent on exactly how regression (2) can produce persistently negative slope coefficients, even in the presence of non-stationary variables, data errors, heteroskedasticity, thick tails, and learning. Further work remains.

5.3. *A new procedure*

Because of the various econometric problems just discussed, we dreamed up a different technique for examining the forecasting properties of the forward rate. Our new test searches over a continuum of forecasting functions and selects the particular function with the most stationary forecast errors. The underlying idea is that the market *should* form forecasts based on all available information. This implies that forecast errors should *not* have a unit root;[10] they should exhibit stationary noise.

To be more specific, consider the forecast error defined as

$$\varphi_t(h) = s_{t+1} - h(f_t, X_t), \tag{5}$$

where $h(.)$ is some function not dependent on time and X_t represents a set of additional (to the forward rate) predictor variables known at t. For a sample of size T, $(t = 1, ..., T)$, we could compute a conditional (on h) unit root test statistic, denoted by $\gamma[\varphi(h)]$, from the $\varphi_t(h)$'s. For example, $\gamma[\varphi(h)]$ might be the augmented Dickey/Fuller test statistic calculated from the time series values, $\{\varphi_1(h), \varphi_2(h), ..., \varphi_T(h)\}$. Then $\gamma[\varphi(h)]$ is a functional which could, in principle, be minimized over all choices of $h(.)$ analogously to a calculus of variations problem. Its solution is a bias function, h^*, which provides the best forecast of the spot exchange rate given the assumption of market informational efficiency.

Notice that this technique finesses potential unit root problems in all variables. It does not conduct a spurious OLS regression such as (1) and, consequently, is immune to the infection of non-stationarity. If there is *any* h^* which rejects the unit root hypothesis, that value provides a satisfactory forecast. If there is no function $h(.)$ that rejects non-stationarity in the forecast errors, that is *per se* evidence of market *in*efficiency. This is so because $h(.)$ would incorporate the effect of a time-varying risk premium.

Although this technique could conceivably include any publicly available predictor variables, it would be interesting to ascertain how well the forward rate fares by itself. Moreover, it is clearly easier, though not necessarily better, to prespecify the bias function. Consider, for example, the forecast error

$$\varphi_t(\delta) = s_{t+1} - \delta_0 - \delta_1 f_t, \tag{6}$$

where $\delta \equiv [\delta_0\ \delta_1]$ is a constant (not a function.) In this special case, it is straightforward, albeit computationally time-consuming, to simply evaluate $\gamma[\varphi(\delta)]$ over a grid of values for δ and select that particular value, δ^*, which provides the greatest rejection probability of a unit root. For example, if it turns out that $\delta^* \approx [0\ 1]$, the best linear forecast of the spot rate is simply the unbiased forecast equal to the forward rate.

A special case is to choose δ as the 'cointegrating' vector of s_{t+1} and f_t. This is usually accomplished by running an OLS regression of s_{t+1} on f_t as in (1) and then testing the OLS regression residuals for a unit root. If the residuals are found to not have a unit root, then s_{t+1} and f_t are said to be cointegrated.

[10] This point was emphasized by Campbell and Shiller (1987). Baillie and Bollerslev (1989) report empirical evidence that the simple forward rate forecast error, $s_{t+1} - f_t$, is stationary. Then they make the following curious statement: '... these findings do not imply the absence of a time-varying risk premium. However, the deviations from the unbiasedness hypothesis are stationary, or transitory, in nature,' (p. 174). How would one distinguish a 'transitory' risk premium from a straightforward stochastic error?

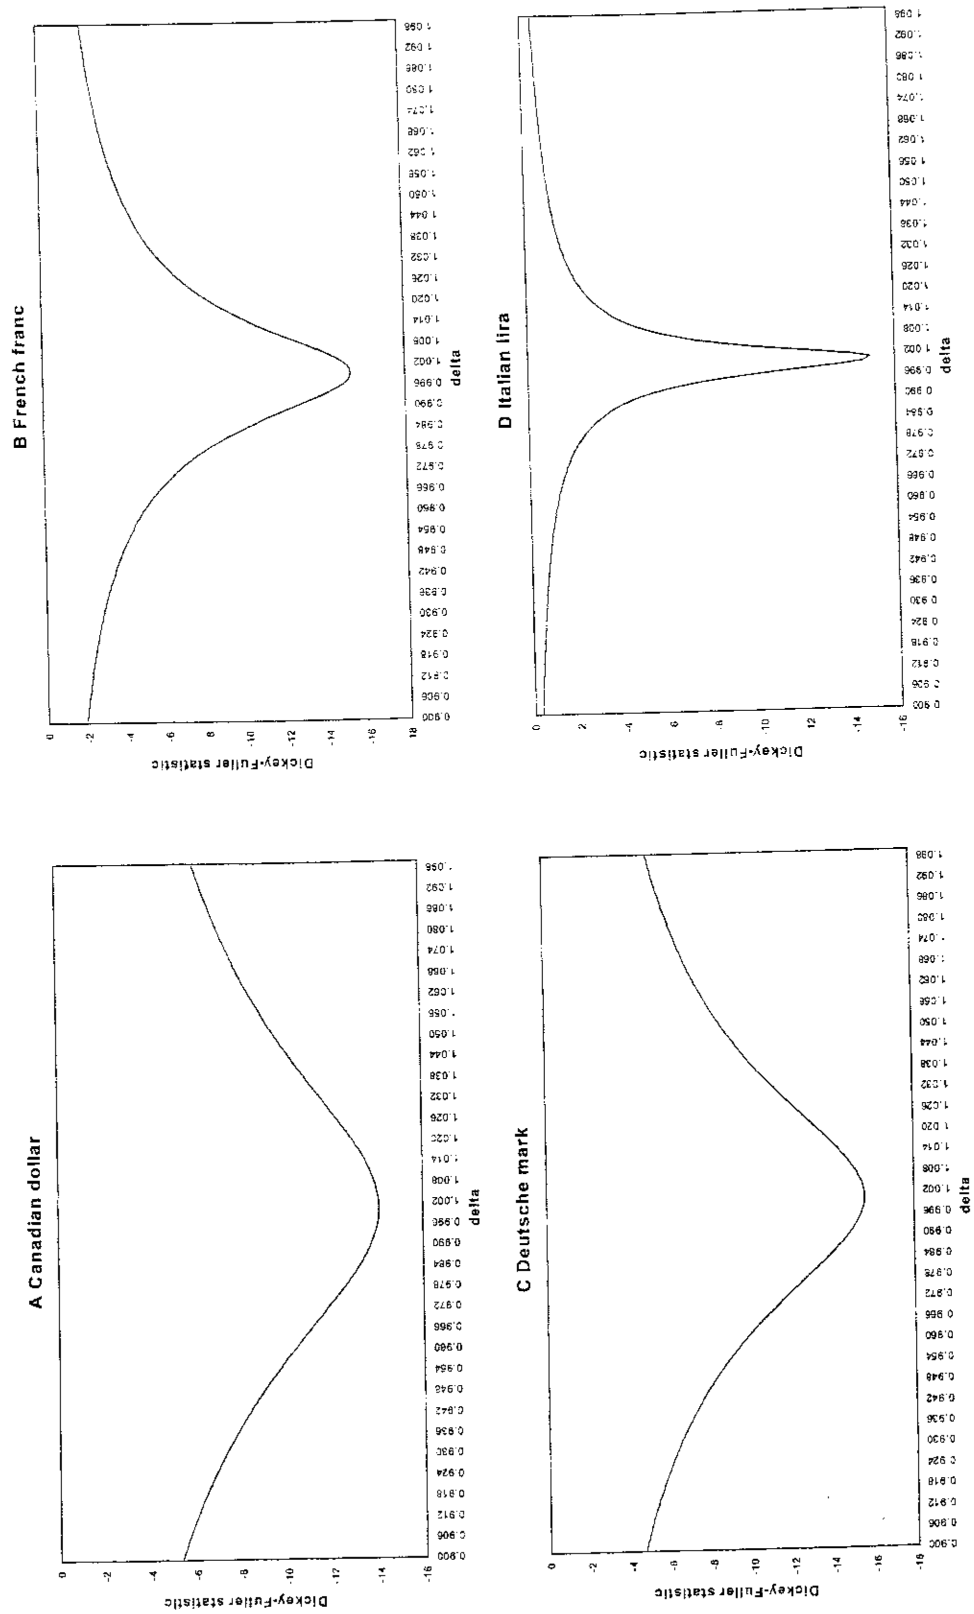

A Canadian dollar
B French franc
C Deutsche mark
D Italian lira
delta
Dickey-Fuller statistic

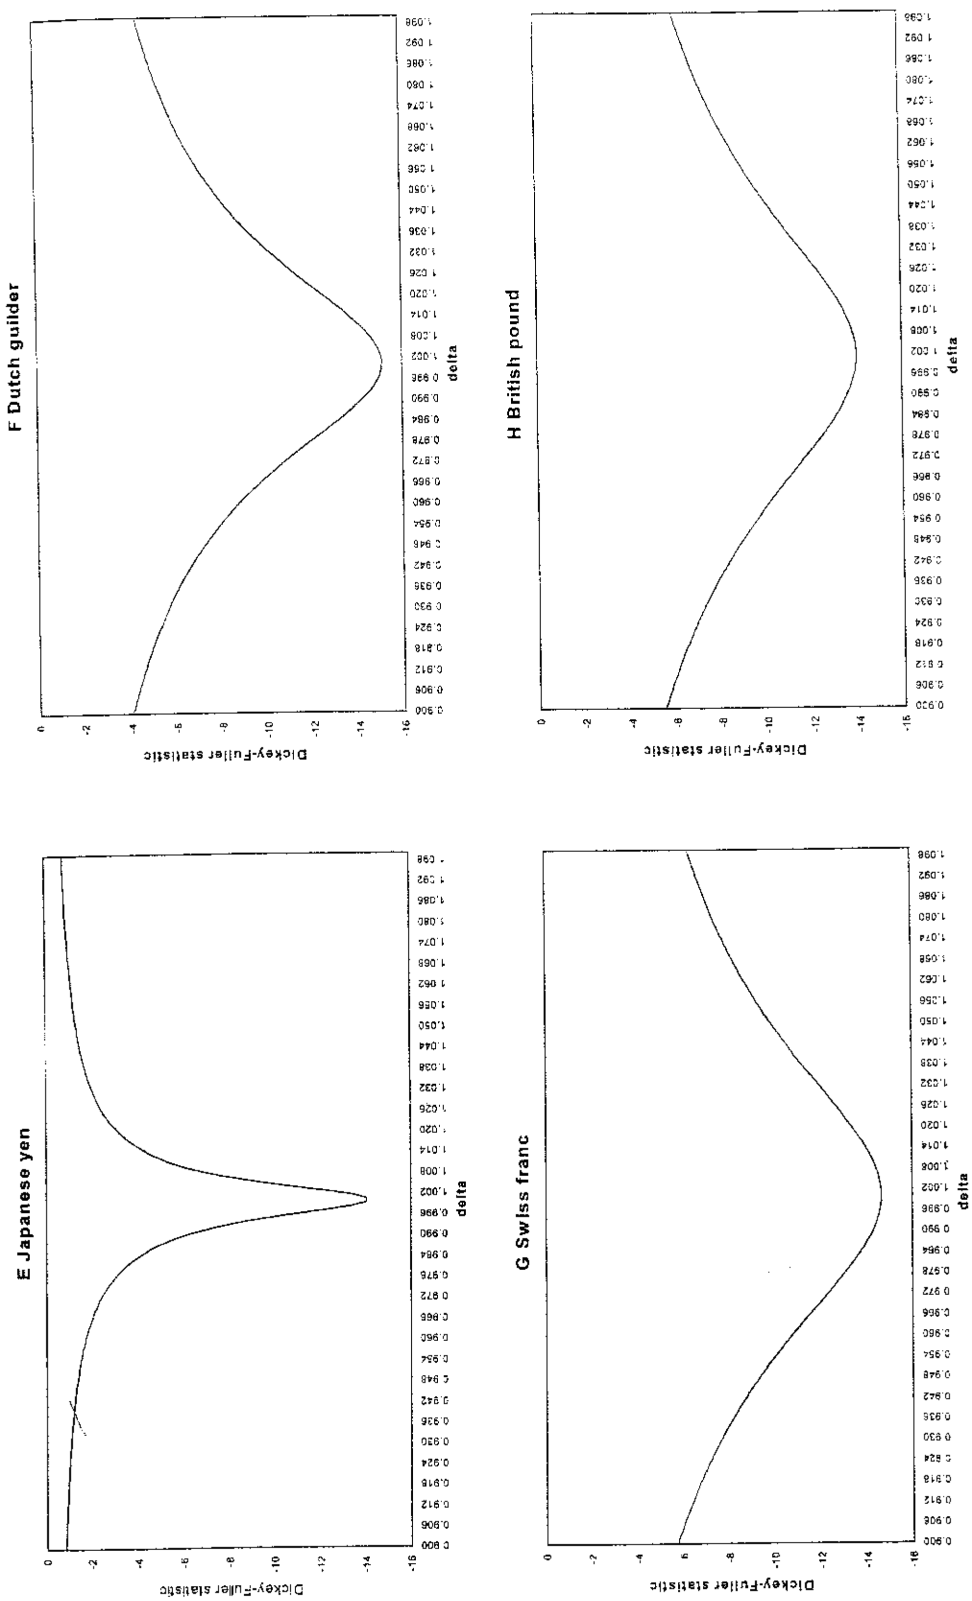

Fig. 1. Dickey-Fuller test statistic for forecast errors: $\varphi_t(\delta) = s_{t+1} - \delta f_t$

The main defect of (6) or its more general form (5) is the possibility that $h(.)$ may be time varying. But the principle of Occam's Razor suggests that we defer this problem for the moment in order to see how well (5) or (6) work when $h(.)$ is a constant.

The panels of Figure 1 plot the simple Dickey/Fuller unit root statistic[11] for a range of δ from the most elementary of all possible bias functions,

$$\varphi_t(\delta) = s_{t+1} - \delta f_t, \tag{7}$$

which is a simplified form of (6) with the constant δ_0 set to zero. The data span our longest Harris sample period, July 1974 through December 1994. As the plots make strikingly evident, the unit root statistic is minimized for every country near the value $\delta = 1.0$. We made calculations while varying δ by increments of 0.001; so, for example, we calculated the test statistic for $\delta = \cdots 0.999,\ 1.000,\ 1.001$, etc. For Canada, Germany, Italy, Japan, The Netherlands, and Switzerland, the minimum Dickey/Fuller statistic was achieved at exactly $\delta^* = 1.000$, while for France, the minimum was located at $\delta^* = 0.999$ and for the UK it was at $\delta^* = 1.001$.

Incredible as it may seem, these data reveal that forward rates provide a virtually unbiased estimates of future spot rates, in most cases to within three significant digits of absolute unbiasedness, (i.e., $\delta = 1.0$). Moreover, at $\delta = 1.0$, the unit root statistic for every country far exceeds in absolute magnitude the critical level corresponding to any reasonable probability that the forecast errors are non-stationary. These forecast errors are decidedly stationary. They reveal a market fully impounding relevant information into the forward rate forecast.

As a check on the robustness of this extraordinary result, we performed the same calculations for the two sub-periods[12] of the Harris Bank data and for the NatWest data. All the results are given in Table 8. As might have been anticipated, there is more variability in the minimizing δ^* across countries during the sub-periods, but its value remains close to 1.0 and the Dickey/Fuller statistic itself is highly significant, thus accepting stationarity in all cases.

Moreover, when the minimizing sample δ^* departs slightly from unity, the function $\gamma[\varphi(\delta)]$ is relatively flat, which suggests that the departure is probably not significant. For example, the largest absolute departure in Table 8 is for Switzerland during the second Harris sub-period when the Dickey/Fuller statistic reaches a minimum of -10.74 at $\delta^* = 0.985$. At $\delta = 1.000$, the statistic remains increases only slightly, to -10.43. This compares to a sharply-peaked curve such as Italy's, where during the same sub-period the Dickey/Fuller statistic's minimum of -10.08 is achieved at $\delta^* = 1.000$, while for $\delta = 0.985$ the statistic is only -2.42. When $\gamma[\varphi(\delta)]$ is relatively flat, there can be some sampling departure from $\delta = 1.0$; but when the function is peaked, the sampling variation appears to be trifling. It both cases, it seems safe to conclude that the forward rate provides a virtually unbiased estimate of the future spot rate.

In retrospect, it might seem odd that so much effort was invested in concocting models of non-stationary risk premiums for horizons of just 1 month. Common sense surely suggests that the forward rate at the beginning of a month *should* incorporate much of what market participants know about the spot rate at the end of the month.

---

[11] I.e. Dickey/Fuller with a single lag and no intercept.

[12] July 1974 through February 1985 is the first Harris sub-period and March 1985 through December 1994 is the second.

Table 8

Forward exchange rate bias functions for unit root statistic minima.

A bias function, $\varphi_t(\delta) = s_{t+1} - \delta f_t$, was computed over a grid of values for δ, where f_t is the log of the forward exchange rate, in currency units per US$, and s_{t+1} is the log of the spot rate one month later. The augmented Dickey/Fuller test was computed for each resulting time series, $\varphi_1(\delta), ..., \varphi_T(\delta)$. The minimum Dickey/Fuller test statistic and its minimizing δ are reported.

	July 74–Dec 94	July 74–Feb 85	March 85–Dec 94	March 85–July 97	July 74–Dec 94	July 74–Feb 85	March 85–Dec 94	March 85–July 97
	Harris			Nat West	Harris			NatWest
	Minimizing δ				Minimum Dickey/Fuller Statistic			
Canada	1.000	1.013	0.995	0.999	−14.10	−10.10	−9.51	−12.72
France	0.999	1.003	0.996	0.994	−15.30	−11.10	−10.80	−13.38
Germany	1.000	1.007	0.988	0.995	−15.60	−11.00	−11.30	−13.48
Italy	1.000	1.000	0.999	0.995	−14.90	−10.50	−10.30	−11.45
Japan	1.000	1.000	0.999	0.993	−14.10	−9.67	−10.40	−11.83
Netherlands	1.000	1.007	0.990	0.994	−15.10	−10.30	−11.30	−13.20
Switzerland	1.000	1.006	0.985	0.995	−14.70	−10.00	−10.80	−13.35
UK	1.001	0.995	1.010	0.995	−14.00	−9.41	−10.70	−13.51

Over such a limited horizon, the risk premium must be small. Remember, it is a risk premium *entirely* dependent upon uncertainty in the two countries' relative inflation rates; this cannot have been momentous during most periods for the large developed countries in our sample. Indeed, we have found it too small to measure, if it exists at all.

6. The delusion of a 'puzzle'

Our empirical sleuthing has left an annoying loose end. If forward exchange rates really are unbiased predictors of future spot rates, why did (2) produce 'significant' negative slope coefficient in some sample periods. Those sample periods were relatively long, often including more than 100 months, and the phenomenon sometimes appeared simultaneously across all sample countries.

New simulations reported in Table 9 help explain why a spurious negative coefficient might appear in a single country. We estimated a regression analogous to (2), where the explanatory variable, x, is non-stationary or nearly non-stationary. If the non-stationarity parameter ρ is exactly unity, x follows a random walk and thus has a unit root. When ρ is slightly less than 1.0, x is actually stationary, but unit root tests have a hard time confirming that fact in modest-sized samples.

The dependent variable, y, is equal to x plus white noise. Consequently, y also is stationary when $\rho < 1$ and is non-stationary when $\rho = 1$; but since the noise is

Table 9

Simulation of regression model with unit root and near unit root variables.

The true model is $y_t = a + bx_t + \varepsilon_t$ where $a = 0$, $b = 1$, ε_t is normally and independently distributed with mean zero and standard deviation σ_ε and x follows a time series process, $x_t = \rho x_{t-1} + \nu_t$ whose disturbance ν_t is normally and independently distributed with mean zero and variance unity. Time series sample sizes are 120 and 240. σ_ε takes values 10, 20, 30 and 40 in the four sections of the table. For the estimated slope, the table reports mean, standard deviation, and the number negative in 10,000 replications.

	$T = 120$			$T = 240$		
ρ	Mean	Standard Deviation	Number negative	Mean	Standard Deviation	Number negative
			$\sigma_\varepsilon = 10$			
1.0	1.002	0.274	17	1.000	0.138	0
0.98	1.002	0.330	46	1.001	0.184	0
0.96	1.003	0.369	67	1.000	0.220	0
0.94	1.003	0.404	95	1.000	0.252	2
0.92	1.004	0.436	118	1.000	0.280	7
0.90	1.004	0.466	166	1.000	0.306	12
0.88	1.005	0.494	205	1.001	0.329	18
0.86	1.005	0.520	254	1.002	0.349	26
0.84	1.005	0.544	313	1.002	0.368	36
0.82	1.005	0.566	383	1.003	0.386	55
0.80	1.006	0.588	455	1.003	0.403	68

(continued)

Table 9

Continued

ρ	$T=120$			$T=240$		
	Mean	Standard Deviation	Number negative	Mean	Standard Deviation	Number negative
			$\sigma_\varepsilon=20$			
1	0.995	0.556	399	1.000	0.274	24
0.98	0.991	0.668	688	1.000	0.377	84
0.96	0.991	0.748	866	0.998	0.452	168
0.94	0.990	0.819	1062	0.996	0.516	293
0.92	0.990	0.884	1256	0.996	0.572	424
0.90	0.989	0.946	1448	0.995	0.621	535
0.88	0.989	1.002	1599	0.995	0.666	655
0.86	0.989	1.055	1739	0.995	0.706	772
0.84	0.989	1.105	1848	0.996	0.743	890
0.82	0.989	1.151	1948	0.996	0.778	975
0.80	0.989	1.195	2033	0.996	0.810	1070
			$\sigma_\varepsilon=30$			
1	0.996	0.824	923	1.000	0.409	134
0.98	0.997	0.991	1389	1.001	0.558	355
0.96	0.999	1.105	1681	1.002	0.664	622
0.94	1.000	1.209	1928	1.002	0.756	883
0.92	1.000	1.306	2141	1.001	0.838	1102
0.90	1.000	1.395	2297	1.000	0.911	1305
0.88	1.000	1.479	2411	1.000	0.978	1504
0.86	1.000	1.557	2536	0.999	1.040	1636
0.84	1.000	1.630	2617	0.998	1.097	1773
0.82	1.000	1.699	2700	0.997	1.150	1868
0.80	0.999	1.763	2793	0.997	1.199	1951
			$\sigma_\varepsilon=40$			
1	0.992	1.090	1526	1.007	0.547	344
0.98	1.006	1.310	1951	1.002	0.758	867
0.96	1.014	1.472	2254	1.002	0.907	1246
0.94	1.016	1.615	2504	1.003	1.035	1537
0.92	1.018	1.747	2703	1.005	1.148	1783
0.90	1.019	1.869	2873	1.006	1.249	1984
0.88	1.019	1.984	3019	1.007	1.340	2196
0.86	1.021	2.091	3132	1.008	1.423	2345
0.84	1.022	2.191	3197	1.008	1.500	2476
0.82	1.023	2.285	3259	1.008	1.571	2583
0.80	1.024	2.373	3325	1.008	1.637	2682

Table 10

Relative volatility: changes in spot exchange rates vs. the forward premium estimated from first-order autoregression residuals.

s_t is the log of the spot exchange rate on date t and f_t is the log of the 1-month forward rate. To purge serial dependence, separate autoregressive models were estimated for $s_{t+1} - s_t$ and for $f_t - s_t$:

$s_{t+1} - s_t = a + b(s_t - s_{t-1}) + \varepsilon_{t+1}$ and $f_{t-1} - s_{t-1} = c + d(f_t - s_t) + \xi_{t+1}$.

The table reports relative volatility of the autoregression residuals. Harris, monthly observations

	$\sigma_\varepsilon/\sigma_\xi$		
	July 74–December 94	July 74–February 85	March 85–December 94
Canada	16.2	13.5	22.7
France	13.9	10.0	32.6
Germany	38.6	27.8	105.1
Italy	10.5	7.0	25.6
Japan	24.3	17.8	68.1
Netherlands	21.6	15.0	96.5
Switzerland	39.8	30.7	80.5
UK	29.3	18.9	81.1

The ratios are much larger in the second sub-period because σ_ε did not change very much while σ_ξ was considerably smaller.

substantial, unit root tests will tend to reckon it stationary in all circumstances, as our previous simulations (Table 5) disclose.

The body of Table 9 reports means, volatilities, and frequencies of negative simulated slope coefficients, b. In what at first might seem a paradox, the results show that a regression such as (2) produces the most reliable results when x is *non-stationary*. In this circumstance, the coefficients a and b are strongly consistent, cf. Hamilton (1994, pp. 586–7). The spurious regression problem with unit root variables, which tends to bias slope coefficients upward when $b < 1$, does not cause trouble when $b=1$. To the contrary, having a unit root reduces estimation error.

In contrast, when x is actually stationary but ρ is not far below unity (i.e. near non-stationarity), the regression is likely to give misleading estimates, especially in the presence of significant random noise. In Table 9's repeated samples of size $T = 120$ and $T = 240$, the overall mean estimate of b is close to 1.0, but the proportion of samples with negative b's is alarming, reaching over 30% for, $T = 120$ and $\sigma_\varepsilon = 40$.

Table 9's parameter values are realistic in the sense that exchange rate data are similar. The forward premium (and the nominal interest differential) are near non-stationary so, for example, a ρ of 0.8–0.9 is within a plausible range. The sample size from the mid-1970s through the mid-1980s was around 120 months at most. Finally, as shown in Table 10, spot exchange rate surprises are orders of magnitude more volatile than movements in the forward premium.[13]

[13] To mitigate the confounding effects of any first-order serial dependence, Table 10 reports the relative volatilities of residuals from autoregressions of the variables.

But if the observed negative b's are spurious, why did many countries display them during the same period? If the eight countries in our sample and in the samples of other researchers were cross-sectionally independent, probably not all would simultaneously display spurious negative coefficients. But, of course, these countries are related and the exchange rates are *all* expressed per US dollar. If US interest rates embark on a quasi-non-stationary trajectory during some period, the forward premium with most other countries is likely to follow. Given the results of our new test, which support the unbiased forward rate hypothesis for all countries during the very period when the slope coefficient estimate in (2) was 'significantly' negative, we believe cross-country dependence is at least a serious contender for the final link in the correct explanation of the puzzle. We could be wrong.

7. Summary and conclusion

The forward premium's negative correlation with later changes in spot exchange rates has been puzzling to scholars and has persuaded some that exchange markets are inefficient. We have found the correlation to be sample-period dependent and have offered an explanation for why it can sometimes seem negative.

Our explanation is based on empirical peculiarities that induce problems in simple regression models.

1. The levels of both spot and forward exchange rates *and* the forward premium are nearly non-stationary.
2. Changes in spot exchange rates are noisy; unexpected changes appear to dominate changes in expectations.
3. Regressions of noisy variables on nearly unit root variables produce spurious estimated coefficients with surprising frequency.

We devised a new test of forecasting ability for a financial variable that depends on market expectations. It exploits the principle that rational forecasts should produce *stationary* forecast errors. Applying our test to monthly exchange rate data, we found that the one-month forward exchange rate, unbiased, produces stationary forecast errors. Over the class of all forecasts of the form δf, where δ is a proportional constant and f is the forward rate, the choice $\delta = 1$ produces forecast errors with the most negative unit root statistic, relative to all other values of δ.

Our results suggest that forward rates are virtually unbiased estimates of future spot rates. Over a 1-month horizon, no time varying risk premium is detectable. Moreover, it is not necessary for explaining the empirical data.

References

Baillie, R. T. and Bollerslev, T., 'Common stochastic trends in a system of exchange rates', *Journal of Finance*, Vol. 44, 1 March 1988, pp. 167–181.

Bilson, J. F. O., 'The "speculative efficiency" hypothesis', *Journal of Business*, Vol. 54, 3, July 1981, pp. 435–451.

Boothe, P. and Longworth, D., 'Foreign exchange market efficiency tests: implications of recent empirical findings', *Journal of International Money and Finance*, Vol. 5, 2, June 1986, pp. 135–152.

Bossaerts, P., 'The econometrics of learning in financial markets', *Econometric Theory*, Vol. 11, 1, March 1995, pp. 151–189.

Campbell and Shiller, 'Cointegration are tests of present value—models', *Journal of Political Economy*, Vol. 95, 5, 1987, pp. 1062–1088.

Corbae, D., Lim, K.-G. and Ouliaris, S., 'On cointegration and tests of forward market unbiasedness', *Review of Economics and Statistics*, Vol. 74, 4, November 1992, pp. 728–732.

Cornell, B., 'The impact of data errors on measurement of the foreign exchange risk premium', *Journal of International Money and Finance*, Vol 8, 1, March 1989, pp. 147–157.

Cumby, R. E. and Obstfeld, M., 'International interest-rate and price-level linkages under flexible exchange rates: a review of recent evidence', in J. F. O. Bilson and R. Marston, (eds), *Exchange Rates: Theory and Practice* (Chicago: University of Chicago Press, 1984).

Campbell, J. Y. and Shiller, R. J., 'Cointegration and tests of present value models', *Journal of Political Economy*, Vol. 95, 5, October 1987, pp. 1062–1088.

Dickey, D. A. and Fuller, W. A., 'Likelihood ratio statistics for autoregressive time series with a unit root', *Econometrica*, Vol. 49, 4, July 1981, pp. 1057–1072.

Dickey, D. A. and Fuller, W. A., 'Distributions of the estimators for auto-regressive time series with a unit root', *Journal of the American Statistical Association*, Vol. 74, 1979, pp. 427–431.

Doukas, J. and Rahman, A., 'Unit root tests: evidence from the foreign exchange futures markets', *Journal of Financial and Quantitative Analysis*, Vol. 22, March 1987, pp. 101–108.

Fama, E., 'Forward and spot exchange rates', *Journal of Monetary Economics*', Vol. 14, 3 November 1984, pp. 319–338.

Goodhart, C. A. E., McMahon, P. C. and Ngama, Y. L., 'Does the forward premium/discount help to predict the future change in the exchange rate?', *Scottish Journal of Political Economy*, Vol. 39, 2, May 1992, pp. 129–140.

Hakkio, C. S. and Rush, M., 'Market efficiency and cointegration: an application to the sterling and deutschemark exchange markets', *Journal of International Money and Finance*, Vol. 8, 1, March 1989, pp. 75–88.

Hamilton, J. D., *Time Series Analysis* (Princeton: Princeton University Press, 1994).

Hansen, L. P. and Hodrick, R. J., 'Forward exchange rates as optimal predictors of future spot rates: an econometric analysis', *Journal of Political Economy*, Vol. 88, 5, October 1980, pp. 829–853.

Hodrick, R. J., *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets* (Amsterdam: Harwood Academic Publishers, 1987).

Hodrick, R. J. and Srivastava, S., 'An investigation of risk and return in forward foreign exchange', *Journal of International Money and Finance*, Vol. 3, 1, April 1984, pp. 1–29.

Huang, R. D., 'Some alternative tests of forward rates as predictors of future spot rates', *Journal of International Money and Finance*, Vol. 3, 1, August 1984, pp. 153–168.

Korajczyk, R., 'The pricing of forward contracts for foreign exchange', *Journal of Political Economy*, Vol. 93, 2, April 1985, pp. 346–368.

Longworth, D., 'Testing the efficiency of the Canadian–US exchange market under the assumption of no risk premium', *Journal of Finance*, Vol. 36, 1 March 1981, pp. 43–49.

MacKinnon, J. G., 'Critical values for cointegration tests', in R. F. Engle and C. W. J. Granger, (eds), *Long-Run Economic Relationships: Readings in Cointegration* (Oxford: Oxford University Press, 1991).

Meese, R. A. and Singleton, K. J., 'On unit roots and the empirical modeling of exchange rates', *Journal of Finance*, Vol. 37, 4, September 1982, pp. 1029–1035.

Perron, P., 'Trends and random walks in macroeconomic time series', *Journal of Economic Dynamics and Control*, Vol. 12, 2/3, June/September 1988, pp. 297–332.

Phillips, P. C. B., 'Time series regression with a unit root', *Econometrica*, Vol. 55, 2, March 1987, pp. 277–301.

Phillips, P. C. B. and Perron, P., 'Testing for a unit root in time series regressions', *Biometrika*, Vol. 75, 1988, pp. 335–346.

Phillips, P. C. B., McFarland, J. W. and McMahon, P. C., 'Robust tests of forward exchange market efficiency with empirical evidence from the 1920s', *Journal of Applied Econometrics*, Vol. 11, 1, January–February 1996, pp. 1–22.

Phillips, P. C. B. and McFarland, J. W., 'Forward exchange market unbiasedness: the case of the Australian dollar since 1984', *Journal of International Money and Finance*, Vol. 16, 6, December 1997, pp. 885–907.