The Forward Volatility Bias in Foreign Exchange*

Pasquale Della Corte[a] Lucio Sarno[b,c] Ilias Tsiakas[a]

a: Warwick Business School

b: Cass Business School

c: Centre for Economic Policy Research (CEPR)

February 2009

Abstract

This paper investigates the relation between spot and forward implied volatility in foreign exchange by formulating and testing the forward volatility unbiasedness hypothesis. This new hypothesis provides a framework for testing the volatility analogue to the extensively researched hypothesis of unbiasedness in forward exchange rates. Using a new data set of spot implied volatility quoted on over-the-counter currency options, we compute the forward implied volatility that corresponds to the forward contract on future spot implied volatility known as a forward volatility agreement. We find statistically significant evidence that forward implied volatility is a systematically biased predictor that overestimates future spot implied volatility. The bias in forward volatility generates high economic value to an investor exploiting predictability in the returns to volatility speculation and indicates the presence of predictable volatility term premiums in foreign exchange.

Keywords: Implied Volatility; Foreign Exchange; Forward Volatility Agreement; Unbiasedness; Volatility Speculation.

JEL Classification: F31; F37; G10; G11.

*__Acknowledgements:__ The authors are indebted for useful conversations or constructive comments to seminar participants at the University of Wisconsin-Madison, HEC Paris, University of Amsterdam, Vienna Graduate School of Finance, Universitat Pompeu Fabra, University of Brescia as well as to John Bilson, Peter Carr, Philippe Jorion, Anthony Neuberger, Adrien Verdelhan and Liuren Wu. We also thank Bilal Hafeez of Deutsche Bank and several foreign exchange traders, particularly Lee Sanders of AXA Investment Managers. We are especially grateful to Deutsche Bank for providing the data on FX implied volatilities. The usual disclaimer applies. *Corresponding author*: Lucio Sarno, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK. E-mail: lucio.sarno@city.ac.uk. *Other authors' contact details:* Pasquale Della Corte, pasquale.dellacorte@wbs.ac.uk; Ilias Tsiakas, ilias.tsiakas@wbs.ac.uk.

1 Introduction

The forward bias puzzle is a widely accepted empirical rejection of the Uncovered Interest Parity (UIP) condition suggesting that forward exchange rates are a biased predictor of future spot exchange rates (e.g., Bilson 1981; Fama, 1984; Backus, Gregory and Telmer, 1993; Engel, 1996; and Backus, Foresi and Telmer, 2001). In practice, this means that high interest rate currencies tend to appreciate rather than depreciate. The forward bias also implies that the returns to currency speculation are predictable, which tends to generate high economic value to an investor designing dynamic allocation strategies exploiting the UIP violation (Burnside, Eichenbaum, Kleshchelski and Rebelo, 2008; and Della Corte, Sarno and Tsiakas, 2009). This is manifested by the widespread use of carry trade strategies in foreign exchange (FX) exploiting the forward bias anomaly (e.g., Galati and Melvin, 2004; and Brunnermeier, Nagel and Pedersen, 2008).

A recent development in FX trading is the ability of investors to engage not only in spot-forward currency speculation but also in spot-forward volatility speculation. This has become possible by trading a contract called the forward volatility agreement (FVA). The FVA is a forward contract on future spot implied volatility, which for each dollar investment delivers the difference between future spot implied volatility and forward implied volatility. Therefore, given today's information, the FVA implicitly determines the forward implied volatility for an interval starting at a future date. Investing in FVAs allows investors to hedge volatility risk and speculate on the level of future volatility.

This is the first paper to investigate the relation between spot and forward implied volatility in foreign exchange by formulating and testing a new hypothesis: the forward volatility unbiasedness hypothesis (FVUH). Our analysis uses a new data set of daily implied volatilities for seven US dollar exchange rates quoted on over-the-counter (OTC) currency options spanning up to 18 years of data.[1] Using the data on spot implied volatility for different maturities, we compute the forward implied volatility that represents the delivery price of an FVA. In order to test the empirical validity of the FVUH, we estimate the volatility analogues to the Fama (1984) predictive regressions. The results provide statistically significant evidence that forward implied volatility is a systematically biased predictor that overestimates future spot implied volatility. This finding is similar to two well-known tendencies: (i) of forward premiums to overestimate the future rate of depreciation (appreciation) of high (low) interest rate currencies; and (ii) of spot implied volatility to overestimate future realized volatility (e.g., Poon and Granger, 2003). Furthermore, the rejection of the forward volatility unbiasedness establishes the presence of non-zero, time-varying and predictable volatility term premiums in foreign exchange.

We assess the economic value of the forward volatility bias in the context of dynamic asset

[1] See, for example, Jorion (1995) for a study of the information content and predictive ability of implied FX volatility derived from options traded on the Chicago Mercantile Exchange.

allocation by designing a volatility speculation strategy. This is a dynamic strategy that exploits predictability in the returns to volatility speculation and, in essence, it implements the carry trade not for currencies but for implied volatilities. The motivation for the "carry trade in volatility" strategy is straightforward: if there is a forward volatility bias, then buying (selling) FVAs when forward implied volatility is higher (lower) than spot implied volatility will consistently generate excess returns over time. Our findings reveal that the in-sample and out-of-sample economic value of the forward volatility bias is high and robust to reasonable transaction costs. Furthermore, the returns to volatility speculation (carry trade in volatility) are largely uncorrelated with the returns to currency speculation (carry trade in currency), which suggests that the source of the forward volatility bias is not related to that of the forward bias. In short, therefore, we find robust statistical and economic evidence establishing the forward volatility bias.

An essential aspect of the analysis is the economic evaluation of departures from forward volatility unbiasedness. A purely statistical rejection of the FVUH does not guarantee that an investor can enjoy tangible economic gains from implementing the carry trade in volatility strategy that exploits predictability in the returns to volatility speculation. This motivates a dynamic asset allocation approach based on standard mean-variance analysis, which is in line with previous studies on volatility timing by West, Edison and Cho (1993), Fleming, Kirby and Ostdiek (2001), Marquering and Verbeek (2004) and Han (2006) among others. The prime objective of the economic evaluation is to measure how much a risk-averse investor is willing to pay for switching from a static portfolio strategy based on forward volatility unbiasedness to a dynamic strategy exploiting the systematic bias in the way the market sets forward implied volatility.

In this context, the main objective of this paper is to analyze the defining characteristics of the empirical relation between forward and spot implied FX volatility. Therefore, a number of questions fall beyond the scope of the analysis. First, we are not testing whether implied volatility is an unbiased forecast of future realized volatility (e.g., Jorion, 1995). Second, we do not aim at offering a theoretical explanation for the forward volatility bias. After all, there is no consensus on the main economic determinants of volatility. Finally, we do not make a conclusive statement on the efficiency of the currency options market. Forward prices may not be equal to expected future spot prices because of transaction costs, information costs and risk aversion (e.g., Engel, 1996). In short, therefore, the main purpose of this paper is confined to establishing the first statistical and economic evidence on the forward volatility bias in foreign exchange.

The remainder of the paper is organized as follows. In the next section we briefly review the literature on the forward unbiasedness hypothesis in exchange rates. Section 3 proposes the forward volatility unbiasedness hypothesis and the empirical results are reported in Section 4. In Section 5 we present the framework for assessing the economic value of departures from forward volatility unbiasedness for an investor with a dynamic carry trade in volatility strategy. The findings on the economic value of the forward volatility bias are discussed in Section 6, followed by robustness checks and further analysis in Section 7. Finally, Section 8 concludes.

2 The Forward Unbiasedness Hypothesis

The forward unbiasedness hypothesis (FUH) in foreign exchange is also known as the speculative efficiency hypothesis (Bilson, 1981). The FUH simply states that the forward exchange rate should be an unbiased predictor of the future spot exchange rate:

$$F_t^k = E_t S_{t+k}, \tag{1}$$

where S_{t+k} is the nominal exchange rate (defined as the domestic price of foreign currency) at time $t+k$, E_t is the expectations operator as of time t, and F_t^k is the k-period forward exchange rate at time t (i.e., the rate agreed now for an exchange of currencies in k periods).

The economic foundation of the FUH lies in the Uncovered Interest Parity (UIP) condition. Assuming risk neutrality and rational expectations, for a k-period horizon, UIP is represented by the following equation:

$$E_t s_{t+k} - s_t = i_t - i_t^*, \tag{2}$$

where $s_{t+k} = \ln(S_{t+k})$, i_t and i_t^* are the k-period domestic and foreign nominal interest rates respectively. In the absence of riskless arbitrage, Covered Interest Parity (CIP) also holds: $f_t^k - s_t = i_t - i_t^*$, where $f_t^k = \ln F_t^k$. It is straightforward to use UIP and CIP to derive the FUH by substituting the interest rate differential $i_t - i_t^*$ in Equation (2) by the forward premium $f_t^k - s_t$.

In performing empirical tests of the FUH, the majority of the literature estimates the following regression, which is commonly referred to as the "Fama regression" (Fama, 1984):

$$s_{t+k} - s_t = a + b\left(f_t^k - s_t\right) + u_{t+k}. \tag{3}$$

The Fama regression is used to determine whether the current forward premium $\left(f_t^k - s_t\right)$ is an unbiased predictor of the future spot exchange rate return $(s_{t+k} - s_t)$. If the FUH holds, we should find that $a = 0$, $b = 1$, and the disturbance term $\{u_{t+k}\}$ is serially uncorrelated.

In his seminal work, Fama (1984) also run a second regression:

$$s_{t+k} - f_t^k = a + c\left(f_t^k - s_t\right) + u_{t+k}. \tag{4}$$

This second regression is used to test whether the current forward premium can predict the return to currency speculation, i.e., the return from issuing a forward contract at time t and converting the proceeds into dollars at the prevailing spot rate at $t+k$, or vice versa (e.g., Backus, Gregory and Telmer, 1993). Even though the empirical finance literature has primarily focused on the first regression (Equation 3), the two Fama regressions contain exactly the same information. Since the left-hand-side of the difference in the two regressions is equal to the predictive variable, i.e., $(s_{t+k} - s_t) - \left(s_{t+k} - f_t^k\right) = f_t^k - s_t$, then it must be that the intercepts (a) are equal, the slope coefficients are related by $b - c = 1$, and the innovations $\{u_{t+k}\}$ in each time period are equal across the two regressions. In other words, if the forward premium is an unbiased predictor of exchange rate returns ($a = 0; b = 1$), then the excess returns to currency speculation should be unpredictable ($a = 0; c = 0$).

Since the contribution of Bilson (1981) and Fama (1984), numerous empirical studies consistently reject the UIP condition (e.g., Hodrick, 1987; Engel, 1996; Engel, Mark and West, 2007). As a result, it is a stylized fact that estimates of b tend to be closer to minus unity than plus unity. This is commonly referred to as the "forward bias puzzle," and implies that high-interest currencies tend to appreciate rather than depreciate, which is the basis of widely used carry trade strategies. In general, attempts to explain the forward bias puzzle using a variety of models have met with mixed success. Therefore, the forward bias remains a puzzle in international finance research.[2]

3 The Forward Volatility Unbiasedness Hypothesis

In this section, we turn our attention to the FX implied volatility (IV) market. In what follows, we set up a framework for testing forward volatility unbiasedness that is analogous to the framework used for testing forward unbiasedness in the traditional FX market.

3.1 Forward Volatility Agreements

The forward implied volatility of exchange rate returns is determined by a forward volatility agreement (FVA). The FVA is a forward contract on future spot implied volatility with a payoff at maturity equal to:

$$\left(\Sigma_{t+k} - \Phi_t^k\right) M, \tag{5}$$

where Σ_{t+k} is the annualized spot implied volatility observed at time $t + k$ and measured over a set interval (e.g., from $t + k$ to $t + 2k$), Φ_t^k is the annualized forward implied volatility determined at time t for the same interval starting at time $t + k$, and M denotes the notional dollar amount that converts the volatility difference into a dollar payoff. For example, setting $k = 3$ months implies that Σ_{t+3} is the observed spot IV at time $t+3$ months for the interval of $t+3$ months to $t+6$ months; and Φ_t^3 is the forward IV determined at time t for the interval of $t+3$ months to $t+6$ months. The FVA

[2] See, for example, Backus, Gregory and Telmer (1993); Bekaert (1996); Bansal (1997); Bekaert, Hodrick and Marshall (1997); Backus, Foresi and Telmer (2001); Bekaert and Hodrick (2001); Lustig and Verdelhan (2007); Brunnermeier, Nagel and Pedersen (2008); Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009); and Verdelhan (2009).

allows investors to hedge volatility risk and speculate on the level of future spot IV by determining the expected value of IV over an interval starting at a future date.

3.2 The Forward Volatility Unbiasedness Hypothesis

The FVA forward contract has zero net market value at entry. No-arbitrage dictates that Φ_t^k must be equal to the risk-neutral expected value of Σ_{t+k}:

$$\Phi_t^k = E_t \Sigma_{t+k}. \tag{6}$$

As the expected value of future spot volatility, Φ_t^k is the forward volatility and should be an unbiased estimator of Σ_{t+k}. Then, Equation (6) above defines the Forward Volatility Unbiasedness Hypothesis (FVUH), which postulates that forward IV, determined conditional on today's information set, is the optimal predictor of future spot IV over the relevant horizon. The FVUH is based on risk neutrality and rational expectations, and therefore, can be thought of as the second-moment analogue of the FUH, which is based on the same set of assumptions.[3]

3.3 Forward Implied Volatility

Forward IV is determined by the term structure of spot IV under the same assumptions of risk neutrality and rational expectations that underpin the FVUH. Define $\Sigma_{t,t+k}$ and $\Sigma_{t,t+2k}$ as the annualized implied volatilities for the intervals t to $t+k$ and t to $2k$, respectively. The forward implied volatility determined at time t for an interval starting at time $t+k$ and ending at $t+2k$ is given by (see, for example, Poterba and Summers, 1986; and Carr and Wu, 2009):

$$\Phi_t^k = \sqrt{2\Sigma_{t,t+2k}^2 - \Sigma_{t,t+k}^2}. \tag{7}$$

Intuitively, Equation (7) indicates that the 6-month spot implied variance is a simple average of the 3-month spot implied variance and the 3-month forward implied variance. This is due to the linear relation between implied variance and time across the term structure.[4] This linear method is widely used by investment banks in setting forward implied volatility. It is also equivalent to the expectations hypothesis of the term structure of implied variance (Campa and Chang, 1995).[5]

[3] At entry, the FVA must have zero value due to no-arbitrage: $E_t\left[e^{-i_t}\left(\Sigma_{t+k} - \Phi_t^k\right)\right] = 0$, where i_t is the k-period domestic interest rate. Hence forward IV is the risk-neutral expected value of future spot IV.

[4] Equation (7) is the only case where we have spot volatility defined over different intervals and therefore we need to use two subscripts to clearly identify the start and end of the interval. From now on, we revert back to using a single subscript where for example Σ_{t+k} is the annualized implied volatility observed at time $t+k$ and measured over a set interval with length k.

[5] The intuition behind Equation (7) is straightforward. By definition, variance is additive in the time dimension, and so is expected variance. Under standard no-arbitrage arguments, implied variance is the risk-neutral expected value of realized variance. It follows that forward implied variance is a linear combination of spot implied variances, which is the basis of the linear method. Furthermore, note that Jensen's inequality causes a convexity bias since by approximation we assume that the square root of expected (implied) variance is equal to expected (implied) volatility. We measure

3.4 Predictive Regressions for Exchange Rate Volatility

In order to test the empirical validity of the FVUH, we estimate the volatility analogues to the two Fama (1984) regressions:

$$\sigma_{t+k} - \sigma_t = \alpha + \beta \left(\varphi_t^k - \sigma_t\right) + \varepsilon_{t+k} \quad (8)$$

$$\sigma_{t+k} - \varphi_t^k = \alpha + \gamma \left(\varphi_t^k - \sigma_t\right) + \varepsilon_{t+k}, \quad (9)$$

where $\sigma_t = \ln(\Sigma_t)$, $\sigma_{t+k} = \ln(\Sigma_{t+k})$, and $\varphi_t^k = \ln\left(\Phi_t^k\right)$. There is a critical difference in the way we measure log-exchange rates in Equations (3) and (4) versus log-volatilities in Equations (8) and (9). The former are observed at a given point in time but the latter are defined over an interval. Our notation is simple and allows for direct correspondence between the FX market and the FX volatility market. Furthermore, note that we run the predictive regressions using the first difference of logs as opposed to levels (e.g., the LHS is $\sigma_{t+k} - \sigma_t$ rather than Σ_{t+k}). Our motivation for this specification is based on the high persistence in the level of FX volatility (e.g., Berger, Chaboud, Hjalmarsson and Howorka, 2008). This is an important consideration since performing OLS estimation on very persistent variables (such as volatility levels) is problematic and can cause spurious results, whereas performing OLS estimation on volatility returns avoids this concern. The same issue arises in the traditional FX market, which explains why the standard Fama regressions are run on exchange rate returns, not exchange rate levels.

There is a natural interpretation for each term in the two regressions above. Specifically, we refer to the term $(\sigma_{t+k} - \sigma_t)$ as the "implied volatility return, " $(\varphi_t^k - \sigma_t)$ as the "forward volatility premium," and $(\sigma_{t+k} - \varphi_t^k)$ as the "excess volatility return." The excess volatility return is the return to volatility speculation and is equal to the return of the FVA.[6]

As in the original Fama regressions, the intercepts (α) in the two regressions above are equal, $\beta - \gamma = 1$, and the innovations $\{\varepsilon_{t+k}\}$ in each time period are equal across the two regressions. Then, $\alpha = 0, \beta = 1$ (no forward volatility bias) implies $\alpha = 0, \gamma = 0$ (no predictability in the excess volatility return). Therefore, a non-zero γ coefficient suggests that the return to volatility speculation (or the return of the FVA) is predictable.

This framework leads to two distinct empirical models for testing the FVUH. The first model simply imposes forward volatility unbiasedness by setting $\alpha = 0, \beta = 1$ in Equation (8) (or equiv-

this bias using a second-order Taylor expansion as in Brockhaus and Long (2000), and find that it is empirically negligible. We also test the FVUH for the forward implied variance (instead of the forward implied volatility) to avoid any convexity bias, and we find no qualitative change in any the findings described below. Further details are available upon request.

[6] For a \$1 investment, the payoff of an FVA is $\Sigma_{t+k} - \Phi_t^k$. Then, the "capital gain" is $\Sigma_{t+k} - \Phi_t^k$, and the FVA "excess return" is $\sigma_{t+k} - \varphi_t^k$, where $\sigma_{t+k} = \ln(\Sigma_{t+k})$ and $\varphi_t^k = \ln\left(\Phi_t^k\right)$. Since delivery takes place at time $t + k$, but Φ_t^k is determined at time t, an investor committing to an FVA can borrow at t an amount $\Phi_t^k e^{-i_t}$ so that the total k-period return is: $\ln(\Sigma_{t+k}) - \ln\left(\Phi_t^k e^{-i_t}\right) = \sigma_{t+k} - \varphi_t^k + i_t$, where i_t is the k-period domestic interest rate. Hence $\sigma_{t+k} - \varphi_t^k$ is the excess return.

alently $\alpha = 0, \gamma = 0$ in Equation (9)). This will be the benchmark model in our analysis and we refer to it as the FVUH model. The second model estimates $\{\alpha, \beta\}$ in Equation (8) (or equivalently $\{\alpha, \gamma\}$ in Equation 9) and uses the parameter estimates to predict the implied volatility returns (or equivalently the excess volatility returns). We refer to the second model as the Forward Volatility Regression (FVR). We assess the statistical and economic significance of possible deviations from the FVUH simply by comparing the performance of the FVUH model with the FVR model under a variety of metrics described later.

4 Empirical Results on Forward Volatility Unbiasedness

4.1 Spot and Forward FX Implied Volatility Data

Our analysis employs a unique data set of daily at-the-money-forward (ATMF) implied volatilities quoted on over-the-counter (OTC) currency options. The data are collected by Reuters from a panel of market participants and were made available to us by Deutsche Bank. These are high quality data involving quotes for contracts of at least $10 million with a prime counterparty. In general, the OTC currency options market is by far the largest and most liquid market of its kind.[7] Therefore, OTC implied volatilities are considered to be of higher quality than those derived from options traded in a particular exchange (e.g., Jorion, 1995).

The IV data sample focuses on seven exchange rates relative to the US dollar: the Australian dollar (AUD), the Canadian dollar (CAD), the Swiss franc (CHF), the Euro (EUR), the British pound (GBP), the Japanese yen (JPY) and the New Zealand dollar (NZD). The end date of the sample is July 11, 2008 for all currencies, but the start date of the sample varies across currencies: January 2, 1991 for AUD and JPY (4416 daily observations), January 2, 1992 for GBP (4162 obs.), January 4, 1993 for CHF (3908 obs.), January 2, 1997 for CAD (2899 obs.), January 16, 1998 for NZD (2637 obs.) and January 4, 1999 for EUR (2396 obs.). Hence the daily data sample ranges from 9.5 to 17.5 years.[8] Finally, our analysis excludes all trading days which occur on a national US holiday.

For each day of the sample, we use information on the 3-month ($3m$), 6-month ($6m$) and 12-month ($12m$) implied volatilities. Using these IV maturities, we construct the forward implied volatilities for $3m$ and $6m$ using Equation (7). Hence our analysis focuses on the relation between spot and forward IV across the $3m$ and $6m$ maturities. For a general discussion of the stylized features of currency option implied volatilities see Jorion (1995) and Carr and Wu (2007).

[7] More generally, the FX market is the largest financial market in the world with an average daily volume of transactions exceeding $3.2 trillion. The average daily turnover of the FX options market is over $200 billion (see Bank of International Settlements, 2007).

[8] A shorter sample of these data starting in September 2001 that is virtually identical for the overlapping period is publicly available on the website of the British Bankers' Association.

Table 1 provides a brief description of the daily spot and forward IV data in annualized percent terms. The mean of the spot and forward IV level is similar across currencies and maturities revolving around 10% per annum with a standard deviation of about 2% per annum. In most cases, IV levels exhibit positive skewness, no excess kurtosis and are highly serially correlated, even at very long lags.[9] Furthermore, the augmented Dickey-Fuller (ADF) statistic indicates that volatility levels are not stationary, which contradicts the widely accepted view that volatility is a highly persistent but stationary process. This apparent inconsistency may be explained by the fact that the ADF statistic has low power and may not reject non-stationarity when applied to a near-unit root process. In contrast, as we will see below, the evidence on the stationarity of volatility returns is unambiguous.

Table 2 reports descriptive statistics for the three types of daily volatility returns: the implied volatility return $(\sigma_{t+k} - \sigma_t)$, the forward volatility premium $(\varphi_t^k - \sigma_t)$, and the excess volatility return $(\sigma_{t+k} - \varphi_t^k)$. The table summarizes the statistics in annualized percent units showing that the mean volatility returns revolve between -10% and $+10\%$ for a high standard deviation in the range of $10\% - 30\%$. In most cases, the volatility returns exhibit low skewness (positive or negative), low excess kurtosis and high serial correlation with short memory. More importantly, the ADF statistic now rejects the null hypothesis of non-stationarity with high confidence. This provides a clear justification for running the predictive regressions (Equations 8 and 9) on volatility returns rather than on volatility levels since there is statistical evidence that the former are stationary but the latter are not.

A first indication of the performance of forward IV as a predictor of future spot IV is illustrated in Figure 1. The figure plots the daily time series of the $3m$ spot and forward IV level for all currencies and makes it visually apparent that the spot and forward IV levels do not move closely with each other. A second indication of the same result can be seen in Figure 2, which shows a scatter plot of the $3m$ IV return relative to the $3m$ forward volatility premium. At first glance, the two variables are far from having a linear relation along the 45-degree line implied by the FVUH.

4.2 Predictive Regression Results

We begin testing the empirical validity of the FVUH by estimating the two forward volatility regressions (Equations 8 and 9). The OLS parameter estimates are reported in Table 3 using volatility returns which are measured over 3-months and 6-months but are observed and estimated daily. This overlapping structure causes the regression errors to have a moving average component. We correct for this effect by computing Newey-West standard errors.

Recall that for the FVUH to hold (and hence for forward IV to be an unbiased estimator of future

[9] It is also interesting to note that on average four currencies display an upward sloping term structure (CHF, EUR, GBP and JPY), whereas three currencies exhibit a downward sloping term structure (AUD, CAD and NZD).

spot IV) three conditions must be met in the FVR regressions: the intercept should be zero ($\alpha = 0$), the slope should be $\beta = 1$ or $\gamma = 0$, and the disturbance term should be serially uncorrelated. We test the FVUH conditions on the parameter estimates both separately with appropriately defined t-statistics and jointly with an F-statistic. The serial correlation in the error term is tested with a Box-Ljung statistic. To facilitate interpretation we also report p-values in all cases.

We first examine the slope estimate of the two forward volatility regressions. For $3m$, we find that the OLS estimates of β are all positive but much lower than unity, ranging from 0.029 (AUD) to 0.541 (JPY). Equivalently, γ ranges from -0.971 to -0.458. For $6m$, the OLS estimates of β range from -0.574 (EUR) to 0.948 (CAD), while γ ranges from -1.574 to -0.052. Overall, in 13 of the 14 cases β is statistically different from unity with at least 99% confidence as indicated by the t-statistics and p-values. The only exception is the $6m$ CAD.

Turning to the intercept of the FVR regression, we find that the value of α consistently revolves around zero (positive or negative) and in most cases it is statistically insignificant. Overall, the F-statistic jointly testing $\{\alpha = 0, \beta = 1\}$ strongly rejects unbiasedness for all but the $6m$ CAD with at least 99% confidence, i.e., the p-values are less than 1% in 13 of 14 cases. Furthermore, the evidence on the serial correlation of innovations is mixed as for only about half of the cases there is significant autocorrelation as shown by the Box-Ljung statistic and the p-values. Finally, the R^2 coefficient of the first FVR regression (Equation 8) is generally low ranging from 1% to 18%, being under 5% two thirds of the time. In contrast, the R^2 for excess volatility returns (Equation 9) is much higher ranging from 3% to 35% with more than half of the time being over 10%.

These empirical results are based on OLS estimation of the predictive regression parameters, which is commonly used in similar studies of the forward bias in the traditional FX market. However, least squares estimation fails to deliver unbiased estimates when the disturbances contain outliers or when the predictive variable is observed with error. These are potentially important issues in determining the reliability of the OLS estimates. Therefore, in addition to OLS, we perform least absolute deviations (LAD) estimation, which is robust to thick-tailed error distributions and is not sensitive to outliers (e.g., Bassett and Koenker, 1978). We also perform errors-in-variables estimation (EIV) based on maximum likelihood and the Kalman filter to estimate the parameters when the explanatory variable is measured with error (e.g., Carr and Wu, 2009). We implement EIV assuming that the forward volatility premium is observed with error and the true value follows an AR(1) process. The OLS, LAD and EIV estimates for β are displayed in Table 4, which shows that β remains very similar in size, sign and statistical significance across the three estimation methods. Hence the rest of our analysis uses the OLS parameter estimates.

In conclusion, the predictive regression results clearly demonstrate that forward IV is a biased predictor of future spot IV. Consequently, the results lead to a firm statistical rejection of the

FVUH suggesting that predictable returns can be generated from FX volatility speculation. In other words, the statistical evidence indicates that in addition to the well established forward bias in the traditional FX market, there is also a forward volatility bias in the implied volatilities quoted on currency options. There is, however, a difference in the bias observed in the two markets. In the FX market, b tends to be negative and is often statistically insignificant. In the FX IV market, β tends to be mildly positive and statistically significant. Hence the bias in forward FX volatility is less severe than the bias in forward exchange rates.

4.3 A Volatility Term Premium Interpretation

Forward implied volatility represents the risk-neutral expected value of future spot implied volatility. The difference between spot and forward implied volatility reflects the volatility term premium defined as the conditional expectation of the return to volatility speculation: $E_t\left[\sigma_{t+k} - \varphi_t^k\right]$. Under the FVUH, the volatility term premium should be equal to zero. However, a rejection of the risk-neutral forward volatility unbiasedness may signify the presence of a premium in the term structure of implied FX volatility.[10]

Our results so far establish two main empirical properties for the volatility term premium in the FX market. First, the descriptive statistics in Table 2 show that unconditionally the average volatility term premium is non-zero and can be either positive or negative. In fact, the volatility term premium exhibits a large cross-sectional variation since for 3-month contracts it ranges from -19.3% to 18.3% per annum, whereas for 6-month contracts it ranges from -8.9% to 10.4% per annum. Second, the predictive regression results demonstrate that the volatility term premium is time-varying and predictable when conditioning on the forward volatility premium. In short, therefore, our analysis indicates the presence of non-zero, time-varying and predictable volatility term premiums in foreign exchange. We will revisit this issue later.

5 Economic Value of Volatility Speculation: The Framework

This section discusses the framework we use in order to evaluate the performance of the carry trade in volatility strategy, which exploits predictability in the returns to FX volatility speculation.

[10] Similarly, Carr and Wu (2009) define the volatility risk premium as the difference between realized and implied volatility. Bollerslev, Tauchen and Zhou (2008) find that this volatility risk premium can explain a large part of the time variation in stock returns. A likely explanation of this finding is that the volatility risk premium is a proxy for time-varying risk aversion. For example, Bakshi and Madan (2006) show that the volatility risk premium may be expressed as a non-linear function of a representative agent's coefficient of relative risk aversion.

5.1 The Carry Trade in Volatility Strategy

We design a dynamic strategy for FX volatility speculation, which implements the carry trade in volatility. Consider a US investor who builds a portfolio by allocating her wealth between the domestic riskless asset and seven FVA contracts. The FVAs are written on seven US dollar nominal exchange rates: AUD, CAD, CHF, EUR, GBP, JPY and NZD. The return from domestic riskless investing is equal to the yield of a US bond proxied by the daily 3-month or 6-month US Eurodeposit rate. Note that the risky assets (buying or selling FVAs) are a zero-cost investment, and hence the investor's net balances stay in the bank and accumulate interest at the domestic riskless rate. This implies that the return from investing in each of the risky assets is equal to the domestic riskless rate plus the excess volatility return for a total return of: $i_t + \sigma_{t+k} - \varphi_t^k$.

The main objective of our analysis is to determine whether there is economic value in predicting the returns to volatility speculation due to a possible systematic bias in the way the market sets forward implied volatility. We consider two strategies for the conditional mean of the returns to volatility speculation: forward volatility unbiasedness (FVUH) and the forward volatility regression (FVR) model. Throughout the analysis we do not model the dynamics of the conditional covariance matrix of the returns to volatility speculation. Therefore, we implicitly assume that the volatility of volatility returns is constant. In this setting, the optimal weights will vary across the two models only to the extent that there are deviations from forward volatility unbiasedness. In particular, the FVR strategy exploits predictability in the returns to volatility speculation in the sense that it provides the forecast $E_t\left(\sigma_{t+k} - \varphi_t^k\right)$, which is also the volatility term premium. In contrast, the FVUH benchmark strategy is equivalent to riskless investing since fixing $\alpha = 0, \beta = 1$ (or equivalently $\alpha = 0, \gamma = 0$) implies that the conditional expectation of excess volatility returns is equal to zero: $E_t\left(\sigma_{t+k} - \varphi_t^k\right) = 0$.

The investor rebalances her portfolio on a daily basis by taking a position on FX volatility over a horizon of three or six months ahead. Hence the rebalancing frequency is not the same as the horizon over which FVA returns are measured. This is sensible for an investor who exploits the daily arrival of FVA quotes defined over alternative maturities. Each day the investor takes two steps. First, she uses the two models (FVUH and FVR) to forecast the returns to volatility speculation. Second, conditional on the forecasts, she dynamically rebalances her portfolio by computing the new optimal weights for the mean-variance strategy described below. This setup is designed to inform us whether a possible bias in forward volatility affects the performance of an allocation strategy in an economically meaningful way. We repeat this exercise for the 3-month and 6-month FVA contracts.[11]

[11] Suppose, for example, that estimation of the forward volatility regressions yields $\alpha = 0$ and $\beta < 1$ (or $\gamma < 0$). Then, on average, when $\varphi_t^k > \sigma_t$ it must be that $\varphi_t^k > \sigma_{t+k}$. Hence a strategy that buys FVAs will consistently generate excess returns over time. Conversely, when $\varphi_t^k < \sigma_t$ the strategy will sell FVAs.

We refer to the dynamic strategy implied by the FVR model as the carry trade in volatility (CTV) strategy. The dynamic CTV strategy can be thought of as the volatility analogue to the traditional carry trade in currency (CTC) strategy studied among others by Burnside *et al.* (2008) and Della Corte, Sarno and Tsiakas (2009). The only risk an investor following the carry trade in volatility strategy is exposed to is FX volatility risk. This has the profound implication that a strategy for volatility speculation may be subject to a fundamentally different source of risk than a strategy for currency speculation, hence suggesting that the returns from two such strategies may be largely uncorrelated. We will empirically explore this issue in more detail later.[12]

5.2 Mean-Variance Dynamic Asset Allocation

Mean-variance analysis is a natural framework for assessing the economic value of strategies which exploit predictability in the mean and variance. We design a maximum expected return strategy, which leads to a portfolio allocation on the efficient frontier. Consider an investor who has a 3-month or 6-month horizon. On a daily basis, the investor constructs a dynamically rebalanced portfolio that maximizes the conditional expected return subject to achieving a target conditional volatility. Computing the dynamic weights of this portfolio requires k-step ahead forecasts of the conditional mean and the conditional covariance matrix. Let r_{t+k} denote the $N \times 1$ vector of risky asset returns; $\mu_{t+k|t} = E_t\left[r_{t+k}\right]$ is the conditional expectation of r_{t+k}; and $V_{t+k|t} = E_t\left[\left(r_{t+k} - \mu_{t+k|t}\right)\left(r_{t+k} - \mu_{t+k|t}\right)'\right]$ is the conditional covariance matrix of r_{t+k}. At each period t, the investor solves the following problem:

$$\max_{w_t} \left\{\mu_{p,t+k|t} = w_t'\mu_{t+k|t} + \left(1 - w_t'\iota\right) r_f\right\}$$
$$\text{s.t.} \quad \left(\sigma_p^*\right)^2 = w_t' V_{t+k|t} w_t, \tag{10}$$

where w_t is the $N \times 1$ vector of portfolio weights on the risky assets, ι is an $N \times 1$ vector of ones, $\mu_{p,t+k|t}$ is the conditional expected return of the portfolio, σ_p^* is the target conditional volatility of the portfolio returns, and r_f is the return on the riskless asset. The solution to this optimization problem delivers the risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}} V_{t+k|t}^{-1}\left(\mu_{t+k|t} - \iota r_f\right), \tag{11}$$

where $C_t = \left(\mu_{t+k|t} - \iota r_f\right)' V_{t+k|t}^{-1}\left(\mu_{t+k|t} - \iota r_f\right)$. The weight on the riskless asset is $1 - w_t'\iota$. Then, the period $t+k$ gross return on the investor's portfolio is:

$$R_{p,t+k} = 1 + r_{p,t+k} = 1 + \left(1 - w_t'\iota\right) r_f + w_t' r_{t+k}. \tag{12}$$

[12] In the traditional CTC strategy, an investor allocates her wealth between a domestic bond and a set of foreign bonds. Then, the return on each risky asset is equal to the riskless rate plus the return to currency speculation for a total return of: $i_t + s_{t+k} - f_t^k$. Note that in the FX market, trading foreign bonds or forward exchange rates leads to the same result because of covered interest parity: $i_t^* + s_{t+k} - s_t = i_t + s_{t+k} - f_t^k$. This is an important distinction between the FX market and the FX volatility market as for the latter there is no condition equivalent to covered interest parity.

We can solve for the dynamic weights of the maximum expected return strategy described above without having to define a particular utility function. Finally, in this mean-variance strategy we assume that the volatility of volatility returns is constant: $V_{t+k|t} = \overline{V}$, where $\overline{V}$ is the unconditional covariance matrix of volatility returns.

5.3 Performance Measures

We evaluate the performance of the carry trade in volatility strategy relative to the FVUH benchmark using the West, Edison and Cho (1993) and Fleming, Kirby and Ostdiek (2001) methodology, which is based on mean-variance analysis with quadratic utility. Quadratic utility is an attractive assumption because mean-variance applies exactly and provides a high degree of analytical tractability.[13] At any point in time, one set of estimates of the returns to volatility speculation is better than a second set if investment decisions based on the first set lead to higher utility. We thus measure the economic value of the forward volatility bias using a certainty equivalent measure for the pair of FVUH and FVR portfolios. Suppose that holding the optimal portfolio based on the FVUH model yields the same average utility as holding the optimal portfolio based on the FVR model that is subject to daily expenses Π. Since the investor would be indifferent between the two strategies, we interpret Π as the maximum performance fee she will pay to switch from the FVUH to the FVR strategy. In other words, this utility-based criterion measures how much a mean-variance investor is willing to pay for conditioning on the forward volatility bias rather than assume unbiasedness.

To estimate the fee, we find the value of Π that satisfies:

$$\sum_{t=0}^{T-k} \left\{ \left(R^*_{p,t+k} - \Pi \right) - \frac{\delta}{2\left(1+\delta\right)} \left(R^*_{p,t+k} - \Pi \right)^2 \right\} = \sum_{t=0}^{T-k} \left\{ R_{p,t+k} - \frac{\delta}{2\left(1+\delta\right)} R^2_{p,t+k} \right\}, \qquad (13)$$

where $R^*_{p,t+k}$ is the gross portfolio return constructed using the expected return and volatility forecasts from the FVR model, $R_{p,t+k}$ is implied by the benchmark FVUH model, and δ is the investor's constant degree of relative risk aversion (RRA). We report Π in annualized basis points.

Our analysis also uses the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-proof performance measure defined as:

$$\Theta = \frac{1}{(1-\delta)} \ln \left(\frac{1}{T} \sum_{t=1}^{T-k} \left[\frac{R^*_{p,t+k}}{R_{p,t+k}} \right]^{1-\delta} \right), \qquad (14)$$

where Θ can be interpreted as the annualized certainty equivalent of the excess portfolio returns. As a manipulation-proof performance measure, Θ is an attractive alternative to the performance fee Π, because it is robust to a number of assumptions such as the distribution of portfolio returns. In

[13] For a mean-variance evaluation of economic models see also Marquering and Verbeek (2004), Han (2006), and Della Corte, Sarno and Thornton (2008). For a discussion of the advantages and disadvantages of using quadratic utility in this framework see Della Corte, Sarno and Tsiakas (2009).

contrast to Π, Θ does not require the assumption of quadratic utility to rank portfolios and thus it is useful to report Θ alongside Π.

5.4 Transaction Costs

The impact of transaction costs is an essential consideration in assessing the profitability of the dynamic carry trade in volatility strategy relative to the static and riskless FVUH strategy.[14] Making an accurate determination of the size of transaction costs is generally difficult, which is reflected in the wide range of estimates used in empirical studies.[15] This issue is more pronounced in the market for FVAs because these are instruments for which there is less information available on the size of transaction costs. We can avoid these concerns by calculating the break-even proportional transaction cost, τ^{BE}, that renders investors indifferent between the two strategies (e.g., Han, 2006). In comparing the dynamic CTV strategy with the static FVUH strategy, an investor who pays transaction costs lower than τ^{BE} will prefer the dynamic strategy. Since τ^{BE} is a proportional cost paid every time the portfolio is rebalanced, we report τ^{BE} in daily basis points.[16]

6 Economic Value of Volatility Speculation: The Results

We assess the economic value of the forward volatility bias by analyzing the performance of dynamically rebalanced portfolios based on the carry trade in volatility (CTV) strategy relative to the FVUH static benchmark. The economic evaluation is conducted both in sample and out of sample. The in-sample period ranges from January 2, 1991 to July 11, 2008. Note, however, that the data sample of implied volatilities does not start on the same date for all currencies. Hence we add risky assets in the portfolio allocation as data on them becomes available. The last currency to be added is the euro for which the data sample starts in January 1999. The out-of-sample period starts at the beginning of the sample (January 1991) and proceeds forward by sequentially updating the parameter estimates of the forward volatility regression day-by-day using a 3-year rolling window.[17]

Our economic evaluation focuses on the performance fee, Π, a US investor is willing to pay for switching from the static benchmark FVUH strategy to the dynamic CTV strategy. We report the estimates of Π as annualized fees in basis points for a target annualized portfolio volatility $\sigma_p^* = 10\%$

[14] For the FVUH strategy, the optimal weights fluctuate only to the extent that domestic interest rates vary over time.

[15] For example, the size of transaction costs depends on the type of investor (e.g., individual vs. institutional investor), the value of the transaction and the nature of the broker (e.g., brokerage firm vs. direct internet trading).

[16] In private communications, foreign exchange traders in large banks have revealed to us that for major currencies the full bid-ask spread on FVAs revolves around 20-30 basis points. Since our analysis is based on the mid-point of the bid-ask spread, a reasonable transaction cost would be half of the full spread, i.e., 10-15 basis points. Note that in the traditional FX market proportional transaction costs are very low for professional investors ranging around 1-2 basis points.

[17] This rolling window approach maximizes the length of the out-of-sample period compared to a recursive approach which needs to start at a much later date so there is a long enough sample on all currencies.

and a degree of RRA $\delta = 6$. The choice of σ_p^* and RRA is reasonable and consistent with numerous empirical studies (e.g., Fleming, Kirby and Ostdiek, 2001; Marquering and Verbeek, 2004; Han 2006). We have experimented with different σ_p^* and RRA values and found that qualitatively they have little effect on the asset allocation results discussed below.

The in-sample fees reported in Table 5 and the out-of-sample fees in Table 6 show that there is very high economic value associated with the forward volatility bias. In particular, switching from the static FVUH to the CTV portfolio gives the following staggering fees: (i) in-sample $\Pi = 1936$ annual basis points (*bps*) for investing in 3-month FVAs and $\Pi = 883$ *bps* for 6-month contracts, and (ii) out-of-sample $\Pi = 2001$ *bps* for $3m$ and $\Pi = 928$ *bps* for $6m$. These results are also reflected in the Sharpe ratios (SR), which for the CTV strategy are as follows: (i) in-sample $SR = 1.91$ for $3m$ and $SR = 1.18$ for $6m$, and (ii) out-of-sample $SR = 1.83$ for $3m$ and $SR = 1.25$ for $6m$. Another way of quantifying the high profitability of the carry trade in volatility strategy is to recognize that a portfolio volatility of around 10% can generate average portfolio returns exceeding 20%.

The portfolio weights on the risky assets (FVAs) required to generate this performance are in fact quite reasonable. Figure 3 illustrates that the average weights for the $3m$ carry trade in volatility strategy revolve from around -0.25 to $+0.25$ in-sample and from -0.45 to $+0.60$ out-of-sample. The figure also displays the 95% intervals of the variation in the weights, which in most cases ranges between -1 and $+1$. In short, therefore, the carry trade in volatility vastly outperforms the FVUH while taking reasonable positions in the FVAs.

We confirm these results by also computing the Goetzmann *et al.* (2007) manipulation-proof performance measure, Θ. Since Θ does not require the assumption of a particular utility function, a comparison of Π with Θ will inform us on the degree to which the performance fees depend on the choice of quadratic utility. Tables 5 and 6 show that both in-sample and out-of-sample Θ is consistently higher than Π by at least a few basis points. Therefore, the high economic value of the forward volatility bias is not driven by quadratic utility. A possible explanation of this result is that the CTV strategy exploits predictability in the returns to volatility speculation but not their volatility. Our analysis assumes that the volatility of volatility is constant, which might mitigate the impact of the utility function in assessing the mean-variance tradeoff of the CTV strategy.

6.1 Transaction Costs

If transaction costs are sufficiently high, the day-to-day fluctuations in the dynamic weights of the CTV strategy will render the strategy too costly to implement relative to the static FVUH benchmark. We address this concern by computing the break-even transaction cost, τ^{BE}, as the daily proportional cost that cancels out the utility advantage (and hence positive performance fee) of the CTV strategy. In comparing the dynamic CTV strategy with the static FVUH strategy, an investor who pays a transaction cost lower than τ^{BE} will prefer the dynamic strategy.

The break-even transaction costs are also reported in Tables 5 and 6, which demonstrate that both in-sample and out-of-sample the values of τ^{BE} are moderately high. For instance, a US investor trading 3-month FVAs will switch back to the FVUH model if she is subject to a proportional transaction cost of 45 *bps* (both in sample and out of sample), which seems to be higher than any actual transaction costs investors can expect to pay. Hence we can conclude that transaction costs will not reverse the high economic value of the forward volatility bias.

7 Robustness and Further Analysis

7.1 Portfolio Rebalancing with Non-Overlapping Returns

This section discusses directions we can follow in assessing the robustness of the results. To begin with, our analysis has so far focused on daily rebalancing, where the investor takes positions every day on 3-month and 6-month ahead implied volatility. An alternative way to evaluate the forward volatility bias is to consider portfolio rebalancing at the much lower frequencies of quarterly for 3-month ahead strategies and semi-annual for 6-month ahead strategies. This is equivalent to rebalancing only after the options determining the implied volatilities have expired. This approach is easier to implement and involves much lower transaction costs but discards most of the IV information that arrives daily. The $3m$ strategy now uses 4 IV observations per year as opposed to 252, whereas the $6m$ strategy only uses 2 IV observations per year. Due to the drastic reduction in the number of portfolio return observations we only show in-sample results in Table 7. Quite simply, there are not enough data to run a reliable out-of-sample exercise with quarterly and semi-annual rebalancing.

The results in Table 7 indicate that there is high economic value in the forward volatility bias even when rebalancing infrequently. For quarterly rebalancing, the CTV strategy delivers $SR = 1.65$ and $\Pi = 1762$ *bps*, whereas for semi-annual rebalancing $SR = 0.98$ and $\Pi = 648$ *bps*. The portfolio performance of the CTV strategy is certainly lower than for daily rebalancing but the CTV still substantially outperforms the FVUH benchmark. As expected, however, the break-even transaction costs are now very high: $\tau^{BE} = 1930$ *bps* for quarterly rebalancing and $\tau^{BE} = 837$ *bps* for semi-annual rebalancing. In short, therefore, there is robust economic value in the carry trade in volatility strategy even when rebalancing at a low frequency.

7.2 Carry Trade in Volatility vs. Carry Trade in Currency: Are their Returns Correlated?

One question that arises naturally from our results is whether the high economic value of the forward volatility bias (CTV strategy) in the FX options market is related to the economic value of the forward bias (CTC strategy) in the traditional FX market. Indeed, it is important to understand whether the returns to volatility speculation are correlated with the returns to currency speculation. If the correlation between these two strategies is high, then the FX market and the FX options market may be potentially driven by the same underlying inefficiency.

We address this issue by designing a dynamic strategy for currency speculation that closely corresponds to the strategy for volatility speculation described in Section 5.1. Specifically, we consider a US investor who builds a portfolio by allocating her wealth between the domestic riskless asset and seven forward exchange rates. The seven forward rates are for the same exchange rates and the same sample range as the volatility speculation strategy investing in the seven FVAs. We then use the original Fama regressions (Equations 3 and 4) and the same mean-variance framework to assess the economic value of predictability in exchange rate returns. In essence, we provide an economic evaluation of the carry trade in currency (CTC) strategy for the same exchange rate sample.

The simplest way of assessing the relation of the CTV strategy with the CTC strategy is to examine the correlation in their portfolio returns (net of the riskless rate). Table 8 shows that this correlation is low. For daily rebalancing, the correlation is 0.06 for $3m$ contracts and 0.14 for $6m$. For quarterly rebalancing ($3m$) it is 0.11, whereas for semi-annual rebalancing ($6m$) it is -0.03.

A more involved way of addressing this issue is to compare the separate portfolio performance of each of the two strategies with that of a combined strategy. The combined portfolio is constructed by investing in the same US bond as before and 14 risky assets: the seven forward volatility agreements plus the seven forward exchange rates. Table 5 presents the (in-sample) results, which are indicative of the low correlation between the CTV and the CTC strategies. In examining the two strategies separately, we clearly observe that the CTV strategy has far superior performance to the CTC strategy. For instance, the 3-month contracts give a Sharpe ratio of 1.91 for the CTV versus 0.88 for the CTC. The performance fees are 1936 *bps* and 571 *bps* respectively.[18]

More importantly, however, the combined strategy performs better than the CTV strategy alone. As we move from the CTV to the combined strategy, the Sharpe ratio rises from 1.91 to 2.18 for $3m$ and from 1.18 to 1.78 for $6m$. The performance fees increase from 1936 *bps* to 2311 *bps* for $3m$ and from 883 *bps* to 1640 *bps* for $6m$. Table 6 demonstrates that the out-of-sample results are qualitatively similar and so are the results in Table 7 for quarterly and semi-annual rebalancing. The clear increase in the economic value when combining CTV with CTC is evidence that there is distinct incremental information in the CTC over and above the information already incorporated in the CTV. Therefore, we can conclude that the forward volatility bias is largely distinct from the

[18]It is worth noting that simple carry trades exploiting the forward bias in the traditional FX market have been very profitable over the years (Galati, Heath and McGuire, 2007; Brunnermeier, Nagel and Pedersen, 2008). Our findings demonstrate that volatility speculation strategies can in fact be even more profitable than currency speculation strategies.

forward bias.

Finally, we turn to Figure 4, which illustrates the rolling Sharpe ratios for the $3m$ and $6m$ out-of-sample CTV and CTC strategies using a three-year rolling window. The figure shows that the SRs tend to be uncorrelated for long periods of time, especially during the last few years of the sample when all assets are available for inclusion in the portfolio. Moreover, it is interesting to note that for the years 2007 and 2008 the SR of the CTV displays a clear upward trend but the SR of the CTC shows a clear downward trend. This indicates that the CTV has done well during the recent credit crunch when the CTC has not. In other words, this is further evidence that the returns to volatility speculation tend to be uncorrelated with the returns to currency speculation even during the recent unwinding of the carry trade in currency.

7.3 Separating the Effect of Each Risky Asset on Portfolio Performance

A strong point of our analysis is that the carry trade in volatility portfolio invests in seven FVAs. In the FX market, this is a relatively large number of risky assets covering the exchange rates of some of the largest developed economies in the world. It would therefore be of interest to determine whether the economic value results are driven by one particular currency in the portfolio. Figure 5 presents the out-of-sample annualized Sharpe ratios with daily rebalancing for the full CTV portfolio as well as for portfolios where one of the risky assets is removed. The figure illustrates that the exclusion of any one FVA from the portfolio has little effect on the SR. This is also true for the combined CTV and CTC portfolio. Hence the main economic findings are not driven by any one of the currencies.

7.4 Is Implied Volatility a Random Walk?

The high Sharpe ratio of the CTV strategy is driven by two main factors: (i) investing in a relatively large set of FVAs and, more importantly, (ii) the high predictability in the returns to volatility speculation. As a measure of profitability, the SR values are specific to a strategy regardless of the benchmark we compare the strategy to. This is not the case in computing the performance fees, which capture the certainty equivalent relative to a particular benchmark. Setting a different benchmark will alter the performance fees. Indeed, the majority of studies of the forward bias in the traditional FX market tend to use the random walk of Meese and Rogoff (1983) as the benchmark model, not forward unbiasedness. Selecting the random walk as the benchmark would almost surely decrease the economic value in any alternative model. This is plausible in our case since the β estimate is much closer to zero (i.e., random walk) than unity (i.e., forward volatility unbiasedness). In our analysis, the choice of the FVUH benchmark is motivated by the objective of providing an empirical evaluation of departures from forward volatility unbiasedness.

Having said that, it would nevertheless be interesting to determine whether in future work the random walk model would be a sensible benchmark for assessing the economic value of predictability in the returns to volatility speculation. As a further robustness check, Table 9 presents the out-of-sample portfolio performance of the random walk with drift (RW) model against the FVUH benchmark for daily rebalancing. The RW model uses the OLS estimate of the intercept (α) of the forward volatility regression but imposes a slope coefficient of $\beta = 0$ (or equivalently $\gamma = -1$). The table shows that the out-of-sample economic value of the RW model is virtually identical to the CTV strategy. For the $3m$ strategies, the CTV generates $SR = 1.83$ and $\Pi = 2001$ *bps*, whereas the RW generates $SR = 1.92$ and $\Pi = 2004$ *bps*. For the $6m$ strategies, the CTV generates $SR = 1.25$ and $\Pi = 928$ *bps*, whereas the RW generates $SR = 1.21$ and $\Pi = 902$ *bps*. These results clearly suggest that the RW model is a useful benchmark to adopt in future studies of forecasting FX implied volatility.

Our final robustness check on the economic value of departures from FVUH involves the RW in an asset allocation framework that is much simpler than mean-variance. Table 9 also presents the economic performance of the RW model for a $1/N$ allocation strategy over the seven FVA returns. In contrast to mean-variance, the $1/N$ strategy does not use an estimate of the unconditional covariance of the FVA returns, has no target portfolio volatility, and can be readily implemented daily out-of-sample without the need to specify a rolling window. Even in this simple case, the results remain strong as the $3m$ Sharpe ratio is 1.29 and the $6m$ is 0.79. The performance fees relative to the FVUH are 1522 *bps* for $3m$ and 470 *bps* for $6m$. Therefore, we can conclude that the economic value of the carry trade in volatility is high even when we implement the simplest model (RW) with the simplest asset allocation strategy ($1/N$).

7.5 Is there a Volatility Term Premium?

We have defined the volatility term premium as $E_t\left[\sigma_{t+k} - \varphi_t^k\right]$. Under the FVUH, the volatility term premium should be equal to zero. Our empirical results have so far established that: (i) the unconditional (sample average) volatility term premium is non-zero and can be either positive or negative (see Table 2); (ii) the volatility term premium is time-varying and predictable when conditioning on the forward volatility premium as shown in the predictive regressions (see Tables 3 and 4); (iii) there is high economic value in predicting the volatility term premium of a mean-variance portfolio leading to a highly profitable carry trade in volatility strategy (see Tables 5-7); and (iv) the economic performance of the random walk model is very similar to that of the carry trade in volatility strategy (see Table 9).

These results motivate a simple strategy designed to provide a more careful examination of the volatility term premium for each individual FVA rather than a portfolio of FVAs. This strategy will be based on the random walk model in the following way. Consider an investor who goes long on an

FVA when $\sigma_t > \varphi_t^k$ and short on an FVA when $\varphi_t^k > \sigma_t$. The conditional return of this strategy is $\left(\sigma_{t+k} - \varphi_t^k\right) \times sign\left(\sigma_t - \varphi_t^k\right)$, which is simply a reformulation of the volatility term premium.

Table 10 shows that in 13 of the 14 cases the volatility term premium is positive with an annualized conditional mean ranging between $10\% - 20\%$ and an annualized standard deviation around $20\% - 30\%$. Not surprisingly, the single exception is the $6m$ CAD for which we have seen that the FVUH holds empirically. The volatility term premium tends to have low skewness (positive or negative), low excess kurtosis, and are highly persistent from day to day whether the horizon is 3 months or 6 months. Finally, it is important to note that the correlation between on the one hand the volatility term premium (defined as $E_t\left[\sigma_{t+k} - \varphi_t^k\right]$), which is the return to volatility speculation, and on the other hand the excess currency return (defined as $E_t\left[s_{t+k} - f_t^k\right]$), which is the return to currency speculation, is very low revolving around zero being positive half of the time. This is further evidence that what causes a violation of the FVUH is uncorrelated with what causes a violation of the FUH. In short, we can conclude that the volatility term premium is non-zero, time-varying and predictable in a carry trade in volatility strategy, and largely uncorrelated to the return to currency speculation.

8 Conclusion

The introduction of the forward volatility agreement (FVA) has allowed investors to speculate on the future volatility of exchange rate returns. An FVA contract determines the forward implied volatility defined over an interval starting at a future date. Forward volatility is by design meant to be an unbiased predictor of future spot volatility for all relevant maturities. However, if there is a bias in the way the market sets forward volatility, then the returns to volatility speculation will be predictable and a carry trade in volatility strategy can be profitable. Still, there is no study to date in the foreign exchange literature on the empirical issues surrounding FVAs. These include the empirical properties of FVAs (e.g., their risk-return tradeoff), the extent to which forward volatility is a biased predictor of future spot volatility, and the economic value of predictability in the returns to volatility speculation.

This paper fills this gap in the literature by formulating and testing the forward volatility unbiasedness hypothesis. Our empirical results are startling. First, we find statistically significant evidence that forward volatility is a systematically biased predictor that overestimates future spot volatility. This is similar to the tendency of forward exchange rates to overestimate the future rate of depreciation of high interest currencies, and the tendency of spot implied volatility to overestimate future realized volatility. Second, the rejection of the forward volatility unbiasedness indicates the presence of non-zero, time-varying and predictable volatility term premiums in foreign exchange. Third, there is very high in-sample and out-of-sample economic value in predicting the returns to volatility speculation in the context of dynamic asset allocation. The economic gains are robust to reasonable transaction costs and largely uncorrelated with the gains from currency speculation strategies. Therefore, the profitability of the carry trade in volatility strategy is distinct from the profitability of the carry trade in currency strategy. In the end, our statistical and economic analysis establishes the forward volatility bias, which we view as a new puzzle in foreign exchange.

To put these findings in context, consider that the empirical rejection of uncovered interest parity leading to the forward bias puzzle has over the years generated an enormous literature in foreign exchange. At the same time, the carry trade has been a highly profitable currency speculation strategy. As this is the first study to establish the volatility analogue to the forward bias puzzle and demonstrate the high economic value of volatility speculation strategies, there are certainly many directions in which our analysis can be extended. These may involve using alternative data sets, improvements in the econometric techniques and the empirical setting, refinements in the framework for the economic evaluation of realistic trading strategies and, finally, the development of theoretical models aiming at explaining these findings and rationalizing the volatility term premium. Having established the main result motivating such extensions, we leave these for future research.

Table 1. Descriptive Statistics on Daily FX Volatility Levels

The table reports descriptive statistics for the daily spot and forward implied volatilities on seven US dollar exchange rates for 3-month, 6-month and 12-month maturities. The sample ends on July 11, 2008, and starts on Jan 2, 1991 for AUD and JPY (4416 obs), Jan 2, 1992 for GBP (4162 obs), Jan 4, 1993 for CHF (3908 obs), Jan 2, 1997 for CAD (2899 obs), Jan 16, 1998 for NZD (2637 obs), and Jan 4, 1999 for EUR (2396 obs). The means and standard deviations are reported in annualized percent units. For the autocorellation, a lag of 1 corresponds to one trading day, 63 to three months and 126 to six months. ADF is the augmented Dickey-Fuller statistic for the null hypothesis of non-stationarity. The asterisks *, **, *** denote significance at the 10%, 5% and 1% level, respectively.

					Autocorrelation			
	Mean	*St.Dev*	*Skew*	*Kurt*	1	63	126	*ADF*
AUD								
3m Implied Vol	9.782	2.057	0.448	2.739	0.993	0.760	0.550	-2.895^{**}
6m Implied Vol	9.652	1.885	0.274	2.627	0.995	0.806	0.596	-2.558
12m Implied Vol	9.509	1.817	0.145	2.481	0.997	0.841	0.644	-2.487
3m Forward Vol	9.499	1.810	0.050	2.526	0.986	0.830	0.625	-2.289
6m Forward Vol	9.349	1.817	-0.019	2.413	0.989	0.859	0.682	-2.058
CAD								
3m Implied Vol	7.188	1.787	0.588	3.338	0.996	0.782	0.544	-2.380
6m Implied Vol	7.087	1.695	0.503	3.200	0.997	0.826	0.608	-2.116
12m Implied Vol	7.034	1.633	0.442	3.056	0.998	0.849	0.651	-2.012
3m Forward Vol	6.976	1.632	0.403	3.118	0.996	0.853	0.660	-1.500
6m Forward Vol	6.979	1.582	0.398	2.914	0.997	0.867	0.691	-1.773
CHF								
3m Implied Vol	10.853	1.881	0.030	3.524	0.988	0.656	0.472	-3.726^{***}
6m Implied Vol	10.956	1.762	-0.254	3.473	0.989	0.735	0.564	-3.117^{**}
12m Implied Vol	11.030	1.691	-0.421	3.473	0.994	0.792	0.627	-2.669^{*}
3m Forward Vol	11.041	1.739	-0.382	3.567	0.963	0.764	0.607	-2.795^{*}
6m Forward Vol	11.094	1.680	-0.486	3.549	0.979	0.810	0.658	-2.481
EUR								
3m Implied Vol	9.872	2.002	0.024	3.220	0.994	0.815	0.680	-2.656^{*}
6m Implied Vol	9.980	1.920	-0.077	3.146	0.996	0.851	0.711	-2.087
12m Implied Vol	10.058	1.850	-0.112	3.119	0.998	0.866	0.722	-1.671
3m Forward Vol	10.081	1.874	-0.160	3.041	0.996	0.871	0.727	-1.664
6m Forward Vol	10.131	1.799	-0.144	3.081	0.998	0.875	0.725	-1.447
GBP								
3m Implied Vol	9.046	1.994	1.106	4.668	0.993	0.729	0.583	-3.348^{**}
6m Implied Vol	9.225	1.823	0.908	3.921	0.996	0.794	0.644	-3.093^{**}
12m Implied Vol	9.382	1.759	0.775	3.303	0.997	0.845	0.716	-2.607^{*}
3m Forward Vol	9.379	1.757	0.800	3.628	0.988	0.788	0.638	-3.166^{**}
6m Forward Vol	9.527	1.745	0.697	2.947	0.996	0.867	0.759	-2.498
JPY								
3m Implied Vol	10.740	2.464	1.130	4.786	0.987	0.729	0.581	-3.162^{**}
6m Implied Vol	10.834	2.389	1.031	4.144	0.993	0.818	0.674	-2.656^{*}
12m Implied Vol	10.913	2.361	0.979	3.899	0.996	0.866	0.735	-2.143
3m Forward Vol	10.907	2.400	0.930	3.706	0.992	0.872	0.739	-2.214
6m Forward Vol	10.982	2.378	0.911	3.589	0.995	0.893	0.774	-2.104
NZD								
3m Implied Vol	11.983	1.848	0.611	2.300	0.991	0.585	0.243	-3.499^{***}
6m Implied Vol	11.843	1.698	0.487	2.172	0.994	0.656	0.296	-3.012^{**}
12m Implied Vol	11.720	1.623	0.398	2.278	0.994	0.720	0.364	-2.588^{*}
3m Forward Vol	11.684	1.665	0.324	2.349	0.992	0.728	0.374	-2.731^{*}
6m Forward Vol	11.583	1.630	0.327	2.485	0.990	0.776	0.450	-2.431

Table 2. Descriptive Statistics on Daily FX Volatility Returns

The table displays descriptive statistics for the daily FX volatility returns on seven US dollar exchange rates for 3-month and 6-month maturities. The *Implied Volatility Return* $(\sigma_{t+k}-\sigma_t)$ is defined as the log future spot IV minus the log spot IV. The *Forward Volatility Premium* $(\varphi_t^k-\sigma_t)$ is defined as the log forward IV minus the log spot IV. The *Excess Volatility Return* $(\sigma_{t+k}-\varphi_t^k)$ is defined as the log future spot IV minus the log forward IV. The means and standard deviations are reported in annualized percent units. For the autocorellation, a lag of 1 corresponds to one trading day, 63 to three months and 126 to six months. ADF is the augmented Dickey-Fuller statistic for the null hypothesis of non-stationarity. The asterisks *, **, *** denote significance at the 10%, 5% and 1% level, respectively.

					Autocorrelation			
	Mean	*St.Dev*	*Skew*	*Kurt*	1	63	126	*ADF*
AUD								
3m Implied Vol Return	0.792	28.769	0.498	3.608	0.975	−0.053	−0.040	−6.421***
3m Forward Vol Premium	−10.993	19.015	−0.491	4.574	0.909	0.472	0.317	−4.602***
3m Excess Vol Return	11.786	34.179	0.626	3.261	0.972	0.410	0.197	−4.693***
6m Implied Vol Return	1.521	24.847	0.572	3.591	0.987	0.525	0.036	−4.165***
6m Forward Vol Premium	−6.711	11.231	−2.468	18.401	0.840	0.544	0.402	−3.458***
6m Excess Vol Return	8.232	28.225	0.602	2.941	0.980	0.664	0.266	−3.473***
CAD								
3m Implied Vol Return	7.194	31.129	1.457	7.653	0.981	−0.085	−0.202	−4.821***
3m Forward Vol Premium	−11.125	14.259	−2.612	16.209	0.915	0.104	0.054	−5.230***
3m Excess Vol Return	18.319	31.591	1.515	7.262	0.980	0.200	−0.047	−4.418***
6m Implied Vol Return	7.942	26.659	1.324	5.675	0.992	0.412	−0.214	−3.625***
6m Forward Vol Premium	−2.417	5.666	−2.840	18.037	0.932	0.300	0.061	−5.091***
6m Excess Vol Return	10.359	26.115	1.486	5.718	0.992	0.477	−0.133	−3.279**
CHF								
3m Implied Vol Return	−1.045	28.299	0.809	3.860	0.966	−0.179	−0.051	−6.667***
3m Forward Vol Premium	8.334	15.128	−1.078	9.923	0.749	0.178	−0.004	−7.681***
3m Excess Vol Return	−9.378	29.729	0.619	3.811	0.944	0.148	0.022	−5.705***
6m Implied Vol Return	−1.398	22.192	0.500	3.736	0.980	0.422	−0.104	−4.249***
6m Forward Vol Premium	3.524	7.775	−1.244	15.229	0.716	0.106	0.033	−5.265***
6m Excess Vol Return	−4.922	22.908	0.523	4.304	0.973	0.518	0.049	−3.604***
EUR								
3m Implied Vol Return	0.943	25.676	0.364	3.041	0.973	−0.037	0.065	−5.700***
3m Forward Vol Premium	9.698	11.288	0.035	2.910	0.934	0.352	0.200	−4.215***
3m Excess Vol Return	−8.755	27.738	0.268	3.087	0.980	0.315	0.202	−4.118***
6m Implied Vol Return	0.936	22.963	0.447	3.617	0.990	0.591	0.148	−2.953**
6m Forward Vol Premium	3.933	5.564	0.165	3.729	0.945	0.455	0.319	−3.914***
6m Excess Vol Return	−2.997	24.368	0.277	3.238	0.993	0.675	0.270	−3.461***
GBP								
3m Implied Vol Return	−2.5472	32.244	0.756	4.095	0.977	−0.196	−0.177	−6.932***
3m Forward Vol Premium	16.751	19.559	−2.714	34.349	0.920	0.276	0.207	−5.615***
3m Excess Vol Return	−19.298	33.397	0.402	5.094	0.973	0.200	−0.030	−7.584***
6m Implied Vol Return	−1.785	24.294	0.243	3.495	0.990	0.320	−0.287	−5.127***
6m Forward Vol Premium	7.073	8.938	0.611	5.292	0.960	0.425	0.242	−5.714***
6m Excess Vol Return	−8.858	24.206	−0.410	4.040	0.991	0.459	−0.067	−4.082***
JPY								
3m Implied Vol Return	−0.324	32.059	0.350	3.282	0.964	−0.235	0.045	−7.584***
3m Forward Vol Premium	7.171	16.776	−0.453	4.436	0.915	0.300	0.227	−7.008***
3m Excess Vol Return	−7.495	31.692	0.190	3.451	0.967	0.185	0.124	−5.452***
6m Implied Vol Return	0.111	23.510	0.055	3.253	0.981	0.448	−0.102	−5.764***
6m Forward Vol Premium	3.478	8.013	−0.273	4.621	0.897	0.267	0.100	−6.403***
6m Excess Vol Return	−3.367	23.426	0.013	3.235	0.984	0.553	0.056	−4.818***
NZD								
3m Implied Vol Return	1.967	27.331	0.416	3.584	0.978	−0.079	−0.055	−5.452***
3m Forward Vol Premium	−10.492	15.417	−0.631	3.068	0.957	0.549	0.416	−4.225***
3m Excess Vol Return	12.459	30.170	0.577	2.944	0.984	0.321	0.083	−3.492***
6m Implied Vol Return	3.060	23.345	0.344	3.097	0.991	0.476	−0.126	−3.794***
6m Forward Vol Premium	−4.880	8.937	−1.121	5.154	0.943	0.603	0.477	−3.252**
6m Excess Vol Return	7.940	24.433	0.670	2.738	0.988	0.592	0.086	−3.542***

Table 3. Predictive Regressions with OLS Estimation

The table presents the ordinary least squares (OLS) estimates for $\sigma_{t+k} - \sigma_t = \alpha + \beta \left(\varphi_t^k - \sigma_t\right) + \varepsilon_{t+k}$ and $\sigma_{t+k} - \varphi_t^k = \alpha + \gamma \left(\varphi_t^k - \sigma_t\right) + \varepsilon_{t+k}$ on seven US dollar exchange rates. The returns are measured over 3-months or 6-months but are observed and estimated daily. The *Implied Volatility Return* $(\sigma_{t+k} - \sigma_t)$ is defined as the log future spot IV minus the log spot IV. The *Excess Volatility Return* $(\sigma_{t+k} - \varphi_t^k)$ is defined as the log future spot IV minus the log forward IV. t^α is the t-statistic for the null hypothesis $\alpha = 0$. t^β is the t-statistic for the null hypothesis $\beta = 1$. t^γ is the t-statistic for the null hypothesis $\gamma = 0$. F is the F-statistic for the joint null hypothesis $\alpha = 0$ and $\beta = 1$. BL is the Box-Ljung statistic for the null hypothesis of no autocorellation in the regression residuals between 64 (127) and 252 trading days for the 3-month (6-month) figures. R^2 is the coefficient of determination. Newey-West asymptotic standard errors are reported in parentheses and p-values in brackets.

	Implied Volatility Return							*Excess Volatility Return*			
	α	β	t^α	t^β	F	BL	R^2	γ	t^γ	BL	R^2
3-month											
AUD	0.0028	0.029	0.425	−15.540	971	329	0.01	−0.971	−15.543	329	0.29
	(0.0065)	(0.062)	[0.671]	[<0.01]	[<0.01]	[<0.01]		(0.062)	[<0.01]	[<0.01]	
CAD	0.0299	0.428	3.453	−6.220	116	265	0.04	−0.571	−6.224	265	0.07
	(0.0087)	(0.092)	[<0.01]	[<0.01]	[<0.01]	[0.279]		(0.092)	[<0.01]	[0.279]	
CHF	−0.0092	0.319	−1.364	−7.569	282	370	0.03	−0.681	−7.569	370	0.12
	(0.0068)	(0.090)	[0.173]	[<0.01]	[<0.01]	[<0.01]		(0.090)	[<0.01]	[<0.01]	
EUR	0.0007	0.068	0.089	−6.128	232	369	0.01	−0.932	−6.128	369	0.14
	(0.0081)	(0.152)	[0.929]	[<0.01]	[<0.01]	[<0.01]		(0.152)	[<0.01]	[<0.01]	
GBP	−0.0232	0.401	−3.070	−5.763	340	300	0.06	−0.599	−5.763	300	0.12
	(0.0075)	(0.104)	[<0.01]	[<0.01]	[<0.01]	[0.021]		(0.104)	[<0.01]	[0.021]	
JPY	−0.0105	0.541	−1.503	−5.456	142	240	0.08	−0.458	−5.456	240	0.06
	(0.0070)	(0.084)	[0.133]	[<0.01]	[<0.01]	[0.699]		(0.084)	[<0.01]	[0.699]	
NZD	0.0090	0.157	1.156	−8.486	327	243	0.01	−0.843	−8.486	243	0.18
	(0.0078)	(0.099)	[0.248]	[<0.01]	[<0.01]	[0.642]		(0.099)	[<0.01]	[0.642]	
6-month											
AUD	0.0005	−0.211	0.063	−13.158	764	338	0.01	−1.211	−13.158	338	0.23
	(0.0085)	(0.092)	[0.950]	[<0.01]	[<0.01]	[<0.01]		(0.920)	[<0.01]	[<0.01]	
CAD	0.0511	0.948	4.754	−0.241	0.194	229	0.04	−0.052	−0.241	229	0.01
	(0.0108)	(0.216)	[<0.01]	[0.809]	[0.824]	[0.846]		(0.216)	[0.809]	[0.846]	
CHF	−0.0111	0.233	−1.352	−5.686	152	249	0.01	−0.767	−5.686	249	0.07
	(0.0082)	(0.135)	[0.176]	[<0.01]	[<0.01]	[0.540]		(0.135)	[<0.01]	[0.540]	
EUR	0.0160	−0.574	1.573	−5.095	210	153	0.02	−1.574	−5.095	153	0.13
	(0.0101)	(0.309)	[0.116]	[<0.01]	[<0.01]	[1.000]		(0.309)	[<0.01]	[1.000]	
GBP	−0.0275	0.527	−3.602	−2.743	83	207	0.04	−0.473	−2.743	207	0.03
	(0.0076)	(0.173)	[<0.01]	[<0.01]	[<0.01]	[0.980]		(0.173)	[<0.01]	[0.980]	
JPY	−0.0087	0.531	−1.097	−3.881	62	163	0.03	−0.469	−3.881	163	0.03
	(0.0079)	(0.121)	[0.273]	[<0.01]	[<0.01]	[1.000]		(0.121)	[<0.01]	[1.000]	
NZD	0.0196	0.175	2.128	−4.594	144	179	0.01	−0.825	−4.594	179	0.09
	(0.0092)	(0.180)	[0.030]	[<0.01]	[<0.01]	[1.000]		(0.180)	[<0.01]	[1.000]	

Table 4. The Forward Volatility Bias under Alternative Estimation Methods

The table presents the β estimates of the regression $\sigma_{t+k}-\sigma_t= \alpha + \beta \left(\varphi_t^k-\sigma_t\right) +\varepsilon_{t+k}$ using three estimation methods: ordinary least squares (OLS), least absolute deviation (LAD) and errors-in-variables (EIV). The returns are measured over 3-months or 6-months but are observed and estimated daily for the FVAs of seven US dollar exchange rates. The LAD estimator minimizes the sum of the absolute value of residuals and is robust to outliers (Bassett and Koenker, 1978). The EIV method uses maximum likelihood jointly with the Kalman filter to estimate the parameters when the explanatory variable is measured with error (Carr and Wu, 2009). t^β is the t-statistic for the null hypothesis $\beta = 1$. Asymptotic standard errors are reported in parentheses and p-values in brackets. OLS uses Newey-West standard errors. LAD uses Weiss (1990) standard errors.

	OLS Estimation		*LAD Estimation*		*EIV Estimation*	
	β	t^β	β	t^β	β	t^β
3-month						
AUD	0.029	−15.540	0.057	−14.361	0.005	−41.270
	(0.062)	[<0.01]	(0.065)	[<0.01]	(0.069)	[<0.01]
CAD	0.428	−6.220	0.442	−5.935	0.445	−11.870
	(0.092)	[<0.01]	(0.096)	[<0.01]	(0.047)	[<0.01]
CHF	0.319	−7.569	0.239	−8.020	0.424	−22.048
	(0.090)	[<0.01]	(0.095)	[<0.01]	(0.026)	[<0.01]
EUR	0.068	−6.128	0.032	−6.012	0.038	−23.270
	(0.152)	[<0.01]	(0.161)	[<0.01]	(0.041)	[<0.01]
GBP	0.401	−5.763	0.368	−5.930	0.410	−29.846
	(0.104)	[<0.01]	(0.107)	[<0.01]	(0.020)	[<0.01]
JPY	0.541	−5.456	0.598	−4.562	0.566	−16.216
	(0.084)	[<0.01]	(0.088)	[<0.01]	(0.027)	[<0.01]
NZD	0.157	−8.486	0.059	−9.029	0.152	−24.047
	(0.099)	[<0.01]	(0.104)	[<0.01]	(0.035)	[<0.01]
6-month						
AUD	−0.211	−13.158	−0.271	−13.065	−0.378	−34.920
	(0.092)	[<0.01]	(0.097)	[<0.01]	(0.039)	[<0.01]
CAD	0.948	−0.241	1.033	0.170	1.013	0.128
	(0.216)	[0.809]	(0.192)	[0.865]	(0.104)	[0.900]
CHF	0.233	−5.686	0.283	−5.041	0.258	−17.322
	(0.135)	[<0.01]	(0.142)	[<0.01]	(0.043)	[<0.01]
EUR	−0.574	−5.095	−1.055	−6.050	−0.656	−22.348
	(0.309)	[<0.01]	(0.400)	[<0.01]	(0.074)	[<0.01]
GBP	0.527	−2.743	0.585	−2.292	0.535	−13.077
	(0.173)	[<0.01]	(0.181)	[0.022]	(0.036)	[0.010]
JPY	0.531	−3.881	0.435	−4.444	0.560	−9.515
	(0.121)	[<0.01]	(0.127)	[<0.01]	(0.046)	[<0.01]
NZD	0.175	−4.594	−0.103	−5.653	0.175	−16.251
	(0.180)	[<0.01]	(0.195)	[<0.01]	(0.051)	[<0.01]

Table 5. In-Sample Economic Value for Daily Rebalancing

The table shows the in-sample portfolio performance of maximum return dynamic strategies against the static benchmark for daily rebalancing. The *CTV* is the carry trade in volatility strategy, which conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and seven forward volatility agreements. The *CTC* is the carry trade in currency strategy, which conditions on the forward bias by building an efficient portfolio investing in a US bond and seven forward exchange rates. The combination of the *CTV* and *CTC* strategies conditions on both the forward volatility bias and the forward bias. The CTV and CTC strategies are based on the OLS estimates from the predictive regressions. The static benchmark is riskless investing implied by both forward volatility unbiasedness and forward unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p, σ_p and SR respectively. The performance fee (Π) denotes the amount an investor with quadratic utility, target portfolio volatility of 10% and a degree of relative risk aversion equal to 6 is willing to pay for switching from the static to the dynamic model. The performance fees are expressed in annual basis points. Θ is the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-free performance measure also expressed in annual basis points. The break-even transaction cost τ^{BE} is defined as the proportional cost that cancels out the utility advantage of the dynamic strategy and is expressed in daily basis points. The in-sample period runs from January 2, 1991 to July 11, 2008.

	μ_p	σ_p	SR	Π	Θ	τ^{BE}
	Carry Trade in Volatility (CTV) Strategy					
3-month	29.5	13.2	1.91	1936	1993	45
6-month	19.3	12.7	1.18	883	1044	17
	Carry Trade in Currency (CTC) Strategy					
3-month	13.8	10.8	0.88	571	605	17
6-month	15.9	10.6	1.09	746	807	32
	Combination of CTV and CTC Strategies					
3-month	33.3	13.3	2.18	2311	2356	29
6-month	26.8	12.6	1.78	1640	1727	13

Table 6. Out-of-Sample Economic Value for Daily Rebalancing

The table shows the out-of-sample portfolio performance of maximum return dynamic strategies against the static benchmark for daily rebalancing. The CTV is the carry trade in volatility strategy, which conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and seven forward volatility agreements. The CTC is the carry trade in currency strategy, which conditions on the forward bias by building an efficient portfolio investing in a US bond and seven forward exchange rates. The combination of the CTV and CTC strategies conditions on both the forward volatility bias and the forward bias. The CTV and CTC strategies are based on the OLS estimates from the predictive regressions. The static benchmark is riskless investing implied by both forward volatility unbiasedness and forward unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p, σ_p and SR respectively. The performance fee (Π) denotes the amount an investor with quadratic utility, target portfolio volatility of 10% and a degree of relative risk aversion equal to 6 is willing to pay for switching from the static to the dynamic model. The performance fees are expressed in annual basis points. Θ is the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-free performance measure also expressed in annual basis points. The break-even transaction cost τ^{BE} is defined as the proportional cost that cancels out the utility advantage of the dynamic strategy and is expressed in daily basis points. The out-of-sample period starts at the beginning of the sample (January 2, 1991) and proceeds forward by sequentially updating the parameter estimates day-by-day using a 3-year rolling window.

	μ_p	σ_p	SR	Π	Θ	τ^{BE}
	Carry Trade in Volatility (CTV) Strategy					
3-month	32.2	15.3	1.83	2001	2122	45
6-month	24.3	15.9	1.25	928	1263	17
	Carry Trade in Currency (CTC) Strategy					
3-month	26.0	14.1	1.53	1496	1561	25
6-month	32.9	17.3	1.65	1508	1952	13
	Combination of CTV and CTC Strategies					
3-month	47.6	17.4	2.49	3277	3400	25
6-month	49.9	22.6	2.01	2659	3114	20

Table 7. Economic Value for Quarterly and Semi-Annual Rebalancing

The table shows the in-sample portfolio performance of maximum return dynamic strategies against the static benchmark for quarterly and semi-annual rebalancing. The portfolios are rebalanced quarterly for the strategies investing in 3-month FVAs and semi-annually for the strategies investing in 6-month FVAs. This is equivalent to rebalancing only after the options determining the implied volatilities have expired and it vastly reduces the sample size of portfolio returns. The CTV is the carry trade in volatility strategy, which conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and seven forward volatility agreements. The CTC is the carry trade in currency strategy, which conditions on the forward bias by building an efficient portfolio investing in a US bond and seven forward exchange rates. The combination of the CTV and CTC strategies conditions on both the forward volatility bias and the forward bias. The CTV and CTC strategies are based on the OLS estimates from the predictive regressions. The static benchmark is riskless investing implied by both forward volatility unbiasedness and forward unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p, σ_p and SR respectively. The performance fee (Π) denotes the amount an investor with quadratic utility, target portfolio volatility of 10% and a degree of relative risk aversion equal to 6 is willing to pay for switching from the static to the dynamic model. The performance fees are expressed in annual basis points. Θ is the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-free performance measure also expressed in annual basis points. The break-even transaction cost τ^{BE} is defined as the proportional cost that cancels out the utility advantage of the dynamic strategy and is expressed in quarterly or semi-annual basis points. The in-sample period runs from January 2, 1991 to July 11, 2008. We do not conduct an out-of-sample analysis due to the small sample size.

	μ_p	σ_p	SR	Π	Θ	τ^{BE}
	Carry Trade in Volatility (CTV) Strategy					
Quarterly	30.5	15.8	1.65	1762	1888	1930
Semi-annual	17.2	12.8	0.98	648	824	837
	Carry Trade in Currency (CTC) Strategy					
Quarterly	12.6	10.9	0.76	438	489	441
Semi-annual	15.7	10.1	1.03	669	724	1090
	Combination of CTV and CTC Strategies					
Quarterly	39.4	16.3	2.15	2597	2718	1233
Semi-annual	36.5	15.3	2.08	2253	2382	492

Table 8. Cross-Correlations in Portfolio Returns across Different Strategies

The table reports the in-sample correlations between the portfolio returns (net of the riskless rate) of the CTV and CTC strategies both for daily and for quarterly/semi-annual rebalancing. With daily rebalancing the portfolio weights are updated every day whether investing in 3-month or 6-month FVAs. With quarterly/semi-annual rebalancing the portfolio weights are updated quarterly when investing in 3-month FVAs and semi-annually when investing in 6-month FVAs. The CTV is the carry trade in volatility strategy, which conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and seven forward volatility agreements. The CTC is the carry trade in currency strategy, which conditions on the forward bias by building an efficient portfolio investing in a US bond and seven forward exchange rates. The CTV and CTC strategies are based on the OLS estimates from the predictive regressions. The in-sample period runs from January 2, 1991 to July 11, 2008.

	Daily Rebalancing	*Quarterly/Semi-annual Rebalancing*
	$Corr(CTV, CTC)$	$Corr(CTV, CTC)$
3-month	0.06	0.11
6-month	0.14	-0.03

Table 9. Economic Value of the Random Walk Model for Daily Rebalancing

The table presents the out-of-sample portfolio performance of the Random Walk model against the static benchmark for daily rebalancing. We consider two dynamic strategies based on the Random Walk: (i) the mean-variance strategy that builds an efficient portfolio maximizing the expected portfolio return for a target portfolio volatility of 10%; and (ii) the simple 1/N strategy. Both dynamic strategies invest in a US bond and seven forward volatility agreements. The static benchmark is riskless investing implied by forward volatility unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p, σ_p and SR respectively. The performance fee (Π) denotes the amount an investor with quadratic utility, target portfolio volatility of 10% and a degree of relative risk aversion equal to 6 is willing to pay for switching from the static to the dynamic model. The performance fees are expressed in annual basis points. Θ is the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-free performance measure also expressed in annual basis points. The break-even transaction cost τ^{BE} is defined as the proportional cost that cancels out the utility advantage of the dynamic strategy and is expressed in daily basis points.

	μ_p	σ_p	SR	Π	Θ	τ^{BE}
	Out-of-Sample Mean-Variance Strategy					
3-month	30.8	13.8	1.92	2004	2048	43
6-month	21.6	14.3	1.21	902	1129	14
	Out-of-Sample 1/N Strategy					
3-month	37.8	26.1	1.29	1522	1530	52
6-month	20.5	20.4	0.79	470	559	24

Table 10. The Volatility Term Premium

The table displays descriptive statistics for the daily *Volatility Term Premium* corresponding to the forward volatility agreements on seven US dollar exchange rates for 3-month and 6-month maturities. The volatility term premium is defined as the conditional expectation $E_t\left[\sigma_{t+k}-\varphi_t^k\right]$. The volatility term premium is computed for the following simple strategy: go long on an FVA when $\sigma_t > \varphi_t^k$ and short on an FVA when $\varphi_t^k > \sigma_t$. The means and standard deviations are reported in annualized percent units. For the autocorellation, a lag of 1 corresponds to one trading day, 63 to three months and 126 to six months. The correlation with FX is defined as the correlation between the returns to volatility speculation with the returns to currency speculation: $Corr\left(E_t\left[\sigma_{t+k}-\varphi_t^k\right], E_t\left[s_{t+k}-f_t^k\right]\right)$.

					Autocorrelation			*Correlation*
	Mean	*St.Dev*	*Skew*	*Kurt*	1	63	126	$(E_t\left[\sigma_{t+k}-\varphi_t^k\right],$ $E_t\left[s_{t+k}-f_t^k\right])$
3-month								
AUD	28.71	31.57	0.14	3.79	0.780	0.097	0.016	0.02
CAD	14.04	32.13	0.80	7.92	0.803	0.094	−0.039	−0.01
CHF	18.63	28.62	−0.03	3.55	0.676	0.037	−0.041	0.13
EUR	20.17	26.21	0.00	2.94	0.810	0.092	−0.088	−0.19
GBP	17.28	33.67	0.40	4.16	0.838	0.142	−0.026	−0.18
JPY	13.51	31.19	0.21	3.21	0.844	0.109	−0.013	0.05
NZD	20.12	29.12	−0.03	3.55	0.849	0.084	−0.109	−0.12
6-month								
AUD	14.88	26.83	0.01	3.56	0.824	0.178	0.089	0.06
CAD	−0.32	27.12	−0.44	6.74	0.880	0.033	−0.098	−0.08
CHF	8.00	22.47	0.23	3.84	0.630	0.160	−0.047	0.02
EUR	11.91	22.96	−0.27	3.61	0.786	0.207	−0.097	−0.11
GBP	7.05	24.50	0.47	4.01	0.903	0.326	−0.023	0.02
JPY	6.08	23.15	−0.04	3.27	0.844	0.132	−0.060	0.11
NZD	10.78	23.88	−0.08	3.50	0.872	0.095	−0.045	−0.12

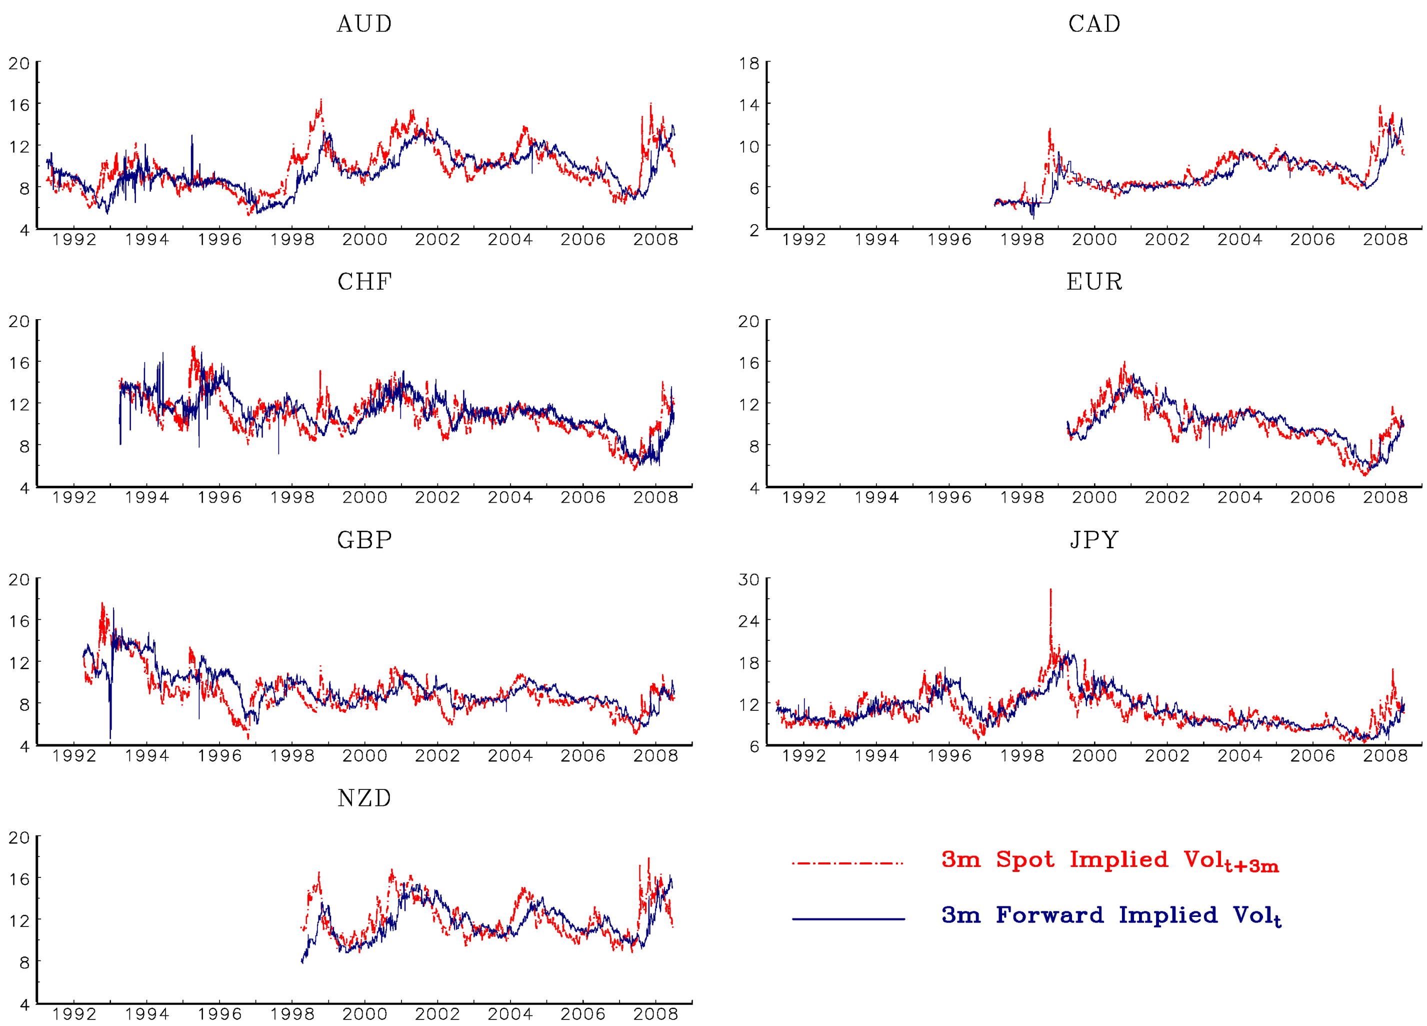

Figure 1. Spot and Forward Implied Volatilities

The figure displays the daily 3-month ahead spot and forward implied volatilities in annualized percent units on seven exchange rates vis-à-vis the US dollar.

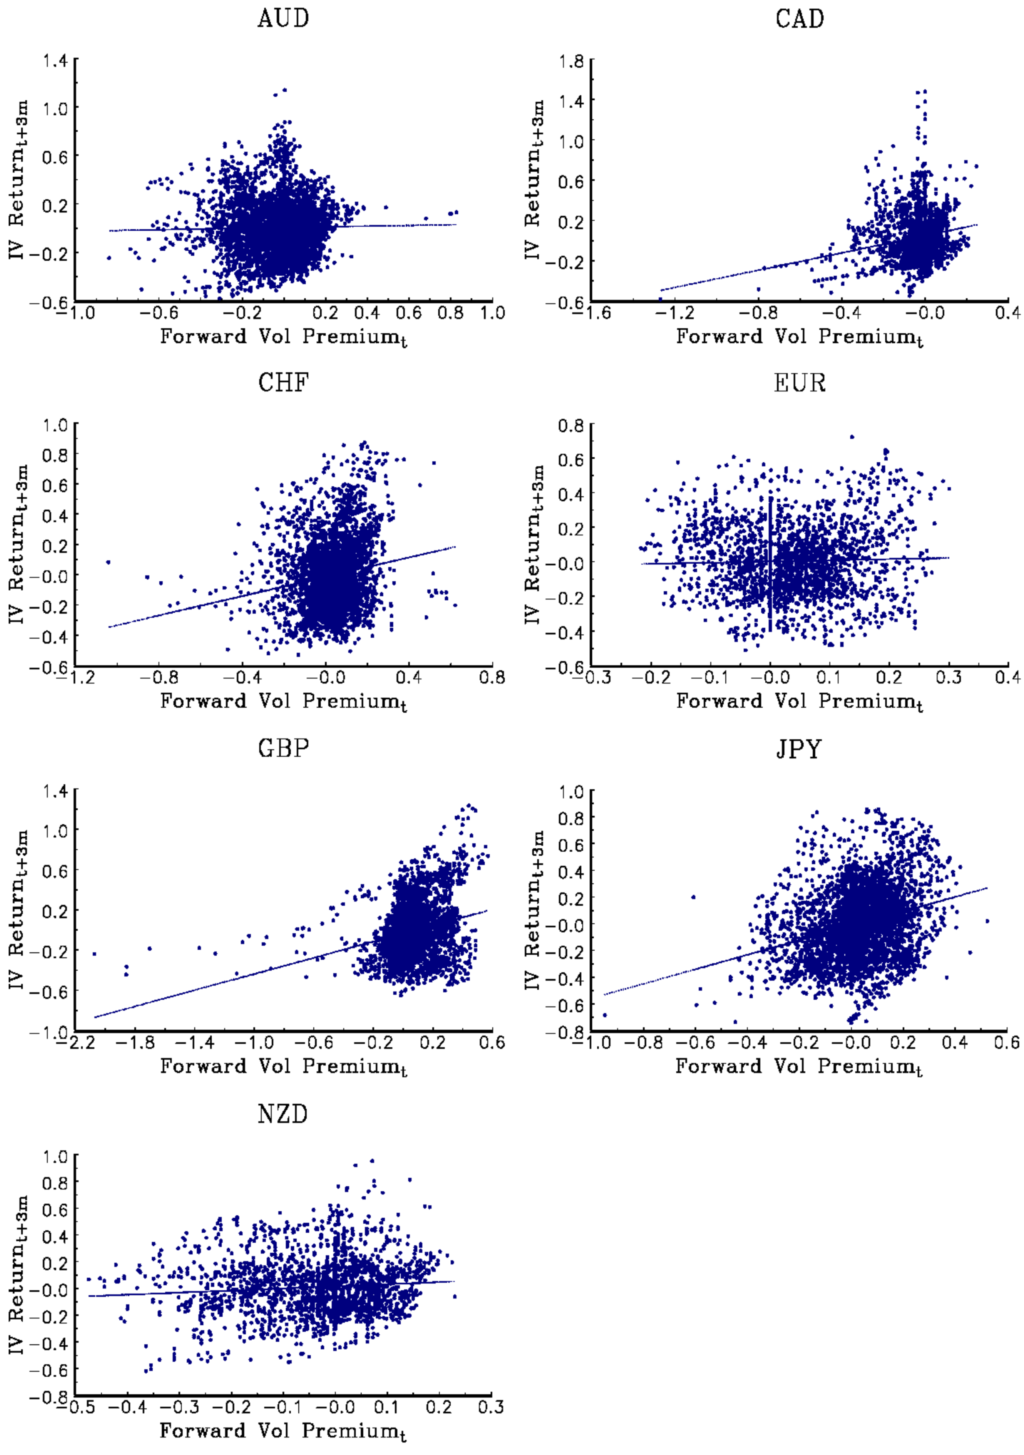

Figure 2. Implied Volatility Returns and Forward Volatility Premiums

The figure displays a scatter plot of the daily 3-month ahead implied volatility (IV) returns and forward volatility premiums on seven exchange rates vis-à-vis the US dollar. The IV return is defined as the log future spot IV minus the log spot IV. The forward volatility premium is defined as the log forward IV minus the log spot IV.

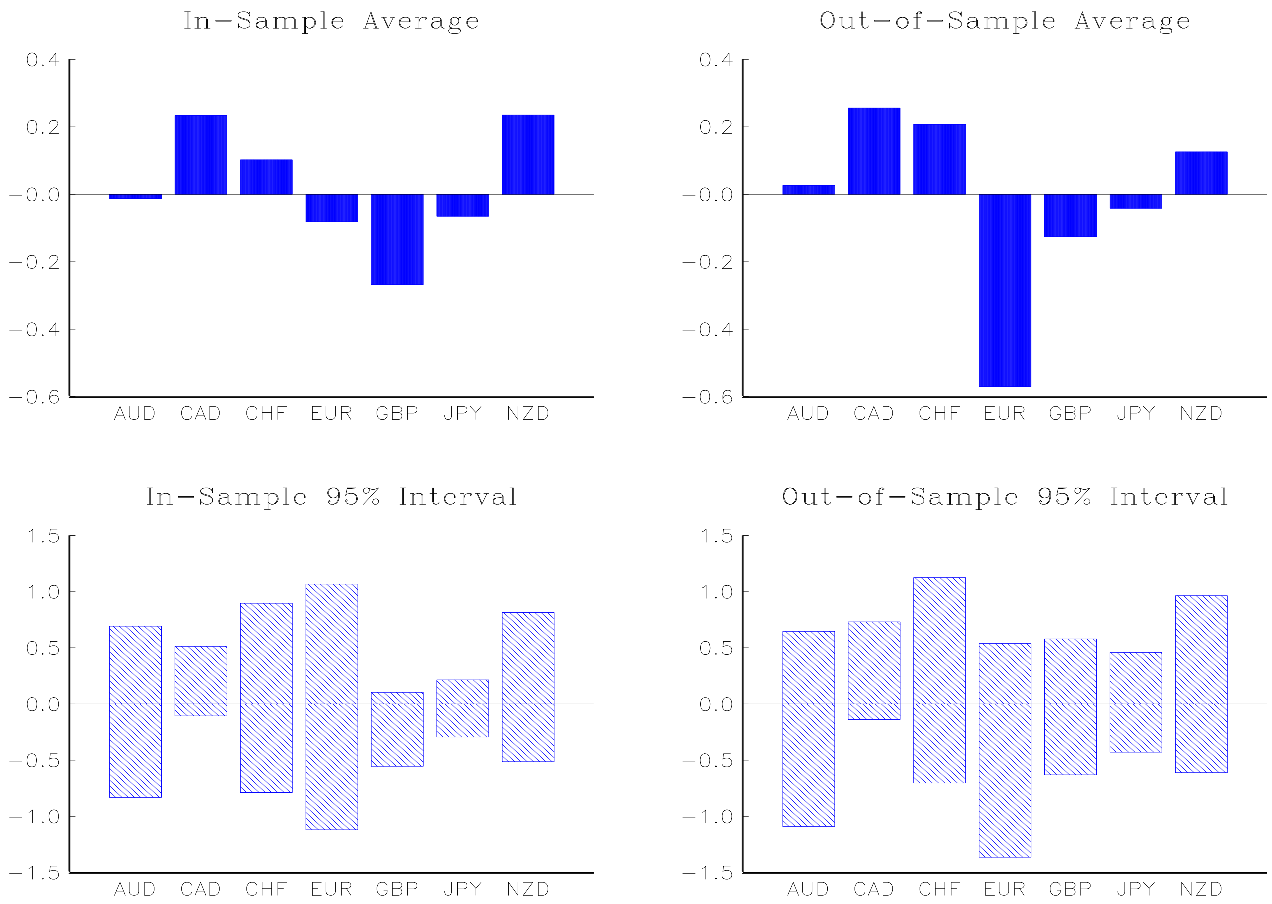

Figure 3. Portfolio Weights for the Carry Trade in Volatility

The figure displays the average portfolio weights and the 95% interval (range) for the 3-month carry trade in volatility. The carry trade in volatility strategy invests in a US bond and seven forward volatility agreements. The top left and bottom left panels are for the in-sample strategy, whereas the top right and bottom right panels are for the out-of-sample strategy.

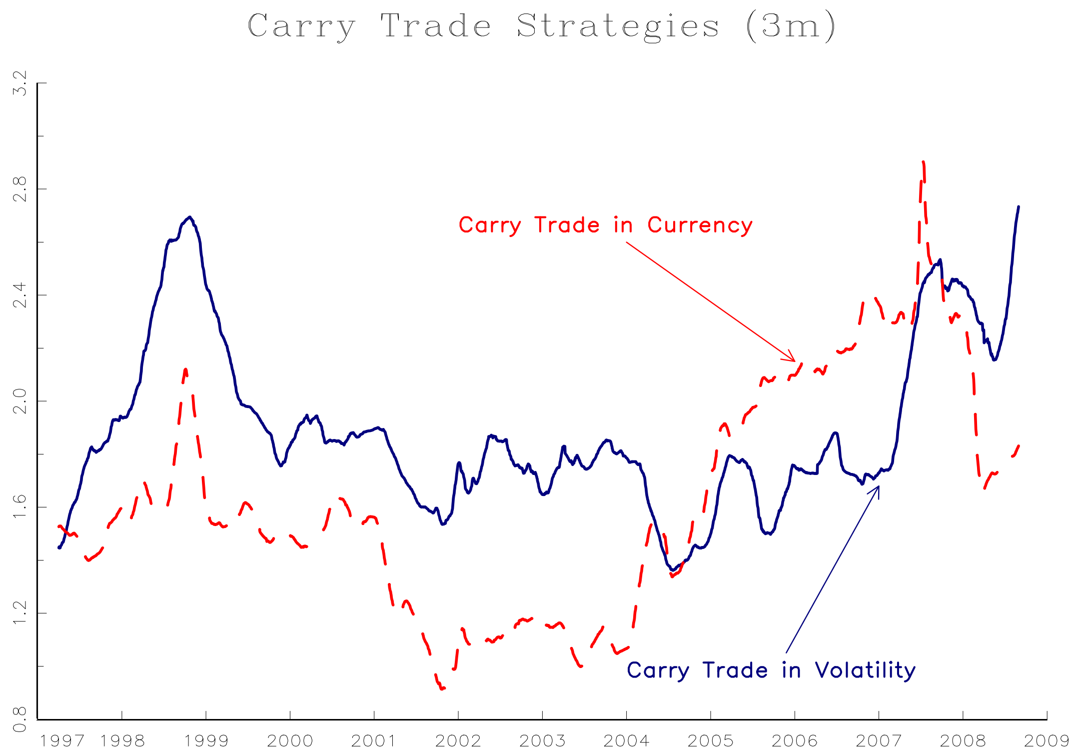

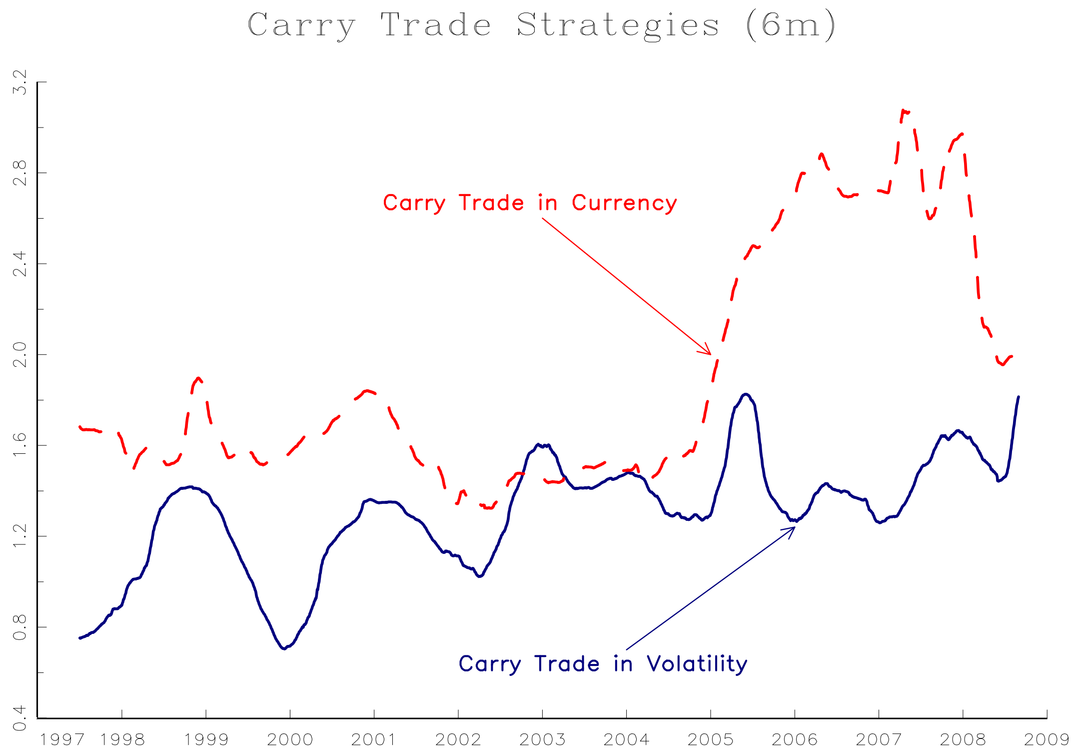

Figure 4. Rolling Out-of-Sample Sharpe Ratios

The figure displays the rolling out-of-sample annualized Sharpe Ratio for the carry trade in volatility and the carry trade in currency. The rolling Sharpe Ratio is computed using a three-year rolling window. The carry trade in volatility strategy invests in a US bond and seven forward volatility agreements. The carry trade in currency strategy invests in a US bond and seven forward exchange rates. The top panel shows the strategies for a 3-month maturity. The bottom panel shows the strategies for a 6-month maturity.

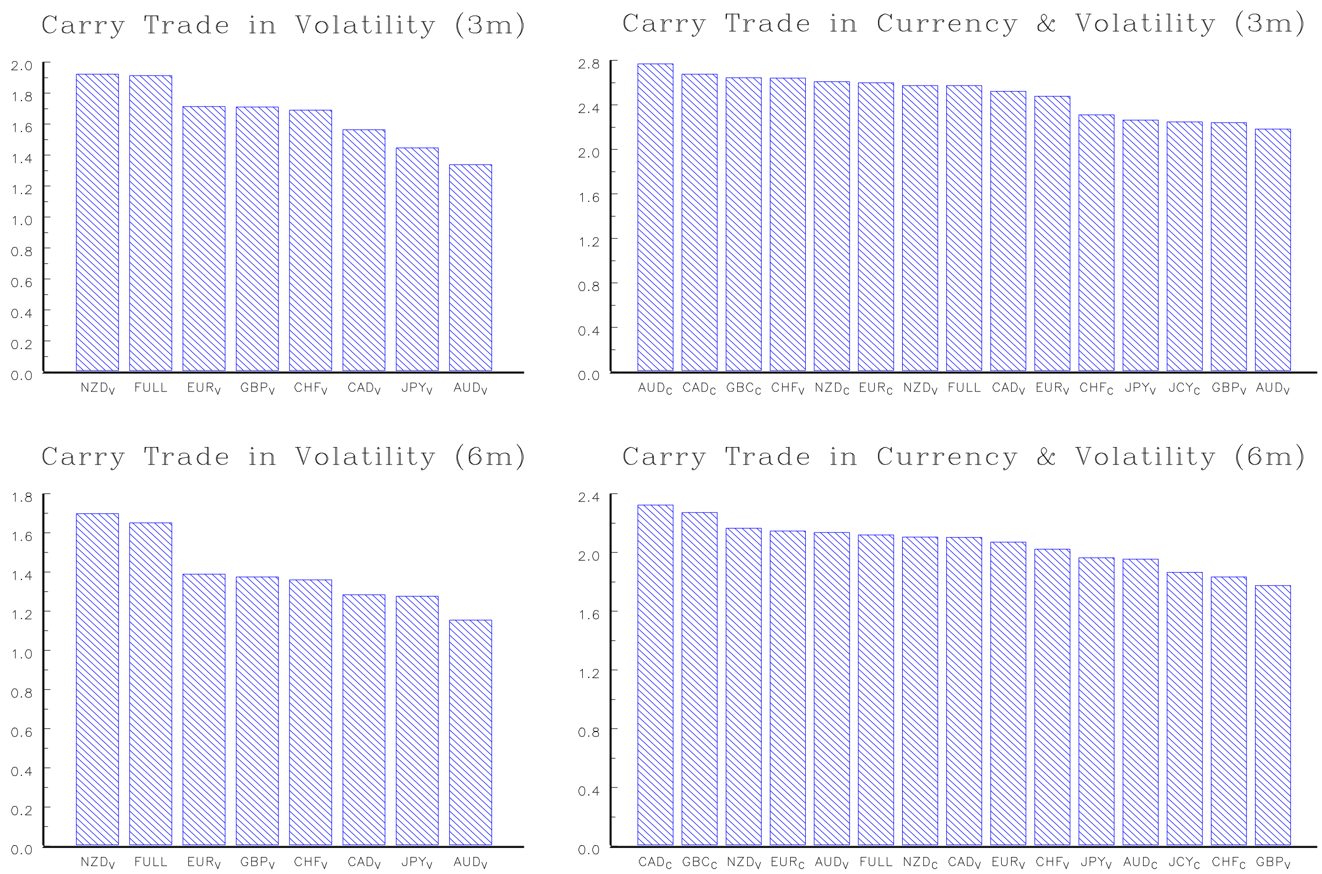

Figure 5. Out-of-Sample Sharpe Ratios

The figure displays the out-of-sample annualized Sharpe Ratio for the full portfolio and for portfolios where one of the risky assets is removed. The top left and bottom left panels show the carry trade in volatility strategy for 3-month and 6-month maturities. The carry trade in volatility strategy invests in a US bond and seven forward volatility agreements. The top right and bottom right panels show the combined carry trade in currency and volatility strategy for 3-month and 6-month maturities. The combined carry trade in currency and volatility strategy invests in a US bond, seven forward volatility agreements and seven forward exchange rates. For example, EUR_V is the portfolio excluding the FVA written on the euro and EUR_C is the portfolio excluding the euro forward exchange rate.

References

Backus, D.K., A.W. Gregory, and C.I. Telmer (1993). "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance* **48**, 1887–1908.

Backus, D.K., S. Foresi, and C.I. Telmer (2001). "Affine Term Structure Models and the Forward Premium Anomaly," *Journal of Finance* **56**, 279–304.

Bank of International Settlements (2007). "Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007," Monetary and Economic Department.

Bakshi, G., and D. Madan (2006). "A Theory of Volatility Spreads," *Management Science* **52**, 1945–1956.

Bansal, R. (1997). "An Exploration of the Forward Premium Puzzle in Currency Markets," *Review of Financial Studies* **10**, 369–403.

Bassett, G., and R. Koenker (1978). "Asymptotic Theory of Least Absolute Error Regression," *Journal of the American Statistical Association* **73**, 618–622.

Berger, D., A. Chaboud, E. Hjalmarsson, and E. Howorka (2008). "What Drives Volatility Persistence in the Foreign Exchange Market?" *Journal of Financial Economics* (forthcoming).

Bekaert, G. (1996). "The Time-Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective," *Review of Financial Studies* **9**, 427–470.

Bekaert, G., and R.J. Hodrick (2001). "Expectations Hypotheses Tests," *Journal of Finance* **56**, 115–138.

Bekaert, G., R.J. Hodrick, and D. Marshall (1997). "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics* **40**, 3–39.

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

Bollerslev, T., G. Tauchen, and H. Zhou (2008). "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies* (forthcoming).

Brockhaus, O., and D. Long (2000). "Volatility Swaps Made Simple," *Risk* **19**, 92–95.

Brunnermeier, M.K., S. Nagel, and L.H. Pedersen (2008). "Carry Trades and Currency Crashes" *NBER Macroeconomics Annual* (forthcoming).

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2008). "Do Peso Problems Explain the Returns to the Carry Trade?" NBER Working Paper 12489.

Campa, J.M, and P.H.K. Chang (1995). "Testing the Expectations Hypothesis of the Term Structure of Volatilities in Foreign Exchange Options," *Journal of Finance* **50**, 529–547.

Carr, P., and L. Wu (2007). "Stochastic Skew in Currency Options," *Journal of Financial Economics* **86**, 213–247.

Carr, P., and L. Wu (2009). "Variance Risk Premiums," *Review of Financial Studies* (forthcoming).

Della Corte, P., L. Sarno, and D.L. Thornton (2008). "The Expectation Hypothesis of the Term Structure of very Short-Term Rates: Statistical Tests and Economic Value," *Journal of Financial Economics* **89**, 158–174.

Della Corte, P., L. Sarno, and I. Tsiakas (2009). “An Economic Evaluation of Empirical Exchange Rate Models.” *Review of Financial Studies* (forthcoming).

Engel, C. (1996). “The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence,” *Journal of Empirical Finance* **3**, 123–192.

Engel, C., N. Mark, and K.D. West (2007). “Exchange Rate Models Are Not As Bad As You Think,” in D. Acemoglou, K. Rogoff and M. Woodford (eds.), *NBER Macroeconomics Annual* **22**, 381–441.

Fama, E.F. (1984). “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* **14**, 319–338.

Farhi, E., S.P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009). “Crash Risk in Currency Markets,” Unpublished working paper, Harvard University.

Fleming, J., C. Kirby, and B. Ostdiek (2001). “The Economic Value of Volatility Timing,” *Journal of Finance* **56**, 329–352.

Galati, G., A. Heath, and P. McGuire (2007). “Evidence of Carry Trade Activity,” *BIS Quarterly Review*, September, 27–41.

Galati, G., and M. Melvin (2004). “Why Has FX Trading Surged? Explaining the 2004 Triennial Survey,” *BIS Quarterly Review*, December, 67–74.

Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch (2007). “Portfolio Performance Manipulation and Manipulation-Proof Performance Measures,” *Review of Financial Studies* **20**, 1503–1546.

Han, Y. (2006). “Asset Allocation with a High Dimensional Latent Factor Stochastic Volatility Model,” *Review of Financial Studies* **19**, 237–271.

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

Jorion, P. (1995). “Predicting Volatility in the Foreign Exchange Market,” *Journal of Finance* **50**, 507–528.

Lustig, H., and A. Verdelhan (2007). “The Cross-Section of Currency Risk Premia and US Consumption Growth Risk,” *American Economic Review* **97**, 89–117.

Marquering, W., and M. Verbeek (2004). “The Economic Value of Predicting Stock Index Returns and Volatility,” *Journal of Financial and Quantitative Analysis* **39**, 407–429.

Meese, R.A., and K. Rogoff (1983). “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics* **14**, 3–24.

Poon, S.-H., and C. Granger (2003). “Forecasting Volatility in Financial Markets: A Review,” *Journal of Economic Literature* **41**, 478–539.

Poterba, J.M., and L.H. Summers (1986). “The Persistence of Volatility and Stock Market Fluctuations,” *American Economic Review* **76**, 1142–1151.

Verdelhan, A. (2009). “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance* (forthcoming).

Weiss, A.A. (1990). “Least Absolute Error Estimation in the Presence of Serial Correlation,” *Journal of Econometrics* **44**, 127–158.

West, K.D., H.J. Edison, and D. Cho (1993). "A Utility-Based Comparison of Some Models of Exchange Rate Volatility," *Journal of International Economics* **35**, 23–45.