Real Exchange Rates, Efficient Markets and Uncovered Interest Parity: A Review

John E. Floyd
University of Toronto[1]

September 25, 2007

[1]I would like to thank my colleagues Gordon Anderson, Miguel Faig, Allan Hynes, Alex Maynard and Angelo Melino for helpful comments on this material.

Contents

1 Introduction

This paper reviews the evidence regarding recent real and nominal exchange rate experience and develops, as far as possible, a coherent explanation of that evidence using the basic principles of economic theory and econometrics and paying particular attention to the recent literature. To make its contents available to students and others not well-versed in technical issues the exposition is much more careful, simplistic and extensive than would otherwise be required. Appendices covering data sources, elementary time-series analysis and other background material in econometrics are included along with extensive references to textbooks and other sources from which appropriate technical background can be obtained.

There is an extensive set of facts to be explained. Real and nominal exchange rates are highly correlated with each other and much more variable around their trends than are the ratios of the respective countries' price levels. Both can be described as near-random-walk variables that differ from each other in trend as a result of differences in countries' inflation rates. Spot and forward exchange rates move nearly in unison but the period-to-period movements in spot rates are much greater than would be predicted by the corresponding forward premia. Typically, the level of the forward rate predicts next period's spot rate quite well, but the forward premium is a very poor predictor of the change in the spot rate between the current period and the next. Covered interest parity—the equality of the domestic/foreign interest rate differential with the forward discount on the domestic currency—tends to hold to a reasonable approximation but uncovered interest parity—the equality, on average, of the domestic/foreign interest rate differential with the actual movement of the nominal exchange rate from the current to next period—does not. The purchasing power parity theory, which states that nominal exchange rate movements should exclusively reflect the underlying movements in the domestic and foreign price levels with real exchange rates constant, is inconsistent with the evidence although over very long periods real exchange rates tend to return toward average levels.

With the addition of data from more recent years, the time period studied here is longer than that covered in most of the literature. Accordingly, apart from the Canadian case, it is now possible to reject on the basis of monthly data for the period 1957 through 2002 the hypothesis that real exchange rate series are random walks in favour of the alternative that there is a small degree of mean reversion. This hypothesis could previously be rejected using annual data extending back for a century or more.

During the past two decades the convention has been to analyse nominal exchange rate movements within the framework of an asset theory of the exchange rate and apply the principles of modern finance to explain them. The risk attached to an asset must thereby be related directly to the covariance of its return with the return to capital, somehow measured, in the economy as a whole. In this tradition, exchange rate movements have generally been interpreted as deviations from some long-run purchasing power parity equilibrium relationship. And foreign exchange market efficiency has tended to be judged in terms of whether forward exchange rates can provide unbiased forecasts of these movements, although it is now recognised that apparent market inefficiency may be explained instead by time-varying risk premia.

Here we explore the implications of defining the real exchange rate as the relative price of domestic output in terms of foreign output, taking a structural view of its determination based on the differential effects of ongoing technological change, economic growth and political developments on countries' relative output prices. This structural interpretation complements rather than replaces the asset market perspective. Nevertheless, once we recognise that changes in the international relative price structure are as unpredictable to agents as they are to economists it becomes unreasonable to expect forward premia to predict future nominal exchange rate movements with any reliability apart from cases where there are continuing long-term differences in countries' inflation rates. In the absence of major inflation rate differences, forward exchange rates will always move in near unison with spot rates because the best prediction of tomorrow's exchange rate tends to be today's exchange rate. These structural aspects of real exchange rate behaviour have important implications for the relationship between observed exchange rate movements and foreign exchange market efficiency. In particular, it is possible to show that market efficiency is consistent with a zero correlation between forward premia and changes in the future spot exchange rates. This enables us to explain major features of the well-known forward premium anomaly—the failure of uncovered interest parity to hold. While we can explain the situation where forward premia show no significant relationship to future changes in spot exchange rates, we still cannot explain why the correlations between forward premia and subsequent changes in the spot rate tend very frequently to be negative, significantly so in the 1980s.

2 Real and Nominal Exchange Rates and Relative Price Levels

Figure 1 plots the real and nominal exchange rates and the price-level ratios of Canada, the United Kingdom and Japan with respect to the United States. Figure 2 presents similar plots for Germany and France with respect to the United States and for France with respect to Germany. The plots relative to the U.S. run from 1957 to 2002,while the France/Germany plot is for the period 1990 through 2002. Apart from the France/Germany case, the real exchange rates are defined as the ratios of the respective countries' consumer price indexes to the U.S. consumer price index after the former have been multiplied by the U.S. dollar price of the domestic currency. In the France/Germany case, both price indexes are multiplied by the U.S. dollar price of home currency and the resulting series for France is divided by the corresponding series for Germany. Nominal exchange rates are expressed as U.S. dollar prices of domestic currency in all cases except France/Germany where the Deutschmark price of the franc is used. The price-level ratios are simply the ratios of the respective consumer price indexes unadjusted for exchange rate changes.

Notice two important regularities in the plots involving the United States. First, the ratios of the various countries' price levels to the U.S. price level are very smooth in comparison to the corresponding real and nominal exchange rates. The Canada/U.S. price level ratio seems more variable than the others at first glance but it is actually less variable—the apparent greater variability is an illusion stemming from the fact that the scale along the vertical axis is less compressed in the Canadian case than in the other cases. Canada's price level only increased relative to the U.S. price level by about 5% over the period 1957–2002 and at the peak it was only 15% above the period's lowest price level ratio. In contrast, the ratio of the U.K. to the U.S. price level rose in excess of 200% along a rather smooth time path. The price levels of Germany and Japan fell rather smoothly relative to the U.S. price level from the late 1960s in the case of Germany and the late 1970s in the case of Japan right through to the end of the period.

The second important regularity is the high degree of correlation between countries' real and nominal exchange rates with respect to the United States. We can express the real exchange rate as

$$Q = \frac{\Pi P}{P^*} \tag{1}$$

with Q being the real exchange rate, Π the nominal exchange rate, defined

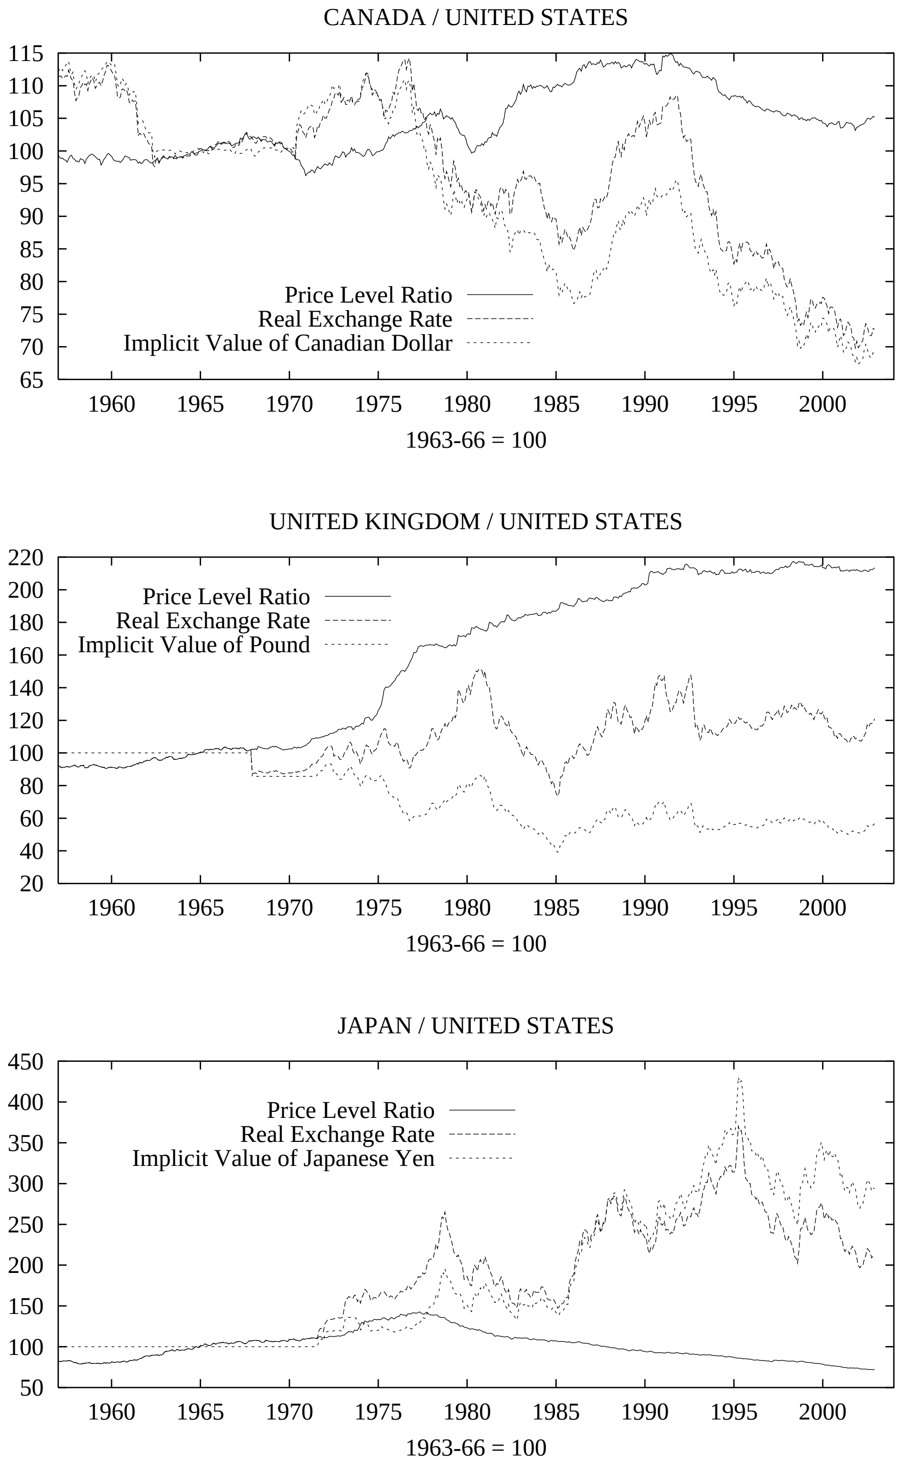

Figure 1: Real exchange rate, price level ratio and nominal exchange rate (price of the domestic currency in U.S. dollars) for Canada, the United Kingdom and Japan, January 1957 through December 2002, 1963–66 = 100. Source: *International Financial Statistics*.

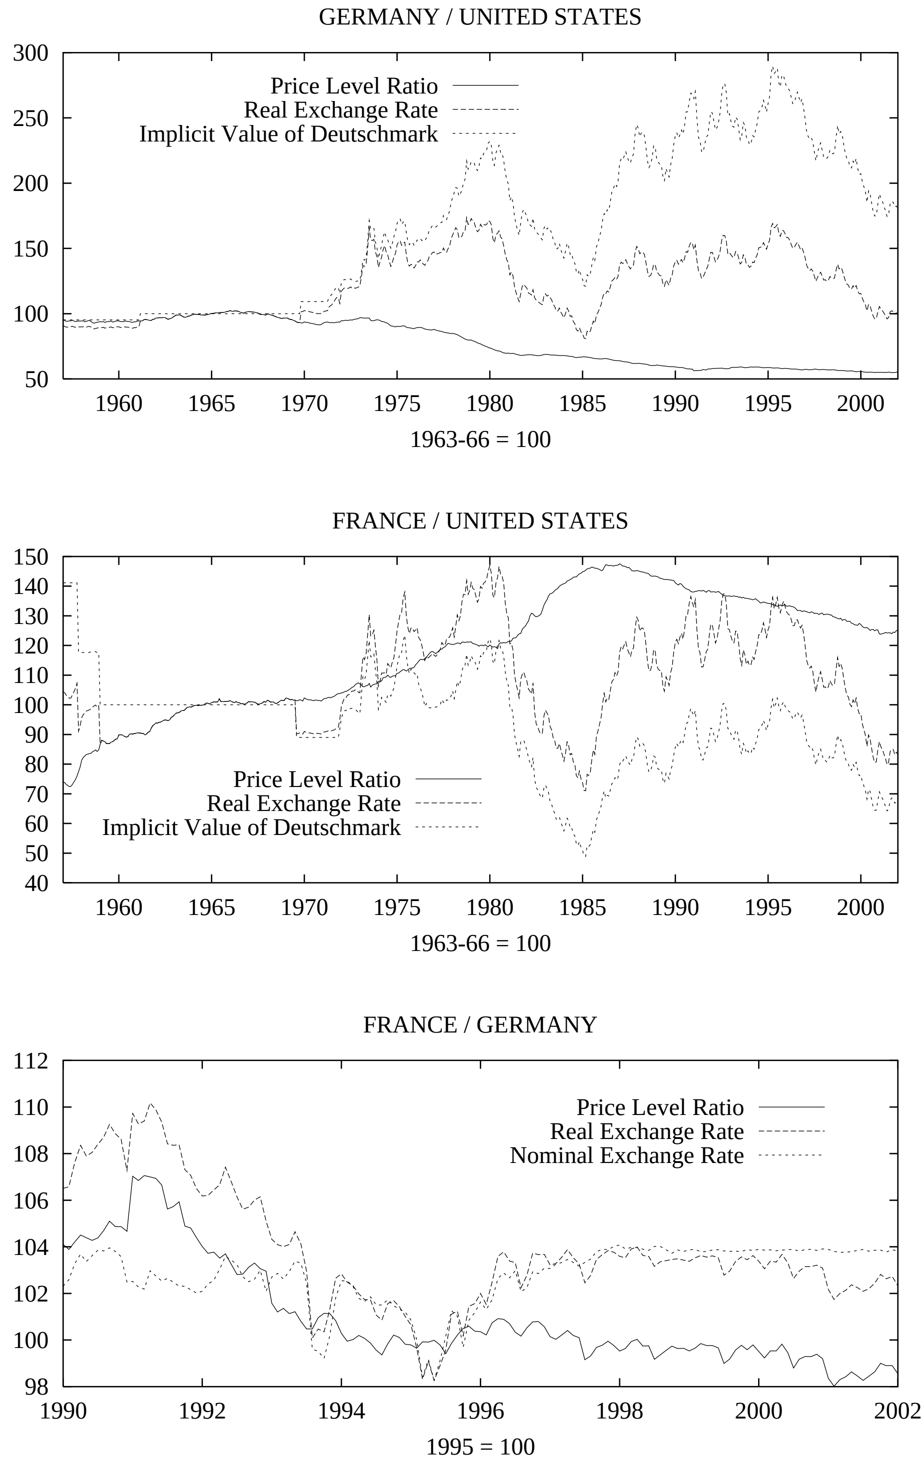

Figure 2: Real exchange rate, price level ratio and nominal exchange rate (price of the domestic currency in units of foreign currency) for Germany and France viz. à viz. the U.S., January 1957 through December 2002, 1963–66 = 100, and France viz. à viz. Germany, January 1995 through December 2002, 1995 = 100. Source: *International Financial Statistics*.

as the foreign currency price of domestic currency, P the domestic price level and P^* the foreign price level. This expression can be rearranged to yield

$$\Pi \quad = \quad \frac{QP^*}{P.} \tag{2}$$

In Figure 1 and the top two panels of Figure 2, most of the short-term (higher-frequency) variations in Π are matched by variations in Q while variations in the ratio of P to P^* are reflected in the trend in Π. In the cases of Canada and Japan, there are also trends in Q that are reflected in Π.

In the France/Germany case plotted in the bottom panel of Figure 2, the patterns are quite different. Except for the period from late 1993 to the beginning of 1996, the real exchange rates and price level ratios appear highly correlated with each other while the nominal exchange rate between the two currencies shows much less variability, having a more or less horizontal trend. During the 1993–1996 period the pattern corresponds to the regularities that appear in the previous plots of the various countries' real and nominal exchange rates and price-level ratios with respect to the United States. The high correlation between the real exchange rate and the price-level ratio France with respect to Germany along with the greater stability of the nominal exchange rate between the two currencies over most of the period reflects the fact that France was attempting to stabilise its nominal exchange rate with respect to Germany in preparation for the European Currency Union that was formally adopted in 1998, after which year the nominal exchange rate series is virtually horizontal. Under these circumstances, as is particularly evident after 1998, real exchange rate movements are reflected in movements in the countries' relative price levels. The high correlation between the nominal and real exchange rates during the two-year period after late 1993 reflects the currency crises and subsequent exchange rate instability that occurred during those years. The real exchange rate fell by about 6 percent between 1991 and 1995. By keeping the franc relatively stable on average in terms of the Deutschmark, the French authorities forced this adjustment almost entirely onto the French price level. A comparison of the top two panels of Figure 2 indicates that the French and German real exchange rates with respect to the U.S. show essentially the same pattern since the break-down of the Bretton-Woods system in 1973. The trend in the nominal exchange rate, however, was downward in France and upward in Germany, reflecting the fact that the German inflation rate was lower, and the French inflation rate higher, than the rate of inflation in the U.S.

3 Real Exchange Rates as Near-Random-Walks

Another important fact that needs explanation is the tendency of real exchange rates to exhibit near-random-walk behaviour. To illustrate the random-walk concept, let us represent a real exchange rate series by the following equation[1]

$$q_t = (1-\rho)\,\bar{q}_t + \rho\, q_{t-1} + \epsilon_t \tag{3}$$

where q_t is the logarithm of the real exchange rate, $\bar{q}_t$ is its trend value in period t, and ϵ_t is a white noise error term. If $\rho = 1$, this equation reduces to

$$q_t = q_{t-1} + \epsilon_t \tag{4}$$

and q_t is a random walk. If $\rho = 0$ the equation reduces to

$$q_t = \bar{q}_t + \epsilon_t \tag{5}$$

and the exchange rate varies randomly around a trend. The parameter ρ can be called the mean reversion parameter—as ρ varies from unity to zero the degree of period-to-period mean reversion goes from zero to complete mean reversion. The persistence of movements of q_t thus depends on ρ. As ρ goes to unity, every movement in the series becomes permanent; as ρ goes to zero, every movement of q_t becomes a deviation from a fixed trend value and the series exhibits no persistence at all.

Two important results emerge in the random-walk case where ρ is equal to unity. First, assuming that the error term ϵ_t is unpredictable, the best prediction of tomorrow's real exchange rate is the level of the real exchange rate today. Second, the real exchange rate will wander far and wide with no tendency to return to any initial level. If ρ is greater than unity, the time path of the real exchange rate will be explosive.

In fact, of course, the time-series properties of real exchange rates—that is, the properties of the equation that best describes their evolution through time—are more complicated than the simple illustration provided by equation (3) above. More appropriate representations would be

[1]Some very elementary principles of time series analysis pertinent to the discussion that follows are outlined in Appendix B. Beyond that, you should read James H. Stock and Mark W. Watson, *Introduction to Econometrics*, Addison Wesley, 2003, Chapter 12, and Walter Enders, *Applied Economic Time Series*, John Wiley and Sons, 1995, Chapter 4.

$$q_t = \alpha + \rho_1\, q_{t-1} + \rho_2\, q_{t-2} + \rho_3\, q_{t-3} + \rho_4\, q_{t-4} + \epsilon_t \tag{6}$$

or perhaps

$$q_t = \alpha + \rho_1\, q_{t-1} + \rho_2\, q_{t-2} + \rho_3\, q_{t-3} + \rho_4\, q_{t-4} + \epsilon_t + \gamma_1\, \epsilon_{t-1} + \gamma_2\, \epsilon_{t-2} \tag{7}$$

where the parameter α performs a role similar but not limited to the role played by $(1-\rho)\, \bar{q}_t$ in (3), and where the included lagged values of q and lagged error-terms need not be restricted to four and two respectively. Equation (6) is an autoregressive process with four lags [AR(4)] while (7) is a autoregressive-moving-average process with four autoregressive lags and a moving-average of two lags of the error term [ARMA(4,2)]. One might also add terms of the form $\delta\, t$ to (6) and (7) to incorporate the possibility that q_t might fluctuate around a deterministic trend—the terms α and $\delta\, t$ are both deterministic in that they do not depend on current or past values of the stochastic process ϵ_t.

By adding and subtracting $\rho_2\, q_{t-1}$, $\rho_3\, q_{t-1}$, $\rho_4\, q_{t-1}$, $\rho_3\, q_{t-2}$, $\rho_4\, q_{t-2}$ and $\rho_4\, q_{t-3}$, rearranging the terms, and expressing $q_t - q_{t-1}$ as Δq_t, equation (6) can be converted into the form

$$\Delta q_t = \alpha - (1-\rho)\, q_{t-1} + \beta_1\, \Delta q_{t-1} + \beta_2\, \Delta q_{t-2} + \beta_3\, \Delta q_{t-3} + \epsilon_t \tag{8}$$

where

$$\begin{aligned} \rho &= \rho_1 + \rho_2 + \rho_3 + \rho_4 \\ \beta_1 &= \rho_2 + \rho_3 + \rho_4 \\ \beta_2 &= \rho_3 + \rho_4 \\ \beta_3 &= \rho_4 . \end{aligned}$$

As in the case of (3), stationarity or mean reversion requires that $\rho < 1$ and the real exchange rate will be a random walk if $\rho = 1$.

It turns out that an equation like (7) that includes moving-average terms can be expressed in the form of a pure autoregressive process like equation (6) containing an infinite number of autoregressive lags [AR(∞)]. Simply reorganise (7) to move ϵ_t to the left of the equality and q_t to the right, lag the resulting equation repeatedly to obtain expressions for ϵ_{t-1}, ϵ_{t-2}, ϵ_{t-3} ... etc. and substitute these expressions successively into (7) and simplify. The resulting infinite order autoregressive process can then be converted into an equation like (8) containing an infinite succession of lags of Δq_t.[2]

[2] See the Enders book, pages 225-227.

Our problem is to determine whether the time-series processes that can reasonably describe the evolution of actual real-world real exchange rates indicate that those series are random walks. If they are not random walks, we need to determine how fast real exchange rates revert to their mean levels. To do this we use ordinary-least-squares to estimate equations like (8) containing an appropriate number of autoregressive lags, and possibly but not necessarily constant and trend terms, to see if we can reject the null hypothesis that $(1-\rho) = 0$. It turns out that, under the null hypothesis that $\rho = 1$, an infinite-order autoregressive process like (8) can be well approximated by a process containing no more than $T^{1/3}$ lags where T is the number of observations.[3] To select the appropriate number of lags to include we can start with an unreasonably large number and progressively drop the longest lag if that lag turns out to be statistically insignificant. Alternatively, we can calculate AIC and BIC information criteria for regressions performed for each number of lags and pick the configuration for which either or both of these statistics are minimised.[4] Of course, all these significance tests and criteria comparisons must apply to regressions estimated from the same number of observations. A constant term and a trend term should be included in the regressions where a plot of the series indicates that a trend appears to be present. The constant term will capture any tendency of the series to 'drift' upward or downward by a constant amount per period, while the inclusion of a trend term δt will allow this drift to increase or decrease at a constant rate through time.

It turns out that the OLS estimates of $(1 - \rho)$ and the coefficients of deterministic regressors such as the constant and trend terms, if those are included, are not distributed according to the standard t-distribution. The appropriate tables of critical values to use in evaluating the significance of these coefficients have been calculated by David Dickey and Wayne Fuller and can be found on pages 223, 419 and 421 of the book by Walter Enders referred to in footnote 1.

A severe problem with these tests is that they have low power when ρ is close to unity—that is, they will lead to rejection of the false null hypothesis only a small proportion of the time. The exchange rate may be stationary in a large fraction of cases, even though we cannot reject non-stationarity.[5]

The tests described thus far, known as Dickey-Fuller tests, assume that

[3] See S. Said and David Dickey, "Testing for a Unit Root in Time Series Regression", *Biometrica*, Vol. 75, No. 2 (June), 1988, 311-40.

[4] See pages 455-457 of the book by Stock and Watson for a discussion of these criteria and the formulas to use in calculating them.

[5] For a more detailed discussion of the power of these tests, see Appendix B.

the errors ϵ_t are statistically independent of each other and have a constant variance. An alternative procedure, developed by Peter Phillips and Pierre Perron, can be used to conduct the tests under the assumption that there is some interdependence of the errors and they are heterogeneously distributed.[6] The following equations are estimated by ordinary-least-squares:

$$q_t = a_0 + a_1\, q_{t-1} + a_2(t - T/2) + u_t \tag{9}$$

$$q_t = \tilde{a}_0 + \tilde{a}_1\, q_{t-1} + v_t \tag{10}$$

$$q_t = \hat{a}_1\, q_{t-1} + w_t \tag{11}$$

where T is the number of observations and u_t, v_t and w_t are error terms. Test statistics are then calculated based on modifications of the conventional t-statistics to allow for heterogeneity and interdependence of the error process. The critical values for the estimated coefficients are the same as those for the corresponding statistics estimated using the Dickey-Fuller approach.

Recent empirical work on real exchange rates has found that ρ is typically not far below unity—the null hypothesis that $\rho = 1$ usually cannot be rejected for short-sample periods at reasonable significance levels but can very often be rejected for long sample periods. The results of large-sample tests, together with the fact that the tests have low power when ρ is close to unity, make it reasonable to conclude that there is generally some mean reversion.[7] Tables 1 and 2 present the results of tests performed on the real exchange rates of Canada vs. the U.S., Canada vs. the U.K., and the U.K. vs. the U.S. using annual data spanning periods longer than 100 years. Tables 3 and 4 present the results of tests using monthly real exchange rates for the period since 1957 for Canada, France, Germany, Japan and the U.K. with respect to the U.S. and for France vs. Germany.

The Dickey-Fuller tests on logarithms of annual real exchange rates in Table 1 indicate that the null hypothesis of non-stationarity of the U.K./U.S. real exchange rate can be clearly rejected for the span of years 1803 to 2002. There is no evidence of any drift or trend—the F-statistics are significant only because the lagged real exchange rate is significant. One lag of the change in the real exchange rate appears to produce residuals closest to white-noise and thereby give the best estimate of ρ. In the Canada/U.S. case we can not reject the null hypothesis of a random walk for the entire span of data available, 1874 to 2002. For a slightly shorter period ending

[6] This procedure is discussed by Enders on pages 239 and 240 of the book previously cited.

[7] See, Kenneth Rogoff, "The Purchasing Power Parity Puzzle," *The Journal of Economic Literature*, Vol. 34, No. 2 (June), 1996, 647-668.

Table 1: Dickey-Fuller Test Results for Real Exchange Rates: Annual Data

Dependent Variable ΔY_t	Drift	Y_{t-1}	ΔY_{t-1}	ΔY_{t-2}	ΔY_{t-3}	Trend	F
U.K. / U.S. 1805–2002	1.62	-0.129**	0.149°°			-0.015	4.647*
	(1.41)	(-3.72)	(2.14)			(-1.50)	
	0.083	-0.105**	0.132°°				5.810**
	(0.158)	(-3.40)	(1.93)				
		-0.105***	0.132°°				
		(-3.41)	(1.93)				
Canada / U.S. 1874–2002	1.22	-0.100				-0.023	2.135
	(1.35)	(-2.19)				(-1.90)	
	-0.31	-0.058					1.366
	(-0.777)	(-1.437)					
		-.0589					
		(-1.461)					
Canada / U.S. 1874–1993	1.22	-0.200**				-0.0156	4.289*
	(1.37)	(-3.563)				(-1.267)	
	0.231	-0.179**					5.601**
	(0.545)	(-3.297)					
		-0.1716**					
		(-3.312)					

Continued on Next Page .....

Table 1: Continued

Dependent Variable ΔY_t	Drift	Y_{t-1}	ΔY_{t-1}	ΔY_{t-2}	ΔY_{t-3}	Trend	F
Canada / U.K. 1877–2002	2.18	-0.0846	0.247$^{\circ\circ\circ}$	0.120°	-0.244$^{\circ\circ\circ}$	-0.040	2.059
	(1.64)	(-2.01)	(2.82)	(1.33)	(-2.65)	(-2.14)	
	-0.388	-0.036	0.238$^{\circ\circ\circ}$	0.101	-0.027$^{\circ\circ\circ}$		0.777
	(-0.671)	(-1.003)	(2.698)	(1.11)	(-2.906)		
		-.0377	0.241$^{\circ\circ\circ}$	0.104	-0.264$^{\circ\circ\circ}$		
		(-1.054)	(2.743)	(1.148)	(2.869)		
Canada / U.K. 1877–1984	2.02	-0.184*	0.330$^{\circ\circ\circ}$	0.190$^{\circ\circ}$	-0.124	-0.0188	3.892*
	(1.573)	(-3.395)	(3.579)	(1.913)	(-1.188)	(-0.977)	
	0.946	-0.175**	0.328$^{\circ\circ\circ}$	0.188$^{\circ\circ}$	-0.124		5.364**
	(1.426)	(-3.274)	(3.567)	(1.890)	(-1.188)		
		-0.1397***	0.317$^{\circ\circ\circ}$	0.165$^{\circ\circ}$	-0.160		
		(-2.933)	(3.443)	(1.672)	(-1.577)		

Notes and Sources: All estimates use the logarithms of the relevant real exchange rates. The numbers in parentheses below the coefficients are the conventional t-statistics. The F statistic in the rightmost column tests the null hypothesis that the coefficients of the lagged level of the real exchange rate Y_{t-1} and any drift and trend terms included in the regression are simultaneously zero. The superscripts *, ** and *** indicate significance at the 10%, 5% and 1% levels, respectively, using the Dickey-Fuller tables and the superscripts $^{\circ}$, $^{\circ\circ}$ and $^{\circ\circ\circ}$ indicate significance at the 10%, 5% and 1% levels according to a standard t-test. For sources see Appendix A.

Table 2: Phillips-Perron Test Results for Real Exchange Rates: Annual Data

$$Y_t = a_0 + a_1 Y_{t-1} + a_2 (t - T/2) + u_t$$

$$Y_t = \tilde{a}_0 + \tilde{a}_1 Y_{t-1} + v_t \qquad Y_t = \hat{a}_1 Y_{t-1} + w_t$$

	$a_0 = 0$	$a_1 = 1$	$a_2 = 0$	$a_0 = 0$ & $a_1 = 1$	$\tilde{a}_1 = 1$	$\hat{a}_1 = 1$
U.K. / U.S. 1805–2002						
Lags = 1	0.087	-3.572**	-0.641	6.464*	-3.501***	-3.515***
Lags = 4	-0.088	-3.495**	-0.721	6.169*	-3.413***	-3.428***
Canada / U.S. 1874–2002						
Lags = 1	-0.727	-2.188	-1.907	2.856	-1.423	-1.472
Lags = 4	-0.685	-2.365	-1.671	2.943	-1.568	-1.636*
Canada / U.S. 1874–1993						
Lags = 1	0.664	-3.562**	-1.269	6.369**	-3.328**	-3.317***
Lags = 4	0.715	-3.703**	-0.994	6.724**	-3.449***	-3.343***
Canada / U.K. 1874–2002						
Lags = 1	0.600	-2.340	-1.453	2.698	-1.534	-1.565
Lags = 4	0.599	-2.344	-1.448	2.691	-1.525	-1.571
Canada / U.K. 1874–1984						
Lags = 1	1.068	-3.073	0.046	4.691	-3.021**	-2.898***
Lags = 4	1.127	-3.223*	0.307	5.172	-3.177**	-3.045***

Notes and Sources: All estimates use logarithms of the relevant real exchange rate series. The superscripts *, ** and *** indicate significance at the 10%, 5% and 1% levels, respectively, using the Dickey-Fuller tables which are also appropriate for the Phillips-Perron test. The statistics in all columns but the fourth from the left are t-based. For sources see Appendix A.

in 1993, however, the null hypothesis of a random walk can be rejected at the 5% level, again with no evidence of drift or trend. Adding lags of the change in the real exchange rate does not appear to improve the fit. The Canada/U.K. results are very similar to those for Canada vs. the United States with two exceptions—the null hypothesis of a random walk can be rejected for a slightly shorter span of years, to 1984 instead of 1993, and three lags of the change in the real exchange rate seem to give the best fit.

Phillips-Perron tests, shown in Table 2, yield the same results as the Dickey-Fuller tests with one exception. With 4 lags the null hypothesis of a random walk can be rejected at the 10% level for the Canada/U.S. real exchange rate for the whole time-span 1874 to 2002, contradicting the Dickey-Fuller test. Truncation lags of 1 and 4 were chosen according to the selections of two commercial econometrics programs, SHAZAM (1 lag) and RATS (4 lags). SHAZAM's default is to select the truncation lag as the highest significant lag order from either the autocorrelation function or the partial autocorrelation function of the first-differenced series. The basis for the RATS default truncation lag is not explained in the program's manual.

It is apparent that it is the behaviour of the Canadan real exchange rate during recent years that is leading to failure to reject the random-walk hypothesis. As can be seen from the top panel of Figure 3, the Canadian real exchange rate with respect to the U.S. has trended downward since the mid-1970's—there appears to have been a shift in the trend about 1974. Using techniques developed by Perron, we can test whether the series is stationary around a structural shift in trend against the null hypothesis that it is a random walk with a shift in the drift term.[8] We construct a trend dummy variable, D_T, equal to zero from 1874 to 1973 and to $(t-1973)$ from 1974 through to 2002, and regress the level of the real exchange rate on trend and the trend dummy for the whole period:

$$q_t \quad = \quad \beta_0 + \beta_1\, t + \beta_2 D_T + \epsilon_t \tag{12}$$

Then we test the residuals, e_t from the above regression for stationarity, fitting the equation

$$\Delta e_t \quad = \quad \rho\, e_{t-1} + \vartheta_t. \tag{13}$$

The resulting t-statistic for ρ is -3.65 and the Durbin-Watson statistic of 1.91 indicates that the residuals are not serially correlated. Given that the

[8] Pierre Perron, "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis," *Econometrica*, Vol. 57, No. 6 (November) 1989. Perron's method is discussed on pages 245–251 of the Enders book cited.

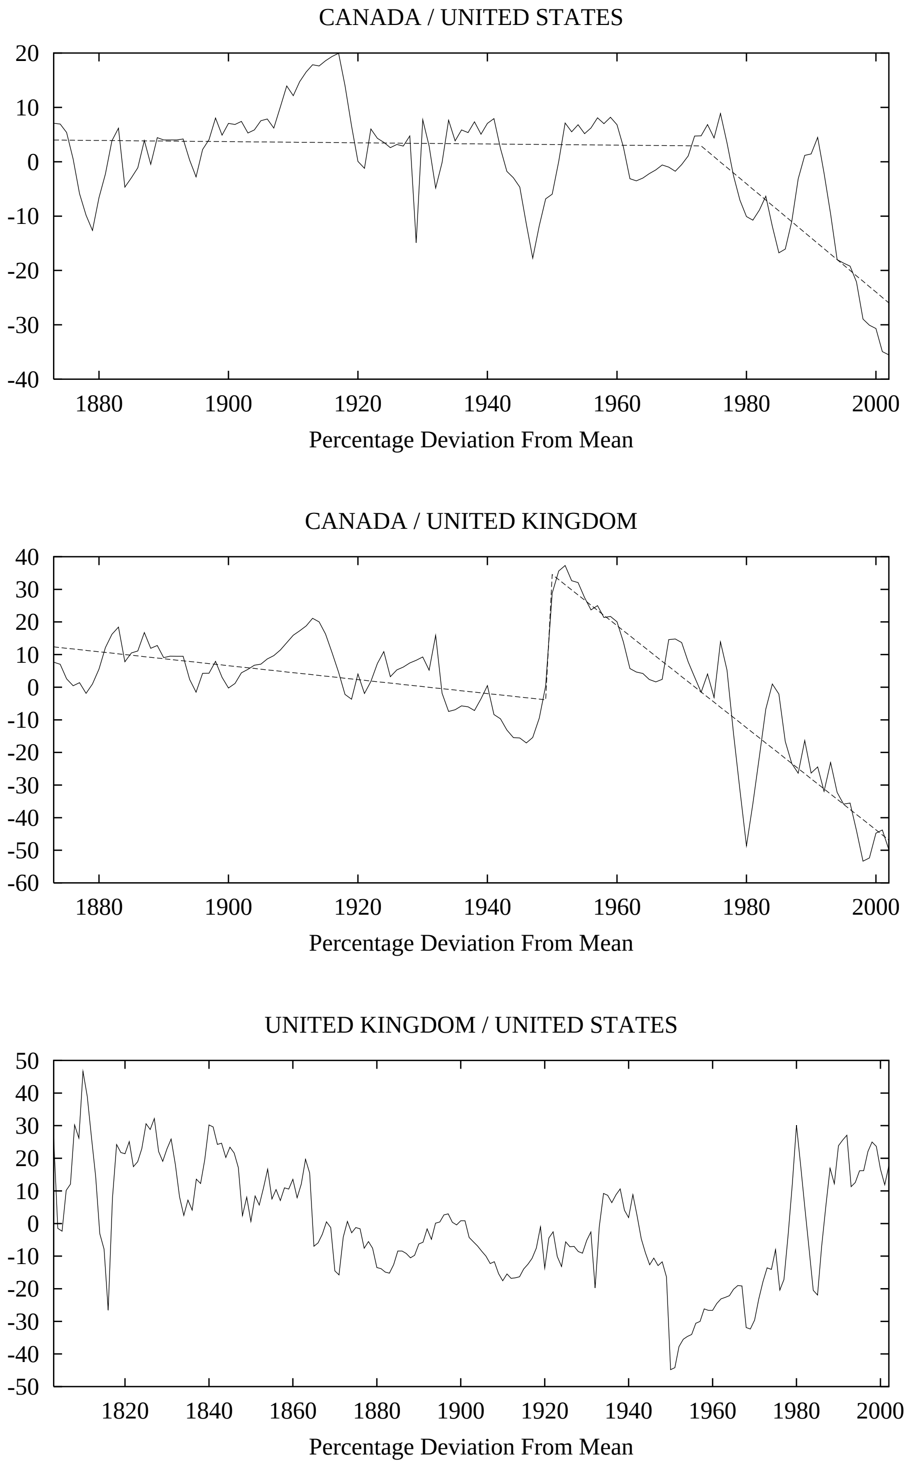

Figure 3: Real exchange rates over long periods: Canada viz. à viz. the U.S. and the U.K., 1873 to 2002 and U.K. viz. à viz. the U.S., 1803 to 2002. For sources see Appendix A:.

break in trend occurred a fraction 0.78 of the distance from the beginning to the end of the sample period, the 10% critical value for the t-statistic for ρ is about -3.52 in the relevant table in Perron's article.[9] We can therefore reject the null hypothesis of a random walk with a shift in drift in favour of stationarity around a breaking trend.

In the Canada/U.K. case, shown in the middle panel of Figure 3, there appears to have been an upward shift of the level of the series in 1950, following the 1949 devaluation of the pound, together with a change in trend after that year. To test, again using Perron's method, whether the Canada/U.K. real exchange rate is stationary around a trend that shifted in level and slope in 1950 as opposed to the null hypothesis of a random walk with a pulse shock in 1950 and a change in drift after that year, we construct a time-dummy, again called D_T, equal to zero from 1873 through 1949 and $(t - 1949)$ from 1950 onward, and a level dummy, D_L, equal to zero from 1873 through 1949 and unity from 1950 onward. We then fit the following two equations by ordinary-least-squares:

$$\begin{aligned} q_t &= \beta_0 + \beta_1\, t + \beta_2 D_L + \beta_3 D_T + \epsilon_t & (14) \\ \Delta e_t &= \rho\, e_{t-1} + \Delta e_{t-1} + \Delta e_{t-2} + \Delta e_{t-3} + \vartheta_t & (15) \end{aligned}$$

where three lags of Δe_t in the second equation are sufficient to eliminate serial correlation in the residuals. In this case the t-statistic for the estimate of ρ is -4.87 as compared with a 2.5% critical value in Perron's table of -4.26.[10] We can easily reject the null-hypothesis of a random walk.

There also appears to be a structural shift of level and trend in the U.K./U.S. real exchange rate series in the bottom panel of Figure 3, but, as indicated in Table 1, we could reject the null hypothesis of a random walk without taking it into account.

Tables 3 and 4 present Dicky-Fuller and Phillips-Perron test results for monthly real exchange rate data for the period 1957 through 2002 for Canada, the U.K., Japan, France and Germany vs. the United States and for France vs. Germany. As in the case of annual data, the logarithms of the real exchange rate series are used. In the cases of France/Germany, France/U.S., Germany/U.S. and U.K./U.S. the null hypothesis of a random walk can be rejected, at the 10% level or better, in favour of slow mean reversion with no trend. For Japan/U.S., though a positive trend is apparent in the bottom panel of Figure 1, a random walk can be rejected in favour of slow mean reversion with no trend at the 10% level in the Dickey-Fuller

[9] See page 1377.

[10] Here the ratio of time until the break to the length of the sample period is about 0.6.

Table 3: Dickey-Fuller Test Results for Real Exchange Rates:
Monthly Data, 1957–2002

Dependent Variable ΔY_t	Drift	Y_{t-1}	Trend	F_{Lags}	$F_{Y_{t-1}\ T}$	$F_{Y_{t-1}\ D}$	F_{All}
Canada / U.S.	0.210	-0.014	-0.001	3.319	3.098		2.405
	(1.746)	(-2.425)	(-2.213)				
	-0.039	-0.004		3.151		1.151	
	(-0.910)	(-1.14)					
		-0.005		3.305			
		(-1.214)					
				18 lags			
U.K. / U.S.	-0.496	-0.035**	0.002	5.005	6.008		4.078
	(-2.017)	(-3.464)	2.379				
	0.046	-0.018*		4.770		3.257	
	(0.502)	(-2.510)					
		-0.018**		4.787			
		(-2.504)					
				18 lags			
Japan / U.S.	-0.740	-0.015	0.003	7.446	2.937		2.328
	(-1.481)	(-2.243)	(1.748)				
	0.114	-0.004		7.143		1.957	
	(1.096)	(-1.676)					
		-0.004*		7.325			
		(-1.647)					
				11 lags			

Continued on Next Page .....

Table 3: Continued

Dependent Variable ΔY_t	Drift	Y_{t-1}	Trend	F_{Lags}	$F_{Y_{t-1}\ T}$	$F_{Y_{t-1}\ D}$	F_{All}
France / U.S.	-0.080	-0.020	0.0003	3.500	3.729		2.486
	(-0.390)	(-2.702)	(0.493)				
	0.009	-0.019**		3.493		3.613	
	(0.088)	(-2.688)					
		-0.019***		3.500			
		(-2.689)					
				lags = 17			
Germany / U.S.	-0.121	-0.014	0.001	5.545	3.090		2.099
	(-0.539)	(-2.415)	(0.788)				
	0.038	-0.012		5.491		2.840	
	(0.389)	(-2.359)					
		-0.012**		5.510			
		(-2.353)					
				lags = 10			
France / Germany	2.243	-0.035	-0.001	3.132	6.066		4.254
	(1.562)	(-3.393)	(-2.029)				
	-0.051	-0.019		2.975		4.297	
	(-0.921)	(-2.823)					
		-0.019***		2.996			
		(-2.783)					
				lags = 14			

Notes and Sources: All the real exchange rate series are expressed in logarithms. The numbers in the brackets below the coefficients are the conventional t-statistics. The subscripts of the F statistics indicate the variables whose coefficients are zero under the relevant null hypotheses, with *All* referring to lagged Y, Trend and Drift and *Lags* referring to lags of the dependent variable under augmented tests. The superscripts *, ** and *** have the same meaning as in Table 3. The lags, although not marked with superscripts, are all significant at the 1% level by conventional standards. For sources see Appendix A.

Table 4: Phillips-Perron Test Results for Real Exchange Rates: Monthly Data, 1957–2002

$$Y_t = a_0 + a_1 Y_{t-1} + a_2 (t - T/2) + u_t$$

$$Y_t = \tilde{a}_0 + \tilde{a}_1 Y_{t-1} + v_t \qquad Y_t = \hat{a}_1 Y_{t-1} + w_t$$

	$a_0 = 0$	$a_1 = 1$	$a_2 = 0$	$a_0 = 0$ & $a_1 = 1$	$\tilde{a}_1 = 1$	$\hat{a}_1 = 1$
Canada/U.S.						
Lags = 1	-1.735	-1.200	-1.224	0.954	-0.055	-0.070
Lags = 4	-1.683	-1.260	-1.113	0.919	-0.111	-0.140
U.K./U.S.						
Lags = 1	0.445	-2.782	0.934	3.317	-2.154	-2.158**
Lags = 4	0.412	-3.002	0.587	3.663	-2.309	-2.314**
Japan/U.S.						
Lags = 1	1.403	-1.140	-0.218	1.202	-1.484	-1.484
Lags = 4	1.281	-1.395	-0.525	1.376	-1.495	-1.500
France/U.S.						
Lags = 1	-0.134	-1.955	-0.018	1.963	-1.979	-1.983*
Lags = 4	-0.124	-2.138	-0.096	2.326	-2.151	-2.154*
Germany/U.S.						
Lags = 1	0.373	-1.618	-0.293	1.693	-1.828	-1.830*
Lags = 4	0.343	-1.793	-0.425	1.947	-1.943	-1.945*
France/Germany						
Lags = 1	-0.883	-3.290*	-1.620	5.707	-2.913**	-2.916***
Lags = 4	-0.860	-3.354*	-1.472	5.715	-2.930**	-2.935***

Notes and Sources: All the real exchange rate series are expressed in logarithms. The superscripts *, ** and *** indicate significance at the 10%, 5% and 1% levels, respectively, using the Dickey-Fuller tables which are also appropriate for the Phillips-Perron test. The statistics in all columns but the fourth from the left are t-based. For sources see Appendix A.

test but not in the Phillips-Perron test. In the Canada/U.S. case the null hypothesis of a random walk cannot be rejected. In the Dickey-Fuller tests appropriate lags of the changes in real exchange rates were selected on the basis of the AIC and BIC criteria. The lags for the Phillips-Perron tests were chosen, as in the case of the annual data, at one and four, based on the default choices by SHAZAM and RATS. The Perron structural change analysis was also applied to the monthly Canadian real exchange rate with respect to the U.S., with the trend shift occurring in January 1974, .37 of the distance through the sample period. The null hypothesis of non-stationarity with a change in drift after 1973 could not be rejected.

On the basis of the tests using annual and monthly data, we can reject the view that real exchange rates are random walks in every case examined except for Canada for the shorter period 1957–2002 and possibly also Japan for the same period. In both these cases there are clear trends, downward in the case of Canada and upward in the case of Japan. We have to conclude from all of this, as Rogoff has done in the article cited in footnote 7, that real exchange rates are not random walks. Rather, they are slowly mean-reverting series, the shocks to which have highly persistent effects. We must keep in mind here that all our tests have low power to reject the null of a random walk when it is not true. While we cannot reject the null-hypothesis of a random walk in the case of the monthly Canadian data, we clearly can do so in the case of annual data when we allow for structural shifts in trend. And if there is no random walk in annual data, there cannot be one in monthly data. Common sense must tell us that the observed downward trend of the Canadian real exchange rate with respect to the United States over the past thirty years is a temporary phenomenon—otherwise the real exchange rate will reach zero some forty years into the future. It makes no sense, in a stable world economy, for the value of a country's output to go to zero and that of its trading partners to become infinite! This argument also applies to the Japanese real exchange rate—the value of Japanese output is unlikely to eventually become infinite! A country's real exchange rate—that is, the per unit value of its output relative to that of its trading partners—will change through time in response to technological change affecting the traded and non-traded components of its output, reallocations of world investment between capital stock employed in the domestic economy and capital stock employed abroad, and changes in the international relative prices of goods produced in the domestic economy and abroad. Shocks may also result from improvements or deteriorations in political stability and the management of economic policy in the home economy relative to the rest of the world. In a stable world no country is going to have all the good, or bad, luck so it is unlikely that actual real exchange rates will wander forever in one or other direction with no tendency to mean revert. On the other hand, most technological and political developments have long lives, so it is reasonable to expect their effects to persist over long periods of time.

Tables 5 and 6 show the degree of mean reversion implied by the annual and monthly tests if we assume that the relevant real exchange rates are stationary. With respect to the annual data, if we ignore the time-spans contaminated by trend breaks it would appear that the half-life of a technological or other shock to the real exchange rate is somewhere between three and five years and the three-quarter life is between seven and twelve years. Apart from the extremes of Canada and Japan vs. the U.S., where adjustment is much slower, and France vs. Germany, where the adjustment is more rapid, the monthly data suggest more or less the same conclusion.

4 Forward Exchange Rates and Covered Interest Parity

The interest parity condition holds that the 1-month and 3-month forward premia on the domestic currency must equal the excess of foreign over domestic interest rates on securities maturing in one and three months respectively, adjusted for risk. That is, letting Π_t' represent the forward exchange rate and Π_t the spot rate, with exchange rates defined as prices of domestic currency in units of foreign currency,

$$i_t^* - i_t \;=\; (\Pi_t' - \Pi_t)/\Pi_t - \theta_t \;=\; \Phi_t - \theta_t \tag{16}$$

where Φ_t is the forward premium and θ_t is the risk premium required to get world asset holders to hold domestic assets under conditions where future changes in exchange rates are fully compensated for. This is what is called the 'country risk premium', as it depends on the security of investments in the two countries and not on movements in the exchange rate—it is the risk premium that would hold if the domestic and foreign economies were part of a single currency area. In the absence of such risk, arbitrage will ensure that the interest differential equals the forward premium—otherwise a sure profit could be obtained by shifting one's portfolio between domestic and foreign assets and purchasing forward exchange to neutralise the effects of any exchange rate changes that might occur over the maturity life of the assets.

In the case where the interest rate differential equals the forward premium, covered interest parity is said to hold. It will never hold exactly,

Table 5: Fraction of Total Response to a Real Exchange Rate Shock Remaining in the Twenty Subsequent Years: Annual Data

Year	U.K. /U.S. 1805–2002	Canada /U.S. 1874–2002	Canada /U.S. 1874–1993	Canada /U.K. 1874–2002	Canada /U.K. 1874–1984
0	1.000000	1.000000	1.000000	1.000000	1.000000
1	0.894484	0.941111	0.828419	0.962268	0.860341
2	0.800102	0.885689	0.686279	0.925960	0.740187
3	0.715679	0.833531	0.568527	0.891022	0.636813
4	0.640164	0.784445	0.470978	0.857402	0.547877
5	0.572617	0.738249	0.390168	0.825050	0.471361
6	0.512197	0.694774	0.323222	0.793919	0.405531
7	0.458152	0.653859	0.267764	0.763963	0.348895
8	0.409810	0.615354	0.221821	0.735137	0.300169
9	0.366569	0.579116	0.183760	0.707399	0.258248
10	0.327890	0.545012	0.152231	0.680708	0.222181
11	0.293292	0.512917	0.126111	0.655023	0.191152
12	0.262346	0.482711	0.104473	0.630308	0.164456
13	0.234664	0.454285	0.086547	0.606525	0.141488
14	0.209903	0.427532	0.071697	0.583640	0.121728
15	0.187755	0.402355	0.059395	0.561618	0.104728
16	0.167944	0.378661	0.049204	0.540427	0.090101
17	0.150223	0.356361	0.040762	0.520036	0.077518
18	0.134373	0.335375	0.033768	0.500414	0.066692
19	0.120194	0.315625	0.027974	0.481532	0.057378
20	0.107512	0.297038	0.023174	0.463363	0.049364

Notes and Sources: Theses statistics are based on the Dickey-Fuller test results in Table 1.

Table 6: Fraction of Total Response to a Real Exchange Rate Shock Remaining in the Twenty Subsequent Years: Monthly Data

Year	Canada /U.S.	Japan /U.S.	U.K. /U.S.	France /U.S.	Germany /U.S.	France /Germany
0	1.000000	1.000000	1.000000	1.000000	1.000000	1.000000
1	0.930621	0.942395	0.768508	0.863465	0.907301	0.667389
2	0.866056	0.888109	0.590605	0.745571	0.823195	0.445409
3	0.805970	0.836950	0.453885	0.643774	0.746885	0.297261
4	0.750053	0.788737	0.348815	0.555876	0.677650	0.198389
5	0.698015	0.743303	0.268067	0.479980	0.614832	0.132403
6	0.649587	0.700485	0.206012	0.414445	0.557838	0.088364
7	0.604520	0.660134	0.158322	0.357859	0.506126	0.058973
8	0.562579	0.622107	0.121672	0.308999	0.459209	0.039358
9	0.523548	0.586271	0.093506	0.266809	0.416641	0.026267
10	0.487225	0.552499	0.071860	0.230380	0.378018	0.017530
11	0.453422	0.520672	0.055225	0.198925	0.342976	0.011610
12	0.421964	0.490679	0.042441	0.171765	0.311183	0.007808
13	0.392688	0.462414	0.032616	0.148313	0.282336	0.005211
14	0.365444	0.435776	0.025066	0.128063	0.256164	0.003478
15	0.340090	0.410674	0.019263	0.110578	0.232418	0.002321
16	0.316495	0.387017	0.014804	0.095480	0.210873	0.001549
17	0.294537	0.364723	0.011377	0.082444	0.191325	0.001034
18	0.274102	0.343713	0.087433	0.071187	0.173589	0.000700
19	0.255085	0.323914	0.067193	0.061468	0.157498	0.000460
20	0.237388	0.305255	0.051638	0.053075	0.142898	0.000307

Notes and Sources: Theses statistics are based on the coefficients of the lagged level of the real exchange rate generated by the Phillips-Perron tests in Table 4. The coefficients were taken from the regressions containing trend and constant terms.

of course, because the risk premium will never be zero, although it will be virtually zero on assets issued by the same company in the two currencies or on assets issued in the two currencies by institutions in third countries in the off-shore market.[11]

The evidence suggests that spot and forward exchange rates move so closely together that they can hardly be distinguished from each other on a plot. This is illustrated for spot and 90-day forward rates for the Canadian dollar in terms of U.S. dollars in the top panel of Figure 4. And, as the bottom two panels indicate, covered interest parity seems to hold approximately in a comparison of the 1-month and 3-month forward premia on the Canadian dollar in terms of the U.S. dollar with the respective interest rate differentials on 1-month and 3-month corporate paper. Despite a rather close fit overall, however, there are some clear and substantial deviations from covered interest parity in certain years after 1995.

It turns out that these deviations are the result of problems with the collection of spot and forward exchange rate data, as is illustrated by the case of the Japanese yen with respect to the U.S. dollar in recent years in Figure 5. Two different monthly estimates of the spot and forward rates were obtained from *Datastream* for 1999 through 2002—the mnemonics for the series are given below the charts in the top two panels. While the spot and forward rates are very similar in each of the two alternative estimates, the resulting 1-month forward premia on the yen in terms of the dollar implied by the estimates, expressed in annual percentage rates, are strikingly different as shown in the bottom panel. There are two reasons for this. First, even slight differences between spot and forward rates have big effects on the forward premia expressed in annual percentage rates. Second, it makes a difference whether the spot and forward exchange rate data pertain to prices asked, prices offered or actual contract prices, and whether the group of transactions that are averaged and the time interval over which they are averaged to obtain noon or closing prices for any given day is large or small.[12] These problems arise in the data for recent years with respect to all the currencies examined here. Indeed, as noted previously

[11]Even in these cases there will be some risk because, although the institution on which repayment depends is the same for both assets, or the assets are liabilities of institutions in third countries, future government intervention could still prevent repayment in one of the currencies.

[12]I would like to thank Alex Maynard for discussions of these issues. For elaboration, see A. Maynard and P.C.B. Phillips, "Rethinking An Old Empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly," *Journal of Applied Econometrics*, Vol. 16, No. 6, 2001, 677-680.

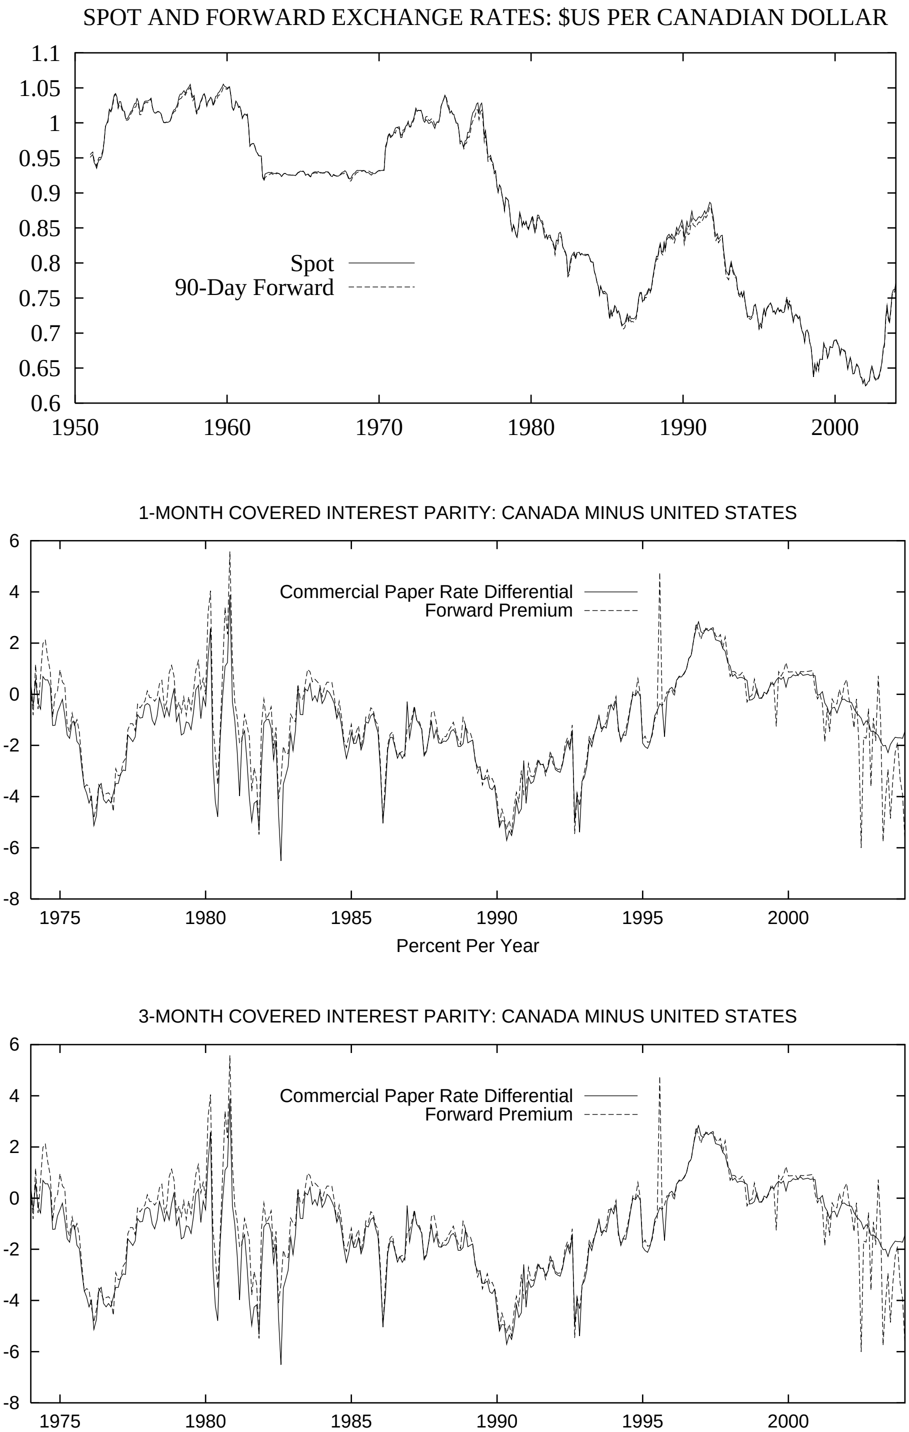

Figure 4: Canada vs. United States: Spot and forward exchange rates, U.S. dollars per Canadian dollar (top panel), 1-Month covered interest parity (middle panel) and 3-month covered interest parity (bottom panel). Source: *Cansim*.

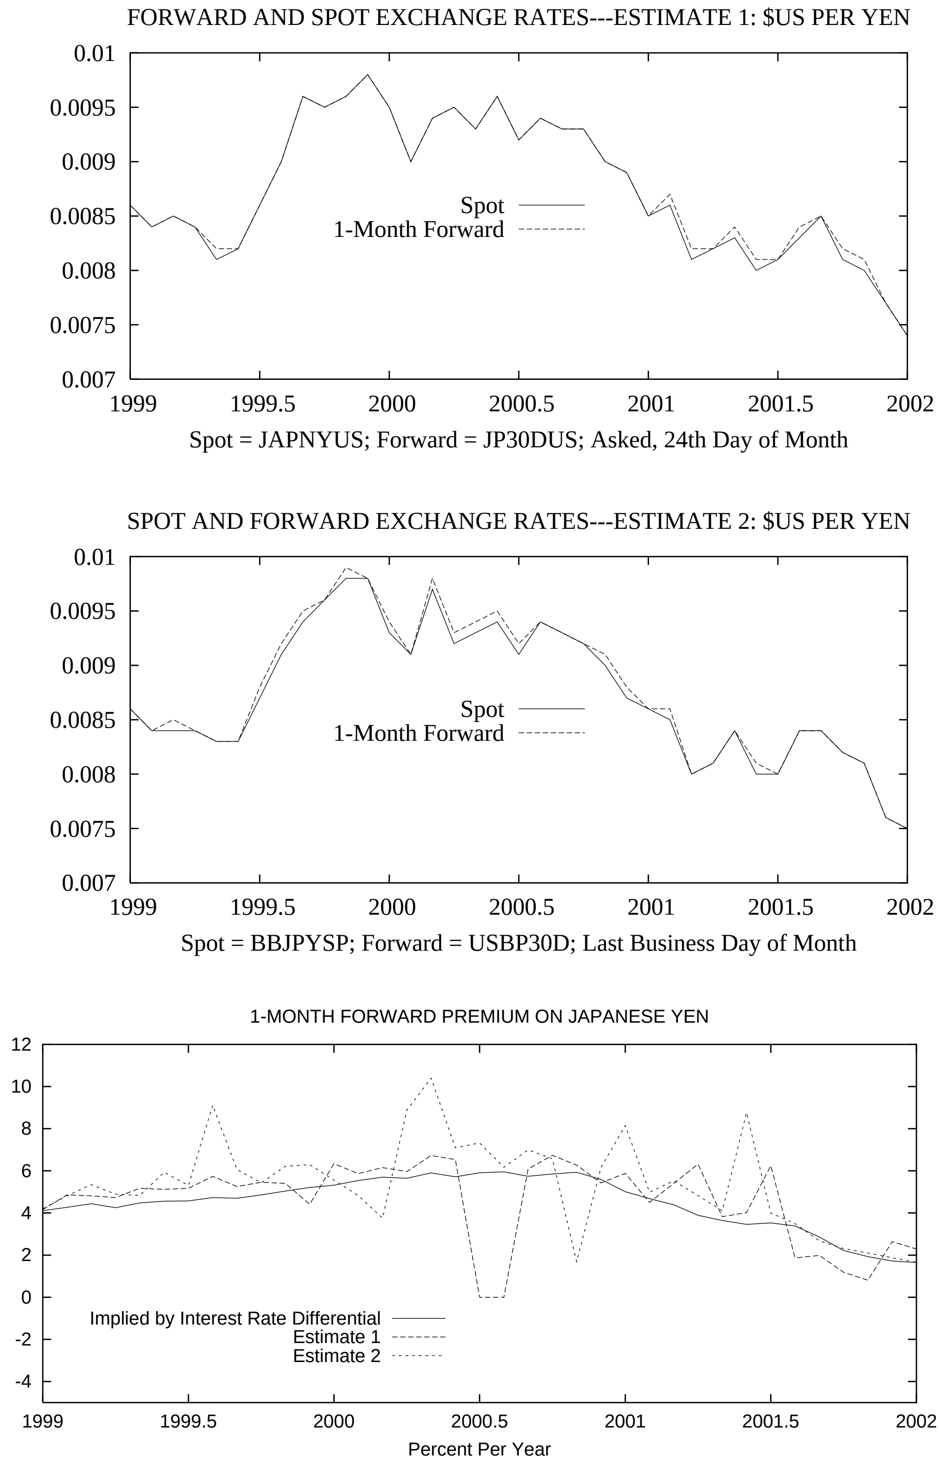

Figure 5: Alternative *Datastream* estimates of the Japanese Spot and 1-month forward exchange rates with respect to the U.S. Dollar, and the corresponding forward premia on the yen. Source: *Datastream*.

with respect to the bottom two panels of Figure 4, the problems also arise in the Canadian exchange rate data which were collected by *Cansim* and not by *Datastream*.[13] For these reasons the 'implicit' forward premia implied by the interest differentials will be used in subsequent empirical analysis along with, and sometimes instead of, the forward premia calculated from the relevant spot and forward exchange rates. The measures of the forward premia implied by interest rate differentials clearly seem superior in the Canadian case shown in Figure 4, and the case is even stronger for the other currencies in terms of the U.S. dollar in that we are able to use off-shore interest rates to calculate the implicit forward premia and thereby minimise the effects of country risk differences.

5 Inflation Differentials and Forward Premia

Foreign exchange market efficiency, as implied by rational use of all available information by investors, implies that

$$\Phi_t \;=\; E_t\{(\Pi_{t+1} - \Pi_t)/\Pi_t\} - \phi_t \;=\; E_\Pi - \phi_t \tag{17}$$

where $E_\Pi = E_t\{(\Pi_{t+1} - \Pi_t)/\Pi_t\}$ is the expected rate of change in the spot exchange rate between this period and next and ϕ_t is a foreign exchange risk premium on the domestic currency. Otherwise, agents could make an expected profit by selling one of the currencies short and purchasing it spot on the delivery date to cover the contract (or, what is the same thing, by purchasing the other currency forward and selling it at the spot rate on delivery). It follows from (1) that

$$E_\Pi \;=\; E_Q + E_{P^*/P} \;=\; E_Q + E_{P^*} - E_P \tag{18}$$

where E_Q is the expected rate of change in the real exchange rate from this period to next and $E_{P^*/P}$ is the expected rate of change in the foreign relative to the domestic price level—that is, the expected rate of foreign inflation, E_{P^*}, minus the expected rate of domestic inflation, E_P. Substitution of (18) into (17) yields, ignoring the time subscript,

$$\Phi \;=\; E_Q + E_{P^*} - E_P - \phi. \tag{19}$$

In addition to the foreign exchange risk premium, the forward premium will depend on the expected rate of change in the real exchange rate and the

[13]For a complete discussion of the data sources, see Appendix A.

expected foreign/domestic inflation rate differential. If the real exchange rate is a random walk and investors cannot forecast the shocks to it, E_Q will be zero. As noted above, however, although the real exchange rate is not a random walk shocks to it are very persistent with a slow rate of mean reversion. We can therefore expect that, unless investors can forecast the underlying shocks, E_Q will tend to be very slightly negative when the real exchange rate is above its long-run average level and positive and relatively small when it is below that level. While the expected rate of inflation, like the expected change in the real exchange rate, is unobserved it is reasonable that investors will anticipate continuing inflation during inflationary periods, so there should be an observed relationship between the difference between foreign and domestic inflation rates and the forward premium.

The forward premia on domestic currencies in terms of the U.S. dollar and the excess of the U.S. minus domestic inflation rates for Canada, France, Germany, the U.K. and Japan are plotted in Figures 6 and 7. As the top panel of Figure 6 illustrates for the U.S. minus France, month-over-month inflation rate differentials are much more variable than year-over-year inflation rate differentials. Accordingly, year-over-year differentials are used in all plots against forward premia. As can be seen from the figures, there is a loose correspondence between the inflation rate differentials and the forward premia. On average, as shown in Table 7, the inflation rate differentials and the forward premia have the same signs but their magnitudes tend to diverge by more than one percentage point per annum in the case of Canada, France and Japan. There are three potential reasons for this divergence. First, there may be differences between the actual and expected domestic minus U.S. inflation rates—greater expected inflation in Canada and France than actually occurred, relative to U.S. inflation, and less expected inflation in Japan relative the the U.S. than actually occurred. Second, there may have been non-zero expectations as to the direction of future movements in the real exchange rates, downward in Canada and France, and upward in the case of Japan. Third, there may have been negative foreign exchange risk premia on the Canadian dollar and the French franc relative to the U.S. dollar and a positive risk premium on the yen. The table also indicates another important fact—that the spot exchange rates fluctuate much more widely than the forward premia, the ratio of their standard deviations being in the neighbourhood of ten to one.

As indicated by the regression results presented in Table 8, year-over-year inflation differentials, here calculated as domestic minus U.S., have a small but statistically significant effect in the right direction on the forward premium—an increase in the domestic inflation rate relative to the U.S. infla-

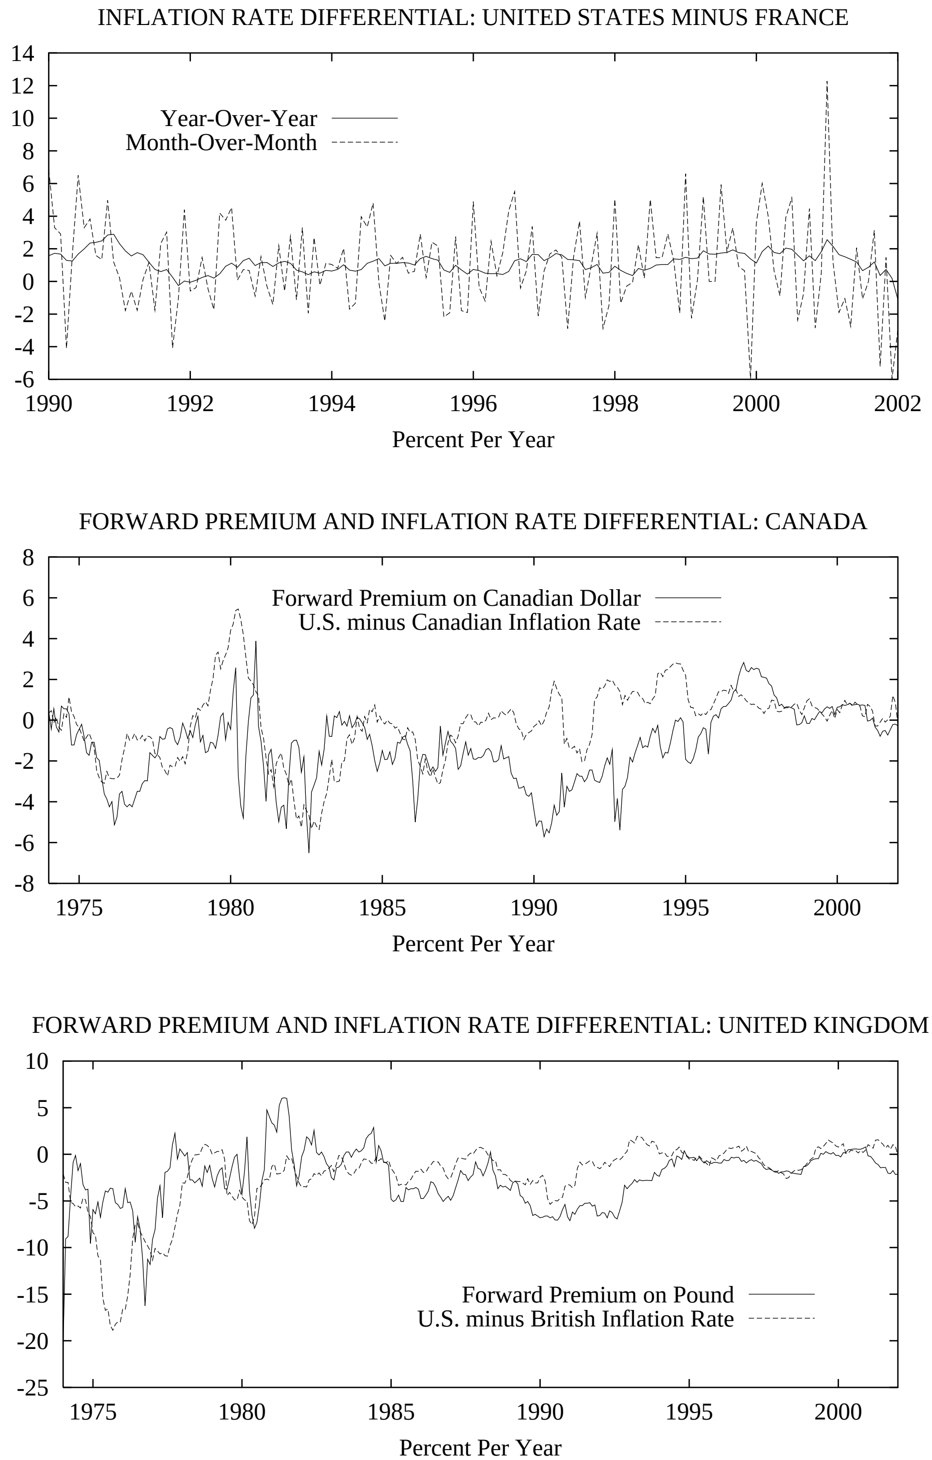

Figure 6: Month-over-month vs. year-over-year inflation differential, U.S. minus France (top panel) and year-over-year inflation differentials vs. the 1-month forward premium on domestic currency, Canada and U.K. vs. the U.S. The forward premia estimates are those implied by 1-month corporate paper rates for Canada/U.S. and 1-month off-shore rates for U.K./U.S. Sources: *Reuters*, *Datastream* and International Monetary Fund *International Financial Statistics*.

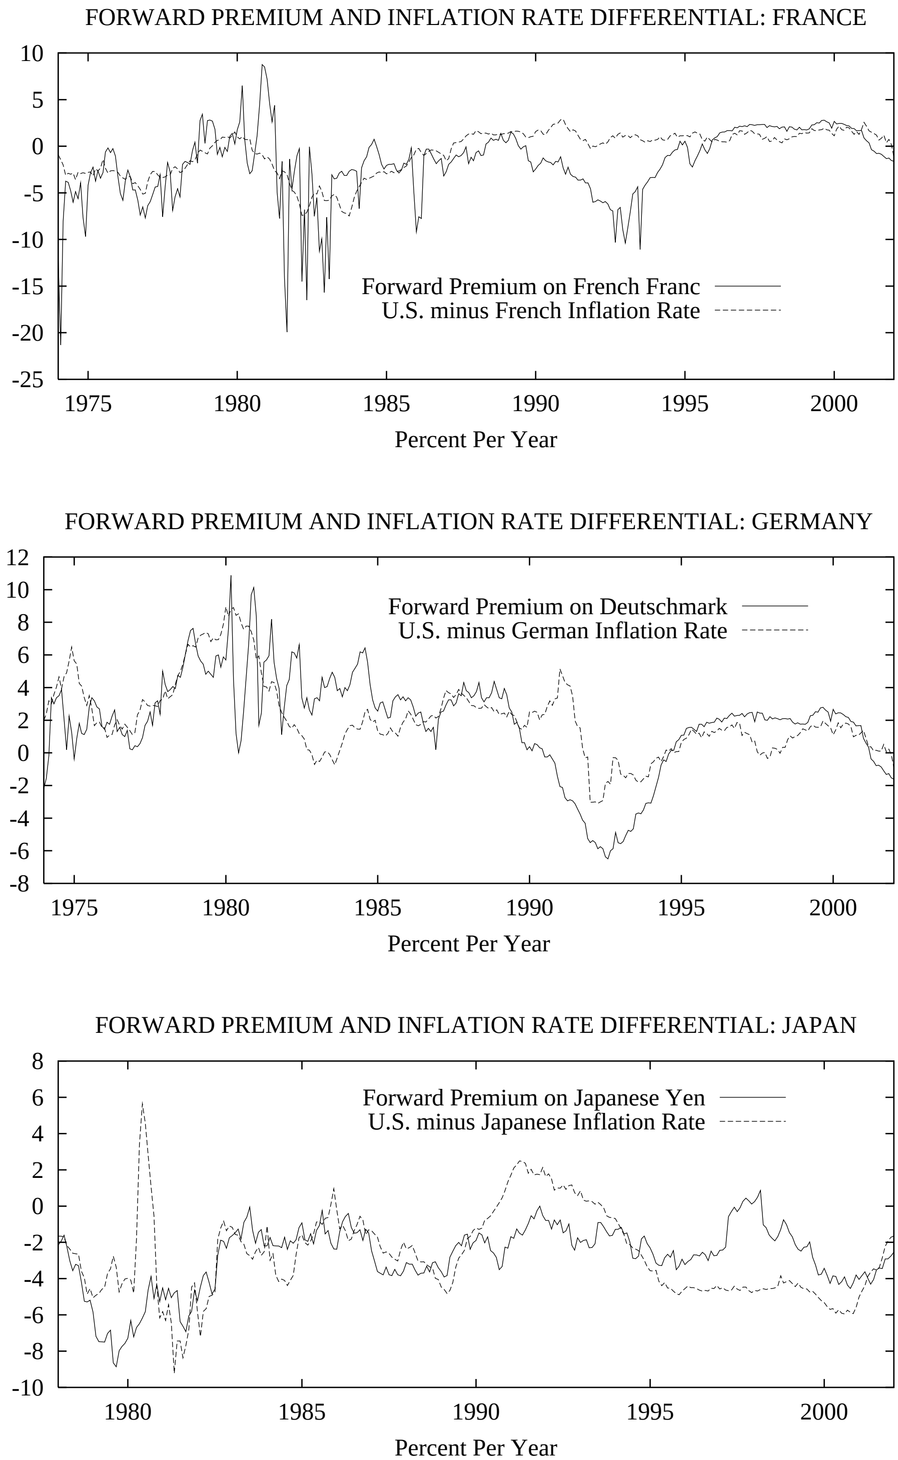

Figure 7: One-Month forward premium on the domestic currency in terms of U.S. dollars vs. the U.S. minus domestic year-over-year inflation differential: France, Germany and Japan. The forward premia are estimated from 1-month offshore interest rate differentials. Sources: *Reuters*, *Datastream* and International Monetary Fund *International Financial Statistics*.

Table 7: Inflation Rates, Spot Exchange Rate Variability and Forward Premia: Canada, France, Germany, the U.K. and Japan vs. the United States

Period Averages	Forward Premium on Domestic Currency	U.S. Minus Domestic Inflation Rate	Standard Deviation of % Change in Spot Rate	Standard Deviation of Forward Premium
1974–2003				
Canada	-1.283	-0.189	13.17	1.79
U.K.	-2.527	-2.122	36.97	3.07
Japan	3.244	1.840	42.30	4.15
1974–1998				
France	-2.198	-0.758	38.55	3.99
Germany	1.938	2.096	39.13	3.11

Notes and Sources: The inflation rate differences and forward premia are expressed as percent per annum. Except for Canada and Japan the forward premia are estimated from interest rate differentials on off-shore 1-month securities. The Canadian forward premium is estimated from interest rates on 1-month corporate paper and the Japanese forward premium is calculated directly from 1-month forward exchange rates. For sources see Appendix A.

Table 8: Regressions of the Forward Premium on the Domestic minus U.S. Inflation Rate Difference and the Change in the Spot Exchange Rate from the Previous Period: Canada, France, Germany, Britain and Japan vs. the United States

	Canada 1974–2002		France 1974–2001		Germany 1974–2001	
Premium Based on	Exch. Rates	Interest Rates	Exch. Rates	Interest Rates	Exch. Rates	Interest Rates
Constant	-0.927	-1.247	-1.292	-1.394	0.331	0.281
	(-4.58)	(-6.16)	(-3.50)	(-3.75)	(0.662)	(0.584)
Inflation Difference	-0.294	-0.299	-0.695	-0.760	-0.802	-0.757
	(-3.42)	(-3.30)	(-5.16)	(-5.54)	(-5.29)	(-5.10)
% Δ in Spot Rate	-0.006	-0.011	-0.018	-0.018	-0.006	-0.007
	(-0.85)	(-1.67)	(-3.10)	(-2.980)	(-1.34)	(-1.48)
NOBS	348	348	336	336	336	336
R-Square	.085	.103	.219	.235	.348	.352
Standard Error	1.77	1.70	3.35	3.47	2.57	2.41
Durbin-Watson	.372	.245	.654	.622	.252	.171

Continued on Next Page ....

Table 8: Continued

	U.K. 1974–2002		Japan 1974–2002	Japan Feb. 1977–2002	
Premium Based on	Exch. Rates	Interest Rates	Exch. Rates	Exch. Rates	Interest Rates
Constant	-1.717 (-5.48)	-1.809 (-5.80)	2.242 (6.813)	1.827 (3.205)	1.433 (2.892)
Inflation Difference	-0.314 (-3.66)	-0.316 (-3.77)	-0.534 (-5.71)	-0.712 (-3.69)	-0.468 (-2.86)
% Δ in Spot Rate	-0.013 (-2.07)	-0.011 (-1.80)	-0.013 (-2.48)	-0.009 (-2.17)	-0.009 (-3.11)
Lag 1				-0.007 (-1.55)	-0.009 (-2.58)
Lag 2				-0.005 (-1.36)	-0.007 (-2.16)
NOBS	348	348	348	311	311
R-Square	.172	.169	.227	.243	.251
Standard Error	2.82	2.81	3.68	2.78	2.08
Durbin-Watson	.330	.291	1.59	.387	.139

Notes and Sources: The dependent variables are the one-month forward discount of the domestic currency in terms of the U.S. dollar, calculated both directly and using the interest rate differential combined with the assumption that the country-risk premium is constant. The percentage changes in exchange rates are month-to-month at annual rates and the inflation rate differential, domestic minus U.S., is year-over-year. The figures in brackets are t-ratios based on heteroscedasticity and serial correlation adjusted standard errors of the coefficients with 5 lags. For sources see Appendix A.

tion rate leads to a decline in the forward premium on the domestic currency in all cases. Also, a decline in the forward premium is associated with a rise (appreciation) of the spot exchange rate in all cases although for Canada and Germany the relationship is not statistically significant at the 5% level in a two-tailed test. This is consistent with mean reversion in response to nominal exchange rate shocks—in all cases the coefficients are tiny, as would be implied by very slow rates of mean reversion. In the period from 1977 to 2002, two lags of the spot exchange rate shocks are also statistically significant, with negative signs, for Japan when the forward premium measure based on off-shore interest rates is used. In view of the significant serial correlation in the residuals in all of the regressions in Table 8, heteroskedasticity and autocorrelation consistent estimators of the standard errors of the regression coefficients were used with the truncation lag parameter set at 5 following the guideline suggested by Stock and Watson in their introductory econometrics text referred to in footnote 1.[14]

6 Errors From Exchange Rate Forecasts Based on Current Spot and Forward Rates

An important implication of the near random walk character of real exchange rate movements is that when shocks to the real exchange rate are unobserved by agents, and unpredictable, the best forecast of next period's real exchange rate tends to be the level of the real exchange rate this period. Indeed, the evidence is that simple random-walk based forecasts are superior to forecasts from sophisticated models.[15] It follows that this period's nominal exchange rate, adjusted to account for any obviously continuing differences between the domestic and the foreign inflation rates, generally provides the best forecast of next period's nominal exchange rate, even in the presence of a slight degree of mean reversion.

Figures 8 and 9 present plots of the distributions of forecast errors—actual minus predicted spot exchange rates—from naive current spot rate

[14]See Chapter 12, page 505, for the guideline with respect to the truncation lag. Stock and Watson incorporate HAC estimation of coefficient standard-errors throughout their basic discussion of regression analysis in Part 2.

[15]The classic study here is by Richard Meese and Kenneth Rogoff, "Empirical Exchange Rate Models of the Seventies: Do They Fit Out-of Sample," *Journal of International Economics*, 14, 1-2 (February), 1983, 3–24. Research in the past two decades has not diminished the force of their conclusions. Situations occur, of course, where one or more of a large group of alternative forecast techniques will do better, but those alternatives will not do systematically or predictably better than the random walk forecast.

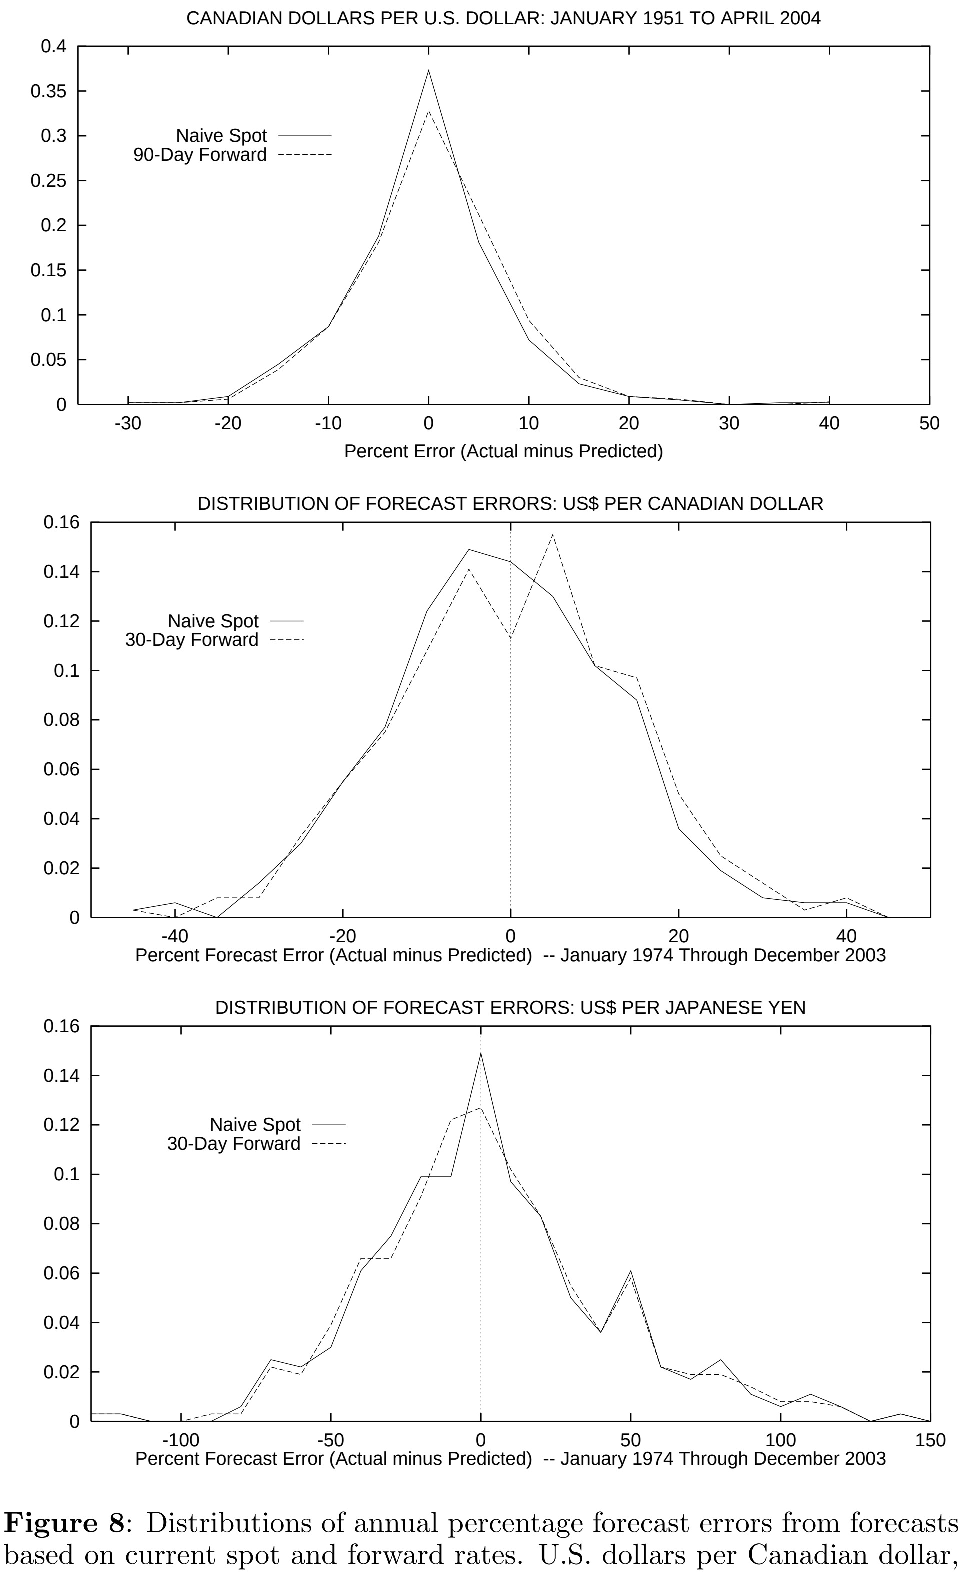

Figure 8: Distributions of annual percentage forecast errors from forecasts based on current spot and forward rates. U.S. dollars per Canadian dollar, 3-month forecasts for January 1951 through April 2004 and 1-month forecasts for 1974–2003, and U.S. dollars per Japanese yen, 1-month forecasts for 1974–2003. Sources: *Cansim* and *Datastream* for Canadian data and *Reuters* and *Datastream* for U.K. data.

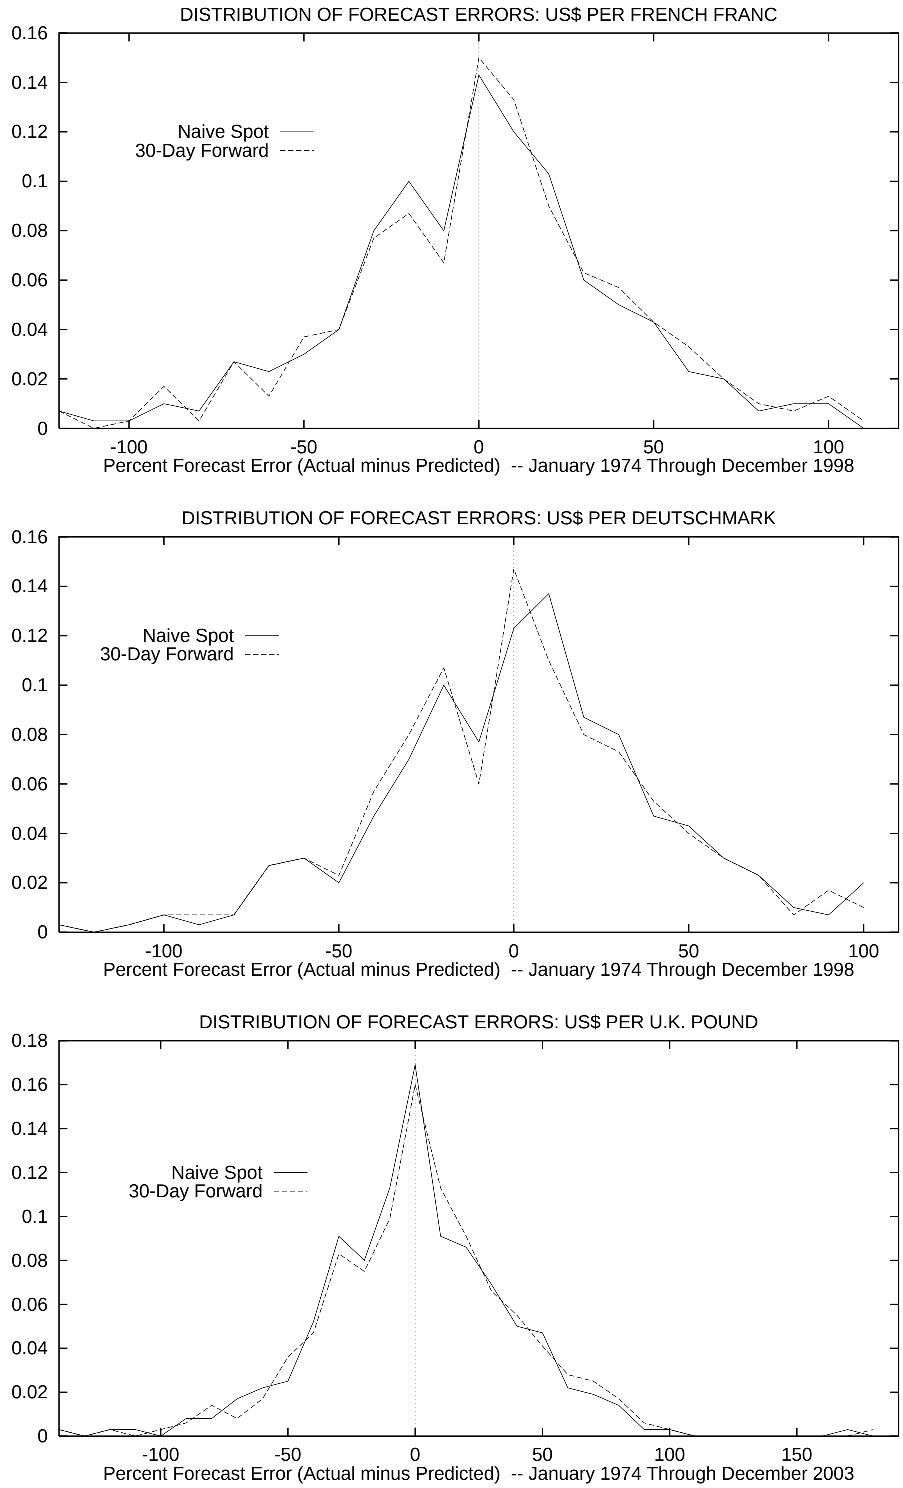

Figure 9: Distribution of annual percentage forecast errors from forecasts based on current spot and forward rates. U.S. dollars per French franc and Deutschmark, 1-month forecasts for 1974–1998, and U.S. dollars per U.K. pound, 1-month forecasts for 1974–2003. Source: *Cansim* and *Datastream* for Canadian data and *Reuters* and *Datastream* for U.K. data.

forecasts of next period's spot rate, with no adjustment for inflation differences, along with the distributions of forecast errors resulting from using the current forward rate as the forecast of next period's spot rate. The curves in the figures simply connect the mid-points of histograms—more sophisticated kernel-density estimates produce smoother curves that tell essentially the same story. The distributions of forecast errors from the two forecast methods seem quite similar.

A more detailed analysis of the forecast errors is shown in Table 9. In all cases except Germany and Japan the naive spot rate forecast error averages less than 1 percentage point on an annual basis. For Germany the percentage error is about 2.5 percent and for Japan it is around 4 percent. In every case but Germany and Japan, the naive spot based forecast errors are negative and the forward rate based forecast errors positive. In Germany and Japan, the fact that the forward rate predicted, on average, appreciations of the home currencies in terms of the dollar is consistent with the lower average rates of inflation in those countries than in the U.S. indicated in Table 7. In France and the U.K., and to a lesser extent in Canada, the forward rate predicted, on average, falls in the spot rate in excess of those that actually occurred.

As indicated in the third column from the right in Table 9, the current spot rate and forward rate based forecast errors were highly correlated, reflecting the fact, indicated in the right-most two columns of Table 7, that spot rates are much more variable, period to period, than the corresponding forward premia. Because of the very high correlation of the percentage forecast errors from naive spot and forward rate based predictions, we can treat the differences as paired differences and thereby measure the standard deviation of the mean differences as the square root of the variance of the difference divided by the sample size. On this basis, as shown in the middle columns of Table 9, we are able in all cases to reject at reasonable significance levels the null hypothesis of zero mean difference.

Finally, as the right-most two columns in Table 9 indicate, the root-mean-squared-errors of forecasts based on the current forward exchange rate are everywhere bigger than those of forecasts based on the naive assumption that next period's spot rate will be the same as the current spot rate. The differences, however, are extremely small.

Table 9: Errors at Annual Percentage Rates From Forecasts of Future Spot Exchange Rates Based on Current Spot and Forward Rates: Canada, France, Germany, the U.K. and Japan vs. the U.S.

	Mean		Spot minus Forward		Correl-ation Spot and Forward	Root Mean Squared Error	
	Spot	Forward	Mean	Std. Dev. of Mean		Spot	Forward
1951–2004 (ends April)							
Canada	-0.358	0.360	-0.718	0.056	.983	7.49	7.74
1974–2003							
Canada	-0.830	0.469	-1.296	0.093	.993	13.98	14.32
U.K.	-0.333	2.205	-2.537	0.160	.997	36.82	37.37
Japan	3.917	0.990	2.927	0.241	.994	41.94	42.00
1974–1998							
France	-0.080	2.130	-2.210	0.230	.995	38.48	38.96
Germany	2.575	0.651	1.924	0.180	.997	39.15	39.39

Notes and Sources: All forecast errors are expressed as percent per annum. For Canada in the period January 1951 to April 2004 the forecasts are 3-month. All the remaining forecasts are 1-month. Except for Japan, all forward rate based forecasts use domestic/U.S. interest rate differentials to estimate the forward premia. Off-shore rates are used in all cases but Canada, where corporate paper rates are used. The forward rate forecast of the Japanese Yen is obtained directly from the forward exchange rate. For sources see Appendix A.

7 Uncovered Interest Parity

The interest parity condition (16) can be combined with the efficient markets condition (17) to yield

$$i_t^* - i_t \;=\; \Phi_t - \theta_t \;=\; E_\Pi - \phi_t - \theta_t. \tag{20}$$

If agents are risk-neutral, so that the risk premia disappear, this reduces to

$$i_t^* - i_t \;=\; \Phi_t \;=\; E_\Pi \tag{21}$$

which is known as the condition of uncovered interest parity—the foreign/ domestic interest rate differential equals the expected rate of appreciation of the domestic currency. It also equals the forward premium since, if uncovered interest parity holds, so must covered interest parity.

If agents form their expectations rationally, in the sense that they take into account all information available to them, and on average correctly anticipate future exchange rate movements, the expected rate of change in the exchange rate will equal the mean of the actual rate of change—prediction errors in the upward and downward directions will be equally likely. Forward exchange rates will be unbiased predictors of future spot rates. Under these conditions, letting s_t and f_t be the logarithms of the spot and forward exchange rates, the regression

$$s_{t+1} \;=\; \alpha + \beta f_t + \epsilon_t \tag{22}$$

should produce estimates of α equal to zero and β equal to unity. If there is a constant risk premium, the estimate of α will differ from zero but that of β will still be unity. Alternatively, the change in the logarithm of the spot rate can be expressed as

$$s_{t+1} - s_t \;=\; \alpha + \beta\,(f_t - s_t) + \epsilon_t \tag{23}$$

which is identical to (22) when $\beta = 1.0$. The term $s_{t+1} - s_t$ is proportional to the percentage rate of change in the spot exchange rate and $(f_t - s_t)$ is proportional to the forward premium. Using (21) it can be seen that $i_t^* - i_t = f_t - s_t$. Equation (22) can be defined as the 'forward rate' version of the 'unbiasedness hypothesis' and equation (23) as the 'forward premium' version.[16] The unbiasedness hypothesis implies uncovered interest parity.

[16] When the exchange rate is defined as the domestic currency price of foreign currency, the term 'forward discount version' should be used.

The unbiasedness hypothesis can be tested by estimating (22) and (23) using ordinary least squares. The results are shown in Tables 10 and 11. As indicated in Table 10, the null hypothesis that $\beta = 1$ in (22) can be rejected at the 5% level in one case, U.K./U.S., and at the 10% level in the cases of France and Japan vs. the United States. For Canada and Germany vs. the U.S. the null cannot be rejected. When 3-month forward rates are used in the Canada/U.S. case it is evident from the Durbin-Watson statistics that there is serial correlation in the residuals. To compensate for this it was necessary to use heteroskedasticity and autocorrelation consistent estimates of the standard-errors of the coefficients. Overall, we have to include that forward-rate unbiasedness occurs to a reasonable approximation.

With respect to forward-premium unbiasedness, the results in Table 11 clearly indicate rejection of the null that $\beta = 1$. Except for the Canada/U.S. case with 3-month forward premia over the period 1951–1973, and U.K./U.S. with 1-month forward premia for the period from August 1990 to the end of 2003, the signs of the coefficients of the lagged forward premia are everywhere negative, implying that β is not only less than unity, but negative. For the period 1974–2003 these negative coefficients are statistically significant only for Canada and the U.K. vs. the United States. For the sub-period January 1978 to July 1990, however, the negative coefficients are statistically significant for Canada, Germany, the U.K. and Japan vs. the U.S., while for the sub-period August 1990 to December 2003 (December 2001 in the case of Germany) the negative coefficient is statistically significant only for Canada vs. the U.S. and Japan vs. the U.S., and then only at the 10% level, in cases where interest rate differentials are used to measure the forward premium. The highly significant negative relationship between the percentage change in the spot rate and that predicted by the forward premium results almost exclusively from the patterns in the data during the period January 1978 to July 1990. As can be seen from Figures 1 and 2, this was a period during which the U.S. dollar temporarily appreciated very substantially in terms of the other currencies and then fell back near to its beginning level. It was also a period of declines in all countries' inflation rates. Note that the R^2 values in all of the above regressions are extremely small.

The fact that the estimated values of β in equation (23) are significantly below unity and negative has come to be known as the 'forward premium anomaly' or 'forward premium puzzle'. A large econometric literature has grown up attempting to explain this puzzle. Most of this literature is referenced in the paper by Maynard and Phillips cited in footnote 12 and in

Table 10: Regressions of Spot Exchange Rates on Lagged 1-Month and 3-Month Forward Rates

		Constant	Lagged Forward Rate	NOBS	R-SQ	Durbin-Watson Statistic
	Canada/U.S. 1951-2003	-0.00046 (-0.288)	0.9919 (0.730)	636	.983	.486
3-Month	Canada/U.S. 1951-1973	0.0008 (0.522)	0.9795 (0.904)	276	.931	.432
	Canada/U.S. 1974-2003	-0.0049 (-1.116)	0.9778 (1.240)	360	.967	.495
	Canada/U.S. 1974-2003	-0.00179 (-1.14)	0.9920 (1.368)	360	.988	1.94
	France/U.S. 1974-2001	-0.02103 (-1.382)	0.9880* (1.392)	336	.975	1.92
1-Month	Germany/U.S. 1974-2001	-0.00879 (-1.399)	0.9888 (1.276)	336	.974	1.88
	Japan/U.S. 1974-2003	-0.03392 (-1.346)	0.9933* (1.359)	360	.991	1.85
	U.K./U.S. 1974-2003	0.01359°° (2.348)	0.9766** (2.216)	360	.960	1.81

Notes and Sources: The figures in brackets are the t-statistics. In the case of the lagged forward rate these pertain to the deviation of the coefficients from unity. The superscripts *, ** and *** indicate one-tailed significance at the 10%, 5% and 1% levels, respectively, while the superscript °° indicates two-tailed significance at the 5% level. In the 3-Month regressions, the obvious serial correlation in the residuals required heteroskedasticity and autocorrelation consistent estimates of the coefficient standard-errors—the truncation lag was set at 6 for the full period and 5 for the split periods. For sources see Appendix A.

Table 11: Regressions of Percentage Changes in Spot Exchange Rates on Lagged 1-Month and 3-Month Forward Premia

	Constant	Lagged Forward Premium Actual	Lagged Forward Premium Implicit	NOBS / R-SQ	F	Durbin-Watson Statistic
3-Month			HAC truncation lag ⇓			
Canada/U.S. 1951-2003	-0.732 (-1.38)	-0.430 (-1.28)	6	636 .006	4.16 [.0418]	.514
Canada/U.S. 1951-1973	0.294 (0.519)	0.295 (0.628)	5	276 .004	0.98 [.3228]	.443
Canada/U.S. 1974-2003	-1.769 (-2.18)*	-0.826 (-2.17)**	5	360 .021	7.80 [.0055]	.541
1-Month						
Canada/U.S. 1974-2003	-2.523*** (-2.57)	-1.576*** (-3.42)		360 .032	11.7 [.0007]	2.01
Canada/U.S. 1974-2003	-2.341** (-2.17)		-1.106** (-2.24)	360 .014	5.01 [.0258]	2.03
Canada/U.S. 1978:1-1990:7	-3.014** (-2.01)	-2.250*** (-3.09)		151 .060	9.56 [.0023]	2.24
Canada/U.S. 1978:1-1990:7	-3.999** (-2.20)		-2.116*** (-2.71)	151 .047	7.34 [.0075]	2.24
Canada/U.S. 1990:8-2003	-2.128 (-1.46)	-1.598** (-2.29)		161 .032	5.26 [.0231]	1.82
Canada/U.S. 1990:8-2003	-1.609 (-1.08)		-0.925 (-1.18)	161 .009	1.39 [.2398]	1.86

Continued on Next Page .....

Table 11: Continued

	Constant	Lagged Forward Premium Actual	Lagged Forward Premium Implicit	NOBS / R-SQ	F	Durbin-Watson Statistic
France/U.S. 1974-2001	-2.678 (-1.17)	-0.607 (-1.10)		336 .004	1.21 [.2723]	2.00
France/U.S. 1974-2001	-2.931 (-1.27)		-0.698 (-1.32)	336 .005	1.74 [.1879]	2.00
France/U.S. 1978:1-1990:7	-1.875 (-0.51)	-0.563 (-0.66)		151 .003	.436 [.5102]	2.11
France/U.S. 1978:1-1990:7	-1.810 (-0.48)		-0.466 (-0.56)	151 .002	.315 [.5752]	2.10
France/U.S. 1990:8-2001	-3.017 (-0.94)	-0.061 (-0.67)		137 .000	.004 [.9469]	1.79
France/U.S. 1990:8-2001	-3.336 (-1.05)		-0.456 (-0.49)	137 .002	.239 [.6259]	1.79
Germany/U.S. 1974-2001	2.197 (0.875)	-0.746 (-1.11)		336 .004	1.23 [.2675]	1.93
Germany/U.S. 1974-2001	2.176 (0.873)		-0.794 (-1.12)	336 .004	1.25 [.2634]	1.93
Germany/U.S. 1978:1-1990:7	20.163*** (2.748)	-4.382*** (-2.76)		151 .048	7.57 [.0067]	2.21
Germany/U.S. 1978:1-1990:7	18.631** (2.508)		-4.256** (-2.48)	151 .040	6.15 [.0143]	2.23
Germany/U.S. 1990:8-2001	-3.046 (-0.97)	-0.438 (-0.42)		137 .001	.175 [.6763]	1.72
Germany/U.S. 1990:8-2001	-3.099 (-0.99)		-0.655 (-0.60)	137 .003	.355 [.5523]	1.73

Continued on Next Page .....

Table 11: Continued

	Constant	Lagged Forward Premium Actual	Lagged Forward Premium Implicit	NOBS / R-SQ	F	Durbin-Watson Statistic
Japan/U.S. 1974-2003	4.0259	-0.282		360	.356	1.91
	(-1.56)	(-0.60)		.001	[.5509]	
Japan/U.S. 1978:1-1990:7	19.073***	-3.348***		151	8.11	2.03
	(2.959)	(-2.85)		.052	[.0050]	
Japan/U.S. 1978:1-1990:7	16.854***		-4.634***	151	7.56	2.07
	(2.831)		(-2.75)	.048	[.0067]	
Japan/U.S. 1990:8-2003	7.918**	-2.059*		161	3.66	1.96
	(1.830)	(-1.91)		.022	[.0577]	
Japan/U.S. 1990:8-2003	7.737*		-2.21*	161	2.91	1.95
	(1.71)		(-1.70)	.017	[.0899]	
U.K./U.S. 1974-2003	-4.436*	-1.452**		360	5.27	1.90
	(-1.79)	(-2.30)		.014	[.0223]	
U.K./U.S. 1974-2003	-4.611*		-1.463**	360	5.24	1.90
	(-1.83)		(-2.29)	.014	[.0227]	
U.K./U.S. 1978:1-1990:7	-8.851**	-4.631***		151	18.1	2.23
	(-2.30)	(-4.26)		.108	[.0000]	
U.K./U.S. 1978:1-1990:7	-9.998**		-4.891***	151	18.7	2.21
	(-2.52)		(-4.33)	.112	[.0000]	
U.K./U.S. 1990:8-2003	2.4425	1.3464		161	1.26	1.84
	(0.676)	(1.124)		.008	[.2629]	
U.K./U.S. 1990:8-2003	2.8091		1.4896	161	1.41	1.84
	(0.752)		(1.188)	.009	[.2364]	

Notes and Sources: The figures in brackets are t-ratios. The superscripts *, ** and *** indicate significance at the 10%, 5% and 1% levels, respectively, using the standard t-test. The actual forward premium is based on actual forward and spot exchange rates; the implicit forward premium is based on the domestic/foreign interest rate differential. For sources see Appendix A.

a recent paper by Maynard.[17] The reasons why the coefficient of β in estimates of (23) is almost uniformly negative rather than plus unity will be addressed in subsequent sections of this paper. First, however, we must address the issue of risk premia.

8 Exchange Rates as Asset Prices: Risk Premia

During the past two decades traditional foreign exchange market analysis has viewed the exchange rate as an asset price. Deviations of the exchange rate from some constant 'purchasing power parity' level were seen as consequences of the evolution of asset prices in the face of policy shocks and other 'news' affecting asset returns.[18]

The fact that the exchange rate is the domestic currency price of agents' holdings of foreign currency clearly makes it an asset price—the flow of services from this asset is the reduction of transactions costs of making international exchange as well as the potential speculative gains from appreciation of the foreign currency. Forward exchange rates represent the price of long or short positions in foreign currency. Forward commitments yield risk-management services and net forward positions are asset holdings on which speculative gains and losses can occur. This suggests that the risks of holding long and short positions in foreign exchange can be analysed using the tools of modern finance.

Risk is the variance of one's portfolio return, and the contribution of any asset to that risk is its contribution to the variance of the return to the portfolio. The only variance of an asset's return that matters is the variance that cannot be diversified away by holding the asset in conjunction with other assets—this is called systematic or non-diversifiable risk. A large group of assets whose returns are uncorrelated will have a nearly constant average return. A risk-free asset is an asset or aggregate of assets whose return is constant—this constant return equals the risk-free interest rate.

---

[17] Alex Maynard, "Testing Forward Rate Unbiasedness: On Regression in Levels and Returns," *Review of Economics and Statistics*, Vol. 85, No. 2 (May), 2003, pp. 313-327.

[18] Robert J. Hodrick, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, New York: Harwood Academic Publishers, 1987, makes a comprehensive presentation of the issues. See also Michael Mussa, "Empirical Regularities in the Behavior of Exchange Rates and Theories of the Foreign Exchange Market," in Karl Brunner and Allan H. Meltzer, eds., *Carnegie-Rochester Conference Series on Public Policy*, Vol. 11, *(Policies for Employment, Prices, and Exchange Rates)*, North Holland, 1979, and Richard M. Levich, "Empirical Studies of Exchange Rates: Price Behavior, Rate Determination and Market Efficiency," in R. W. Jones and P. B. Kenen, eds., *Handbook of International Economics*, Vol. II, Amsterdam: North Holland, 1985.

Imagine a representative agent who has an horizon of T periods and in each period maximises the expected discounted value of the utility levels achieved in the current and subsequent $(T-1)$ periods:[19]

$$E\left\{\sum_{t=0}^{T-1}(1+\mu)^{-1}U(c_t)\right\} \tag{24}$$

where $U(c_t)$ is the utility of consumption c_t in the t-th period and μ is the subjective discount rate. Suppose that at each point in time t this individual can hold her wealth in any of n risky assets having net stochastic rates of return z_{it}, $i = 1 \ldots n$, and in a riskless asset, with a rate of return r_t. If the individual has chosen an anticipated consumption path for which her expected utility is maximised, she will not be able to increase her expected utility by shifting a unit of consumption from any period to any other period by purchasing additional units of any of the assets in her portfolio. This means that for every risky asset,

$$U'(c_t) \;=\; (1+\mu)^{-1}E\{U'(c_{t+1})(1+z_{it})\} \tag{25}$$

and for the riskless asset

$$\begin{aligned} U'(c_t) &= (1+\mu)^{-1}E\{U'(c_{t+1})(1+r_t)\} \\ &= (1+\mu)^{-1}(1+r_t)E\{U'(c_{t+1})\} \end{aligned} \tag{26}$$

where $U'(c_t)$ is the marginal utility of consumption in the t-th period. By shifting a unit of consumption from period t to period $t+1$, for example, the individual would give up $U'(c_t)$ units of utility in period t in order to obtain $(1+z_{it})$ units of output in period $t+1$. These units of output would yield an expected utility of $E\{U'(c_{t+1})(1+z_{it})\}$ in period $t+1$. The discounted value of this expected $t+1$ utility in period t is obtained by multiplying by $(1+\mu)^{-1}$. Since the individual will adjust her consumption path until it does not pay to shift consumption in this fashion, the equalities in (25) and (26) must hold in equilibrium. Substituting (26) into (25) to eliminate $U'(c_t)$ and multiplying both sides by $(1+\mu)$ we obtain

$$\begin{aligned} E\{U'(c_{t+1})(1+z_{it})\} &= (1+r_t)E\{U'(c_{t+1})\} \\ E\{U'(c_{t+1})\} + E\{U'(c_{t+1})(z_{it})\} &= E\{U'(c_{t+1})\} + E\{U'(c_{t+1})(r_t)\} \\ E\{U'(c_{t+1})(z_{it}-r_t)\} &= 0 \end{aligned} \tag{27}$$

[19] The analysis here follows that in Olivier Jean Blanchard and Stanley Fischer, *Lectures on Macroeconomics*, MIT Press, 1989, Ch. 10, 506-509.

Next, note from the constancy of r_t in the above equation and the definition of covariance that[20]

$$\begin{aligned} Cov\{U'(c_{t+1})\, z_{it}\} &= Cov\{U'(c_{t+1})\,(z_{it} - r_t)\} \\ &= E\{U'(c_{t+1})(z_{it} - r_t)\} - E\{U'(c_{t+1})\}E\{z_{it} - r_t\} \\ &= -E\{U'(c_{t+1})\}E\{z_{it} - r_t\} \\ &= -E\{U'(c_{t+1})\}E\{z_{it}\} + E\{U'(c_{t+1})\}E\{r_t\} \qquad (28) \end{aligned}$$

which, noting that $E\{r_t\} = r_t$, can be manipulated to yield

$$E\{z_{it}\} = r_t - \frac{Cov\{U'(c_{t+1}),\, z_{it}\}}{E\{U'(c_{t+1})\}} \qquad (29)$$

The risk premium on the i-th risky asset is therefore

$$-\frac{Cov\{U'(c_{t+1}),\, z_{it}\}}{E\{U'(c_{t+1})\}} \qquad (30)$$

The risk premium on each risky asset is inversely related to the covariance of the return on that asset with the marginal utility of consumption. When consumption is low and the marginal utility of consumption is therefore high, the gain from a positive increment to income and consumption is larger than would be the case when consumption is high and the marginal

[20]The derivation here also uses the facts that, given two random variables x and y and a constant a,

$$\begin{aligned} Cov\{x, y\} &= E\{(x - E\{x\})(y - E\{y\})\} \\ &= E\{xy - E\{x\}y - E\{y\}x + E\{x\}E\{y\} \\ &= E\{xy\} - E\{x\}E\{y\} - E\{x\}E\{y\} + E\{x\}E\{y\} \\ &= E\{xy\} - E\{x\}E\{y\} \end{aligned}$$

and

$$\begin{aligned} Cov\{x, y + a\} &= E\{(x - E\{x\})((y + a) - E\{y + a\})\} \\ &= E\{(x - E\{x\})((y - E\{y\}) + (a - E\{a\}))\} \\ &= E\{(x - E\{x\})((y - E\{y\})\} \\ &= Cov\{x, y\}. \end{aligned}$$

utility of consumption is low. Hence, a variable asset return that is highly inversely correlated with the marginal utility of consumption—i.e., is high when the marginal utility of consumption is low and low when the marginal utility of consumption is high—will be less valuable than one that is less inversely correlated, or positively correlated, with the marginal utility of consumption. The positive shock to income will occur when consumption is already high and the negative shock will occur when consumption is already low. As a result, that asset will have to yield a higher expected return to get people to hold it—it will have a higher risk premium. A negative risk premium will require that the positive shocks to the asset return occur when consumption is low and the marginal utility of consumption is high and negative shocks occur when consumption is high and the marginal utility of consumption is low—in this case the representative agent will be willing to hold the asset at an expected return lower than the risk-free rate of interest.

Imagine now a composite asset, m, whose return is perfectly positively correlated with consumption and hence perfectly negatively correlated with the marginal utility of consumption. This asset can be thought of as a market portfolio consisting of every asset in the economy weighted in proportion to its share of the country's wealth—its return is the return to capital in the economy as a whole. Letting z_{mt} be the return to the market portfolio, we have

$$U'(c_{t+1}) \quad = \quad -\gamma\, z_{mt}. \tag{31}$$

where γ is the constant of proportionality. It follows that for any risky asset z_{it}[21]

$$Cov\{U'(c_{t+1}),\, z_{it}\} \quad = \quad Cov\{-\gamma z_{mt}, z_{it}\} \;=\; -\gamma\, Cov\{z_{mt}, z_{it}\} \tag{32}$$

Using equation (29) to characterise the expected return from holding the market portfolio $E\{z_{mt}\}$ and then substituting (32), we can express the difference between the expected return on the market portfolio and the return

[21]Here we use the fact that

$$\begin{aligned} Cov\{ax, y\} &= E\{(a\,x - E\{a\,x\})(y - E\{y\})\} \\ &= E\{(a\,x - a\,E\{x\})(y - E\{y\})\} \\ &= a\,E\{(x - E\{x\})(y - E\{y\})\} \\ &= a\,Cov\{x, y\}. \end{aligned}$$

on the risk-free asset as follows:[22]

$$
\begin{aligned}
E\{z_{mt}\} &= r_t + \gamma \frac{Cov\{z_{mt}, z_{mt}\}}{E\{U'(c_{t+1})\}} \\
E\{z_{mt}\} - r_t &= \gamma \frac{Var\{z_{mt}\}}{E\{U'(c_{t+1})\}} \qquad (33)
\end{aligned}
$$

This equation can be rearranged to obtain an expression for γ:

$$
\gamma = [E\{z_{mt}\} - r_t] \frac{E\{U'(c_{t+1})\}}{Var\{z_{mt}\}} \qquad (34)
$$

Now, substituting (32) into (29) and using (34) to eliminate γ we obtain

$$
E\{z_{it}\} - r_t = \frac{Cov\{z_{it}, z_{mt}\}}{Var\{z_{mt}\}} [E\{z_{mt}\} - r_t]. \qquad (35)
$$

The i-th asset will be more risky than the market portfolio when its return is positively correlated with the return to the market portfolio and its covariance with the return to the market portfolio exceeds the variance of the return to that portfolio—that is, when its return varies directly with and more widely than the return to capital in the economy as a whole. Wealthowners will require an expected return above the expected return on the market portfolio to make it worth their while to hold this asset and will thereby bid the price of the asset down appropriately relative to its flow of earnings. Assets whose returns are positively correlated with the return to the market portfolio but fluctuate less than it will be less risky than the market portfolio. Asset holders will be willing to hold these assets at an expected return below the expected return on the market portfolio, bidding their prices up relative to their flows of earnings. If the variations in the return to an asset are uncorrelated with variations in the return to the market portfolio no risk premium will be required to get people to hold the asset—its effect on the variance of any individual's portfolio can be completely diversified away. Finally, if the return to an asset covaries negatively with the return to capital in the economy as a whole it provides a hedge against the risk of the market portfolio and wealthowners will be willing to hold it at an expected rate of return lower than the riskless rate.

The above analysis suggests that the risk premium on an uncovered forward position in foreign exchange should depend on how the return on that forward position covaries with the return to capital in the economy

[22] This uses the fact that $Cov\{z_{mt}, z_{mt}\} = Var\{z_{mt}\}$.

as a whole, and hence, how much of the variance in that return the asset holder can diversify away. This is a useful insight as to how to think about foreign exchange risk, but it is as yet impossible to implement in practice. To obtain direct estimates of foreign exchange risk using these principles we need a measure of the return to capital in the economy as a whole—i.e., to the market portfolio. The relevant market portfolio here is presumably one containing every asset in the country, including the human capital of the entire population. How are we to calculate the return on that portfolio? One possibility is to use aggregate real consumption but consumption is an endogenous variable that is affected by people's savings decisions, so variations in it may reflect wealthowners choices and not the earnings flow on their wealth.[23] Also, we have to deal with the fact that a sizable fraction of the world capital stock is non-tradeable human capital embodied in the person of its owners. These human capital assets have no market price. No attempt will be made to grapple with these issues in this presentation—it will simply be assumed that some aggregate 'portfolio' lies behind the scenes and that there will exist risk premia related to the covariance structure of asset returns with this unobservable aggregate.

Although foreign currency holdings are obviously assets whose value is represented by the exchange rate, exchange rates have a much less direct role in pricing bonds, equities and other assets whose earnings are denominated in a currency foreign to their owner. As in the case of pure forward positions in foreign currency, changes in the exchange rate may signify capital gains and losses. But, unlike that case, exchange rate changes may also reflect rather than cause changes in capital values. Imagine the situation faced by a New York resident who owns capital in California. Gains and losses of the value of that capital will occur in the ordinary course of business. If one were then to imagine that California is given its own currency, in which all assets in the State are then denominated, what difference would it make to the New York resident? Presumably none once we allow for price flexibility because the value of the California dollar in terms of the U.S. dollar will adjust until everything in California is worth the same in U.S. dollars as if the latter were the medium of exchange in California—money is a veil. A change in the California dollar price of the U.S. dollar can occur without there being any change in the real value of the New Yorker's holdings of a particular California asset if the exchange rate and the asset price in California dollars

[23] Robert E. Lucas, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 4 (December) 1978, 1426–1445, explores the implications of truly exogenous consumption for asset pricing by modelling an economy in which output in each period is exogenous and perishable and in which consumption is equal to output because no saving is possible.

move in unison. The real value of a New Yorker's holdings of California assets may or may not be correlated with movements of the exchange rate.

Suppose that the real exchange rate is a constant. Then movements in the nominal exchange rate simply reflect differences in the movements of the domestic and foreign price levels. If there is unexpected inflation abroad, capital losses will be experienced on assets whose nominal earnings are fixed in foreign currency. This will be reflected in a devaluation of the foreign currency but the exchange rate movement is simply a reflection of the foreign inflation. Assets whose earnings are fixed in real terms abroad will experience no capital loss, measured in units of either foreign or domestic output. Their nominal prices will rise with the rise in the price level abroad but this will be exactly compensated for by the decline in the domestic currency value of foreign currency. If there is anticipated inflation abroad then an inflation premium will be added to interest rates in the foreign economy to compensate lenders.

This rise in the foreign nominal interest rate relative to the domestic nominal interest rate will be exactly matched by a rise in the forward premium on domestic currency as indicated in equation (16). Again, the exchange rate movements reflect the anticipated inflationary conditions in the foreign economy—they play no independent role. If we substitute the efficient markets condition (17) into the interest parity condition (16) we obtain

$$i_t^* - i_t \quad = \quad E_\Pi - \phi_t - \theta_t. \tag{36}$$

Then, substituting into this equation the Fisher equations

$$i_t = r_t + E_P$$

and

$$i_t^* = r_t^* + E_{P^*},$$

where r_t and r_t^* are the domestic and foreign real interest rates, we obtain

$$r_t^* + E_{P^*} - r_t - E_P \quad = \quad E_\Pi - \phi_t - \theta_t \tag{37}$$

which reduces to

$$\begin{aligned} r_t^* - r_t \quad &= \quad (E_P + E_\Pi - E_{P^*}) - \phi_t - \theta_t \\ &= \quad E_Q - \phi_t - \theta_t \end{aligned} \tag{38}$$

where E_Q ($= E_P + E_\Pi - E_{P^*}$) is the expected rate of change in the real exchange rate. Since the real exchange rate is constant in the above discussion, E_Q will equal zero and the foreign inflation will have no effect on the domestic/foreign real interest rate differential—E_{P^*} and E_{Π} both rise by the same amount.

Now suppose that both countries experience zero inflation and that the real exchange rate changes. An increase in the real exchange rate represents a rise in the price of domestic output in units of foreign output. This implies that domestically employed capital, whose service flow is measured in units of domestic output, is now more valuable in units of foreign output. The owners of capital employed in the domestic economy receive a capital gain. When P and P^* are constant the domestic currency must appreciate (Π must rise). But the capital gain is fundamentally unrelated to the nominal exchange rate movement in the sense that had the nominal exchange rate been fixed (or had the two countries had a common currency) the real exchange rate would still have changed and the capital gain on domestically employed capital would still have been received. The domestic price level would have then risen relative to the foreign price level by the increase in the relative value of domestic goods in terms of foreign goods. The nominal exchange rate is playing a passive role.

If an increase in the real exchange rate is anticipated (by the residents of both countries) we can see from (38) that the domestic real interest rate will decline relative to the foreign real interest rate. This happens because the owners of domestically employed capital anticipate a future capital gain and are willing to hold that capital at a lower (net of capital gain) real interest rate than previously. This interest rate effect would also occur independently of whether the nominal exchange rate is fixed or flexible—the capital gain has to do with technological or other real-sector developments in the domestic relative to the foreign economy and is independent of the currency system.

The fundamental issue of concern here is the relationship between exchange rates and the risk premia on countries' assets. The issues involved are exceedingly complex. Assume for the moment that the real exchange rate is constant and that in the absence of international capital flows the return to the underlying unobservable market portfolio in the domestic economy varies more widely than but is highly correlated with the return to the underlying unobservable market portfolio abroad. Domestic residents will be living with greater risk than residents abroad and, assuming that domestic and foreign residents have the same risk aversion, the premium of market interest rates over the risk-free rate will be greater in the domestic economy than abroad. When international trade in securities is then allowed, we can imagine this risk differential persisting, although to the extent that domestic and foreign asset returns are not perfectly correlated with each other there will be a gain to portfolio diversification that will induce each countries' residents to hold a fraction of their portfolio in ownership claims to capital employed in the other country. This international pooling of asset holdings will increase if domestic and foreign residents have different aversions to risk. The relevant market portfolio will now be the 'world' market portfolio and interest rates on domestic equities will be above the rate of return on this portfolio and interest rates on foreign equities will be below it. The domestic/foreign interest rate differential will depend on the amplitude of the variation of the return to domestic equities relative to the amplitude of variation of the return to foreign equities as compared to the variation of the return on the underlying unobservable international market portfolio.

Things change when we introduce the possibility of variations in the real exchange rate. Suppose for the moment that real exchange rate movements are completely uncorrelated with movements in the return to the world market portfolio and with the returns to the separate domestic and foreign market portfolios that would exist in the absence of international trade in equities. A rise (fall) in the real exchange rate will create a capital gain (loss) on foreign residents holdings of domestic assets and a capital loss (gain) on domestic residents holdings of foreign assets. This assumes that the residents of every country measure their wealth in units of that country's output. These variations in wealth can only be diversified away by domestic and foreign residents by holding very small fractions of their portfolios in equities of the other country. Even if a country's asset holders diversified their foreign holdings across assets in all other countries, the resulting diversified asset should not have a weight in excess of that of any of the individual home country assets in the overall portfolio—otherwise the diversifiable risk would not be diversified away. The introduction of real exchange rate variability, uncorrelated with the returns to capital in the domestic and foreign economies measured in terms of units of their own output, would thus cause asset holders to hold high fractions of their wealth in their own country's assets. This is a potential explanation of the well-known 'home bias portfolio puzzle'.[24] A positive covariance of these real exchange rate shocks with the returns on domestic assets, on the other hand, will act as a hedge against domestic asset returns, moderating both the increases and decreases in wealth over the course of the business cycle and leading domestic asset holders to hold more foreign assets in their portfolios than they otherwise would. It

[24]For a discussion of this puzzle and references to the literature, see Maurice Obstfeld and Kenneth Rogoff, "The Six Major Puzzles in International Macroeconomics: Is there a Common Cause?" in Ben Bernanke and Kenneth Rogoff (eds.), *NBER Macroeconomics Annual 2000*, Cambridge: NBER and the MIT Press.

would have the reverse effect on the desired portfolio holdings of domestic assets by foreign asset holders. The exact nature of these effects becomes less clear when we allow for the fact that residents of each country consume goods obtained abroad—capital gains and losses should then be measured in terms of a weighted average of domestic and foreign output units with the weight depending on the fraction of consumption (or perhaps better, absorption) falling on goods produced outside the country.

Some evidence relevant to the above discussion is presented in Figures 10 through 14. The top panel in each figure plots the year-over-year growth rates of real GDP for an individual country and the United States. The coefficient of correlation of the two series is given below the chart. The middle panel in each figure gives compares the same country's year-over-year CPI inflation rate with that of the U.S. Again, the correlation coefficient is given below the chart. Finally, the bottom panel gives the percentage deviation of the country's real exchange rate with respect to the United States from its trend, along with the percentage deviations of the country's real GDP and the U.S. real GDP from their respective trends. The correlations between the country's percentage deviation of real GDP and that of the U.S. and between the percentage deviations of its real exchange rate and its real GDP are indicated below the chart. The time periods of the plots were chosen on the basis of availability of data uncontaminated by regime changes and the correlations for Germany exclude the run-up to and the period following German reunification.

There is clearly a rough correspondence of real GDP growth of the countries examined with that of the U.S. suggesting a loose similarity of the business cycle across the industrial countries.[25] Even a tight correspondence of real GDP growth across countries, however, would not necessarily indicate a similar tight correspondence between their underlying unobservable mar-

[25] Correlations of roughly the magnitude obtained here where found by David K. Bachus, Patrick J. Keyhoe and Finn E. Kydland, "International Business Cycles: Theory and Evidence", in T. Cooley, (ed) *Frontiers of Business Cycle Research*, Princeton University Press, 1995. These authors detrended the output variables using the Hodrick-Prescott filter. Harris Dellas, in "A Real Model of the World Business Cycle," *Journal of International Money and Finance*, Vol. 5, No. 3 (September) 1986, 381-94, used spectral analysis to obtain a pair-wise coherence coefficient of .9 for the outputs of the U.S. and Germany at a frequency of 2.5 years and of .7 for the outputs of the U.S. and the U.K. at a frequency of nine quarters. The coherence coefficients for comparisons of Japan with the U.S., the U.K. and Germany ranged between .5 and .7 at frequencies of between 3.5 and 5 years. Stefan Gerlach, in "World Business Cycles Under Fixed and Flexible Exchange Rates," *Journal of Money, Credit and Banking*," Vol. 20, No. 4 (November), 1988, 620-630, also found that the output movements of Belgium, Canada, France, Germany, Italy, Netherlands, Norway and Sweden were correlated in the business-cycle frequency band.

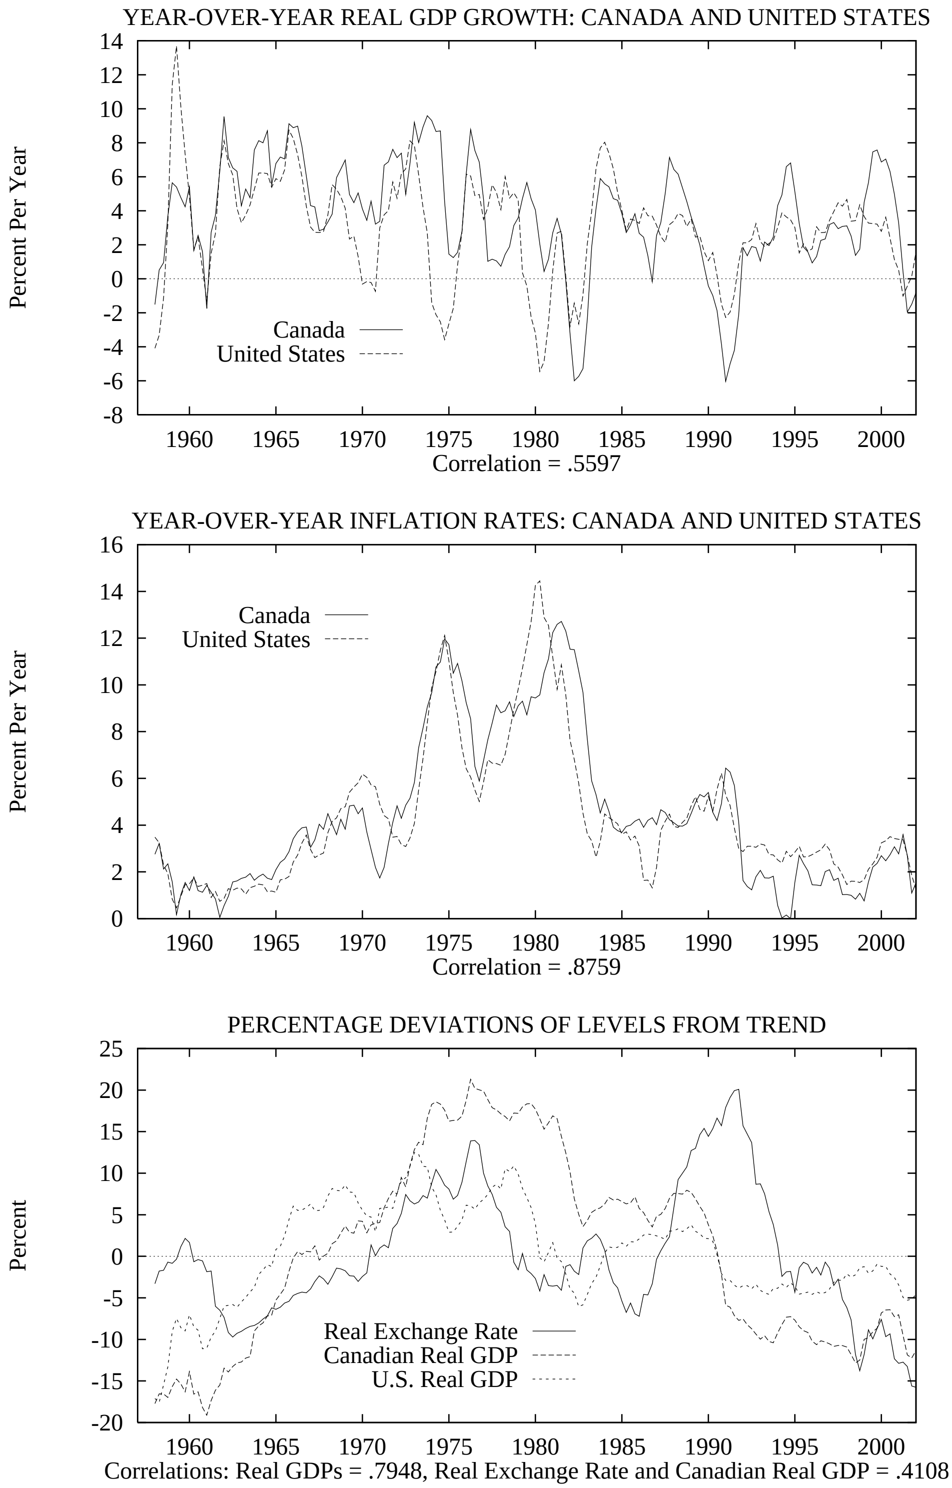

Figure 10: Year-over-year Canadian and U.S. real GDP growth rates and CPI inflation rates, 1958-2002, and percentage deviations of Canadian and U.S. real GDPs and the Canadian real exchange rate with respect to the U.S. from trend, 1957-2002. Source: International Monetary Fund, *International Financial Statistics*.

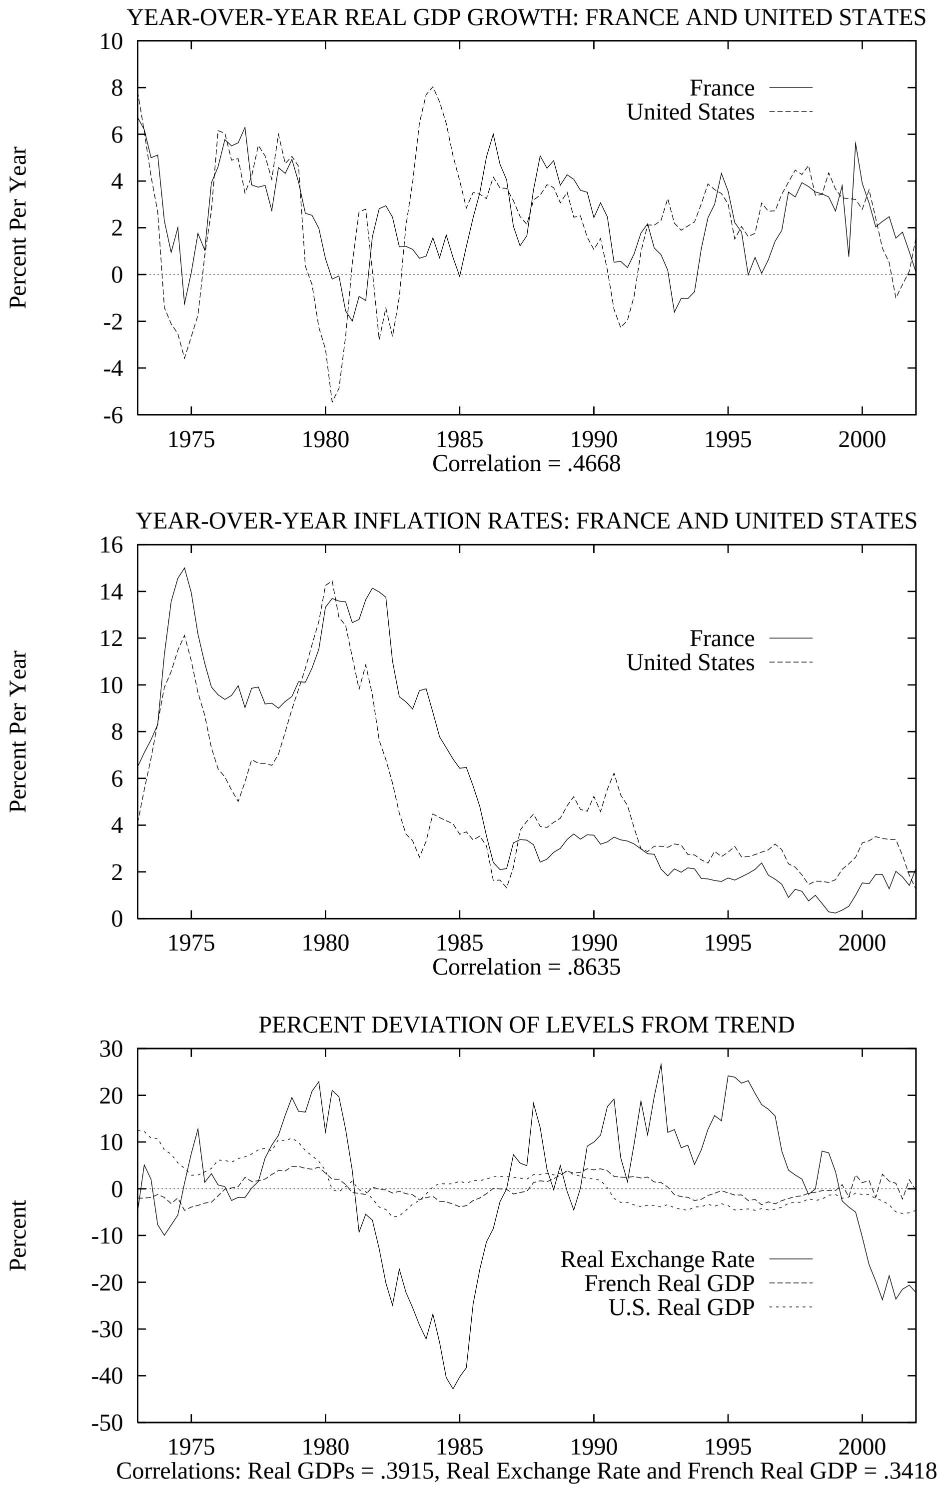

Figure 11: Year-over-year French and U.S. real GDP growth rates and CPI inflation rates and percentage deviations of French and U.S. real GDPs and the French real exchange rate with respect to the U.S. from trend, 1973-2002. Source: International Monetary Fund, *International Financial Statistics*.

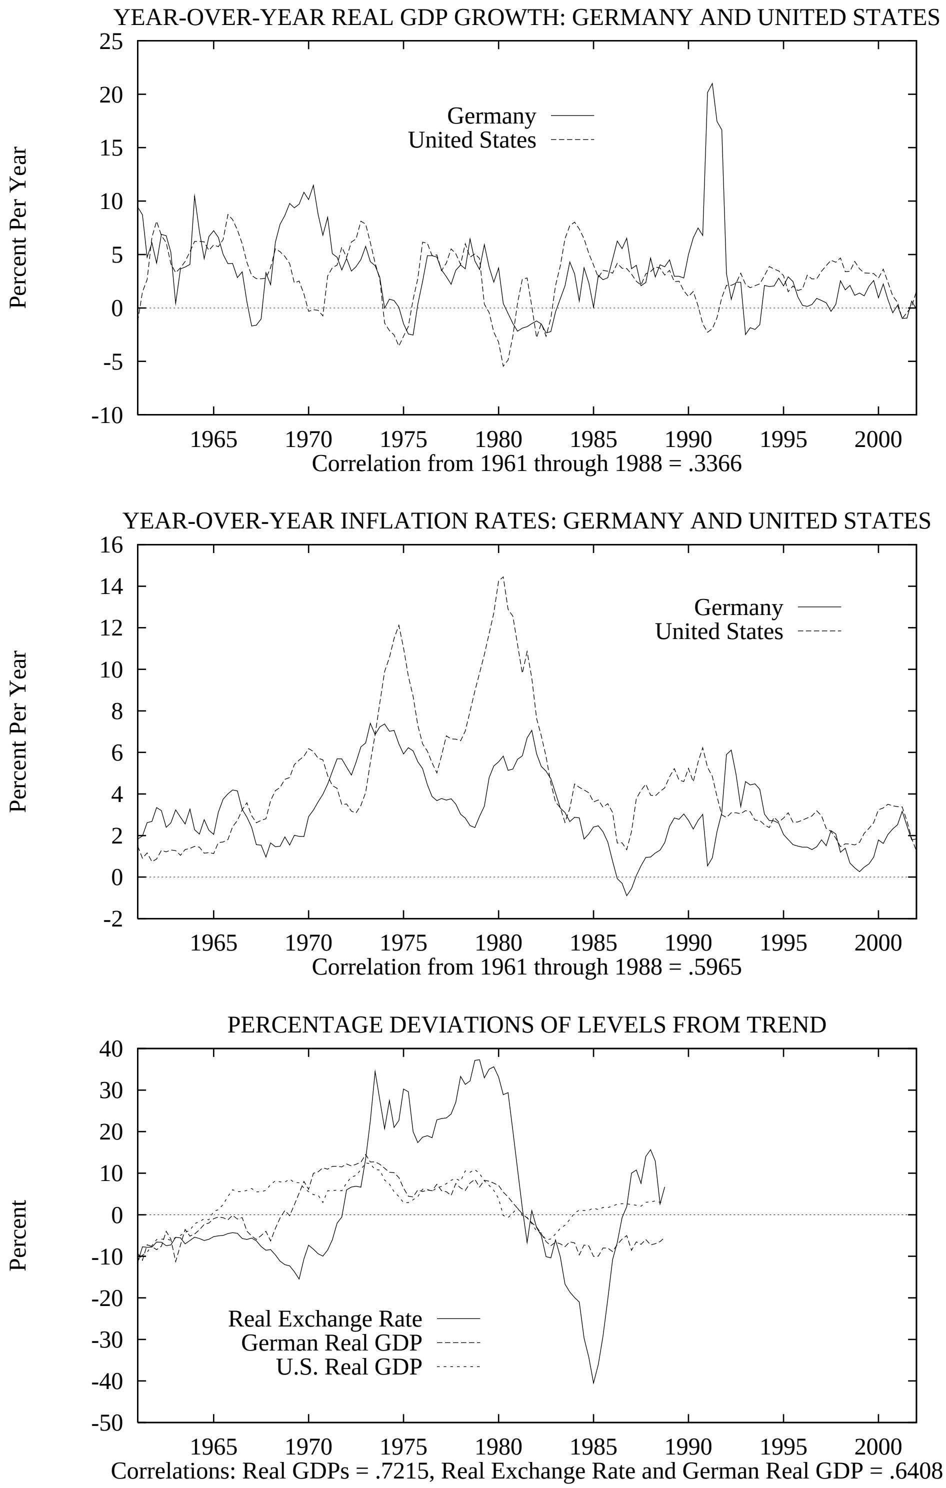

Figure 12: Year-over-year German and U.S. real GDP growth rates and CPI inflation rates, 1961-2002, and percentage deviations of German and U.S. real GDPs and the German real exchange rate with respect to the U.S. from trend, 1961-1988. Source: International Monetary Fund, *International Financial Statistics*.

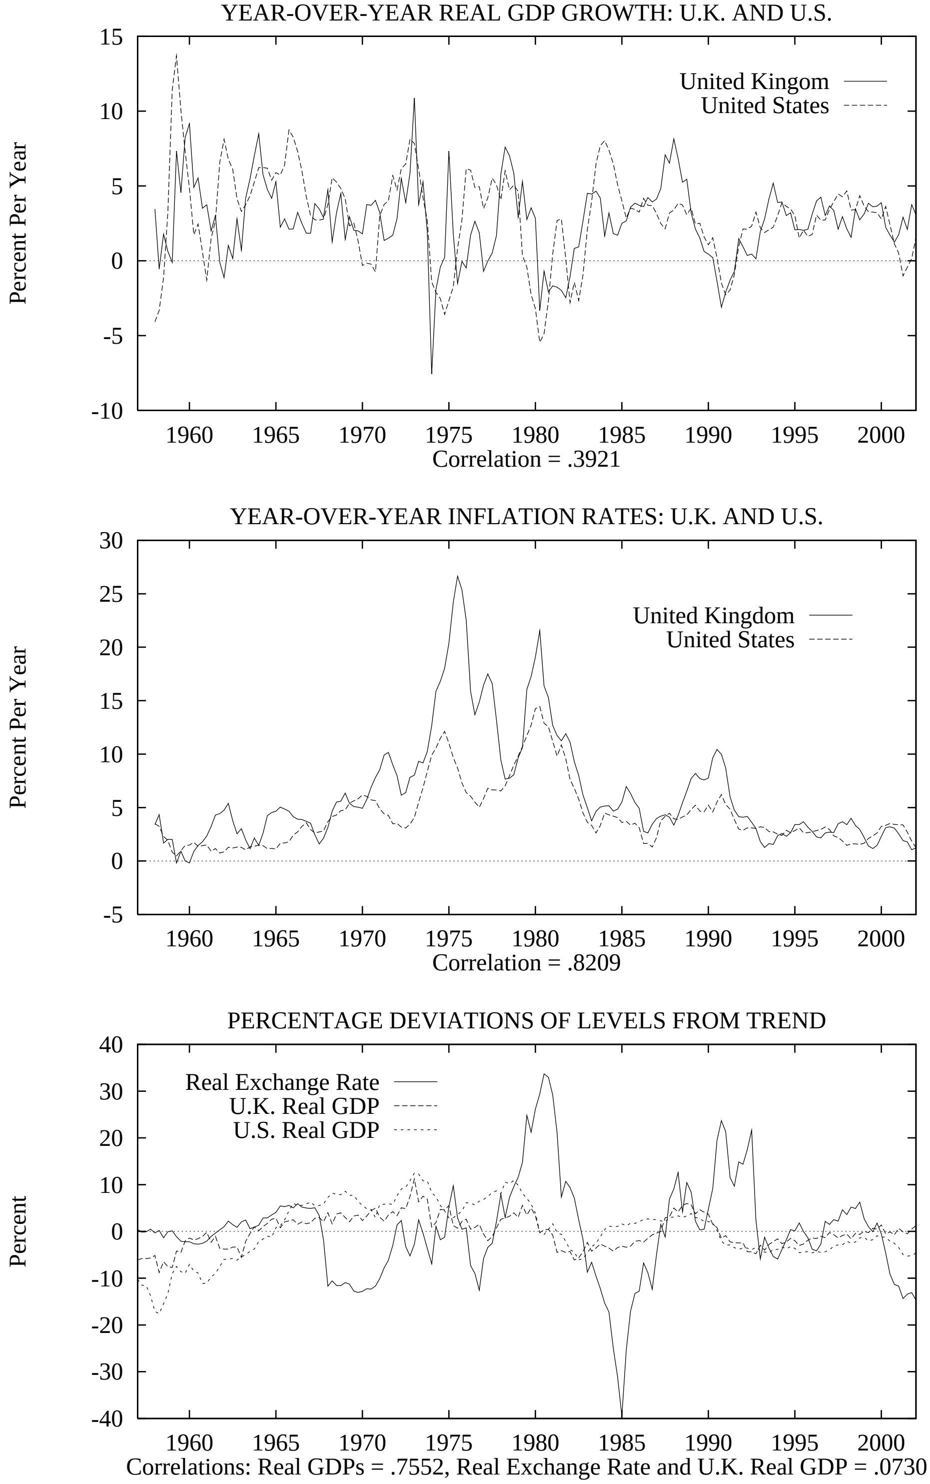

Figure 13: Year-over-year U.K. and U.S. real GDP growth rates and CPI inflation rates, 1958-2002, and percentage deviations of U.K. and U.S. real GDPs and the U.K. real exchange rate with respect to the U.S. from trend, 1957-2002. Source: International Monetary Fund, *International Financial Statistics*.

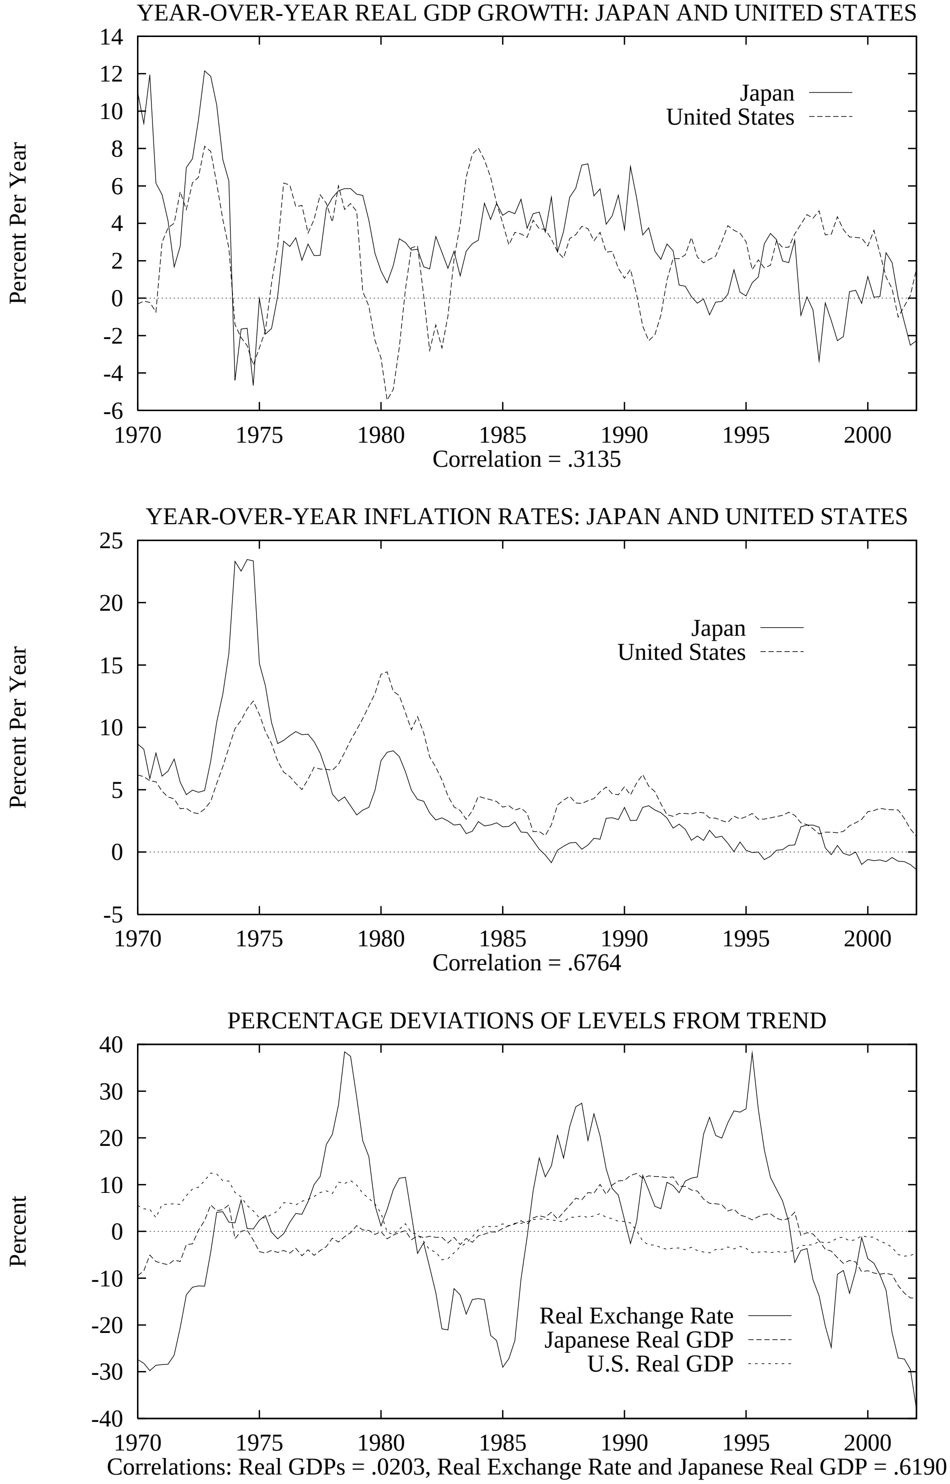

Figure 14: Year-over-year Japanese and U.S. real GDP growth rates and CPI inflation rates and percentage deviations of Japanese and U.S. real GDPs and the Japanese real exchange rate with respect to the U.S. from trend, 1970-2002. Source: International Monetary Fund, *International Financial Statistics*.

Table 12: Standard Deviations of Percentage Deviations from Trend: Real GDP and Real Exchange Rate vs. the United States

	Period	Real GDP	Real Exchange Rate
U.S.	1958-2002	5.92	
Canada	1958-2002	11.25	7.77
France	1973-2002	2.37	15.99
Germany	1961-1988	7.29	17.56
U.K.	1958-2002	3.34	10.71
Japan	1970-2002	6.47	17.77

Source: International Monetary Fund, *International Financial Statistics*.

ket portfolios referred to in the arguments above. The evidence suggests a relationship but we cannot determine its strength. The correspondence of the inflation rates across countries is much stronger than that of real GDP growth. This again suggests similarity of these countries' business cycles and of their monetary policies through time. Except for the U.K. there is a positive correlation between the percentage deviations of the countries' real exchange rates from trend and the percentage deviations of their real GDPs from trend.

In comparison to the percentage deviations of real GDP from trend, the percentage deviations of the countries' real exchange rates with respect to the United States from trend are, with the exception of Canada, enormous. The summary statistics are shown in Table 12. For Germany, Britain and Japan the standard deviations of the real exchange rate variable are between 2 and 3 times the standard deviations of the country's real GDP. In the case of France the ratio exceeds 6.5. The fact that the standard deviation of the percentage deviations of Canada's real exchange from trend is less than half that of the other countries seems consistent with a greater integration of the Canadian economy with that of the U.S. as compared to the European countries and Japan. By comparison, the standard deviation of the percentage deviation from trend of the French real exchange rate with respect to Germany (not shown in the Table) is 4.42 percent during the 1973-1988 period, 1.5 during 1992-2002 and only 0.42 percent during 1998-2002.

A full analysis of the country risk premia on equities is beyond the scope of this paper—the above discussion and evidence presented is little more than a suggested direction for further research. That having been said, a brief discussion of issues concerning the country risk premia on bonds and other assets of fixed nominal redemption value is also in order.

Although the market value of bonds could be correlated with the return on the underlying unobserved market portfolio in association with the fluctuation of market interest rates, the main risk from holding assets of fixed nominal redemption value is the probability of default. Although default is more likely to occur when times are bad and the return on the market portfolio is low, investors will presumably have in mind an estimated probability that default will occur sometime during the life of the asset, with the chance of future bad times subsumed in this probability. A risk factor relevant to bond holdings and not to equities is the effects of unanticipated inflation and deflation. It is probably the case that unanticipated changes in the inflation rate are positively correlated with variations in the return to the market portfolio thereby causing bonds to be somewhat of a hedge—that is, when times are good and the return to the market portfolio is above average, inflation will tend to be high and the real returns to bond holdings below average and when times are bad the return to the market portfolio will fall and the return to holding bonds will be above average. Movements in the real exchange rate also represent a risk factor for holdings of securities of fixed nominal redemption value because a rise in a country's real exchange rate reduces the real value in units of domestic output of bonds held abroad. Again, however, we have to keep in mind that this same rise in the real exchange rate may be positively related to domestic wealth and this wealth effect could offset part of, or even more than offset, the real capital loss on bonds held outside the country.

The above analysis has dealt only with country risk. Under conditions of price flexibility, the movements of the real exchange rate will be the same whether the nominal exchange rate is fixed or flexible. Foreign exchange risk arises solely from gains and losses on short-term uncovered positions in foreign exchange—these can only arise from fluctuations in the nominal exchange rate in relation to levels predicted by the previous period's forward rate. Investors' horizons in the case of these assets are typically one to three months but may extend as long as a year. Long-term positions in foreign assets are not covered in the forward market because the movements of the real exchange rate over the long run will be the same regardless of the level of the nominal exchange rate—if the nominal exchange rate is fixed, for example, then the real exchange rate movements will be reflected in the relative domestic and foreign price levels.

Figure 15 plots the prediction errors from using the forward exchange rate to predict the future spot rate for 1-month and 3-month positions in the Canadian dollar vs. the U.S. dollar and for 1-month positions in the British pound vs. the U.S. dollar. These prediction errors, actual minus

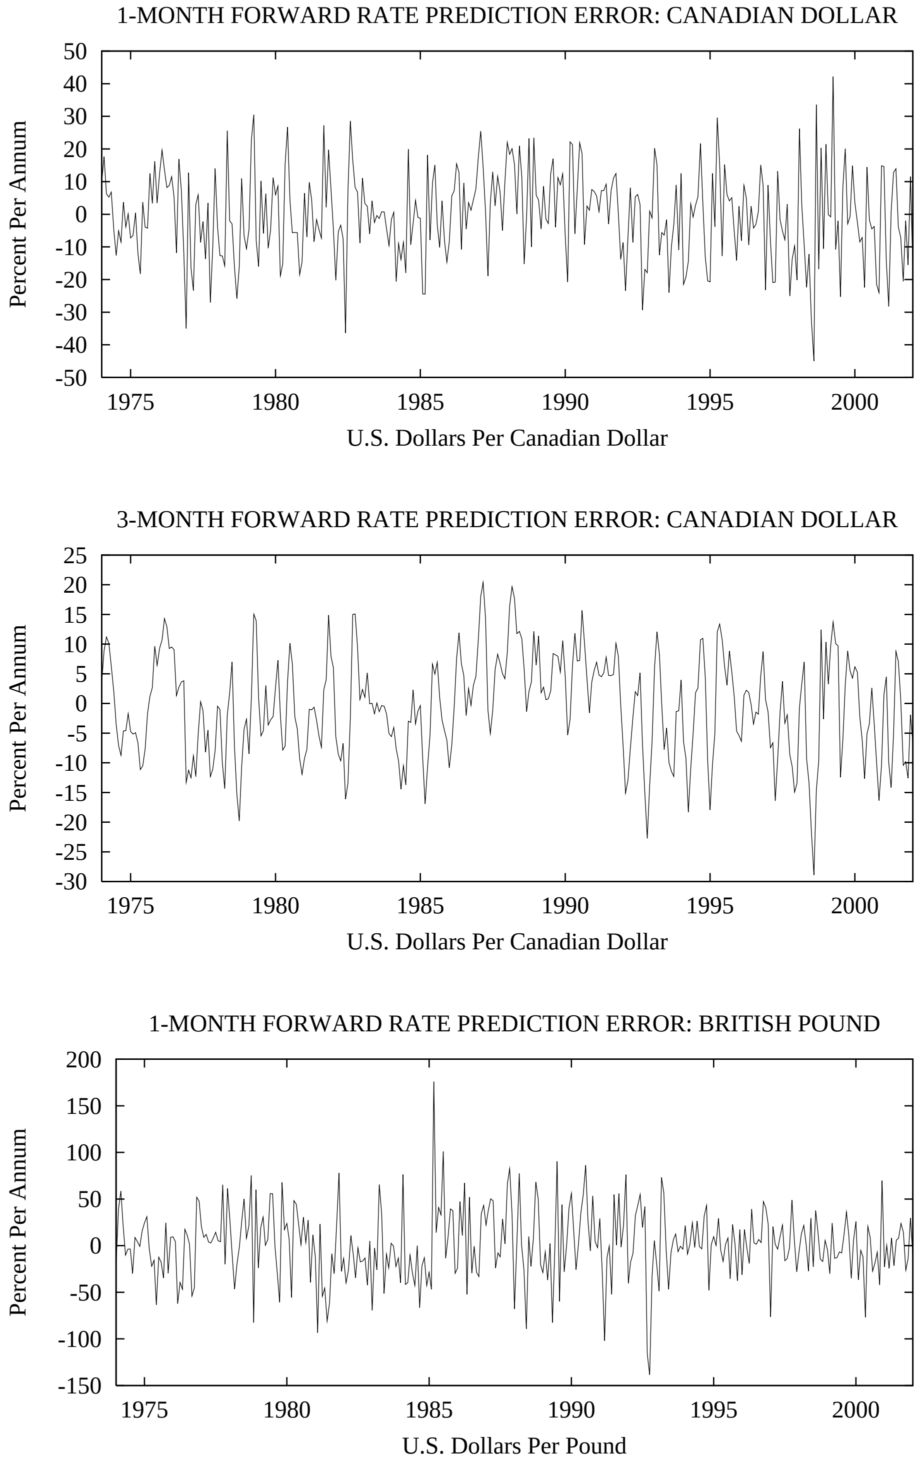

Figure 15: Prediction errors from using the forward exchange rate to forecast the future spot exchange rate. The exchange rates are defined as U.S. dollars per Canadian dollar and per U.K. pound. Sources: *Cansim*, *Reuters* and *Datastream*.

predicted, give the gains (losses) from short (long) positions in U.S. dollars by Canadian and British residents. A glance at the Figure suggests that the variations are predominately short-term noise that should be uncorrelated with underlying and unobservable market portfolios. The general patterns in corresponding plots (not shown) for the Deutschmark, French franc and Japanese yen are indistinguishable to the native eye from those in Figure 15. This suggests that restricting short-term uncovered foreign exchange positions to small fractions of portfolios, all this foreign-exchange risk can be diversified away—a seemingly easy accomplishment. But this intuition may well be wrong. As can be seen in Table 13, the one-month forecast error for the U.K. pound is significantly positively correlated with the percentage deviations of British real GDP from trend at the 5% level and a five-month centred moving average of these one-month forecast errors is significantly positively correlated with that output measure at the 1% level. The year-over-year growth rate of British GDP is also significantly correlated with the five-month centered moving average of the one-month forecast errors. The centered moving average of the 1-month forward rate forecast error for the Canadian dollar is significantly correlated with the percentage deviations of Canadian real GDP from trend at the 1% level. The 1-month forward rate forecast errors and the five-month centered moving average of those forecast errors for the Deutschmark are both significantly correlated with the year-over-year growth rate of German real GDP. These correlations suggest that, apart from the Japan/U.S. case, we cannot rule out the possibility that the returns from short-term uncovered foreign exchange positions are positively correlated with the relevant underlying market portfolios.

9 Explaining the Forward Premium Puzzle

Given the relationships between real exchange rates, nominal exchange rates and domestic relative to foreign price levels, it makes little sense to analyse exchange rate movements exclusively as an asset pricing problem—that is, using models in which asset pricing adjustments are the sole driving force. The practice has been justified by the interpretation of real exchange rate movements as deviations from equilibrium purchasing-power-parity levels. As an examination of the real exchange rate movements in Figures 1 and 2 makes clear, it is inappropriate to assume that the real exchange rate is constant. But one might nevertheless adopt the assumption that the equilibrium real exchange rate is constant (after all, it is often asserted, arbitrage should imply that every good must have the same real price in

Table 13: Correlations of Percentage Deviations of Real GDP from Trend and Real GDP Growth with Errors from Forward Rate Based Forecasts of Spot Exchange Rates

	Percent Deviation of Real GDP from Trend with		Year-over-Year Real GDP Growth with	
	1-Month Forecast Error	5-Month Moving Average Forecast Error	1-Month Forecast Error	5-Month Moving Average Forecast Error
Canada/U.S. 1974:Q3 to 2002:Q2	.1625 (.087)	.2445 (.009)	-0.0369 (.699)	-0.0486 (.611)
France/U.S. 1974:Q3 to 2001:Q2	.0752 (.439)	.1134 (.243)	.1333 (.169)	.2583 (.007)
Germany/U.S. 1974:Q3 to 1988:Q4	.1301 (.326)	.0881 (.507)	.3186 (.014)	.5594 (.000)
U.K./U.S. 1974:Q3 to 2002:Q2	.1901 (.044)	.3350 (.000)	.1539 (.104)	.3036 (.001)
Japan/U.S. 1977:Q4 to 2002:Q2	.0380 (.667)	.1479 (.142)	.0121 (.906)	.0930 (.357)

The figures in brackets are P-Values. The moving average figures are five-month centered moving averages of the one-month forward rate based forecasts of the spot exchange rates. For sources see Appendix A.

all countries) and treat observed parallel movements in real and nominal exchange rates as disequilibrium deviations from purchasing power parity that will be corrected with time. Once these disequilibria are corrected, the nominal exchange rate movements associated with them will be reversed. It is then but a short further step to interpret these real and nominal exchange rate shocks as the result of foreign exchange market speculation fuelled by differences in countries' monetary policies and 'news' about future policy developments. As in the case of other asset prices, perceived changes in the future path of the exchange rate can then be thought of as having a leverage effect on its current level.

A weakness of this approach is that it imposes *ad hoc* price level stickiness. A speculative shift out of domestic currency into foreign currency, or a shock to the money supply, leads to a change in both the real and nominal exchange rates only because the price levels cannot adjust in response to these monetary forces—were they to do so, the nominal exchange rate adjustment would simply reflect the underlying equilibrium price level changes and all real factors, including the real exchange rate, would be unaffected. With price levels that are sticky but adjust eventually, the deviations from purchasing power parity will be temporary until equilibrium is restored. The problem is that the very slow mean reversion of real and nominal exchange rates implies that complete price level adjustments to nominal shocks must take many years—much longer than reasonable. Another weakness of the approach is that the assumption that purchasing power parity holds in long-run equilibrium has no particular theoretical basis.

9.1 Real Exchange Rates as Relative Output Prices

The real exchange rate is defined as the relative price of domestic output in terms of foreign output. This makes it natural to interpret real exchange rate movements as international relative price adjustments. In a world where countries differ in their natural resource endowments and technological change is ongoing, where labour is not internationally mobile, and where there are goods that are produced in significant part by inputs that are not internationally traded, one would be surprised if purchasing power parity held. Constancy of the real exchange rate through time would be no more plausible than constancy of the price of wheat, automobiles, or TV sets. And the first place one would look to explain real exchange rate movements would be at the factors affecting equilibrium international relative prices.

These factors determining real exchange rate movements can be incorporated in a standard general equilibrium model of domestic and foreign production and consumption of traded and non-traded output components. The derivations are technically straight-forward but extremely messy and not very informative beyond what can be seen on the basis of intuition alone. The problem is that we end up with a reduced form equation and cannot unravel the underlying structural parameters without a much better understanding of how to define and model technological change than is now possible.[26] For example, observed domestic and foreign real income movements may be either positively or negatively related to movements of the real exchange rate because increases in income typically arise from increases in production of non-traded (home) and traded (home and foreign) output components and have at the same time income effects on the demands for these components of output. The effect on the relative price of home in terms of foreign output will depend on whether the growth of technology favours traded vs. non-traded components on the supply side and on the magnitudes of the income effects on the demands for the two types of output components. Shifts of aggregate demand from the private to the public sector might be expected to shift demand to domestically produced output from foreign produced output and thereby raise the real exchange rate. Oil, agricultural, and metals prices influence the terms of trade and thereby directly affect the price of domestic output in terms of foreign output. At the same time, these changes in the international valuation of domestic output affect the distribution of world income and the demand for traded and non-traded output components both at home and abroad. This, combined with the fact that the terms of trade are necessarily related to supply-side forces

[26] A substantial literature has developed on these and related issues beginning with Bela Balassa, "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, Vol. 72, No. 6 (December), 1964, 584–96, Paul A. Samuelson, "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, Vol. 46, No. 2 (May) 1964, 145–54, and later, Lawrence H. Officer, *Purchasing Power Parity and Exchange Rates: Theory, Evidence and Relevance*, London and Greenwich Connecticut: JAI Press, 1982. More recently, the emphasis has been on representative agent models—see Allan C. Stockman, "A Theory of Exchange Rate Determination," *Journal of Political Economy*, Vol. 88, No. 4 (August), 1980, and "Real Exchange Rates Under Alternative Exchange Rate Regimes," *Journal of International Money and Finance*, Vol. 3, No. 2 (June), 1983, 147–66, Allan C. Stockman and Lars E. O. Svensson, "Capital Flows, Investment and Exchange Rates," *Journal of Monetary Economics*, Vol. 19, No. 2 (March) 1987, 171–202, Elhanan Helpman, "An Exploration in the Theory of Exchange Rate Regimes," *Journal of Political Economy*, Vol. 89, No. 5 (October), 1981, 865–890, Elhanan Helpman and Assaf Razin, "Dynamics of a Floating Exchange Rate Regime," *Journal of Political Economy*, Vol. 90, No. 4 (August), 1982, and Sebastian Edwards, *Real Exchange Rates, Devaluation, and Adjustment*, MIT Press, 1989.

affecting the relative prices of non-traded in terms of traded components in either or both the domestic and foreign economies, produces important indirect relationships between commodity prices and the real exchange rate.

Although political forces and resulting government policy actions will frequently play a direct role, probably the most important driving force here is world technology, which is impossible to model. As the world economy grows some natural resources are likely to be better candidates for development than others, and capital will flow to the regions in which these favoured resources are present. This will cause the real exchange rates in these regions to rise to effect the inward transfer of real capital. At the same time, the differential impact of world cyclical and other transitory factors across areas is likely to induce intertemporal consumption smoothing, which will temporarily raise the real exchange rates of borrowing regions. These may or may not be the regions that are receiving longer-term injections of rest-of-world savings for purposes of resource development. Finally, countries' permanent savings rates may change.

It is not surprising that these real forces should lead to real exchange rate series best described as near-random-walk processes. In any period the 'hit' to the real exchange rate resulting from on-going technological change and the response of the world economy to it, as well as political developments affecting local savings and investment opportunities, would seem just as likely to be in one direction as in the other. This will result in real exchange rate series that wander widely from any initial values. At the same time, it would seem improbable that any country or region should have more than its share of the bad (or good) luck over a long period of, say, 100 years. An area's resource endowment will perhaps be favoured by world technology for a span of years but then new resources that the region does not possess will become the focus of development. Political instability may reduce investment opportunities and savings rates for a time but then things will eventually change as conflicts are resolved, migration occurs and national borders are redrawn. Virtually zero trends of countries' real exchange rates over many decades, and the long-term mean reversion that this implies, should not be a surprise. This suggests that the effects of individual shocks to the real exchange rate, equally likely to be in either direction when they occur, eventually dissipate with time although the time required for their effects to disappear is typically very long and varies enormously from shock to shock. There are, of course, technological, political and other factors that explain all movements in real exchange rates, but we do not have the capacity to model and forecast them. So real exchange rates appear to us as random walks in the short-run with mean reversion in the very long run.[27]

Market participants have no more information than economists. To forecast real exchange rate movements one has to know the parameters of the international relative price structure and predict how the forcing variables will evolve in the future—an impossible task given the current state of knowledge. Accordingly, it is not surprising that the best predictor of tomorrow's real exchange rate tends to be today's.

9.2 Implications for the Forward Premium: Equal Domestic and Foreign Inflation Rates

The near-random-walk behaviour of the real exchange rate has important implications for determining the forward premium. To explore these it is useful to abstract for the moment from differences between the domestic and foreign inflation rates.

Suppose that the real exchange rate is a random-walk process, domestic and foreign inflation rates are the same and these facts are known by all agents. Two possible scenarios arise according to whether agents have (or think they have) information about the innovations in this random walk process. If agents knowingly have no information then the forward exchange rate will equal the current spot rate and the β coefficient in equation (23) will be undefined. The β-coefficient in (22) will be unity.[28] [29]

The alternative scenario is that agents believe that they have information about the innovations in the process determining the real exchange rate. If they have perfect information, they will forecast the real exchange rate accurately each period and the β-coefficients in both (22) and (23) will be

---

[27] I would like to thank Angelo Melino for helping me clarify my thinking on this point.

[28] If the real exchange rate is a random walk and there is no inflation,

$$s_{t+1} = s_t + \epsilon_t.$$

Since $f_t = s_t$ we can write this as

$$s_{t+1} = \beta\, f_t + \epsilon_t$$

where $\beta = 1$.

[29] Scott W. Barnhart and Andrew C. Szakmary, in "Testing the Unbiased Forward Rate Hypothesis: Evidence on Unit Roots, Co-Integration, and Stochastic Coefficients", *Journal of Financial and Quantitative Analysis*, Vol. 26, No. 2 (June), 1991, construct a purely econometric analysis of the relationship between (22) and (23), arguing that the correct analysis requires an error correction model that combines the "long-run" effects in (22) with the "short-run" effects that are poorly captured in (23). Their empirical findings are fully consistent with the argument presented here.

unity. If they act on information they think they have, but that information is worthless, the β-coefficient in (23) will be zero since the random variation in $(f_t - s_t)$ will be uncorrelated with the variation in $(s_{t+1} - s_t)$. At the same time, the 'forward rate' representation (22) will yield an estimate of β close to unity as long as the forward rate tracks the spot rate well through time. This will happen when, as is in fact the case, the variance of the forward premium is small in relation to the variance of the innovations to the spot rate. Suppose that, on average, agents correctly predict some fraction of the innovations in the real exchange rate. Then there will be a positive correlation between $(f_t - s_t)$ and $(s_{t+1} - s_t)$ in (23) and a positive estimated value for β. It turns out that β will equal the ratio of the variance of the correctly predicted component of the real exchange rate innovations to the variance of the forward premium. This is a standard errors-in-variables interpretation of (23).[30] If agents act only on 'hard news' and nearly always interpret it correctly, β will approach unity; if they act on 'soft news' and, say, 10% of the variance of their forward premium settings is actually reflected in changes in the spot rate between this period and next, β will equal 0.10. The short-fall of β from unity is thus a measure of the quality of the information that is acted upon and the accuracy of agents' interpretation of that information.

But this ignores the possibility that agents may learn from their own forecast errors. If agents can fit equation (23) to their own forecasts and calculate the resulting values for β, they can scale their forecasts so that, on average, they will not underpredict future movements in the spot rate. The estimated β values for the scaled forecasts will then be unity and equations (22) and (23) describing the behaviour of agents in the aggregate will then both yield unitary estimates of β.

There are two reasons why the aggregate forecasts may not be optimal in the sense described above. First, some agents may take speculative positions only infrequently, on occasions when they 'think' they have good information. These agents may not have a sufficiently long forecast history to estimate their β-coefficients. Second, the forecast error process will not be the same at all points in time. Since different technological and other forces affect real exchange rates at different times, the errors agents make in predicting future exchange rate movements on the basis of the available 'information' may be different during some periods than others in ways that

[30] A brief discussion of the errors-in-variables idea is presented in Appendix C. For a further discussion, see G. S. Maddala, *Introduction to Econometrics*, New York: MacMillan, 1988, pages 380-382.

they have an insufficiently long forecast history to estimate. In a changing world, past forecast errors will not necessarily be a good predictor of future forecast errors. The estimated β in equation (23) will thus contain an historic average prediction error that will give little guidance in making current forecasts. When agents behave rationally, the estimated β will incorporate whatever optimisation they are able to effect.

Suppose that the major influence in the forward market is agents who adopt naive forecasts but that activist activity is nevertheless present on the part of a significant fringe of agents who from time to time think they have information worth acting on. Suppose further that, despite occasional mammoth profits by some individuals in this group, the forecasts of these activist agents are, on average, worthless. The random noise in the observed forward premium will generate a defined estimate of β in equation (23)—that estimate will be zero. The differences between the spot and forward exchange rates resulting from the activities of this second group of agents must be small in relation to the transactions costs facing the naive-forecasting agents—otherwise, the latter agents will profit from the ineptness of the former. In addition, of course, even if all agents view the exchange rate as a random walk some variance in the forward premium will result from the fact that small differences between spot and forward exchange rates resulting from imbalances of hedging pressure will not be worth taking uncovered positions to correct because of the risks and transactions costs involved.

9.3 The Role of Inflation Rate Differences

The role of the domestic and foreign expected inflation rates must now be incorporated into the argument. From the definition of the real exchange rate we obtained equation (18) which is repeated below for convenience:

$$E_\Pi \;=\; E_Q + E_{P^*/P} \;=\; E_Q + E_{P^*} - E_P. \tag{39}$$

Under the assumption that the real exchange rate is perceived by agents to be a random walk, E_Q will be zero and the forward premium will equal

$$\Phi \;=\; E_{P^*} - E_P - \phi, \tag{40}$$

the difference between the expected rates of domestic and foreign inflation between the periods t and $t+1$, minus the foreign exchange risk premium.

Casual knowledge about the stability of inflation rates might suggest that a simple projection of the current domestic/foreign inflation rate difference will provide an adequate inflation-component to the one-month forward premium, but this is not the case. For example, it was well-known in 1995 that the Canadian inflation rate was roughly zero and the U.S. inflation rate was around three percent. Suppose that the Canadian inflation rate for a particular month, calculated on an annual basis, was -0.11 percent—a seemingly trivial deviation from zero—and the U.S. inflation rate calculated on an annual basis was 3.3 percent, less than a third of a percentage point above its projected level. The inflation rate difference is -0.41 percent, a non-trivial magnitude. Being able to take into account inflation rate differences, and attempts to take them into account, could result in substantial deviations of forward exchange rates relative to current spot rates. Even though, as shown in Table 7, agents get the sign of the inflation rate differential correct in their forward forecasts, it is more difficult for them to adopt a naive forecasting rule for the future inflation rate differential than for the future *level* of the real exchange rate because the innovations in the inflation rate differential as a proportion of its current value are much greater than the innovations in the real exchange rate as a proportion of its current level.

The evidence presented in Figures 6 and 7 suggests that agents set the inflation component of the forward premium on the basis of past history with an adjustment based on 'news'. Since current 'facts' are subject to differing interpretations, random variability of the forward premium is inevitable. Even if every agent, realizing the insurmountable problems in forecasting real exchange rates, naively projects last period's real exchange rate forward, there will be variability in the forward premium arising from variation in forecasts of the inflation differential not reflected in future innovations in the spot exchange rate. Widespread belief in the near random-walk nature of real exchange rates together with unpredictability of real exchange rate innovations and inflation differentials will therefore result in estimated β-coefficients in equation (23) well below unity and quite possibly in the neighbourhood of zero. The very low variability of forward premia relative to the month-to-month innovations in nominal exchange rates results in the forward rate tracking the spot rate closely through time, with the result that the estimated β-coefficients in equation (22) should not be far from unity.

Nothing in the above analysis suggests that markets are inefficient in the sense that agents do not use all the information available to them. Agents' problem is rather that the information is imperfect and subject to differing interpretations and that their reactions in response to these interpretations depend upon their attitudes toward risk. They form their expectations rationally in the sense that they use, as best they can, all information available to them. They do not have rational expectations in the sense in which that term is often used in the technical literature to mean that agents act rationally *and* fully understand how the economy functions *and* observe the shocks affecting it. Here is is assumed that agents do not have any information about the shocks affecting the real exchange rate and that they have very limited information about the shocks to the domestic and foreign inflation rates.

The treatment of the real exchange rate as the relative price of domestic in terms of foreign output *as well as* an asset price leads to a major change in the way the β-coefficients in equations (22) and (23) are interpreted. This view of the real exchange rate suggests that agents' expectations about future exchange rates will be dominated by a naive projection of current rates while the expectations generation process in the traditional asset theoretic approach implies that, on average, agents' expectations of next period's nominal exchange rate will equal the actual realization. This latter specification is not an unreasonable one in a world where all movements of the real exchange rate around some purchasing power parity level represent asset-market-related disequilibrating shocks that agents should be expected to know something about. But it imposes a heavy load on the unbiasedness hypothesis in reconciling the estimates of β in equations (22) and (23) in a world where real exchange rate innovations are a reflection of on-going changes in the world equilibrium relative price structure together with intertemporal consumption smoothing in response to world business cycles. Once we take account of the fact that we know very little about the process generating *equilibrium* real exchange rates and recognise that agents know no more than economists, a framework more congenial to understanding the reason why the unbiasedness hypothesis fails emerges. We can now easily explain why estimates of β in equation (23) are less than unity—we would not expect them to be much different from zero. Our problem is to explain why they tend to be negative.

9.4 Why are Estimates of β Negative?

We must first note from Table 11 that for France vs. the U.S. the coefficients of the lagged forward premium are negative but everywhere insignificant—the F-statistics for the regressions indicate no relationship. And for Germany, the U.K. and Japan vs. the U.S., statistically significant negative coefficients arise entirely as a result of the relationships among the variables during the period 1978 through mid-1990.[31] For Canada vs. the U.S., significant negative coefficients are found for the period 1974 through 2003, for

[31] The exact boundaries of this period were chosen to conform to the period analysed by Bennett McCallum in 'A Reconsideration of the Uncovered Interest Parity Relationship," *Journal of Monetary Economics*, Vol. 33, No. 1 (February), 1994, 105–132.

the period from 1978 through mid-1990 and, in the case where the actual forward premium rather then the implicit one based on interest rate differentials is used, for the period from mid-1990 to the end of 2003. The coefficient is positive but insignificant for the Canada vs. U.S. case for the period 1951 through 1973. It is also positive and insignificant for the U.K. vs. the U.S. for the period following mid-1990. It is also noteworthy that for all countries vs. the U.S. the magnitudes of the coefficients are much smaller when the coefficients are statistically insignificant than when they are significant. In summary, with the exception of Canada, significant negative coefficients arise entirely as a result of developments in the late 1970s and the 1980s. Canada's real exchange rates with respect to the U.S. follows an obvious trend, downward since the mid-1970s.

As can be seen from Figures 1 and 2, the period between 1978 and 1991 was one in which the U.S. real exchange rates with respect to the other countries examined rose very substantially—by 50% or more—and then fell back to near beginning levels. It was also a period during which world inflation rates fell very substantially as can be seen from the centre panels of Figures 10 through 14.

Finally, it must be noted that the R^2s in Table 11 for the cases in which significant relationships exist are typically between .03 and .06 and only in the U.K. case do they reach .10. This is evident from the scatter plots shown in Figures 16 and 17.

Negative estimates of β might suggest that agents' predictions are not merely inaccurate, but are dominated by agents who get the sign wrong. Following the line of argument developed by De Long, Shleifer, Summers and Waldman[32] we could make a distinction between arbitragers and noise traders, explaining negative observed β-coefficients as the result of market activity dominated by poorly informed agents who gamble on the basis of their assessment of existing information, getting the sign wrong more than half the time but occasionally in individual cases making enormous profits. The question is then why it does not pay more professional and better-informed agents, called arbitragers, to eliminate this perverse relationship between the forward premium and next period's innovation in the spot rate

[32] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldman, 'The Size and Incidence of Losses from Noise Trading," *Journal of Finance*, Vol. 44, No. 3 (July) 1989, 681–696, 'Noise Trader Risk in Financial Markets," *Journal of Political Economy*, Vol. 98, No. 4 (August), 1990, and 'Positive Feedback Investment Strategies and Destabilising Rational Speculation," *Journal of Finance*, 1990. For a useful summary of these ideas, see Andrei Shleifer and Lawrence H. Summers, 'The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, Vol. 4, No. 2 (Spring), 1990, 19–33.

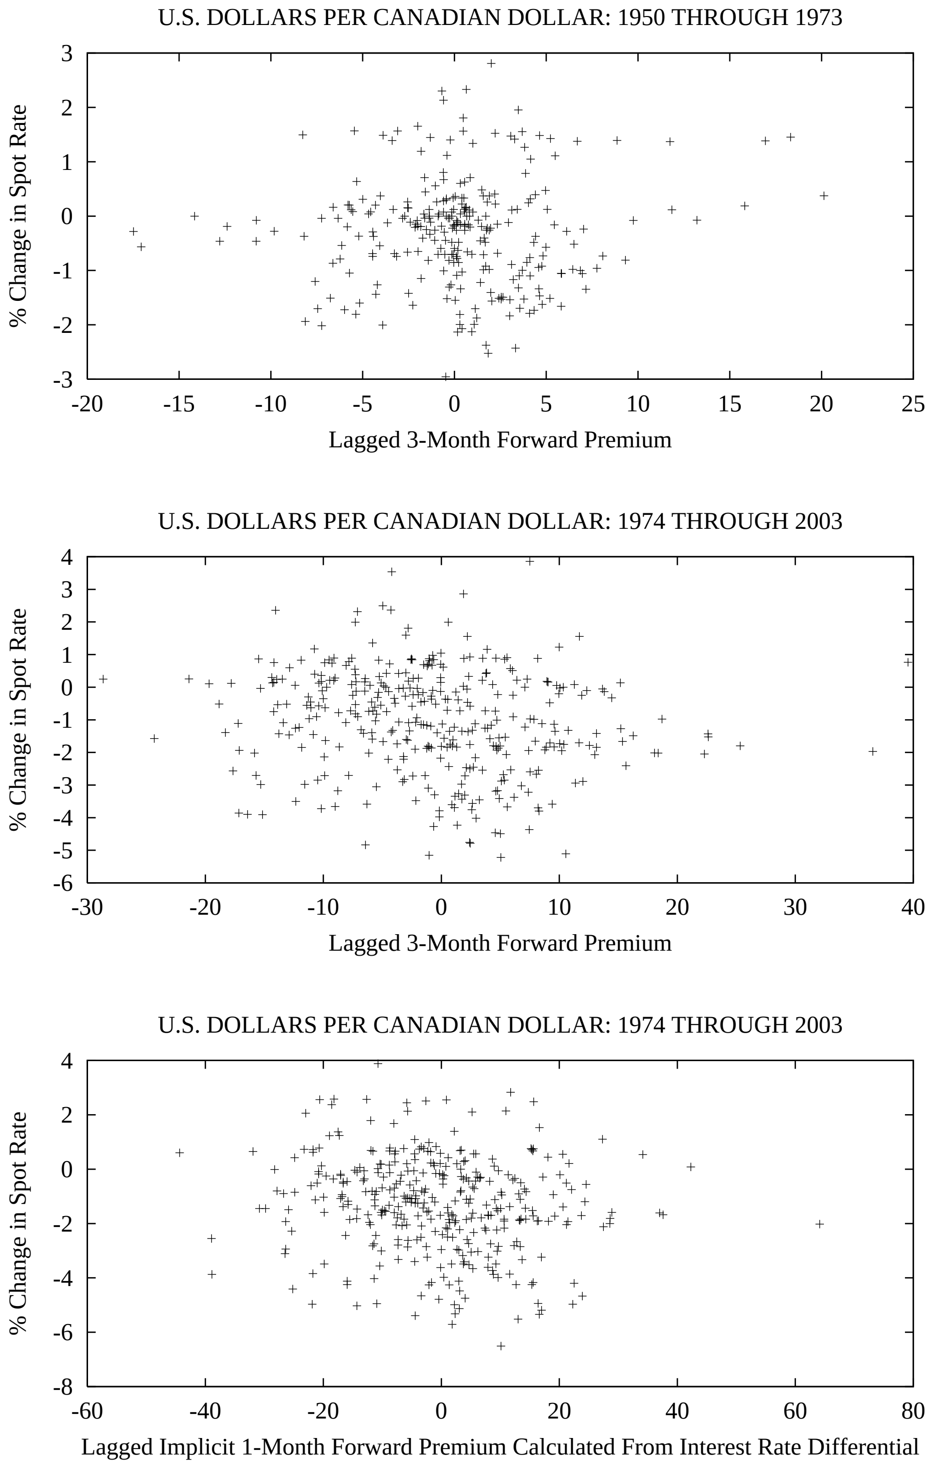

Figure 16: Percentage change in the spot rate vs. the change predicted by lagged 3-month and 1-month forward premia, United States dollars per Canadian dollar. Source: *Cansim*.

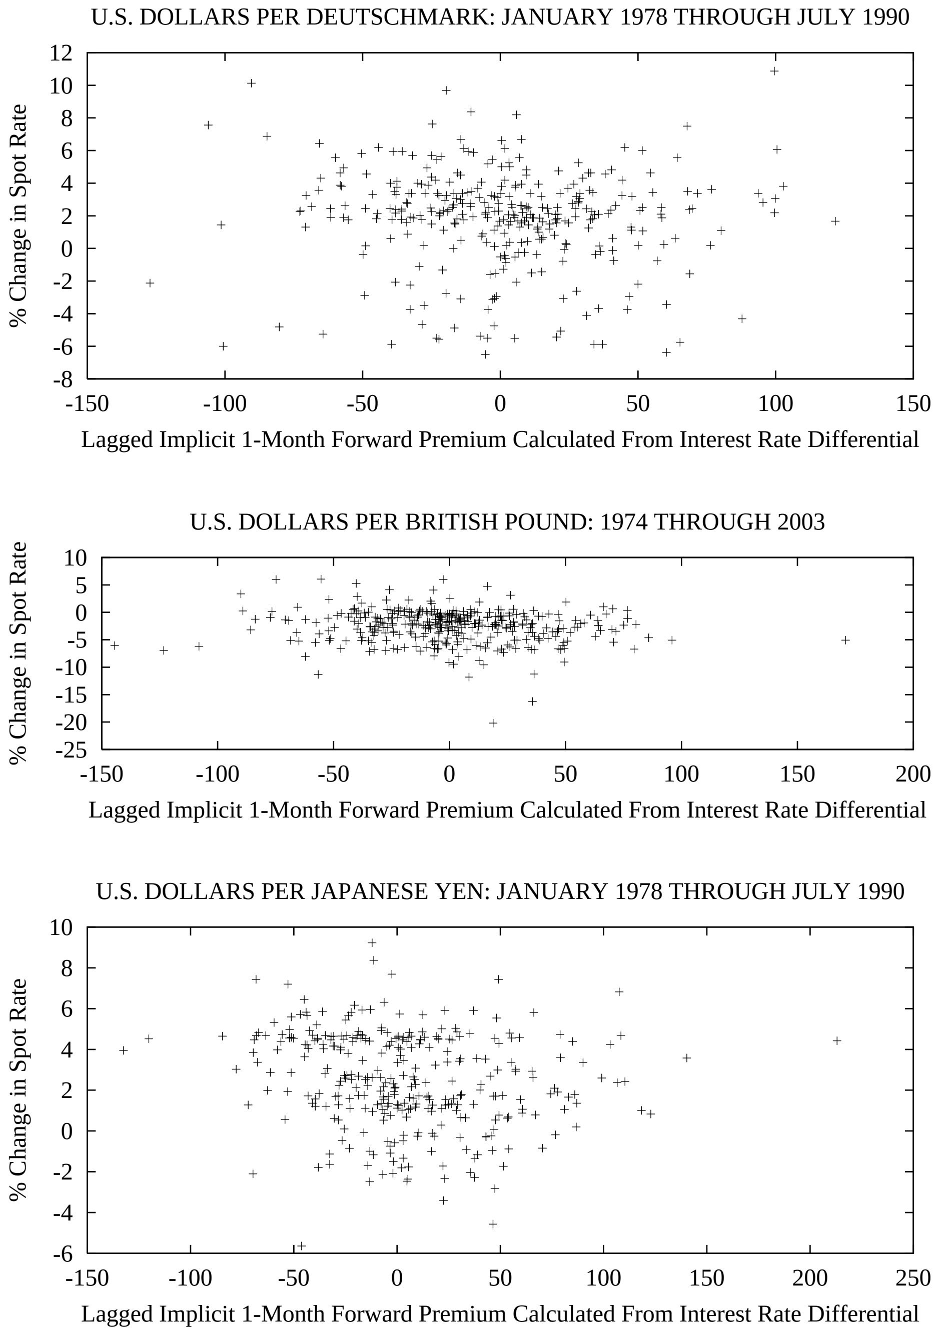

Figure 17: Percentage change in the spot rate vs. that predicted by 1-month forward premia, United States dollars per deutschmark, French franc and Japanese Yen. Source: *Reuters* and *Datastream*.

by continually taking forward positions on the basis of naive random-walk forecasts. The distinction between noise traders and arbitragers here is not a rigid one—one could think of a continuum ranging from pure noise traders at one end to pure arbitragers at the other, with the positions of particular traders along the continuum depending upon the quality of information they are able to obtain and willing to act upon and their attitudes toward risk. A look at the scatter plots in Figures 16 and 17 and the R^2s in Table 11, combined with the fact that, apart from a sub-period of the sample, the relationship between changes in the spot rate and lagged forward rates is not statistically significant, and not always negative, suggests that it is unlikely that an arbitrage-based trading rule could lead to a significant profit.

We have no difficulty explaining why the coefficient of the forward premium in a regression with the percent change in the spot rate as the dependent variable is less than unity—we have little reason to expect it to be bigger than zero. Moreover, the fact that forward premia are very small relative to percentage changes in the spot rate explains the tendency of the β-coefficient in a regression of the logarithm of the spot rate on the logarithm of the forward rate to be close to but less than unity—every period, the forward rate rises to approximately the level of the spot rate that period. Accordingly, as can be seen from the top panel of Figure 4, there can be no doubt that the coefficient will be near unity even when the coefficient in a regression of the change in the spot rate on the forward premium turns out to be negative. But we are unable to explain the preponderance of, albeit insignificant, negative coefficients in the forward premium regressions, even though we have little reason to expect them to be significantly or predominantly positive.

10 Summary and Conclusions

You should have learned a number of things from working carefully through this paper:

- Real and nominal exchange rates tend to move in step with each other, differing for the most part only in trend and showing substantially more variability than the ratios of countries' price levels, which tend to be smooth, trend-like series.

- Real exchanges rates appear as near random-walks—the hypothesis of stationarity can usually be rejected using low power tests on short sample periods but there is evidence of mean reversion over periods as long as

45 years using monthly data, and over more than a century using annual data.

- The weight of the evidence on predicting exchange rates is that, for practical purposes, the best forecast of tomorrow's exchange rate is today's rate adjusted for ongoing inflation rate differences.

- Spot and forward exchange rates move so closely together that they can barely be distinguished from each other when plotted. Yet the relative first difference of the spot rate has a variance many times greater than the variance of the forward premium—the spot rate typically moves by much more than the forward premium would predict.

- Exchange rates are asset prices in that they represent the prices of uncovered foreign-currency positions and exchange rate changes lead to capital gains on foreign exchange reserve positions. At the same time, real exchange rates are the relative prices of different countries' outputs and vary in response to technological and political developments in different parts of the world. In this context nominal exchange rates reflect movements in real exchange rates in conjunction with differences in the countries' inflation rates.

- There is reason to believe that the risk associated with uncovered foreign exchange positions is, like that on other assets, determined by the covariance of the asset's return with the return to capital in the economy as a whole—all the remaining variance can be diversified away. Actual measurement along these lines is stymied, however, by difficulties in measuring the return to capital in the economy as a whole in the presence of non-traded human capital and by difficulties in understanding how portfolios are diversified in the presence of swings in real exchange rates.

- The purchasing power parity theory, which postulates that real exchange rates are constant, is inconsistent with the data on real exchange rates although it is often argued that the theory holds in the long-run with observed violations in the short run due to disequilibria resulting from price level stickiness. This interpretation of real exchange rate movements as disequilibrium deviations from purchasing power parity is neither useful, necessary nor correct in explaining the observed data—the data are quite consistent with constant variation of equilibrium real exchange rates as a result of ongoing technological change and economic growth.

- Covered interest parity holds to a reasonable approximation. This implies, if agents are behaving rationally, that security specific or country specific risk premia unrelated to nominal exchange rate movements tend to be small.

- While forward exchange rates tend to be good predictors of future spot rates because the two move so closely together, forward premia tend to be poor, and often perverse, predictors of the future movements in spot rates. This means that uncovered interest parity does not hold—current domestic/foreign interest rate differentials (which approximately equal the forward premia) do not explain future movements in the exchange rate, except in situations where there are substantial permanent differences between the domestic and foreign inflation rates.

- Conventional asset-theoretic views of exchange rate changes as deviations from long-run purchasing power parity equilibrium hold that, barring time varying risk premia or violations of market efficiency, the forward premium should bear a one-to-one relationship, on average, to realized future movements of the spot rate. When real exchange rate movements are viewed as changes in equilibrium levels in response to ongoing world technological change, economic growth, and political developments rather than as disequilibrium departures from some purchasing power parity level, less stringent demands are placed on the explanatory power of the efficient markets hypothesis. Asset market efficiency is fully consistent with a zero correlation between the forward premium and the future change in the spot rate when real exchange rates are viewed by agents as near random walks and the countries' inflation rates do not differ substantially.

- Across a wide range of studies the data indicate a weak negative (perverse) correlation between the forward premium and future movements in the spot rate that, by and large, is statistically significant only for a sub-period surrounding the 1980s. This negative relationship, though as yet unexplainable, is sufficiently weak that there is reason to doubt that above normal profits could be made by taking advantage of the known superiority of naive inflation-adjusted current spot-rate forecasts.

Appendix A: Data Sources

Annual Data

The exchange rate of the U.K. pound with respect to the U.S. dollar was obtained back to 1803 from L.H. Officer, "Dollar-Sterling Mint Parity and Exchange Rates, 1791-1834," *Journal of Economic History*, Vol. 43, No. 3 (September) 1983, 579-616. For the years 1834 to 1868 it was obtained from B.O. Michell, *European Historical Statistics, 1750-1975*, Cambridge University Press, 1981, and for the period 1869 to 1975 from Milton Friedman and Anna J. Schwartz, *Monetary Trends in the United States and the United Kingdom*, NBER, 1982. For the years since 1975 the data were obtained from the International Monetary Fund, *International Financial Statistics*, series 112/RF. Annual Canadian nominal exchange rate data were obtained back to 1871 from M.C. Urquhart, K.A.H. Buckley, and F.H. Leacy, *Historical Statistics of Canada*, Statistics Canada and Social Science Federation of Canada, 2nd Edition, 1983, and for recent years from *International Financial Statistics*, series 156/RF.

The consumer price index series for Canada was obtained back to 1873 from *Historical Statistics of Canada* and for recent years from *International Financial Statistics*, series 156/64. For the United States, the consumer price index was obtained from *Historical Statistics of the United States*, U.S. Department of Commerce, 1975, and from *International Financial Statistics*, series 112/64. A series for the United States GDP deflator back to 1803 was obtained by my late colleague Trevor J.O. Dick from Thomas Senior Berry, *Production and Population Since 1789: Revised GNP Series in Constant Dollars*, Bostwick Paper No. 6, The Bostwick Press, Richmond Virginia, 1988. For the years 1869-1975 the series was obtained from the Friedman and Schwartz volume noted above, and from 1975 onward, series 111/99A.R from *International Monetary Statistics* was used. For the United Kingdom, a GDP deflator series back to 1803 was obtained, again by Trevor Dick, from Simon Kuznets, *Secular Movements in Production and Prices*, Houghton Mifflin, New York, 1930. These data were based on retail price indices. To 1975 the U.S. GDP deflator series was obtained from the above Friedman and Schwartz volume and from 1975 to the present, *International Financial Statistics* series 112/99A.R was used.

From these data, the Canada to U.S. and U.K. to U.S. price level ratios were constructed by taking the ratios of the respective CPI price indexes for Canada/U.S. and the GDP deflator series for U.K./U.S. And the real exchange rates were then constructed by multiplying these price level ratios by the U.S. dollar prices of the Canadian dollar and British pound, respectively. The real exchange rate and price level ratio series were put on a 1950=100 base. Nominal exchange rate series, indexed on a 1950=100 base, for the the Canadian dollar and British pound in terms of the U.S. dollar were also created.

All these annual data are contained in the Rats data file `anndata.rat`, in the Lisp file `anndata.lsp` and in the spreadsheet files `anndata.xls` and `anndata.wk1`. These files, along with a text file `anndata.cat` that catalogs the descriptors in `anndata.rat` can be obtained from the Internet at a location noted below. These files also contain real GDP series for Canada, obtained from *Historical Statistics of Canada* and *International Financial Statistics*, series 156/99B.R, and for the United States and United Kingdom, obtained historically from the same sources as the GDP deflators and from *International Financial Statistics* series 112/99B and 111/99B. In addition, real GDP growth rate and inflation rate series for the three countries, computed from the above-mentioned series, are included in these files.

Quarterly Data

Quarterly exchange rate series giving the domestic currency prices of the U.S. dollar were obtained from the International Monetary Fund, *International Financial Statistics* for the countries studied here. The series are 156/RF for Canada, 132/RF for France, 134/RF for Germany, 112/RH for the U.K. and 158/RH for Japan. Consumer price index series for the above countries plus the U.S. were obtained from the same source—111/64 for the U.S. and 156/64, 132/64, 134/64, 112/64 and 158/64—the first number being the country code used in obtaining the exchange rate series, and the number 64 being the mnemonic for the CPI. All CPI series were put on a 1985=100 base. The nominal GDP series for the U.K. was obtained over the Internet from the *U.K. Office of National Statistics*. Nominal GDP series for the other countries were obtained from *International Financial Statistics*—the series mnemonic is 99B.C. All the GDP series are seasonally adjusted. And all the above series, exchange rate, CPI and GDP, run from 1957Q1 to 2002Q4.

The above series are contained in the Rats file `qrtdata1.rat` and the spreadsheet files `qrtdata1.xls` and `qrtdata1.wk1`. They are also spread over a number of Lisp files that are collected together in the self-extracting zip file `qlspdat1.exe` and the tar-zip file `qlspdat1.tar.gz`. All these files, including a catalog text file `qrtdata.cat` containing the descriptors in the file `qrtdata.rat`, can be obtained from the Internet location noted below.

The files also contain real exchange rate series with respect to the U.S. for the five other countries—these are calculated as the domestic CPI divided by the product of the corresponding exchange rate (domestic currency price of the U.S. dollar) and the U.S. CPI . Like the CPI series, these were also put on a 1985=100 base. In addition, real GDP series in billions of 1985 home currency units for each of the six countries are included in the data files. These series were obtained by deflating the respective countries' nominal GDP figures by their CPIs.

Finally, series giving the percentage errors, actual minus predicted, from forecasts using the previous periods' forward exchange rates with respect to the U.S. dollar are included for each of the five other countries. These are quarterly averages of the monthly figures described in the monthly data below.

Monthly Data

Consumer price indexes, monthly from 1957 to the end of 2002, were obtained for all six countries from *International Financial Statistics*. The mnemonics are those given for the quarterly series, which are quarterly averages of the monthly data.

A series for the spot price of the Canadian dollar in terms of U.S. dollars was obtained back to November 1950 from *Cansim II*, series V47426. The corresponding 90-Day forward exchange rate series was obtained from *Cansim I*, series B3401, through 1999 and from the Bank of Canada web-site to February 2004. Also from *Cansim II*, 1-month and 3-month Canadian corporate paper rates, series V122509 and V122491 respectively, were obtained for the period from 1956 to February 2004. Comparable corporate paper rate series for the United States were also available from *Cansim II* back to April 1971, series V122144, for the 1-month rate and to 1962, series V122141, for the 3-month rate.

For the monthly analysis pertaining to the years 1974 through 2003, exchange rates obtained by my colleague Alex Maynard for all six countries from *Reuters* were used for the years through February 1999—the mnemonics for Canada, France, Germany, Japan and the U.K. are, respectively, CAD, FRF, DEM, JPY and GBP for the spot rates and CAD1M, FRF1M, DEM1M, JPY1M and GBP1M for the 1-month forward rates. For the period since February 1999 the data were obtained from *Datastream*—the mnemonics are CDNDLUS, FRNFRUS, WGMRKUS, JAPYNUS and BRITPUS for the spot rates and CD30DUS, FF30DUS, WG30DUS, JP30DUS and BP30DUS for the 1-Month (30-Day) forward rates. Alternate exchange rate series for the Deutschmark, pound and yen were constructed using the data published in the paper by Bennett McCallum, "A Reconsideration of the Uncovered Interest Parity Relationship," *Journal of Monetary Economics*, Vol. 33, No. 1 (February) 1994, 105-132, supplemented beyond July 1990 with series from *Datastream* different from the ones accessed above—namely, BBDEMSP and BBJPYSP for the spot rates on the Deutschmark and the yen, and USWG30D, USBP30D and BBJPJPY1F for the forward rates on the Deutschmark, pound and yen. These alternative series turned out to contain more noise in the later years than the previous ones noted above. Because of the erratic variability of the forward rates, Euro-currency deposit rates were also used to estimate the relevant forward premia. These were obtained from *Reuters*, again by Alex Maynard, for the period prior to March 1999 and from *Datastream* thereafter. The *Reuters* series are USD1MD, FRF1MD, DEM1MD, JPY1MD, and GBP1MD, and the *Datastream* series are GSUSD1M, GSFR1M, GSDEM1M, ECJAP1M and GSGBP1M—the country associated with each series should be obvious. The nominal exchange rates obtained quarterly from *International Financial Statistics* were also obtained monthly—the mnemonics are the same as for the quarterly data. These and the CPI series can be used to calculate the series plotted in Figures 1 and 2. They were also used to create real exchange rate series REXMCAUS, REXMFRUS, REXMGRUS, REXMJNUS, REXMUKUS and REXMFRGR, which are expressed as percentage deviations from mean values.

Additional interest rate series were obtained from *International Financial Statistics* for the U.S., Germany, Japan and the U.K.—the series mnemonics are 60B, the call money rate, for Japan and 60C, treasury bill yields, for the other three countries.

All the above monthly series are contained in `mondata1.rat`, and the worksheet files `mondata1.wk1` and `mondata1.xls`. These series are spread over a number of lisp files which are collected together in the self-extracting zip file `mlspdat1.exe` and the the tar-zip file `mlspdat1.tar.gz`. The contents of these included files should be obvious from the file names. A catalog text file `mondata1.cat` based on the descriptors in `mondata1.rat` is also included. The data files also contain implicit forward premium series constructed from the interest rate differentials as well as the 1-month and 3-month commercial paper rate differentials for Canada vs. the U.S. Month-over-month and year-over-year inflation rate differentials with respect to the U.S. for all five other countries are also included.

Finally the percentage errors from using the current spot and forward rates to forecast next period's spot rate are calculated from the spot exchange rate and interest rate differential series for all currencies with respect to the U.S. dollar. Forward predictions of each exchange rate, calculated as an intermediate step in computing the forward rate forecast errors, are also included in the files. All the forecast errors are based on actual percentage differences, not logarithmic differences. It may be desirable to alternatively use logarithmic differences instead of actual percentage differences to calculate the forward premia. Such series, while not included in the data files, can be calculated using included data.

Statistical Program Files

All of the above data files can be obtained from my web-site at

`www.economics.utoronto.ca/floyd/simprex.html`

along with the statistical program files used to analyse these data. The statistical results were obtained using, alternatively, three different statistical programs, Rats, Ox, and X-LispStat. Rats is a commercial program that can be purchased at

`www.estima.com`.

Ox is also a commercial program but a console version is freely available for academic use. It can be obtained at

`www.doornik.com/download.html`

and an introductory manual can be obtained at

`www.doornik.com/ox/OxIntro.pdf`.

To use the console version of Ox for the statistical analysis here you will need a set of functions I wrote for that purpose contained in the file `newfuncs.ox` which is in the zip file noted below containing the code and output files.

X-LispStat is a completely free platform for statistical computing written by Luke Tierney, a professor of statistics at the University of Minnesota. It can be obtained, along with a minimal manual that I wrote for beginning students, from my home page at

`www.economics.utoronto.ca/floyd/intstat.html` or by itself from

`www.stat.uiowa.edu/~luke/xls/xlsinfo/xlsinfo.html`.

To do sophisticated work with X-LispStat you would need to consult Luke Tierney's book.[33] For the analysis here the book should not be necessary, but you will need a set of basic functions I wrote for econometric work. These are contained in the file `newfuncs.lsp` which is made available in the same zip file as the program code and output.

After downloading the relevant data files from my web-site, download the code files and resulting output files created using the above statistical programs. These files are contained in the zip files `ratfile1.exe`, `oxfile1.exe` and `lspfile1.exe` and in the alternative set of tar-zip files `ratfile1.tar.gz`, `oxfile1.tar.gz` and `lspfile1.tar.gz`. The required function files for X-LispStat and Ox, `newfuncs.lsp` and `newfuncs.ox`, are zipped in with these code and output files. You can look through the statistical results by examining the output files without actually downloading the programs to your computer. The easiest output files for the general reader to understand are those created with XLispStat. Of course, you will need the appropriate program to run the code files. The output files will be recognised by their `.rou` (RATS), `.oou` (Ox) and `.lou` (X-LispStat) file extensions. The program code files have the respective extensions `.prg`, `.ox` and `lsp`. The statistical calculations and results for Table 1 through Table 6 can be found in X-LispStat and Ox files with the root name `unitroot`. None of these calculations were done with RATS because the Dickey-Fuller and Phillips-Perron unitroot-test functions I wrote for the other two programs are more complete than the ones available with RATS. Also, the calculations and results for Table 7 and Table 9, found in the files with the root name `fprmanal`, are more simply done in Ox and X-LispStat than RATS so the latter program was not used for these calculations. Program code and output files are present for all three statistical programs for the remaining tables. For Table 8 the calculations and results are in files with the root name `fprmregs`. Files with the root name `sfxrregs` contain the program code and results for Table 10 and Table 11. Finally, the materials for Table 12 and Table 13 are contained in the files with the root name `corrstdq`.

---

[33] Luke Tierney, *Lisp-Stat: An Object-Oriented Environment for Statistical Computing and Dynamic Graphics*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, 1990.

Appendix B: Some Basics of Time Series Analysis

An excellent place to learn about time series analysis is from Walter Enders' book[34] For purposes of understanding what is going on here, it is recommended that you read from page 211 to page 225. The discussion below is to introduce you to enough of the basic principles to allow you to begin reading at page 211.

Time Series Processes

A time series is a series of numbers indexed by time that portrays the time path of some variable. It is often convenient to imagine that these numbers are generated by some mathematical process. For example the series y_t might be generated by the equation

$$y_t \quad = \quad a\, y_{t-1} + \epsilon_t. \tag{41}$$

This is a first order autoregressive process—first order because there is only one lag of y_t on the right-hand side and autoregressive because the y_t are autocorrelated in the sense that the level of the variable in each period depends on its level in a previous period. The term ϵ_t is a white noise process, defined as a series of drawings of a zero-mean, constant-variance non-autocorrelated random variable.

Lagging (41) repeatedly, we obtain

$$\begin{aligned} y_{t-1} &= a\, y_{t-2} + \epsilon_{t-1} \\ y_{t-2} &= a\, y_{t-3} + \epsilon_{t-2} \\ y_{t-3} &= a\, y_{t-4} + \epsilon_{t-3} \\ \cdots &= \cdots \\ y_1 &= a\, y_0 + \epsilon_1. \end{aligned}$$

Repeated substitution then yields

$$\begin{aligned} y_t \quad = \quad & a^t\, y_0 + \epsilon_t + a\, \epsilon_{t-1} + a^2 \epsilon_{t-2} + a^3\, \epsilon_{t-3} + \cdots \\ & + a^4\, \epsilon_{t-4} + \cdots\cdots\cdots + a^{t-1}\, \epsilon_1 \end{aligned} \tag{42}$$

The time path of y_t depends critically on the parameter a. If this parameter equals zero then

$$y_t \quad = \quad \epsilon_t. \tag{43}$$

[34] Walter Enders *Applied Econometric Time Series*, John Wiley and Sons, 1995.

and y_t is itself a white noise process. The variance of y_t will equal the variance of ϵ_t which we will denote by σ^2. If $a = 1$, y_t becomes (utilising (41) and (42))

$$\begin{aligned} y_t &= y_{t-1} + \epsilon_t \\ &= \epsilon_t + \epsilon_{t-1} + \epsilon_{t-2} + \cdots + \epsilon_1 + y_0. \end{aligned} \tag{44}$$

The series is a random walk. It wanders without limit, with y_t moving either up or down in each period relative to its previous value y_{t-1} by some random amount ϵ_t, as can be seen by rewriting (41) as

$$y_t - y_{t-1} = \epsilon_t. \tag{45}$$

The variance of y_t will equal $(\sigma^2 + \sigma^2 + \sigma^2 + \cdots\cdots)$,[35] which will grow in proportion to the number of periods over which the variance is being calculated.

In the case where $a = 1$ the series is said to be non-stationary or have a unit root (the root of (41) equals a). Its expected level at any point in time is its current level and its variance in the limit is infinity. Its future path need never pass through the level at which it started or any other level previously achieved, although there is no reason why it could not return to these levels. When $a > 1$ the series explodes, with the values of $|y_t|$ getting larger and larger with time. This can be seen from the fact that a^t will get bigger and bigger as t increases when $a > 1$. If a is negative the series oscillates around zero, doing so explosively if $a < -1$.

When $-1 < a < 1$ the series is stationary as can be seen from the fact that a^t gets smaller in (42) as t increases. If the ϵ_t are zero beyond some point, y_t will approach zero as t increases, with the speed of approach being greater, the smaller is a. The effects of each ϵ_t shock will thus dissipate with time. The variance of y_t will equal $[1 + (a)^2 + (a^2)^2 + (a^3)^2 + \cdots\cdots\cdots]\,\sigma^2$ which will be finite in the limit as t increases. In the case were $a = 0.9$, for example, this variance will equal $5.26\sigma^2$ in the limit.[36] The series will vary around zero with a persistence that will be greater the greater is a.

[35] This follows from the fact that $Var\{x + y\} = Var\{x\} + Var\{y\}$ when x and y are uncorrelated variables.

[36] Here we use the relationship in the previous footnote plus the facts that $Var\{a\,x\} = a^2\,Var\{x\}$ and

$$1 + b + b^2 + b^3 + b^4 + \cdots\cdots = \frac{1}{1-b},$$

and then set b equal to a^2.

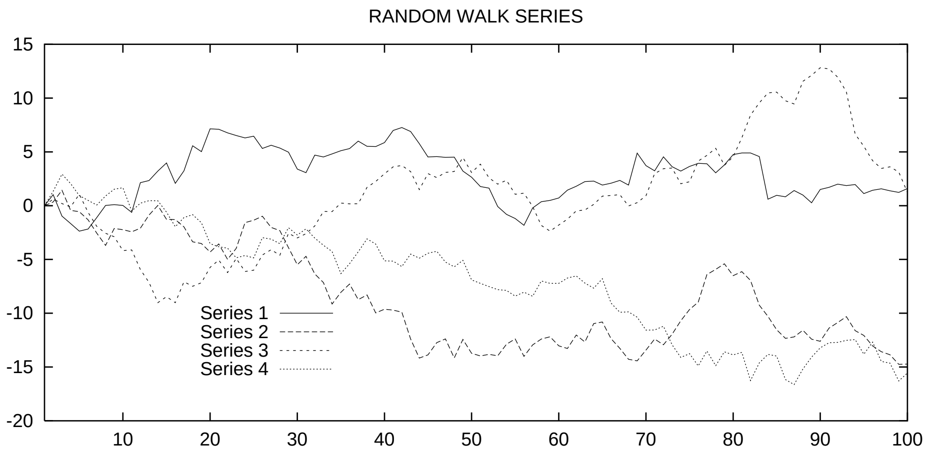

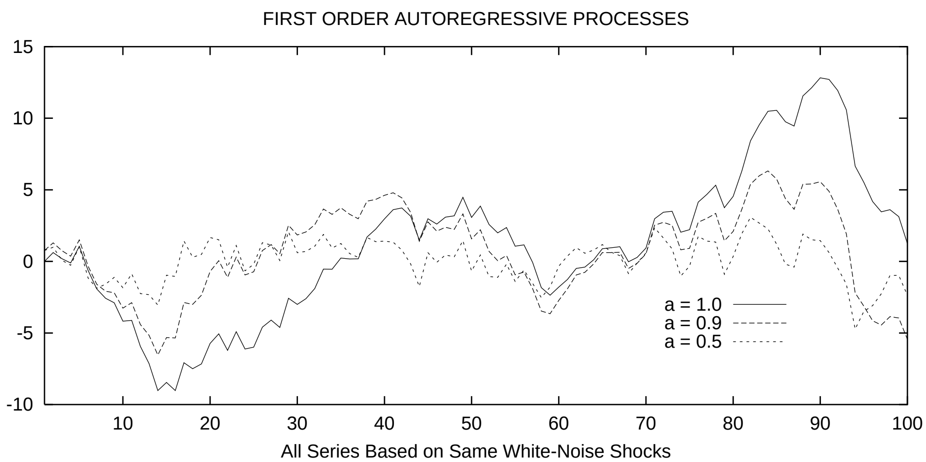

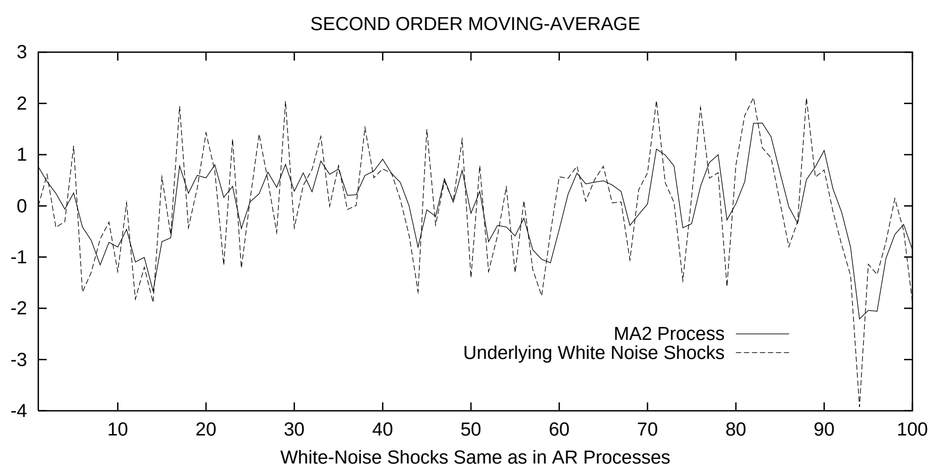

Figure B1: Random-walk, Autoregressive and Moving Average Processes: Some Examples.

Several examples of non-stationary random-walk processes with $a = 1$ are plotted in the top panel of Figure B1. The middle panel plots a single random-walk process along with two stationary processes based on the same ϵ_t series, having $a = 0.9$ and $a = 0.5$ respectively.

Testing for Stationarity

Tests for stationarity involve determining whether parameters like a in (41) are statistically significantly less than unity in absolute value. The standard procedure is to subtract y_{t-1} from both sides of (41) to obtain

$$\begin{aligned} y_t - y_{t-1} &= -(1-a)\, y_{t-1} + \epsilon_t. \\ &= \delta\, y_{t-1} + \epsilon_t. \end{aligned} \tag{46}$$

where $\delta = -(1 - a)$ and then test whether δ is less than zero (which is the same as testing whether $a < 1$).[37] This test involves simply running the least-squares regression indicated by (46) and examining the t-statistic of the coefficient of y_{t-1}. The problem here, as you will note from reading Enders, is that under the null-hypothesis that $\delta = 0$ the estimator of δ is not distributed according to the $t-$distribution. One must use the table of critical values constructed by David Dickey and Wayne Fuller instead of the standard t-tables.

A major problem here is that the test procedure just outlined has poor ability—statisticians use the term low power—to detect stationarity when the true value of δ is negative and close to zero (a is less than but close to unity). It is easy to see why. When we test the null-hypothesis that δ equals zero we, in effect, use the Dickey-Fuller table to determine an appropriate critical value of $\hat{\delta}$, the estimated value of δ. This critical value will be the relevant entry in the Dickey-Fuller table multiplied by the standard error of $\hat{\delta}$. It will be some negative number $\hat{\delta}_1$ below which $\hat{\delta}$ has some small probability, say .05, of lying if δ is really zero. So if δ in fact equals zero there is only a 5% chance that we will reject the null-hypothesis of non-stationarity and conclude that y_t is stationary. This means that there is a 95% chance that we will conclude that y_t is non-stationary. Now suppose that a equals .999999 so that the true value of δ is -.000001. This means that y_t is in fact stationary. But application of the test will nevertheless lead us to conclude that it is non-stationary almost 95% of the time because $\hat{\delta}$ will still fall below $\hat{\delta}_1$ only very slightly more than 5% of the time. If a is

[37] Note that the symbol δ reused here with a different meaning than in the main text.

0.8 or 0.9 and the true value of δ is therefore -0.1 or -0.2, the estimate $\hat{\delta}$ will still lie above $\hat{\delta}_1$ a high percentage of the time leading us to conclude that y_t is non-stationary when in fact it is stationary. So we run a small risk, 5% in the example above, of concluding that y_t is stationary when it is not, and a very high risk of concluding that y_t is non-stationary when it is stationary at true values of a not far below unity. These tests must therefore be viewed with caution.

ARIMA Processes

Equation (41) is a first-order autoregressive process. A second-order autoregressive process would be represented by

$$y_t = a_1 \, y_{t-1} + a_2 \, y_{t-2} + \epsilon_t, \tag{47}$$

with two lags of y_t, and third and higher order processes can be similarly defined.

Time series are not all autoregressive processes. They can be moving average processes. An example would be an equation generating y_t of the form

$$y_t = b_0 \, \epsilon_t + b_1 \, \epsilon_{t-1} + b_2 \, \epsilon_{t-2} \tag{48}$$

which is a second-order moving average process—second-order because it contains two lagged error terms. A second-order moving average process with $b_0 = 0.4$, $b_1 = 0.3$, and $b_2 = 0.3$ is presented in the bottom panel of Figure B1 along with the ϵ_t process used to generate it. This white-noise process is the same one that was used to generate the three autoregressive processes in the middle panel of Figure B1.

Time series can also be combined autoregressive and moving average (ARMA) processes. Consider the equation

$$y_t = a_1 \, y_{t-1} + a_2 \, y_{t-2} + b_0 \, \epsilon_t + b_1 \, \epsilon_{t-1} + b_2 \, \epsilon_{t-2}. \tag{49}$$

This defines an ARMA(2,2) process—a process that is second-order autoregressive and second-order moving average. In general, an ARMA(p, q) process has p autoregressive lags and q moving average lags.

We can go a step further and assume that y_t above is a stationary process that is actually the first difference of another series z_t—i.e., $y_t = z_t - z_{t-1}$— and that the process z_t is a non-stationary autoregressive-moving average process that has to be differenced once to produce the stationary ARMA(2,2) process. It is said to be integrated of order 1 because it has to be differenced once to produce a stationary process. If it had to be differenced twice to produce a stationary process it would be integrated of order 2, and so forth. The process z_t is thus an autoregressive-integrated-moving average ARIMA(2,1,2) process—differencing it once produces an ARMA(2,2). In general, an ARIMA(p, d, q) process is one whose d-th difference is a stationary autoregressive-moving average process with p autoregressive lags and q moving average lags.

Autocorrelation and Partial Autocorrelation Functions

The autocorrelation function of a time-series process is the set of correlations between the t-th term and each of n $(t$-$j)$-th terms, where $j = 1, 2, \cdots n$. The autocorrelation function for any of the above processes consists of the set of correlations

$$\begin{gathered} r\{y_t, y_{t-1}\} \\ r\{y_t, y_{t-2}\} \\ r\{y_t, y_{t-3}\} \\ r\{y_t, y_{t-4}\} \\ r\{y_t, y_{t-5}\} \\ r\{y_t, y_{t-6}\} \\ r\{y_t, y_{t-7}\} \\ r\{y_t, y_{t-8}\} \end{gathered}$$

where $r\{x, y\}$ is the simple correlation coefficient between x and y and n is set equal to 8. One can calculate this autocorrelation function by creating eight lags of the data series y_t, with each lag being a separate series, and then calculating the simple correlations of the y_t series with each of the eight lag-series.[38] The autocorrelation function can be plotted as a histogram as shown for the AR(1) process (41) in the top panel of Figure B2 and the MA(2) process (48) in the third panel from the top. Such plots are called correlograms.

Consider the correlation between y_t and y_{t-2} in the correlogram shown in the top panel of Figure B2. The question arises as to whether the high degree

[38] Most statistical computer programs do these calculations to an approximation in order to conserve computing resources.

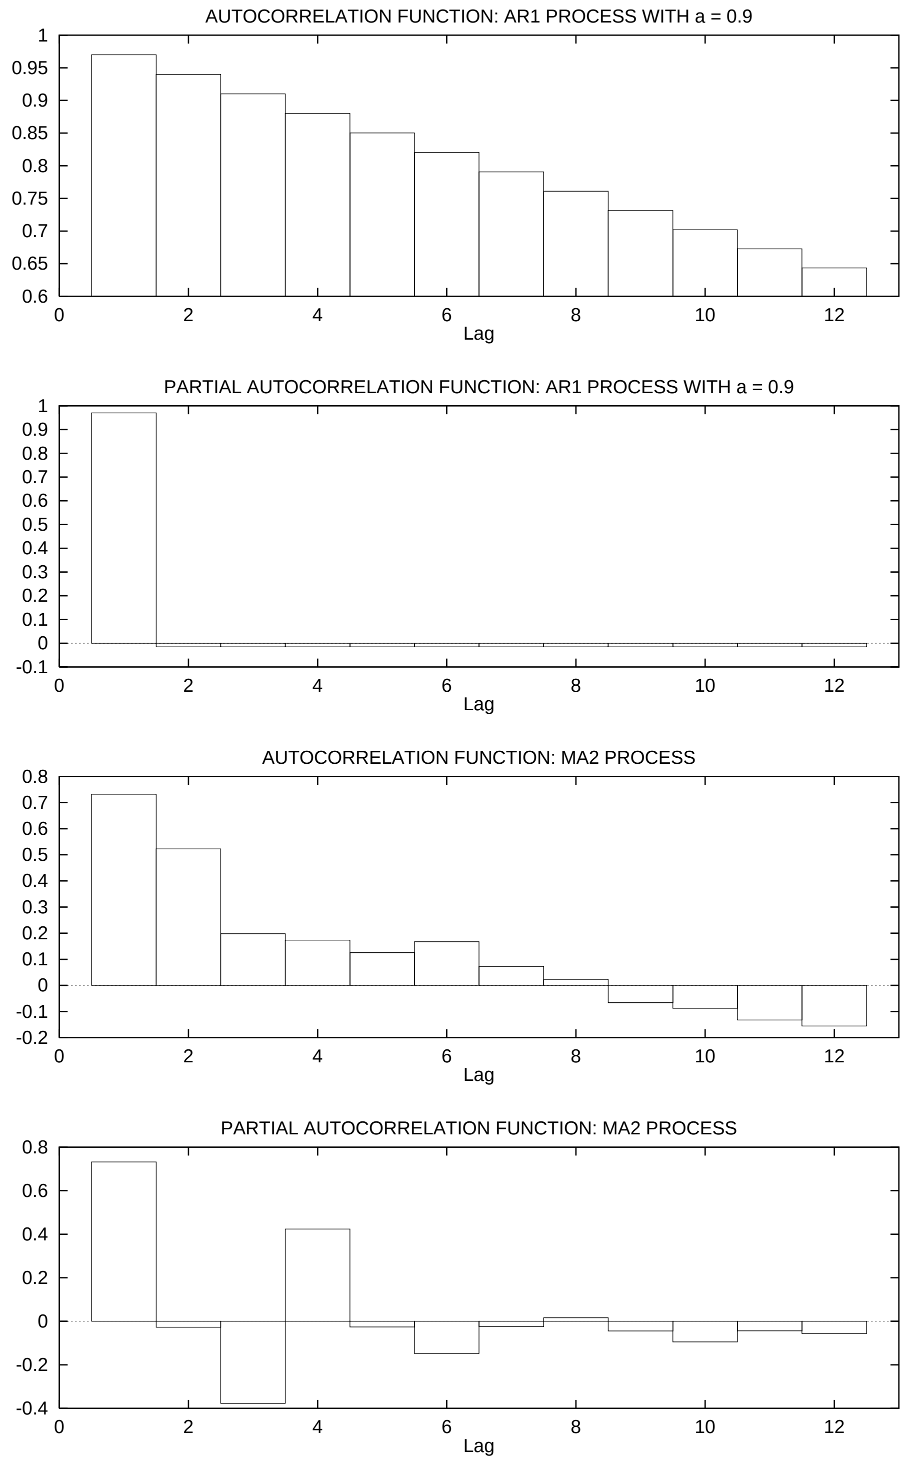

Figure B2: Correlograms for Autoregressive and Moving Average Processes: Some Examples.

of observed correlation arises because the twice-lagged values are directly correlated with the original series, or because they are correlated with the once-lagged values which are directly correlated with the original series. We can answer this question by calculating the partial autocorrelation function. This can be done by running a single regression of y_t on all n lags of y_t and extracting the regression coefficients.[39] These coefficients give the partial correlation of each of the j lagged series with the original series, holding all the other lagged series constant. The partial correlogram for the AR(1) series in the top panel of Figure B2 is shown in the second panel from the top, and the partial correlogram for the MA(2) process in the third panel from the top is shown in the bottom panel. As we might expect, all the partial correlations for the AR(1) series beyond the first are nearly zero—this occurs because a direct correlation only exists between the original series and its first lag, it being a first-order process. The second-lag series is correlated with the original series only because of its correlation with the first-lag series which is directly correlated with the original series.

Time series econometricians use the autocorrelation and partial autocorrelation functions of a series, along with their correlograms, to try to figure out the numbers of autoregressive and moving average lags, p and q in the ARIMA(p, d, q) process generating it. They determine d by testing the series for stationarity. If the original series is stationary then $d = 0$ and the process is integrated of order zero and is an ARMA(p, q) process. If the original series is non-stationary but its first difference is stationary then $d = 1$ and the process is integrated of order one—i.e., an ARIMA$(p, 1, q)$. If the original series has to be differenced twice to produce a stationary series then it is ARIMA$(p, 2, q)$, and so forth. Once d has been established, the correlograms are examined to attempt to determine appropriate values for p and q.

[39] Again, conventional computer packages perform these calculations to a convenient approximation.

Appendix C: Errors in Variables

Suppose that the truth is given by the equation

$$y_t \;=\; k\,x_t + \epsilon_t \tag{50}$$

but that we estimate

$$y_t \;=\; \alpha + \beta\,z_t + v_t \tag{51}$$

where z_t is an error-ridden proxy for x_t,

$$z_t \;=\; x_t + u_t. \tag{52}$$

This model applies to a variety of situations. For example, we can interpret y_t as consumption, x_t as permanent income, and z_t as current income, with k representing the fraction of permanent income consumed and β the marginal propensity to consume out of current income. Alternatively, in the analysis of this paper we interpret y_t as the change in the spot exchange rate between periods t and $t+1$, given by $(s_{t+1} - s_t)$ in equation (23), x_t as the forecast based on a correct interpretation of all information available to agents at time t, and z_t as the forward discount in period t, given by $(f_t - s_t)$ in equation (23). The forward discount is agents' actual forecast rather than the 'ideal' one with u_t representing the forecast error. In this interpretation, k equals unity, with the error term ϵ_t representing 'news' unavailable at the time of the forecast, and β has the same meaning as the β appearing in equation (23).

Suppose now that we estimate (51) using ordinary least-squares. The the estimate of β will be

$$\hat{b} \;=\; \frac{\Sigma(y-\bar{y})(z-\bar{z})}{\Sigma(z-\bar{z})^2} \tag{53}$$

Substituting (52) into the above, we obtain

$$\begin{aligned}
\hat{b} \;&=\; \frac{\Sigma(y-\bar{y})(x-\bar{x})}{\Sigma((x+u)-(\bar{x}+\bar{u}))^2} + \frac{\Sigma(y-\bar{y})(u-\bar{u})}{\Sigma((x+u)-(\bar{x}+\bar{u}))^2} \\
&=\; \frac{\Sigma(y-\bar{y})(x-\bar{x})}{\Sigma((x-\bar{x})^2 + 2(x-\bar{x})(u-\bar{u}) + (u-\bar{u})^2)} \\
&\qquad + \frac{\Sigma(y-\bar{y})(u-\bar{u})}{\Sigma((x-\bar{x})^2 + 2(x-\bar{x})(u-\bar{u}) + (u-\bar{u})^2)}.
\end{aligned} \tag{54}$$

When we divide the numerator and denominator of the above by the degrees of freedom (the number of observations less two) the terms become the sample variances and covariances as follows:

$$\begin{aligned}\hat{b} &= \frac{Cov\{y,x\}}{Var\{x\} + 2\,Cov\{x,u\} + Var\{u\}} \\ &\qquad + \frac{Cov\{y,u\}}{Var\{x\} + 2\,Cov\{x,u\} + Var\{u\}}\end{aligned} \tag{55}$$

In the limit, as the sample size gets larger and larger, the variances and covariances converge upon their population values, which we denote as σ_i^2 and $\sigma_{i,j}$, where $(i = x, y, u)$ and $(j = x, y, u)$, and $\hat{b}$ converges on β.

$$\beta = \frac{\sigma_{y,x}}{\sigma_x^2 + 2\,\sigma_{x,u} + \sigma_u^2} + \frac{\sigma_{y,u}}{\sigma_x^2 + 2\,\sigma_{x,u} + \sigma_u^2} \tag{56}$$

Now by assumption the forecast errors u_t are uncorrelated with the perfect-forecast values x_t—that is, agent's forecasts are unbiased. Furthermore, the forecast errors are independent of the future change in the spot rate. Hence, $\sigma_{x,u}$ and $\sigma_{y,u}$ are both zero. Equation (56) thus reduces to

$$\beta = \frac{\sigma_{y,x}}{\sigma_x^2 + \sigma_u^2} \tag{57}$$

which can be manipulated to yield

$$\beta = \frac{\sigma_{y,x}}{\sigma_x^2} \, \frac{\sigma_x^2}{\sigma_x^2 + \sigma_u^2}. \tag{58}$$

Since

$$\frac{\sigma_{y,x}}{\sigma_x^2} = k,$$

which equals unity, (58) reduces to

$$\beta = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_u^2}. \tag{59}$$

If a significant amount of information is available to agents and they interpret it perfectly, making no forecast errors, $\sigma_u^2 = 0$ and $\beta = 1$ as the traditional uncovered interest parity condition maintains. If agents' have no information about future exchange rate movements, both x_t and its variance will equal zero. If agents make no forecast, and hence make no errors, u_t

and its variance will also be zero. In this case, β will be undefined—the forward discount, z_t will be zero at all points in time. Suppose that agents have no information available to them about the innovations in the random-walk process that defines the exchange rate so that x_t and σ_x^2 are both zero, but they nevertheless mistakenly believe that they have some information, making forecasts that are purely random error. Then σ_u^2 will be positive and β will be zero. The forward discount will be uncorrelated with future movements in the spot rate.

Another way of interpreting (59) is to replace $\sigma_x^2 + \sigma_u^2$ with σ_z^2 to obtain

$$\beta \quad = \quad \frac{\sigma_x^2}{\sigma_z^2}. \tag{60}$$

We can view β as the ratio of the variance of the perfect-forecast, given the information available to agents, to the variance of the actual forecast implicit in the forward discount.

The fundamental question, of course, is why agents would persist in making erroneous forecasts, especially when their forecasts are dominated by error. The answer comes in part from Figure 6—the forward discount tends to be miniscule relative to the period-to-period movements of the spot rate. Even if agents have no information about the future course of the spot rate, forward discounts (or premia) of some magnitude are inevitable. If the exchange rate is perceived to be a random walk and no one wants to speculate, the forward rate will deviate from the spot rate (and expected future spot rate) on account of hedging pressure. At some point agents will induced by the size of the forward premium or discount to go unhedged or actively seek a forward position in the expectation that this period's spot rate will be a good forecast of next period's spot rate. This will limit the size of the forward discount or premium, but not reduce it to zero. So even if σ_x^2 is effectively zero, σ_u^2 will not be. When there is persistent inflation in one of the countries involved, agents will adjust the forward discount to reflect it, and their capacity to correctly forecast the future spot rate will improve even if the real exchange rate is known to be a random walk. In this event, σ_x^2 will increase relative to σ_u^2 and β will become significantly greater than zero. This is consistent with the evidence that, as inflation differentials increase, the forward discount becomes a better predictor of the future change in the spot rate.

References

Bachus, David K., Patrick J. Keyhoe and Finn E. Kydland (1995), "International Business Cycles: Theory and Evidence," in T. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press.

Balassa, Bela (1964), "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, Vol. 72, No. 6 (December), 584-96.

Barnhart, Scott W. and Andrew C. Szakmary (1991) "Testing the Unbiased Forward Rate Hypothesis: Evidence on Unit Roots, Co-Integration and Stochastic Coefficients," *Journal of Financial and Quantitative Analysis*, Vol. 26, No. 2 (June).

Berry, Thomas Senior (1988), *Production and Population Since 1789: Revised GNP Series in Constant Dollars*, Bostwick Paper No. 6, The Bostwick Press.

Blanchard, Olivier Jean and Stanley Fischer (1989) *Lectures on Macroeconomics*, MIT Press.

Dellas, Harris (1986), "A Real Model of the World Business Cycle," *Journal of International Money and Finance*, Vol. 5, No. 3 (September), 381-94.

De Long, Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldman (1989) "The Size and Incidence of Losses from Noise Trading," *Journal of Finance*, Vol. 44, No. 3 (July), 681-696.

De Long, Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldman (1990) "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, Vol. 98, No. 4 (August).

De Long, Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldman (1990) "Positive Feedback Investment Strategies and Destabilising Rational Speculation," *Journal of Finance*.

Edwards, Sebastian (1989) *Real Exchange Rates, Devaluation, and Adjustment*, MIT Press.

Enders, Walter (1995), *Applied Economic Time Series*, John Wiley and Sons.

Friedman, Milton and Anna J. Schwartz (1982) *Monetary Trends in the United States and the United Kingdom*, New York: National Bureau of Economic Research.

Gerlach, Stefan (1988), "World Business Cycles Under Fixed and Flexible Exchange Rates," *Journal of Money, Credit and Banking*, Vol. 20, No. 4 (November), 620-30.

Helpman, Elhanan (1981), "An Exploration in the Theory of Exchange Rate Regimes," *Journal of Political Economy*, Vol. 89, No. 5 (October), 865-890.

Helpman, Elhanan and Assaf Razin (1982), "Dynamics of a Floating Exchange Rate Regime," *Journal of Political Economy*, Vol. 90, No. 4 (August).

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

Kuznets, Simon (1930) *Secular Movements in Production and Prices*, New York: Houghton Mifflin.

Levich, Richard M. (1985), "Empirical Studies of Exchange Rates: Price Behavior, Rate Determination and Market Efficiency," in R.W. Jones and P.B. Kenen, eds., *Handbook of International Economics*, Vol. 2, North Holland.

Lucas, Robert (1978), "Asset Prices in an Exchange Economy," *Econometrica*, Vol. 46, No. 4 (December), 1426-45.

Maddala, G.S. (1988) *Introduction to Econometrics*, MacMillan.

Maynard, A. (2003), "Testing Forward Rate Unbiasedness: On Regression in Levels and Returns," *Review of Economics and Statistics*, Vol. 85, No. 2 (May), 313-327.

Maynard, A. and P.C.B. Phillips (2001), "Rethinking an Old Empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly," *Journal of Applied Econometrics*, Vol. 15, No. 6, 677-680.

McCallum, Bennett T. (1994), "A Reconsideration of the Uncovered Interest Parity Relationship," *Journal of Monetary Economics*, Vol. 33, No. 1 (February), 105-32.

Meese, Richard and Kenneth Rogoff (1983), "Empirical Exchange Rate Models of the Seventies: Do They Fit Out-of Sample," *Journal of International Economics*, Vol. 14, No. 1-2 (February), 3-24.

Michell, B.O. (1981) *European Historical Statistics, 1750-1975*, Cambridge University Press.

Mussa, Michael (1979), "Empirical Regularities in the Behavior of Exchange Rates and Theories of the Foreign Exchange Market," in Karl Brunner and Allan H. Meltzer, eds., *Carnegie-Rochester Conference Series on Public Policy*, Vol. 11, (*Policies for Employment, Prices, and Exchange Rates*), North Holland.

Obstfeld, Maurice and Kenneth Rogoff (2000), "The Six Major Puzzles in International Macroeconomics: Is there a Common Cause?" in Ben Bernanke and Kenneth Rogoff eds., *NBER Macroeconomics Annual 2000*, Cambridge: NBER and the MIT Press.

Officer, Lawrence H. (1982), *Purchasing Power Parity and Exchange Rates: Theory, Evidence and Relevance*, London and Greenwich Connecticut: JAI Press.

Officer, Lawrence H. (1983), "Dollar-Sterling Mint Parity and Exchange Rates, 1791-1834," *Journal of Economics History*, Vol. 43, No.3 (September), 570-616.

Perron, Pierre (1989), "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis," *Econometrica*, Vol. 57, No. 6 (November), 1361-1401.

Rogoff, Kenneth (1996), "The Purchasing Power Parity Puzzle," *The Journal of Economic Literature*, Vol. 34, (June), 647-668.

Said, S and David Dickey (1988), "Testing for a Unit Root in Time Series Regressions", *Biometrica*, Vol. 75, No. 2 (June), 311-40.

Samuelson, Paul A. (1964), "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, Vol. 46, No. 2 (May), 145-54.

Shleifer, Andrei and Lawrence H. Summers, (1990), "The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, Vol. 4, No. 2 (Spring), 19-33.

Stock, James H. and Mark W. Watson (2003), *Introduction to Econometrics*, Addison Wesley.

Stockman, Allan C. (1980), "A Theory of Exchange Rate Determination," *Journal of Political Economy*, Vol. 88, No. 4 (August).

Stockman, Allan C. (1983), "Real Exchange Rates Under Alternative Exchange Rate Regimes," *Journal of International Money and Finance*, Vol. 3, No. 2 (June), 147-66.

Stockman, Allan C. and Lars E.O. Svensson (1987), "Capital Flows, Investment and Exchange Rates," *Journal of Monetary Economics*, Vol. 19, No. 2 (March), 171-202.

Urquhart, M.C., K.A.H. Buckley and F.H. Leacy (1983) *Historical Statistics of Canada*, Statistics Canada and Social Science Federation of Canada, 2nd. Edition.