Towards a Solution to the Puzzles in Exchange Rate Economics: Where Do We Stand?

Lucio Sarno

Finance Group, Warwick Business School, University of Warwick

First version: September 2004 - Revised: February 2005

Abstract

This paper provides a selective overview of puzzles in exchange rate economics. We begin with the forward bias puzzle: high interest rate currencies appreciate when one might guess that investors would demand higher interest rates on currencies expected to fall in value. We then analyze the purchasing power parity puzzle: the real exchange rate displays no (strong) reversion to a stable long-run equilibrium level. Finally, we cover the exchange rate disconnect puzzle: the lack of a link between the nominal exchange rate and economic fundamentals. For each puzzle, we critically review the literature and speculate on potential solutions.

JEL classification: F31.

Keywords: exchange rates; interest rate parity; purchasing power parity; forecasting.

1 Introduction

Exchange rate economics is characterized by a number of anomalies, or puzzles, which we struggle to explain on the basis of either sound economic theory or practical thinking. Put more simply, the international finance profession has not yet been able to produce theories and, as a consequence, empirical models which allow us to explain the behavior of exchange rates with a reasonable degree of accuracy. This failure is witnessed by a variety of phenomena, and this paper analyses three specific ones.

The first puzzle analyzed is the 'forward bias puzzle,' relating to the fact that the foreign exchange market is not only informationally inefficient, but it appears to be so inefficient that the forward market–capturing market expectations of future movements in exchange rates–may systematically predict future exchange rate movements in the wrong direction (Fama, 1984). The second puzzle relates to the often documented lack of any strong tendency of exchange rates to move in sync with relative prices, which is what one would expect if purchasing power has to remain constant across countries over long periods of time in a world with international arbitrage in goods markets–this is usually termed the 'purchasing power parity' puzzle (Rogoff, 1996). The third puzzle, which in some respects encompasses the previous two, is the missing link between nominal exchange rates and the menu of economic or financial fundamentals that international economics theory suggests should drive exchange rates–this phenomenon is termed the 'exchange rate disconnect' puzzle (Obstfeld and Rogoff, 2000). In essence, fundamentals appear to be unable to explain both the actual level of exchange rates–not only daily, but even monthly, quarterly and annually–and their volatility.

This paper summarizes the present author's reading of the research relating to the above three puzzles in exchange rate economics–what we have learned, which aspects of the puzzles we have solved and which remain, and where further research progress is most likely to be made. While the paper will be of use to specialists in international finance and macroeconomics, given the importance of the relevant issues discussed in this article, essentially relating to understanding why exchange rates move the way they do, we also hope that our assessment of the central questions motivating our analysis will be of interest to a wider audience of economists, policy makers and practitioners. In the remainder of the article we tackle each of the three puzzles described above in separate sections, and we then briefly summarize and record our conclusions in a final section.

2 The forward bias puzzle

2.1 What is the forward bias and why do we care?

In an efficient speculative market, prices should fully reflect information available to market participants and it should be impossible for a trader to earn excess returns to speculation. The Uncovered Interest Parity (UIP) condition is the cornerstone parity condition for foreign exchange market efficiency:

$$\Delta_k s^e_{t+k} = i_{t,k} - i^*_{t,k}, \tag{1}$$

where s_t denotes the logarithm of the spot exchange rate (domestic price of foreign currency) at time t; $i_{t,k}$ and $i^*_{t,k}$ are the nominal interest rates available on similar domestic and foreign securities respectively (with k periods to maturity); $\Delta_k s_{t+k} \equiv s_{t+k} - s_t$; and the superscript e denotes the market expectation based on information at time t. In its simplest form, the efficient markets hypothesis can be reduced to a joint hypothesis that foreign exchange market participants are, in an aggregate sense, (a) endowed with rational expectations and (b) risk-neutral.

Most often, analyses of foreign exchange market efficiency have taken place in the context of the relationship between spot and forward exchange rates under the assumption that covered interest parity (CIP) holds: $f^k_t - s_t = i_{t,k} - i^*_{t,k}$, where f^k_t is the logarithm of the k-period forward rate (i.e. the rate agreed now for an exchange of currencies k periods ahead). Indeed, CIP is a reasonably mild assumption, given the extensive empirical evidence suggesting that CIP holds (for a survey of this evidence, see e.g. Sarno and Taylor, 2003, Ch. 2). Note that, unlike CIP, UIP is not an arbitrage condition since one of the terms in the UIP equation (1), namely the expected exchange rate, is unknown at time t and, therefore, non-zero deviations from UIP do not necessarily imply the existence of arbitrage profits due to the foreign exchange risk associated with future exchange rate movements.

Using CIP and replacing the interest rate differential $i_{t,k} - i^*_{t,k}$ with the forward premium (or forward discount) $f^k_t - s_t$, a number of researchers have tested UIP by estimating a regression of the form:

$$\Delta s_{t+1} = \alpha + \beta \left(f^1_t - s_t \right) + \upsilon_{t+1}, \tag{2}$$

where we have assumed that $k = 1$ for simplicity, and υ_{t+1} is a disturbance term. Under UIP, $\alpha = 0$, the slope parameter β must equal unity, and the disturbance term υ_{t+1} (the rational expectations forecast error) must be uncorrelated with information available at time t (e.g. Fama, 1984).

Empirical studies based on the estimation of equation (2), for a large variety of currencies and time periods, generally report results which reject UIP and the efficient markets hypothesis (e.g. see the references in the survey of Hodrick, 1987; Lewis, 1995; Engel, 1996). Indeed it constitutes a stylized fact that estimates of β, using exchange rates against the dollar, are often statistically insignificantly different from zero and generally closer to minus unity than plus unity (Froot and Thaler, 1990).[1] The stylized fact of a negative β coefficient in this regression implies that the more the foreign currency is at a premium in the forward market, the less the home currency is predicted to depreciate.[2] The negative value of β is the central feature of the forward bias puzzle and, following much previous literature, we shall refer to equation (2) as the 'Fama regression.'[3]

2.2 How has the forward bias been addressed?

The rejection of the simple, risk-neutral efficient markets hypothesis may be due to risk-aversion of market participants or to a departure from the rational expectations hypothesis, or both of these reasons. If foreign exchange market participants are risk averse, the UIP condition may be distorted by a risk premium, ρ_t say, because agents demand a higher rate of return than the interest differential in return for the risk of holding foreign currency.[4]

Note that the vast majority of studies in this context estimate the Fama regression using ordinary least squares (OLS). This can be problematic in the presence of an omitted risk premium in the regression, in which case OLS would yield biased and inconsistent estimates of β (Fama, 1984; Liu and Maddala, 1992). Recently, Barnhart, McNown and Wallace (1999) have shown that two conditions are needed for this problem to arise: (i) the forward rate must be a function of an unobservable omitted variable, such as predictable excess returns; (ii) the term containing the forward rate in the estimated regression must be stationary or, if nonstationary, can be normalized to a stationary variable. Under these conditions, Barnhart, McNown and Wallace document the severity of this problem in a variety of spot-forward regressions, concluding that most common tests of UIP are non-informative in the presence of an omitted risk premium.

McCallum (1994) suggests that the negativity of the estimated UIP slope coefficient may be the result of simultaneity induced by the existence of a monetary policy reaction function where the interest rate differential is set in order to avoid large current exchange rate movements as well as to smooth interest rate movements. This may be seen as a special case of the general point made by Fama (1984) that negativity of estimates of β require a negative covariation between the risk premium and the expected rate of depreciation.

More recently Chinn and Meredith (2004) have extended the analysis of McCallum (1994) to tighten the link between monetary policy and the behavior of UIP deviations. Chinn and Meredith start from showing that, empirically, while the forward bias is very robust when using short-horizon data, estimates of β in long-horizon UIP regressions–i.e. using longer-maturity bonds–have the correct sign (positive) and are generally closer to unity than to zero. They reconcile the difference in the estimates of β at short and long horizons using a macroeconomic model that enriches the framework of McCallum (1994) by incorporating a reaction function that causes interest rates to respond to innovations in output and inflation (as opposed to the exchange-rate targeting assumption used by McCallum). Stochastic simulations of the model generate artificial data with similar moments to the actual data and, more importantly, estimation of UIP regressions on these data generates forward bias at short horizons but not at long horizons, consistent with the empirical work on actual data. Intuitively the long-horizon results differ sharply from the results at short horizon because the model's "fundamentals" play a more important role over longer horizons, while interest rate differentials are biased predictors of exchange rate movements in the short-term due to the behavior of the authorities in "leaning against the wind" in the face of exchange rate shocks via their effect on output and inflation. These results are encouraging since the empirical work suggests that the forward bias may be confined to short maturity assets and, hence, be less pervasive that previously thought. At the theoretical level, however, explaining the differences in the estimates of β for short and long horizons in the model used by Chinn and Meredith (2004) still requires the existence of underlying shocks in exchange markets of a size that is hard to imagine it could be generated by risk premia. Shocks of large size are necessary for the model to generate the observed exchange rate volatility. In addition, the model also assumes (a) that exchange rate forecasts are unbiased predictions of the future spot exchange rates, in contrast with the empirical evidence on survey data analyses (e.g. Froot and Ito, 1989); and (b) that the expectations hypothesis of the term structure of interest rates holds, which is also a controversial assumption (e.g. Campbell, Lo and MacKinlay, 1997). Nevertheless, the finding of Chinn and Meredith that the forward bias characterizes primarily, if not exclusively, short-horizon UIP regressions is bound to have an impact in redirecting some of the future research in this area.

What is clear from this analysis, however, is that a time-varying risk premium will confound simple efficiency tests of the kind outlined above. This should not be surprising since, maintaining rational expectations, we can always define the risk premium, without loss of generality, as the difference between the two sides of equation (1). While this would allow us to study some of the properties of the risk premium by examining its projection on available information, there is no reason to expect that this implicitly defined risk premium will behave in a manner consistent with our economic intuition.

An earlier strand of the literature attempted to understand foreign exchange risk focusing on simple extensions of the static version of the capital asset pricing model (see e.g. Adler and Dumas, 1983; Frankel, 1982; Giovannini and Jorion, 1989; Engel, 1992). Most of these studies provide evidence that the risk aversion parameter is very large but often not significantly different from zero and also that the restrictions imposed in the model are rejected.

A subsequent literature has built on this research to analyze the role of the risk premium in a dynamic general equilibrium context. The chief example is the dynamic, two-country, general equilibrium model of Lucas (1982). This model shows that, under canonical assumptions, there is a wedge between the spot and forward rate that is driven by the risk premium. For the risk premium to explain a significant chunk of the forward rate forecast error or excess returns, either there must be a very large coefficient of relative risk aversion ϕ, or consumption must be highly correlated with the exchange rate. The intuition for the fact that high correlation between consumption and the exchange rate raises the risk premium is that forward exchange positions provide less of a hedge against variations in consumption the greater is this covariation. The fact is, however, that consumption tends to be fairly smooth in any advanced economy, while the nominal exchange rate–at least under floating rate regimes–is typically a lot more volatile, so that this covariation will be quite small. As a consequence, tests of the implications of these models relating forward exchange rates and expected spot exchange rates from the first order conditions of the Lucas model or similar general equilibrium models have led to rejections of the models (e.g. Mark, 1985; Bekaert and Hodrick, 1992; Bekaert, 1994). The message which emerges from the empirical analysis of risk-premium models–in general or partial equilibrium–is that it is hard to explain excess returns in forward foreign exchange by an appeal to risk premia alone: either ϕ, the coefficient of relative risk aversion, must be incredibly large, or else the conditional covariance of consumption and the spot rate must be incredibly high.

An alternative explanation of the rejection of the simple efficient markets hypothesis is that there is a failure, in some sense, of the rational expectations component of the joint hypothesis underlying the notion of market efficiency. The literature identifies in this group several possibilities: rational bubbles; learning about regime shifts (Lewis, 1989a,b) or about fundamentals, such as learning about the interest rate process (Gourinchas and Tornell, 2004); the "peso problem" originally suggested by

Rogoff (1979); or inefficient information processing, as suggested, for example, by Bilson (1981). This literature is well covered in several surveys (e.g. Lewis, 1995; Engel, 1996; Sarno and Taylor, 2003, Ch. 2), to which the interested reader is referred.

Rational bubbles, learning and peso problems all imply departures from rational expectations that generate non-zero and potentially predictable excess returns even when agents are risk-neutral. A problem with admitting peso problems, bubbles or learning into the class of explanations of the forward bias is that, as noted above, a very large number of econometric studies–encompassing a very large range of exchange rates and sample periods–have found that the direction of the bias is the same under each scenario, i.e. the estimated UIP slope parameter, β is generally negative and closer to minus unity than plus unity. For example, Lewis (1989a), in her study of the relationship of the early 1980s dollar appreciation with learning about the US money supply process, notes a degree of persistence in the forward rate errors which, in itself, is *prima facie* evidence against the learning explanation: agents cannot forever be learning about a once-for-all regime shift. Similarly, the peso problem is essentially a small-sample phenomenon; it cannot explain the fact that estimates of β are *generally* negative.

A limitation with much of the empirical literature on the possible rationalizations of the rejection of the simple, risk-neutral efficient markets hypothesis is that in testing one leg of the joint hypothesis, researchers have typically had no alternative but to assume that the other leg is true. For instance, the search for a stable empirical risk premium model has generally been conditioned on the assumption of rational expectations (see e.g. Fama, 1984; Hodrick and Srivastava, 1984). Other studies assume, however, that investors are risk-neutral and hence that the deviation from the unbiasedness hypothesis would suggest rejection of the rational expectations hypothesis (e.g. Bilson, 1981; Cumby and Obstfeld, 1984).

The availability of survey data on exchange rate expectations–e.g. from the American Express Bank, the *Economist* and *Money Market Services*–has allowed researchers to conduct tests of each component of the joint hypothesis. Once exchange rate expectations are available no need exists to impose any assumption regarding the expectations formation mechanism of market agents. Important contributions in this area includes the work by Frankel and Froot (1987) and Froot and Frankel (1989). This line of research has established that *both* risk aversion *and* departures from rational expectations are responsible for the rejection of the simple efficient markets hypothesis.

2.3 Rays of hope in the search for a solution to the forward bias puzzle

Figure 1 shows weekly time series for the dollar-sterling log-spot exchange rate as well as four log-forward rates for maturities of 1, 3, 6 and 12 months over the sample period 1984-2003. Simple eyeball econometrics clearly suggests that the spot rate moves closely together with each of the forward exchange rates throughout the sample period. Indeed, the co-movement is so strong that the differences between the spot and forward rates appear miniscule. This informal eyeball analysis is in stark contrast with the stylized fact discussed above that the forward rate is not only a biased predictor of the spot exchange rate but also it may systematically mispredict its direction. A recent strand of research has contributed to this literature beginning from an analysis of the dynamic relationship between spot and forward rates that, while assuming that UIP does not hold (the forward rate is a biased predictor), measures the predictive power of the forward rate in a richer characterization of the spot-forward relationship.

2.3.1 The spot exchange rate and the term structure of forward rates

Under the assumptions that (i) each of s_t and f_t^k are well described by unit root processes and that (ii) departures from the risk-neutral efficient markets hypothesis–namely expected foreign exchange excess returns, $f_t^k - E_t\left(s_{t+k}|\Omega_t\right)$, defined with respect to a given information set Ω_t–are stationary, it is straightforward to derive an expression which implies that the forward premium, $f_t^k - s_t$ is stationary (Clarida and Taylor, 1997). In turn, this result implies that forward and spot exchange rates have a common stochastic trend and are cointegrated with cointegrating vector $[1, -1]$. Since this is true for any k, if we consider the vector of forward rates of tenor 1 to m periods, together with the current spot rate, $[s_t, f_t^1, f_t^2, f_t^3, \ldots, f_t^m]'$, then this must be cointegrated with m unique cointegrating vectors, each given by a row of the matrix $[-\iota, I_m]$, where I_m is an m-dimensional identity matrix and ι is an m-dimensional column vector of ones. Finally, by the Granger Representation Theorem this vector of forward and spot rates must possess a VECM representation in which the term structure of forward premia plays the part of the equilibrium errors. This linear VECM may be written as follows:

$$\Delta y_t = \nu + \textstyle\sum_{d=1}^{p-1} \Gamma_d \Delta y_{t-d} + \Pi y_{t-1} + u_t \qquad (3)$$

where $y_t = [s_t, f_t^4, f_t^{13}, f_t^{26}, f_t^{52}]'$, with the superscript denoting the number of weeks corresponding to the maturity of the forward contract; $\Pi = \alpha\beta'$ is the long-run impact matrix whose rank determines the number of cointegrating vectors linking spot and forward rates (equal to four in this specific VECM with our definition of y_t as a 5×1 vector with one spot rate and four forward rate time series); and u_t is a vector of Gaussian error terms.

Clarida and Taylor (1997) exploit this linear VECM representation to show that sufficient information may be extracted from the term structure in order to forecast the spot dollar exchange rate during the recent floating exchange rate regime. Their dynamic out-of-sample forecasts suggest that the linear VECM is superior to a range of alternative forecasts, including a random walk and standard spot-forward regressions. In short, this evidence suggests that the term structure of forward rates provides satisfactory predictions of the future spot exchange rate.[5]

Clarida, Sarno, Taylor and Valente (CSTV, 2003) then generalize the linear VECM in equation (3) to a multivariate Markov-switching framework and examine the performance of such a model in out-of-sample exchange rate forecasting. This generalized term structure model was inspired by encouraging results previously reported in the literature on the presence of nonlinearities (and particularly by the success of Markov-switching models) in the context of exchange rate modelling. Using weekly data on major spot and forward dollar exchange rates over the period 1979 through 1995, CSTV report evidence of the presence of nonlinearities in the term structure and forecast dynamically out of sample over the period 1996 through to 1998. The results suggest that the Markov-switching VECM (MS-VECM) forecasts are strongly superior to the random walk forecasts at a range of forecasting horizons up to 52 weeks ahead, using standard forecast accuracy criteria. Moreover, the MS-VECM also outperforms a linear VECM for spot and forward rates in out-of-sample forecasting of the spot rate, although the magnitude of the gain, in point forecasting, from using an MS-VECM relative to a linear VECM is rather small at short horizons (about 10% on average at the 4-week forecast horizon).

It is possible, however, that traditional measures of forecast accuracy mask somehow the potential superiority of nonlinear models (Satchell and Timmermann, 1995; Granger, 2003). The vast majority of studies on exchange rate forecasting has traditionally focused on accuracy evaluations based on point forecasts, such as the mean absolute error (MAE) and the root mean square error (RMSE). Several authors have recently emphasized the importance of evaluating the forecasting ability of economic models on the basis of density, as opposed to point, forecasting performance (see, *inter alia*, the survey by Tay and Wallis, 2000, and the references therein). In a decision-theoretical context, the need to consider the predictive density of a time series–as opposed to considering only its conditional mean and variance–seems fairly accepted in the light of the argument that economic agents may not have loss functions that depend symmetrically on the realizations of future values of potentially non-

Gaussian variables. In this case, agents are interested in knowing not only the mean and variance of the variables in question, but their full predictive densities. In various contexts in economics and finance–among which the recent boom in financial risk management represents an obvious case–there is an increasingly strong need to provide and evaluate density forecasts. These issues are particularly important in the context of nonlinear models since these models may provide highly non-normal densities.

Several researchers have proposed methods for evaluating density forecasts–e.g. see Diebold, Gunther and Tay (1998), Granger and Pesaran (1999) and Berkowitz (2001). Sarno and Valente (2005) re-examine the short-horizon forecasting performance of the MS-VECM of the term structure using weekly data for eight US dollar exchange rates during the recent floating exchange rate regime. On the basis of density forecasting tests, Sarno and Valente document that the MS-VECM produces very satisfactory one-week-ahead density forecasts and outperform its more parsimonious linear counterpart as well as the standard benchmark in the exchange rate forecasting literature, namely the random walk model.

Sarno and Valente then illustrate the practical importance of the density forecasts for the purpose of risk management. In recent years, trading accounts at large financial institutions have shown a dramatic growth and become increasingly more complex. Partly in response to this trend, major trading institutions have developed large-scale risk measurement models designed to manage risk. These models generally employ the Value-at-Risk (VaR) methodology. VaR may be defined as the expected maximum loss over a target horizon within a given confidence interval (Jorion, 2001). More formally, VaR is an interval forecast, typically a one-sided 95 or 99 percent interval of the distribution of expected wealth or returns. Users of the VaR methodology generally assume that expected returns are normally or t-distributed. However, this assumption contrasts with the large amount of empirical evidence suggesting that the distribution of exchange rate returns is not standard. Point forecast analysis and testing procedures based upon it do not take into account these features, so that VaR analysis often relies on dubious parametric distributional assumptions. Sarno and Valente (2005) investigate the implications of density forecasts for a risk manager who has to quantify the risk associated with a simple internationally diversified portfolio over a one-week horizon. In this simple application it is shown how density forecasts can help us discriminating among competing exchange rate models. The random walk model and the linear exchange rate model produce forecasts that do not capture satisfactorily the higher moments of the predictive distribution of the exchange rate, generating VaRs that poorly estimate the probability of large losses. However, the Markov-switching VECM, which does better at matching the higher moments of the predictive distribution of exchange rates, produced VaRs that are generally in line with the target violation rate.

Overall, these findings highlight how better density forecasts of exchange rates, obtained using the term structure of forward rates, can potentially lead to substantial improvements in risk management and, more precisely, to better estimates of downside risk. This is the case even though the forward rate is not an optimal predictor of the future spot exchange rate, i.e. even though there is a forward bias.

2.3.2 Is the forward bias 'economically' important?

Sarno, Valente and Leon (2004) approach the forward bias puzzle from a different angle. Specifically, they start from noting that prior empirical research in this area has generally relied on linear frameworks in analyzing the properties of UIP deviations. However, several authors have argued that the relationship between expected exchange rates and interest rate differentials may be nonlinear for a variety of reasons, including transactions costs (see, *inter alia*, Baldwin, 1990; Dumas, 1992; Hollifield and Uppal, 1995; Sercu and Wu, 2000), central bank intervention (e.g. Mark and Moh, 2002), and the existence of limits to speculation (e.g. Lyons, 2001, pp. 206-220). In particular, the limits to speculation hypothesis is based on the idea that financial institutions only take up a currency trading strategy if this strategy is expected to yield an excess return per unit of risk (or a Sharpe ratio) that is higher than the one implied by alternative trading strategies, such as, for example, a simple buy-and-hold equity strategy. This argument effectively defines a band of inaction where the forward bias does not attract speculative capital and, therefore, does not imply any glaring profitable opportunity and will persist until it generates Sharpe ratios that are large enough to attract speculative capital away from alternative trading strategies (Lyons, 2001).

Although the literature has already documented that normal values of the forward premia may impact on future exchange rates differently from extreme values (e.g. Bilson, 1981; Flood and Taylor, 1996; Huisman, Koedijk, Kool and Nissen, 1998) and that the response of dollar exchange rate changes may be different for positive and negative values of the interest rate differential (e.g. Bansal, 1997; Bansal and Dahlquist, 2000), and some authors have investigated the role of nonlinearities in the term structure of forward premia for exchange rate forecasting (e.g. CSTV, 2003), the potential importance of nonlinearities to shed light on the forward bias puzzle remains largely under-researched. Sarno, Valente and Leon (2004) build an empirical framework that provides a characterization of the UIP condition which allows us to test some of the general predictions of the limits to speculation hypothesis and to assess its potential to explain the forward bias puzzle and the excess returns predictability documented in the literature. Their empirical results, obtained using five major US dollar exchange rates since 1985 and considering forward rates with 1- and 3-month maturity, are as follows. First, there is strong evidence that the relationship between spot and forward exchange rates is characterized by significant nonlinearities. While the detection of nonlinearities in this context is not novel *per se*, this empirical model proves especially useful for understanding the properties of deviations from UIP. In particular, consistent with Lyons' limits-to-speculation hypothesis, in the neighborhood of UIP, expected excess returns and hence the forward bias are statistically significant and persistent but economically too small to attract speculative capital, while for expected excess returns which are large enough to attract speculative capital the spot-forward relationship reverts rapidly towards the UIP condition.

Given these findings, Sarno, Valente and Leon (2004) carry out a battery of Monte Carlo experiments to demonstrate that, if the true data generating process (DGP) governing the relationship between spot and forward exchange rates were of the nonlinear form they consider, it is possible to replicate the empirical results generally reported in the literature. In particular, estimation of the conventional linear spot-forward regressions would lead us to reject both the validity of UIP and the hypothesis of no predictability of foreign exchange excess returns with parameters estimates that are very close to the ones observed using actual data. However, the failure of UIP and the findings of a forward bias and predictability of excess returns are features that the DGP has only in one regime, which is the regime where deviations from UIP are tiny enough to be economically unimportant and unlikely to attract speculative capital.

A plausible interpretation of this evidence is that the stylized fact that the UIP condition is statistically rejected by the data may not be indicative of substantial market inefficiencies. Indeed, the inefficiencies implied by this rejection appear to be very tiny and it is not clear that they are economically important. Further research is awaited to shed light on the economic significance of the forward bias.

3 The purchasing power parity puzzles

3.1 The search for purchasing power parity

The purchasing power parity (PPP) hypothesis states that national price levels should be equal when expressed in a common currency. Although very few economists would believe that this simple proposition holds at each point in time, a large literature in international finance has examined empirically the validity of PPP over the long-run either by testing whether nominal exchange rates and relative prices move together or by testing whether the real exchange rate has a tendency to revert to a stable equilibrium level over time. The latter approach is motivated by the fact that the real exchange rate may be defined as the nominal exchange rate adjusted for relative national price levels. More formally, the real exchange rate, q_t, may be expressed in logarithmic form as

$$q_t \equiv s_t - p_t + p_t^* \tag{4}$$

where p_t and p_t^* denote the logarithms of the domestic and foreign price levels respectively. The real exchange rate, q_t may thus be interpreted as a measure of the deviation from PPP and must be stationary for long-run PPP to hold (see the surveys of Froot and Rogoff, 1995; Rogoff, 1996; Sarno and Taylor, 2002; Taylor and Taylor, 2004).

Although long-run PPP is a very simple proposition about exchange rate behavior, it has attracted the attention of researchers for decades. Indeed, whether long-run PPP holds or whether the real exchange rate is stationary has important economic implications on a number of fronts. In particular, the degree of persistence in the real exchange rate can be used to infer the principal impulses driving exchange rate movements. For example, if the real exchange rate is highly persistent or close to a random walk, then the shocks are likely to be real-side, principally technology shocks, whereas if it is not very persistent, then the shocks must be principally to aggregate demand, such as, for example, innovations to monetary policy (Rogoff, 1996). Further, from a theoretical perspective, if PPP is not a valid long-run international parity condition, this casts doubts on the predictions of much open-economy macroeconomics that is based on the assumption of long-run PPP. Indeed, the implications of open economy dynamic models are sensitive to the presence or absence of a unit root in the real exchange rate (e.g. Lane, 2001; Sarno, 2001). Finally, estimates of PPP exchange rates are often used for practical purposes such as determining the degree of misalignment of the nominal exchange rate and the appropriate policy response, the setting of exchange rate parities, and the international comparison of national income levels. These practical uses of the PPP concept, and in particular the calculation of PPP exchange rates, would obviously be of very limited use if PPP deviations contain a unit root.

Regardless of the great interest in this area of research, manifested by the large number of papers on PPP published over the last few decades, and regardless of the increasing quality of data sets utilized and of the econometric techniques employed, the validity of long-run PPP and the properties of PPP deviations remain the subject of an ongoing controversy. Specifically, earlier cointegration studies generally reported the absence of significant mean reversion of the real exchange rate for the recent floating experience (Mark, 1990), but were supportive of reversion toward PPP for the gold standard period (McCloskey and Zecher, 1984; Diebold, Husted and Rush, 1991), for the interwar float (Taylor and McMahon, 1988), for the 1950s US-Canadian float (McNown and Wallace, 1989), and for the exchange rates of high-inflation countries (Choudhry, McNown and Wallace, 1991).

An important point to note is that, in testing for mean reversion in real exchange rates, most studies in the literature have examined real exchange rates constructed using official price indices. If real exchange rate adjustment towards the PPP equilibrium is driven by arbitrage in international goods markets, however, the appropriate price index to be used in implementing PPP is of crucial importance.[6] In particular, all commonly used price indices include some proportion of nontradable goods, for which arbitrage does not occur. An influential attempt in the literature to construct appropriate price indices for the real exchange rate has been carried out by Summers and Heston (1991), although their data is not of great help in practice for time series econometricians since it is constructed at infrequent and long time intervals. This is the main reason why economists typically use price indices made available by official sources when constructing the real exchange rate, despite their limitations for the purpose of testing the validity of long-run PPP. However, some work on PPP looks at the cost of production of a basket of goods–producer choices–rather than the cost of a basked of goods in terms of consumer choices. For example, Lafrance, Osakwe and Normandin (1998) and Sarno and Chowdhury (2003) provide evidence that PPP works much better if it is based on costs of production–essentially unit labor costs–or on indices made only of tradable goods, rather than consumer price indices from official sources.

One well-documented explanation for the inability to find clear-cut evidence of PPP is the low power of conventional statistical tests to reject a false null hypothesis of a unit root in the real exchange rate or no cointegration between the nominal exchange rate and relative prices with a sample span corresponding to the length of the recent float (Frankel, 1986, 1990; Froot and Rogoff, 1995; Lothian and Taylor, 1997).[7] Researchers have sought to overcome the power problem in testing for mean reversion in the real exchange rate either using long span studies (e.g. Lothian and Taylor, 1996; Taylor, 2002) or panel unit root tests (e.g. Abuaf and Jorion, 1990; Frankel and Rose, 1996; O'Connell, 1998; Papell, 1998; Sarno and Taylor, 1998; Taylor and Sarno, 1998). However, whether or not the long-span or panel-data studies do in fact answer the question whether PPP holds in the long run remains contentious. As far as the long-span studies are concerned, as noted in particular by Frankel and Rose (1996), the long samples required to generate a reasonable level of statistical power with standard univariate unit root tests may be unavailable for many currencies (perhaps thereby generating a 'survivorship bias' in tests on the available data) and, in any case, may potentially be inappropriate because of differences in real exchange rate behavior both across different historical periods and across different nominal exchange rate regimes (e.g. Baxter and Stockman, 1989; Taylor, 2002). As for panel-data studies, these provide mixed evidence. While, for example, Abuaf and Jorion (1990), Frankel and Rose (1996) and Taylor and Sarno (1998) find results favorable to long-run PPP, O'Connell (1998) rejects it on the basis of his empirical evidence.

In light of the evidence provided by this literature, there remain several unresolved puzzles, among which two are prominent. First, it is still controversial whether long-run PPP is valid during the recent floating exchange rate regime. Second, it is puzzling why the majority of studies which favor long-run PPP find empirical estimates of the persistence of PPP deviations that are too high–the half-life of shocks ranges between three and five years–to be explained in light of conventional nominal rigidities and to be reconciled with the large short-term volatility of real exchange rates (Rogoff, 1996).

A source of potentially important bias in estimates of the half life is caused by cross-sectional aggregation in moving from the law of one price for individual goods to PPP deviations based on price indices. Imbs, Mumtaz, Ravn and Rey (2005) demonstrate analytically how such bias is bound to be present in estimates of the real exchange half life and then provide empirical evidence that the bias is upwards and substantial. Crucini, Telmer and Zachariadis (2005) take a similar approach to understanding the behavior of deviations from the law of one price and PPP by examining micro-data on absolute prices of goods. They study good-by-good deviations from the law of one price for over 5,000 goods and services between European Union countries for the years 1975, 1980, 1985 and 1990, reporting that between most countries there are roughly as many overpriced goods as there are underpriced goods so that PPP holds to a good approximation, particularly after controlling for wealth differences.

It is instructive to graph the real exchange rate and its components over a long span of time to speculate on its low-frequency properties. The top panel of Figure 2 plots the time series for log-prices in the UK and the US as well as the log-nominal dollar-sterling exchange rate over the sample period 1791-2000.[8] It is quite interesting how the price series move together–even without adjusting by the exchange rate to express prices in a common currency–over such a long period. It is also apparent how the biggest and more persistent wedge between the two prices seems to occur in the post-Bretton Woods period, essentially from the 1970s onwards. This wedge also coincided with the beginning of a corresponding trend in the nominal exchange rate, exactly as one would expect under PPP. The bottom panel of Figure 2 then graphs the log-real exchange rate constructed from these time series (in deviation from the mean). It is interesting how the real exchange rate appears to have a tendency to return to its long-run mean (a feature of a stationary process), although the mean is crossed only 20 times in over 200 years of data, indicating a remarkable degree of persistence. Further, the real exchange rate appears to be more persistent when it is in the proximity of the long-run mean, whereas reversion towards the mean happens more rapidly when the absolute size of the PPP deviation is large. This eyeball analysis of 200 years of real dollar-sterling therefore suggests that this real exchange rate may be stationary, albeit persistent, and that it is very persistent in the neighborhood of PPP, while being mean-reverting at a faster speed when the deviation from PPP gets larger. This is consistent with the existence of nonlinear dynamics in the real exchange rate, implying that the speed of mean reversion is state dependent. We now move from such a simplistic analysis of the time properties of the real exchange rate to a more formal treatment of the theoretical rationale and the empirics of nonlinear dynamics of real exchange rates.

3.2 Nonlinear dynamics in real exchange rates: rationale and implications

In the procedures conventionally applied to test for long-run PPP, the null hypothesis is usually that the process generating the real exchange rate series has a unit root, while the alternative hypothesis is that all of the roots of the process lie within the unit circle. The maintained hypothesis in the conventional framework assumes a linear autoregressive process for the real exchange rate, which means that adjustment is both continuous and of constant speed, regardless of the size of the deviation from PPP. However, the presence of transactions costs may imply a nonlinear process, which has important implications for the conventional unit root tests of long-run PPP. A number of authors have developed theoretical models of nonlinear real exchange rate adjustment arising from transactions costs in international arbitrage (e.g. Benninga and Protopapadakis, 1988; Dumas, 1992; Sercu, Uppal and Van Hulle, 1995; Obstfeld and Rogoff, 2000; Anderson and van Wincoop, 2004). In most of these models, proportional or 'iceberg' transport costs ('iceberg' because a fraction of goods are presumed to 'melt' when shipped) create a band for the real exchange rate within which the marginal cost of arbitrage exceeds the marginal benefit. Assuming instantaneous goods arbitrage at the edges of the band then typically implies that the thresholds become reflecting barriers.

Drawing on recent work on the theory of investment under uncertainty, some of these studies show that the thresholds should be interpreted more broadly than as simply reflecting shipping costs and trade barriers *per se*, but also as resulting from the sunk costs of international arbitrage and the resulting tendency for traders to wait for sufficiently large arbitrage opportunities to open up before entering the market (see in particular Dumas, 1992; Obstfeld and Rogoff, 2000; O'Connell and Wei, 2002).

Recently, Taylor (2001) has shown that empirical estimates of the half life of shocks to the real exchange rate may be biased upwards because of two empirical pitfalls. The first pitfall identified by Taylor relates to temporal aggregation in the data. Using a model in which the real exchange rate follows an AR(1) process at a higher frequency than that at which the data is sampled, Taylor shows analytically that the degree of upward bias in the estimated half life rises as the degree of temporal aggregation increases–i.e. as the length of time between observed data points increases. The second pitfall highlighted by Taylor concerns the possibility of nonlinear adjustment of real exchange rates. On the basis of Monte Carlo experiments with a nonlinear artificial data generating process, Taylor shows that there can also be substantial upward bias in the estimated half life of adjustment from assuming linear adjustment when in fact the true adjustment process is nonlinear. The time aggregation problem is a difficult issue for researchers to deal with since, as discussed above, long spans of data are required in order to have a reasonable level of power when tests of nonstationarity of the real exchange rate are applied, and long spans of high-frequency data do not exist. On the other hand, Taylor also shows that the problem becomes particularly acute when the degree of temporal aggregation exceeds the length of the actual half life, so that this source of bias may be mitigated somewhat if the researcher believes that the true half life is substantially greater than the frequency of observation.

Overall, these models suggest that the exchange rate will become increasingly mean reverting with the size of the deviation from the equilibrium level. A characterization of nonlinear adjustment, which allows for smooth rather than discrete adjustment, is in terms of a smooth transition autoregressive (STAR) model (Granger and Teräsvirta, 1993). In the STAR model, adjustment takes place in every period but the speed of adjustment varies with the extent of the deviation from parity. A STAR model for the real exchange rate q_t may be written as follows:

$$[q_t - \mu] = \sum_{j=1}^{p} \beta_j [q_{t-j} - \mu] + \left[\sum_{j=1}^{p} \beta_j^* [q_{t-j} - \mu]\right] \Phi[\theta; q_{t-d} - \mu] + \varepsilon_t \tag{5}$$

where $\varepsilon_t \sim iid(0, \sigma^2)$; The transition function $\Phi[\theta; q_{t-d} - \mu]$ determines the degree of mean reversion and is itself governed by the parameter $\theta > 0$, which effectively determines the speed of mean reversion, and the parameter μ which is the equilibrium level of q_t; the integer $d > 0$ denotes a delay parameter. A simple transition function suggested by Granger and Teräsvirta (1993) is the exponential function:

$$\Phi[\theta; q_{t-d} - \mu] = 1 - \exp\left[-\theta^2 [q_{t-d} - \mu]^2\right] \tag{6}$$

in which case (5) would be termed an exponential STAR or ESTAR model. The exponential transition function is bounded between zero and unity, $\Phi : \Re \rightarrow [0, 1]$, has the properties $\Phi[0] = 0$ and $\lim_{x \rightarrow \pm\infty} \Phi[x] = 1$, and is symmetrically inverse–bell shaped around zero. These properties of the ESTAR model are attractive in the present context because they allow a smooth transition between regimes and symmetric adjustment of the real exchange rate for deviations above and below the equilibrium level. The transition parameter θ determines the speed of transition between the two extreme regimes, with lower absolute values of θ implying slower transition. The inner regime corresponds to $q_{t-d} = \mu$, when $\Phi = 0$ and (5) becomes a linear AR(p) model: $[q_{t-d} - \mu] = \sum_{j=1}^{p} \beta_j [q_{t-j} - \mu] + \varepsilon_t$. The outer regime corresponds, for a given θ, to $\lim_{[q(t-d)-\mu] \rightarrow \pm\infty} \Phi[\theta; q_{t-d} - \mu]$, where (5) becomes a different AR(p) model: $[q_{t-d} - \mu] = \sum_{j=1}^{p} (\beta_j + \beta_j^*) [q_{t-j} - \mu] + \varepsilon_t$, with a correspondingly different speed of mean reversion so long as $\beta_j^* \neq 0$ for at least one value of j.

It is also instructive to reparameterize the STAR model (5) as

$$\begin{aligned} \Delta q_t \quad &= \quad \alpha + \rho q_{t-1} + \sum_{j=1}^{p-1} \phi_j \Delta q_{t-j} + \\ &\quad \left\{\alpha^* + \rho^* q_{t-1} + \sum_{j=1}^{p-1} \phi_j^* \Delta q_{t-j}\right\} \Phi[\theta; q_{t-d}] + \varepsilon_t \end{aligned} \tag{7}$$

where $\Delta q_{t-j} \equiv q_{t-j} - q_{t-j-1}$. In this form, the crucial parameters are ρ and ρ^*. Our above discussion of the effect of transactions costs suggests that the larger the deviation from PPP the stronger will be the tendency to move back to equilibrium. This implies that while $\rho \geq 0$ is admissible, we must have $\rho^* < 0$ and $(\rho + \rho^*) < 0$. That is, for small deviations q_t may be characterized by unit root or even explosive behavior, but for large deviations the process is mean reverting. This analysis has implications for the conventional test for a unit root in the real exchange rate process, which is based on a linear AR(p) model, written below as an augmented Dickey-Fuller regression:

$$\Delta q_t = \alpha' + \rho' q_{t-1} + \sum_{j=1}^{p-1} \phi_j' \Delta q_{t-j} + \varepsilon_t. \tag{8}$$

Assuming that the true process for q_t is given by the nonlinear model (7), estimates of the parameter ρ' in (8) will tend to lie between ρ and $(\rho+\rho^*)$, depending upon the distribution of observed deviations from the equilibrium level μ. Hence, the null hypothesis $H_0 : \rho' = 0$ (a single unit root) may not be rejected against the stationary linear alternative hypothesis $H_1 : \rho' < 0$, even though the true nonlinear process is globally stable with $(\rho + \rho^*) < 0$. Thus, failure to reject the unit root hypothesis on the basis of a linear model does not necessarily invalidate long-run PPP.

Note that the arguments made here to rationalize mean reversion in the real exchange rates are based on ideas that relate to the law of one price in the sense that refer to tradable goods only. However, we argue that this is reasonable given that Engel (1999), in a study that measures the proportion of dollar real exchange rate movements that can be accounted for by movements in the relative prices of nontradable goods, finds that relative prices of nontradable goods appear to account for essentially none of the movement of dollar real exchange rates. Hence, much of the explanation for the time series properties of PPP deviations is likely to reside in the behavior of deviations from the law of one price–i.e. movements in the relative prices of tradable goods.

3.3 The empirics of nonlinear reversion to PPP: where do we stand?

We now turn to the empirical evidence on nonlinear mean reversion in real exchange rates.[9] Michael, Nobay and Peel (1997) apply the ESTAR model to monthly interwar data for the French franc-US dollar, French franc-UK sterling and UK sterling-US dollar as well as for the Lothian and Taylor (1996) long span data set. Their results clearly reject the linear framework in favor of an ESTAR process. The systematic pattern in the estimates of the nonlinear models provides strong evidence of mean-reverting behavior for PPP deviations, and helps explain the mixed results of previous studies. However, the periods examined by MNP are ones over which the validity of long-run PPP is uncontentious (Taylor and McMahon, 1988; Lothian and Taylor, 1996).

Using data for the recent float, however, Taylor, Peel and Sarno (TSP) (2001) provide strong confirmation that four major real bilateral dollar exchange rates are well characterized by nonlinearly mean reverting processes. For example, the estimated model for dollar-sterling over the 1973-1996 sample period is as follows:

$$\widehat{q}_t = q_{t-1} - [1 - \exp\{-0.452\{q_{t-1} + 0.149\}^2\}][q_{t-1} + 0.149] \qquad (9)$$
$$(-2.771) \qquad (4.274) \qquad (4.274)$$

where the hat denotes the fitted value, figures in parentheses are t-ratios. The recorded R^2 for this simple nonlinear AR(1) is 0.94. This estimated model, which may be seen as representative of the results reported by TPS (2001), implies an equilibrium level of the real exchange rate in the neighborhood of which the behavior of the log-level of the real exchange rate is close to a random walk, becoming increasingly mean reverting with the absolute size of the deviation from equilibrium, consistent with the recent theoretical literature on the nature of real exchange rate dynamics in the presence of international arbitrage costs.

TPS also estimated the impulse response functions corresponding to their estimated nonlinear real exchange rate models by Monte Carlo integration.[10] By taking account of nonlinearities, TPS find the speed of real exchange rate adjustment to be typically much faster than the very slow speeds of real exchange rate adjustment hitherto recorded in the literature. For example, the estimated half lives (in months) for dollar-sterling and dollar-yen are the following

Shock (%):	40	30	20	10	5	1
dollar-sterling	10	20	22	26	29	32
dollar-yen	14	18	24	32	38	42

where in the first row we report the size of the shock (in percentage terms) to the log-level of the real exchange rate. The estimated half lives of these major real dollar exchange rates illustrate the nonlinear nature of the response to shocks, with larger shocks mean reverting much faster than smaller shocks. The dollar-sterling rate displays quite fast mean reversion, ranging from a half life of under one year for the largest shocks of forty percent to just under three years for small shocks of one percent; for shocks of five to ten percent, the half lives are just over two years. The dollar-yen displays higher persistence, with half lives ranging from fourteen to forty-two months.

These results therefore seem to shed some light on the PPP puzzles. Only for small shocks occurring when the real exchange rate is near its equilibrium do nonlinear models consistently yield half lives in the range of three to five years, which Rogoff (1996) terms 'glacial.' For dollar-sterling, even small shocks of one to five percent have a half life under three years; for larger shocks, the speed of mean reversion is even faster. However, while these estimates of the speed of reversion to

PPP appear more intuitively appealing than the estimates obtained from linear models, some authors maintain that they are still too long to be explained by conventional general equilibrium models with nominal rigidities (notably, see Rogoff, 2002).[11]

4 The exchange rate disconnect puzzle

4.1 The robust failure of exchange rate models

A logical way of examining the empirical ability of exchange rate models is to examine their out-of-sample forecasting performance. In a seminal paper, Meese and Rogoff (1983) compare the out-of-sample forecasts produced by various exchange rate models with forecasts produced by a random walk model, the forward exchange rate, a univariate regression of the spot rate, and a vector autoregression. They use rolling regressions to generate a succession of out-of-sample forecasts for each model and for various time horizons. The conclusion which emerges from this study is that, on a comparison of root mean square errors (RMSEs), none of the exchange rate models outperforms the simple random walk. This is the case even though actual future values of the right-hand-side variables are allowed in the dynamic forecasts (thereby giving the models a very large informational advantage).

A variant of the Meese-Rogoff approach involves employing a time-varying parameter model. In fact, the poor forecasting performance noted by Meese and Rogoff may be due to the fact the parameters in the estimated equations are unstable. This instability may be rationalized on a number of grounds, in response to policy regime changes as an example of a Lucas critique problem (Lucas, 1976), or because of implicit instability in the money demand or PPP equations, or else because of agents' heterogeneity leading to different responses to macroeconomic developments over time. For example, Schinasi and Swamy (1989) use a Kalman-filter maximum likelihood estimation technique to estimate time-varying parameter models which are found to outperform the random walk model of the exchange rate for certain time periods and currencies.

A general finding in this literature is that researchers have found that one key to improving forecast performance based on economic fundamentals lies in the introduction of equation dynamics. This has been done in various ways: by using dynamic forecasting equations for the forcing variables in the forward-looking, rational expectations version of the flexible-price monetary model, by incorporating dynamic partial adjustment terms into the estimating equation, by using time-varying parameter estimation techniques, and by using dynamic equilibrium correction forms (e.g. Koedijk and Schotman, 1990; MacDonald and Taylor, 1994). However, it remains true that most studies which claim to have beaten the random walk in out-of-sample forecasting turn out to be fragile in the sense that it is generally hard to replicate the superior forecasting performance for alternative periods and alternative currencies. Put another way, the Meese-Rogoff findings are extremely robust.

A related approach, due in the context of foreign exchange market analysis originally to Mark (1995), considers long-horizon predictability through analysis of equations of the form:

$$\Delta_k s_{t+k} = \alpha + \beta_k (z_t - s_t) + u_{t+k} \tag{10}$$

where z_t is an exchange rate fundamentals term, for example that suggested by the monetary class of models, $z_t \equiv [(m_t - m_t^*) - (y_t - y_t^*)]$, with m and y denoting money and income respectively; and u_{t+k} is a disturbance term. If the fundamentals in question help forecast the exchange rate, then we should find $\beta_k > 0$ and significantly different from zero. In a series of forecasting tests over long horizons for several quarterly dollar exchange rates, Mark finds that equation (10) may be able to predict the nominal exchange rate only at fairly long horizons, such as the four-year horizon. Moreover, both the goodness of in-sample fit and the estimated value of β_k rise as the horizon k rises. Mark interprets this as evidence that, while quarter-to-quarter exchange rate movements may be noisy, systematic movements related to the fundamentals become apparent in long-horizon changes.

In general, long-horizon regressions have been used extensively in the literature, but with mixed success (see Kilian, 1999). One reason may be that previous research has focused on linear models. In fact, in a linear world, it can be argued that there is no rationale for conducting long-horizon forecast tests. The problem is that under linearity k-step ahead forecasts are obtained by linear extrapolation from 1-step ahead forecasts. Thus, by construction there cannot be any gain in power at higher horizons (see Berkowitz and Giorgianni, 2001). However, the mounting evidence of nonlinear exchange rate dynamics provides a new rationale for the use of long-horizon regression tests, which remains under-researched.

Another relevant area of research which has, until recently, been under-researched involves the role of world commodity prices in determining exchange rates. Recently, Chen and Rogoff (2003) tackle the disturbing fact that standard exchange rate models cannot explain the high volatility and persistence observed in real exchange rates by investigating the determinants of real exchange rate movements for three economies (Australia, Canada, and New Zealand) where primary commodities constitute a significant share of their exports. The idea is that, because commodity products are transacted in highly centralized global markets, an exogenous source of terms of trade fluctuations can be identified for these major commodity exporters. For Australia and New Zealand especially, Chen and Rogoff find that the US dollar price of their commodity exports has a strong and stable influence on their floating real rates. However, after controlling for commodity price shocks, there is still a PPP puzzle in the residual. These results offer insight to developing commodity-exporting countries as they liberalize their capital markets and move towards floating exchange rates but also, more generally, to policy makers and central banks' staff interested in the behavior of exchange rates. Cashin (2003) discusses a variety of efforts made at the International Monetary Fund to study the stylized facts and economic consequences of movements in commodity prices and the terms of trade (e.g. Cashin, Cespedes and Sahay, 2004), while the 'Bank of Canada' equation supports the idea that there is a close long-run link–i.e. a cointegrating relationship–between commodity prices and the value of the Canadian dollar-US dollar exchange rate. Further work by Chen (2003) also suggests that, for primary commodity producers, nominal exchange rates exhibit a robust response to movements in the world prices of their corresponding commodity exports and that incorporating commodity export prices into standard exchange rate models can generate a marked improvement in their in-sample performance. In terms of out-of-sample forecasting, however, while commodity-price-augmented specifications offer some evidence of exchange rate predictability, no single specification is found to provide a consistent forecast improvement over a random walk at all horizons and across all currency pairs. Overall, the empirical evidence shown in this line of research has established that there is a close link between commodity prices and (nominal and real) exchange rates, and the potential for this link to shed light–theoretically and empirically–on the exchange rate disconnect puzzle may not have been realized yet.

4.2 Why is there a disconnect puzzle? Some speculation

One obvious problem is that three of the building blocks of the monetary model–money demand equations, PPP and UIP–do not work very well. Money demand equations have proven unstable, especially in the US (Friedman and Kuttner, 1992), but changing the numeraire currency doesn't seem to help the monetary model much. PPP, which predicts that differences in countries' inflation rates should be reflected in changes in the exchange rate, doesn't describe exchange rate behavior very well at short horizons (Sarno and Taylor, 2002). And UIP has performed even more poorly (Engel, 1996).

But that again begs the question as to why PPP and UIP perform so poorly, which we partly addressed in the previous sections. Why are floating exchange rates so volatile and unrelated to prices and interest differentials? Many researchers have claimed that volatile expectations or departures from rationality are likely to account for the failure of exchange rate models. For example, Frankel (1996) argues that exchange rates are detached from fundamentals by swings in expectations about future values of the exchange rate. Four pieces of evidence suggest that expectations are to blame for such behavior. (i) Survey measures of exchange rate expectations are very poor forecasters and the expectations, themselves, are frequently not internally consistent (Frankel and Froot, 1987; Sarno and Taylor, 2003). (ii) Failure of expectations to be rational is often blamed for the failure of UIP (Engel, 1996). (iii) Trend following trading rules appear to make risk-adjusted excess returns, in apparent violation of the efficient markets hypothesis (Neely, 1997; Neely, Weller, and Dittmar, 1997). (iv) Switching from a fixed exchange rate to a floating rate—which changes the way expectations are formed—changes the behavior of nominal and real exchange rates and the ability of UIP to explain exchange rate changes (Neely and Sarno, 2002).

This latter point requires some explanation. Fixed exchange rates anchor investor sentiment about the future value of a currency because of the government's commitment to stabilize its value. If expectations are based on fundamentals, rather than irrationally changing expectations, then the relationship between fundamentals and exchange rates should be the same under a fixed exchange rate regime as it is under a floating regime. This is not the case. Countries that move from floating exchange rates to fixed exchange rates experience a dramatic change in the relationship between prices and exchange rates. Specifically, real exchange rates are much more volatile under floating exchange rate regimes, where expectations are not tied down by promises of government intervention (Mussa, 1986). This result suggests that, contrary to the efficient markets hypothesis, swings in investor expectations may detach exchange rates from fundamental values in the short run. Similarly, UIP seems to do such a poor job explaining floating exchange rates while doing a pretty good job with semi-fixed rates such as those found in the European Monetary System (Flood and Rose, 1996).

4.3 Economic value versus statistical significance

Prior research on the ability of monetary-fundamentals models to forecast exchange rates relies on *statistical* measures of forecast accuracy, like RMSEs. Surprisingly little attention has been directed, however, to assessing whether there is any *economic* value to exchange rate predictability (i.e., to using a model where the exchange rate is forecast using economic fundamentals). An important exception is the study of West, Edison and Cho (1993), who focus on a utility-based metric of forecast evaluation rather than conventional statistical criteria. However, their focus is on exchange rate volatility, whereas the typical focus of the literature inspired by Meese and Rogoff (1983) is on the level of the exchange rate. Specifically, West, Edison and Cho examine various time series models for the conditional variance of the exchange rate change. They don't analyze the relationship between exchange rates and macroeconomic fundamentals (say, money and income) implied by exchange rate determination theory, focusing instead on weekly dollar exchange rates and the corresponding pairs of Eurodeposit rates.[12] Nevertheless, their paper is different from the standard literature on exchange rate forecasting in that they propose departing from standard statistical measures of forecast accuracy.

Building on these ideas, Abhyankar, Sarno and Valente (2005) investigate the ability of a monetary-fundamentals model to predict exchange rates by measuring the economic or utility-based value to an investor who relies on this model to allocate her wealth between two assets that are identical in all respects except the currency of denomination. They focus on two key questions. First, as a preliminary to the forecasting exercise, they ask how exchange rate predictability affects optimal portfolio choice for investors with a range of horizons up to ten years. Second, and more importantly, they ask whether there is any additional economic value to a utility-maximizing investor who uses exchange rate forecasts from a monetary-fundamentals model relative to an investor who uses forecasts from a naive random walk model. The economic value of predictability is quantified in a Bayesian framework that allows the researchers to account for uncertainty surrounding parameter estimates in the forecasting model. Indeed, parameter uncertainty or 'estimation risk' is likely to be of importance, especially over long horizons.

The results, obtained using three major US dollar exchange rates–namely the Canadian dollar, the UK sterling and the Japanese yen–during the recent float and considering forecast horizons from 1 to 10 years, are as follows. First, exchange rate predictability substantially affects, both quantitatively and qualitatively, the choice between domestic and foreign assets for all currencies and across different levels of risk aversion. Specifically, exchange rate predictability can generate optimal weights to the foreign asset that are substantially different (in magnitude and, sometimes, in sign) from the optimal weights generated under a random walk model. Second, the main result is that there is evidence of economic value to exchange rate predictability across all exchange rates examined and for a wide range of plausible levels of risk aversion. In particular, the realized end-of-period wealth achieved by a US investor over a ten-year horizon using a monetary fundamentals-exchange rate model for forecasting the exchange rate is higher than the corresponding end-of-period wealth obtained by an investor who acts as if the exchange rate were a random walk. These results show that the economic value of predictability can be substantial also over relatively short horizons and across different levels of risk aversion. These findings may be viewed as suggesting that the case against the predictive power of monetary fundamentals may be overstated.

4.4 Insights from the microstructure approach to exchange rates

The market microstructure approach to exchange rate economics represents a drastic departure from the conventional macroeconomic approach to exchange rate economics (Lyons, 2001). In fact, the microstructure approach has relaxed some of the most controversial assumptions underlying traditional exchange rate economics and, specifically, has given emphasis to the role of heterogeneity of agents, to the fact that public information is not the only source of information that matters for exchange rate determination, and to the importance of the institutional details within which foreign exchange trading takes place.

An interesting literature in empirical microstructure has shown that time-aggregated order flow variables may have (much) more explanatory power than macroeconomic variables in explaining exchange rate behavior. 'Order flow' in this context is taken to be a variant of the more familiar concept of 'net demand,' and measures the net of buyer-initiated orders and seller-initiated orders. As noted by Lyons (2001), it is a variant of, rather than a synonym for, 'net demand' because in equilibrium order flow does not necessarily equal zero. A notable study in this literature is due to Evans and Lyons (2002), who provide a model which sheds light on the role of order flow in determining exchange rates. In their model, order flow is a proximate determinant of prices since it aggregates disperse information that currency markets need to aggregate–anything pertaining to the realization of uncertain demands (differential interpretation of news, shocks to hedging demands, shocks to liquidity demands, etc.). Using four months of daily data on signed order flow for Deutche mark-dollar and yen-dollar from the Reuters dealing system, Evans and Lyons provide evidence that order flow is a significant determinant of some major bilateral exchange rates, obtaining coefficients of determination substantially larger than the ones usually obtained using standard macroeconomic models of nominal exchange rates. Essentially, the R^2 increases from 1-5% for a regression of the exchange rate change on interest rate differentials to 40-60% in a regression which also uses order flow to explain the daily variation in exchange rates.

In a simplistic micro-macro dichotomy, one may view the standard macro approach to exchange rates as based on the view that only public information matters for exchange rates, and the micro approach as based on the view that private information is crucial for understanding exchange rates. However, neither of these extreme views is likely to be correct, whereas an hybrid view seems much more plausible.

The finding that order flow has more explanatory power than macro variables in explaining exchange rate behavior is interesting and has a fairly clear interpretation in terms of expectations formation mechanisms. Specifically, this finding does not necessarily imply that order flow is the underlying driver of exchange rates. Indeed, it may well be that macroeconomic fundamentals are an important underlying driving force, but that conventional measures of the macroeconomic fundamentals are so imprecise that an order-flow "proxy" performs better in estimation. This interpretation as a proxy is particularly plausible with respect to expectations–that is, even if macro variables fully describe the true model, when implemented empirically these variables may provide a poor measure of *expected* future fundamentals. Thus, it may be that order flow provides a more precise proxy for variation in these expectations. In this sense, unlike expectations measured from survey data, order flow represents a willingness to back one's beliefs with real money (see the discussion by Lyons, 2002), and the Evans-Lyons results may be seen as suggesting that both public and private information matter for exchange rate determination and that information impacts on prices not only directly but also via order flow (see Evans and Lyons, 2004). In this sense, the Evans-Lyons story is one where traditional macro analysis is augmented with simple price determination microeconomics.

5 Concluding remarks

Exchange rate economics is alive and continues to attract the attention of academics, policy-makers, and practitioners. One reason this field receives attention is the large number of puzzles the profession has been unable to resolve. This paper has reviewed three fundamental puzzles in the literature on exchange rate economics. The common feature of these puzzles is the missing link between exchange rates and some of the key variables which are expected, in theory, to explain or predict them–interest rates, forward rates, relative prices, money and income.

We begun the paper with a review of the literature testing UIP as a means to test foreign exchange market efficiency. This literature has established simply too clearly that the forward exchange rate is anything but an unbiased predictor of the future exchange rate, resulting in the conventional wisdom that the foreign exchange market is characterized by massive inefficiency. There are, however, results that suggest these rejections may not be as strong as this vast body of literature suggests. In particular, while the forward rate is likely to be a biased predictor of the future nominal exchange rate, the term structure of forward premia appears to contain some valuable information about future exchange rate movements that may be easily captured using conventional partial adjustment models. Moreover, there is an increasingly large body of evidence suggesting that bilateral exchange rates contain important nonlinearities, which are theoretically motivated and may be capable of enhancing the explanatory and predictive power of forward rates. Indeed, exchange rates appear to be linked nonlinearly to forward premia in the context of a model for UIP deviations which allows for time-variation in the forward bias and nonlinear reversion towards UIP, consistent with the presence of limits to speculation in the foreign exchange market. These findings raise the question of whether the statistical rejection of UIP recorded by the literature is really indicative of major inefficiencies in the foreign exchange market and are likely to spur further research attempting to distinguish between the statistical and the economic significance of the rejection of the market efficiency hypothesis.

With respect to the behavior of the real exchange rate, the theoretical literature focusing on the importance of international trade costs has led researchers to consider nonlinear models of real exchange rates that give a role to the size of the deviation from the PPP equilibrium in determining the speed of mean reversion in the real exchange rate, a feature that cannot be captured within a linear framework. The empirical evidence provided by these studies suggests that major real bilateral dollar exchange rates are well characterized by nonlinearly mean reverting processes over the floating rate period since 1973 and over century-long periods of time. These models imply an equilibrium level of the real exchange rate in the neighborhood of which the behavior of the real exchange rate is close to a random walk, becoming increasingly mean reverting with the absolute size of the deviation from equilibrium. In addition, the half lives of shocks to the real exchange rates implied by these models suggest much faster real exchange rate adjustment than typically recorded in the literature, hence shedding some light on Rogoff's (1996) PPP puzzle.

A number of questions still remain, however. In particular, further work on real exchange rate behavior might usefully be addressed to unravel the relative contribution of prices and nominal exchange rates to movements in real exchange rates (see, for example, Engel and Morley, 2001; Cheung, Lai and Bergman, 2004; Sarno and Valente, 2004). This might be done, for example, in the context of nonlinear vector equilibrium correction models of the nominal exchange rate and domestic and foreign prices and other variables. Such a framework might also be extended to allow for the relative impact of monetary and fiscal policy as well as relative productivity on real exchange rate movements to be isolated. Also, the implications of nonlinearities in real exchange rate movements for exchange rate forecasting and, in turn, the influence of official exchange rate intervention in generating exchange rate nonlinearites, have yet to be fully examined. Finally, it seems tantalizing to investigate the existence of possible interactions between the nonlinearities in UIP deviations and PPP deviations in order to address whether the large misalignments that are presumably responsible for generating reversion towards these parity conditions contain common features.

With respect to the ability of empirical exchange rate models to explain and forecast the nominal exchange rate, the literature is still somewhat in the dark. About 20 years after Meese and Rogoff's (1983) paper, their findings that empirical exchange rate models are not able to beat a random walk have not been convincingly overturned. Economic fundamentals typically suggested by open-economy macro theory do not appear to contain sufficient information as to provide satisfactory out-of-sample forecasts of the exchange rate, especially at short horizons. However, as noted above, the information contained in the term structure of forward premia appears to be more useful, especially in empirical models that allow for nonlinearities. This calls for further research trying to shed light on the nature of the information embedded in forward premia and on the way forward markets aggregate such information. The growing literature linking world commodity prices to nominal and real exchange rates also highlights another empirically important determinant of exchange rate dynamics–namely world commodity prices–which was unexplored until recently.

This literature has been dominated by the use of two crucial elements: (i) a random walk as the key benchmark to beat; and (ii) a symmetric function of the forecasts errors and, more generally, conventional statistical criteria for testing point forecast accuracy to decide whether an empirical exchange rate model forecasts satisfactorily. While (i) is clearly an important element in this research since the behavior of nominal exchange rates is close to a random walk, the importance of (ii) is perhaps more questionable. In fact, some researchers (West, Edison and Cho, 1993; Abhyankar, Sarno and Valente, 2005) have departed from standard statistical criteria in favor of measures of economic or utility-based value of an exchange rate forecasting model. In some sense, the point of this strand of research is that, if one knows the reason why a forecasting model is employed in the first place, then it seems logical to assess its validity on the basis of the specific objective function that captures that reason, rather than relying merely on a statistical criterion. This research suggests that the link between fundamentals and exchange rates is indeed strong and gets stronger as time goes by, in contrast with the stylized facts leading to the exchange rate disconnect puzzle.

Moreover, although the relevant literature has traditionally focused on accuracy evaluations based on point forecasts, several authors have recently emphasized the importance of evaluating the forecast accuracy of economic models on the basis of density—as opposed to point—forecasting performance. Especially when evaluating nonlinear models, which are capable of producing highly non-normal forecast densities, it would seem appropriate to consider a model's density forecasting performance. Some recent evidence attempts to discriminate between competing forecasting models in terms of density forecast accuracy, and the potential use of this framework in the context of exchange rate forecasting seems appealing especially when models are highly nonlinear since the benefits of using a nonlinear model relative to a linear one is likely to be found in the forecasts of the higher moments of the exchange rate distribution. In other words, not only we care about forecasting the level of the exchange rate, but also about measuring accurately the uncertainty surrounding these forecasts, which requires density forecasts. This is particularly important for risk management and also for the policy-maker interested in forecasting the probability of large changes in the exchange rate (depreciations or appreciations). The probability of large exchange rate movements risks to be poorly estimated when relying on linear models.

The microstructure literature has then added new insights into the exchange rate debate. One important finding of this literature is the existence of a strong link between daily exchange rates and order flow. The latter variable can explain fractions of the variation in exchange rates which simply cannot be explained using fundamentals-based models. Understanding the information driving order flow is therefore an important avenue for future research.

Overall, we end this paper with some cautionary degree of optimism. While the profession is still facing a number of puzzles, the research covered in this article should encourage us that there are rays of hope in a variety of contexts that reflect some improvement in our understanding of exchange rate movements.

References

Abhyankar, Abhay, Sarno, Lucio, and Giorgio Valente (2005) 'Exchange Rates and Fundamentals: Evidence on the Economic Value of Predictability,' *Journal of International Economics*, forthcoming.

Abuaf, Niso, and Philippe Jorion (1990) 'Purchasing Power Parity in the Long Run,' *Journal of Finance*, 45, 157-74.

Adler, Michael, and Bernard Dumas (1983) 'International Portfolio Choice and Corporation Finance: A Synthesis,' *Journal of Finance*, 38, 925-984.

Anderson, James E., and Eric van Wincoop (2004) 'Trade Costs,' Boston College and University of Virginia, mimeo.

Backus, David K., Gregory, AllanW., and Chris I. Telmer (1993) 'Accounting for Forward Rates in Markets for Foreign Currency,' *Journal of Finance*, 48, 1887-1908.

Baldwin, Richard (1990) 'Re-Interpreting the Failure of Foreign Exchange Market Efficiency Tests: Small Transaction Costs, Big Hysteresis Bands,' *National Bureau of Economic Research Working Paper N. 3319*.

Bansal, Ravi (1997) 'An Exploration of the Forward Premium Puzzle in Currency Markets,' *Review of Financial Studies*, 10, 369-403.

Bansal, Ravi and Dahlquist, Magnus (2000) 'The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies,' *Journal of International Economics*, 51, 115-44.

Barnhart, Scott W., McKnown, Robert, and Myles S. Wallace (1999) 'Non-Informative Tests of the Unbiased Forward Exchange Rate,' *Journal of Financial and Quantitative Analysis*, 34, 265-291.

Baxter, Marianne, and Alan C. Stockman (1989) 'Business Cycles and the Exchange Rate System,' *Journal of Monetary Economics*, 23, 377-400.

Bekaert, Geert (1994) 'The Time Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective,' *Review of Financial Studies*, 9, 427-470.

Bekaert, Geert, and Robert J. Hodrick (1992) 'Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets,' *Journal of Finance*, 47, 467-509.

Bekaert, Geert, and Robert J., Hodrick (2001) 'Expectations Hypotheses Tests,' *Journal of Finance*, 56, 1357-94.

Berkowitz, Jeremy (2001) 'Testing Density Forecasts, with Applications to Risk Management,' *Journal of Business and Economic Statistics*, 19, 465-74.

Berkowitz, Jeremy, and Lorenzo Giorgianni (2001) 'Long-Horizon Exchange Rate Predictability?', *Review of Economics and Statistics*, 83, 81-91.

Benninga, Simon, and Aris A. Protopapadakis (1988) 'The Equilibrium Pricing of Exchange Rates and Assets When Trade Takes Time,' *Journal of International Economics*, 7, 129-49.

Bilson, John F.O. (1981) 'The 'Speculative Efficiency' Hypothesis,' *Journal of Business*, 54, 435-451.

Campbell, John Y., Lo, Andrew W., and A. Craig MacKinlay (1997) *The Econometrics of Financial Markets*, Princeton, New Jersey: Princeton University Press.

Cashin, Paul (2003) 'Commodity Prices and the Terms of Trade,' *International Monetary Fund Research Bulletin*, 4, 1-4.

Cashin, Paul, Céspedes, Luis, and Ratna Sahay (2004) 'Commodity Currencies and the Real Exchange Rate,' *Journal of Development Economics*, 75, 239-268.

Cassel, Gustav (1922) *Money and Foreign Exchange After 1914*, London: Constable.

Chen, Yu-Chin (2003) 'Exchange Rates and Fundamentals: Evidence from Commodity Economies,' University of Washington, mimeo.

Chen, Yu-Chin, and Rogoff, Kenneth (2003) 'Commodity Currencies,' *Journal of International Economics*, 60, 133-60.

Chinn, Menzie D., and Guy Meredith (2004) 'Monetary Policy and Long-Horizon Uncovered Interest Parity,' *International Monetary Fund Staff Papers*, 51, 409-30.

Cheung, Yin-Wong, Lai, Kon S., and Michael Bergman (2004) 'Dissecting the PPP Puzzle: The Unconventional Roles of Nominal Exchange Rate and Price Adjustments,' *Journal of International Economics*, 64, 135-150.

Choudhry, Taufiq, McNown, Robert, and Myles S. Wallace (1991) 'Purchasing Power Parity and the Canadian Float in the 1950s,' *Review of Economics and Statistics*, 73, 558-563.

Clarida, Richard H., Sarno, Lucio, Taylor, Mark P., and Giorgio Valente (2003) 'The Out-of-Sample Success of Term Structure Models as Exchange Rate Predictors: a Step Beyond,' *Journal of International Economics*, 60, 61-83.

Clarida, Richard H., and Mark P. Taylor (1997) 'The Term Structure of Forward Exchange Premiums and the Forecastability of Spot Exchange Rates: Correcting the Errors,' *Review of Economics and Statistics*, 79, 353-361.

Crucini, Mario, Telmer, Chris, and Marios Zachariadis (2005) 'Understanding European Real

Exchange Rates,' *American Economics Review*, forthcoming.

Cumby, Robert E., and Maurice Obstfeld (1984) 'International Interest Rate and Price Level Linkages under Flexible Exchange Rates: A Review of Recent Evidence,' in Bilson, John F.O. and Richard C. Marston (eds.) *Exchange Rate Theory and Practice*, National Bureau of Economic Research Conference Report, Chicago and London: University of Chicago Press, 121-151.

Diebold, Francis X., Gunther, Todd A., and Anthony S. Tay (1998) 'Evaluating Density Forecasts with Applications to Financial Risk Management,' *International Economic Review*, 39, 863-83.

Diebold, Francis X., Husted, Steven, and Mark Rush (1991) 'Real Exchange Rates under the Gold Standard,' *Journal of Political Economy,* 99, 1252-1271.

Dumas, Bernard (1992) 'Dynamic Equilibrium and the Real Exchange Rate in Spatially Separated World,' *Review of Financial Studies*, 5, 153-180.

Engel, Charles (1992) 'On the Foreign Exchange Risk Premium in a General Equilibrium Model,' *Journal of International Economics*, 32, 305-319.

Engel, Charles (1996) 'The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence,' *Journal of Empirical Finance*, 3, 123-192.

Engel, Charles (1999) 'Accounting for US Real Exchange Rate Changes,' *Journal of Political Economy*, 107, 507-538.

Engel, Charles, and James C. Morley (2002) 'The Adjustment of Prices and the Adjustment of the Exchange Rate,' University of Wisconsin and Washington University, mimeo.

Fama, Eugene F. (1984) 'Forward and Spot Exchange Rates,' *Journal of Monetary Economics*, 14, 319-38.

Flood, Robert B., and Andrew K. Rose (1996) 'Fixes: Of the Forward Discount Puzzle,' *Review of Economics and Statistics*, 78, 748-52.

Flood, Robert B., and Andrew K. Rose (2002) 'Uncovered Interest Parity in Crisis,' *International Monetary Fund Staff Papers*, 49, 252-66.

Flood, R.B., and Taylor, M.P. (1996) 'Exchange Rate Economics: What's Wrong with the Conventional Macro Approach?', in Frankel, Jeffrey A., Galli, Giampaolo, and Alberto Giovannini (eds.) *The Microstructure of Foreign Exchange Markets*, National Bureau of Economic Research Conference Report series. Chicago and London: University of Chicago Press, 261-94.

Frankel, Jeffrey A. (1982) 'A Test of Perfect Substitutability in the Foreign Exchange Market,' *Southern Economic Journal*, 49, 406-416.

Frankel, Jeffrey A. (1986) 'International Capital Mobility and Crowding out in the US Economy: Imperfect Integration of Financial Markets or Goods Markets?', in Hafer, Rick W. (ed.) *How Open is the US Economy?*, Lexington: Lexington Books.

Frankel, Jeffrey A. (1990) 'Zen and the Art of Modern Macroeconomics: The Search for Perfect Nothingness,' in Haraf, W., and Willett, T. (eds.) *Monetary Policy for a Volatile Global Economy*, Washington DC: American Enterprise Institute.

Frankel, Jeffrey A. (1996) 'How Well Do Foreign Exchange Markets Work: Might a Tobin Tax Help?' in Mahbub ul Haq, Inge Kaul, and Isabelle Grunberg, (eds.) *The Tobin Tax: Coping with Financial Volatility*, New York and Oxford: Oxford University Press.

Frankel, Jeffrey A., and Kenneth Froot (1987) 'Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations,' *American Economic Review*, 77, 133-153.

Frankel, Jeffrey A., and Andrew K. Rose (1996) 'A Panel Project on Purchasing Power Parity: Mean Reversion Within and Between Countries,' *Journal of International Economics*, 40, 209-24.

Friedman, Benjamin M., and Kenneth N. Kuttner (1992) 'Money, Income, Prices, and Interest Rates,' *American Economic Review*, 82, 472-92.

Froot, Kenneth A., and Jeffrey A. Frankel (1989) 'Forward Discount Bias: Is It an Exchange Risk Premium?', *Quarterly Journal of Economics*, 104, 139-161.

Froot, Kenneth, and Kenneth Rogoff (1995) 'Perspectives on PPP and Long-run Real Exchange Rates,' in Rogoff, K., and G. Grossman (eds.) *Handbook of International Economics*, Amsterdam: North Holland.

Froot, Kenneth, and Takatoshi Ito (1990) 'On the Consistency of Short-Run and Long-Run Exchange Rate Expectations,' *Journal of International Money and Finance*, 48, 487-510.

Froot, Kenneth A., and Richard H. Thaler (1990) 'Foreign Exchange,' *Journal of Economic Perspectives*, 4, 179-92.

Giovannini, Alberto, and Philippe Jorion (1989) 'The Time Variation of Risk and Return in the Foreign Exchange and Stock Markets,' *Journal of Finance*, 44, 307-325.

Granger, Clive W.J. (2003) 'Time Series Concepts for Conditional Distributions,' *Oxford Bulletin of Economics and Statistics*, 65, 689-701.

Granger, Clive W.J., and Hashem H. Pesaran (1999) 'A Decision Theoretic Approach to Forecast Evaluation,' in: Chan, Wai-Sum, Lin, Wai Keung and Howell Tong (eds.) *Statistics and Finance: An Interface*, London: Imperial College Press.

Granger, Clive W.J., and Timo Teräsvirta (1993) *Modelling Nonlinear Economic Relationships*, Oxford; New York; Toronto and Melbourne: Oxford University Press.

Gourinchas, Pierre-Olivier, and Aaron Tornell (2004) 'Exchange Rate Puzzles and Distorted Beliefs,' *Journal of International Economics*, 63, 303-33.

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

Hodrick, Robert J., and Sanjay Srivastava (1984) 'An Investigation of Risk and Return in Forward Foreign Exchange,' *Journal of International Money and Finance*, 3, 5-29.

Hollifield, Burton, and Raman Uppal (1997) 'An Examination of Uncovered Interest Parity in Segmented International Commodity Markets,' *Journal of Finance*, 52, 2145-2170.

Huisman, Ronald, Koedijk, Kees G., Kool, Clemens, and Francois Nissen (1998) 'Extreme Support for Uncovered Interest Parity,' *Journal of International Money and Finance*, 17, 211-28.

Imbs, Jean, Mumtaz, Haaron, Ravn, Morten O., and Helene Rey (2005),'PPP Strikes Back: Aggregation and the Real Exchange Rate,' *Quarterly Journal of Economics*, forthcoming.

Jorion, Philippe (2001) 'Value at Risk: The New Benchmark for Managing Financial Risk,' London and New York: McGraw-Hill.

Kilian, Lutz (1999) 'Exchange Rates and Monetary Fundamentals: What Do We Learn from Long-Horizon Regressions?,' *Journal of Applied Econometrics*, 14, 491-510.

Koedijk, Kees G., and Peter Schotman (1990) 'How to Beat the Random Walk: An Empirical Model of Real Exchange Rates,' *Journal of International Economics*, 29, 311-32.

Koedijk, Kees G., Tims, Ben, and Mathijs A. van Dijk (2004) 'Purchasing Power Parity and the Euro Area,' *Journal of International Money and Finance*, 23, 1081-107.

Lafrance, Robert, Osakwe, Patrick, and Michel Normandin (1998) 'Evaluating Alternative Measures of the Real Effective Exchange Rate,' Bank of Canada, Working Paper 98-20.

Lane, Philip R. (2001) 'The New Open Economy Macroeconomics: A Survey,' *Journal of International Economics*, 54, 235-66.

Lewis, Karen K. (1989a) 'Can Learning Affect Exchange-Rate behaviour? The Case of the Dollar in the Early 1980's,' *Journal of Monetary Economics*, 23, 79-100.

Lewis, Karen K. (1989b) 'Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange,' *American Economic Review*, 79, 621-36.

Lewis, Karen K. (1995) 'Puzzles in International Financial Markets,' Grossman, G. M., and Rogoff,

K. (eds.) *The Handbook of International Economics*, vol. 3. Amsterdam; New York and Oxford: Elsevier, North-Holland, 1913-71.

Liu, Peter C., and Maddala, G.S. (1992) 'Rationality of Survey Data and Tests for Market Efficiency in the Foreign Exchange Markets,' *Journal of International Money and Finance*, 11, 366-81.

Lothian, James R., and Mark P. Taylor (1996) 'Real Exchange Rate behavior: The Recent Float from the Perspective of the Past Two Centuries,' *Journal of Political Economy*, 104, 488-510.

Lothian, James R., and Mark P. Taylor (1997) 'Real Exchange Rate behavior: The Problem of Power and Sample Size,' *Journal of International Money and Finance*, 16, 945-954.

Lucas, Robert E. Jr. (1976) 'Econometric Policy Evaluation: A Critique,' *Journal of Monetary Economics*, 1, Supplementary Series, 19-46.

Lucas, Robert E. Jr. (1982) 'Interest Rates and Currency Prices in a Two-Country World,' *Journal of Monetary Economis*, 10, 335-359.

Lyons, Richard K. (2001) *The Microstructure Approach to Exchange Rates*, Cambridge and London: MIT Press.

Lyons, Richard K. (2002) 'The Failure of Empirical Exchange Rate Models: Room for Optimism?', *Economic Policy*, Web Essay.

MacDonald, Ronald and Mark P. Taylor (1994) 'The Monetary Model of the Exchange Rate: Long-Run Relationships, Short-Run Dynamics and How to Beat a Random Walk,' *Journal of International Money and Finance*, 13, 276-90.

Mark, Nelson C. (1985) 'On Time Varying Risk Premia in the Foreign Exchange Market: An Econometric Analysis,' *Journal of Monetary Economics*, 16, 3-18.

Mark, Nelson C. (1990) 'Real and Nominal Exchange Rates in the Long Run: An Empirical Investigation,' *Journal of International Economics*, 28, 115-36.

Mark, Nelson C. (1995) 'Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability,' *American Economic Review*, 85, 201-18.

Mark, Nelson C., and Young-Kyu Moh (2002) 'Continuous-Time Market Dynamics, ARCH Effects, and the Forward Premium Anomaly,' University of Notre Dame and Tulane University, mimeo.

McCallum, Bennett T. (1994) 'A Reconsideration of the Uncovered Interest Parity Relationship,' *Journal of Monetary Economics*, 33, 105-132.

McCloskley, Deidre N., and J. Richard Zecher (1984) 'The Success of Purchasing Power Parity: Historical Evidence and Its Implications for Macroeconomics,' in Bordo, Michael D., and Anna J.

Schwartz (eds.) *A Retrospective on the Classical Gold Standard, 1821-1931*, Chicago: Chicago University Press, 121-150.

McNown, Robert, and Myles S. Wallace (1989) 'National Price Levels, Purchasing Power Parity, and Cointegration: A Test of Four High Inflation Economies,' *Journal of International Money and Finance*, 8, 533-545.

Meese, Richard A., and Kenneth Rogoff (1983) 'Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?', *Journal of International Economics*, 14, 3-24.

Michael, Panos, Nobay, Avelino R., and David A. Peel (1997) 'Transactions Costs and Nonlinear Adjustment in Real Exchange Rates: An Empirical Investigation,' *Journal of Political Economy*, 105, 862-879.

Mussa, Michael (1986) 'Nominal Exchange Rate Regimes and the Behavior of Real Exchange Rates: Evidence and Implications,' in Brunner, Karl and Meltzer, Allan H. (eds.) *Carnegie-Rochester Conference Series on Public Policy: Real Business Cycles, Real Exchange Rates, and Actual Policies*, 25, 117-214.

Neely, Christopher J. (1997) 'Technical Analysis in the Foreign Exchange Market: A Layman's Guide,' *Federal Reserve Bank of St. Louis Review*, 79, 23-38.

Neely, Christopher J., and Lucio Sarno (2002) 'How Well Do Monetary Fundamentals Forecast Exchange Rates?', *Federal Reserve Bank of St. Louis Review*, 84, 51-74.

Neely, Christopher J., Weller, Paul, and Robert Dittmar (1997) 'Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach,' *Journal of Financial and Quantitative Analysis*, 32, 405-26.

Obstfeld, M., and Rogoff, K. (2000) 'The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?', in Bernanke, B. and Rogoff, K. (eds.) *NBER Macroeconomics Annual,* Cambridge: MIT Press, 339-390.

O'Connell, Paul G.J. (1998) 'The Overvaluation of Purchasing Power Parity,' *Journal of International Economics*, 44, 1-19.

O'Connell, Paul G.J., and Shang-Jin Wei (2001) 'The Bigger They Are, The Harder They Fall: Retail Price Differences across U.S. Cities,' *Journal of International Economics*, 56, 21-53.

Papell, David H. (1998) 'Searching for Stationarity: Purchasing Power Parity Under the Current Float,' *Journal of International Economics*, 43, 313-332.

Rogoff, Kenneth (1979) *Expectations and Exchange Rate Volatility*, unpublished PhD thesis, Massachusets Institute of Technology.

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

Rogoff, Kenneth (2002) 'The Failure of Empirical Exchange Rate Models: No Longer New, But Still True,' *Economic Policy*, Web Essay.

Samuelson, Paul (1964) 'Theoretical Notes on Trade Problems,' *Review of Economics and Statistics,* 46, 145-54.

Sarno, Lucio (2001) 'Toward a New Paradigm in Open Economy Modeling: Where Do We Stand?', *Federal Reserve Bank of St. Louis Review*, 83, 21-36.

Sarno, Lucio, and Chowdhury, Ibrahim (2003) 'The Behaviour of the Real Exchange Rate: Evidence from an Alternative Price Index,' *Economic Notes*, 32, 295-333.

Sarno, Lucio, and Mark P. Taylor (1998) 'Real Exchange Rates Under the Recent Float: Unequivocal Evidence of Mean Reversion,' *Economics Letters*, 60, 131-37.

Sarno, Lucio, and Mark P. Taylor (2002) 'Purchasing Power Parity and the Real Exchange Rate,' *International Monetary Fund Staff Papers*, 49, 65-105.

Sarno, Lucio, and Mark P. Taylor (2003) *The Economics of Exchange Rates*, Cambridge: Cambridge University Press.

Sarno, Lucio, and Giorgio Valente (2004) 'Deviations from Purchasing Power Parity Under Different Exchange Rate Regimes: Do They Revert and, If So, How?,' University of Warwick, mimeo.

Sarno, Lucio, and Giorgio Valente (2005) 'Empirical Exchange Rate Models and Currency Risk: Some Evidence from Density Forecasts,' *Journal of International Money and Finance*, forthcoming.

Sarno, Lucio, Valente, Giorgio, and Hyginus Leon (2004) 'The Forward Bias Puzzle and Nonlinearity in Deviations from Uncovered Interest Parity: A New Perspective,' University of Warwick and International Monetary Fund, mimeo.

Satchell, Stephen, and Allan Timmermann (1995) 'An Assessment of the Economic Value of Non-Linear Foreign Exchange Rates Forecasts,' *Journal of Forecasting*, 14, 477-97.

Schinasi, Garry J., and Swamy, P.A.V.B. (1989) 'The Out-of-Sample Forecasting Performance of Exchange Rate Models When Coefficients Are Allowed to Change,' *Journal of International Money and Finance*, 8, 375-390.

Sercu, Piet, Uppal, Raman, and Cynthia Van Hulle (1995) 'The Exchange Rate in the Presence of Transactions Costs: Implications for Tests of Purchasing Power Parity,' *Journal of Finance*, 50, 1309-19.

Sercu, Piet, and Xueping Wu, (2000) 'Uncovered Interest Arbitrage and Transaction Costs: Errors-in-Variables versus Hysteresis Effects,' University of Leuven and City University of Hong Kong, mimeo.

Summers, Robert, and Alan Heston (1991) 'The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950-1988,' *Quarterly Journal of Economics*, 106, 327-68.

Tay, Anthony S., and Kenneth F. Wallis (2000) 'Density Forecasting: A Survey,' *Journal of Forecasting*, 19, 235-54.

Taylor, Alan M. (2001) 'Potential Pitfalls for the Purchasing Power Parity Puzzle? Sampling and Specification Biases in Mean Reversion Tests of the Law of One Price,' *Econometrica*, 69, 473-98.

Taylor, Alan M. (2002) 'A Century of Purchasing-Power Parity,' *Review of Economics and Statistics*, 84, 139-50.

Taylor, Alan M., and Mark P. Taylor (2004) 'The Purchasing Power Parity Debate,' *Journal of Economic Perspectives*, 18, 135-58.

Taylor, Mark P., Peel, David A., and Lucio Sarno (2001) 'Nonlinear Mean-Reversion in Real Exchange Rates: Toward a Solution to the Purchasing Power Parity Puzzles,' *International Economic Review*, 42, 1015-42.

Taylor, Mark P., and Lucio Sarno (1998) 'The Behavior of Real Exchange Rates During the Post-Bretton Woods Period,' *Journal of International Economics*, 46, 281-312.

Taylor, Mark P., and Patrick C. McMahon (1988) 'Long-Run Purchasing Power Parity in the 1920s,' *European Economic Review*, 32, 179-97.

West, Kenneth D., Edison, Hali J., and Dongchul Cho (1993) 'A Utility-Based Comparison of Some Models of Exchange Rate Volatility,' *Journal of International Economics*, 35, 23-45.

Lead footnote: This work was partly undertaken while Lucio Sarno was a Visiting Scholar at the International Monetary Fund and at the Central Bank of Norway. The paper is based on Sarno's *State-of-the-Art Lecture*, given at the 2004 Annual Conference of the Canadian Economics Association (CEA) Toronto, 4-6 June 2004. The research reported in this paper was partially funded by a research grant from the Economic and Social Research Council (ESRC Grant No. RES-000-22-0404). The author is grateful to an anonymous referee, Farooq Akram, David Laidler, Dagfinn Rime and Barbara Spencer, for comments or useful conversations at various stages of the writing up of this paper. The author alone is responsible for any errors that may remain. Correspondence: Prof. Lucio Sarno, Finance Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, UK. Tel: +44 2476 528219. Email: lucio.sarno@warwick.ac.uk

Notes

[1] Exceptions include Bansal and Dahlquist (2000), who document that the forward bias is largely confined to developed economies and to countries for which the US interest rate exceeds foreign interest rates; Bekaert and Hodrick (2001), who, paying particular attention to small-sample distortions of tests applied to UIP and expectations hypotheses tests, provide a 'partial rehabilitation' of UIP; and Flood and Rose (2002), who report that the failure of UIP is less severe during the 1990s and for countries which have faced currency crises over the sample period investigated.

[2] Equivalently, via the CIP condition, these findings indicate that the more domestic interest rates exceed foreign interest rates, the more the domestic currency tends on average to appreciate over the holding period, not to depreciate so as to offset on average the interest differential in favor of the home currency.

[3] It is also worth noting that the literature has investigated the predictability of UIP deviations (or foreign exchange excess returns) using the forward premium as a predictor variable in a linear model obtained from reparameterizing equation (2) as follows: $ER^1_{t+1} = \alpha + \beta^\tau \left(f^1_t - s_t\right) + \upsilon_{t+1}$, where the excess returns $ER^1_{t+1} \equiv \Delta s_{t+1} - \left(f^1_t - s_t\right) \equiv s_{t+1} - f^1_t$ and $\beta^\tau \equiv \beta - 1$. This regression was investigated, for example, by Bilson (1981), Fama (1984) and Backus, Gregory and Telmer (1993) and was shown to generate strong predictability of excess returns (deviations from UIP) on the basis of the lagged forward premium. Specifically, while β^τ should be zero under UIP, the evidence, consistent with a negative estimate of β in equation (2), is that β^τ is negative and massively statistically significant.

[4] Our use of the term 'premium' rather than 'discount' is again arbitrary and follows standard usage in the literature; risk premia can, however, be negative. Note that $\rho_t = i_{t,k} - i^*_{t,k} - \Delta_k s^e_{t+k}$.

[5] Nucci (2003) also finds that the term structures of forward premia of other currencies have incremental information

content in addition to the currency's own forward premia. This result is rationalized on the basis of the abundant evidence of the co-movements of excess returns from investing in different currencies.

[6] It should be clear from our discussion that we take a commodity-arbitrage view to rationalize the validity of long-run PPP, as popularized, for example, by Samuelson (1964). An older view of PPP is the Cassellian view, according to which the appropriate definition of the price level for implementing PPP is the general price level and whether or not the price level samples nontradable goods is irrelevant (Cassell, 1922). Cassel's idea is that PPP refers to the internal value of the currencies concerned, which can only be measured using the general price level. While the Cassellian view is still somewhat extant in the profession, our reading of the recent theoretical and empirical literature on PPP is, however, that most experts in international finance who believe in the validity of long-run PPP take the commodity-arbitrage view (e.g. Froot and Rogoff, 1995; Rogoff, 1996; Sarno and Taylor, 2002, and the references therein).

[7] In an illustrative Monte Carlo exercise, Sarno and Taylor (2002) show how, given the observed persistence properties of the real exchange rate, standard unit root tests may be unable to detect stationarity of the real exchange, indeed having less than a 50/50 chance to do so for samples spanning up to 100 years.

[8] The data are the same as in Lothian and Taylor (1996), updated from 1991 to 2000 using consumer price indices and nominal exchange rate data taken from the *International Financial Statistics* of the International Monetary Fund.

[9] In doing so, we select a few representative studies in this sub-section. For a full review of the relevant nonlinear literature, see Sarno and Taylor (2002) and Taylor and Taylor (2004).

[10] Note that, because of the nonlinearity, the half-lives of shocks to the real exchange rates vary both with the size of the shock and with the initial conditions.

[11] An interesting experiment in terms of gauging the extent to which market integration and the reduction of trade costs impacts on the degree of mean reversion in real exchange rates is provided by the advent of the euro in 1999. In a recent study, Koedijk, Tims and van Dijk (2004) provide empirical evidence that the introduction of the euro and, more generally, the process of economic integration in Europe has accelerated convergence to PPP, consistent with a transactions-costs goods-market arbitrage view of the mean reversion properties of the real exchange rate. Future research on these data is warranted to refine the estimates of Koedijk, Tims and van Dijk and test for nonlinearities.

[12] The main empirical result in West, Edison and Cho (1993) is that, of homoskedastic, GARCH, autoregressive and nonparametric models for the conditional variance of exchange rate movements, GARCH models tend to produce the highest utility, on average. A mean squared error criterion also favors GARCH, but not as sharply.

Figure 1: Spot and Forward Dollar-Sterling Rates

Figure 2: 200 years of prices and exchange rates

UK and US prices and the spot dollar-sterling rate (in logs)

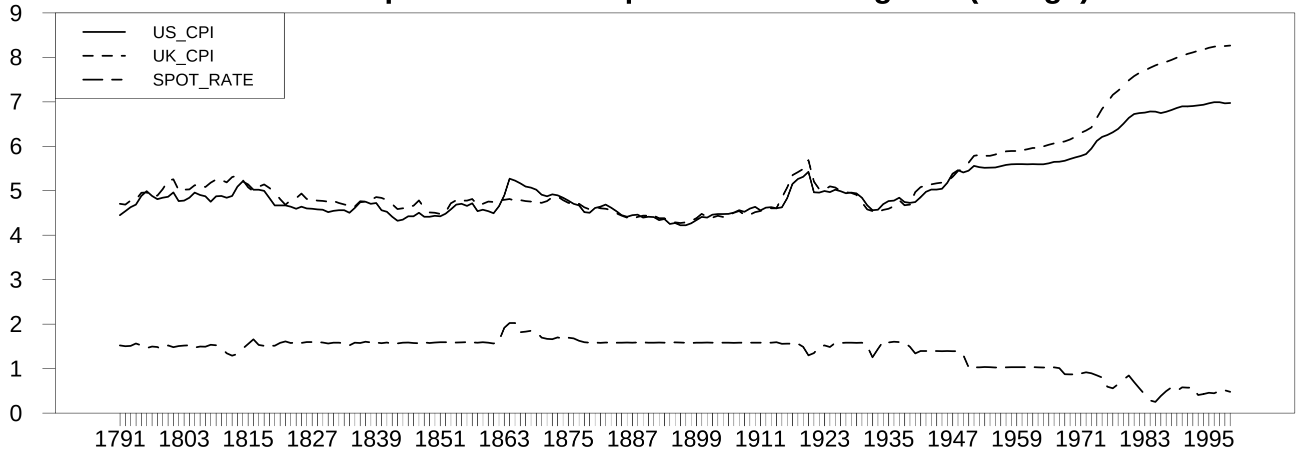

The real dollar-sterling exchange rate (in logs)