Journal of International Money and Finance 27 (2008) 1284–1302

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Is the forward bias economically small? Evidence from European rates[☆]

Piet Sercu[a,*], Martina Vandebroek[a,1], Xueping Wu[b,2]

[a] KU Leuven, Graduate School of Business Studies, Naamsestraat 69, B-3000 Leuven, Belgium
[b] City University of Hong Kong, Department of Economics and Finance, Tat Chee Avenue, Kowloon, Hong Kong

JEL classification:
F31

Keywords:
Forward bias
Transaction costs
Trading rule
EMS
ERM

A B S T R A C T

For the purpose of testing uncovered interest parity (UIP), rates of European currencies against the Mark offer a distinct advantage: the admissible band of the Exchange Rate Mechanism (ERM) induces statistically significant mean-reversion in weekly rates. Thus, unlike for freely floating rates, there is an expectation signal that has non-trivial variation and is sufficiently traceable for research purposes. When running the standard regression tests of the unbiased-expectations hypothesis at the one-week horizon, we nevertheless obtain essentially zero coefficients for intra-EMS exchange rates (and the familiar negative coefficients for extra-EMS rates). Even more puzzlingly, lagged exchange-rate changes remain significant when added to the regression, a feature that seems hard to explain as a missing-variable effect. The deviation from UIP is significant not just statistically but also economically: trading-rule tests reveal that for sufficiently large filters the average profit per trade exceeds transaction costs, and that cumulative gains can be quite impressive. The size of the profits and the patterns from buy versus sell decisions also allow us to reject the hypotheses of either a risk premium or peso issues about realignments as sufficient explanations.

☆ The authors thank Marco Lyrio, Hans Dewachter and Hans Degryse for useful comments and criticisms, and likewise many participants during presentations at KU Leuven, Universitaet Zuerich, and Konstanz Universitaet. The comments by two anonymous referees from this journal have been very valuable too. Any remaining errors are ours.
* Corresponding author. Tel.: +32 16 32 6756; fax: +32 16 32 6632.
E-mail addresses: piet.sercu@econ.kuleuven.ac.be (P. Sercu), martina.vandebroek@econ.kuleuven.ac.be (M. Vandebroek), efxpwu@cityu.edu.hk (X. Wu).
1 Tel.: +32 16 32 6975; fax: +32 16 32 6632.
2 Tel.: +852 2788 7577; fax: +852 2788 8806.

0261-5606/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jimonfin.2008.06.002

1. Introduction

In their review paper on the bias in the interest-rate differential as a predictor of exchange-rate changes, Froot and Thaler (1990) end on a somewhat ambivalent note. Having documented the widespread failure of the standard regression tests, they note that the amount of exploitable predictability seems to be tiny, the uncertainty from taking advantage of the bias huge, and the profits small. Roll and Yan (2000) add that the significance of the standard regression test is overstated because the regressor has very long memory; taking into account the likely true standard errors, the classical regression evidence becomes totally uninformative. We contribute to the debate in four ways. First, we study a new sample: managed intra-European rates against the DEM instead of the usual USD rates. This sample is unusual in that there is very strong predictability in the exchange-rate changes, with R^2's of 20–35%; yet our unbiased-expectations regressions do as badly as usual. Second, we offer statistical evidence that does not suffer from long-memory problems and strongly rejects unbiased expectations. Third, we show that the profits from trading on reversal are huge and the risks are small. In an average year we would be bearish, bullish and neutral about equal amounts of time, and given a buy or sell signal we'd have conditional annualized returns of about 20%, netting us about 14% *p.a.* extra. In addition, the risk is low, which provides us very high Sharpe ratios. Lastly, none of the standard rational theories seems to be able to explain a bias that behaves like ours, swinging from hugely positive to zero or hugely negative in a matter of days. In short, we think the bias is not a marginal phenomenon, and remains as challenging as ever.

The background is as follows: the forward bias is a long-standing puzzle in international finance: the forward premium or interest-rate differential between two currencies systematically mispredicts future exchange-rate changes, and does so by far more than conventional risk theories predict.[3] Fama (1984) infers that there must be a missing variable – possibly a risk premium – whose variation over time must be even larger than the time-series variance of these conditional expected changes.

It is, however, not always obvious how momentous some of these findings are. For example, Fama's moment condition can be met by a missing variable that merely consists of misalignments too small to matter relative to trading costs, combined with expectations that hardly change over time. In the same vein, Froot and Thaler (1990) argue that the expected profits are economically small, especially relative to the residual uncertainty. Bansal and Dahlquist (2000) find that the negative association between forward premium and realized exchange-rate change may be a 1970s–1980s OECD phenomenon; in wider and more recent data the picture looks better. De Grauwe (1989) adds that the negative coefficients disappear when the base currency is not the USD. Lastly, the statistical significance of the findings is, at best, unclear: Roll and Yan (2000) show that the near-unit-root characteristics of the forward premium invalidate the usual standard deviations, and Schotman et al. (1997) document a similar phenomenon related to thick tails and outliers in the regressor.

In this paper, we study a new data set, *viz.* 1985–1998 weekly changes in exchange rates against the German Mark. In Section 1 we document a non-trivial short-run predictability of exchange rates created by affiliation to the European Monetary System's Exchange Rate Mechanism (ERM), whether this affiliation be a formal membership like for the Guilder or a market-perceived loose association as for the Swiss Franc. If Uncovered Interest Parity (UIP) holds, the forward premium should pick up this strong predictability. However, our results from standard UIP regression tests on one-week Mark data are not substantially better, and in a sense even worse than in other data sets, as Section 2 shows. First, in our Cumby–Obstfeld–Fama (henceforth COF) regression tests of UIP the slopes for European monies are still far from unity, and very sensitive to data period and estimation technique. Second, the autocorrelations are not affected by the introduction of the forward premium as a regressor – evidence that positively contradicts UIP, is hard to explain as a missing-variables problem, and does not suffer from a near-unit-root effect. Even more interesting is the economic significance of these combined phenomena (predictable rates, and un-aligned forward premia): the trading-rule

[3] The seminal papers are Cumby and Obstfeld (1984) and Fama (1984); Froot and Thaler (1990) offer an excellent early survey; Hollifield and Uppal (1997) bring up a general-equilibrium analysis predicting a small bias; Bansal (1997) shows the bias may be large and asymmetric. The more recent kernel-based models are discussed in the concluding section.

tests of Section 3 reveal that the cumulative 11-year returns from exploiting the deviations from UIP, over and above the average risk premium, turn out to be a few hundred percentages, in one case even reaching 600%. The total size and the details of the returns then allow us to conclude, in Section 4, that none of the usual suspects is likely to be singly and fully responsible for such a result. Rather, we need at least a combination, possibly reinforced by a more subtle pricing effect (like the transaction-cost-induced hysteresis features stressed by Baldwin (1990)), or a market inefficiency. In short, we think the puzzle has deepened: the amount of money left on the table was quite easy to spot and uncannily large.

Our results are different from earlier work on forex technical trading[4]: we study managed rates not floating ones (against the Mark rather than the dollar) and, accordingly, the trading rule is based on reversal rather than momentum. There are also a number of papers that, like ours, are more UIP-focused when testing trading rules, but their buy or sell decisions are based on the risk-free rates rather than on exchange-rate forecasts.[5] As, in the short and medium run, interest rates are largely independent of expected exchange-rate movements, picking high-yield currencies does pay. But our exchange-rate-related signals are more frequent and more profitable than yield-based ones.

2. Predictability in intra-European exchange rates

The data we start from are daily and weekly London exchange rates against Sterling (GBP)[6] and one-week spot interest rates, both from Datastream (midpoint Barclays quotes, or for USD/BEF, National Westminster quotes). There is nothing unusual in our results for the BEF, so the source seems to have little impact. Our data cover almost 13 years, starting in 1985/6/1 (the date where Datastream's coverage is expanded) to 1998/4/1 (the date Euroland rates became quasi-fixed). In our tests, these GBP rates were re-expressed into DEM per unit of foreign currency. This must have introduced small deviations relative to simultaneous quotes made immediately in DEM, but it is hard to see how this could have influenced the outcome of the tests in any significant way.

The 10 resulting exchange rates against the DEM are mostly European: four core-ERM currencies (BEF, DKK, FRF, and NLG) and three intermittently or informally associated with the ERM (CHF, ITL, and as the weakest affiliate, GBP).[7] To verify whether the findings are indeed typical for European exchange rates we also include three major outside currencies (USD, JPY, and CAD).

Exchange-rate changes against the USD are commonly accepted to exhibit a slight but statistically clear positive autocorrelation, leading to the profitable momentum-based trading results cited in the introduction. For ERM-member rates, tied to each other by a narrow band, one naturally expects also cross-correlations: currencies that for some reason did not follow the pack immediately must catch up later. In exploratory research that led to this paper we studied daily data against the USD, and found large cross-correlations among ERM-member rates, and weaker cross-links with quasi-members. These links are significant for one- and two-day lags, but rarely beyond that. Also, we found that, from the USD point of view, the DEM was the bellwether currency, with the other European currencies following the DEM's movements against the USD, partly with a lag. Switching data to a DEM basis, we accordingly expect reversal rather than continuation patterns. We indeed found clear negative autocorrelations at lags 1 and 2 for ERM-member rates, and traces of negative autocorrelation for the other European currencies. These results are available on request.

[4] For results from daily data see Levich and Thomas (1993), Neely et al. (1997), Sweeney (1986), Taylor (1994), and Surajaras and Sweeney (1992). For weekly data see Kho (1996). Okunev and White (2003) use monthly data. Intraday data do not seem to work, see Neely and Weller (2003) and Raj (2000).

[5] The seminal article is probably Robinson and Warburton (1980); see also Green (1992), who also uses the inflation rate, and LeBaron (1998) or Villanueva (2005).

[6] GBP = Great Britain, Pound; USD = US, Dollar; BEF = Belgium, Franc; DEM = Deutschland, Mark; DKK = Denmark, Crown; FRF = France, Franc; NLG = Netherlands, Guilder; CHF = Confederatio Helvetica (Switzerland), Franc; ITL = Italy, Lira; CAD = Canada, Dollar; JPY = Japan, Yen.

[7] Although the Swiss central bank denies intervention, the CHF is widely seen as informally linked to the DEM and, now, the Euro. The ITL was an ERM member but with an unusually wide band. The GBP unilaterally tracked the ECU in 1990–1991 as a prelude to formal ERM membership, in the spring of 1992, but dropped out in September 1992.

Tests of whether that predictability is reflected in interest-rate differentials are hampered by the fact that for most currencies Datastream offers no long histories of one-day spot interest rates from which forward premia can be reconstructed;[8] but coverage of spot one-week rates is adequate. Thus, in Table 1 we document also the autocorrelation pattern for one-week exchange-rate movements rather than one-day returns, first estimated in the regular simple way ("OLS"), and then taking into account the substantial fluctuations in uncertainty ("GARCH", using a GARCH(1,1) variance model and an AR(1) mean equation). We observe strong autocorrelations at lag one for the first four currencies, the ERM core members, many times larger than their standard deviations, and basically no autocorrelation for the other currencies. We also look at subperiods: the period before September 1992 with a narrow ERM band and many formal realignments, the turbulent September 1992–December 1993 period, and the more quiet wide-band period as of 1994 that ended in the fixed rates for Euro-currencies. The negative autocorrelations remain clearly present in the first and last subperiods. Unsurprisingly, in the turbulent transition period they are less pronounced, statistically as well as economically. Note that for the ITL, with its wide band, the autocorrelations are weaker, but for core members the coefficients are large not just statistically but also economically, with averages in excess of −0.20 and individual cases up to −0.50. We conclude that there was a substantial predictability in daily as well as weekly exchange-rate changes.

3. Regression tests of UIP, weekly data

3.1. Standard tests

According to the UIP hypothesis, expected exchange-rate changes should be offset by differentials in the interest earned, or, equivalently, by the forward premium.[9] Formally, the hypothesis is $E_{t-1}(\tilde{s}_t) = \mathrm{FP}_{t-1}$, where $\tilde{s}_t$ denotes $[S_t - S_{t-1}]/S_{t-1}$, the simple percentage change[10] in the exchange rate, and FP_{t-1} denotes the forward premium set at time $t-1$ for delivery at t.[11] A familiar test is to run the Cumby–Obstfeld–Fama (COF) regression of exchange-rate changes on forward premia, possibly augmented by other variables X_{t-1} known at $t-1$:

$$E_{t-1}\left(\tilde{s}_t\right) = \kappa_0 + \kappa_1 \mathrm{FP}_{t-1} + \kappa_2 X_{t-1}. \tag{1}$$

The UIP hypotheses predict $\kappa_0 = 0 = \kappa_2$ and $\kappa_1 = 1$.

One simple test is to derive an unconditional version of Eq. (1), which can then be tested by a cross-sectional regression with, as the left-hand side variable, mean realized rates of change in the spot rates computed from long time series and, as the regressor, long-term mean forward premia. When we do so in our smallish database (10 currencies, 13 years) we find an OLS slope coefficient of 0.57 and an R^2 ($\overline{R}^2$) of 0.42 (0.35). Larger databases – Lothian and Wu (2005) even look at two centuries of data – usually provide a slope that is closer to unity. So in the long-run, static buy-and-hold strategies seem to have similar payoffs across currencies, with capital gains by and large offsetting low interest rates and *vice versa*.

8 Overnight and/or tomorrow/next interest rates are sometimes available, but in daily tests we'd need one-day spot (that is, second/third working day) because spot forex is delivered on the second working day.

9 In the computations we computed forward premia from LIBOR interest rate, the rates against which also banks set their forward premia.

10 We use simple percentages because asset pricing theory makes predictions about prices or simple percentage changes therein. Statements about logs can be made only at the cost of unnecessary and unrealistic additional assumptions like lognormality. We are aware of Siegel's paradox. But the paradox serves to remind us that UIP cannot be true, and trying to lessen its empirical impact does not seem desirable to us.

11 The time subscripts reflect the moment the variable is known: t for s, and $t-1$ for FP or for the risk-free rates for investments between times $t-1$ and t. In the exposition, we use the terms "forward premium" and "interest-rate differential" interchangeably. In the tests we compute forward premia using one-week LIBOR rates – the banks' standard – taking into account the two-working-days delivery rule in spot and forward markets, the 365-days-per-year convention for GBP and (pre-1999) BEF interbank money markets, and the 360-days-per-year convention for other currencies.

Table 1
First-order autocorrelations, weekly
$E_{t-1}(\tilde{s}_t) = \kappa_0 + \rho_1 \tilde{s}_{t-1}$

	Autocorrelation coefficients ρ_1 for individual currencies										Averages	
	BEF	DKK	FRF	NLG	CHF	ITL	GBP	JPY	CAD	USD	ERM	Other
Total period (1985/6–1998/3), $\sigma(\rho) = 0.039$												
OLS	*−0.28	*−0.15	*−0.17	*−0.48	−0.03	0.11	0.01	−0.01	0.01	0.02	−0.22	0.03
GARCH	*−0.38	*−0.22	*−0.11	*−0.44	−0.00	−0.01	0.01	−0.01	0.04	0.01	−0.23	0.01
Early ERM *(tight band, 1985/6–1992/8), $\sigma(\rho) = 0.052$*												
OLS	*−0.36	*−0.30	*−0.13	*−0.51	−0.03	0.03	0.00	−0.03	0.03	0.01	−0.26	−0.01
GARCH	*−0.39	*−0.28	*−0.17	*−0.52	0.02	0.12	−0.01	−0.04	0.00	−0.01	−0.27	0.01
September 1992–end 1993 (turbulence, 1992/9–1993/12), $\sigma(\rho) = 0.12$												
OLS	0.02	0.01	*−0.24	*−0.45	0.02	0.13	0.04	0.02	−0.08	0.08	−0.13	0.04
GARCH	−0.21	0.01	*−0.26	*−0.32	−0.00	0.13	0.07	0.08	−0.21	0.04	−0.16	−0.02
Late ERM *(wide band, 1994/1–1998/3), $\sigma(\rho) = 0.066$*												
OLS	*−0.39	−0.09	*−0.17	*−0.26	−0.08	0.11	−0.03	−0.13	0.09	−0.01	−0.20	0.01
GARCH	*−0.42	*−0.18	*−0.16	*−0.26	−0.05	−0.08	−0.03	−0.13	0.11	−0.01	−0.21	−0.03

The variables s_t are weekly percentage changes in the exchange rate against the DEM. Autocorrelations are estimated using OLS and GARCH(1,1) specifications for the variance. The averages shown are for the first and second sets of five currencies, respectively, labeled somewhat inaccurately "ERM" and "other". Standard deviations are shown in the panel headers, and an asterisk denotes significance at the 1% level (one-sided).

Long-run, static investments are quite different from our day-trading strategies which, as we show in Section 3, do generate spectacular return differences across currencies. But also regressionwise the picture already changes substantially if we test Eq. (1) as a set of time-series relations. In Table 2 we summarize our results for Eq. (1) on weekly data against the DEM. We use OLS and two system estimators, Full Information Maximum Likelihood (FIML) assuming normality and constant moments, and Generalized Method of Moments (GMM), using all 10 forward premia and 10 lagged changes as instruments. Estimates are provided for the entire data set and for each of the three subperiods defined before, except for GMM where the second subperiod is too short to handle 210 moment conditions. Since the forward premium is almost a unit-root process, the regular t-tests vastly overstate the significance; so our (two-sided) probabilities against H_0: $\kappa_1 = 1$ are obtained from the Lagrange multiplier test, a statistic which fares much better in the Bekaert and Hodrick (2001) experiments on UIP tests. But our main conclusions do not rest on just this statistic. One alternative test is a significance check on the results of a trading rule, see Section 3. But another test is already possible within the current regression framework: add s_{t-1} an additional regressor X_{t-1} in Eq. (1). There is no unit-root problem with these variables; the standard deviation for its coefficient, κ_2, is given in the header of each panel in the table.

In terms of the slope coefficient for the forward premium, the picture is qualitatively not better than in the extant literature despite the predictability we just documented. True, in Panel A of Table 2 there are some promising averages: at 0.71, the mean OLS slope for the forward premium is rather good. But the median slopes are already substantially lower than the simple averages, and the two system estimators tend to come up with much lower slopes. Upon closer inspection, the positive equation-by-equation slopes are entirely due to the tumultuous middle period, where the size of the estimates is, in fact, bizarre. In the larger subsamples (Panels A2 and A4), negative slopes dominate, and the recent figures are worse than the early ones. The evidence from the COF coefficients is, in short, not reassuring: DEM-based test results do not provide any better support for UIP, and there is no improvement over time either.

While we observe no clear difference between intra-European versus other rates in Panel A, there is a sharp divide in Panel B, where ERM members still show massive negative first-order autocorrelation. The standard t-test for this variable is reliable, and most coefficients remain significant even though the forward premium has been added as a regressor. Also algebraically the κ_2s in Table 2 are strikingly similar to the ρ_1s in Table 1. Thus, forward premia do not at all pick up the predictability inherent in the negative autocorrelation of s. The significance of the lagged changes is puzzling: the predictability

Table 2
Cumby–Obstfeld–Fama tests of UIP
$E_{t-1}(\tilde{s}_t) = \kappa_0 + \kappa_1 \mathrm{FP}_{t-1} + \kappa_2 \tilde{s}_{t-1}$

	Coefficients for individual currencies										Central values	
	BEF	NLG	DKK	FRF	ITL	CHF	GBP	JPY	CAD	USD	Avg	Med
Panel A: COF slope coefficient (κ_1)												
A1. Total period (1985/6–1998/3)												
OLS	−0.05	−0.18	0.38	0.69	1.17	−0.33	2.07	0.30	0.26	2.75	0.71	0.34
p (lm)	0.77	0.90	0.82	0.93	0.96	0.77	0.71	0.87	0.79	0.52		
FIML	−0.34	−0.36	0.17	0.46	0.32	−0.33	0.97	−1.06	−0.48	1.95	0.13	−0.08
p (lm)	0.00	0.00	0.00	0.07	0.25	0.04	0.87	0.10	0.00	0.26		
GMM	−0.11	−0.59	0.27	0.38	0.53	−0.38	1.57	−0.56	0.10	2.08	0.33	0.18
p (lm)	0.00	0.00	0.00	0.00	0.03	0.00	0.01	0.03	0.00	0.01		
A2. Early ERM (tight band, 1985/6–1992/8)												
OLS	−0.37	−0.33	−0.21	1.71	0.38	−0.44	0.80	−3.96	−1.90	−0.42	−0.47	−0.35
p (lm)	0.76	0.92	0.78	0.89	0.87	0.77	0.97	0.42	0.48	0.70		
FIML	−0.58	−0.67	−0.28	1.49	0.17	−0.28	−0.62	−4.42	−1.48	−0.96	−0.76	−0.60
p (lm)	0.00	0.00	0.00	0.17	0.01	0.06	0.04	0.00	0.00	0.09		
GMM	−0.46	−0.75	−0.17	1.33	0.28	−0.41	0.93	−4.52	−1.37	0.30	−0.49	−0.29
p (lm)	0.00	0.00	0.00	0.04	0.00	0.00	0.84	0.00	0.00	0.16		
A3. September 1992–end 1994 (turbulence, 1992/9–1993/12)												
OLS	1.74	2.48	0.89	−2.42	9.31	−12.53	3.44	57.92	9.06	63.30	13.32	2.96
p (lm)	0.97	0.98	0.99	0.79	0.40	0.81	0.84	0.21	0.83	0.08		
FIML	0.57	−0.62	−0.04	−3.78	7.13	−3.81	5.35	19.29	−5.43	12.10	3.08	0.27
p (lm)	0.80	0.71	0.19	0.00	0.20	0.62	0.12	0.21	0.25	0.24		
A4. Late ERM (wide band, 1994/1–1998/3)												
OLS	−1.11	−1.55	−4.31	−2.35	−5.52	−3.90	−7.23	−19.06	−2.66	−7.20	−5.49	−4.10
p (lm)	0.94	0.95	0.77	0.76	0.46	0.84	0.29	0.52	0.65	0.30		
FIML	−2.18	−1.85	−4.18	−2.22	−4.15	−2.71	−1.74	−23.64	−5.08	0.31	−4.74	−2.47
p (lm)	0.04	0.00	0.00	0.00	0.01	0.35	0.01	0.05	0.00	0.80		
GMM	0.07	−1.63	−4.28	−2.12	−5.30	−3.57	−6.80	−17.82	−2.58	−6.79	−5.08	−3.92
p (lm)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		
											Averages	
											ERM	Other
Panel B: Autoregression coefficient (κ_2)												
B1. Total period (1985/6–1998/4), $\sigma(\kappa_2) = 0.039$												
OLS	*−0.28	*−0.48	*−0.16	*−0.18	*0.10	−0.04	0.01	0.01	−0.01	0.02	−0.20	0.00
FIML	*−0.29	*−0.43	*−0.17	*−0.21	0.08	−0.03	0.02	0.02	−0.01	0.02	−0.20	0.00
GMM	*−0.28	*−0.45	*−0.16	*−0.17	0.09	−0.02	0.01	0.00	−0.01	0.02	−0.19	0.00
B2. Early ERM (tight band, 1985/6–1992/8), $\sigma(\kappa_2) = 0.052$												
OLS	*−0.35	*−0.51	*−0.30	*−0.14	−0.03	−0.03	0.00	−0.05	0.03	−0.01	−0.27	−0.01
FIML	*−0.35	*−0.46	*−0.29	*−0.12	0.01	−0.03	0.04	−0.02	0.04	0.02	−0.24	0.01
GMM	*−0.33	*−0.48	*−0.29	*−0.14	−0.03	−0.04	0.02	−0.03	0.04	0.02	−0.26	0.00
B3. September 1992–end 1993 (turbulence, 1992/9–1993/12), $\sigma(\kappa_2) = 0.12$												
OLS	0.01	*−0.44	−0.05	−0.22	−0.09	−0.03	0.06	−0.12	0.03	0.08	−0.16	0.01
FIML	−0.13	*−0.39	−0.09	*−0.31	−0.06	0.05	0.08	−0.04	−0.04	0.05	−0.20	0.02
B4. Late ERM (wide band, 1994/1–1998/3), $\sigma(\kappa_2) = 0.066$												
OLS	*−0.39	*−0.27	−0.14	*−0.18	0.09	−0.09	−0.05	0.09	−0.14	−0.02	−0.18	−0.04
FIML	*−0.37	*−0.25	*−0.16	*−0.16	0.05	−0.10	−0.02	0.07	−0.09	0.01	−0.18	−0.02
GMM	*−0.39	*−0.26	−0.13	*−0.18	0.09	−0.09	−0.05	0.08	−0.13	−0.02	−0.18	−0.04

The regressands s_t are weekly percentage changes in the exchange rate against the DEM, the regressors are the beginning-of-period one-week forward premium (FP_t) and the lagged regressand. The estimation methods are OLS, Full Information Maximum Likelihood (FIML; normality, constant moments), and Generalized Method of Moments (GMM; 2×10 instruments). The *p* values are two-sided and based on the Lagrange multiplier test. The two central values shown in the rightmost columns are, in Panel A, the mean and the median for κ_1, and in Panel B, the means of κ_2 for the ERM core currencies (BEF, NLG, DKK, FRF, CHF, the latter an informal member) and the non-ERM ones, respectively.

was easy to spot, there are statistical doubts about the significance, and it is hard to imagine a risk premium that always mirrors the AR(2) prediction.

3.2. Tests on "extreme" observations

One potential explanation of the poor results is that, because of transaction costs or risk premia, the forward premia are noisy estimators of the true conditional expectation, creating the standard errors-in-the-regressor bias towards zero. To attenuate this, Nissen (1997) and Huisman et al. (1998) look at subsamples where forward premia are large, hoping that for these observations the expected changes are unusually large, too.[12] If so, the better signal-to-noise ratio should increase the COF regression coefficients back towards unity.

We adopt the original Bilson (1981) formulation, a variant of Eq. (1):

$$E_{t-1}\left(\tilde{s}_t\right) = \left[\kappa_0 + \lambda_0 I(P)_{t-1}\right] + \left[\kappa_1 + \lambda_1 I(P)_{t-1}\right]\mathrm{FP}_{t-1}. \tag{2}$$

In Eq. (2), the indicator $I(P)_{t-1}$ equals unity if, on day $t-1$, the forward premium is "large", that is, belongs to the top P% of the ranked observations, with P being set equal to 40, 20, 10, or 5%; otherwise, $I(P)_{t-1}$ equals zero.[13]

Table 3 provides the crucial results. These are weakly in line with the hypothesis in the sense that 6–8 of the 10 coefficients are positive in the large-FP sample, against 4–6 in the modal-FP sample. Likewise, seven or eight of the estimated differences (κ_1) have the right sign, positive. But only one positive difference is significant (ITL), and in the 60/40 split the CHF goes the wrong way. Nor do we see larger differences the more unbalanced the split is. True, there is such a pattern in the mean κ_1, but this is entirely due to the bizarre outlier for the USD; the median does not conform to the hypothesis. Results for GARCH, not shown, are even less in favor. Thus, the simple errors-in-variables-like effect Nissen or Huisman et al. had found in their monthly data against the USD are not strongly present in our numbers.

4. A trading-rule test

If intra-ERM rates are to a non-trivial extent predictable and one-week interest rates do not pick this up, average returns from trading must be positive, at least before transaction costs. To discover whether the potential gains are large, before and after costs, we formally test a trading rule. Ideally, we would like to take transaction costs into account at every trade. While Datastream provides bid–ask rates, these are indicative quotes whose spread substantially overstates the spreads in any individual market maker's binding quotes and *a fortiori* overstates the difference between the market's best bid and ask. To still obtain an approximate answer to the question whether the returns exceed the two-way costs, we keep track of the number of trades and compute the average before-cost profit per trade, as described in the next section. This average profit per trade can then be judged against the limited information we have about transaction costs.

4.1. The trading rule

The purpose of this section is not to find the optimal trading rule, unlike many of the papers cited in the introductory section. Rather, we want to show that even the simple linear relation picked up in our UIP tests provides information with clear economic significance. The fact that more sophisticated rules would have done even better just reinforces that point.

[12] See Sercu and Vandebroek (2005) for the distributional characteristics of expectations v noise required to make the technique work.

[13] Huisman et al. work with panel data and a restriction of equal betas across currencies. Their high/low criterion is, accordingly, based on the ranked dispersions of the forward premia within each cross-section. We do not follow their approach as, in our sample, the cross-currency constraint is rejected.

Table 3
Tests of UIP on extreme versus modal FPs
$E_{t-1}(\tilde{s}_t) = [\kappa_0 + \lambda_0 I(P)_{t-1}] + [\kappa_1 + \lambda_1 I(P)_{t-1}]\mathrm{FP}_{t-1}$.

Coefficient	Regression coefficients for individual currencies										Central values	
	BEF	NLG	DKK	FRF	ITL	CHF	GBP	JPY	CAD	USD	Avg	Med
Split at the 60th percentile												
κ_0	−0.06	−0.13	0.14	−0.15	−3.64	7.04	0.18	−2.00	4.49	−1.95	0.65	0.44
κ_1	−0.19	0.29	0.09	1.00	5.47	−7.57	2.47	3.02	−5.33	3.77	−0.08	0.36
p (lm)	0.87	0.88	0.90	0.36	*0.01	*0.02	0.35	0.41	0.06	0.56	0.43	0.40
$\kappa_0+\kappa_1$	−0.25	0.17	0.24	0.85	1.83	−0.53	2.66	1.02	−0.84	1.83	0.57	0.80
Split at the 80th percentile												
κ_0	−0.69	0.00	0.15	−0.10	−2.01	1.10	1.74	−1.77	1.05	2.26	−0.06	0.27
κ_1	0.09	−0.51	0.07	1.52	4.03	−1.43	0.91	3.75	−1.74	2.69	0.74	1.03
p (lm)	0.93	0.73	0.91	0.14	*0.00	0.45	0.68	0.19	0.50	0.66	0.50	0.47
$\kappa_0+\kappa_1$	−0.59	−0.51	0.23	1.42	2.03	−0.33	2.65	1.98	−0.69	4.95	0.69	1.30
Split at the 90th percentile												
κ_0	−0.60	−0.74	0.11	0.05	−1.49	0.27	2.05	−0.08	0.62	0.09	0.02	0.10
κ_1	−0.02	1.95	0.37	2.88	3.23	−0.60	1.40	2.37	0.24	67.79	1.31	8.84
p (lm)	0.99	0.24	0.59	*0.01	*0.02	0.71	0.62	0.42	0.94	0.19	0.51	0.42
$\kappa_0+\kappa_1$	−0.61	1.21	0.48	2.93	1.74	−0.33	3.45	2.29	0.85	67.87	1.33	8.94
Split at the 95th percentile												
κ_0	0.16	−0.45	0.03	−0.03	−1.17	−1.84	1.83	0.15	0.70	1.04	−0.07	0.03
κ_1	−1.36	1.95	0.67	4.47	4.76	2.26	1.45	3.87	−0.78	76.97	1.92	10.62
p (lm)	0.26	0.32	0.35	*0.00	*0.01	0.27	0.73	0.31	0.82	0.43	0.34	0.36
$\kappa_0+\kappa_1$	−1.21	1.50	0.69	4.44	3.59	0.42	3.28	4.01	−0.08	78.01	1.85	10.65

The regressands (s_t) are weekly percentage changes in the exchange rate (against the DEM), the regressors are the beginning-of-period one-week forward premiums (FP_t). The intercept and slope coefficient are allowed to differ depending on whether the premium is extreme or modal, *i.e.* whether |FP| belongs to the *P*% highest of the ranked |FP|, where $P = 5, 10, 20, 40\%$. Estimation is by OLS; p values are from the Lagrange multiplier test.

Our test assumes daily trading. As we have mentioned, at the daily horizon European rates against the DEM tend to exhibit negative first- and second-order autocorrelation. The linear trading rule that tries to measure the potential gains from this phenomenon is implemented in calendar time[14] and works as follows: we take the DEM as the home currency. Using two years of past data, we estimate the first- and second-order partial correlation coefficients. Periodically these estimates are updated using the most recent 2-year sample. At any date $t-1$, then, our contrarian forecast for the next-day change will be a partial reversal of the recent changes:

$$\left(s_t - \widehat{\mathrm{FP}}_{t-1}\right) = \rho_{-1}(s_{t-1} - \mathrm{FP}_{t-2}) + \rho_{-2}(s_{t-2} - \mathrm{FP}_{t-3}). \tag{3}$$

Whenever the predicted net return against the DEM is positive and larger than a pre-set filter ϕ, we receive a "buy FX" signal; likewise we get a "sell FX" message whenever the predicted movement is more negative than $-\phi$.

We consider three (related) trading rules, the first of which is a "pooled" trading rule, that is, one where both buying and selling are possible. Under this rule, upon a "buy" signal we go long foreign exchange (FX) and short DEM; after the "sell" recommendation we switch to long DEM and short FX; and when there is neither a buy nor a sell signal we do nothing.[15] Formally, in the pooled trading rule we set $D_{\mathrm{b},t-1} = 1$ whenever there is a buy recommendation and 0 otherwise, and we likewise set $D_{\mathrm{s},t-1} = -1$ whenever there is a sell recommendation and 0 otherwise. Returns on long and short positions are cumulated separately, in accounts indicated by subscripts A (asset) and L (liability). In the

[14] See Bjerring et al. (1983) for the difference between event-time and calendar-time studies.

[15] It can be verified that, in the equations below, "doing nothing" is shown as holding DEM both long and short, so that the net return is zero. We need this seemingly Byzantine twist for the annualization, *infra*, where both the long and short sides need to be fully invested all the time.

equations below, the first two show the raw daily returns on the long and short accounts, respectively; note how, each day, the Ds activate either an FX or a DEM position. The last equation defines the net return (NR) as the return from the long account minus the return from the short account:

$$R_{A,t} = D_{b,t-1}(s_t + r^*_{t-1} + s_t r^*_{t-1}) + (1 - D_{b,t-1})r_{t-1}, \quad \text{where } D_b = \{0,1\}, \tag{4}$$

$$R_{L,t} = (1 + D_{s,t-1})r_{t-1} - D_{s,t-1}(s_t + r^*_{t-1} + s_t r^*_{t-1}), \quad \text{where } D_s = \{0,-1\}, \tag{5}$$

$$\text{NR}_t = R_{A,t} - R_{L,t} = (D_{b,t-1} + D_{s,t-1})\left[(s_t + r^*_{t-1} + s_t r^*_{t-1}) - r_{t-1}\right], \tag{6}$$

where r_{t-1} (r^*_{t-1}) is the one-day return on a DEM (FX) money-market investment and between dates $t-1$ and t and s_t the percentage exchange-rate change over the same period. We recognise $[(s_t + r^*_{t-1} + s_t r^*_{t-1}) - r_{t-1}]$ as the return, measured in DEM, on a one-day FX investment, in excess of the DEM return; multiplying by $(D_{b,t-1} + D_{s,t-1})$, this excess return is "weighted" by $+1$ in case of a buy signal, -1 in case of a sell, and a zero in case of no signal.

We also want to test whether the gains are symmetric. For this purpose we test a separate "buy" strategy, one where $D_{s,t-1}$ always remains zero (that is, upon bad news we liquidate a long FX position but we never actually shortsell FX); and we likewise implement a "sell" strategy where $D_{b,t-1}$ is never switched to unity: upon a buy signal we stop shortselling FX (if we were short at all) but we never actually go long FX. These one-directional results are relevant to a liquidity trader, who is typically interested in a transaction of a particular type. But these results can also provide insights related to the forward puzzle, and more specifically whether peso-type problems connected with possible realignments are likely to have been a major factor or not.[16] Given the reputation of the home currency, the DEM, for ERM members only one type of parity change is conceivable: a devaluation of FX. Peso-type realignment risk, if any, should show up especially after a "buy FX" signal (which follows a drop of FX), and less after a "sell FX" signal because then the FX has picked up, recently. In short, a positive difference between the average "buy" and "sell" returns would be consistent with an extra peso-type realignment risk for ERM members.

4.2. Economic and statistical significance measures

For the purpose of interpreting the average daily returns for various filters and currencies, it is useful to correct for two obvious sources of differences in risk across currencies. First, the average risk premium may differ between, say, ITL and NLG. Second, there can be substantial differences as to the times spent long and short each currency. We can correct for these by considering a static control strategy. In this control strategy we take into account how much of the time the contrarian trader is long or short by simply adopting a static strategy with the same average exposure and, therefore, with the same risk premium, provided the latter is constant. Specifically, for each currency's time series of signals, $D_b + D_s$, we compute the average and then define a static control strategy where we hold a constant position equal to $\overline{D}_b + \overline{D}_s$, instead of the time-varying one in Eq. (6). If our signals are useless, in the sense of being uncorrelated with subsequent returns, then our trading strategy has the same expected return as the control strategy. Formally, define the excess net return (XNR) as the net return in excess of the net return on the control strategy:

$$\text{XNR}_t = \left[(D_{b,t-1} + D_{s,t-1}) - (\overline{D}_b + \overline{D}_s)\right]\left[(s_t + r^*_{t-1} + s_t r^*_{t-1}) - r_{t-1}\right]. \tag{7}$$

The expected value is zero when the signal is uncorrelated with the subsequent net return:

$$E(\text{XNR}) = \text{cov}[(D_b + D_s), (s + r^* + sr^* - r)]. \tag{8}$$

[16] A peso risk, in the original, narrow sense, is a huge potential event with a very small probability. The event is so huge that it affects the true expectation, but the probability is so small that the researcher often does not observe the event and, therefore, mis-estimates the expectation. More generally, peso risk refers to deviations between *ex ante* and *ex post* distributions.

This also means that the significance of the mean excess return can be tested in the same way as the significance of a covariance or a simple regression coefficient. Specifically, one regresses the net return $(s + r^* + sr^* - r)$ on the trading signal $(D_b + D_s)$. This test can be done separately for the pooled, buy, and sell rules. We use OLS with Newey–West's corrected *t*-statistics.

The above tests bear on mean excess returns per day. To judge the economic significance we complement this in two ways. First, to see whether returns are sufficient to recover transaction costs, we compute average returns per run of signals of the same sign. Second, to assess the likelihood of competing explanations of the forward puzzle, we cumulate the returns over long periods and compute *p.a.* returns. The details are as follows.

In the first variant we compute the average return per run of consecutive signals *D* of the same sign. Thus, in computing this average, we assume that the trader does not automatically close out at the end of each day, but waits until the signal either becomes zero or changes sign. Within a given run of similar signals, the returns on the long and short sides are cumulated separately[17] and the net value is computed at the end of the run. To the speculator, this average return per run is more relevant than the average return per buy or sell day since the speculator probably can hold positions longer than one day and incurs the costs only once.

In the second variant we cumulate the returns over long periods and compute *p.a.* returns. Again, returns on long positions are cumulated separately from returns from short positions, and we compute the final net value as

$$\text{net final value} = \prod_{t=t_1}^{t_n}(1+R_{A,t}) - \prod_{t=t_1}^{t_n}(1+R_{L,t}). \quad (9)$$

For each strategy, the *p.a.* average net return is obtained by computing average returns for each of the legs separately:

$$\text{average p.a. net return} = \sqrt[N]{\prod_{t=t_1}^{t_n}(1+R_{A,t})} - \sqrt[N]{\prod_{t=t_1}^{t_n}(1+R_{L,t})}, \quad (10)$$

where N is the time between starting date t_1 and end date t_n measured in years rather than trading days. The *p.a.* excess return then is computed as the above net return minus the analogous net return on the static control strategy.

4.3. Risk-free investments

As mentioned, in Datastream the one-day spot rates are unavailable for most of the period. We use returns or capital gains from holding a 7-day CD for one day,

$$r_t = \frac{1/[1+i_{t+1,t+6}(6/\text{ND})]}{1/[1+i_{t,t+7}(7/\text{ND})]} - 1, \quad (11)$$

where $i_{t,t+7}$ is a *p.a.* simple interest rate for a 7-day investment (LIBOR, in practice); and following interbank conventions: ND equals 365 for BEF or GBP, and 360 for other currencies. The numerator and denominator show the values, on days $t+1$ and t, of a CD that has 7 days to go on day t. To implement this, we need to assume that $i_{t+1,t+6} = i_{t+1,t+7}$, which, traders confirm, is acceptable. True, there may have been a Hicksian maturity risk premium associated with holding seven-day paper instead of one-day; but recall that the control strategy would also benefit from any such premium so that it would be largely neutralized in the excess net return.

While we would have been happier with perfect data, it is reassuring that the results turn out to be robust to this approximation. In an exploratory run, for instance, we simply took Datastream's

[17] Directly compounding the net returns, or *a fortiori* excess net returns, would be hard to interpret since there is no implementable strategy that provides this cumulative payoff.

"representative short-term" interest data to be one-day spot rates (even though, in reality, they may be three-month rates) and found essentially the same results. Indeed, the picture hardly changes even if we set all rates equal to zero. (These results are available on request.) The reason for the near-irrelevance of interest rates is their independence of exchange-rate changes, plus the fact that, via the control strategy, any maturity-related bias largely disappears.

4.4. Empirical results

Table 4, Panel A shows total-period *p.a.* net excess returns and *t*-tests on the average daily net excess return, for various values of the filter. We also provide means for ERM and non-ERM currencies. There is a classification problem with the CHF (not a member, but widely viewed as *de facto* pegged) and the ITL (a member, but with a much wider band). Somewhat arbitrarily, we put the CHF into the ERM-group and the ITL not. This is conservative in the sense that it blurs the differences between the subsamples; it also avoids the hindsight-selection bias that would have arisen from grouping just the four strictly-ERM currencies that, in the end, all turned out to stay in the system or even merge with the DEM.

We note that across all filters the returns are unanimously positive for the continental currencies, while they are patchy and algebraically small for the other currencies. The returns are systematically significant for core-ERM currencies, most of the time also for the CHF and ITL, but never significant for the GBP and the control group in the strictest sense (CAD, JPY, USD). It clearly pays to decrease the filter from 10 bp to 5 bp or even 1 bp, but further refinements no longer add anything substantial. The trading rule produces rather impressive excess returns: the ERM-group average for the 1 bp rule is 14%, and individual-currency results range between 8 and 19% (between 300 and 650% cumulative over eleven years,[18] that is). But even the widest filter, 10 bp, still pays out a respectable excess average return of 8%.

To verify the intertemporal stability and internal validity of the results, and to further reduce potential problems of heteroskedasticity, we run all tests also on subperiods, using again the tranquillity-versus-turbulence criterion. To preserve space we just show these subperiod results for the worst-performing filter, 10 bp. The summary is in Panel B of Table 4. Despite the much smaller sample sizes, only one cell loses its significance rating (CHF in the first period). The numbers remain in remarkable agreement with the total-sample results: almost unanimously clear and positive returns per subperiod for the core-ERM currencies, less impressive but positive returns for the CHF and ITL, and essentially random results for the GBP and the control currencies JPY, CAD, and USD. For the other filters the results are generally stronger but otherwise show the same patterns across currencies and periods.

In Panel C of the table we check whether the profits stem mostly from the long positions, or the short ones. Results for the 0.5- and 0.1-bp filters, being too similar to the 1-bp results, are omitted. It turns out that the returns from the buy-only and sell-only rules are quite comparable. Each is significant in its own for the core-ERM members, often so for ITL and CHF, and never so for the outside currencies. Judging by *t*-tests, the difference between buy and sell is never significant at the 1-% level except in one case (out of 18 – *i.e.* three trading rules $\times$ six continental currencies).[19]

From these buy- or sell-only results we conclude that, for liquidity traders whose transaction costs are irrelevant, there would have been substantial gains from using the trading rule. Of course, being cumulative returns before transaction costs, the net excess payoffs reported in Table 4C are representative only for liquidity traders that have to do a transaction every day (and of a similar size every day). To provide an idea of the profitability in a more general situation, Table 5 reports the mean profit per run and per signal for various filters, and for the best- and worst-performing filters also the subperiod results. We see that the large filters, which produced the lowest *p.a.* net excess returns, actually provide the highest profits per signal; thus, in terms of cumulative returns the problem with the large filters obviously is that there are not enough of these trading opportunities. The issue is whether these average returns per trade exceed the transaction costs is reserved for the concluding section.

[18] We have 13 years of data but lose two in estimating the ARIMA model.

[19] True, another case is almost significant (with $t = 2.07$), but both the significant and the almost-significant *t*s occur for the same currency, the NLG; and if we want to cite this as evidence of peso-type problems connected with realignments, then the NLG, commonly regarded as a DEM clone, is surely the least convincing candidate.

Table 4
Excess *p.a.* returns for all filters

		Coefficients for individual currencies										Averages	
		BEF	NLG	DKK	FRF	CHF	ITL	GBP	JPY	CAD	USD	ERM	Other
Panel A: Pooled (buy and sell), total period (1985/6–1987/3)													
ϕ (bp)	10 bp	9.44	10.4	8.42	9.94	2.19	1.3	−0.18	3.49	0.76	−0.92	8.08	0.89
	t-Test	*12.61	*11.03	*10.23	*16.81	*3.33	0.97	−0.16	1.12	0.64	−1.17		
	5 bp	15.1	13.35	10.12	13.29	3.51	4.41	0.9	4.13	0.1	1.47	11.07	2.20
	t-Test	*16.54	*12.55	*10.08	*17.45	*3.93	*2.64	0.50	0.98	0.08	1.04		
	1 bp	19.35	15.11	11.9	15.62	8.72	6.46	0.28	2.14	1.17	3.83	14.14	2.78
	t-Test	*16.71	*12.19	*9.35	*16.05	*6.78	*3.20	0.09	0.43	0.42	1.60		
	0.5 bp	19.20	15.61	12.44	15.78	8.48	5.74	−1.05	3.63	1.34	4.11	14.30	2.75
	t-Test	*16.54	*12.55	*10.08	*17.45	*3.93	*2.64	0.50	0.98	0.08	1.04		
	0.1 bp	19.26	15.28	12.49	15.75	9.48	5.14	−2.9	5.11	1.73	4.12	14.45	2.64
	t-Test	*15.82	*11.38	*9.21	*15.23	*6.50	*2.38	−0.90	0.92	0.56	1.46		
Panel B: Pooled (buy and sell), three subperiods, $\phi = 10$bp													
smpl	1985–	8.84	11.42	4.09	6.63	0.56	0.02	−0.15	5.19	0.72	−0.10	6.31	1.14
	t-Test	*5.15	*7.92	*5.85	*11.60	0.69	0.17	−0.12	1.30	0.50	−0.29		
	1992–	15.65	28.09	18.20	23.52	1.36	9.83	6.61	18.04	5.84	3.85	17.36	8.83
	t-Test	*8.58	*8.62	*8.25	*10.86	*4.42	*6.06	1.09	1.25	1.06	1.24		
	1994–	8.31	5.71	8.49	8.54	3.27	0.01	−1.80	−1.05	−0.41	−2.49	6.86	−1.15
	t-Test	*10.01	*4.69	* 6.50	*10.05	*3.16	0.04	−1.41	−0.26	−0.27	−2.28		
Panel C: Total period, buy versus sell, various filters													
rule	10buy	4.68	6.07	4.52	5.00	1.28	0.68	−0.56	2.23	−0.62	−0.77	4.31	0.19
	t-Test	*7.85	*8.14	*7.65	*13.32	*4.24	0.65	−0.66	1.24	−0.71	−1.60		
	10sell	4.76	4.32	3.89	4.94	0.91	0.61	0.39	1.26	1.38	−0.15	3.76	0.70
	t-Test	*11.12	*10.94	*8.51	*13.83	1.72	1.14	0.39	0.40	1.64	−0.47		
	t_{b-s}	−0.11	2.07	0.84	0.12	0.61	0.06	−0.72	0.27	−1.65	−1.07		
	5buy	8.13	7.74	5.41	6.85	2.08	2.83	−0.17	3.44	−1.35	0.95	6.04	1.14
	t-Test	*12.61	*10.00	*8.26	*14.62	*4.03	2.25	−0.08	1.51	−0.93	1.17		
	5sell	6.97	5.61	4.71	6.44	1.43	1.58	1.07	0.69	1.44	0.52	5.03	1.06
	t-Test	*13.58	*11.64	*9.02	*15.36	*2.39	2.23	0.85	0.10	1.09	0.45		
	t_{b-s}	1.29	0.91	0.98	1.36	0.73	1.08	−0.04	−0.01	0.19	0.28		
	1buy	10.23	7.99	6.43	8.30	4.73	4.09	−0.15	−0.59	0.86	2.19	7.54	1.28
	t-Test	*16.04	*10.79	*8.73	*15.47	*6.77	*3.24	0.01	0.00	0.64	1.69		
	1sell	9.12	7.12	5.48	7.31	3.99	2.37	0.44	2.74	0.31	1.64	6.60	1.50
	t-Test	*15.86	*11.74	*8.78	*14.81	*5.50	*2.45	0.16	0.82	0.12	1.10		
	t_{b-s}	1.41	*2.34	0.84	0.65	0.82	0.87	−0.50	0.38	−1.42	0.30		

We use daily data, June 14, 1985 to April 1, 1998 with the DEM as reference currency. A buy signal ($D = 1$) occurs whenever the exchange-rate change, as predicted by the AR(2) equation, exceeds the filter size, and likewise for the sell signal ($D = -1$). On $D = 1$ (-1) we go long FX (DEM) and short DEM (FX), while on $D = 0$ both the short and long side are DEM. Returns are from holding a 7-day DEM or FX risk-free investment for one day. In the control strategy, we follow a static rule of always investing $\overline{D}$ units (where $\overline{D}$ is the average position for that currency over the entire sample period). The *t*-tests are on the mean net return (long minus short) per day in excess of the net return from the trading strategy, and an asterisk indicates significance at 1% one-sided. The daily returns themselves have been compounded and annualized. The two average returns in the rightmost columns are computed over the first and second sets of five currencies, respectively. Subperiods are (i) is until September 1, 1992; (ii) September 1992 till end 1993; (iii) as of 1994. In the third part of the table, signals $D = -1$ are set to $D = 0$ in the buy-only rule, and signals $D = 1$ are set to $D = 0$ in the sell-only rule. The rule "10buy" refers to a buy-only game with filter 10, and so on. The *t*-test on the difference between the buy and sell returns is indicated as t_{b-s}.

Table 5
Average Abnormal Returns per run and per signal, subperiod results for the largest and smallest filters

Filter	Strategy	BEF	DKK	FRF	NLG	CHF	ITL	GBP	USD	JPY	CAD
Panel A: average profit per run, %											
0.100	Buy	0.20	0.24	0.23	0.24	0.19	0.08	0.06	0.04	0.13	0.08
	Sell	0.23	0.21	0.22	0.22	0.10	0.07	0.01	0.07	0.04	0.12
0.050	Buy	0.17	0.19	0.15	0.17	0.12	0.09	0.03	0.04	0.07	0.07
	Sell	0.18	0.18	0.16	0.16	0.08	0.10	0.03	0.01	0.05	0.03
0.010	Buy	0.14	0.15	0.11	0.12	0.08	0.08	0.02	0.01	0.04	0.02
	Sell	0.14	0.14	0.11	0.12	0.06	0.05	0.02	0.02	0.04	0.01
0.005	Buy	0.13	0.14	0.10	0.12	0.07	0.09	0.02	0.02	0.03	0.01
	Sell	0.14	0.14	0.10	0.12	0.05	0.05	0.02	0.00	0.03	0.01
Panel B: average profit per signal, %											
0.100	Buy	0.16	0.19	0.17	0.19	0.17	0.07	0.05	0.03	0.12	0.07
	Sell	0.18	0.17	0.17	0.18	0.08	0.06	0.01	0.07	0.04	0.11
0.050	Buy	0.12	0.13	0.10	0.12	0.09	0.06	0.02	0.04	0.06	0.06
	Sell	0.14	0.13	0.12	0.12	0.06	0.09	0.03	0.01	0.05	0.03
0.010	Buy	0.08	0.08	0.06	0.07	0.04	0.05	0.02	0.01	0.03	0.01
	Sell	0.09	0.08	0.06	0.07	0.03	0.03	0.01	0.02	0.02	0.00
0.005	Buy	0.07	0.07	0.06	0.07	0.04	0.05	0.01	0.02	0.02	0.00
	Sell	0.08	0.07	0.06	0.07	0.03	0.03	0.01	0.00	0.02	0.01

There is a buy signal ($D=1$) whenever the absolute value of the exchange-rate change, as predicted by AR(2), exceeds the filter; and there is a sell signal ($D=1$) when the predicted return is below minus the filter. The trade, if any, is a one-working-day forward sale or purchase depending on the sign of the predicted change. The table shows average excess returns (AERs) per trade and per run, for strategies where both selling and buying is allowed ("pooled"), as well as for strategies with either just buying or just selling. The AERs are averages of gross returns over the entire period (June 14, 1985–April 1, 1998), in excess of the return from a static trading strategy consisting of always buying forward $\overline{D}$ units (where $\overline{D}$ is the average position over the entire sample). For the AER per signal we assume that the trader liquidates every day. For the AER per run we assume that when a signal D is followed by another one of the same signs, then the trader rolls over her position until the run of identical Ds is finished.

Prior to that, we provide a last performance measure, the Sharpe ratio for a portfolio of trading strategies.[20] We provide three versions, one using simple mean–variance optimisation in combining the five currencies' trading rules, and two that follow a naive weighting scheme. In the first go, we start from the time series of net returns, as defined in Eq. (6), on each of the pooled time series for the five "ERM" currencies; we compute the variance–covariance matrix Ω and the vector E of five mean excess returns, and identify the mean–variance optimal weights $w_{5\times 1} = \frac{\Omega^{-1}E}{J'\Omega^{-1}E}$ for a fully invested portfolio. (J denotes the 5×1 unit vector.) The Sharpe ratio follows:

$$\text{Sharpe ratio} = \frac{\overline{\text{NR}}_p}{\text{std(NR)}}, \quad \text{where } \text{NR}_p = \sum_{j=1}^{5} w_j \text{NR}_j. \tag{12}$$

Table 6 provides these Sharpe ratios for each of the three filters (10, 5, and 1 bp), along with the weights for each of the five currencies. The weights are all below 0.5 and (up to one minor exception) positive; that is, our ratios are not driven by extremely large positive and negative positions, a feature that plagues many applications of the Sharpe criterion. The NLG is the most desirable currency, it seems; the BEF, despite its higher forward premia, is too risky to come first, but still beats the FRF and DKK.

[20] We followed up on this suggestion because of the ratio's popularity among practitioners. Theoreticians may, correctly, object that it should be applied just to total-wealth returns, not to a small part of one's wealth. Also, our portfolio returns are far from normal, with an especially anomalous probability peak at the zero return as the main source of non-normality, and right-skewness as the second one.

Table 6
Sharpe ratios and *ex post* μ–σ-optimal weights

Filter	10 bp	5 bp	1 bp
Mean–variance optimal portfolios			
w_{BEF}	0.27	0.31	0.28
w_{DKK}	0.14	0.12	0.11
w_{FRF}	0.13	0.10	0.09
w_{NLG}	0.46	0.49	0.49
w_{CHF}	0.00	−0.02	0.04
μ–σ optimized Sharpe ratio	0.35	0.43	0.46
Equal-investment portfolios			
Naive Sharpe ratio 1	0.32	0.37	0.40
Naive Sharpe ratio 2	0.30	0.39	0.41

For each filter size we start from the time series of net returns, as defined in Eq. (6), on each of the pooled time series for the five "ERM" currencies. From these five series we compute the variance–covariance matrix and the vector of five mean excess returns, from which we identify the mean–variance optimal weights for a fully invested portfolio. This provides a mean excess return, a standard deviation, and a Sharpe ratio. Under the first alternative scheme, on every day the trader invests equal amounts (say, DEM 1) in each "active" currency, that is, the weights are $w_{i,t} = D_{\text{b},i,t} + D_{\text{s},i,t}$; a position in the DEM money market is then added so that the total net investment is DEM 1. This provides an excess portfolio return for every date, from which the mean, the standard deviation, and the Sharpe ratio follow. In a second implementation the weights are set as $w_{i,t} = [D_{\text{b},i,t} + D_{\text{s},i,t}]/\sqrt{[\sum_{j=1}^{5}(D_{\text{b},i,t} + D_{\text{s},i,t})^2]}$ if at least one position is non-zero, and $w_{i,t} = 0$ otherwise; the DEM money-market account then again serves to bring the sum of the weights up to unity.

In-sample computations like these suffer from data snooping. To reduce this problem, we also offer two sets of ratios based on naive equal-weighting schemes.[21] From Table 6 we see that the alternative new return/risk ratios do not drop very far from the level obtained after *ex post* mean–variance optimisation.

5. Discussion and conclusion

Froot and Thaler (1990) advance three possible explanations of the forward bias: a regular risk premium, peso risk, and learning. We review these in light of the new evidence, along with a fourth view, the transaction cost avenue proposed by *e.g.* Baldwin (1990).

Lewis (1989) points out that if the change in the Fed's monetary policy was not initially recognized by the market, this would have produced forecast errors that are negatively correlated with the forward premium. But, in general, the forward bias is too widespread over time and across currencies to be explicable by episodes like this one. In the case of the ERM, more specifically, there was no pronounced shift in monetary policy. Also, the system was well known by the time our data start, so there was little to learn.

The regular risk premium is a covariance with a stochastic discount factor which, in turn, is usually specified as the return on a portfolio or a combination of portfolios. We first focus on the risk premium as specified by the standard CAPM or its international version (InCAPM). In these models the portfolios are, in case of the CAPM, the market portfolio and, for the InCAPM, a combination of the world market and the various countries' T-bills. In the CAPM, the spot rate's beta can never be large because the R^2 of an exchange-rate market-model regression is invariably tiny[22] and the regressor standard deviation larger than the regressand one:

$$E(s_i + r_i^* - r) = \beta_i E(r_\text{m} - r) \quad \text{with}\; \beta_i = \rho_{i,\text{m}} \frac{\sigma(s)}{\sigma(r_\text{m})} \ll 1,$$

[21] In a first implementation, every day the trader invests the same total amount (say DEM 1), and allocates also equal amounts (DEM 1) in each currency. That is, the risky weights are $w_{i,t} = D_{\text{b},i,t} + D_{\text{s},i,t}$, and a position in the DEM money market is then added so that the net investment is DEM 1. This generates a realized excess portfolio return for every date, from which the mean, the standard deviation, and the Sharpe ratio follow. In this version, the total risk of the portfolio changes drastically over time, as the number of active positions can vary from zero to five. To test the robustness to this effect, we consider a second implementation, where the weights are set as $w_{i,t} = [D_{\text{b},i,t} - D_{\text{s},i,t}]/\sqrt{\sum_{j=1}^{5}(D_{\text{b},i,t} + D_{\text{s},i,t})^2}$ if at least one position is non-zero, and $w_{i,t} = 0$ otherwise; the DEM money-market account then again serves to bring the sum of the weights up to unity. The motivation is that this scheme provides a first-pass correction for the fluctuating number of active positions.

[22] See for instance, Allayannis (1996) for a review.

where r_m is the market return, ρ is the correlation coefficient, and σ indicates a standard deviation. Thus, with such a small beta, even if the risk premium on the market would be 10 or 20% we would never get anywhere near a 14-% average for an exchange rate.

Actually, we need to explain not just the 14% extra, but also the fact that 7.5% is obtained from being long and 6.5% from being short, with positions changing quite rapidly and lasting one or two days. Obviously, in the CAPM logic this requires a time-varying beta. It turns out that the required swings in beta are too large to be credible. For the 1-bp filter, for instance, one is long about 37% of the time, short about 29%, and out of the market 34%. (These are averages across the five ERM currencies.) Thus, to get 7.5% from being long or out of the market – the buy rule – one needs about a 20% annualized return on days one receives a +1 signal. Similarly, to get 6.6% from the sell rule in 29% of the time, the annualized extra return on negative-signal days must be over minus 22%. While betas can and do change over time, it is hard to imagine that the characteristics of two economies would change so fast, so drastically, and so often, that CAPM-equilibrium expected returns swing from +20 to −20% *p.a.* in a matter of days.

If the CAPM cannot explain an average excess return of 15%, might the InCAPM do so? In the two-country version an own-variance risk premium is added (Sercu, 1980):

$$E(s_i + r_i^* - r) = A_w \mathrm{cov}(s_i, r_w) + (1 - A_i)\frac{W_i}{W_w}\mathrm{var}(s_i),$$

where r_w is the world-market return; A_i (A_w) is relative risk aversion for the country-i (world) investor; and W_i (W_w) is country i's (the world's) invested wealth. But again it is implausible that returns of 10 or 20% *p.a.* would be explained by the new risk premium: the consensus is that relative risk aversion exceeds unity, making the additional risk premium negative. In addition, a variance cannot change sign, so in the InCAPM the variation in the last term should have come from the $1 - A$ term, requiring investors to swing quickly from strongly risk-preferent to highly risk-avert in a matter of days.

More generalized asset pricing model works with pricing kernels, which can specialise into marginal utilities, consumption growth rates or market returns depending on the additional assumptions. (See *e.g.* Brandt et al., 2006; Brennan and Xia, 2006; Dai and Singleton, 2002; Inci and Lu, 2004; Ahn, 2004). If, in such a setting, the market is complete, then the exchange rate must equal the ratio of the marginal utilities V_c abroad versus at home.[23] This implies a relation, shown below, between the exchange-rate change and the percentage changes in the marginal utilities. But the expectation percentage growth in marginal utility equals (minus) the risk-free rate, in each country. The final result is that the forward bias is the second-order term in $\mathrm{d}(V_{c^*}/V_c)/(V_{c^*}/V_c)$:

$$\begin{aligned} \frac{\mathrm{d}(V_c^*/V_c)}{V_c^*/V_c} &= \frac{\mathrm{d}V_c^*}{V_c^*} - \frac{\mathrm{d}V_c}{V_c} - \rho\sigma\sigma^*\,\mathrm{d}t + \sigma^2\,\mathrm{d}t; \\ E(\tfrac{\mathrm{d}S}{S}) &= r\,\mathrm{d}t - r^*\,\mathrm{d}t - \rho\sigma\sigma^*\,\mathrm{d}t + \sigma^2\,\mathrm{d}t; \\ \Rightarrow E\left(\frac{\mathrm{d}S}{S}\right) - (r\,\mathrm{d}t - r^*\,\mathrm{d}t) &= -\rho\sigma\sigma^*\,\mathrm{d}t + \sigma^2\,\mathrm{d}t, \end{aligned} \tag{13}$$

where σ (σ^*) is the *p.a.* standard deviation of the domestic (foreign) price kernel and ρ their correlation. All moments could be changing over time, so it is plausible that the difference between the two second moments flips from positive to negative. Still, it remains puzzling why the difference of the two moments would shift from +20 to −20% p.a. in a matter of days. There is a second contradiction: Brandt et al. (2006), inferring properties of kernels from stock markets, find that they do not fit exchange rates in the sense that exchange rates should be even more volatile than they already are, to be consistent with market completeness and stock-market behavior. As they

[23] This version assumes Von Neuman–Morgenstern investors; the final result does not strictly require this, however; the percentage changes in the marginal utilities can be replaced by kernels as they emerge from a no-arbitrage condition.

study floating rates, the anomaly should be even greater for controlled rates like the ERM data we study.

De Roon et al. (2000) advance the idea of hedging pressure. They extend the CAPM with non-marketable positions and obtain an expanded version similar to Mayers' (1976) model with human wealth:

$$E\left(s_i + r_i^* - r\right) = \beta_i E(r_{\mathrm{m}} - r) + A_{\mathrm{w}} \mathrm{cov}(e_i, r_{\mathrm{n}}),$$

where e_i is the residual of the regression of s on r_{m}, and r_{n} is the return on the aggregate position in non-tradeable assets $k = 1, \ldots, K$ held by the private investors, that is, $r_{\mathrm{n}} = \sum_{k=1}^{K} (q_k p_k / Y_{\mathrm{m}}) r_k$ with q_k (p_k) the beginning-of-period quantity (price) of non-tradeable asset k as held by all investors taken together, r_k the return on the non-tradable asset, and Y_{m} the beginning-of-period total invested marketable wealth. De Roon et al. stretch the Mayers model somewhat and look at positions of active professional traders in related currencies rather than positions of the aggregate private investor.[24] The attractive part of this interpretation is that professional traders' positions q_k can change sign over time indeed, and quite fast. But the covariances obtained by De Roon et al. are of the order of 0.10–0.25%.[25] Thus, to explain a 15-% return as a reward for bearing covariance with non-tradeables risk, the product of A_{w} and $q_k p_k / Y$ would have to be of order of magnitude of 60–150, which is impossible.

A last potential explanation in the literature is the peso effect. Peso problems in the wider sense are present in any empirical study involving expectations and second moments; for instance, we have not observed a collapse of the Maastricht agenda, even though this scenario must have been part of the *ex ante* distribution. Our discussion of peso-type risks here is restricted to the more narrow issue of realignments and especially devaluations. These are *a priori* not likely to explain the near-zero slopes for intra-European rates. For one thing, while there must have been more-than-occasional episodes of realignment fears in the sample period, such realignments have actually occurred fairly frequently. Also, they were small, *ex post*: Sercu and Uppal (1995, pp. 364–365) report an average devaluation jump size of 4%. Nor can one argue that individual *ex ante* jumps might have been much bigger: the understanding in the ERM was that a devaluation should just undo the cumulative change in the real rate since the previous realignment. Thus, the idea that 14% is the product of a small probability times a huge jump size is utterly implausible.

A second way to refute peso issue about realignments explanation is based on the observation that our COF regression coefficients for non-European rates are at least as bad as the ERM ones. If this is to be explained by peso risk, the floating-rate data would have to be more exposed to that risk – but it is hard to conceive what type of catastrophic event could occur in the case of floating rates. Third, if peso risk is present, it should have shown up in our trading-rule results: going long FX after it has dropped should be the position with comparatively more peso risk, not going long DEM after a rise of FX. That is, the returns from buy-only should exceed those from sell-only. While this algebraically true, the observed differences are pairwise insignificant, and never amount to more than a fraction of the 14% we have to explain away.

Our work also responds to claims that the gains are small and risky, and wiped out by transaction costs. For a committed exporter or importer who is trading very frequently, judicious exercising of the option to postpone a purchase or a sale would increase revenue by several percentage points. Detailed subperiod results (available on request) show that the trading rules never produced negative payoffs over any two-year interval. To further illustrate this point, we compare our Sharpe ratios to what could be obtained in the stock market. Like ours, the benchmark ratio we use is fully in-sample. Specifically, we look at a comparable trading strategy, implemented in Ang and Bekaert (2004), who consider various national stock markets and try to determine, for each country's market, which of two regimes

[24] The most important non-tradable asset for individuals would be human wealth. Other important assets that are difficult to trade include appliances or furniture, etc. and, in many countries, real estate; housing is, in addition, often indivisible and therefore difficult to diversify. But it is again hard to see how the exchange rate would have strong (and erratically fluctuating) correlations with returns from human wealth, white goods, and real estate.

[25] The regression coefficients in their Table III, Panel D are or order of magnitude 1. To obtain the covariance one needs to undo a multiplication by 100 (see key to their table), scale up from a bi-weekly to a *p.a.* basis (*i.e.* multiply by 25) and multiply by var(s), a number below 0.01 *p.a.*

holds – bull or bear. His data period is the rather exceptional 1990s boom. Depending on the estimation method and the way of mean–variance optimisation (which takes place each period, not marginally as in our experiment), his Sharpe ratios range from 0.07 to 1.07 per annum. Thus, the best ratio translates into $1.07/\sqrt{255} = 0.067$ *per diem*. For our trading rule, the ratios ranged between 0.30 and 0.46 *per diem*, 5–6 times better than the best ratio obtained in an exceptional stock-market decade.

Let us, finally, address transaction costs. An exchange rate is quoted in at least four digits (think of a below-par EUR/USD) and often in five digits (like the USD/DEM, of old). In the mid-1990s, binding quotes for USD/DEM were about three ticks apart, and the market spread between best bid and best ask was occasionally as low as two or one tick (Lyons, 1998; see Rosenstreich, 2005; Weithers, 2006). Assuming, conservatively, a market spread of three ticks and a DEM-per-USD rate of 15,000 ticks, the two-way cost was at most 2 bp. Of course, our data include lower-volume currencies, too, but 6 bp is about the maximum for the rates we look at. Interest spreads should be added, but they are comparatively tiny. Even a spread of 25 basis points per annum, which for a long time was the norm, translates into an implicit commission of just 25/2/360 = 0.035 basis points (not percentages!) for a one-day spot deposit.

From Table 5 it is unclear whether the small filters would have been profitable to round-trip speculators, but the larger filters, which pay off 10–25% per trade, have been profitable even after transaction costs. Thus, one cannot claim that transaction costs would have wiped out the gains we have observed.

Is the only remaining explanation, then, a market inefficiency, especially in light of the fact that gains exceed the transaction costs? Baldwin (1990) argues that it is not. Specifically, he points out, the observed return differentials should clearly exceed the costs before the arbitrageurs would rationally move in. True, the holder of DEM has the option to shift her funds into FX as soon as the expected total return on FX exceeds the return on DEM. However, she will normally wait until that option is sufficiently far in the money. The first reason is one that holds for ordinary options, too: the holder of an American-style option will not exercise as soon as it is in the money by a minute amount, but will wait until the payoff is big enough: immediate exercise would kill the (roughly) 50% probability that the exercise value increases over the next time interval, and the gains from moving deeper into the money outweigh the losses from moving back towards the money. The second reason for delayed exercise is that, if and when the investor has shifted her funds into FX, still sooner or later the return on DEM is bound to become more attractive again. But due to the transaction cost, she will then have to live with below-normal returns on FX until the loss is sufficiently big to warrant switching back to DEM. In short, to trigger trades, the return differential has to compensate for the transaction cost of moving into and out of FX, plus the possible regret from moving in too soon, plus the likely cost of reaping below-normal returns before moving out again. Whether this type of reasoning is enough to explain the non-disappearance of trading profits is far from certain: Baldwin's thoughts are far from an operational test or trading rule. Still, the jury is still out on whether or not transaction costs justify inaction to the extent we have witnessed in our data.

We conclude with a brief wrap-up. One issue that hampers tests of UIP is that the market's expectations are unobservable. In this paper we are able to identify clear non-zero expectations from the exchange-rate data themselves, for the simple reason that the ERM (or perceived dirty floating) induces statistically significant mean-reversion in daily and weekly exchange rates against the DEM. UIP hypothesizes that this predictability should be picked up by the one-week forward premium, but when running the Cumby and Obstfeld (1984), Fama (1984) regression tests of the unbiased-expectations hypothesis at the one-week horizon, we obtain essentially zero coefficients for intra-ERM exchange rates (and the familiar negative coefficients for extra-ERM rates). These slopes are also quite unstable across estimation methods and periods. Lastly, lagged exchange-rate changes remain significant when added to the regression, that is, forward premia seem to essentially ignore this source of predictability in exchange rates. Especially this last finding seems damaging to UIP: the mean-reversion is quite simple to detect, the lagged change is not at all a near-stationary regressor, and the idea that, over time, the risk premium would perfectly mirror the lagged changes looks far-fetched. Nor can one maintain that the phenomenon is economically trivial. We indeed test a trading rule, and find that for a sufficiently large filter the average profit per trade is larger than transaction costs. The size of the profits and the patterns from buy versus sell decisions allow us to reject the regular risk premium and the peso hypotheses too, at least as being sufficient in themselves to explain the results. To figure out whether all this implies a market inefficiency, however, we need more insight into the effect of transaction costs on optimal trading.

References

Ahn, H.-H., 2004. Common factors and local factors: implications for term structures and exchange rates. Journal of Financial and Quantitative Analysis 30 (1), 69–102.

Allayannis, G., 1996. Exchange Rate Exposure Revisited. Working Paper, Darden GSB, University of Virginia.

Ang, A., Bekaert, G., 2004. How do regimes affect asset allocations. Financial Analysts Journal 60, 86–99.

Baldwin, E.R., 1990. Reinterpreting the Failure of the Foreign Exchange Market Efficiency Tests: Small Transaction Costs, Big Hysteresis Bands. CEPR Discussion Papers, 407 and NBER Working Paper, 3319.

Bansal, R., 1997. An exploration of the forward premium puzzle in currency markets. Review of Financial Studies 10, 369–403.

Bansal, R., Dahlquist, M., 2000. The forward premium puzzle: different tales from developed and emerging economies. Journal of International Economics 5, 115–144.

Bekaert, G., Hodrick, R.J., 2001. Expectations hypotheses tests. The Journal of Finance 56, 1357–1394.

Brennan, M., Xia, Y., 2006. International capital markets and foreign exchange risk. Review of Financial Studies 19, 753–795.

Bilson, J., 1981. The speculative efficiency hypothesis. Journal of Business 54, 435–451.

Bjerring, J.H., Lakonishok, J., Vermaelen, T., 1983. Stock prices and financial analysts' recommendations. Journal of Finance 38, 187–204.

Brandt, M.W., Cochrane, J.H., Santa-Clara, P., 2006. International risk sharing is better than you think, or exchange rates are too smooth. Journal of Monetary Economics 53, 671–698.

Cumby, R., Obstfeld, M., 1984. International interest rate and price level linkages under flexible exchange rates: a review of recent evidence. In: Bilson, J., Marston, R. (Eds.), Exchange Rate Theory and Practice. University of Chicago Press, Chicago, IL.

Dai, Q., Singleton, K.J., 2002. Expectation puzzles, time-varying risk premia, and affine models of the term structure. Journal of Financial Economics 63 (March), 415–441.

De Grauwe, P., 1989. On the Nature of Risk in the Foreign Exchange Markets: Evidence from the Dollar- and the EMS-markets. CEPR Discussion Paper Series 352.

De Roon, F., Nijman, T., Veld, J., 2000. Hedging pressure effects in futures markets. Journal of Finance 55 (June), 1437–1456.

Fama, E., 1984. Forward and spot exchange rates. Journal of Monetary Economics 14, 319–338.

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

Green, P., 1992. Is currency trading profitable? Exploiting deviations from uncovered interest rate parity. Financial Analysts Journal 48, 82–86.

Hollifield, B., Uppal, R., 1997. An examination of uncovered interest parity in segmented international commodity markets. Journal of Finance 52, 2145–2170.

Huisman, R., Koedijk, K., Kool, C., Nissen, F., 1998. Extreme support for uncovered interest parity. Journal for International Money and Finance 17, 211–228.

Inci, A.C., Lu, B., 2004. Exchange rates and interest rates: can term structure models explain currency movements? Journal of Economic Dynamics and Control 28, 1595–1624.

Kho, V., 1996. Time varying risk premia, volatility, and technical trading profits. Journal of Financial Economics 41, 249–290.

LeBaron, B., 1998. A Dynamic Trading Strategy Approach to Deviations from Uncovered Interest Parity. Working Paper, University of Wisconsin at Madison.

Levich, R.M., Thomas, L.R., 1993. The significance of technical trading-rule profits on the Forex market. Journal of International Money and Finance 12, 451–474.

Lewis, K.K., 1989. Changing beliefs and systematic rational forecast errors with evidence from foreign exchange. American Economic Review 79, 621–636.

Lothian, J.R., Wu, L., 2005. Uncovered Interest Rate Parity over the Past Two Centuries. Graduate School of Business, Fordham University.

Lyons, R., 1998. Profits and position control: a week of FX dealing. Journal of International Money and Finance 17, 97–115.

Mayers, D., 1976. Nonmarketable assets, market segmentation, and the level of asset prices. Journal of Financial and Quantitative Analysis 11, 1–12.

Neely, C.J., Weller, P., 2003. Intraday technical trading in the foreign exchange market. Journal of International Money and Finance 22, 223–237.

Neely, C.J., Weller, P., Dittmar, R., 1997. Is technical trading profitable in foreign exchange markets? Journal of Financial and Quantitative Analysis 32, 405–426.

Nissen, F., 1997. International Financial Markets Dynamics, Unpublished Doctoral Dissertation, Maastricht University.

Okunev, J., White, D., 2003. Do momentum-based strategies still work in foreign currency markets? Journal of Financial and Quantitative Analysis 38, 425–448.

Raj, M., 2000. Transaction data tests of efficiency in the singapore futures markets. Journal of Futures Markets 20, 687–704.

Robinson, B., Warburton, P., February 1980. Managing currency holdings: lessons from the floating period. London Business School Economic Outlook, 18–27.

Roll, R., Yan, S., 2000. An explanation of the forward premium "Puzzle. European Financial Management 6, 121–148.

Rosenstreich, P., 2005. Forex Revolution: an Insider's Guide to the Real World of Foreign Exchange Trading. Prentice Hall Books. 304pp.

Schotman, P., Straetmans, S., De Vries, C., 1997. Big News in Small Samples. Tinbergen Institute Discussion Paper 97-083/2.

Sercu, P., 1980. A generalization of the international asset pricing model. Journal the L'Association Française de Finance 1, 91–135, Reprinted in Stulz, Ren M., Karolyi G. Andrew (Eds.), 2003. International Capital Markets, Volume 1. Edward Elgar Publishing Ltd., pp. 28–76.

Sercu, P., Uppal, R., 1995. International Financial Markets and the Firm. International Thompson, Cincinnati.
Sercu, P., Vandebroek, M., 2005. What UIP tests on extreme samples reveal about the missing variable. Journal of International Money and Finance 24, 1237–1260.
Surajaras, P., Sweeney, R.J., 1992. Profit-Making Speculation in Foreign Exchange Markets. Westview Press.
Sweeney, R.J., 1986. Beating the foreign exchange market. Journal of Finance 41, 163–182.
Taylor, S.J., 1994. Trading futures using a channel rule: a study of the predictive power of technical analysis with currency examples. Journal of Futures Markets 14, 215–235.
Villanueva, M., 2005. FX dynamics, limited participation, and the forward bias anomaly. The Financial Review 40, 67–93.
Weithers, T., 2006. Foreign Exchange: a Practical Guide to the FX Markets. Wiley Finance. 336pp.