

Breedon, Francis; Rime, Dagfinn; Vitale, Paolo

Working Paper

A Transaction Data Study of the Forward Bias Puzzle

Working Paper, No. 2010/26

Provided in Cooperation with:

Norges Bank, Oslo

Suggested Citation: Breedon, Francis; Rime, Dagfinn; Vitale, Paolo (2010) : A Transaction Data Study of the Forward Bias Puzzle, Working Paper, No. 2010/26, ISBN 978-82-7553-583-0, Norges Bank, Oslo, <https://hdl.handle.net/11250/2497435>

This Version is available at:

<https://hdl.handle.net/10419/209971>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/4.0/deed.no>

2010 | 26

Working Paper

Research Department

A transaction data study of the forward bias puzzle

Francis Breedon, Dagfinn Rime and Paolo Vitale

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:

servicesenter@norges-bank.no

Fra 1999 og fremover er publikasjonene tilgjengelig på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatterens regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail:

servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)

ISBN 978-82-7553-583-0 (online)

A Transaction Data Study of the Forward Bias Puzzle*

Francis Breedon

Queen Mary, University of London[†]

Dagfinn Rime

Norges Bank and NTNU[‡]

Paolo Vitale

University d'Annunzio[§]

November 2010

*Vitale is grateful to the Imperial College Business School, where part of this research was undertaken, for its hospitality. Financial support from the Risk Lab at Imperial College Business School is gratefully acknowledged. We also wish to thank Martin Evans, Steinar Holden, Hanno Lustig, Michael Moore, Carol Osler, Lucio Sarno, Adrien Verdelhan and participants at seminars at the University of Bologna, the University of Modena, the EIEF Research Center in Rome, Luiss University, Warwick Business School, Cass Business School, Norges Bank, the Norwegian University of Science and Technology, the University of Sassari, the 2009 AEA Meeting in San Francisco, and the 6th Annual Central Bank Workshop on Microstructure for helpful comments. Hong Xu and Filip Zikes provided excellent research assistance. The authors alone are responsible for the views expressed in the paper and for any errors that may remain.

[†]School of Economics and Finance, Queen Mary, University of London, Mile End Road, London E1 4NS (United Kingdom); telephone: ++44-(0)20-7882-8845; webpage: <http://www.econ.qmul.ac.uk>; e-mail: f.breedon@qmul.ac.uk

[‡]Norges Bank and Norwegian University of Science and Technology (NTNU). Bankplassen 2, PO box 1179 Sentrum N-0107, Oslo (Norway); telephone: ++47-2231-6757; webpage: <http://www.norges-bank.no/research/rime/>; e-mail: dagfinn.rime@norges-bank.no

[§]Faculty of Economics, Università d'Annunzio, Viale Pindaro 42, 65127 Pescara (Italy); telephone ++39-085-453-7647; webpage: <http://www.ch.unich.it/~vitale>; e-mail: p.vitale@unich.it

A Transaction Data Study of the Forward Bias Puzzle

ABSTRACT

Using ten years of FX transactions data we demonstrate that a large share of the FX forward discount bias can be accounted for by order flow. A simple microstructure-based decomposition suggests that order flow creates a time-varying risk premium that is correlated with the forward discount. The order flow related risk premium is particularly important in currency pairs traditionally associated with carry trade activity, as for these crosses it accounts for more than half of the forward bias (with the rest accounted for by systematic forecasting errors). We also find evidence that order flow is partly driven by carry trade activity, which is itself driven by expectations of carry trade profits. However, carry trading increases currency-crash risk in that the carry-induced order flow generates negative skewness in FX returns.

JEL Nos.: F31, G14 and G15.

Keywords: Forward Discount Puzzle, FX Microstructure, Carry Trade, Survey Data.

*Come l'araba Fenice,
che vi sia ciascun lo dice,
ove sia nessun lo sa^a
Metastasio, Demetrio*

^aLike the Arabian Phoenix,
everyone swears it exists,
but no one knows where

1 Introduction

The uncovered interest rate parity (UIP) condition states that, under risk neutrality, the gain from borrowing a low interest rate currency and investing in a higher interest rate one will, in equilibrium, be matched by an equally large expected cost in the form of a depreciation of the high interest rate currency. Combined with the hypothesis of rational expectations, this condition implies that the forward rate should be an unbiased estimator of the corresponding future spot rate. The empirical literature, Bilson (1981), Fama (1984), Froot and Frankel (1989) and Burnside, Eichenbaum, and Rebelo (2007, 2009) among (many) others, suggests that the forward rate unbiasedness (FRU) condition is systematically violated.¹ This is termed the forward discount bias and represents one of the longest standing puzzles in international finance. Despite the large range of alternative explanations put forward, there is no general consensus on the reasons why violations of FRU persist. Much like the whereabouts of the mythological Phoenix in Metastasio's citation, the forward discount bias arguably remains an unresolved puzzle.

Since this systematic bias means that high interest rate currencies tend to offer excess returns relative to low interest rate currencies, the forward discount puzzle results in a simple trading strategy widely used by practitioners that takes long positions in a high-yielding currency and short position in a currency offering a lower yield. Such a strategy is called the carry trade and has increasingly become the subject of popular debate and discussion. This debate has proved informative since it reveals how market participants themselves view the processes underlying the forward bias. For example, the following quote from the Wall Street Journal from May 2007 is indicative of market commentary on carry trading: "The carry trade has lifted currencies linked to high interest rates to their most overvalued level in 25 years,

¹See Lewis (1995) and Engel (1996) for surveys of research on this topic.

increasing the risk of a potentially damaging selloff, industry experts warn”². This quotation suggests that market participants (i) believe it is the very activity of carry trading that “lifts” high interest rate currencies; (ii) believe that carry trading raises high interest rate currencies to “overvalued” levels relative to low interest rate ones; and (iii) also realize that there is significant risk of a dramatic reversal connected to carry trading.

In this paper we explore all three of these themes. By combining data on FX order flow³ with information on market participants’ expectations of future currency values, we decompose the forward discount bias into two parts, one associated with time-varying risk premia as a function of order flow, the other with forecast errors. Furthermore, using actual transactions from the most liquid segment of the foreign exchange market, we study the importance and implications of carry trading.

Overall, in line with previous studies, we find that forecast errors seem to play a role in the forward bias. More interestingly, we also find an equally important role for an order flow related risk premium. Such a role is particularly pronounced for currency pairs typically associated with carry trading. For these currencies we find that the forward discount generates order flow consistent with carry trading. In addition, we see that carry trading is sustained by expectations of carry trade profits, but that the trade imbalance it induces brings about skewness in FX returns. This means that carry traders expect profits from their activity but that this trading also increases the risk of a sudden crash in exchange rates.

Our empirical approach combines the Reuters survey of market participants’ forecasts of future currency values and FX transactions data from Electronic Broking Services (EBS), for euro, yen and sterling against US dollars, over a period of ten years between January 1997 and April 2007. Although the main focus of this study is to combine these data sets, it is worth noting that individually they are arguably superior to most data sets previously used in the literature. For example, whereas Burnside et al. (2009) refer to indicative bid-ask quotes released by a large FX dealer, we have access to data on actual transactions completed on the main electronic trading platform which currently dominates spot FX markets for the major crosses. With respect to the work using survey data, e.g. Bacchetta, Mertens, and van Wincoop (2008), our survey of exchange rate forecasts, while shorter in length, focuses almost entirely on financial institutions and contains information on all individual forecasts rather than sample averages.

²From “Carry Trade Prompts Warnings” WSJ 18 May 2007.

³Order flow is the net buying pressure for foreign currency and is signed positive or negative according to if the initiating party in a transaction is buying or selling (Lyons, 2001).

This paper is organized as follows. In the next Section, we provide a brief literature review. Section 3 describes the data set on trade imbalance and survey forecasts and shows how the forward discount bias is large and significant and only partially due to forecast errors. Based on these preliminary results, Section 4 decomposes the forward bias into two components, one related to forecast errors, the other to a risk premium that is related to order flow. In Section 5 we investigate the role of carry trade activity in generating order flow in FX markets and how the carry trade increases the risk of large reversals. In the last Section we offer some final remarks and suggest further lines of research. An Appendix contains summary statistics from our data set and an illustrative microstructure model, alongside some robustness checks for our empirical analysis.

2 A Brief Literature Review

The FRU condition is a cornerstone in the study of the FX market. This condition states that in a risk-neutral efficient market, when agents are rational, the gain from borrowing at a low rate in one currency and lending at a high rate in another equals, on average, the loss on the exchange rate. Via the covered interest rate parity condition (CIP), this implies that the forward rate f_t at time t for delivery in period $t+1$ is the rational forecast for the corresponding spot rate s_{t+1} . Following Fama (1984) the FRU condition is usually tested by regressing FX returns, $s_{t+1} - s_t$, on the forward discount, $fd_t = f_t - s_t$, (the so-called Fama regression) and checking if $\alpha = 0$ and $\beta = 1$.⁴

$$s_{t+1} - s_t = \alpha + \beta fd_t + \epsilon_{t+1}, \quad (2.1)$$

However, in a multitude of studies (Lewis, 1995; Engel, 1996; Bacchetta et al., 2008; Burnside et al., 2009, among others), Fama's β is found to be significantly smaller than 1 and usually negative. Thus, Froot and Thaler (1990) indicate that the average value of the coefficient β across 75 published estimates is -0.88. Given the importance of the UIP hypothesis both in theoretical models and models used for policy, e.g. recent DSGE models used by many central banks, it seems essential to understand how breaches of the underlying assumptions contribute to the forward bias.

⁴The CIP condition states that $f_t - s_t = (i_t - i_t^*)\Delta t$ where i and i^* denote domestic and foreign interest rates, while Δt is the time interval, in years, between periods t and $t+1$. Akram, Rime, and Sarno (2008) show that the CIP condition holds for the purposes of this paper.

Some of the strongest earlier results on the forward discount puzzle have come from the analysis of market expectations derived from survey data. Froot and Frankel (1989) were amongst the first to investigate the role of forecast errors in explaining the failure of the FRU condition. They examined exchange rate forecasts for the US dollar against the Deutsche mark, British pound, French franc, Swiss franc and Japanese yen over several horizons, collected in the early and mid-1980s by *AMEX*, *The Economist* and *MMS*. Pooling together forecasts for different exchange rates, they estimate the contribution of forecast errors on Fama's β to lie between -6.07 and -0.52 depending on the survey data and the horizon of the forecasts.

Froot and Frankel's analysis has been extended by several authors, such as Frankel and Chinn (1993), Chinn and Frankel (2002), Cavaglia, Verschoor, and Wolff (1994), Bacchetta et al. (2008), who have considered alternative survey data, covering longer periods and more currency pairs. For example, Bacchetta et al. (2008) employ monthly surveys of 3-, 6- and 12-month forecasts for seven exchange rates over the period between August 1986 and July 2005. The estimated contribution from forecast errors to the coefficient β range from -3.62 to -0.76 across the seven exchange rates and the three horizons.

Although systematic forecast errors may seem irrational, these errors can also be due to either learning or a peso problem, as shown by Lewis (1989a,b) and Evans and Lewis (1995). In addition, slow reaction to news, through either ambiguity aversion (Ilut, 2009) or infrequent portfolio adjustments, induced by rational inattention combined with random walk expectations (Bacchetta and van Wincoop, 2009), may also generate forecast errors and a negative Fama's beta. Unfortunately, there is no consensus among researchers on the correct explanation for the presence of systematic errors in exchange rate forecasts. Importantly, even after allowing for forecast errors, the majority of these studies still find a statistically significant deviation from the UIP condition, indicating a role for time-varying risk premia (Jongen, Verschoor, and Wolff, 2008).

If perfect capital substitutability does *not* hold, a risk premium enters into the uncovered interest rate relationship. If this *time-varying* risk premium is *negatively* correlated with the forward discount, then Fama's β can turn out to be smaller than 1. Detecting such risk premia has been a very active, but arguably unsuccessful, research area. Cumby (1988), Hodrick (1989), and Bekaert, Hodrick, and Marshall (1997) find that *implausible* degrees of risk aversion are required to obtain a negative β in Fama's regression. However, recently there has been some progress. Lustig and Verdelhan (2007) find a role for consumption risk, whilst

Bansal and Shaliastovich (2007), Verdelhan (2010), and Moore and Roche (2010) all find some success explaining the puzzle with non-standard preferences.

The microstructure approach to foreign exchange, and order flow based models in particular, has recently made progress on exchange rate determination. Thus, Evans and Lyons (2002) and Berger, Chaboud, Chernenko, Howorka, and Wright (2008) find that trade imbalance in FX markets has *large* explanatory power for exchange rate returns. Payne (2003), Bjønnes and Rime (2005), Daniélsson and Love (2006), Killeen, Lyons, and Moore (2006) provide evidence that order flow has a *significant, large* and *persistent* impact on exchange rate returns. In addition, Evans and Lyons (2005), Froot and Ramadorai (2005) and Rime, Sarno, and Sojli (2010) show how order flow *anticipates* movements in exchange rate fundamentals. Finally, Breedon and Vitale (2010) and Breedon and Ranaldo (2008) suggest that order flow could be an important element of the FX risk premium through standard portfolio-balance effects.

Given their relative success in explaining the exchange rate disconnect puzzle, microstructure-based models have also been applied to shed light on the forward bias. For example, Burnside et al. (2007) suggest a mechanism whereby the forward bias arises through adverse selection mechanisms. Burnside et al. (2009) propose that transaction costs, while not necessarily explaining the puzzle, make it less obvious that the excess returns it implies can actually be achieved in practice. Ranaldo and Sarkar (2008) also find a role for illiquidity and volatility in explaining the puzzle. In a similar vein Bacchetta and van Wincoop (2009) suggest that infrequent portfolio adjustment could generate forward bias.

Our study follows this line of research, as we link the time-varying risk premium to trade imbalance in FX markets and attempt to measure it directly using transaction data. The intuition behind our contribution is that, as market participants are risk averse in FX markets, order flow gives rise to changes in risk premia, irrespective of whether order flow is driven by informational differences or not. Following such intuition with our study we aim at plugging a gap in the existing literature and providing some new insights on the origin of the forward discount bias.

3 Data and Preliminary Analysis

3.1 The Data

This study employs two innovative data sets to explore the link between expectations, risk premia and order flow. The first is a detailed transactions data set from EBS, created on a one-second timeslice basis, covering the period from the beginning of 1997 to April 2007 for trading in euro (dmark prior to 1999), yen and pound against the US dollar (i.e., the EUR/USD, USD/JPY, and GBP/USD exchange rates). For EUR/USD and USD/JPY we estimate that EBS covers close to half of all spot transactions, though for GBP/USD its coverage is poor (less than 5%). To our knowledge this is the longest data set of order flow from the foreign exchange market to date.

The second dataset is a detailed monthly survey of FX forecasts. At the beginning of each month (generally the first Tuesday of the month), Reuters call about 50 market participants to provide their forecasts of the same exchange rates at the 1-, 3-, 6- and 12-month horizon. Besides offering a meticulous archive of individual forecasts (the longest uninterrupted sample available), the Reuters survey has a number of advantages over other FX forecast surveys such as those undertaken by Consensus Economics, WSJ, ZEW, Blue Chip and Forecasts Unlimited (formerly the FT currency forecasts and the Currency Forecast Digest). First, since it is conducted by the key FX news provider, it is very much focussed on FX market participants, whereas other surveys often include many other forecasters such as professional forecast firms, corporations and academic institutions. This is important since, as Ito (1990) finds, these other forecasters are not comparable with those actively trading in foreign exchange. Second, the pool of forecasters is relatively constant. Other surveys have both gaps in coverage (missing individuals, months and in some cases years) and a relatively rapid turnover of contributors. Third, it is the only survey that collects forecasts for 1, 3, 6 and 12 months ahead, thus offering the most complete short-term coverage. Fourth, Reuters publish a ranking of forecasters each month that is widely followed and quoted by market participants, and the contributors thus have a strong incentive to take the survey seriously. See the Appendix for some descriptive statistics of our data on FX transactions and survey data on FX forecasts.

In addition, we also have data on interest rates and (at-the-money forward) implied volatilities for the same horizons as the forecasts. We construct monthly data (the frequency of the survey forecasts) by measuring all market prices (spot exchange rates, interest rates and implied volatilities) at the date of the survey compilation. Monthly order flow is then the

aggregate order flow since the previous forecast date. This gives us 124 observations at the monthly frequency.

3.2 The Forward Discount Bias

The starting point for almost all studies of the forward discount bias is Fama’s forward discount regression. In Panel A of Table 1 we show GMM estimates of Fama style regressions on monthly observations of spot returns on forward discounts for four different horizons (1 month, 3 months, 6 months and one year) for EUR/USD, USD/JPY, and GBP/USD,

$$s_{t+1} - s_t = \alpha + \beta fd_t + \epsilon_{t+1}, \quad (3.1)$$

where $fd_t = f_t - s_t$, and f_t is the log of the forward rate observed at the beginning of period t for maturity in period $t + 1$ and s_t is the log spot rate.

The results reported in Panel A in Table 1 are in line with previous studies: the estimated slope coefficient, β , is always negative and usually (particularly at the long horizons) significantly smaller than 1 (indicated by †), the value consistent with FRU. The Table suggests that, as found elsewhere, a profitable speculative strategy in these FX markets between 1997 and 2007 would have been that of betting against the forward discount, in that currencies with higher interest rates have tended to appreciate (for $fd_t > 0$, $s_{t+1} - s_t$ is on average negative) and *vice versa*.

In Panel B, we follow Froot and Frankel (1989) and report results from similar regressions using the expected return, $r_{e,t} = s_{t,e} - s_t$, constructed from the Reuters survey, as a dependent variable. As in previous studies, we find a substantial difference between Panel A and Panel B. Almost all coefficients are in fact larger in Panel B (except the one for USD/JPY 1 month), indicating that the forward discount is linked to market expectations of future exchange rates. However, all coefficients are still smaller than one, the value predicted by the UIP, and some, pertaining to the EUR/USD and USD/JPY exchange rates, are significantly so. This suggests that part of the forward discount bias is not explained by forecast errors, leaving room for an expected risk premium.⁵

⁵Indeed, most other studies of survey data find that in most cases the hypothesis of perfect substitutability (i.e. the restriction $\alpha = 0$ and $\beta = 1$ in the regression of $r_{e,t}$ on fd_t) is violated. (see (Cavaglia et al., 1994))

Table 1
Fama's Regression: Monthly Data

Panel A presents results from GMM estimates of β^k from the regression

$$r_t^k = \alpha^k + \beta^k f d_t^k + \epsilon_{t+k},$$

where $r_t^k = s_{t+k} - s_t$ is the return over the next k months, $f d_t^k = f_t^k - s_t$ is the corresponding forward discount, while f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t . Panel B presents results from GMM estimates of β_{er}^k from the regression

$$r_{e,t}^k = \alpha_{er}^k + \beta_{er}^k f d_t^k + \epsilon_{t,k}^{er},$$

where $r_{e,t}^k = s_{t,e}^k - s_t$ is the expected return over the next k months the interval $(t, t+k)$ and $s_{t,e}^k$ denotes the median value in month t of the k months ahead exchange rate forecasts contained in the Reuters survey. The maturity k is equal to 1, 3, 6 and 12, while t -statistics are reported in brackets. Coefficient values indicated by † are significantly smaller than 1 at the 5%-level. Sample: Jan 1997 - Apr 2007.

	1 Month	3 Month	6 Month	12 Month
Panel A: Realized return				
EUR/USD	-4.810† (-2.59)	-4.920† (-3.13)	-5.076† (-4.29)	-5.254† (-6.02)
USD/JPY	-1.874 (-1.19)	-1.608 (-1.09)	-1.761† (-1.48)	-1.854† (-2.34)
GBP/USD	-2.514 (-1.30)	-2.040 (-1.23)	-1.950† (-1.36)	-2.186† (-1.90)
Panel B: Expected return				
EUR/USD	-3.603† (-1.87)	-0.766† (-1.10)	0.316 (0.74)	0.642 (1.86)
USD/JPY	-2.870† (-1.80)	-1.404† (-1.75)	-0.432† (-0.72)	-0.036† (-0.09)
GBP/USD	-1.351 (-0.64)	0.007 (0.01)	0.333 (0.76)	0.474 (1.47)

4 A Microstructure-Based Decomposition of the Forward Bias

The deviation of Fama's β from unity can be due to violations of the key assumptions underlying the FRU condition, namely risk neutrality and rational expectations, leading to omitted variables in the Fama regression. If these omitted variables are negatively correlated with the forward discount then the estimates of β from the Fama regression will be downward-biased. This is the key idea behind the suggestion by Froot and Frankel (1989) to decompose the β -coefficient into its hypothesized value of 1 and deviations caused by the existence of risk premia and forecast errors that are correlated with the forward discount.

The log excess return on the FX currency is given by $er_{t+1} = s_{t+1} - s_t - (i_t - i_t^*) \Delta t$, where i and i^* are domestic and foreign nominal interest rates measured on a yearly basis. The FX risk premium is then defined as the expected excess return, $\rho_t = E er_{t+1}$, where the expectation is conditioned on period t information. Hence, we have

$$E_t [s_{t+1}] - s_t = (i_t - i_t^*) \Delta t + \rho_t. \quad (4.1)$$

Considering the definition of forecast error, $u_{t+1} = s_{t+1} - E_t[s_{t+1}]$, and using the CIP condition, $(i_t - i_t^*)\Delta t = f_t - s_t = fd_t$, we get

$$\Delta s_{t+1} = fd_t + \rho_t + u_{t+1}. \quad (4.2)$$

Thus, as suggested by Froot and Frankel (1989), the β -coefficient in Fama's regression, $\Delta s_{t+1} = \alpha + \beta fd_t + \epsilon_{t+1}$, differs from 1, the value consistent with the FRU condition, if the forward discount is correlated either with the forecast error, u_{t+1} , or the risk premium ρ_t .

Unfortunately, until recently it has proven difficult to find variables that enable us to measure directly FX risk premia and estimate in any satisfactory way their contribution to the forward bias. However, the recent microstructure approach to FX suggests that FX risk premia can be related to the trading process in the inter-dealer FX market. Specifically, Lyons (1997), Killeen et al. (2006) and Breedon and Vitale (2010) formulate market microstructure models of the FX market where spot rates are set by risk averse FX dealers and where order flow affects currency values via a portfolio-balance effect. Thus, as customer orders collected in the direct market are unwound into the inter-dealer market, a risk premium on the foreign currency emerges since FX dealers must be compensated for the extra risk they are forced to bear. In these models it is shown that the FX risk premium, ρ_t , is an increasing function of

inter-dealer order flow and a decreasing function of the capacity of FX dealers to absorb risky assets. Such capacity is proportional to their risk tolerance and their conditional precision of the future spot rate.

In line with such market microstructure models, in our empirical implementation we assume that the time-varying risk premium, ρ_t , is given by the ratio o_t/ν_t , where o_t is inter-dealer order flow, and ν_t is the conditional precision of next period spot rate.⁶ Then, as we measure the risk premium ρ_t via risk-adjusted order flow, o_t/ν_t , the following decomposition of Fama's beta applies,

$$\beta = 1 + \beta_o + \beta_u, \text{ where}$$

$$\beta_o = \frac{\text{cov}\left(\frac{1}{\nu_t}o_t, fd_t\right)}{\text{var}(fd_t)} \quad \text{and} \quad \beta_u = \frac{\text{cov}(u_{t+1}, fd_t)}{\text{var}(fd_t)}.$$

While analogous to the decomposition of Fama's β provided by Froot and Frankel (1989), ours gives more substance to the interpretation of the time-varying risk premium, which is now a function of order flow, o_t , and the conditional precision ν_t . Thus, unlike traditional attempts to explain the forward discount bias via the portfolio-balance approach, we use transaction data to measure directly deviations from the UIP condition and pin down their impact on Fama's beta.

4.1 Decomposing Fama's Beta

With our transaction and forecast data we can now estimate the contribution from risk premia - the coefficient β_o - and forecast errors - the coefficient β_u - on Fama's β (see equation(4.3)). The coefficient β_o can be estimated by running a linear regression of order flow, o_t , on the forward discount, fd_t , which allows us to identify the relationship between the risk premium related to order flow imbalance and the forward premium. Similarly, if we let $s_{t,e}$ denote the median value of the forecasts of professional FX traders for period $t + 1$ exchange rate formulated at time t , β_u can be estimated by running a linear regression of the forecast error,

⁶In the appendix we present a sketch of a model where these market microstructure mechanisms are incorporated and where ρ_t is shown to be equal to o_t/ν_t .

$s_{t+1} - s_{t,e}$, on the forward discount. We estimate these jointly in the following system which gives us one overidentifying restriction,

$$s_{t+1} - s_t = \alpha + (1 + \beta_o + \beta_u) fd_t + \epsilon_{t+1}, \quad (4.3)$$

$$o_t = \alpha_o + \beta_o fd_{t-1} + \epsilon_t^o, \quad (4.4)$$

$$s_{t+1} - s_{t,e} = \alpha_u + \beta_u fd_t + \epsilon_{t+1}^u. \quad (4.5)$$

To be consistent with the framework outlined above, and to have an order flow measure that matches the maturity of the forward contract, we aggregate order flow over the preceding interval $(t-1, t)$. In addition, since a given order flow imbalance will create a greater risk premium the more uncertain investors are about the future, we also multiply the aggregated order flow by an estimate of the *average* conditional variance of the exchange rate s_t across FX investors at time $t-1$. As a proxy of this conditional variance we employ the implied volatility from FX options of the appropriate maturity observed at the beginning of period $t-1$.⁷

The results from GMM estimation of the system above are presented in Table 2. The first column reports the implied Fama beta-coefficient, $1 + \beta_o + \beta_u$. In square brackets below the coefficients we report p -values for the J -test of the overidentifying restriction in our system. The reported values show the restriction $\beta = 1 + \beta_o + \beta_u$ is never rejected, confirming the validity of our decomposition and suggesting that we capture a significant share of the bias. In addition, the estimated values for the forecast error and the order flow coefficients, β_u and β_o , suggest the following: on the one hand, the forecast errors contribute significantly to a negative bias in the forward discount for EUR/USD and GBP/USD, but not for USD/JPY. On the other hand, order flow contributes significantly to a negative bias for EUR/USD and USD/JPY but not for GBP/USD.

Indeed, taking average values of the coefficients across the four horizons, we see that for EUR/USD risk-adjusted order flow explains roughly half of the deviation of β from 1, ie. half of the forward discount bias, while the other half is explained by the forecast error (see Table 3). For USD/JPY an even stronger conclusion is reached, as nearly all the bias is explained by risk-adjusted order flow. By contrast, for GBP/USD the proportion explained by risk-adjusted order flow is less than 10%. The poor results for GBP/USD may well reflect the fact

⁷As an alternative estimate we consider the conditional variance of the next period exchange rate forecasts collected by Reuters at the beginning of period $t-1$. These results are discussed in the Appendix.

Table 2
Decomposition of Fama's Beta

The Table presents the coefficient values of β_o^k and β_u^k (with t -statistics below) from GMM estimation of the system

$$\begin{aligned}
 s_{t+k} - s_t &= \alpha^k + \left(1 + \beta_o^k + \beta_u^k\right) f d_t^k + \epsilon_{t+k}, \\
 o_{t,k} &= \alpha_o^k + \beta_o^k f d_{t-k}^k + \epsilon_{t,k}^o, \\
 s_{t+k} - s_{t,e}^k &= \alpha_u^k + \beta_u^k f d_t^k + \epsilon_{t+k}^u.
 \end{aligned}$$

The order flow variable $o_{t,k}$ is cumulated between month $t-k$ and t , and is also pre-multiplied by the k months ahead exchange rate variance, measured by squared implied volatility at the end of month $t-k$, $f d_t^k = f_t^k - s_t$ is the forward discount; f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t ; $s_{t,e}^k$ denotes the median value in month t of the k months ahead exchange rate forecasts contained in the Reuters survey. The column "Implied" reports the implied Fama's β ($1 + \beta_o^k + \beta_u^k$) and in squared brackets is the p -value from the J -test of the overidentifying restriction (that the implied β is equal to Fama's beta). A † indicates that $\beta_o^k + \beta_u^k$ is not significantly different from zero at the 5% level (i.e. UIP cannot be rejected). Sample: Jan 1997 - Apr 2007.

	EUR/USD			USD/JPY			GBP/USD		
	Implied	OF	ExpE	Implied	OF	ExpE	Implied	OF	ExpE
1 Month	-4.31 [0.71]	-4.25 (-3.44)	-1.06 (-1.28)	-2.04 [0.82]	-3.97 (-3.21)	0.92 (1.04)	-0.57† [0.16]	-0.10 (-0.08)	-1.47 (-2.09)
3 Month	-4.73 [0.13]	-2.63 (-2.55)	-3.11 (-2.18)	-2.31 [0.31]	-2.70 (-2.26)	-0.61 (-0.39)	-1.82† [0.23]	-0.54 (-0.72)	-2.27 (-1.52)
6 Month	-5.43 [0.13]	-2.81 (-3.14)	-3.62 (-2.37)	-2.99 [0.17]	-3.23 (-2.70)	-0.76 (-0.48)	-2.09† [0.37]	-0.35 (-0.58)	-2.74 (-1.61)
12 Month	-6.01 [0.15]	-2.80 (-3.43)	-4.21 (-2.81)	-3.37 [0.14]	-4.23 (-4.44)	-0.14 (-0.12)	-2.82 [0.22]	0.29 (0.64)	-4.11 (-2.72)

that EBS has a very small market share for that cross (see Table 8 in the appendix) so that our transaction data are not representative.

Table 3
Share of Forward Bias Explained by Order Flow

The Table presents estimates of the overall forward bias ($\beta_u^k + \beta_o^k$) and the share explained by order flow $\beta_o^k / (\beta_u^k + \beta_o^k)$ derived from our GMM estimates presented in Table 2.

	EUR/USD		USD/JPY		GBP/USD	
	Forward bias	OF share	Forward bias	OF share	Forward bias	OF share
1 Month	-5.31	0.80	-3.04	1.30	-1.57	0.06
3 Month	-5.73	0.46	-3.31	0.81	-2.82	0.19
6 Month	-6.43	0.44	-3.99	0.81	-3.09	0.11
12 Month	-7.01	0.40	-4.37	0.97	-3.82	-0.08
Mean	-6.12	0.52	-3.68	0.97	-2.82	0.07

5 Carry Trades and the Forward Discount Bias

5.1 Carry Trades and the Decomposition

Interestingly, our decomposition can offer some insights into the impact of carry trades in FX markets and their role in generating the forward discount bias. Certainly, results from Table 2 indicate that the role of the time-varying risk premium in explaining the forward discount bias is more pronounced for USD/JPY, which is the archetypal carry trade cross.

Galati, Heath, and McGuire (2007), Burnside et al. (2009, 2007), and Jylhä and Suominen (2010), Lustig, Roussanov, and Verdelhan (2009) find *positive* returns for carry trade. Carry trade profitability is a direct consequence of the failure of the FRU condition, as indeed, contrary to its prediction, high interest rate currencies tend to appreciate vis-a-vis low interest rate currencies.

Several explanations for the apparent profitability of the carry trade have been proposed. Thus, recent studies suggest that carry trade profits are mitigated by transaction costs (Burnside et al., 2009), are associated with volatility and illiquidity (Ranaldo and Sarkar, 2008; Jylhä and Suominen, 2010), are counter-cyclical (Lustig et al., 2009) and subject to reversal risk (Breedon, 2001; Brunnermeier, Nagel, and Pedersen, 2009).

Plantin and Shin (2008) show that in the presence of liquidity constraints expectations of carry trade profitability are self-fulfilling. In their model, when carry traders short a low interest rate currency to buy a high interest rate currency they drive down the value of the former and drive up that of the latter, so that their expectations are fulfilled. This happens because in Plantin and Shin’s model trade imbalance has a positive impact on exchange rate returns, as suggested by recent empirical evidence from the market microstructure approach to exchange rates and by our results here.

Our market microstructure perspective can accommodate carry trade activity and show how it contributes to the forward discount bias. Thus, consider that according to the models of Lyons (1997), Killeen et al. (2006) and Breedon and Vitale (2010) FX dealers unwind into the inter-dealer market the orders they receive from their customer base. This means that inter-dealer order flow, o_t , is positive function of the flow of customer orders in the direct section of the FX market, c_t . Simplifying, we impose the following equality

$$o_t = c_t. \tag{5.6}$$

Although the customers of FX dealers trade for a large variety of reasons, anecdotal evidence suggests that in several FX markets a significant component of customer trading activity is motivated by carry trading. Thus, let us assume that in the presence of a negative forward discount, $(i_t - i_t^*)\Delta t = fd_t < 0$, these customers expect positive profits from a long carry trade strategy on the foreign currency and so purchase it.⁸ To capture our *carry trade hypothesis* we assume that customer order flow respects the following formulation,

$$c_t = -\mu fd_t + n_t.$$

Here μ is a positive constant, so that carry traders sell the foreign currency if this is a low interest rate currency (and vice versa if it is the high interest rate currency), while n_t is a second component of customer order flow not related to the forward discount.

In the presence of such carry trade activity, and using the dealer-customer condition (5.6), we derive a negative covariance between order flow and the forward discount, $\text{cov}[o_t, fd_t] < 0$.

⁸Although our data set does not include forward transactions explicitly, the standard carry trade involves a spot transaction at initiation and completion and so is captured by our data.

This implies that β_o takes a negative value and hence that Fama's beta is smaller than 1. Specifically, for ν_t time-invariant, we find that

$$\beta_o = \frac{\text{cov}\left(\frac{1}{\nu_t}o_t, fd_t\right)}{\text{var}(fd_t)} = -\frac{\mu}{\nu}.$$

Assuming that the FX dealers are rational, so that $\beta_u = 0$, we conclude that

$$\beta = 1 - \frac{\mu}{\nu}.$$

In brief, according to our carry trading hypothesis, Fama's beta is smaller than 1. Moreover, if such activity is particularly intensive, i.e. if μ is large, β can actually take a negative value, as found in many empirical studies on the forward discount bias.

5.2 Order Flow, the Forward Discount and the Time-Varying Risk Premium

The negative correlation between order flow and forward discount is clearly documented for the USD/JPY and EUR/USD rates in Table 4. Here, we report the results of regressing order flow over the past period on the past interest rate differential. For these two exchange rates we find a strong and significant impact of interest rate differentials on order flow. As US interest rates rise relative to those in Japan or in the euro area, market participants subsequently buy more USD. The negative coefficient for EUR/USD is due to a positive interest rate differential giving rise to negative order flow since EUR is the base currency, while the negative coefficient for USD/JPY is due to a negative interest rate differential giving rise to a positive order flow since USD is the base currency in USD/JPY. The large explanatory power for EUR/USD and USD/JPY, given by the \bar{R}^2 , confirms that in these markets carry trading generates a significant proportion of total trade imbalance.

Results for GBP/USD in Table 4 give a different picture. The coefficient β_o is neither negative nor significant, while the explanatory power is an order of magnitude smaller, indicating that carry trading does not generate much order flow in this market. There are two main explanations for the weak results obtained for GBP/USD. First, as discussed above, EBS is not the dominant electronic trading platform for this cross and our order flow measure is thus significantly less representative in this case. Second, GBP/USD is not often considered a carry trading cross and the carry trade activity that we find to be important in the case of USD/JPY

Table 4
The Impact of the Forward Discount on Order Flow

This Table reports estimates of a linear regression of order flow, $o_{t,k}$, on the forward discount, fd_t^k ,

$$o_{t,k} = \alpha_o^k + \beta_o^k fd_{t-k}^k + \epsilon_{t,k}^o,$$

with $k = 1, 3, 6, 12$ months. The order flow variable $o_{t,k}$ is cumulated between month $t - k$ and t , and is also pre-multiplied by the k months ahead exchange rate variance, measured by squared implied volatility at the end of month $t - k$; the forward discount is $fd_t^k = f_t^k - s_t$, where f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t . Sample: Jan 1997 - Apr 2007.

Currency	Horizon	β_o^k	t -stat	adj. R^2
EUR/USD	1	-0.037	-3.58	0.17
	3	-0.039	-3.90	0.21
	6	-0.039	-3.94	0.22
	12	-0.039	-3.59	0.23
USD/JPY	1	-0.047	-2.23	0.06
	3	-0.055	-2.68	0.11
	6	-0.058	-3.09	0.14
	12	-0.064	-3.99	0.21
GBP/USD	1	0.005	0.39	-0.01
	3	0.006	0.50	0.00
	6	0.005	0.48	0.00
	12	0.010	1.40	0.07

in particular is thus less relevant for GBP/USD. As a result, we drop GBP/USD from the rest of our analysis.

For carry trading to be a significant explanation of the forward discount puzzle, three conditions must hold. First, traders expect carry trade activity to generate positive profits. This is the case when, in the face of a negative (positive) forward discount, the expected excess return on a long (short) carry trade position is positive, i.e. if

$$\begin{aligned} \text{for } i_t < i_t^* &\Rightarrow E_t[s_{t+1}] - s_t + (i_t^* - i_t) \Delta t > 0 \text{ and} \\ \text{for } i_t > i_t^* &\Rightarrow E_t[s_{t+1}] - s_t + (i_t^* - i_t) \Delta t < 0. \end{aligned}$$

This condition holds if in the regression of the expected return on the foreign currency, $r_{e,t} = s_{t,e} - s_t$, on the forward discount, fd_t ,

$$r_{e,t} = \alpha_{er} + \beta_{er} fd_t + \epsilon_t^{er},$$

β_{er} is smaller than one. Results reported in panel B of Table 1 indicate that such a condition holds for USD/JPY, as the slope coefficient is significantly smaller than one, the value consistent with the UIP, across all maturities. Results for EUR/USD are less supportive as the slope coefficient, while always smaller than 1, is significantly so only for the 1- and 3-month horizons. This might be interpreted as indicating that carry traders mostly concentrate their speculative positions on EUR/USD over shorter horizons.

Second, expectations of carry trade profitability generate trade imbalance. In particular, for $E_{t-1}[s_t] - s_{t-1}$ positive (negative), FX customers purchase the foreign (domestic) currency for the domestic (foreign) one, i.e. order flow in the interval $(t-1, t)$ is positive (negative). To test this condition we run a regression of the risk-adjusted order flow in the interval $(t-1, t)$, o_t , on the expected return at time $t-1$, $r_{e,t-1}$,

$$o_t = \alpha_o + \lambda_o r_{e,t-1} + \epsilon_t^o,$$

to see whether expectations of an appreciation (depreciation) of the foreign currency, and hence expectations of profits from a long (short) carry trade position on the foreign currency, generate corresponding flows. This is the case if λ_o is positive. GMM estimates of this regression for the EUR/USD and USD/JPY rates are in Table 5. The results are clearly supportive. In fact, the slope coefficient is positive for all maturities and rates. In addition, most values are

significantly larger than zero, indicating that when FX customers expect profits from a long (short) position on the foreign currency, they purchase (sell) it.

Table 5
The Impact of the Expected Return on Risk-Adjusted Order Flow

The Table reports results of GMM estimates of the regression of risk-adjusted order flow on the expected return on the foreign currency,

$$o_{t,k} = \alpha_o^k + \lambda_o^k r_{e,t-k}^k + \epsilon_{t,k}^o,$$

where $k = 1, 3, 6, 12$ months. The order flow variable $o_{t,k}$ is cumulated between month $t - k$ and t , and is also pre-multiplied by the k months ahead exchange rate variance, measured by squared implied volatility at the end of month $t - k$. The expected return on the foreign currency is $r_{e,t}^k = s_{t,e}^k - s_t$, where $s_{t,e}^k$ denotes the median value in month t of the k months ahead exchange rate forecasts contained in the Reuters survey; the forward discount is $fd_t^k = f_t^k - s_t$, where f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t . The Table contains the estimates of the slope coefficient λ_o^k (with the corresponding t -statistics in brackets). Sample: Jan 1997 – Apr 2007.

	1 Month	3 Month	6 Month	12 Month
EUR/USD	0.00 (0.02)	0.26 (2.78)	0.56 (2.85)	0.80 (2.46)
USD/JPY	0.11 (0.97)	0.90 (3.10)	1.42 (2.51)	1.64 (1.66)

Thirdly, trade imbalance in the FX markets affects expected risk premia. In Table 6 we investigate if order flow is a determinant of the expected risk premia, defined as $s_{t,e} - s_t - fd_t$. Results in Table 6 are clear: for most horizons and exchange rates there is a positive and significant impact of order flow on expected risk premia, consistent with our analytical framework. An example may clarify the effect: when the USD is expected to appreciate against the JPY, and the US interest rate is higher than the Japanese one, the expected risk premium is positive. The results in Table 6 indicate that this occurs when there has been a period with net buying of USD against JPY (positive order flow). This would be the case e.g. if market participants are following carry trade strategies, i.e. borrowing in JPY and lending in USD.

Indeed, the thesis that the impact of order flow on expected risk premia is related to carry trades is supported by the relatively large explanatory power of order flow for the USD/JPY rate, i.e. for a currency pair on which carry trade activity is usually intense. In fact, while not reported in Table 6, for this rate the adjusted coefficient of multiple determination, \bar{R}^2 , in the regressions of the expected risk premium on order flow ranges from 1% to 48%.

Table 6
The Impact of Order Flow on Expected Risk Premia

The Table reports GMM estimates of the coefficient γ_{ep}^k in the regression of the expected risk premium on order flow,

$$s_{t,e}^k - f_t^k = \alpha_{ep}^k + \gamma_{ep}^k o_{t,k} + \epsilon_{t,k}^{ep},$$

with $k = 1, 3, 6, 12$ months. The order flow variable $o_{t,k}$ is cumulated between month $t - k$ and t and is also pre-multiplied by the k months ahead exchange rate variance, measured by squared implied volatility at the end of month $t - k$. t -statistics in brackets. Sample: Jan 1997 - Apr 2007.

	1 Month	3 Month	6 Month	12 Month
EUR/USD	0.198 (1.04)	0.212 (2.82)	0.100 (2.24)	0.027 (0.92)
USD/JPY	0.054 (0.35)	0.154 (3.54)	0.150 (3.32)	0.153 (5.89)

All in all, the evidence provided in Tables 1 (panel B), 5 and 6 supports our carry trade hypothesis for the EUR/USD and USD/JPY rates, suggesting that for these rates the component of the forward discount bias associated with the time-varying risk premium is generated by carry trade activity. In fact, we see that shifts in the forward discount induce expectations of carry trade profitability and generate trade imbalance accordingly. In turn, order flow affects expected risk premia and brings about a Fama's β smaller than 1.

5.3 Carry Trade Activity and Currency Crash Risk

The evidence from our analysis and earlier studies suggesting that carry trading is profitable is puzzling, in that one may wonder why smart traders would not under-cut FX dealers' quotes and eliminate the excess returns such investors enjoy. However, as suggested by Brunnermeier et al. (2009), such activity is subject to crash risk, in that movements in currency returns consistent with carry trade profitability may suddenly change direction and entail large losses for carry trading positions.

In our sample, EUR/USD daily returns display pronounced positive skewness, whereas the opposite holds for USD/JPY. This reflects the fact that the USD generally is an investment currency in the carry trade strategy whereas the EUR and the JPY are funding currencies. For carry trading in the EUR/USD cross to be profitable the USD must appreciate vis-a-vis the EUR (and hence the EUR/USD rate must decrease). However, *positive* skewness of

EUR/USD returns indicates the risk of a currency crash, in that the appreciation of the US currency is subject to sudden and deep reversals, which cause carry traders to suffer speculative losses. A similar argument holds for the USD/JPY cross, as for carry trading to be profitable the USD/JPY rate must increase. In this case, currency crash risk translates into *negative* skewness of the return on USD/JPY.

Brunnermeier et al. (2009) claim that currency reversals are the result of the sudden unwinding of carry trades when these speculators hit liquidity constraints. An empirical implication of such a thesis is that the trade imbalance provoked by carry trading *per se* augments the risk of currency reversals (carry crashes). They provide some weak evidence of such an effect, but argue that their trade imbalance data (based on CFTC FX futures positions) is problematic.

A way to test their empirical implication using our data on FX transactions consists of regressing the skewness of FX returns on order flow. In particular, as speculators accumulate USD vis-à-vis the EUR a negative order flow in the EUR/USD market should translate into larger positive skewness for the corresponding FX return, if carry trading increases currency crash risk. Similarly, as the same investors purchase USD vis-à-vis the JPY, a positive order flow in the USD/JPY market should now translate into larger negative skewness of the FX return. In both cases, regressing the skewness of the FX returns on order flow should yield a negative slope coefficient. In order to compare our results with those of Brunnermeier et al. (2009), we also include the forward discount and implied volatility as controls.

Results of the regression of the realized skewness of daily FX returns in the period $(t-1, t)$, ζ_t , on lagged order flow, forward discount and implied volatility (all of the appropriate maturity) are reported in Table 7. The coefficient on order flow is correctly signed (negative) and significant for all horizons except the 12-month for EUR/USD and the 1-month for USD/JPY. The forward discount is correctly signed for both exchange rates (in the sense that positive carry is a predictor of currency crashes), but significant only for USD/JPY. Interestingly, we find a significant relationship between implied volatility and skewness at the 3-month horizon for EUR/USD and at the 3-month horizon and above for USD/JPY. This result is interesting since for both crosses it implies that low volatility is a predictor of carry crashes, which is seemingly at odds with Brunnermeier et al. (2009), who suggest that carry crashes and volatility are positively related. The main explanation for this difference is that Brunnermeier et al. (2009) look at the *contemporaneous* relationship between volatility and skewness, while we undertake a predictive regression. As admittedly tenuous, out-of-sample evidence, it is intriguing to note

Table 7
The Impact of Order Flow on the Skewness of FX Returns

The Table reports GMM estimates of the coefficients from the regression of the average skewness of daily FX returns in the period $(t - k, t)$, ζ_t^k ,

$$\zeta_t^k = \alpha_{sk}^k + \gamma_{sk}^k o_{t,k} + \beta_{sk}^k fd_{t-k}^k + \delta_{sk}^k ImpVol_{t-k}^k + \epsilon_{t,k}^{sk},$$

with $k = 1, 3, 6, 12$ months. The order flow variable $o_{t,k}$ is cumulated between month $t - k$ and t ; the forward discount is $fd_t^k = f_t^k - s_t$, where f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t ; $ImpVol_t^k$ denotes the k months ahead implied volatility at the end of month t . t -statistics in parenthesis. Sample: Jan 1997 - Apr 2007.

	EUR/USD			USD/JPY		
	OF	FD	IV	OF	FD	IV
1 Month	-13.78 (-1.85)	37.88 (0.78)	-1.14 (-0.40)	-3.07 (-0.58)	48.10 (2.08)	-0.19 (-0.10)
3 Month	-5.97 (-1.88)	9.40 (0.50)	-8.55 (-2.58)	-4.25 (-2.21)	34.92 (4.50)	4.54 (2.69)
6 Month	-3.94 (-2.01)	4.72 (0.45)	-7.03 (-1.82)	-2.90 (-4.08)	22.78 (5.83)	6.86 (4.28)
12 Month	-1.27 (-1.16)	5.85 (0.99)	-6.10 (-1.41)	-1.29 (-4.03)	7.97 (5.40)	4.41 (2.99)

that for both EUR/USD and USD/JPY implied volatility reached multi-year lows in mid-2007 just before the financial crisis and a significant carry crash for both currency pairs.

In brief, we conclude that while carry traders can expect profits from their speculative activity in the EUR/USD and JPY/USD markets, they also face significant crash risk, which is at its highest when carry trading has resulted in significant order flow imbalance and when the interest rate differential is high and/or volatility is low.

6 Concluding Remarks

A large body of research has been devoted to the forward discount bias and the profitability of carry trades. Our study contributes to this literature by analyzing the information contained in Reuters survey data of exchange rate forecasts and in EBS transaction data. We exploit this information to decompose the forward discount bias into two parts, one due to forecast errors and the other due to an order flow related time-varying risk premium.

In common with most other studies, our results suggest that forecast errors only partially explain the forward discount bias, as when using expected returns in lieu of actual returns the coefficient on the forward discount is still smaller than 1, the value consistent with uncovered interest rate parity. However, we find evidence, particularly strong for the EUR/USD and USD/JPY crosses, that order flow affects expected risk premia and that these condition realized returns, indicating that microstructural mechanisms contribute to the forward discount puzzle. Thus, according to our decomposition of Fama's beta, the portfolio-balance effect of trade imbalance explains roughly 50 percent of the forward discount bias for the EUR/USD rate and more than 90 percent of the bias for the USD/JPY rate. We do not find any similar importance of order flow for the GBP/USD forward bias, and we argue that this is partly because the EBS trading platform is not a widely used trading platform for this cross.

In addition, our results suggest that carry trade activity may actually generate part of the forward discount bias. Thus, we find that: i) movements in interest rate differentials generate order flow imbalance in FX markets in line with carry trading; ii) such activity is sustained by expectations of carry trade profits; and iii) it affects expected risk premia resulting in the appreciation of high interest rate currencies. Finally, we see that carry trading activity does not represent a *free lunch*, in that the positive profits it is expected to gain are offset, to some extent, by the currency crash risk it provokes.

Appendix

A Descriptive Statistics of the Data

FX TRANSACTIONS: Our FX transactions data set, created on a one-second timeslice basis, comes from EBS, the dominant electronic broker for the EUR and JPY rates, but not for the GBP rate. Table 8 contains summary statistics on the FX transaction data including estimates of EBS market share in each cross. In the paper we use DEM data instead of EUR for the period to the introduction of the EUR in 1999. EBS suggest that they were also the dominant trading platform for DEM over that period and our overall results are similar if we exclude those two years.

Table 8
EBS Turnover Data Summary Statistics

This Table presents summary statistics for our sample of EBS turnover data. We show estimates of EBS share of electronic inter-dealer trading and overall FX turnover. We also show average trade size (2000-2007) and average bid-ask spread (1997-2007) for all active trading hours (i.e. hours in which at least one trade took place). The share of electronic inter-dealer broking is derived from a comparable sample of EBS and Reuters Dealing-2002 (the other electronic interdealer broking platform) from August 2000 to January 2001 (Breedon and Vitale, 2010). Overall market share is estimated from the 1998, 2001, 2004 and 2007 BIS surveys by assuming that all trading between reporting dealers is electronic. This is likely to be an overestimate at the start of the sample (as other trading methods were used) but an underestimate at the end of the sample (as EBS is now being used by some customers such as hedge funds).

	EUR/USD	USD/JPY	GBP/USD
EBS share of electronic	81%	95%	7%
Electronic share of total	54%	50%	54%
EBS share of total	44%	48%	4%
Average Trade Size	\$4.49 mln.	\$3.87 mln.	\$3.57 mln.
Average Bid-Ask Spread	0.017%	0.018%	0.056%

FX FORECASTS: Our forecast data set is based on the full set of forecasts that make up the Reuters survey of FX forecasts. At the beginning of each month (generally the first Tuesday of the month), Reuters call about 50 market participants to obtain their forecasts of future exchange rates. The forecast horizons are set at one month, three months, six months, and twelve months respectively. Table 9 contains summary statistics for the FX forecasts. Note

that, in common with other forecast surveys, the median forecast does not outperform a naive, random walk forecast (i.e. Theil statistics are greater than 1).

B A Microstructure Model of the Inter-dealer Market

In this section we present a simple analytical framework for the FX market with the objective of deriving a time-varying risk premium which depends on the trading imbalance in the inter-dealer FX market. This illustrative framework, which is based on the model of the FX market put forward by Breedon and Vitale (2010), is designed to represent the trading features of the electronic trading platforms, such as *Electronic Broking Services* (EBS) and *Reuters* indirect dealing systems, which dominate the inter-dealer markets.

As on these platforms transactions are completed via a centralized *limit order book*, where subscribers can at any time either add/delete *limit orders* or hit outstanding limit orders with *market orders* of opposite sign, we assume that in the inter-dealer market a single foreign currency is traded for the currency of a large domestic economy. Trades are completed according to a sequence of Walrasian auctions which are intended to represent the Reuters and EBS electronic trading platforms. Hence, we assume that in any period t , FX traders simultaneously enter either market or limit orders and then a clearing price (exchange rate) for the foreign currency is established.

We follow Breedon and Vitale (2010) in assuming that FX dealers select portfolios of domestic and foreign deposits by maximizing the expected CARA utility of their end-of-period wealth. Then, under normality they choose to submit the following aggregate demand for the foreign currency in period t

$$o_t \equiv \nu_t \left(\bar{E}_t^1 [s_{t+1}] - s_t + (i_t^* - i_t) \Delta t \right), \quad (\text{A.1})$$

where ν_t is the *aggregate* trading intensity of the population of FX dealers, given by the risk-tolerance weighted average of their conditional precision of next period spot rate, and $\bar{E}_t^1 [s_{t+1}]$ is the weighted average of the expected value of next period spot rate across all FX dealers, where the individual FX dealers' weights are given by their trading intensities.

While the assumptions behind its derivation are specific to the current formulation, the demand function in equation (A.1) holds under alternative specifications. Thus, equation

Table 9
Foreign-exchange Forecasts Summary Statistics

This Table presents summary statistics for our sample of foreign exchange forecasts. For each forecasting horizon, we show the maximum, average and minimum number of individual forecasts each month, the maximum, average and minimum standard deviation of those forecasts (expressed as a percentage of the average forecast) and the Theil statistic (RMSE of the average forecast divided by the RMSE of a random walk forecast) Notice that one forecaster consistently only provided one-month forecasts.

		EUR/USD	USD/JPY	GBP/USD
Panel A: One-month horizon				
No. of forecasts	Max no.	66	66	65
	Ave. no.	52.1	51.2	51.0
	Min. no	30	30	30
Forecast dispersion	Max stdev.	2.9	13.4	2.1
	Ave. stdev.	1.7	3.1	1.3
	Min stdev.	0.9	1.1	0.8
Forecast accuracy	Theil stat.	1.00	1.04	1.03
Panel B: Three-month horizon				
No. of forecasts	Max no.	67	67	66
	Ave. no.	52.5	51.9	51.5
	Min. no	29	29	29
Forecast dispersion	Max stdev.	4.5	6.9	4.0
	Ave. stdev.	2.9	2.9	2.2
	Min stdev.	1.5	1.4	1.5
Forecast accuracy	Theil stat.	1.07	1.15	1.01
Panel C: Six-month horizon				
No. of forecasts	Max no.	66	66	65
	Ave. no.	52.3	51.7	51.2
	Min. no	29	29	29
Forecast dispersion	Max stdev.	6.0	14.6	4.9
	Ave. stdev.	4.1	3.1	3.1
	Min stdev.	2.3	1.7	2.1
Forecast accuracy	Theil stat.	1.13	1.15	1.02
Panel D: One-year horizon				
No. of forecasts	Max no.	66	66	65
	Ave. no.	51.8	51.4	50.7
	Min. no	29	29	29
Forecast dispersion	Max stdev.	9.0	7.8	5.9
	Ave. stdev.	5.6	3.7	4.2
	Min stdev.	3.3	1.4	3.0
Forecast accuracy	Theil stat.	1.13	1.21	0.98

(A.1) can be derived from a mean-variance portfolio choice model, or from an OLG portfolio model, as in Bacchetta and van Wincoop (2006), or even from an inter-temporal portfolio choice problem, as in Evans and Hnatkovska (2007).

As the (net) demand for foreign currency on the part of the FX dealers is entered on the centralized platform, o_t will correspond to order flow, ie. the difference between buyer and seller initiated transactions for the foreign currency. This order flow will then be absorbed by broker-dealers which trade in the inter-dealer market on behalf of traders who do not have direct access to the centralized platform.

Rearranging equation (A.1) we obtain a modified UIP equation,

$$\bar{E}_t^1 [s_{t+1}] - s_t = (i_t - i_t^*) \Delta t + \frac{1}{\nu_t} o_t. \quad (\text{A.2})$$

Equation (A.2) implies that, thanks to the FX dealers' risk-aversion, *uncovered interest rate parity* does not hold. Indeed, the interest rate differential, $i_t - i_t^*$, is *proportional* to the difference between the average expected devaluation of the domestic currency in period t and a risk premium on the foreign currency the FX dealers collectively require to hold foreign assets. This is a *time-varying* risk premium, given by the product of the total demand of foreign assets the FX dealers have to share and the inverse of their aggregate trading intensity, ν_t (which measures the dealers' capacity to hold risky assets). In other words, the larger the average risk tolerance of our population of FX dealers, the smaller the risk premium imposed on the foreign currency. Likewise, the smaller the perceived uncertainty of the currency return, measured by the inverse of the average precision, the less the perceived risk of the foreign currency and so the smaller the required risk premium.

C Some Robustness Checks: Sensitivity to Volatility

According to our analytical framework, the variable o_t is obtained by multiplying the cumulative order flow between $t - 1$ and t by an estimate of the average conditional variance of the exchange rate s_t across FX investors at time $t - 1$. As a measure of this conditional variance we have employed the implied volatility observed at the beginning of period $t - 1$. However, as an alternative estimate we can use the cross-section variance of the individual FX forecasts in period $t - 1$ of the exchange rate at time t contained in Reuters survey. This definition captures the concept of differences in beliefs that is found to be important in FX markets by Beber,

Breedon, and Buraschi (2010). In Table 10 we report the results of the regressions using the cross-section variance of Reuters individual forecasts in lieu of the implied volatility.

Table 10
Decomposition of Fama’s Beta: Dispersion of Forecasts as Measure of Uncertainty

The Table presents the coefficient value of β_o^k and β_u^k (with t -statistics below) from GMM estimation of the system. The column “Implied” reports the implied Fama’s β ($1 + \beta_o^k + \beta_u^k$) and in squared brackets is the p -value from the J -test of the overidentifying restriction (that the implied β is equal to Fama’s beta). A † indicates that $\beta_o^k + \beta_u^k$ is not significantly different from zero at the 5% level (i.e. UIP cannot be rejected). Sample: Jan 2000 – Apr 2007.

	EUR/USD			USD/JPY			GBP/USD		
	Implied	OF	ExpE	Implied	OF	ExpE	Implied	OF	ExpE
1 Month	-5.07 [0.90]	-4.00 (-2.77)	-2.07 (-2.46)	-1.42† [0.69]	-3.12 (-2.20)	0.70 (0.84)	-1.83 [0.82]	-0.51 (-0.45)	-2.32 (-3.00)
3 Month	-4.71 [0.11]	-1.97 (-1.55)	-3.74 (-2.45)	-1.12† [0.42]	-1.89 (-2.01)	-0.23 (-0.15)	-2.00† [0.40]	-0.46 (-0.61)	-2.55 (-1.63)
6 Month	-5.36 [0.27]	-2.23 (-1.55)	-4.13 (-2.49)	-1.83 [0.75]	-1.48 (-1.93)	-1.36 (-0.83)	-2.33 [0.36]	-0.27 (-0.44)	-3.06 (-1.90)
1 Year	-6.17 [0.14]	-2.60 (-1.96)	-4.57 (-2.69)	-2.77 [0.18]	-1.77 (-1.75)	-2.00 (-1.45)	-3.31 [0.20]	0.18 (0.30)	-4.49 (-3.64)

Since our order flow measure in the regressions is multiplied with volatility one may wonder whether the contribution from this variable comes from variation in the transaction data or from the volatility measure. We address this by extending our system for estimating the decomposition, given by equations (4.3)-(4.5), with an equation for volatility, making it a four-equation system, and at the same time dropping order flow’s dependence on volatility. This way we measure the separate role of each variable on the bias. Results are reported in Table 11, where only the implied β from the system and the contribution of order flow is reported in order to save space. The implied beta coefficient is again of a similar magnitude as the unconstrained estimate in Table 1, while the contribution from order flow is equally strong and significant.

Table 11
Decomposition of Fama's Beta: Sensitivity to Volatility

The Table presents GMM estimates of the following 4-equation system,

$$\begin{aligned}
 s_{t+k} - s_t &= \alpha^k + \left(1 + \beta_o^k + \beta_u^k + \beta_v^k\right) f d_t^k + \epsilon_{t+k}, \\
 o_{t,k} &= \alpha_o^k + \beta_o^k f d_{t-k}^k + \epsilon_{t,k}^o, \\
 s_{t+k} - s_{t,e}^k &= \alpha_u^k + \beta_u^k f d_t^k + \epsilon_{t+k}^u, \\
 ImpVol_{t-k}^k &= \alpha_v + \beta_v f d_{t-k}^k + \epsilon_{t-k,k}^v.
 \end{aligned}$$

In this system order flow and volatility have separate equations and effects. The order flow variable $o_{t,k}$ is cumulated between month $t - k$ and t ; the forward discount is $f d_t^k = f_t^k - s_t$, where f_t^k and s_t are the log of the forward rate (for maturity k) and the spot rate observed at the beginning of month t ; $ImpVol_t^k$ denotes the k months ahead exchange rate variance at the end of month t . Sample: Jan 2000 – Apr 2007. The column "OF" reports the coefficient β_o^k (with t -statistics below); the column "Implied" reports the implied Fama's β ($1 + \beta_o^k + \beta_u^k + \beta_v^k$) and in squared brackets is the p -value from the J -test of the overidentifying restriction (that the implied β is equal to Fama's beta). A † indicates that $\beta_o^k + \beta_u^k + \beta_v^k$ is not significantly different from zero at the 5% level (i.e. UIP cannot be rejected).

	EUR/USD		USD/JPY		GBP/USD	
	Implied	OF	Implied	OF	Implied	OF
Panel A: Implied Volatility						
1 Month	-4.46 [0.78]	-4.30 (-3.55)	-0.79 [0.37]	-1.18 (-1.45)	-1.26† [0.37]	-0.64 (-0.57)
3 Month	-4.73 [0.14]	-3.19 (-3.45)	-1.30 [0.56]	-1.59 (-1.89)	-1.87† [0.54]	-0.74 (-1.03)
6 Month	-5.33 [0.18]	-3.29 (-4.66)	-1.76 [1.00]	-1.16 (-3.75)	-2.09† [0.55]	-0.41 (-0.70)
12 Month	-5.79 [0.16]	-3.16 (-3.90)	-1.98 [0.36]	-1.25 (-6.60)	-2.96 [0.28]	0.21 (0.52)
Panel B: Forecast Dispersion						
1 Month	-4.23 [0.65]	-4.26 (-3.18)	-0.86† [0.30]	-0.97 (-1.18)	-1.13† [0.35]	-0.60 (-0.53)
3 Month	-4.72 [0.14]	-3.34 (-3.37)	-1.12 [0.34]	-1.89 (-2.05)	-1.95† [0.57]	-0.68 (-0.89)
6 Month	-5.29 [0.18]	-3.35 (-4.60)	-1.66 [0.78]	-1.36 (-3.54)	-2.12† [0.53]	-0.42 (-0.69)
12 Month	-5.78 [0.16]	-3.18 (-3.86)	-1.96 [0.38]	-1.26 (-6.80)	-3.00 [0.29]	0.22 (0.54)

References

- Akram, Q.F., Rime, D., Sarno, L., 2008. Arbitrage in the foreign exchange market: Turning on the microscope. *Journal of International Economics* 76, 237–253.
- Bacchetta, P., Mertens, E., van Wincoop, E., 2008. Predictability in financial markets: What do survey expectations tell us? *Journal of International Money and Finance* 28, 406–426.
- Bacchetta, P., van Wincoop, E., 2006. Can information heterogeneity explain the exchange rate determination puzzle? *American Economic Review* 96, 552–576.
- Bacchetta, P., van Wincoop, E., 2009. Infrequent portfolio decisions: A solution to the forward discount puzzle. *American Economic Review* Forthcoming.
- Bansal, R., Shaliastovich, I., 2007. Risk and return in bond, currency and equity markets. typescript, Duke University.
- Beber, A., Breedon, F., Buraschi, A., 2010. Differences in beliefs and currency risk premiums. *Journal of Financial Economics* 98, 415–438.
- Bekaert, G., Hodrick, R.J., Marshall, D., 1997. The implications of first-order risk aversion for asset market risk premiums. *Journal of Monetary Economics* 40, 3–39.
- Berger, D.W., Chaboud, A.P., Chernenko, S.V., Howorka, E., Wright, J.H., 2008. Order flow and exchange rate dynamics in electronic brokerage system data. *Journal of International Economics* 75, 93–109.
- Bilson, J.F.O., 1981. The “speculative efficiency” hypothesis. *Journal of Business* 54, 435–451.
- Bjønnes, G.H., Rime, D., 2005. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75, 571–605.
- Breedon, F., 2001. Market liquidity under stress: Observations from the fx market. Bank of International Settlements: Conference Proceedings.
- Breedon, F., Ranaldo, A., 2008. Time-of-day patterns in fx returns and order flow. mimeo, Imperial College Business School.
- Breedon, F., Vitale, P., 2010. An empirical study of portfolio-balance and information effects of order flow on exchange rates. *Journal of International Money and Finance* 29, 504–524.

- Brunnermeier, M.K., Nagel, S., Pedersen, L.H., 2009. Carry trades and currency crashes. In: D. Acemoglu, K. Rogoff, M. Woodford (Eds.), *NBER Macroeconomics Annual 2008*, vol. 23, 313–347. Cambridge, MA: MIT Press.
- Burnside, C., Eichenbaum, M., Rebelo, S., 2007. The returns to currency speculation in emerging markets. *American Economic Review Papers and Proceedings* 97, 333–338.
- Burnside, C., Eichenbaum, M.S., Rebelo, S., 2009. Understanding the forward premium puzzle: A microstructure approach. *American Economic Journal: Macroeconomics* Forthcoming.
- Cavaglia, S.M.F.G., Verschoor, W.F.C., Wolff, C.C.P., 1994. On the biasedness of forward foreign exchange rates: Irrationality or risk premia? *Journal of Business* 67, 321–343.
- Chinn, M., Frankel, J.A., 2002. Survey data on exchange rate expectations: More currencies, more horizons, more tests. In: W. Allen, D. Dickinson (Eds.), *Monetary Policy, Capital Flows and Financial Market Developments in the Era of Financial Globalization: Essays in Honor of Max Fry*. London: Routledge.
- Cumby, R.E., 1988. Is it risk? explaining deviations from uncovered interest rate parity. *Journal of Monetary Economics* 22, 279–300.
- Daniélsson, J., Love, R., 2006. Feedback trading. *International Journal of Finance and Economics* 11, 35–53.
- Engel, C., 1996. The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of Empirical Finance* 3, 123–191.
- Evans, M.D.D., Lewis, K.K., 1995. Do long-term swings in the dollar affect estimates of the risk premia? *Review of Financial Studies* 8, 709–742.
- Evans, M.D.D., Lyons, R.K., 2002. Order flow and exchange rate dynamics. *Journal of Political Economy* 110, 170–180.
- Evans, M.D.D., Lyons, R.K., 2005. Meese-rogooff redux: Micro-based exchange-rate forecasting. *American Economic Review Papers and Proceedings* 95, 405–414.
- Evans, M.D., Hnatkovska, V.V., 2007. International financial integration and the real economy. *IMF Staff Papers* 54, 220–269.
- Fama, E.F., 1984. Forward and spot exchange rates. *Journal of Monetary Economics* 14, 319–338.

- Frankel, J.A., Chinn, M.D., 1993. Exchange rate expectations and the risk premium: Tests for a cross section of 17 countries. *Review of International Economics* 1, 136–144.
- Froot, K.A., Frankel, J.A., 1989. Forward discount bias: Is it an exchange risk premium? *Quarterly Journal of Economics* 104, 139–161.
- Froot, K.A., Ramadorai, T., 2005. Currency returns, intrinsic value, and institutional-investor flows. *Journal of Finance* 60, 1535–1566.
- Froot, K.A., Thaler, R.T., 1990. Anomalies: Foreign exchange. *Journal of Economic Perspectives* 4, 179–192.
- Galati, G., Heath, A., McGuire, P., 2007. Evidence on carry trade activity. *BIS Quarterly Review* 27–41.
- Hodrick, R.J., 1989. Risk, uncertainty and exchange rates. *Journal of Monetary Economics* 23, 433–459.
- Ilut, C.L., 2009. Ambiguity aversion: Implications for the uncovered interest rate parity puzzle. typescript, Northwestern University.
- Ito, T., 1990. Foreign exchange expectations: Micro survey data. *American Economic Review* 80, 434–439.
- Jongen, R., Verschoor, W.F., Wolff, C.C., 2008. Foreign exchange rate expectations: Survey and synthesis. *Journal of Economic Surveys* 22, 140–165.
- Jylhä, P., Suominen, M., 2010. Speculative capital and currency carry trade. *Journal of Financial Economics* Forthcoming.
- Killeen, W.P., Lyons, R.K., Moore, M.J., 2006. Fixed versus flexible: Lessons from EMS order flow. *Journal of International Money and Finance* 25, 551–579.
- Lewis, K.K., 1989a. Can learning affect exchange rate behavior? the case of the dollar in the early 1980s. *Journal of Monetary Economics* 23, 79–100.
- Lewis, K.K., 1989b. Changing beliefs and systematic rational forecast errors with evidence from foreign exchange. *American Economic Review* 79, 621–636.
- Lewis, K.K., 1995. Puzzles in international financial markets. In: G.M. Grossman, K. Rogoff (Eds.), *Handbook of International Economics*, vol. 3, chap. 37, 1913–71. Amsterdam: North Holland.

- Lustig, H., Verdelhan, A., 2007. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97, 89–117.
- Lustig, H.N., Roussanov, N.L., Verdelhan, A., 2009. Common risk factors in currency markets. typescript, UCLA.
- Lyons, R.K., 1997. A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics* 275–298.
- Lyons, R.K., 2001. *The Microstructure Approach to Exchange Rates*. Cambridge, MA: MIT Press.
- Moore, M.J., Roche, M.J., 2010. Solving exchange rate puzzles with neither sticky prices nor trade costs. *Journal of International Money and Finance* 29, 1151–1170.
- Payne, R., 2003. Informed trade in spot foreign exchange markets: An empirical investigation. *Journal of International Economics* 61, 307–329.
- Plantin, G., Shin, H.H., 2008. Carry trades and speculative dynamics. mimeo, Princeton University.
- Ranaldo, A., Sarkar, A., 2008. Exchange rate risk, transactions costs and the forward bias puzzle. typescript, Swiss National Bank.
- Rime, D., Sarno, L., Sojli, E., 2010. Exchange rate forecasting, order flow and macroeconomic information. *Journal of International Economics* 80, 72–88.
- Verdelhan, A., 2010. A habit-based explanation of the exchange rate risk premium. *Journal of Finance* 65, 123–146.