Foreign Exchange Implied Variance and the Forward Premium Puzzle

Efthymios Pavlidis[a], Mark B. Shackleton[b] and Nikolaos Voukelatos[c*]

[a] *Department of Economics, Lancaster University, Lancaster, UK*

[b] *Department of Accounting & Finance, Lancaster University, Lancaster, UK*

[c] *Kent Business School, University of Kent, Canterbury, UK*

Abstract

We explore the hypothesis that Jensen's Inequality is related to the magnitude of the commonly observed difference between forward rates and the subsequent realizations of spot exchange rates. Compared to the standard specification, it is shown that using the option-implied variance of the spot rate as an additional regressor in the unbiased forward specification results in slope coefficients that are closer to their theoretical value of unity. Furthermore, implied variance is found to have a higher explanatory power over future spot returns compared to that of the forward premium. Our empirical findings are consistent with the hypothesis that the time-varying risk-premium documented in previous studies contains a Jensen's term of the future spot variance.

JEL Classifications: C22; F31

Keywords: Forward premium; Implied variance; Jensen's inequality

1. Introduction

* Correspondence to: Nikolaos Voukelatos, Kent Business School, University of Kent, Canterbury CT2 7PE, UK. E-mail: n.voukelatos@kent.ac.uk

The forward premium puzzle refers to the widely observed rejection of the forward premium as a conditionally unbiased predictor of future spot exchange rate returns. When exchange rate returns are regressed on the lagged forward premium/discount, interest rate parity predicts a slope coefficient equal to one. However, a large body of related literature reports a coefficient less than the theoretical value of unity and, in many cases, significantly negative. For instance, McCallum (1994), using Yen, Mark and Pound rates against the Dollar for the period 1978-1990, estimates a slope coefficient of -4.[1] According to the *Uncovered Interest Rate Parity* (UIP) condition, this translates to an expectation of appreciating currency for the country with the higher nominal interest rate.

This paper examines the forward premium puzzle using data on the British Pound/U.S Dollar exchange rate, with particular emphasis on the role of option implied information in explaining this widely documented anomaly. Given that previous regressions of future spot rates (returns) on forward rates (forward premium) have resulted in estimated slopes that are significantly different from the value of unity that UIP predicts (and even negative in certain cases), the hypothesis is explored that an omitted future spot variance is at least partly driving the results.

Assuming lognormally distributed exchange rates, *Jensen's Inequality* results in a time-varying risk-premium in the forward markets by incorporating a term that refers to the future variance of the spot rate. The above correction has received relatively little attention in the existing literature, with a number of researchers arguing that its overall effect in accounting for deviations from UIP is not significant. The motivation for this study stems from the fact that previous studies have used *historical* measures of volatility or simulations to proxy for the spot rate's variance. However, since the variable in the extended regression specification refers to *future* volatility, it could be the case that the observed failure of *Jensen's Inequality Term* (JIT) in estimating a forward slope closer to unity might be due to the use of a poor proxy for future variance.

The subsequent analysis attempts to correct this potential limitation by estimating forward-looking volatility proxies, namely volatilities implied from currency option

[1] Other studies that report negative coefficients include Froot and Frankel (1989), Byers and Peel (1991), Backus, Gregory and Telmer (1993), and Mark, Wu and Hai (1997). Also, see Hodrick (1987) and Engel (1996) for comprehensive reviews of the forward premium puzzle literature.

prices, which have been shown to have significant forecasting power over future volatility in foreign exchange markets (Pong, Shackleton, Taylor and Xu, 2004) and to be directly related to the forward premium's persistence (Kellard and Sarantis, 2008). Our study is also related to Menkhoff, Sarno, Schmeling and Schrimpf (2012) who examine the excess returns of carry trades and report that global FX volatility risk, proxied by innovations in implied volatility, constitutes a driver of foreign exchange risk premia.

The variety of explanations that have been suggested to account for the forward premium anomaly indicates that researchers so far have failed to reach a consensus with respect to the source of this potentially negative correlation between the forward rate differential and exchange rate returns. In general, three types of explanations have been proposed. The first stream has examined the presence of a time-varying risk-premium which is assumed to be negatively correlated with expected spot rate returns. Although the models developed by Boyer and Adams (1988) and McCallum (1994) were able to produce a negative value of the forward premium's slope, Engel (1996) argues that these models are unlikely to be reconciled with existing models of risk-averse behaviour. The second stream suggests that the way investors form expectations about future levels of exchange rates results in a systematic forecast error in the forward premium due to e.g. peso problems (Lewis, 1994, and Evans, 1995), misinterpretation of interest rate shocks as transitory (Gourinchas and Tornell, 2004), or delayed overshooting (Bacchetta and Wincoop, 2007, 2009). Finally, a third stream of literature points to monetary policy interventions as a potential driver of the forward premium puzzle (McCallum, 1994).[2]

In order to examine the predictive power of the forward rate, some of the earlier studies estimated a simple model where the log of the future spot rate is regressed on the log of the forward rate through Ordinary Least Squares (OLS) minimization. Subsequent research, though, on the time-series properties of the above variables has demonstrated that this model is potentially mis-specified. More recent surveys examining the forward

[2] A number of more recent studies have documented that the predictive ability of the forward rate on future spot rates depends on the forecasting time-horizon. Chaboud and Wright (2005) and Bernoth, Hagen and Vries (2010) show that the forward rate's slope starts close to its theoretical value of unity at relatively short horizons, and slowly turns negative as maturity approaches the monthly level. At the other end of the spectrum, Chinn and Meredith (2004) demonstrate that the slope is significantly positive and closer to one over multi-year horizons.

premium anomaly have explored the typical regression model in which exchange rate returns are the dependent variable and the lagged forward premium/discount is the only explanatory variable. However, it has been noted that under the mild assumption of log-normal distribution for exchange rates, forward prices and price levels, two correction terms must be added to the regression specification. These terms are related to the variance of the spot rate and its covariance with the price level, and are commonly referred to as *Jensen's Inequality Terms* (JIT). While this correction is dictated by theory, its empirical effect has been frequently questioned. For instance, Bekaert and Hodrick (1993) report that including an historical variance correction term does not result in a slope coefficient that is consistent with theoretical predictions, while similar conclusions have been reached by Cumby (1988), Hodrick (1989), Baillie and Bollerslev (1990), and Backus, Gregory and Telmer (1993).

In contrast to previous studies, to proxy for the future variance of the spot rate in the JIT framework we use the *variance implied by the prices of options* written on exchange rates. More specifically, we examine options written on the British Pound/U.S Dollar exchange rate and we compute the rate's option-implied variance as the model-free estimate, following the methodology of Britten-Jones and Neuberger (2000). We perform rolling estimations of UIP and we consider the proportion of forward premium slopes that lie within two standard errors from the theoretical value of unity as a measure of the validity of UIP for our sample.

Overall, the results show that foreign exchange implied variance is to an extent driving future spot returns, with its explanatory power even higher than that of the forward premium when compared in a univariate setting. Furthermore, including an expectation of the spot rate's future variance in an extended specification provides forward slopes that are closer to one compared to the standard univariate model. When the standard specification is estimated through OLS, 30% of the resulting betas are statistically different from unity, whereas including the option-implied JIT results in the proportion of betas rejecting UIP falling to 23%. Re-estimating the parity specification using Least Absolute Deviations (LAD) instead of OLS further reduces the number of forward premium slopes that violate parity to 18%.

Our empirical results indicate that the spot rate's expected variance accounts for some part of the observed deviations from parity, within Jensen's Inequality framework. Consequently, we argue that previous empirical findings that document an insignificant impact of this additional term to UIP are more likely to be a result of using an inefficient proxy for future variance, rather than a fundamental quality of the term itself. Additionally, the spot rate's option-implied variance can be interpreted as a significant and trackable part of the time-varying risk-premium in foreign exchange markets that has been reported in previous studies. However, it has to be noted that the option-implied JIT cannot fully explain parity violations in our sample, since the extended specification still produces a significant number of betas (18%) that reject the null of unity, albeit to a lower extent compared to the standard specification.

The remaining of the paper is organized as follows: Section 2 gives an overview of the economic relationships leading to the *Uncovered Interest Parity* condition that is empirically tested. Section 3 describes the data used in the empirical analysis, while Section 4 presents the results. Finally, Section 5 concludes.

2. Economics of the Forward Premium Puzzle

Throughout this paper, S_t denotes the spot exchange rate at time t, while $F_t^{t+\tau}$ refers to the forward exchange rate at time t for delivery at time $t+\tau$. Corresponding logarithmic values are denoted by the lower case variables s_t and $f_t^{t+\tau}$, respectively. Both rates use the US Dollar as the numeraire currency. Furthermore, i_t refers to the risk-free interest rate applicable for US investors, while i_t^* denotes the foreign risk-free rate.

The *Covered Interest Rate Parity* condition (CIP) states that the difference between the forward rate and the spot rate at time t must be equal to the interest rate differential between the two countries. There is strong empirical evidence demonstrating that, ignoring transaction costs, CIP generally holds (see for instance Bahmani-Oskooee and Das, 1985, and Clinton, 1988).

$$f_t^{t+\tau} - s_t = i_t^* - i_t \tag{1}$$

The *Uncovered Interest Rate Parity* condition (UIP) then states that the expectation of spot rate returns must be equal to the interest rate differential. Taking into account (1), UIP can be expressed as follows,

$$E_t[s_{t+\tau} - s_t] = f_t^{t+\tau} - s_t = i_t^* - i_t \tag{2}$$

where $E_t[\cdot]$ is a (risk-neutral) expectation operator conditional on information available at time t. This study focuses on the prediction that the expected future spot rate $E_t[s_{t+\tau}]$ must be equal to the current forward rate $f_t^{t+\tau}$, or equivalently, that expected spot returns $E_t[\Delta s_{t+\tau}]$ must be equal to the current forward premium $f_t^{t+\tau} - s_t$.

In order to derive UIP, one must jointly assume rational risk-neutral agents, free capital mobility and the absence of taxes on capital transfers. From the same set of assumptions, it is also implied that risk-neutral expected returns from trading in the forward market must be zero

$$E_t[(F_t^{t+\tau} - S_t)/P_{t+\tau}] = 0 \tag{3}$$

where $P_{t+\tau}$ denotes the domestic dollar price level at time $t+\tau$. Assuming that all three variables in (3) are lognormally distributed and by using a Taylor series expansion to second order, equation (4) is derived

$$E_t[s_{t+\tau}] - f_t^{t+\tau} = -0.5\mathrm{var}_t(s_{t+\tau}) + \mathrm{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{4}$$

where $p_{t+\tau}$ is the logarithm of the price level $P_{t+\tau}$. The above two conditional second moment terms are usually referred to as *Jensen's Inequality Terms* (JIT).[3] Equation (4) can be rewritten, in terms of returns, as

$$E_t[s_{t+\tau}] - s_t = (f_t^{t+\tau} - s_t) - 0.5\mathrm{var}_t(s_{t+\tau}) + \mathrm{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{5}$$

which describes the models that have been examined in past surveys. In this paper, the emphasis is on the effect of $var_t(s_{t+\tau})$, i.e. the spot rate's variance at $t+\tau$ conditional at information available at t, in explaining future spot levels, not taking into account the covariance between spot rates and the price level. Although theory predicts that the latter variable will have some explanatory power in predicting future exchange rates, it has

[3] A detailed discussion on the derivation of equation (4) under a *stochastic discount factor* framework is provided in Engel (1999) and Azar (2008). The above *Jensen's Inequality Terms* are shown to be directly related to the foreign exchange risk-premium not necessarily in terms of a rigorous theoretical framework, but rather as a result of a mathematical paradox, i.e. Siegel's paradox, which is based on the convexity property of exchange rates as ratios and the concavity property of the logarithmic function (see also Beenstock, 1985).

been argued that its measurement is relatively problematic, reducing its actual explanatory power. More specifically, price levels, like many other economic variables, are reported at relatively low frequencies. The resulting smoothing and averaging complicates its inclusion in the model, especially considering the fact that the dependent variable as well as the remaining explanatory variables are estimated at a daily frequency.

3. Data

3.1 Data Sources

This study examines the exchange rate of the British Pound vis-à-vis the US Dollar. We focus on a sample period where reported deviations from parity are relatively more pronounced, with our sample running from January 1988 to June 2001, for a total of 3,330 trading days. Daily spot exchange rates were obtained from DataStream. Monthly forward rates at a daily frequency are proxied by exchange-traded futures rates and were also obtained from DataStream.

The original options dataset consisted of a total of 191,249 options written on the GBP/USD exchange rate, with the dataset containing, among other fields, option prices, strike prices, time-to-maturity, implied volatilities and trading volume. Prices of foreign exchange options are calculated as the mid-point of the best bid and the best ask quote at the end of the trading day, while option implied volatilities per contract are calculated using the Black and Scholes (1973) option pricing formula.

Similarly to previous studies, several filters were introduced. First, all options with prices lying close to zero or outside the theoretical bounds were removed from the sample. Second, options that expired within a trading week (five trading days) were removed. Finally, observations with less than five traded contracts were also dropped to avoid illiquidity concerns. The above filtering resulted in a reduced dataset, with the final sample consisting of 132,787 options.

The continuously compounded risk-free rate of interest is proxied by the LIBOR offered to US investors, whilst the “dividend yield” of the underlying asset, i.e. of the spot exchange rate, is proxied by the UK LIBOR. A significant part of the related

literature suggests that this is a reasonable proxy of the dividend yield for an investor buying a currency option which gives her the right to buy the British Pound using US Dollars. The main intuition behind this choice is the fact that, had she instead bought the underlying, she would have been able to receive a return equal to the UK risk-free rate by investing in UK government bonds. The US and UK LIBOR rates were obtained from DataStream.

3.2 Time-Series Properties

Many studies have examined the time-series properties of foreign exchange rates, typically finding s_t to follow a unit root process, making foreign exchange returns ($s_{t+\tau} - s_t$) an I(0) process. However, results on the order of integration of the forward rate $f_t^{t+\tau}$ and of the forward premium ($f_t^{t+\tau}$ - s_t) have been less than conclusive. For example, Mark, Wu and Hai (1997) support the stationarity of the forward premium in their empirical investigation of three main exchange rates, namely the Pound/Dollar, French Franc/Dollar and Yen/Dollar rates. On the other hand, Crowder (1994) contradicts these results. Examining monthly observations for the Pound, Mark and Canadian Dollar relative to the US Dollar from January 1974 to December 1991, he finds that the null hypothesis of non-stationarity of ($f_t^{t+\tau}$ - s_t) cannot be rejected.

As can be seen from Figure 1, logarithmic spot and monthly forward rates in our sample exhibit very high, positive, slow-decaying autocorrelations at roughly the first 100 lags. The correlograms of these two variables are very similar, although autocorrelations in the forward rate are at a slightly lower level than those of the spot rate. By subtracting s_t, exchange rate returns have significantly lower serial correlations after the first few lags. However, autocorrelations in the forward premium ($f_t^{t+\tau}$ - s_t) are much higher than those of ($s_{t+\tau}$ - s_t), suggesting a more pronounced deviation from stationarity.

The order of integration for the variables of interest is directly examined by performing the Augmented Dickey-Fuller test (ADF; Said and Dickey, 1984), the Kwiatkowski, Phillips, Schmidt and Shin test (KPSS; Kwiatkowski *et al*, 1992), and the Geweke and Porter-Hudak test (GPH; Geweke and Porter-Hudak, 1983), with the results

presented in Table 1. The ADF tests the null that the series examined is a unit root I(1) against the alternative of the series being stationary, and the number of lags is selected using the Schwartz Information Criterion. As can be seen from the Table, the results support the presence of a unit root in both the spot exchange rate (t-stat = -2.45) and the forward rate (t-stat = -2.43). The results of the KPSS test, however, highlight an important difference between spot returns and the forward premium. More specifically, the KPSS tests the null that the series examined is stationary I(0) against the alternative of a unit root, with the number of lags selected by Newey-West bandwidth using Bartlett kernel spectral levels. Spot rate returns are found to be stationary at the 5% level (t-stat = 0.05), which is to be expected given that they are the first difference of a unit root process, while the null of stationarity is rejected for the forward premium (t-stat = 3.67). Finally, the GPH test estimates that the order of integration is relatively close to unity for s_t (d = 0.81) and $f_t^{t+\tau}$ (d = 0.76), and are thus not inconsistent with the ADF results. The GPH results are also consistent with KPSS with respect to the series' first differences, since spot rate returns are shown to be effectively stationary (d = 0.05), while the forward premium might be characterized by a fractionally integrated process (d = 0.62).

Before proceeding with examining the forward unbiasedness hypothesis, attention should be drawn to the issue of potential data contamination in the time-series of spot and forward rates. This concern has been extensively discussed by Maynard and Phillips (2001), who report that forward premia series that are obtained from five alternative sources exhibit significant differences, tending to '… obfuscate the true time series properties of the forward premium, creating a clear (finite sample) bias in favour of stationarity'. Furthermore, they observe that the forward premium in their sample is characterized by large one-day fluctuations that are not present in the interest rate differential. This pattern is also found in the data used in this study as evidenced by the time-series of the nominal interest rate differential between the US and the UK, and of the forward premium, plotted in Figure 2. This implies a significant deviation from CIP, which predicts that the forward premium should be identical to the nominal interest rate differential.

3.3 Implied Variance

The main variable of interest $var_t(s_{t+\tau})$ refers to the future period $t+\tau$ and is, therefore, not observable at time t. One methodology that has been used in the related literature involves fitting an historical model to past data and estimating future volatility through this model's parameters. However, since the JIT requires the variance at $t+\tau$, a forward-looking measure, such as implied variance, might be a more appropriate proxy for future variance.

In order to obtain a forward-looking measure of $var_t(s_{t+\tau})$, implied variances are extracted from a large set of options written on foreign exchange using the model-free methodology developed by Britten-Jones and Neuberger (2000). So-called 'model-free' implied variances have been widely used in the related literature due to their appealing properties, especially the fact that volatility estimates do not depend on the validity of a particular option pricing model. In addition to being valid under a wide range of stochastic processes, model-free volatilities have also been shown to more efficiently forecast future volatility levels compared, for instance, to Black and Scholes implied volatilities (Jiang and Tian, 2005). The detailed methodology for extracting model-free estimates of the spot rate's future variance is discussed in the Appendix.

Figure 3 plots the annualized 1-month implied variance of the GBP/USD exchange rate across the sample period January 1988 to June 2001. The standardized monthly implied variance is obtained by estimating the model-free implied variance for the two nearest expiration dates, and then linearly interpolating to the standardized one month horizon. With respect to its time-series properties, the last row of Table 1 suggests that implied variance is best described by a fractionally integrated process I(d), similarly to the forward premium discussed in the previous Section. More specifically, the ADF test strongly rejects the null of a unit root in the time-series of implied variance (t-stat = -49.39), the KPSS test rejects the null of stationarity (t-stat = 1.57), while the GPH test estimates the order of integration of implied variance to be roughly $d = 0.44$.

Finally, it has to be noted that option implied variance reflects the market's expectations of future variance under the *risk-neutral* probability measure. Risk-neutral

variances will almost certainly differ from *real-world* expected variances, since they reflect risk preferences as well as the market's subjective expectations. This volatility risk-premium that is incorporated in risk-neutral implied variance is associated with an upwards biased proxy of $var_t(s_{t+\tau})$ and, depending on the risk-premium's time-variation, this is likely to be reflected in the residual errors when estimating UIP regressions. However, even if upwards biased, risk-neutral implied variances represent more timely proxies for future variance compared to historical estimates, due to their inherently forward-looking nature.

4. Empirical Results

Previous studies of the forward unbiasedness hypothesis have mainly focused on two types of specifications, based on exchange rate levels and exchange rate returns.[4] The first type is specified by a regression of future spot rate levels on current forward rates, described in (6), while the second specification involves regressing exchange rate returns on the forward premium and it is given in equation (7).

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \varepsilon_t \tag{6}$$

$$s_{t+\tau} - s_t = \alpha + \beta(f_t^{t+\tau} - s_t) + \varepsilon_t \tag{7}$$

Similarly to the more recent stream of the literature on the forward premium puzzle, and addressing the time-series complications that characterize level UIP regressions, we test the standard specification where the forward premium is the only explanatory variable of future spot rate returns. Before examining the effect of expected variance, the standard specification in (7) is estimated for the period 1988-2001 through OLS regressions and statistical inference is based on Newey-West HAC standard errors, with the results presented in Panel A of Table 2. The estimated slope for the entire sample is found to be less than the value of unity that is predicted by UIP. At a level of 0.1312, the forward premium's beta is positive but, with an OLS standard error of 0.32, statistically indistinguishable from zero. A beta that is statistically different from one is

[4] A third type refers to Error-Correction Models (ECM) which include lagged differences between the spot and the forward rates as additional explanatory variables of exchange rate returns.

consistent with findings of previous studies that have examined a wide set of exchange rates at various sampling frequencies. Furthermore, the low Adj.R^2 (0.03%) suggests that the forward premium can explain a very small part of subsequent rate changes when we examine a point estimate across the entire sample.

In addition to regressing equation (7) for the entire sample period, rolling 2-year estimations are performed in order to examine the evolution of observed betas across time.[5] The methodology of evaluating UIP through rolling regressions, where β follows a MA process, is motivated by the findings of Baillie and Bollerslev (2000), among other studies, who report a significant time-variation of the slope coefficient in forward premium regressions, indicating that the rejection of the forward unbiasedness hypothesis is probably dependent on the period examined rather than a universal characteristic of foreign exchange markets.

As can be seem from Figure 4, rolling estimates of the slope coefficient indeed exhibit significant variability throughout the period examined. Starting from a level of around -2, beta follows an upward trend until 1993, when it experiences a sharp decline. Despite the significant volatility of the estimated slope, which ranges from a minimum of -2.43 to a maximum of 3.43, beta remains mostly positive after 1990. More importantly, roughly 30% of the coefficients are statistically different from the theoretical value of unity at the 5% confidence level, highlighting a significant deviation from the predictions of the UIP.

As has been discussed in Section 2, the specification in (7) fails to incorporate the JIT correction for the expected variance of the future spot rate, considering the risk-premium to be equal only to the difference between $f_t^{t+\tau}$ and $s_{t+\tau}$. Although previous studies have shown that the contribution of this term in explaining violations of UIP is not significant, estimating $var_t(s_{t+\tau})$, conditional on information available at time t, is not straightforward. Therefore, its documented inability to account for deviations of the forward rate's slope from its theoretical value of unity might be attributed, at least partly, to measurement error rather than to a fundamental quality.

[5] Further empirical analysis has been conducted under a wider set of rolling windows as a robustness check. Our findings are consistent across all chosen windows and are, therefore, not reported for brevity.

We address this concern by using option-implied variance as an alternative proxy for the exchange rate's future variance. Pong *et al* (2004) document that implied currency volatilities provide efficient forecasts of future volatility, particularly at the one-month horizon. In addition, we find that foreign exchange implied variance also has explanatory power over future spot returns. More specifically, we examine the relationship between future spot returns ($s_{t+\tau}$ - s_t) and implied variance $\text{var}_t(s_{t+\tau})$ by estimating OLS regressions of (8), both across the full sample as well for rolling 2-year windows, with the results presented in Panel C of Table 2.

$$s_{t+\tau} - s_t = \alpha + \gamma \text{var}_t(s_{t+\tau}) + \varepsilon_t \tag{8}$$

Our results suggest that implied variance is negatively related to future exchange rate returns (γ = -0.0386). Moreover, when $\text{var}_t(s_{t+\tau})$ is assumed to be the only factor driving spot returns in (8), the Adj.R^2 is found to be higher than the one obtained from the univariate UIP specification in (7), increasing from 0.03% to 0.06% and suggesting that, in the short term, implied variance has a higher explanatory power over spot returns compared to the forward premium. The rolling 2-year estimations produce similar results, with the mean Adj.R^2 increasing from 0.91% to 1.07% as we move from the standard UIP specification in (7), where the forward premium is the only dependent variable, to estimating (8) where spot returns are assumed to be driven solely by implied variance.

Given the previously reported explanatory power of foreign exchange implied variance over future spot returns, which is found to be higher than that of the forward premium in a univariate setting, we re-examine the forward unbiasedness hypothesis by testing the extended version in (9).

$$s_{t+\tau} - s_t = \alpha + \beta(f_t^{t+\tau} - s_t) + \gamma \text{var}_t(s_{t+\tau}) + \varepsilon_t \tag{9}$$

When running the OLS regressions, the three parameter vector [α, β, γ] is simultaneously estimated, instead of restricting γ to its theoretical value of -0.5, in order to allow for a more flexible framework. As can be seen from Panel A of Table 2, including an option implied proxy for the future variance of the spot rate fails to improve the predictive power of the forward rate when we run a single OLS regression. More specifically, when the extended specification in (9) is estimated for the entire sample period, the resulting forward beta is in fact slightly lower (0.1037) than the one estimated

in (7). On the other hand, the intercept term and the respective t-statistic decrease, while the Adj.R^2 increases to 0.10%. Although the estimate of the variance's slope is negative, γ is statistically indistinguishable from zero. Overall, the point estimates of the extended specification represent only marginal changes compared to those of the standard specification, and therefore they cannot be considered as being more consistent with UIP.

However, our results are substantially different when we move from point to rolling estimates. As can be seen from Figure 4 and Panel A of Table 2, introducing the JIT in rolling 2-year estimations results in the proportion of betas that fail to reject UIP increasing to 77%. This increase of more than 7% with respect to the standard specification is attributed to slope estimates moving closer to their theoretical value of unity, given that standard errors remain at the same level (mean = 0.46). Finally, the average Adj.R^2 of the rolling regressions increases from 0.91% to 1.96% when the JIT is included.

The methodology of estimating (7) through OLS regressions has been widely adopted in the forward bias literature. However, the resulting slopes and their statistical inference are not free of certain limitations. More specifically, least squares estimation fails to produce unbiased estimates of the forward premium's beta in the presence of outliers or if the explanatory variable is observed with error, with both cases constituting valid concerns for our sample. Therefore, in addition to OLS, we estimate the standard specification in (7) as well as the JIT extension in (9) through Least Absolute Deviations (LAD), which is more robust to outliers in the regressors and in the errors, and statistical inference is based on the kernel methodology proposed by Powell (1989).

Panel B of Table 2 reports the results from the LAD estimation of the forward unbiasedness hypothesis in (7) and its extension in (9). Regressing spot returns only on the current forward premium for the entire sample produces a negative slope coefficient (-0.1375) which is, unsurprisingly, statistically different from the theoretical value of one and indistinguishable from zero. Considering the point estimate for the entire sample period, the inclusion of a proxy for $var_t(s_{t+\tau})$ as a correction term does not lead to a β coefficient closer to unity, similarly to the OLS results.

However, when we take into account beta's time variation in rolling regressions, LAD results appear to be more consistent with UIP compared to the previously obtained OLS results. With respect to the standard specification, 75.67% of LAD betas cannot reject the null of unity, compared to 69.62% when OLS is used. This increase in the proportion of betas that do not reject UIP is of similar magnitude to that obtained when we move from the standard specification to its extension using OLS. Furthermore, estimating the extended JIT specification in (9) through LAD provides coefficients that are even more consistent with UIP, since 82.21% of the resulting forward premium betas are statistically indistinguishable from one.

5. Conclusion

This paper has examined the forward premium anomaly, i.e. the widely reported finding that when spot returns are regressed on the forward premium, the resulting slope coefficients deviate from one and, in many cases, fall below zero. A negative forward slope represents a significant violation of *Uncovered Interest Parity* and implies that, not only does the forward rate fail to predict the future level of exchange rates, but that it effectively predicts spot returns of the wrong sign.

Our empirical results confirm previous findings where the use of shorter samples in rolling regressions produces slope coefficients that are relatively disperse, change signs and, for some 2-year periods, do not reject the null of unity. Furthermore, the LAD estimator results in slopes that are more consistent with UIP compared to OLS estimates, since the former technique is able to address some common econometric concerns related to forward bias regressions, such as the presence of outliers in the explanatory variable and in the errors.

More importantly, we find that the spot rate's future variance has a higher explanatory power over spot returns than the forward premium, and that it goes some way into explaining the previously reported deviations from UIP. Adding the option-implied variance of the spot rate as an additional regressor in the forward unbiasedness specification results in a significantly higher proportion of slopes that are statistically

indistinguishable from their theoretical value of unity. Admittedly, introducing the typically omitted JIT cannot be considered a complete solution to the forward premium puzzle since a significant proportion of slopes still reject UIP in the extended specification. However, our results indicate that the existing treatment of JIT as insignificant in explaining deviations from parity can be mainly attributed to the use of a poor proxy rather than a fundamental quality of the term itself. Overall, our findings suggest that the spot rate's future variance can explain deviations from UIP in certain sub-samples and, in general, constitutes a significant and trackable part of the risk-premium observed in foreign exchange markets.

Appendix: Estimating Model-Free Implied Variance

Britten-Jones and Neuberger (2000) demonstrate that the future variance of asset returns can be calculated without the restriction of assuming a specific model for the return distribution. More specifically, they demonstrate that volatility implied by options prices can be estimated as the expected sum of squared returns under the risk-neutral measure, showing that in the time interval [0, T] it is completely specified by a set of out-of-the-money (OTM) options expiring at T.

(A1)

where V_T is the integrated squared volatility of the asset, A_t is the asset's spot price at time t, and $F_{0,T}$ is the forward price at time 0 for delivery at time T. Moreover, $p(K,T)$ and $c(K,T)$ are the prices of OTM put and call options, respectively, with strike K and expiring at T.

In order to derive equation (A1), the only assumption is that the stochastic processes of the underlying and its volatility are continuous. The Britten-Jones and Neuberger (2000) method requires option prices being quoted for a continuum of strikes. In reality, however, empirical estimation of squared expected returns can only be done using a finite set of discrete strikes. Carr and Wu (2004) and Jiang and Tian (2005) relax

the assumption of continuity and provide discrete versions of this model. This study follows the methodology adopted, among others, by Taylor, Yadav and Zhang (2010) of using a finite set of OTM options written on an asset to estimate the asset's integrated variance until the options' expiration. The discrete version of (A1) is then given as follows:

$$var_{MF} = \frac{2}{T}\sum_{i=1}^{M} \frac{\Delta K_i}{K_i^2} e^{rT} Q_i - \frac{1}{T}\left[\frac{F_{0,T}}{K_0} - 1\right]^2 \qquad \text{(A2)}$$

where var_{MF} is the model-free expectation of the future variance, M is the number of strike prices used, and Q_i is the option's market price at strike K_i. Since K_0 denotes the strike price used to select either call or put options in the formula, the option price Q_i refers to calls when $K_i \geq K_0$, and to puts otherwise. Finally, ΔK_i is centred on $\frac{K_{i+1} - K_{i-1}}{2}$.

As can be easily seen from the above equation, the value of $\frac{1}{T}\left[\frac{F_{0,T}}{K_0} - 1\right]^2$ depends on the strike K_0 that is chosen to separate call or put prices in the summation term. However, this methodology uses a small number of actual option prices to infer a risk-neutral density and, therefore, to create a significantly large number of artificial option-strike combinations. This allows for K_0 to be set equal to $F_{0,T}$, so that the final term in equation (A2) disappears. Consequently, in estimating var_{MF} only OTM calls and puts are used, with K_0 denoting the at-the-money (ATM) strike price.

Although the discrete version of model-free implied variance in (A2) can be empirically estimated, a significantly large set of options is needed to obtain an accurate measure of the underlying's future variance. Since options are actually quoted at a relatively limited number of strikes, the methodology described by Malz (1997) is employed to construct implied volatility curves using a small set of market-traded options.

Within this framework, the implied volatility curve is fitted as a function of option deltas, as opposed to a function of option strikes. Malz (1997) argues that this methodology ensures that volatilities of options that are further from the money (OTM and in-the-money (ITM) contracts) are grouped more closely together than those of near-

the-money options. Taylor, Yadav and Zhang (2010) also mention that '… extrapolating a function of delta provides sensible limits for the magnitude of implied volatility curves'. Following this line of thought, a quadratic function of implied volatility is fitted with respect to option delta. In addition to capturing the 'volatility smile', the quadratic specification has the advantage of requiring a minimum of only three options to be estimated.

$$IV_i = \alpha_0 + \alpha_1 \Delta_i + \alpha_2 \Delta_i^2 \qquad \text{(A3)}$$

Equation (A3) describes the quadratic function used to construct the implied volatility curve, where IV_i is the implied volatility of option i, and Δ_i is the option's delta. Moreover, IV_i is the simple Black and Scholes implied volatility of option i, while Δ_i is the sensitivity of option i to changes in the value of the underlying, measured by the option's delta as the first derivative of the Black and Scholes formula with respect to changes in the underlying's price A_0. It should be noted that, when calculating the model-free implied variance, option deltas are expressed as a function of σ^*, which is a constant measure of volatility used across all options (see also Bliss and Panigirtzoglou, 2002, and Taylor, Yadav and Zhang, 2010, for the use of σ^*).

Call deltas range from zero for deep OTM contracts with high strikes to e^{-rT} for deep ITM ones with low strikes. The respective put deltas range from $-e^{-rT}$ (deep ITM puts with high strikes) to zero (deep OTM puts with low strikes).

The parameter vector $\Phi = [\alpha_0, \alpha_1, \alpha_2]$ of the quadratic function is estimated by minimizing the weighted sum of squared differences between observed volatilities IV and fitted volatilities $\widehat{IV}(\Delta_j, \Phi)$ with respect to Φ, as given in (A4):

$$\min_{\Phi} \sum_{j=1}^{M} w_j \left(IV_j - \widehat{IV}(\Delta_j, \Phi) \right)^2 \qquad \text{(A4)}$$

where M is the number of observed strikes and w_j is the weight of option j's delta. The weight w_j of option j is equal to $\Delta_j(1-\Delta_j)$ and the minimization is subject to the constraint of fitted volatilities being strictly positive, $\widehat{IV}(\Delta_j, \Phi) > 0$. The above weighting scheme ensures that deviations of fitted volatilities from observed levels are more heavily penalized for the nearest-the-money options, i.e. calls with deltas close to 0.50, compared

to further-from-the-money contracts. Placing more weight on near-the-money options is compatible with the stylized fact that these options are more heavily traded, thus reducing the effect of possible outliers of illiquid ITM and OTM contracts.

After fitting the implied volatility curve, a large set of artificial option prices is created using the vector Φ. More specifically, 1,000 equally spaced deltas ranging from 0 to e^{-rT} are used to extract the corresponding strikes. Then, option prices (both calls and puts) are estimated using the Black and Scholes formula with the respective combinations of strike price and volatility. The OTM contracts are identified and used in estimating the asset's integrated variance in equation (A2), i.e. puts with strikes in the range $[0, F_{0,T})$ and calls with strikes in the range $[F_{0,T}, \infty)$. Finally, standardized 30-day estimates are calculated by linearly interpolating between the nearest and second-nearest variances.

References

Azar, Samih Antoine, 2008, Jensen's Inequality in Finance, *International Atlantic Economic Society* 14, 433-440

Bacchetta, Philippe, and Eric van Wincoop, 2009, Infrequent Portfolio Decisions: A Solution to the Forward Premium Puzzle, *American Economic Review* 100, 837-869

Bacchetta, Philippe, and Eric van Wincoop, 2007, Random Walk Expectations and the Forward Discount Puzzle, *American Economic Review, Papers and Proceedings* 97, 346-350

Backus, David, Allan Gregory, and Chris Telmer, 1993, Accounting for Forward Rates in Markets for Foreign Currency, *Journal of Finance* 48, 1887-1908

Bahmani-Oskooee, Mohsen, and Satya Das, 1985, Transaction Costs and the Interest Parity Theorem, *Journal of Political Economy* 93, 793-799

Baillie, Richard, and Tim Bollerslev, 1990, A Multivariate Generalized GARCH Approach to Modeling Risk Premia in Forward Foreign Exchange Markets, *Journal of International Money and Finance* 9, 309-324

Baillie, Richard, and Tim Bollerslev, 2000, The Forward premium Anomaly is not as Bad as you Think, *Journal of International Money and Finance* 19, 471-488

Beenstock, Michael, 1985, Forward Exchange Rates and "Siegel's Paradox", *Oxford Economic Papers* 37, 298-303

Bekaert, Geert, and Robert Hodrick, 1993, On Biases in the Measurement of Foreign Exchange Risk Premiums, *Journal of International Money and Finance* 12, 115-138

Bernoth, Kerstin, Juergen von Hagen, and Casper de Vries, 2010, The Forward Premium Puzzle: New Evidence from Futures Contracts, *DNB Working Paper No. 246*

Black, Fisher, and Myron Scholes, 1973, The pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-654

Bliss, Robert, and Nikolaos Panigirtzoglou, 2002, Testing the Stability of Implied Probability Density Functions, *Journal of Banking and Finance* 26, 381-422

Boyer, Russell, and Charles Adams, 1988, Forward Premia and Risk Premia in a Simple Model of Exchange Rate Determination, *Journal of Money, Credit and Banking* 20, 633-644

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance* 55, 839-866

Byers, David, and David Peel, 1991, Some Evidence on the Efficiency of the Sterling-Dollar and Sterling-Franc Forward Exchange Rates in the Interwar Period, *Economics Letters* 35, 317-322

Carr, Peter, and Liuren Wu, 2004, Time-Changes Levy Processes and Option Pricing, *Journal of Financial Economics* 71, 113-141

Chaboud, Alain, and Jonathan Wright, 2005, Uncovered Interest Parity, It Works, but not for Long, *Journal of International Economics* 66, 349-362

Chinn, Menzie, and Guy Meredith, 2004, Monetary Policy and Long Horizon Uncovered Interest Parity, *IMF Staff Papers* 51, 409-430

Clinton, Kevin, 1988, Transactions Costs and Covered Interest Arbitrage: Theory and Evidence, *Journal of Political Economy* 96, 358-370

Crowder, William, 1994, Foreign Exchange Market Efficiency and Common Stochastic Trends, *Journal of International Money and Finance* 13, 551-564

Cumby, Robert, 1988, Is it Risk? Explaining Deviations from Uncovered Interest Parity, *Journal of Monetary Economics* 22, 279-299

Engel, Charles, 1999, On the Foreign Exchange Risk Premium in Sticky-Price General Equilibrium Models, *International Tax and Public Finance* 6, 491-505

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

Evans, Martin, 1995, Peso Problems: Their Theoretical and Empirical Implications, *Handbook of Statistics: Statistical Methods in Finance*, North-Holland, Amsterdam

Froot, Kenneth, and Jeffrey Frankel, 1989, Forward Discount Bias: Is it an Exchange Risk Premium?, *Quarterly Journal of Economics* 104, 139-161

Geweke, John, and Susan Porter-Hudak, 1983, The Estimation and Application of Long Memory Time Series Models, *Journal of Time Series Analysis* 4, 221-238

Gourinchas, Pierre-Olivier, and Aaron Tornell, 2004, Exchange Rate Puzzles and Distorted Beliefs, *Journal of International Economics* 63, 303-333

Hodrick, Robert, 1987, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, Harwood

Hodrick, Robert, 1989, U.S. International Capital Flows: Perspectives from Rational Maximizing Models, *Carnegie-Rochester Conference Series on Public Policy* 30, 231-288

Jiang, George, and Yisong Tian, 2005, Model-Free Implied Volatility and its Information Content, *The Review of Financial Studies* 18, 1305-1342

Kellard, Neil and Nicholas Sarantis, 2008, Can Exchange Rate Volatility Explain Persistence in the Forward Premium?, *Journal of Empirical Finance* 17, 714-728

Koenker, Roger and Jose Machado, 1999, Goodness of Fit and Related Inference Processes for Quantile Regression, *Journal of the American Statistical Association* 94, 1296-1310

Kwiatkowski, Denis, Peter Phillips, Peter Schmidt and Yongcheol Shin, 1992, Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root: How Sure are we that the Economic Time Series Have a Unit Root?, *Journal of Econometrics* 54, 159-178

Lewis, Karen, 1994, Puzzles in International Financial Markets, *The Handbook of International Economics*, North-Holland, Amsterdam

Malz, Allan, 1997, Option-Implied Probability Distributions and Currency Excess Returns, *Federal Reserve Bank of New York Staff Report* 32, November 1997

Mark, Nelson, Yangru Wu, and Weike Hai, 1997, Understanding Spot and Forward Exchange Rate Regressions, *Journal of Applied Econometrics* 12, 715-734

Maynard, Alex, and Peter Phillips, 2001, Rethinking an Old empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly, *Journal of Applied Econometrics* 16, 671-708

McCallum, Bennett, 1994, A Reconsideration of the Uncovered Interest Parity Relationship, *Journal of Monetary Economics* 33, 105-132

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling and Andreas Schrimpf, 2012, Carry Trades and Global Foreign Exchange Volatility, *Journal of Finance* 67, 681-718

Pong, Shiuyan, Mark B. Shackleton, Stephen J. Taylor, and Xinzhong Xu, 2004, Forecasting Currency Volatility: A Comparison of Implied Volatilities and AR(FI)MA Models, *Journal of Banking and Finance* 28, 2541-2563

Powell, James, 1989, Estimation of Monotonic Regression Models Under Quantile Restrictions, in

Nonparametric and Semiparametric Methods in Econometrics, Cambridge University Press
Said, Said and David Dickey, 1984, Testing for Unit Roots in Autoregressive-Moving Average Models of Unknown Order, *Biometrica* 71, 599-607
Taylor, J. Stephen, Pradeep Yadav, and Yuanyuan Zhang, 2010, The Information Content of Implied Volatilities and Model-Free Volatility Expectations: Evidence from Options Written on Individual Stocks, *Journal of Banking and Finance* 34, 871-881

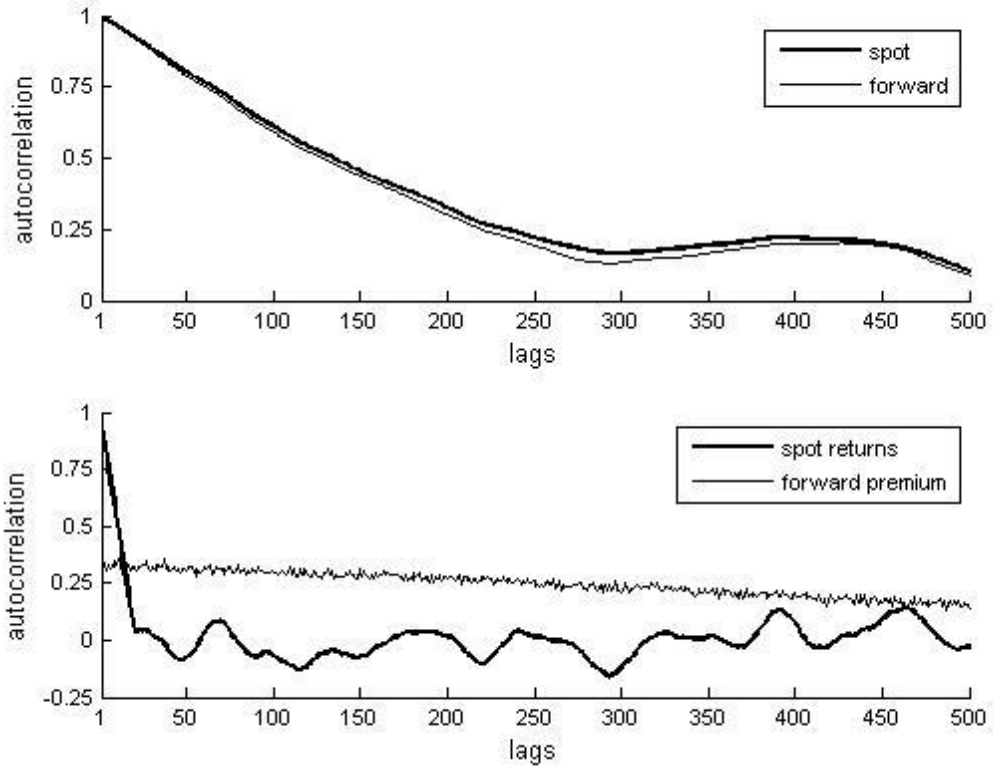

Fig. 1. Correlograms
The upper graph plots the autocorrelation function of the spot and forward rates, while the lower graph plots the autocorrelation function of spot returns and the forward premium. The sample runs from January 1988 to June 2001.

Table 1
Unit-root and Stationarity Test Results

	ADF (t-stat)	KPSS (t-stat)	GPH (d)
s_t	-2.45*		0.81
$f_t^{t+\tau}$	-2.43*		0.76
$s_{t+\tau} - s_t$		0.05*	0.05
$f_t^{t+\tau} - s_t$		3.67	0.62
$var_t(s_{t+\tau})$	-49.39	1.57	0.44

Note: This Table tabulates the results of tests for the order of integration of the main variables, namely the spot exchange rate, the forward rate, spot rate returns, the forward premium and the future spot variance. Panel A describes the results of the Augmented Dickey-Fuller (ADF) test, the null of which is that the series examined is a unit root. Panel B describes the results of the Kwiatkowski, Phillips, Schmidt and Shin (KPSS) test, the null of which is that the series examined is stationary. Panel C reports the variables' order of integration (*d*) based on the Geweke and Porter-Hudak (GPH) test. The number of lags in the ADF test is selected using the Schwartz Information Criterion. The number of lags in the KPSS test is selected automatically by Newey-West bandwidth using Bartlett kernel spectral levels. The 5% critical values are -2.87 and 0.46 for the ADF and KPSS tests, respectively. Finally, * denotes statistical significance at the 5% level.

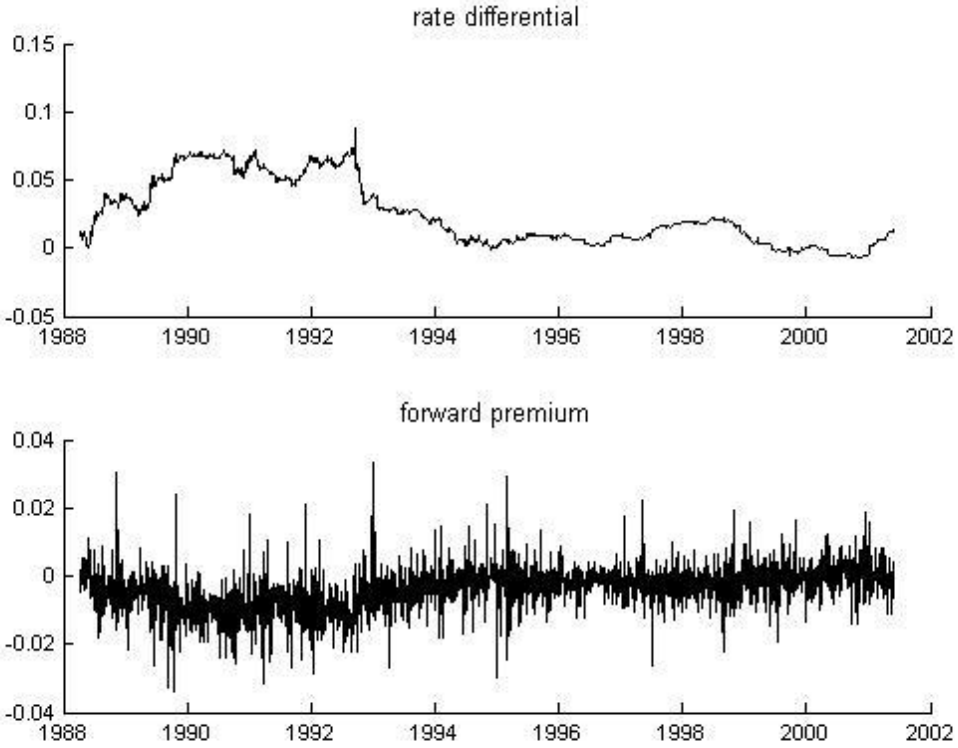

Fig. 2. Interest Rate Differential vs Forward Premium
The upper graph plots the difference between the GBP and USD interest rates, from January 1988 to June 2001. The lower graph plots the forward premium of the GBP/USD exchange rate, from January 1988 to June 2001.

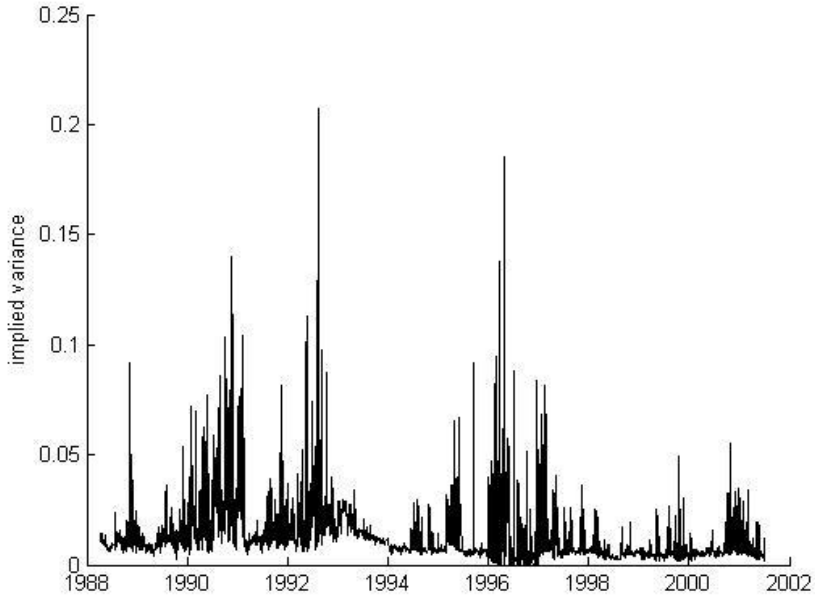

Fig. 3. Implied Variance
This Figure plots the standardized 30-day Implied Variance of the GBP/USD exchange rate, from January 1988 to June 2001.

Table 2
UIP Regression Results

Panel A: OLS estimation of UIP

	point estimate α	point estimate β	point estimate γ	point estimate Adj.R^2	rolling betas non-rejecting H_0: $\beta = 1$	mean rolling Adj.R^2
Standard Specification	-0.0015 (-1.12)	0.1312 (-2.76)		0.03%	69.62%	0.91%
JIT Extended Specification	-0.0010 (-0.75)	0.1037 (-2.96)	-0.0366 (12.77)	0.10%	76.98%	1.96%

Panel B: LAD estimation of UIP

	point estimate α	point estimate β	point estimate γ	point estimate goodness-of-fit	rolling betas non-rejecting H_0: $\beta = 1$	mean rolling goodness-of-fit
Standard Specification	-0.0008 (-1.41)	-0.1375 (-6.74)		0.03%	75.67%	0.58%
JIT Extended Specification	-0.0093 (-1.16)	-0.1321 (-6.62)	0.0098 (20.87)	0.05%	82.21%	1.33%

Panel C: OLS regression of spot returns on implied variance

point estimate α	point estimate γ	point estimate Adj.R^2	mean rolling Adj.R^2
-0.0012 (-0.88)	-0.0386 (11.72)	0.06%	1.07%

Note: Panels A and B tabulate the results of testing the forward rate unbiasedness hypothesis through OLS and LAD regressions, respectively. Panel C reports the results of regressing spot rate returns against the spot rate's implied variance. The t-statistics (in brackets) refer to statistical difference from 0, 1 and -0.5, for α, β and γ, respectively. OLS t-statistics are based on Newey-West HAC standard errors. LAD t-statistics are estimated through the kernel-based methodology proposed by Powell (1989). Goodness-of-fit for the LAD estimator is evaluated through the Koenker and Machado (1999) metric.

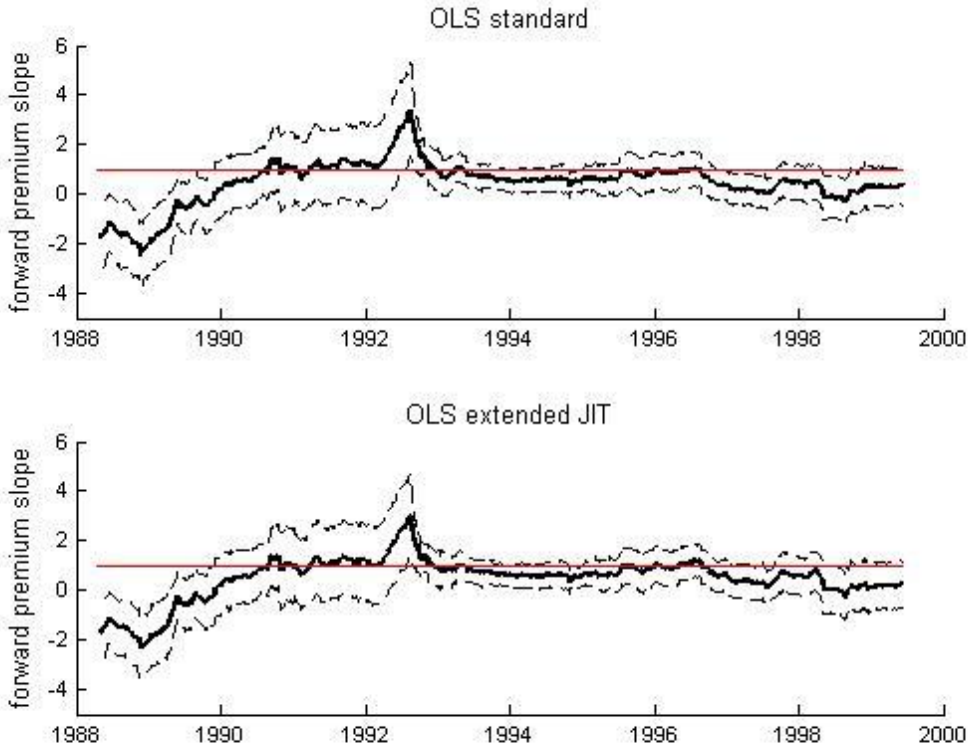

Fig. 4. OLS Forward Premium Slopes.
This figure plots the slopes of rolling 2-year OLS unbiasedness regressions for the GBP/USD exchange rate. The sample period runs from January 1988 to June 2001. The upper graph plots the slopes of the standard specification, while the lower graph plots the slopes of the extended JIT specification. The dashed lines represent the conventional two OLS standard error confidence bands. The red line shows the theoretical value of unity.

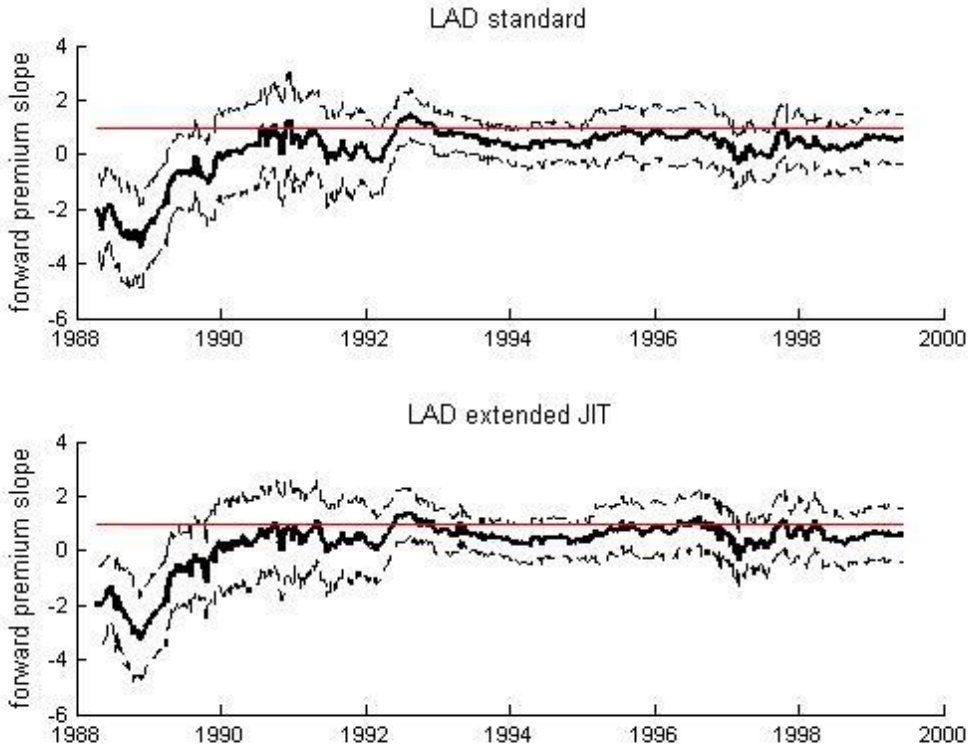

Fig. 5. LAD Forward Premium Slopes.
This figure plots the slopes of rolling 2-year LAD unbiasedness regressions for the GBP/USD exchange rate. The sample period runs from January 1988 to June 2001. The upper graph plots the slopes of the standard specification, while the lower graph plots the slopes of the extended JIT specification. The dashed lines represent the conventional two LAD standard error confidence bands. The red line shows the theoretical value of unity.