ELSEVIER Journal of International Economics 66 (2005) 349–362

Journal of INTERNATIONAL ECONOMICS

www.elsevier.com/locate/econbase

Uncovered interest parity: it works, but not for long

Alain P. Chaboud, Jonathan H. Wright*

International Finance Division, Board of Governors of the Federal Reserve System, 20th and C Streets NW, Washington, DC 20551, USA

Received 23 June 2003; received in revised form 23 June 2004; accepted 7 July 2004

Abstract

If an investor borrows in a low interest currency and invests in a high interest currency, the interest differential accrues in a lumpy manner, formally just like the dividend payments on a stock. The investor will receive the interest differential discretely at the point when a position is rolled over from one day to the next. A position that is not held open overnight receives no interest differential because intradaily interest rates are zero. Using a large data set of intradaily exchange rate data, we run uncovered interest parity (UIP) regressions over different short time intervals taking careful account of the settlement rules in the spot foreign exchange market. We find results that are supportive of the uncovered interest parity hypothesis over very short windows of data that span the time of the discrete interest payment. However, adding even a few hours to the span of the window destroys the positive uncovered interest parity results.
Published by Elsevier B.V.

Keywords: Uncovered interest parity; Settlement; Foreign exchange; High frequency data

JEL classification: C22; F31; F39

1. Introduction

The empirical failure of the simple uncovered interest parity (UIP) relation has been a puzzle to economists working in international finance ever since the work of Hansen and

* Corresponding author. Tel.: +1 202 452 3605; fax: +1 202 736 5638.
E-mail address: jonathan.h.wright@frb.gov (J.H. Wright).

0022-1996/$ - see front matter. Published by Elsevier B.V.
doi:10.1016/j.jinteco.2004.07.004

Hodrick (1980), Bilson (1981) and Baillie et al. (1983). The UIP relation postulates that the interest differential between two countries should equal the expected exchange rate change. As such, a regression of exchange rate returns on the interest differential should give an intercept of zero and a slope coefficient of unity. This hypothesis has however been consistently and decisively rejected by the data. Most often, the estimated slope coefficient is negative, meaning that the currency with the *higher* interest rate tends to *appreciate*. A carry trade (in which the investor borrows in the currency with the low interest rate and invests in the currency with a high interest rate) is profitable on average. Many comprehensive surveys exist (e.g. Froot and Thaler, 1990; Lewis, 1995; Engel, 1996) discussing reasons for the empirical failure of UIP. We do not attempt to review these, other than to note that they include the existence of a risk-premium (defined as the ex ante expected profit on the carry trade) that is time-varying and correlated with the interest differential.

While the UIP regression is usually run over horizons from a month to a year, Meredith and Chinn (1998) and Fujii and Chinn (2001) ran the UIP regression over longer horizons. They regressed exchange rate returns from t to $t+m$ on the difference in yields on m-period government bonds at time t, where the horizon, m, is as long as 10 years. They found that as the horizon m increases, the rejection of the UIP hypothesis becomes less decisive.

In this paper, we are going in exactly the opposite direction, examining UIP over extremely short horizons. We will exploit the fact that interest is only paid on overnight positions—no interest is paid on intraday positions. Lyons and Rose (1995) is the only extant paper that has exploited this fact in looking at the relationship between interest differentials and exchange rates at high frequency, to the best of our knowledge. They considered pairs of currencies in the now-defunct European Monetary System (EMS), and found that currencies which were under attack but in fact stayed within the band[1] actually appreciated intraday. Lyons and Rose argue that this intraday appreciation is compensation for the risk of devaluation that might have occurred, but did not. Investors can be compensated for the risk of devaluation only by intraday appreciation, not by interest differentials, as there are effectively no interest differentials intraday.

We focus instead on the flip-side of the argument of Lyons and Rose (1995). Instead of looking at high frequency exchange rate movements over the intraday period when no interest is paid, we instead consider the overnight period when interest does accrue. Interest is paid on positions that are open at a particular point in time −17:00 New York time. Absent all transactions costs, we would expect to see a jump in the exchange rate to offset the interest differential at this instant. Otherwise, the investor could gain the interest differential while being exposed to exchange rate risk for an arbitrarily short period of time.

Transactions costs are of course important, but by appealing to a model used in the literature on stock price behavior as stocks go ex dividend, we argue that we might nevertheless see such a jump in the exchange rate even if the transactions costs are large

[1] Concretely, Lyons and Rose (1995) considered the French franc-mark and lira-mark bilateral exchange rates, and defined the franc or lira as being under attack on those days on which the interest differential was in the top decile.

relative to the interest differential. An intraday UIP regression over a short period that spans 17:00 New York time might therefore yield results that are more favorable to the UIP hypothesis, as the full overnight interest differential that accrues in such a window is offset by a jump in the exchange rate. We test whether or not this is the case, and find remarkably consistent affirmative results, using 15 years of high frequency exchange rate data on dollar bilateral exchange rates relative to the yen, mark/euro, Swiss franc, and pound.

The plan for the remainder of this paper is as follows. Section 2 lays out the implications of the discrete timing of interest payments for high frequency UIP regressions. Section 3 contains the empirical work. Section 4 concludes.

2. Implications of the discrete timing of interest payments

Let $s(t,h)$ denote the log exchange rate (price of foreign currency in dollars) on day t at time h. Intraday interest rates are zero: only balances held overnight attract interest. A position in a currency that is held at a certain cutoff time is deemed to be held overnight and so attracts interest. This cutoff time is 17:00 New York time.[2] We adopt the convention that this time is the end of day t and the start of day $t+1$. Note that settlement in the spot foreign exchange market is $t+2$ for all currency pairs considered in this paper.[3]

UIP is a prediction about the relationship between holding period interest differentials and expected exchange rate changes that is well defined over any holding period. In this section, we describe conditions under which UIP might hold over a sufficiently short window from late on day t to early on day $t+1$, even if agents are risk averse.

Let us first suppose that there are no transactions costs. Consider the strategy of going short the domestic currency on day t at time h_1, investing the proceeds in the foreign currency, and unwinding the position the next day at time h_2. Let h^* denote the span of time for which this position is open, i.e. the span between time h_1 on day t and time h_2 on day $t+1$. The returns from this transaction will be $s(t+1,h_2)-s(t,h_1)-(i_{t,h1}-i^*_{t,h1})$ where $i_{t,h1}$ and $i^*_{t,h1}$ denote the domestic and foreign interest rates agreed to on day t at time h_1 for loans between the day of settlement for day t trades and the day of settlement for day $t+1$ trades. We shall discuss in Section 3 exactly how to measure these interest rates. For notational convenience, we will henceforth drop the dependence of these interest rates on h_1, but it is understood that they are rates as of time h_1. The ex ante expected return from this strategy is the risk premium $\zeta(t,h_1;t+1,h_2)\equiv E_{t,h_1}s(t+1,h_2)-s(t,h_1)-(i_t-i_t^*)$. The

[2] The use of 17:00 New York time (21:00 GMT during daylight savings time, 22:00 GMT at other times) as the standard cutoff time from one value date to the next was widespread before electronic trading became prevalent and is a rigid convention for EBS, the major electronic brokerage system. In any event, our empirical work does not depend crucially on an exact cutoff at 17:00—what really matters is that trades at 16:30 and 19:00 are for different value dates.

[3] Many papers in the UIP literature have ignored the settlement lag in the spot foreign exchange market. The issue is of course unimportant for low frequency analysis but is crucial for a high frequency analysis of the sort that we are undertaking in this paper. Baillie and Osterberg (1997, 2000) take account of the settlement lag issue in UIP regressions at the daily frequency. Bekaert and Hodrick (1993) also take account of the settlement lag issue. They consider UIP over a 30-day horizon, taking careful account of the settlement lag, but still reject UIP.

difference between the actual and expected return from this strategy is $s(t+1,h_2)-E_{t,h_1}s(t+1,h_2)$. Ruling out stochastic jumps in the exchange rate, in the limit as h^* goes to zero, the exchange rate risk borne by the investor gets small,[4] and the expected return should go to zero. Since the interest differential is the same no matter how small h^* is, if the expected return goes to zero in the limit as the span h^* goes to zero, then the exchange rate should be expected to jump by the amount of the interest differential as the value date rolls.[5] One could indeed argue that the expected return from investing in a foreign currency over any arbitrarily short span of time should be small, but what is special about this short interval is that the whole overnight interest differential nevertheless accrues in this interval.

Now transactions costs are important and the overnight interest differential is usually extremely small. But the situation is quite similar to an argument in the literature on stocks going ex dividend (see Lakonishok and Vermaelen, 1986; Elton and Gruber, 1970). This model supposes that the marginal investor is buying or selling stocks for reasons unrelated to the dividend. Equilibrium prices are set by investors who are not deciding whether to buy or sell, but when to buy or sell, and the transactions cost is a fixed cost that they will pay no matter when they buy or sell. Under these conditions, the stock price will jump by the full amount of the dividend at the point when the stock goes ex dividend, even if transactions costs are large.

The lumpy accrual of interest makes a currency just like a stock that goes ex dividend at the very end of each day, so we can make exactly the same argument for why a currency should jump by the amount of the overnight interest differential at the moment that the value date rolls, even in the face of relatively large transactions costs. Specifically, suppose that an investor wants to buy (or sell) foreign currency but is deciding between doing so at the end of day t (at time h_1) or at the start of day t+1 (at time h_2). She faces transactions costs including the bid-ask spread, but we suppose that these are equal on day t at time h_1 and on day t+1 at time h_2. If such an investor is the marginal investor, if time h_1 on day t and time h_2 on day t+1 are sufficiently close that exchange rate risk is small and the investor is indifferent between trading at these two times, then it must be the case that $E_{t,h_1}s(t+1,h_2)+i_t^* \simeq s(t+1,h_2)+i_t^* \simeq s(t,h_1)+i_t$, in equilibrium trading. Otherwise, there would be no trade as all such investors would want to trade either on one side of the 17:00 cutoff time or the other. This means that there should be a jump in the exchange rate at 17:00 that exactly offsets the interest differential that accrues at that instant, and the risk premium $\zeta(t,h_1;t+1,h_2)=E_{t,h_1}s(t+1,h_2)-s(t,h_1)-(i_t-i_t^*)$ should be arbitrarily small if the two times are sufficiently close. Meanwhile, the interest differential is the same no matter what times h_1 and h_2 are.

Now we turn to the implications of this for the UIP regression, estimated over short windows of time. Consider the algebraic identity

$$s(t+1,h_2)-s(t,h_1)=\zeta(t,h_1;t+1,h_2)+u_{t+1} \tag{1}$$

[4] Concretely, the conditional variance of $s(t+1,h_2)$, conditional on the information set at time h_1 on day t goes to zero in the limit as h^* goes to zero.

[5] If the exchange rate were not to jump, the investor could earn the whole overnight interest differential while being exposed to exchange rate risk only over an arbitrarily short span of time.

where u_{t+1} is an error term that must be orthogonal to anything in the information set on day t at time h_1, including the interest differential, $i_t - i_t^*$. We can rewrite this as

$$s(t+1, h_2) - s(t, h_1) = \alpha + \beta(i_t - i_t^*) + \varepsilon_{t+1} \tag{2}$$

where α=0, β=1 and $\varepsilon_{t+1} = \zeta(t,h_1;t+1,h_2) + u_{t+1}$. This is a UIP regression from time h_1 on day t to time h_2 on day t+1. The bias in the OLS estimate of the slope coefficient β is $\mathrm{Cov}(\varepsilon_{t+1}, i_t - i_t^*)/\mathrm{Var}(i_t - i_t^*)$. If time h_1 on day t is close to time h_2 on day t+1, then $\mathrm{Cov}(\varepsilon_{t+1}, i_t - i_t^*)$ should be small because the risk premium $\zeta(t,h_1;t+1,h_2)$ should be small while u_{t+1} is uncorrelated with the interest differential. Crucially, no matter how close the two times are, the variance of $i_t - i_t^*$ will be the same. This gives an argument for why the slope coefficient in the UIP regression from time h_1 on day t to time h_2 on day t+1 might be close to 1 if these times are sufficiently close together. This would of course be in contrast to the results in UIP regressions at more conventional frequencies, but at these lower frequencies there is no reason to think of the risk premium as being small relative to the interest differential.

In reality, there may be stochastic jumps in exchange rates meaning that exchange rate risk and the risk premium will not to go zero in the limit as h^* goes to zero. But the exchange rate risk may still be especially small over a short span of time as the value date rolls meaning that UIP nearly holds over this span of time while the full overnight interest accrues. We are not trying to “prove” on theoretical grounds that we must necessarily observe positive UIP results around the time that the value date rolls. We have simply presented an argument as to why, even in the presence of transactions costs, an OLS estimate of β close to 1 might be observed in a short window spanning the discrete accrual of interest. To find out whether or not this is the case, we turn to the empirical analysis of the next section.

3. Empirical work

Our spot exchange rate data consist of the high frequency bilateral Japanese yen, German mark/euro, Swiss franc and pound sterling exchange rates viz-a-viz the US dollar provided by Olsen and Associates over the years 1988–2002, inclusive. To construct these data, Olsen and Associates record all Reuters quotes, average the bid and ask, and then linearly interpolate the resulting series to get prices every 5 min.[6] We discard weekends, defined as the time from 23:00 GMT on Friday to 22:55 GMT on Sunday, because there is virtually no foreign exchange trading during this time.

Measurement of the interest differential is tricky, although in practice very short term interest rates are highly correlated with each other. We want to measure the interest differential agreed to on day t that applies to a loan between the value dates for day t and day t+1 trades. The appropriate interest rates to be using are spot/next interest rates.

[6] These are based on Reuters indicative quotes, not transaction prices. Danielsson and Payne (2002) compare Reuters indicative quotes and transactions prices, and find that the five-minute returns on the two series are very highly correlated.

A spot/next loan agreed to on day t is a loan for one business day, starting in two business days, which allows the investor to fix on day t the interest rate that is to apply between the value dates for day t and day t+1 trades.[7] We obtained spot/next interest rates for Japan, Germany, Switzerland, the UK and the US, from the BIS.[8] These are expressed at annualized rates. In our regressions, we then divide the annual foreign–US interest rate differentials by 360, and scale this by the number of days between settlement for day t trades and settlement for day t+1 trades (to allow for weekends). For example, if the interest differential is 3.6 percentage points at an annual rate, then this is a one basis point differential at the daily frequency if day t+1 trades settle one calendar day after day t trades.

In our regressions, for each US–foreign country pair, we omit days for which day t is a holiday in either the US or the foreign country, and all of the three preceding weekdays. We do this because foreign exchange trading is unusually light on holidays, and because the spot/next interest differential may not in fact be the correct interest differential between the value dates for day t and day t+1 trades if there is a holiday coming up.[9]

Although we believe that these spot/next interest rates are the appropriate interest rates to use, we nevertheless also redid all the empirical work in this paper using one-week eurocurrency interest rates on day t (scaled by the number of days between the day t and t+1 value dates) instead. These gave very similar results (available from the authors on request), which is not surprising since very short term interest rates are highly correlated with each other.

3.1. The "Wednesday effect"

We divide the sample into days when the interest differential that will accrue on an overnight position is for only 1 day (single-day interest differential days) and days when the interest differential that will accrue is for more than 1 day (multi-day interest differential days). The value dates for Wednesday and Thursday trades are Friday and the following Monday, respectively, so three times the daily interest differential will accrue between Wednesday and Thursday. Because we have cut out holidays, all multi-day interest differential days are Wednesdays. The institutional settlement rules mean that the interest differential is unusually large (in absolute magnitude) on Wednesday nights,

[7] In money markets, an overnight loan is a loan from today until the next business day, a tomorrow/next loan is a loan from day t+1 to day t+2, and a spot/next loan is a loan from day t+2 to day t+3.

[8] Strictly, we want the interest rates as of time h_1 on day t. The BIS collects spot/next interest rates at around 11:45 am Swiss time each day, which will not be time h_1 in our empirical work. Conceivably, this approximation could introduce some measurement error. However, we believe that the variance of this measurement error is very small relative to the variance of interest differentials.

[9] The complicated rules for settlement of spot foreign exchange with holidays in the US/foreign country are discussed by Loopesko (1984), Stigum (1990) and Walmsley (2000). As an example of the kind of issue that can arise, suppose that Thursday is a holiday in Japan but not the US. For Monday trades, dollar–yen settlement will be on Wednesday. For Tuesday trades, dollar–yen settlement will be on Friday. So the interest differential that we want would be from Wednesday to Friday. But a dollar-denominated spot/next loan agreed to on Monday will instead consist of a loan from Wednesday to Thursday.

Table 1
UIP regressions for daily data on US–foreign bilateral exchange rate: single-day interest differential days

Currency	Intercept	Slope	n	R^2	L
Swiss franc	0.007	−4.86	2462	0.26%	0.39
	(0.018)	(2.03)			
Mark/euro	−0.017	−3.45	2406	0.16%	0.22
	(0.014)	(1.97)			
Pound	−0.028	−1.07	2484	0.01%	0.25
	(0.016)	(2.19)			
Yen	0.009	−3.46	2243	0.11%	0.27
	(0.019)	(2.16)			

This table reports the intercept and slope coefficient estimates in the UIP regression with daily data, measured at 16:30 New York time each day, i.e. the regression in Eq. (3) with h set to 16:30 New York Time. White standard errors are reported in parentheses. The number of observations, n, the percentage R^2, and the structural stability test statistic of Nyblom (1989), L, are also reported for each regression. This stability test is a one-sided test which rejects the null of stability if the statistics is greater than the critical value (0.607, 0.748 and 1.074 at the 10%, 5% and 1% levels, respectively). Stability is not rejected even at the 10% level in any case.

though we find it hard to imagine a reason why the risk premia would be unusually important on Wednesday nights. This is the reason why we split out Wednesdays from the rest of the week.

3.2. Basic results

We first considered the ordinary UIP regression at daily frequency. We ran the regression,

$$s(t+1,h) - s(t,h) = \alpha + \beta(i_t - i_t^*) + \varepsilon_{t+1} \tag{3}$$

for single-day interest differential days where h is a fixed time on both days that we set to 16:30. Here and henceforth throughout this paper, all times are New York times. The results are reported in Table 1. For all currencies, the estimated slope coefficient is negative, and is significantly different from one for all currencies except the pound.[10] This rejection of UIP is similar to that found in the literature, in lower frequency regressions.

We next turned to running the proposed regression, as in Eq. (2), over a window from time h_1 to time h_2, for single-day interest differential days. In theory, as discussed in Section 2, we would like to select these times so as to construct the smallest possible window around 17:00. In practice, however, we also want the markets to be reasonably liquid at these times, and trading activity is light right around 17:00. Also, we want to select the times h_1 and h_2 so that the data we use (based on Reuters quotes) can be thought of as referring to prices before and after the rollover time as unambiguously as possible.

[10] Our standard errors here and throughout this paper are heteroskedasticity-robust White standard errors. We do not use autocorrelation-robust standard errors because the windows over which the UIP regressions are run are non-overlapping.

Table 2
UIP regressions from 16:30 to 21:00 New York time: single-day interest differential days

Currency	Intercept	Slope	n	R^2	L
Swiss franc	−0.002	0.79	2462	0.07%	0.47
	(0.006)	(0.69)			
Mark/euro	−0.007	1.02	2406	0.13%	0.60
	(0.005)	(0.73)			
Pound	−0.013	1.44	2484	0.23%	0.27
	(0.004)	(0.72)			
Yen	−0.008	−1.00	2243	0.06%	0.18
	(0.008)	(0.84)			

This table reports the intercept and slope coefficient estimates in the UIP regression with intradaily data, from 16:30 to 21:00 New York time each day, i.e. the regression in Eq. (2) with h_1 and h_2 set to 16:30 and 21:00 New York Time, respectively. White standard errors are reported in parentheses. The number of observations, n, the percentage R^2, and the structural stability test statistic of Nyblom (1989), L (described above), are also reported for each regression. Stability is not rejected at the 10% level in any case.

We therefore set h_1 to 16:30, and h_2 to 21:00 (morning trading in Tokyo).[11] The results, reported in Table 2, are very different from those in the ordinary UIP regression. In all cases except the yen, the slope coefficient is estimated to be positive and not significantly different from one at the 5% significance level.[12] Even for the yen, the estimated coefficient is higher than in the ordinary UIP regression with daily data. This is quite consistent with the idea of the risk premium being small in short windows around the time of the discrete interest payment, as discussed in Section 2.

Importantly, the regression is also much more precisely estimated over the period from 16:30 to 21:00. The reason why is simple. Over this period, the variance in the regressor (the interest differential) is the same as in the daily regression. The variance in the error term is, however, much lower. The signal-to-noise ratio is thereby more favorable to precise inference using the judiciously chosen intradaily interval that spans the actual interest payment.

The slope coefficients in Eq. (2) setting h_1 to 16:30, and using various values of h_2 from 19:00 to 16:30 the next day are plotted in Fig. 1. A distinctive pattern can be seen in these plots, whereby the later h_2 is in the day, the lower the coefficient estimate is. Setting h_2 to 16:30 just gives a standard daily-frequency UIP regression. The pattern is especially dramatic for the mark/euro and Swiss franc, but can also be observed for the pound and to a lesser extent for the yen. This figure shows the central finding in the paper graphically—on average, currencies do indeed move in the direction predicted by UIP in short windows

[11] Throughout the paper we set h_1 to 16:30. A potential concern is that bid-ask spreads are relatively elevated by 16:30 and that trading activity is reduced at this time. As a robustness check, we redid all the results in this paper setting h_1 to 15:30 (when trading is more active) instead. The results, available from the authors on request, are substantively unchanged.

[12] As a robustness check, we re-ran the results in Table 2 dropping Fridays. This reduces the sample size, but also eliminates observations with the relatively large change in the exchange rate across the weekend. Deleting Fridays, the hypothesis that the slope coefficient is equal to one is not rejected at the 5% significance level for any of the four currencies, slightly strengthening our results (though it is still rejected at the 10% level for the yen).

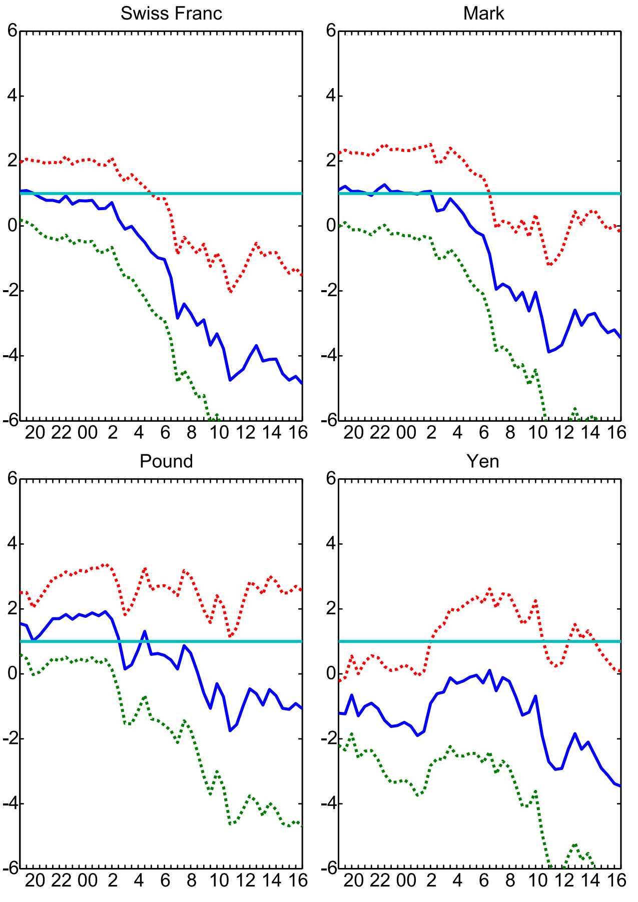

Fig. 1. Estimate of β in Eq. (2), with 90% confidence interval, from 16:30 to time h_2, plotted against h_2. Estimated over single-day interest differential days.

around the time of the discrete interest payment, but then move back the other way, and the latter effect dominates at the daily frequency. The results are weakest for the yen, perhaps because the flow of yen-related news is more intense than news related to other currencies in early Asian trading (e.g. Japanese GDP data come out at 8:50 in the morning Tokyo time, which is either 18:50 or 19:50 in New York).

We repeated this analysis for the multi-day interest differential days. Tables 3 and 4 give the results from the UIP regressions over the 16:30–16:30 and 16:30–21:00 windows on multi-day interest differential days. The slope coefficients setting h_1 to 16:30, and using various values of h_2 from 19:00 to 16:30 the next day for multi-day interest differential days are plotted in Fig. 2. For all four currencies, the coefficient estimates are shifted up relative to those obtained in Fig. 1. The hypothesis that the slope coefficient is equal to one is not rejected using either the 16:30–21:00 window or the daily window, though the standard errors are much larger in the latter case. The slope coefficient is significantly

Table 3
UIP regressions for daily data on US–foreign bilateral exchange rate: multi-day interest differential days [coefficient estimates (standard errors in parentheses)]

Currency	Intercept	Slope	n	R^2	L
Swiss franc	0.036	1.58	585	0.22%	0.11
	(0.040)	(1.50)			
Mark/euro	0.037	2.60	569	0.71%	0.16
	(0.030)	(1.33)			
Pound	0.088	2.70	591	0.87%	0.16
	(0.033)	(1.40)			
Yen	0.037	1.26	544	0.12%	0.11
	(0.046)	(1.51)			

As for Table 1.

positive over the 16:30–21:00 window for all currencies except the yen. These results are again consistent with the idea that UIP is most likely to work over short windows of time when the interest differential is unusually large (in absolute magnitude) relative to any risk premium. The interest differential is unusually large on multi-day interest differential days purely because of the institutional fact that the value date jumps 3 days between Wednesdays and Thursdays.

The standard errors for the multi-day interest differential days (Fig. 2) are smaller than for the single-day interest differential days (Fig. 1), even though the sample size for the single-day interest differential days is much larger. This is simply because the variance of the regressor is so much bigger on the multi-day interest differential days.

Tables 1–4 also report stability tests (Nyblom (1989)) and the R^2 for each regression. The null of stability is not rejected in any case. The R^2 in each regression is tiny: but with our long span of 15 years of data our estimates are nevertheless precise enough to be useful.

Table 5 reports the standard deviation of interest differentials (on single- and multi-day interest differential days) and of the exchange rate changes (over 16:30–21:00 and 16:30–16:30 windows). The variability of interest differentials is extremely small.

Table 4
UIP regressions from 16:30 to 21:00 New York time: multi-day interest differential days [coefficient estimates (standard errors in parentheses)]

Currency	Intercept	Slope	n	R^2	L
Swiss franc	−0.001	0.87	585	1.50%	0.31
	(0.008)	(0.31)			
Mark/euro	−0.000	1.02	569	2.18%	0.28
	(0.007)	(0.29)			
Pound	−0.001	0.94	591	1.82%	0.18
	(0.007)	(0.36)			
Yen	0.016	0.59	544	0.27%	0.11
	(0.013)	(0.46)			

As for Table 2.

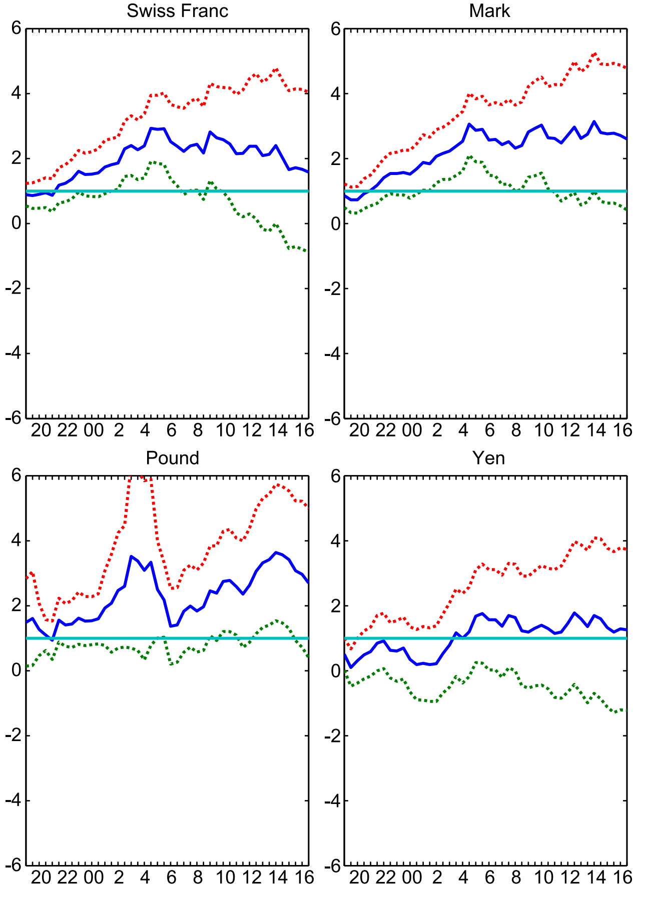

Fig. 2. Estimate of β in Eq. (2), with 90% confidence interval, from 16:30 to time h_2, plotted against h_2. Estimated over multi-day interest differential days.

Nevertheless, with our long sample, we are able to observe results that are favorable to UIP over short windows of high frequency data that span the time of the discrete interest payment.

Table 5
Standard deviations (basis points) of interest differentials and exchange rates

Currency	Interest differentials		Exchange rate returns	
	Single-day	Multi-day	16:30–21:00	16:30–16:30
Swiss franc	0.767	2.294	19.93	73.73
Mark/euro	0.773	2.288	20.97	68.13
Pound	0.645	2.066	18.58	61.51
Yen	0.685	2.062	27.47	73.19

The standard deviation of the exchange rate returns in basis points means the standard deviation of 10,000 times the log change in the exchange rate over the relevant horizon.

Table 6
Percentage of single-day interest differential days when UIP predicts the correct sign of the exchange rate change (standard errors from the variance of a binomial in parentheses)

Currency	16:30–16:30 Change	16:30–21:00 Change
Swiss franc	48.74	50.13
	(1.01)	(1.01)
Mark	47.46	52.11
	(1.02)	(1.02)
Pound	50.14	57.02
	(1.00)	(1.00)
Yen	49.18	48.37
	(1.06)	(1.06)

This table reports the fraction of single-day interest differential days in which the currency with the lower interest rate appreciates, as predicted by UIP, both in the period from 16:30 to 16:30 the next day, and in the period from 16:30 to 21:00.

3.3. UIP and the direction of the exchange rate change

As a simple robustness check, and to guard against results being driven by outliers, we calculated the proportion of times that the exchange rate change had the correct sign as predicted by UIP, both for single- and multi-day interest differentials, over the period from 16:30 to 16:30 the next day and from 16:30 to 21:00 the next day. Where the foreign interest rate is greater than the US interest rate, UIP would call for the dollar to appreciate, and vice versa. The percentage of days on which this prediction is in fact correct is reported in Tables 6 and 7, for single- and multi-day interest differential days, respectively. Confidence intervals are included, constructed by the formula for the variance of a binomial distribution.

For the Swiss franc, mark/euro and yen, using data at the daily frequency (16:30–16:30) on single-day interest differential days (Table 6), the estimated proportion of times that UIP predicts the correct sign is less than half. In other words, UIP gets the sign wrong more often than it gets it right. However, over the short window from 16:30 to 21:00, UIP gets the sign right more than half the time for all currencies except the yen, and significantly more than half the time for the mark/euro and pound. On multi-day interest

Table 7
Percentage of multi-day interest differential days when UIP predicts the correct sign of the exchange rate change (standard errors from the variance of a binomial in parentheses)

Currency	16:30–16:30 Change	16:30–21:00 Change
Swiss franc	51.71	55.15
	(2.07)	(2.06)
Mark	49.38	54.48
	(2.10)	(2.09)
Pound	49.58	57.65
	(2.06)	(2.04)
Yen	50.83	55.43
	(2.15)	(2.13)

As for Table 6, except for multi-day interest differential days.

differential days (Table 7), the proportion of times that UIP calls the sign of the exchange rate movement correctly is significantly greater than 0.5 over the 16:30 to 21:00 window for all four currencies, at the 10% significance level.

4. Conclusion

No interest is paid on intradaily positions. Rather, interest is paid discretely, at the point when a position is rolled over from one day to the next. The common rollover time is determined by market convention. This practice has potential implications for high-frequency exchange rate movements that we have discussed and exploited in this paper.

UIP is both central to theoretical models, and an enormous empirical failure over conventional horizons. But over very short windows of high frequency data that span the time of the discrete interest payment, we find that the slope coefficient in the UIP regression is close to one, and fairly precisely estimated. Relatively large discrete interest payments accrue on positions held between Wednesday and Thursday, as the value date advances across the weekend, and especially positive results are obtained on these multi-day interest differential days. Thus, over the shortest windows with the largest interest differentials, the interest differential appears to swamp any risk premium. An interpretation of this sort says nothing about the nature of the risk premium, and it remains a puzzle that it is sufficiently large and time-varying as to lead to the rejection of UIP at conventional frequencies. The positive uncovered interest parity results disappear quickly if we add even just a few hours to the span of the window, except on the multi-day interest differential days, and this remains very much a puzzle. Nevertheless, if we thought that the failure of UIP represented irrational behavior on the part of market participants, then we might expect it to fail even over these very short horizons with discrete interest differentials.

Acknowledgements

We are grateful to Charles Engel, three anonymous referees, Richard Adams, Kathryn Dominguez, Jon Faust, Dale Henderson, Avner Kalay, Mike Leahy, Richard Lyons, Chris Neely, Richard Olsen and John Rogers for their helpful comments. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any person associated with the Federal Reserve System. The spot/next interest rate data used in this paper are available from the authors on request.

References

Baillie, R.T., Osterberg, W.P., 1997. Central bank intervention and risk in the forward premium. Journal of International Economics 43, 483–497.

Baillie, R.T., Osterberg, W.P., 2000. Deviations from daily uncovered interest parity and the role of intervention. Journal of International Financial Markets, Institutions, and Money 10, 363–379.

Baillie, R.T., Lippens, R.E., McMahon, P.C., 1983. Testing rational expectations and efficiency in the foreign exchange market. Econometrica 51, 553–563.

Bekaert, G., Hodrick, R.J., 1993. On biases in the measurement of foreign exchange risk premiums. Journal of International Money and Finance 12, 115–138.

Bilson, J.F.O., 1981. The "Speculative Efficiency" hypothesis. Journal of Business 54, 435–451.

Danielsson, J., Payne, R., 2002. Real trading patterns and prices in spot foreign exchange markets. Journal of International Money and Finance 21, 203–222.

Elton, E.J., Gruber, M.J., 1970. Marginal stockholder tax rates and the clientele effect. Review of Economics and Statistics 52, 68–74.

Engel, C., 1996. The forward discount anomaly and the risk premium: a survey of recent evidence. Journal of Empirical Finance 3, 123–192.

Froot, K.A., Thaler, R.H., 1990. Anomalies: foreign exchange. Journal of Economic Perspectives 4, 179–192.

Fujii, E., Chinn, M., 2001. Fin de siècle real interest parity. Journal of International Financial Markets, Institutions, and Money 11, 289–308.

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

Lakonishok, J., Vermaelen, T., 1986. Tax-induced trading around ex-dividend days. Journal of Financial Economics 16, 287–319.

Lewis, K.K., 1995. Puzzles in international financial markets. In: Grossman, G.M., Rogoff, K.Handbook of International Economics vol. 3. Elsevier, Amsterdam, pp. 1913–1949.

Loopesko, B.E., 1984. Relationships among exchange rates, intervention, and interest rates. Journal of International Money and Finance 3, 257–277.

Lyons, R.K., Rose, A.K., 1995. Explaining forward exchange bias. . .intraday. Journal of Finance 50, 1321–1329.

Meredith, G., Chinn, M., 1998. Long-horizon uncovered interest rate parity. National Bureau of Economic Research Working 6797.

Nyblom, J., 1989. Testing for the constancy of parameters over time. Journal of the American Statistical Association 84, 223–230.

Stigum, M., 1990. The Money Market. McGraw Hill, New York.

Walmsley, J., 2000. The Foreign Exchange and Money Markets Guide. Wiley, New York.