



City Research Online

City, University of London Institutional Repository

Citation: Della Corte, P., Sarno, L. and Tsiakas, I. (2011). Spot and forward volatility in foreign exchange. *Journal of Financial Economics*, 100(3), pp. 496-513. doi: 10.1016/j.jfineco.2011.01.007

This is the unspecified version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/3303/>

Link to published version: <http://dx.doi.org/10.1016/j.jfineco.2011.01.007>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Spot and Forward Volatility in Foreign Exchange*

Pasquale Della Corte^a Lucio Sarno^{b,c} Ilias Tsiakas^{a,d}

a: Warwick Business School

b: Cass Business School

c: Centre for Economic Policy Research (CEPR)

d: University of Guelph

June 2010

Abstract

This paper investigates the empirical relation between spot and forward implied volatility in foreign exchange. We formulate and test the forward volatility unbiasedness hypothesis, which may be viewed as 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 delivery price of a forward contract on future spot implied volatility. This contract is known as a forward volatility agreement. We find strong evidence that forward implied volatility is a systematically biased predictor that overestimates movements in future spot implied volatility. This 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:** This paper is forthcoming in the *Journal of Financial Economics*. The authors are indebted for useful conversations or constructive comments to Bill Schwert (editor), an anonymous referee, Federico Bandi, John Bilson, Michael Brandt, Michael Brennan, Damiano Brigo, Andrea Buraschi, Peter Carr, Ales Cerny, Giovanni Cespa, Ines Chaieb, Peter Christoffersen, Gregory Connor, Stewart Hodges, Ravi Jagannathan, Philippe Jorion, Andrew Karolyi, Kan Li, Stewart Myers, Anthony Neuberger, Richard Payne, Fulvio Pegoraro, Giorgio Questa, Matthew Spiegel, Hassan Tehrani, Adrien Verdelhan, Liuren Wu and Shaojun Zhang as well as to participants at the 2009 European Finance Association Conference in Bergen, 2009 Northern Finance Association Conference in Niagara-on-the-Lake, 2009 Econometric Society European Meeting in Barcelona, 2009 Global Asset Management Conference at McGill University, 2009 China International Conference in Finance in Guangzhou, 2009 INQUIRE UK Autumn Workshop, 2009 Global Finance Academy Conference at University College Dublin, 2009 QASS Conference on Financial Econometrics at Queen Mary University of London; and seminars at the University of Wisconsin-Madison, HEC Paris, University of Amsterdam, Vienna Graduate School of Finance, Universitat Pompeu Fabra, University of Guelph, University of Manchester, University of Leicester, University of Brescia, IESEG School of Management, Audencia Nantes School of Management, the International Monetary Fund, the Deutsche Bundesbank, the Banque de France and the Central Bank of Hungary. We especially thank Bilal Hafeez of Deutsche Bank, Stephane Knauf of Citi and Nicola Mai of JP Morgan for providing some of the data and for numerous discussions on spot and forward implied volatility markets. Financial support from INQUIRE is gratefully acknowledged. The usual disclaimer applies. *Corresponding author:* Ilias Tsiakas, Warwick Business School, University of Warwick, Coventry CV4 7AL, UK. E-mail: ilias.tsiakas@wbs.ac.uk. *Other authors' contact details:* Pasquale Della Corte, pasquale.dellacorte@wbs.ac.uk; Lucio Sarno, lucio.sarno@city.ac.uk.

1 Introduction

The forward bias in foreign exchange (FX) arises from the well-documented empirical rejection of the Uncovered Interest Parity (UIP) condition, which suggests that the forward exchange rate is a biased predictor of the future spot exchange rate (e.g., Bilson, 1981; Fama, 1984; Engel, 1996). In practice, this means that high interest rate currencies tend to appreciate rather than depreciate. The forward bias implies that the returns to currency speculation are predictable, which generates high economic value to an investor designing a strategy exploiting the UIP violation, commonly referred to as the “carry trade” (Burnside, Eichenbaum, Kleshchelski and Rebelo, 2008; and Della Corte, Sarno and Tsiakas, 2009). Indeed, the carry trade is one of the most popular strategies in international asset allocation (e.g., Galati and Melvin, 2004; and Brunnermeier, Nagel and Pedersen, 2009).

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 a one dollar notional delivers the difference between future spot implied volatility and forward implied volatility. Therefore, given today’s information, the FVA determines the expected 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 paper investigates the empirical relation between spot and forward implied volatility in foreign exchange by formulating and testing the forward volatility unbiasedness hypothesis (FVUH). The FVUH postulates that forward implied volatility conditional on today’s information is an unbiased predictor of future spot implied volatility. Our analysis employs a new data set of daily implied volatilities for nine US dollar exchange rates quoted on over-the-counter (OTC) currency options spanning up to 14 years of data.¹ Using data for the implied volatility of options with different strikes, we compute the “model-free” implied volatility as in Britten-Jones and Neuberger (2000), Jiang and Tian (2005) and Carr and Wu (2009). The term structure of model-free spot implied volatility then allows for direct calculation of 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 analogue to the Fama (1984) predictive regression.

The results provide strong evidence that forward implied volatility is a systematically biased predictor that overestimates movements in future spot implied volatility. This is a new finding that 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., Jorion, 1995; Poon and Granger, 2003). Furthermore,

¹See, 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.

the rejection of forward volatility unbiasedness indicates the presence of conditionally positive, time-varying and predictable volatility term premiums in FX.

We assess the economic significance of the forward volatility bias in the context of dynamic asset 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 lower (higher) than current spot implied volatility will consistently generate excess returns over time. The framework for implementing the carry trade in volatility strategy is 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. 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) tend to be uncorrelated with the returns to currency speculation (carry trade in currency), which suggests that the source of the forward volatility bias may be unrelated to that of the forward bias.

As the objective of this paper is to provide an empirical investigation of the relation between spot and forward implied FX volatility, a number of questions fall beyond the scope of the analysis. First, we are not testing whether implied volatility is an unbiased predictor of future realized volatility (e.g., Jorion, 1995). As a result, we do not examine the volatility risk premium documented by the literature on the implied-realized volatility relation (e.g., Coval and Shumway, 2001; Bakshi and Kapadia, 2003; Low and Zhang, 2005; Carr and Wu, 2009; and Christoffersen, Heston and Jacobs, 2010). Instead, we focus on the spot-forward implied volatility relation and the volatility term premium that characterizes this distinct relation.² Second, we do not aim at offering a theoretical explanation for the forward volatility bias. In short, therefore, the main purpose of this paper is confined to establishing robust statistical and economic evidence on the forward volatility bias in the FX market.

An emerging literature indicates that volatility and the volatility risk premium are correlated with the equity premium. In particular, Ang, Hodrick, Xing and Zhang (2006) find that aggregate volatility risk, proxied by changes in the VIX index, is priced in the cross-section of stock returns as stocks with high exposure to innovations in aggregate market volatility earn low average future returns. Duarte and Jones (2007) focus on the volatility risk premium in the cross-section of stock

²The distinction between the volatility risk premium and the volatility term premium is well understood by practitioners. For example, Deutsche Bank has established two separate indices: (i) the Impact FX Volatility Index, which trades volatility swaps exploiting the volatility risk premium, and (ii) the FX Volatility Harvest Index, which trades FVAs exploiting the volatility term premium.

options and find that it varies positively with the VIX. Correlation risk is also priced in the sense that assets which pay off well when market-wide correlations are higher than expected earn negative excess returns (e.g., Driessen, Maenhout and Vilkov, 2009; Krishnan, Petkova and Ritchken, 2009). Turning to the FX market, recent research shows that global FX volatility is highly correlated with the VIX, and the VIX is correlated with the returns to the carry trade (e.g., Brunnermeier, Nagel and Pedersen, 2009). Finally, Knauf (2003) provides an excellent introduction to the FX volatility market and the use of FVAs as a convenient way of taking a view on FX volatility and exploiting the volatility curve. While Knauf (2003) is an important precursor to this paper, our setting is different and more general in that we analyze the relation between spot and forward implied volatility in the context of an unbiasedness condition, which we formally test in terms of both statistical and economic significance.

The remainder of the paper is organized as follows. In the next section we briefly review the literature on the forward unbiasedness hypothesis in FX. Section 3 formulates the FVUH, 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 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 the FX market, also known as the speculative efficiency hypothesis (Bilson, 1981), simply states that the forward exchange rate should be an unbiased predictor of the future spot exchange rate:

$$E_t S_{t+k} = F_t^k, \quad (1)$$

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 agreed at time t for an exchange of currencies at $t+k$.

The FUH can be equivalently represented as:

$$\frac{E_t S_{t+k} - S_t}{S_t} = \frac{F_t^k - S_t}{S_t}, \quad (2)$$

$$\frac{E_t S_{t+k} - F_t^k}{S_t} = 0, \quad (3)$$

where $\frac{E_t S_{t+k} - S_t}{S_t}$ is the expected spot exchange rate return, $\frac{F_t^k - S_t}{S_t}$ is the forward premium, and $\frac{E_t S_{t+k} - F_t^k}{S_t}$ is the expected return to currency speculation, which captures the return from issuing

a forward contract at time t and converting the proceeds into dollars at the spot rate prevailing at $t + k$, or vice versa (e.g., Hodrick and Srivastava, 1984; Backus, Gregory and Telmer, 1993). Equation (2) is the Uncovered Interest Parity (UIP) condition, which assumes risk neutrality and rational expectations and provides the economic foundation of the FUH. Under UIP, the forward premium is an unbiased predictor of the future rate of depreciation or, equivalently, the expected return to currency speculation in Equation (3) is equal to zero.³

Empirical testing of the FUH involves estimation of the following regression, which is commonly referred to as the “Fama regression” (Fama, 1984):

$$\frac{S_{t+k} - S_t}{S_t} = a + b \left(\frac{F_t^k - S_t}{S_t} \right) + u_{t+k}. \quad (4)$$

If the FUH holds, we should find that $a = 0$, $b = 1$, and the disturbance term $\{u_{t+k}\}$ is serially uncorrelated.⁴

Since the contribution of Bilson (1981) and Fama (1984), numerous empirical studies consistently reject the UIP condition (e.g., Hodrick, 1987; Engel, 1996; Sarno, 2005). 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,” which implies that high-interest currencies tend to appreciate rather than depreciate and forms the basis of the widely-used carry trade strategies in active currency management. In general, attempts to explain the forward bias using a variety of models have met with mixed success. Therefore, the forward bias continues to be heavily scrutinized in international finance research.⁵

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.

³In fact, the UIP condition is defined as $\frac{E_t S_{t+k} - S_t}{S_t} = \frac{i_t - i_t^*}{1 + i_t^*}$, where 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) implies: $\frac{F_t^k - S_t}{S_t} = \frac{i_t - i_t^*}{1 + i_t^*}$. It is straightforward to use these two equations to derive the version of the UIP condition defined in Equation (2).

⁴Note that the majority of the FX literature estimates the Fama regression in logs because it avoids the Siegel paradox (Siegel, 1972) and the distribution of returns may be closer to normal.

⁵See, 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 (2009); Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009); and Verdelhan (2009).

3.1 Forward Volatility Agreements

The forward IV of exchange rate returns represents the delivery price of a forward volatility agreement (FVA). The FVA is a forward contract on future spot IV with a payoff at maturity equal to:

$$\left(SV_{t+k} - FV_t^k\right) M, \quad (5)$$

where SV_{t+k} is the annualized spot IV observed at time $t+k$ and measured over the interval from $t+k$ to $t+2k$; FV_t^k is the annualized forward IV 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 $k = 1$ month implies that SV_{t+1} is the observed spot IV at time $t+1$ month for the interval of $t+1$ month to $t+2$ months; and FV_t^1 is the forward IV determined at time t for the interval of $t+1$ month to $t+2$ months. The FVA 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.⁶

3.2 Forward Implied Volatility

We begin our discussion of how we compute forward implied volatility by first determining the forward implied variance using a simple identity. By definition, variance is additive across time under i.i.d. innovations, and so is expected variance. In particular, the integrated variance between the current date t and a future date $t+2k$ for a risk-neutral exchange rate process S_τ can be decomposed as follows:

$$2k \int_t^{t+2k} \left(\frac{dS_\tau}{S_\tau}\right)^2 = k \int_t^{t+k} \left(\frac{dS_\tau}{S_\tau}\right)^2 + k \int_{t+k}^{t+2k} \left(\frac{dS_\tau}{S_\tau}\right)^2. \quad (6)$$

Taking the expectation at time t and simplifying gives:

$$2E_t \left[\int_t^{t+2k} \left(\frac{dS_\tau}{S_\tau}\right)^2 \right] = E_t \left[\int_t^{t+k} \left(\frac{dS_\tau}{S_\tau}\right)^2 \right] + E_t \left[\int_{t+k}^{t+2k} \left(\frac{dS_\tau}{S_\tau}\right)^2 \right]. \quad (7)$$

Britten-Jones and Neuberger (2000) demonstrate that the risk-neutral expectation of the integrated variance between two arbitrary dates is given by the “model-free” implied variance determined from the set of option prices expiring on these two dates. Hence we can replace the expected integrated variance by the model-free implied variance, which we define later. Equation (7) leads to the following relation for implied variances:

$$2SV_{t,t+2k}^2 = SV_{t,t+k}^2 + E_t [SV_{t+k,t+2k}^2] \quad (8)$$

$$= SV_{t,t+k}^2 + \left(FV_t^k\right)^2, \quad (9)$$

⁶It is straightforward to combine an FVA with a standard volatility swap in order to trade on the forward realized volatility for an interval starting in the future. This paper focuses on forward implied volatility.

where $SV_{t,t+k}^2$ and $SV_{t,t+2k}^2$ are the annualized implied variances for the intervals t to $t+k$ and t to $t+2k$, respectively, and $E_t [SV_{t+k,t+2k}^2] = (FV_t^k)^2$ is the forward implied variance determined at time t for the interval starting at time $t+k$ and ending at $t+2k$. Then, the forward implied variance is simply a linear combination of the spot implied variances:

$$(FV_t^k)^2 = 2SV_{t,t+2k}^2 - SV_{t,t+k}^2. \quad (10)$$

This approach is widely used in the literature (see, among others, Poterba and Summers, 1986; and Carr and Wu, 2009) and by investment banks in setting forward IV. For example, Equations (6)–(10) indicate that the 2-month spot implied variance is a simple average of the 1-month spot implied variance and the 1-month forward implied variance. The linear relation between implied variance and time across the term structure is also equivalent to the expectations hypothesis of the term structure of implied variance (Campa and Chang, 1995).

Our analysis focuses on forward implied volatility rather than forward implied variance, i.e. we are interested in $FV_t^k = E_t [SV_{t+k,t+2k}] = E_t [\sqrt{SV_{t+k,t+2k}^2}] \leq \sqrt{E_t [SV_{t+k,t+2k}^2]} = \sqrt{(FV_t^k)^2}$. Hence, Equation (10) implies:

$$FV_t^k \leq \sqrt{2SV_{t,t+2k}^2 - SV_{t,t+k}^2}. \quad (11)$$

This inequality is due to the convexity bias arising from Jensen's inequality since expected (implied) volatility is generally less than the square root of expected (implied) variance. For simplicity, we set:

$$FV_t^k = \sqrt{2SV_{t,t+2k}^2 - SV_{t,t+k}^2}, \quad (12)$$

and hence our empirical analysis is subject to the convexity bias. However, we deal with this approximation in two ways. First, we measure the convexity bias using a second-order Taylor expansion as in Brockhaus and Long (2000) and find that for our data it is empirically small.⁷ More importantly, we also provide empirical results showing that the spot-forward implied variance relation is qualitatively identical to the spot-forward implied volatility relation. Hence the convexity bias has no discernible effect on our results and the approximation in Equation (12) works well in our framework, which explains why it is widely used by practitioners (e.g., Knauf, 2003). We discuss these results in more detail later.

Equations (6)–(12) are cases where we have implied variances or implied volatilities defined over intervals of different length, and therefore we need to use two subscripts to clearly identify the start and end of the interval. From this point on, we revert back to using a single subscript, where for example SV_{t+k} is the annualized IV observed at time $t+k$ and measured over a set interval with length k .

⁷Brockhaus and Long (2000) show that $FV_t^k = E_t [\sqrt{SV_{t+k,t+2k}^2}] = \sqrt{E_t [SV_{t+k,t+2k}^2]} - \frac{\text{Var}_t [SV_{t+k,t+2k}^2]}{8\sqrt{E_t [SV_{t+k,t+2k}^2]}^3}$.

3.3 The Forward Volatility Unbiasedness Hypothesis

As any forward contract, the FVA's net market value at entry must be equal to zero. Therefore, its exercise price FV_t^k represents the risk-neutral expected value of SV_{t+k} (e.g., Carr and Wu, 2009):

$$E_t SV_{t+k} = FV_t^k. \quad (13)$$

This equation defines the Forward Volatility Unbiasedness Hypothesis (FVUH), which postulates that forward IV conditional on today's information set should be an unbiased predictor of future spot IV over the relevant horizon. The FVUH is based on risk neutrality and rational expectations, and can be thought of as the second-moment analogue of the FUH, which is based on the same set of assumptions.

The FVUH can be equivalently represented as:

$$\frac{E_t SV_{t+k} - SV_t}{SV_t} = \frac{FV_t^k - SV_t}{SV_t}, \quad (14)$$

$$\frac{E_t SV_{t+k} - FV_t^k}{SV_t} = 0, \quad (15)$$

where we define $\frac{E_t SV_{t+k} - SV_t}{SV_t}$ as the expected "implied volatility change," $\frac{FV_t^k - SV_t}{SV_t}$ as the "forward volatility premium," and $\frac{E_t SV_{t+k} - FV_t^k}{SV_t}$ as the expected "excess volatility return" from issuing an FVA contract at time t with maturity at time $t + k$.

The expected IV change has been studied by a large literature (Stein, 1989; Harvey and Whaley, 1991, 1992; Kim and Kim, 2003) and has a clear economic interpretation. Specifically, given that volatility is positively related to the price of an option, predictability in IV changes allows us to devise a profitable option trading strategy (regardless of whether this predictability is due to the forward volatility premium or not); for instance, if volatility is predicted to increase the option is purchased and vice versa (Harvey and Whaley, 1992).

The expected excess volatility return in Equation (15) can be interpreted as the expected return to volatility speculation. An FVA contract delivers a payoff at time $t + k$, but FV_t^k is determined at time t . Consider an investor who at time t buys a k -period FVA and saves in her bank account an amount $FV_t^k / (1 + i_t)$, where i_t is the k -period domestic nominal interest rate. At time $t + k$ the FVA matures and the investor withdraws the amount FV_t^k from her bank account and pays this amount in order to receive SV_{t+k} . This means that at time $t + k$ the investor will earn a total volatility return of $\frac{SV_{t+k} - SV_t}{SV_t}$ and an excess volatility return of $\frac{SV_{t+k} - FV_t^k}{SV_t}$.⁸ Under the FVUH, the

⁸The total return from investing in an FVA is $\frac{SV_{t+k} - FV_t^k / (1 + i_t)}{FV_t^k / (1 + i_t)}$, whereas the excess return is $\frac{SV_{t+k} - FV_t^k / (1 + i_t)}{FV_t^k / (1 + i_t)} - i_t = \frac{SV_{t+k} - FV_t^k}{FV_t^k / (1 + i_t)}$. Since under the FVUH, $SV_t = FV_t^k / (1 + i_t)$, the total return is equal to $\frac{SV_{t+k} - SV_t}{SV_t}$ and the excess return is equal to $\frac{SV_{t+k} - FV_t^k}{SV_t}$.

excess volatility return should be equal to zero. Equivalently, a rejection of the FVUH reflects the presence of a premium in the term structure of FX implied volatility.⁹

3.4 Model-Free Implied Variance

This section discusses the relation between volatility swaps and FVAs with particular reference to model-free implied variance. Specifically, the FVA is similar in structure to a volatility swap. While the FVA studied in this paper is a forward contract on future spot implied volatility, typically a volatility swap is a forward contract on future realized volatility. Variance and volatility swaps are valued by a replicating portfolio and hence this is also the case for FVAs. We first focus our discussion on variance swaps as they can be replicated more precisely than volatility swaps. The valuation of variance swaps will determine the fair delivery (exercise) price that makes the no-arbitrage initial value of the swap equal to zero. It can be shown that a variance swap can be replicated by the sum of (i) a dynamically adjusted constant dollar exposure to the underlying, and (ii) a combination of a static position in a portfolio of options and a forward that together replicate the payoff of a “log contract” (e.g., Demeterfi, Derman, Kamal and Zou, 1999; Windcliff, Forsyth and Vetzal, 2006; Broadie and Jain, 2008).¹⁰ The replicating portfolio strategy captures variance exactly provided that the portfolio of options contains all strikes in the appropriate weights to match the log payoff, and that the price of the underlying evolves continuously with constant or stochastic volatility but without jumps.

A key concept in understanding the pricing of variance swaps is model-free implied variance. Using no-arbitrage conditions under the assumption of a diffusion for the underlying price, Britten-Jones and Neuberger (2000) derive a model-free implied variance, which is fully specified by the set of option prices expiring on the future date. Jiang and Tian (2005) further demonstrate that the model-free implied variance is valid even when the underlying price exhibits jumps and also show that the approximation error is small in calculating the model-free implied variance for a limited range of strikes. More importantly, Jiang and Tian (2007) prove that the exercise price of a variance swap (i.e., the fair value of future variance developed by Demeterfi, Derman, Kamal and Zou, 1999) is exactly equal to the model-free implied variance formulated by Britten-Jones and Neuberger (2000). Therefore, computing and using model-free implied variance is equivalent to using the strike of a variance swap implied by the replicating portfolio.

⁹Similarly, Carr and Wu (2009) define the volatility risk premium as the difference between realized and implied volatility. Bollerslev, Tauchen and Zhou (2009) find that the 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.

¹⁰The log contract is an option whose payoff is proportional to the log of the underlying at expiration (Neuberger, 1994).

The implied volatility of currency options is a U-shaped function of moneyness, leading to the well-known volatility smile. The smile tends to increase the value of the fair variance above the at-the-money-forward (ATMF) implied variance level and the size of the increase will be proportional to factors such as time to maturity and the slope of the skew (e.g., Demeterfi, Derman, Kamal and Zou, 1999; Carr and Wu, 2007; and Bakshi, Carr and Wu, 2008). Using the model-free implied variance accounts directly for the volatility smile since its computation uses information on both ATMF IVs and IVs for alternative strikes.

Even though variance emerges naturally from hedged options, it is volatility that participants prefer to quote. Indeed, our empirical analysis focuses on forward volatility agreements not forward variance agreements. Volatility swaps are more difficult to replicate than variance swaps, as their replication requires a dynamic strategy involving variance swaps. The main complication in valuing volatility swaps is the convexity bias we have discussed above, which arises from the fact that the strike of a volatility swap is not equal to the square root of the strike of a variance swap due to Jensen's inequality. The convexity bias leads to misreplication when a volatility swap is replicated using a buy-and-hold strategy of variance swaps. Simply, the payoff of variance swaps is quadratic with respect to volatility, whereas the payoff of volatility swaps is linear. It can be shown that the replication mismatch is also affected by changes in volatility and the volatility of future volatility (e.g., Demeterfi, Derman, Kamal and Zou, 1999). Since our empirical analysis focuses on forward volatility agreements rather than forward variance agreements, it is subject to the convexity bias, which our empirical analysis will explicitly address in more detail later.

The implied volatilities we use in our empirical analysis are computed as the model-free implied volatilities of currency options. As we will see in the data section below, the availability of IV data is limited to five points, which is standard in the FX IV market (Carr and Wu, 2007): ATMF, 10-delta call, 10-delta put, 25-delta call and 25-delta put. We compute the model-free implied volatility by fitting a cubic spline around these five points. This interpolation method is standard in the literature (e.g., Bates, 1991; Campa, Chang and Reider, 1998; and Jiang and Tian, 2005). Curve-fitting using cubic splines has the advantage that the IV curve is smooth between the maximum and minimum available strikes, beyond which we extrapolate implied volatility by assuming it is constant as in Jiang and Tian (2005) and Carr and Wu (2009). This extrapolation method introduces an approximation error, which is shown by Jiang and Tian (2005) to be small in most empirical settings.¹¹

¹¹In recent years, IV indices are widely used among researchers and practitioners. For example, in stock markets the VIX index is based on the 1-month IV of the S&P 500, while in the FX market the VXY is based on the 3-month IV of options of the G-7 currencies, and the VXY-EM index is based on the 3-month IV of options of emerging market currencies. Finally, the Deutsche Bank FX Volatility Harvest Index is based on 6-month FVAs.

3.5 Predictive Regression for Exchange Rate Volatility

In order to test the empirical validity of the FVUH, we estimate the volatility analogue to the Fama regression:

$$\frac{SV_{t+k} - SV_t}{SV_t} = \alpha + \beta \left(\frac{FV_t^k - SV_t}{SV_t} \right) + \varepsilon_{t+k}. \quad (16)$$

Under the FVUH, $\alpha = 0, \beta = 1$ and the error term $\{\varepsilon_{t+k}\}$ is serially uncorrelated. It is straightforward to show that no bias in forward volatility implies no predictability in the excess volatility return.

There is a critical difference in the way we measure exchange rates in regression (4) versus volatilities in regression (16). 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 currency market and the volatility market. Note that the predictive regression (16) uses volatility changes as opposed to levels (i.e., the left-hand-side is $\frac{SV_{t+k} - SV_t}{SV_t}$ rather than SV_{t+k}) due to the high persistence in the level of FX volatility (e.g., Berger, Chaboud, Hjalmarsson and Howorka, 2009). This is an important consideration since performing ordinary least squares (OLS) estimation on very persistent variables (such as volatility levels) can cause spurious results, whereas OLS estimation on volatility changes avoids this concern. The same issue arises in the traditional FX market, which explains why the standard Fama regression is estimated using exchange rate returns, not exchange rate levels.¹²

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 regression (16). 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 regression (16) and uses the parameter estimates to predict the IV changes (from which we can also determine the excess volatility returns). We refer to the second model as the Forward Volatility Regression (FVR). We assess the significance of deviations from the FVUH simply by comparing the performance of the FVUH model with the FVR model under a variety of metrics, as described later.

4 Empirical Results on Forward Volatility Unbiasedness

4.1 Spot and Forward FX Implied Volatility Data

The OTC currency options market differs from an exchange-listed options market due to specific trading conventions. Currency options trade in terms of IV at a fixed delta rather than in terms

¹²We also estimate the volatility analogue to the log version of the Fama regression. Using logs makes the distribution of IV changes closer to normal. We find, however, that the predictive regression results for log IV changes are very similar to those for discrete IV changes. Hence our analysis focuses on the discrete version of the Fama regression (Equation 16).

of an option premium at a fixed strike price.¹³ The invoice price is then computed according to the Garman-Kohlhagen formula (Black-Scholes adjusted for the foreign interest rate). Specifically, IV quotes are available at five deltas in the form of delta-neutral straddle IV, 10-delta and 25-delta risk reversals, and 10-delta and 25-delta butterfly spreads. A straddle is a portfolio of a call and a put option with the same strike price and maturity. For a delta-neutral straddle (ST), the strike price needs to be sufficiently close to the forward price. This quote is referred to as ATMF IV ($IV_{0\delta ST} = IV_{ATMF}$). The risk reversal (RR) measures the difference in IV between an out-of-the-money call and an out-of-the-money put option with symmetric delta. The butterfly spread (BF) is equal to the average IV of an out-of-the-money call and an out-of-the-money put with symmetric delta minus the delta-neutral straddle IV. For example, $IV_{25\delta RR} = IV_{25\delta Call} - IV_{25\delta Put}$ and $IV_{25\delta BF} = 0.5 \times (IV_{25\delta Call} + IV_{25\delta Put}) - IV_{ATMF}$. From these quotes, it is straightforward to derive the implied volatilities at the five levels of delta. For further details on the currency option market, see Malz (1997), Campa, Chang and Reider (1998), and Carr and Wu (2007).

Our analysis employs a new data set of daily spot IVs for the 1-month and 2-month maturities quoted on OTC currency options for five strikes: ATMF, 10-delta call, 10-delta put, 25-delta call and 25-delta put. The data are collected from a panel of market participants and were made available to us by JP Morgan. These are high quality data involving quotes for contracts of at least \$10 million with a prime counterparty. Since the OTC currency options market is a very large and liquid market, OTC IVs are considered to be of higher quality than those derived from options traded in a particular exchange (e.g., Jorion, 1995).¹⁴

The IV data sample focuses on nine 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), the Norwegian kroner (NOK), the New Zealand dollar (NZD) and the Swedish kronor (SEK). The data sample begins in January 1996 and ends in September 2009 (3571 observations), except for EUR that begins in January 1999 (2804 observations). The analysis excludes all trading days that occur on a national US holiday. For each day of the sample, we calculate the model-free 1-month spot and forward IV as described in Section 3.

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 revolving around 10%

¹³The delta of a currency option represents its sensitivity to changes in the spot exchange rate. A 25-delta call corresponds to a call option with a Garman-Kohlhagen delta of 0.25, and a 25-delta put corresponds to a put option with a Garman-Kohlhagen delta of -0.25. In absolute value, delta is related to the probability that the option will expire in the money. For instance, a 25-delta call option is an out-of-the-money option while a 10-delta call option is a deep out-of-the-money option. The delta of an at-the-money option is equal to 0.50.

¹⁴More 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 for International Settlements, 2007).

per annum with a standard deviation of about 3% per annum. IV levels exhibit positive skewness, high kurtosis and are highly serially correlated, even at very long lags. The augmented Dickey-Fuller (ADF) statistic in most cases rejects the null of non-stationarity, although for some IV series this is not the case, confirming the strong persistence in IV.

Table 2 reports descriptive statistics for the implied volatility change $((SV_{t+k} - SV_t) / SV_t)$, the forward volatility premium $((FV_t^k - SV_t) / SV_t)$, and the excess volatility return $((SV_{t+k} - FV_t^k) / SV_t)$. The table shows that the mean annualized volatility changes revolve mostly between -20% and $+20\%$ for a high standard deviation in the range of 20% - 50% . In most cases, the time series exhibit low skewness (positive or negative) and moderate kurtosis. Moreover, the ADF statistic now strongly rejects the null hypothesis of non-stationarity with high confidence. This provides a clear justification for running the predictive regression (16) on volatility changes rather than on volatility levels since there is stronger statistical evidence rejecting the non-stationarity of the former than the latter. In short, FVUH tests in changes are likely to be better behaved than in levels.

4.2 A Simple FVA Example

We now turn our attention to a concrete FVA example and consider an investor who on September 25, 2007 enters a 1-month FVA written on the dollar price of the euro (EUR) with a notional of $M = 1,000,000 USD$. Note that for this example we go back to using two subscripts to clearly identify the start and end of the volatility interval. Table 3 lists the Garman–Kohlhagen 1-month and 2-month IVs available on this date from Bloomberg at five fixed deltas: 10-delta put, 25-delta put, ATMF, 25-delta call and 10-delta call. The 1-month spot IV ($SV_{t,t+1}$) covers the period of September 25, 2007 to October 25, 2007, and the 2-month spot IV ($SV_{t,t+2}$) covers the period of September 25, 2007 to November 25, 2007. Given these quotes, we compute the model-free 1-month and 2-month spot IVs as in Jiang and Tian (2005) and Carr and Wu (2009), which turn out to be $SV_{t,t+1} = 6.930\%$ and $SV_{t,t+2} = 6.895\%$. It is then straightforward to plug these values into Equation (12) to compute the model-free 1-month forward IV (FV_t^1) that is known on September 25, 2007 and covers the period of October 25, 2007 to November 25, 2007. The model-free forward IV is the fair delivery price of the FVA and is equal to $FV_t^1 = 6.860\%$.

These figures suggest a downward-sloping volatility curve. In a real trade, the FVA delivery price is quoted with a bid-ask spread, which typically revolves around 0.5% for major currencies such as EUR. If the trader goes long the FVA, the contract will expire on October 25, 2007 and deliver a payoff equal to $(SV_{t+1,t+2} - FV_t^1 - 0.5\%) \times M$, where $SV_{t+1,t+2}$ is the model-free 1-month spot IV computed on October 25, 2007 that covers the period of October 25, 2007 to November 25, 2007. As seen in Table 3, it turns out that $SV_{t+1,t+2} = 7.750\%$. Hence, the FVA is cash-settled with a payoff of $390,000 USD$, corresponding to a 1-month excess return of $(SV_{t+1,t+2} - FV_t^1 - 0.5\%) / SV_{t,t+1} = 5.628\%$ and

a 1-month total return of $i_t + (SV_{t+1,t+2} - FV_t^1 - 0.5\%) / SV_{t,t+1} = 6.027\%$, where i_t is the 1-month US nominal interest rate.¹⁵

4.3 Predictive Regression Results

We test the empirical validity of the FVUH by estimating the forward volatility regression (FVR) in equation (16). Table 4 presents the results. The OLS parameter estimates are for IV changes that are measured over 1-month 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 and West (1987) standard errors. We also provide regression results for non-overlapping observations later in this section. The table presents results for the full sample of January 1996 to September 2009. The start of the sample coincides with the period when trading of volatility derivatives surged.¹⁶

Recall that for the FVUH to hold (and hence for forward IV to be an unbiased expectation of future spot IV) three conditions must be met in the FVR: the intercept must be zero ($\alpha = 0$), the slope must be unity ($\beta = 1$), and the disturbance term must be serially uncorrelated. We test the FVUH conditions on each parameter separately with appropriately defined t -statistics. The serial correlation in the error term is tested with a Ljung-Box (1978) statistic, which is applied to the regression residuals between 21 and 252 trading days to eliminate overlapping observations. To facilitate interpretation we also report p -values in all cases.

We first focus on the slope estimate of the FVR. We find that the OLS estimates of β are all positive but much lower than unity, ranging from 0.141 for AUD to 0.668 for SEK. Overall, the FVUH is rejected for eight of nine currencies, the only exception being the SEK. Turning to the intercept of the FVR, we find that the value of α consistently revolves around zero, and in most cases it is not significantly different from zero. Furthermore, the Ljung-Box (1978) statistic indicates that the regression residuals are highly serially correlated. Finally, the R^2 coefficient of the FVR ranges from 1% to 5%.¹⁷

¹⁵We are grateful to Stephane Knauf for providing insights and information on this FVA example.

¹⁶The first variance swap was reportedly traded in 1993 by UBS (see Carr and Lee, 2009). Trading in volatility derivatives took off in the aftermath of the LTCM meltdown in late 1998, when implied stock index volatility levels rose to unprecedented levels (e.g., Gatheral, 2006). Note that the Deutsche Bank FX Volatility Harvest index investing in FVAs is available since the end of 1996. Carr and Wu (2009) also start their empirical analysis of volatility swaps in 1996. However, in a previous draft of this paper we also report results for a subsample ranging from October 2003 to September 2009 as well as for different data sets obtained from Bloomberg and Deutsche Bank respectively. These results (available upon request) are qualitatively identical to the results reported in this paper.

¹⁷For robustness purposes, we obtain estimates of β using three alternative estimation methods. First, we account for small sample bias by computing a bias-corrected estimator, which is based on the moving blocks bootstrap (e.g., Gonçalves and White, 2005). Second, we perform least squares estimation using a 99% winsorized sample, which replaces the 1% largest outliers by the closest value in the sample. This estimator is robust to outliers and does not assume a symmetric distribution (e.g., Hasings, Mosteller, Tukey and Winsor, 1947). Third, following Carr and Wu (2009), we also carry out errors-in-variables estimation assuming that forward IV is observed with error and the true

In conclusion, the predictive regression results clearly demonstrate that forward IV is a biased predictor of future spot IV, leading to a firm statistical rejection of the FVUH. 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 IVs quoted on currency options.¹⁸

4.4 Robustness of the Predictive Regression Results

4.4.1 Non-Overlapping Observations

Our analysis has so far focussed on predictive regressions estimated on daily data. Using daily data maximizes the number of available observations but generates serial correlation in the error term due to the overlapping nature of IV changes. As mentioned before, for the results in Table 4 we do the following: (i) estimate the predictive regressions by OLS, which is unbiased in the presence of overlapping observations; and (ii) compute standard Newey-West (1987) standard errors that account for the serial correlation in IV changes.

It is important to note that the effect of overlapping observations on inference remains an open issue in the literature. On the one hand, Richardson and Smith (1991) show analytically the gain from using overlapping observations in simple regressions due to the reduction in the standard error of the estimator. On the other hand, Christensen and Prabhala (1998) show that overlapping observations for implied or realized volatility can possibly lead to unreliable and inconsistent OLS estimates. We assess the importance of these issues in our framework by reporting results for predictive regressions using non-overlapping monthly IV changes. Table 5 has the results.

Specifically, we compare the predictive regression results in Table 4 based on daily overlapping observations for 1-month IV changes to the results in Table 5 based on non-overlapping monthly observations for the same data. We find that the β coefficients are very similar for overlapping and non-overlapping observations. Notably, the p -values for the null of $\beta = 1$ increase for CHF (0.105) and EUR (0.090), while the FVUH continues to be supported for the SEK. Overall, however, there is still strong evidence rejecting the FVUH since six of nine currencies reject it with 95% confidence plus one more with 90% confidence. This suggests that the statistical evidence against the FVUH is mitigated, but cannot be fully explained, by the overlapping nature of the main data set we use.

value follows an AR(1) process. The results, not reported to conserve space, confirm that the FVUH is still rejected for eight of nine currencies, the only exception being the SEK.

¹⁸The IV quotes typically come from a poll of dealers. Averaging IV quotes across dealers is a source of measurement error, which is potentially severe in the presence of large outliers. We directly account for the effect of the distribution of volatilities across dealers on testing the FVUH by using a separate data set on IV quotes from six individual dealers. The data are taken from Bloomberg and are for six US dollar exchange rates over the shorter sample of December 2005 to September 2009. In unreported results, we find that the OLS estimates of β across dealers are very close to each other in size, sign and statistical significance. In light of this evidence, the forward volatility bias is unlikely to be explained by possible measurement error due to averaging of quotes across dealers.

4.4.2 Implied Variance Results

Most of our analysis focuses on implied volatilities rather than implied variances. Forward implied volatilities are computed as the square root of forward implied variances and hence are subject to the convexity bias due to Jensen’s inequality. We can eliminate this bias by testing for unbiasedness in the spot-forward implied variance relation using the same predictive regression framework.

The results in Table 6 demonstrate that forward implied variances are also biased predictors of future spot implied variances in similar magnitudes to implied volatilities: the β estimate ranges from 0.031 for AUD to 0.640 for SEK. Forward variance unbiasedness is again rejected in eight of nine cases, with SEK still being the single exception. These results indicate that the convexity bias is unlikely to affect the bias in the spot-forward volatility relation.

5 Economic Value of Volatility Speculation: The Framework

This section describes the framework we use in order to evaluate the performance of an asset allocation strategy that exploits predictability in the returns to FX volatility speculation.

5.1 The Carry Trade in Volatility Strategy

Consider a US investor who builds a portfolio by allocating her wealth between the domestic riskless asset and nine FVA contracts. The FVAs are written on nine US dollar nominal exchange rates: AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD and SEK. Note that the risky assets (i.e., the 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 giving a total return of $i_t + (SV_{t+k} - FV_t^k) / SV_t$ (which is also equal to $(SV_{t+k} - SV_t) / SV_t$). The return from domestic riskless investing is proxied by the daily 1-month US Eurodeposit rate.

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 IV. We consider two strategies for the conditional mean of the returns to volatility speculation based on the FVUH model and the FVR model. Throughout the analysis we do not model the dynamics of the conditional covariance matrix of the returns to volatility speculation. In this setting, the optimal weights will vary across the two models only to the extent that there are deviations from the FVUH. In particular, the FVR model exploits predictability in the returns to volatility speculation in the sense that we can use the predictive regression to provide the forecast $(E_t SV_{t+k} - FV_t^k) / SV_t$. In contrast, the FVUH benchmark model is equivalent to riskless investing since fixing $\alpha = 0, \beta = 1$ implies that the conditional expectation of excess volatility returns is equal

to zero: $(E_t SV_{t+k} - FV_t^k) / SV_t = 0$.

The investor rebalances her portfolio on a daily basis by taking a position on FX volatility over a horizon of one month 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.¹⁹

We refer to the dynamic strategy implied by the FVR model as the carry trade in volatility (CTV) strategy. The CTV strategy can be thought of as the volatility analogue to the traditional carry trade in currency (CTC) strategy studied, among others, by Burnside, Eichenbaum, Kleshchelski and Rebelo (2008) and Della Corte, Sarno and Tsiakas (2009). The only risk an investor following the CTV strategy is exposed to is FX volatility risk.

5.2 Mean-Variance Dynamic Asset Allocation

Mean-variance analysis is a natural framework for assessing the economic value of strategies that 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 on a daily basis 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; $\mu_{t+k|t} = E_t[r_{t+k}]$ 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:

$$\begin{aligned} \max_{w_t} \quad & \left\{ \mu_{p,t+k|t} = w_t' \mu_{t+k|t} + (1 - w_t' \iota) r_f \right\} \\ \text{s.t.} \quad & (\sigma_p^*)^2 = w_t' V_{t+k|t} w_t, \end{aligned} \tag{17}$$

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

¹⁹Normally, with daily rebalancing the portfolio from the previous day should be marked to market, which is not possible in our context since there are no quotes for 1 month minus 1 day. However, in the rebalancing framework described above this issue does not arise because the investor essentially trades 21 portfolios per month, one for each day, and holds the FVAs until expiry. Finally, we also examine a monthly rebalancing exercise using non-overlapping returns based on the predictive regressions reported in Table 5, and find very similar economic value results to the case of daily rebalancing.

the portfolio returns, and r_f is the return on the riskless asset. The solution to this optimization problem delivers the risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}} V_{t+k|t}^{-1} \left(\mu_{t+k|t} - \nu r_f \right), \quad (18)$$

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

$$R_{p,t+k} = 1 + r_{p,t+k} = 1 + (1 - w_t') r_f + w_t' r_{t+k}. \quad (19)$$

Note that we assume that $V_{t+k|t} = \bar{V}$, where \bar{V} is the unconditional covariance matrix of IV changes.

5.3 Performance Measure

We evaluate the performance of the CTV strategy relative to the FVUH benchmark using the Goetzmann, Ingersoll, Spiegel and Welch (2007) manipulation-proof performance measure defined as:

$$\Theta = \frac{1}{(1 - \gamma)} \ln \left[\frac{1}{T} \sum_{t=1}^{T-k} \left(\frac{R_{p,t+k}^*}{R_{p,t+k}} \right)^{1-\gamma} \right], \quad (20)$$

where $R_{p,t+k}^*$ is the gross portfolio return implied by the FVR model, $R_{p,t+k}$ is implied by the benchmark FVUH model, and γ may be thought of as the investor's degree of relative risk aversion (RRA).

As a manipulation-proof performance measure, Θ is attractive because it is robust to the distribution of portfolio returns and does not require the assumption of a utility function to rank portfolios. In contrast, the widely-used certainty equivalent return (e.g., Kandel and Stambaugh, 1996; Pastor and Stambaugh, 2000) and the performance fee (e.g., Fleming, Kirby and Ostdiek, 2001) assume a particular utility function. Θ can be interpreted as the annualized certainty equivalent of the excess portfolio returns and hence can be viewed as the maximum performance fee an investor will pay to switch from the FVUH to the FVR strategy. In other words, this criterion measures the risk-adjusted excess return an investor enjoys for conditioning on the forward volatility bias rather than assuming unbiasedness. We report Θ in annualized basis points (*bps*).

6 Economic Value of Volatility Speculation: The Results

We assess the economic value of the forward volatility bias by analyzing the performance of a dynamically rebalanced portfolio based on the CTV strategy relative to the FVUH benchmark. The economic evaluation is conducted both in sample and out of sample. The in-sample period ranges

from January 1996 to September 2009, except for EUR that starts in January 1999; the out-of-sample period starts at the beginning of the sample and proceeds forward by sequentially updating the parameter estimates of the predictive regression (16) day-by-day using a 3-year rolling window.²⁰

Our economic evaluation focuses on the manipulation-proof performance measure, Θ , which is reported in annualized *bps* for a target annualized portfolio volatility $\sigma_p^* = 10\%$ and $\gamma = 6$. The choice of σ_p^* and γ is reasonable and consistent with numerous empirical studies (e.g., Fleming, Kirby and Ostdiek, 2001; Marquering and Verbeek, 2004; Della Corte, Sarno and Thornton, 2008). We have experimented with different σ_p^* and γ values and found that qualitatively they have little effect on the asset allocation results discussed below.

In assessing the profitability of the dynamic CTV strategy, the effect of transaction costs is an essential consideration. For instance, if the bid-ask spread in trading FVAs is sufficiently high, the CTV strategy may be too costly to implement. We assess the effect of transaction costs on the economic value of volatility speculation by directly accounting for the quoted FVA bid-ask spread. In particular, we use 160 *bps* as the quoted FVA bid-ask spread throughout the sample. This corresponds to the highest average spread for a currency over this period. In general, the average bid-ask spread ranges from about 45 to 160 *bps*, but for major currencies it is about 50 *bps*.²¹

It is well-documented that the effective spread is generally lower than the quoted spread, since trading will take place at the best price quoted at any point in time, suggesting that the worse quotes will not attract trades (e.g., Mayhew, 2002; De Fontnouvelle, Fische and Harris, 2003; Battalio, Hatch and Jennings, 2004). Following Goyal and Saretto (2009), we consider effective transaction costs in the range of 50% to 100% of the quoted spread. We then follow Marquering and Verbeek (2004) by deducting the transaction cost from the excess volatility returns *ex post*. This ignores the fact that dynamic portfolios are no longer optimal in the presence of transaction costs but maintains simplicity and tractability in our analysis.

Table 7 reports the in-sample and out-of-sample portfolio performance. The results show that there is very high economic value associated with the forward volatility bias. We focus on the case when the effective spread is 75% of the quoted spread, which is a rather realistic case. Switching from the static FVUH to the CTV portfolio provides the following performance: (i) in-sample $\Theta = 1103$ annual *bps* and (ii) out-of-sample $\Theta = 1166$ *bps*. These results are also reflected in the Sharpe ratio net of transaction costs (SR), which for the CTV strategy is as follows: (i) in-sample $SR = 1.25$, and (ii) out-of-sample $SR = 1.30$. The economic value of volatility speculation remains high even

²⁰Note that we use a rolling estimate of the unconditional covariance matrix \bar{V} as we move through the out-of-sample period, conditioning only on information available at the time that forecasts are formed. This implies that the out-of-sample period starts in January 1999.

²¹The bid-ask spread will likely vary over time. However, as we only have data on the midquote of IVs we base our analysis on average bid-ask spread values, which were provided to us by Deutsche Bank.

when the effective spread is equal to the full quoted spread.

The portfolio weights on the risky assets (FVAs) required to generate this performance are quite reasonable. Figure 1 illustrates that the average weights for the CTV strategy revolve from around -0.30 to $+0.20$ in sample and from -0.50 to $+0.30$ out of sample. The figure also displays the 95% interval of the variation in the weights, which in most cases ranges between -1 and $+1$. In short, therefore, the CTV strategy vastly outperforms the FVUH while taking reasonable positions in the FVAs.

7 Robustness and Further Analysis

7.1 Carry Trade in Volatility vs. Carry Trade in Currency

This section discusses the robustness of the economic value results. To begin with, 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. In other words, it is interesting to determine 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 forward bias in the FX market and the FX options market may be potentially driven by the same underlying cause.

We address this issue by designing a dynamic CTC strategy 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 nine forward exchange rates. The nine forward rates are for the same exchange rates and the same sample period as the volatility speculation strategy investing in the nine FVAs. We then use the original Fama regression (Equation 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 CTC strategy for the same exchange rate sample. Note that for the CTC strategy we use the quoted bid-ask spread of 10 *bps*. We believe that this is reasonable (or perhaps slightly conservative) since professional investors face an average bid-ask spread of about 1-3 *bps*.

The simplest way of assessing the relation of the CTV strategy with the CTC strategy is to examine the correlation between their portfolio returns (net of the riskless rate). We compute this correlation and we find that in sample it is -0.02 , while the out-of-sample correlation is 0.01 .²² This suggests that the returns to the CTV and CTC strategies seem to be largely uncorrelated.

The time variation in the correlations between the CTV and CTC strategies is displayed in Figure 2. These correlations are computed using a three-year out-of-sample rolling estimation window. The

²²These values are for the case when the effective spread is 75% of the quoted spread but remain largely unchanged when the effective spread changes.

correlation is on average close to zero, although it varies noticeably over time. In the early 2000s it is significantly positive, in the mid-2000s it is close to zero and statistically insignificant, and in the late 2000s it is significantly negative.

A more involved way of addressing this issue is to compare the separate portfolio performance of each of the CTV and CTC strategies with that of a combined strategy. The combined portfolio is constructed by investing in the same US bond as before and 18 risky assets: the nine FVAs plus the nine forward exchange rates. Table 7 presents the results, which are indicative of the low correlation between the CTV and the CTC strategies. We focus on the out-of-sample results for a 75% effective spread. In examining each strategy separately, we observe that the CTV strategy has superior performance to the CTC strategy. The CTV strategy gives an out-of-sample Sharpe ratio of 1.30 versus 1.15 for the CTC strategy. The performance measure is 1166 *bps* and 999 *bps* respectively.²³ More importantly, however, the combined strategy performs better than the CTV strategy alone. As we move from the CTV strategy to the combined strategy, the Sharpe ratio rises from 1.30 to 1.94 and the performance measure increases from 1166 *bps* to 2211 *bps*. The substantial increase in economic value when combining CTV with CTC is evidence that there is distinct incremental economic value in the CTC over and above the economic value already incorporated in the CTV. We conclude that the forward volatility bias is largely distinct from the forward bias.

Finally, we turn to Figure 3, which illustrates the annualized out-of-sample Sharpe ratios for the CTV and CTC strategies. The figure shows that the Sharpe ratios tend to be uncorrelated for long periods of time. The CTV strategy tends to perform better at the beginning and end of the sample, whereas the CTC is better in the middle period. Moreover, it is interesting to note that for the last two years of the sample the Sharpe ratio of the CTV strategy is rising but that of the CTC is falling. This indicates that the CTV strategy 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 do not tend to be positively correlated with the returns to currency speculation even during the recent unwinding of the carry trade in currency.

7.2 Is Implied Volatility a Random Walk?

Given that the β estimate is much closer to zero (i.e., spot IV is a random walk) than unity (i.e., forward volatility unbiasedness), it would be interesting to determine whether in future work the random walk (RW) model for IV would be a sensible benchmark for assessing the economic value

²³It is worth noting that simple carry trades exploiting the forward bias in the traditional FX market have been very profitable over the years (e.g., Galati and Melvin, 2004; and Brunnermeier, Nagel and Pedersen, 2009). Our findings demonstrate that volatility speculation strategies can in fact be even more profitable than currency speculation strategies.

of predictability in the returns to volatility speculation.²⁴ The RW model is consistent with a simpler version of the CTV strategy, where the investor goes long on FVAs when spot IV is higher than forward IV and vice versa rather than using the estimates of the predictive regression to form forecasts of future spot IV.²⁵

The portfolio performance of the RW without drift model which sets $\alpha = \beta = 0$ in the FVR (Equation 16) is presented in Table 8. The table shows that the in-sample and out-of-sample economic value of the RW model is virtually identical to the CTV strategy. For example, consider the out-of-sample results when the effective spread is equal to 75% of the quoted spread. Then the RW generates $SR = 1.30$ and $\Theta = 1165$ *bps*, whereas the CTV strategy generates $SR = 1.30$ and $\Theta = 1166$ *bps*. Hence the economic value of the CTV strategy is practically indistinguishable from that of the RW suggesting that the RW is a useful benchmark to adopt in future studies of forecasting FX implied volatility.

7.3 Time-Varying Leverage

The excess volatility return to the CTV strategy $((SV_{t+1} - FV_t^1) / SV_t)$ has time-varying leverage because the FVA payoff $(SV_{t+1} - FV_t^1)$ is scaled by the initial implied volatility, SV_t . For example, consider the case where $SV_t = 10\%$ and leverage is 10. Then, if $SV_{t+1} = 20\%$, leverage drops to 5. In other words, for any particular FVA payoff at time $t + 1$, the return at time $t + 1$ also depends on SV_t because of scaling. This raises the question of whether part of the CTV profits presented in this paper are due to this time-varying leverage effect.

In order to address this issue, we carry out a robustness check where we avoid the scaling by working with payoffs instead of returns. In this case, we estimate the predictive regression $SV_{t+1} - SV_t = \alpha + \beta (FV_t^1 - SV_t) + \varepsilon_{t+1}$ and use these predictions to build an unscaled CTV strategy ($CTV_{Unscaled}$). We then compare the results to the scaled CTV analyzed until now (CTV_{Scaled}), which is based on estimating the predictive regression (16). Figure 4 plots the rolling Sharpe ratios of the scaled and unscaled CTV strategies and shows that there are minor differences in their performance mostly in the early 2000s. However, the two Sharpe ratios move closely together over the full sample and are virtually identical on average. In short, while the leverage effect due to scaling has some impact on the time-variation of excess volatility returns, it is not a key driver of the economic value of the CTV strategy.

²⁴Indeed, the majority of studies in the traditional FX market tend to use the random walk of Meese and Rogoff (1983) as the benchmark model, not forward unbiasedness.

²⁵According to the RW for spot IV, the best predictor of SV_{t+k} is SV_t . Consider an investor who goes long on an FVA when $SV_t > FV_t^k$ and short on an FVA when $FV_t^k > SV_t$. The conditional return of this strategy is $(SV_{t+k} - FV_t^k) / SV_t \times \text{sign}(SV_t - FV_t^k)$. If spot IV follows a RW, this will be a profitable strategy.

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, which is the expectation of spot implied volatility for an interval starting at a future date. However, if there is a bias in the way the market sets forward implied volatility from quotes of spot implied volatility across the term structure, then the returns to volatility speculation will be predictable and a carry trade in volatility strategy can be profitable. Still, little is known about the empirical issues surrounding FVAs. These include the empirical properties of FVAs (e.g., their risk-return tradeoff), the extent to which forward implied volatility is a biased predictor of future spot implied 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 provide several insights. First, we find clear statistical evidence that forward implied volatility is a systematically biased predictor that overestimates movements in future spot implied volatility. This is similar to the tendency of the forward premium 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 forward volatility unbiasedness indicates the presence of conditionally positive, time-varying and predictable volatility term premiums (excess volatility returns) in foreign exchange. Third, there is 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.

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. The present study establishes the volatility analogue to the forward bias puzzle and demonstrates 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 Implied Volatility

The table reports descriptive statistics for the daily model-free spot and forward implied volatility (IV) on nine US dollar exchange rates for 1-month and 2-month maturities. The means and standard deviations are reported in annualized percent units. ρ_l is the autocorrelation coefficient for a lag of l trading days. ADF is the augmented Dickey-Fuller statistic for the null hypothesis of non-stationarity. The superscripts *a*, *b*, and *c* indicate statistical significance at the 10%, 5%, and 1% level, respectively. The sample ranges from January 1996 to September 2009 for all currencies, except for EUR that starts in January 1999.

	<i>Mean</i>	<i>Std</i>	<i>Skew</i>	<i>Kurt</i>	ρ_1	ρ_{21}	ρ_{63}	ρ_{126}	ρ_{252}	<i>ADF</i>
<i>AUD</i>										
<i>1m Spot IV</i>	11.93	4.95	2.91	15.69	0.99	0.87	0.63	0.38	0.23	-4.08 ^c
<i>2m Spot IV</i>	11.79	4.59	2.69	13.59	0.99	0.90	0.68	0.43	0.24	-3.73 ^c
<i>1m Forward IV</i>	11.63	4.23	2.42	11.32	0.99	0.92	0.73	0.49	0.26	-3.31 ^b
<i>CAD</i>										
<i>1m Spot IV</i>	8.23	3.78	2.33	10.42	0.99	0.93	0.72	0.57	0.59	-2.97 ^b
<i>2m Spot IV</i>	8.15	3.61	2.22	9.55	0.99	0.94	0.76	0.61	0.63	-2.69 ^a
<i>1m Forward IV</i>	8.07	3.45	2.12	8.81	0.99	0.95	0.80	0.66	0.66	-2.34
<i>CHF</i>										
<i>1m Spot IV</i>	10.98	2.47	1.58	8.92	0.98	0.79	0.60	0.36	0.09	-3.50 ^c
<i>2m Spot IV</i>	11.06	2.34	1.42	8.52	0.99	0.83	0.64	0.41	0.09	-3.23 ^b
<i>1m Forward IV</i>	11.12	2.23	1.23	7.97	0.99	0.86	0.68	0.45	0.09	-2.89 ^b
<i>EUR</i>										
<i>1m Spot IV</i>	10.76	3.38	1.93	8.96	0.99	0.89	0.65	0.38	0.12	-2.88 ^b
<i>2m Spot IV</i>	10.83	3.22	1.78	8.26	0.99	0.91	0.69	0.42	0.12	-2.73 ^a
<i>1m Forward IV</i>	10.89	3.08	1.62	7.57	0.99	0.92	0.72	0.47	0.13	-2.61 ^a
<i>GBP</i>										
<i>1m Spot IV</i>	9.19	3.31	3.21	16.05	0.99	0.88	0.64	0.34	0.13	-3.40 ^b
<i>2m Spot IV</i>	9.28	3.13	3.13	15.11	0.99	0.90	0.67	0.37	0.13	-3.10 ^b
<i>1m Forward IV</i>	9.36	2.97	3.04	14.20	0.99	0.92	0.70	0.40	0.14	-2.92 ^b
<i>JPY</i>										
<i>1m Spot IV</i>	11.67	3.80	1.82	8.44	0.98	0.80	0.63	0.47	0.32	-3.75 ^c
<i>2m Spot IV</i>	11.62	3.54	1.54	6.52	0.99	0.84	0.68	0.53	0.36	-3.32 ^b
<i>1m Forward IV</i>	11.56	3.30	1.24	4.71	0.99	0.88	0.74	0.58	0.40	-2.84 ^a
<i>NOK</i>										
<i>1m Spot IV</i>	11.61	3.75	2.51	11.33	0.99	0.88	0.66	0.43	0.21	-3.15 ^b
<i>2m Spot IV</i>	11.62	3.51	2.45	10.91	0.99	0.90	0.69	0.47	0.22	-3.12 ^b
<i>1m Forward IV</i>	11.63	3.28	2.37	10.38	0.99	0.92	0.73	0.52	0.22	-2.89 ^b
<i>NZD</i>										
<i>1m Spot IV</i>	12.83	4.65	1.89	9.12	0.99	0.87	0.70	0.47	0.29	-3.54 ^c
<i>2m Spot IV</i>	12.71	4.37	1.73	8.06	0.99	0.90	0.75	0.53	0.33	-3.12 ^b
<i>1m Forward IV</i>	12.57	4.13	1.55	7.07	0.99	0.93	0.80	0.60	0.38	-2.54
<i>SEK</i>										
<i>1m Spot IV</i>	11.62	3.84	2.51	10.78	0.99	0.89	0.68	0.48	0.22	-2.99 ^b
<i>2m Spot IV</i>	11.63	3.58	2.48	10.62	0.99	0.90	0.71	0.50	0.22	-2.95 ^b
<i>1m Forward IV</i>	11.62	3.32	2.43	10.35	0.99	0.92	0.74	0.53	0.22	-2.76 ^a

Table 2. Descriptive Statistics on Daily FX Implied Volatility Changes

The table displays descriptive statistics for the daily model-free spot and forward implied volatility changes on nine US dollar exchange rates for 1-month maturity. The *Implied Volatility Change* is defined as $(SV_{t+1} - SV_t)/SV_t$, where SV_t is the 1-month spot IV over the period t to $t+1$. The *Forward Volatility Premium* is defined as $(FV_t^1 - SV_t)/SV_t$, where FV_t^1 is the 1-month forward IV determined at time t for the period $t+1$ to $t+2$. The *Excess Volatility Return* is defined as $(SV_{t+1} - FV_t^1)/SV_t$. The means and standard deviations are reported in annualized percent units. ρ_l is the autocorrelation coefficient for a lag of l trading days. ADF is the augmented Dickey-Fuller statistic for the null hypothesis of non-stationarity. The superscripts *a*, *b*, and *c* indicate statistical significance at the 10%, 5%, and 1% level, respectively. The sample ranges from January 1996 to September 2009 for all currencies, except for EUR that starts in January 1999.

	<i>Mean</i>	<i>Std</i>	<i>Skew</i>	<i>Kurt</i>	ρ_1	ρ_{21}	ρ_{63}	ρ_{126}	ρ_{252}	<i>ADF</i>
<i>AUD</i>										
<i>Implied Volatility Change</i>	17.26	55.23	2.74	19.24	0.94	0.08	-0.06	-0.11	0.08	-7.36 ^c
<i>Forward Volatility Premium</i>	-16.26	21.58	-0.32	3.79	0.96	0.66	0.33	0.17	0.33	-5.19 ^c
<i>Excess Volatility Return</i>	33.52	58.18	3.01	20.79	0.96	0.32	0.05	-0.05	0.09	-6.31 ^c
<i>CAD</i>										
<i>Implied Volatility Change</i>	17.58	48.47	1.89	10.86	0.94	0.09	-0.10	-0.14	0.04	-7.38 ^c
<i>Forward Volatility Premium</i>	-11.53	16.79	-0.95	5.38	0.94	0.57	0.16	0.01	0.03	-5.34 ^c
<i>Excess Volatility Return</i>	29.11	49.65	2.23	12.27	0.96	0.31	-0.03	-0.13	0.06	-6.58 ^c
<i>CHF</i>										
<i>Implied Volatility Change</i>	9.87	46.59	1.26	6.39	0.92	-0.20	0.06	-0.07	0.02	-9.30 ^c
<i>Forward Volatility Premium</i>	21.31	16.34	-0.26	3.92	0.93	0.49	0.23	0.08	0.12	-5.74 ^c
<i>Excess Volatility Return</i>	-11.44	45.89	1.48	7.08	0.94	0.06	0.11	-0.02	0.01	-7.48 ^c
<i>EUR</i>										
<i>Implied Volatility Change</i>	10.16	45.42	1.49	7.87	0.93	0.00	-0.02	-0.07	0.02	-7.02 ^c
<i>Forward Volatility Premium</i>	22.26	15.95	-0.06	4.84	0.94	0.51	0.14	-0.01	0.23	-5.02 ^c
<i>Excess Volatility Return</i>	-12.10	45.83	1.89	10.81	0.95	0.23	0.06	-0.03	0.01	-5.89 ^c
<i>GBP</i>										
<i>Implied Volatility Change</i>	16.23	54.61	2.06	11.13	0.95	-0.02	-0.02	-0.09	0.05	-8.19 ^c
<i>Forward Volatility Premium</i>	33.68	18.84	0.51	5.22	0.95	0.52	0.12	0.06	0.01	-6.63 ^c
<i>Excess Volatility Return</i>	-17.46	55.09	2.19	12.89	0.96	0.18	0.02	-0.11	0.07	-7.35 ^c
<i>JPY</i>										
<i>Implied Volatility Change</i>	19.25	63.30	1.86	11.30	0.91	-0.17	-0.01	0.02	-0.01	-8.85 ^c
<i>Forward Volatility Premium</i>	0.37	21.73	-0.46	4.07	0.95	0.59	0.33	0.29	0.23	-5.87 ^c
<i>Excess Volatility Return</i>	18.89	62.89	2.35	13.74	0.93	0.06	0.07	0.05	0.02	-7.21 ^c
<i>NOK</i>										
<i>Implied Volatility Change</i>	16.09	52.74	2.92	20.41	0.94	-0.03	-0.05	-0.08	0.04	-8.45 ^c
<i>Forward Volatility Premium</i>	14.25	18.39	-0.01	5.46	0.95	0.54	0.22	0.00	0.13	-6.19 ^c
<i>Excess Volatility Return</i>	1.84	52.42	3.09	21.30	0.95	0.18	-0.02	-0.07	0.04	-7.27 ^c
<i>NZD</i>										
<i>Implied Volatility Change</i>	19.17	52.82	1.82	10.87	0.94	-0.06	-0.03	-0.09	0.01	-8.16 ^c
<i>Forward Volatility Premium</i>	-13.59	22.33	-0.08	3.65	0.96	0.63	0.28	0.10	0.23	-5.15 ^c
<i>Excess Volatility Return</i>	32.76	54.20	2.05	11.67	0.96	0.22	0.09	-0.04	0.02	-6.46 ^c
<i>SEK</i>										
<i>Implied Volatility Change</i>	16.25	51.41	2.18	13.79	0.94	-0.03	0.01	-0.08	0.04	-8.54 ^c
<i>Forward Volatility Premium</i>	12.79	17.62	-0.36	4.34	0.94	0.49	0.18	0.01	0.11	-6.37 ^c
<i>Excess Volatility Return</i>	3.46	50.39	2.56	16.10	0.95	0.17	0.03	-0.06	0.05	-7.60 ^c

Table 3. An FVA Example

The table displays the implied volatility values underlying a 1-month EUR forward volatility agreement (FVA). The contract is entered on September 25, 2007 (*time t*) and expires on October 25, 2007 (*time t+1*). $SV_{t,t+1}$ is the 1-month spot implied volatility covering the period from September 25, 2007 to October 25, 2007. $SV_{t,t+2}$ is the 2-month spot implied volatility covering the period from September 25, 2007 to November 25, 2007. FV_t^1 is the 1-month forward implied volatility covering the period from October 25, 2007 to November 25, 2007. $SV_{t+1,t+2}$ is the 1-month spot implied volatility covering the period from October 25, 2007 to November 25, 2007. $SV_{t,t+1}$, $SV_{t,t+2}$ and FV_t^1 are known on September 25, 2007, whereas $SV_{t+1,t+2}$ is known on October 25, 2007. The implied volatility values are Garman–Kohlhagen quotes at fixed deltas, which are taken from Bloomberg and are expressed in annualized percent units. The model-free implied volatilities are computed as in Jiang and Tian (2005) and Carr and Wu (2009). The payoff of a long FVA is computed as $(SV_{t+1,t+2} - FV_t^1 - 0.5\%)$ assuming a notional amount of 1,000,000 USD and a bid-ask spread of 0.5%.

	<i>September 25, 2007</i>			<i>October 25, 2007</i>
	$SV_{t,t+1}$	$SV_{t,t+2}$	FV_t^1	$SV_{t+1,t+2}$
<i>Garman–Kohlhagen</i>				
<i>10-delta put</i>	7.050	7.125		8.350
<i>25-delta put</i>	6.650	6.650		7.950
<i>ATMF</i>	6.700	6.625		7.425
<i>25-delta call</i>	7.150	7.050		7.300
<i>10-delta call</i>	7.550	7.525		7.700
<i>Model-Free</i>	6.930	6.895	6.860	7.750
<i>FVA payoff</i>				390,000

Table 4. Predictive Regressions

The table presents the ordinary least squares estimates of the predictive regression (16) using model-free spot and forward implied volatility changes on nine US dollar exchange rates for 1-month maturity. t^β is the t -statistic for the null hypothesis $\beta = 1$. LB is the Ljung-Box (1978) statistic for the null hypothesis of no autocorrelation in the regression residuals between 21 and 252 trading days. R^2 is the coefficient of determination. Newey-West (1987) standard errors are reported in parentheses and p -values in brackets, which are computed using a lag equal to the number of overlapping periods plus 1. The superscripts a , b , and c indicate statistical significance at the 10%, 5%, and 1% level, respectively. The sample period comprises daily observations from January 1996 to September 2009, except for EUR that starts in January 1999. The data are obtained from JP Morgan.

	α	β	t^β	LB	R^2
<i>AUD</i>	0.016 ^a (0.009)	0.141 (0.227)	- 3.79 [<0.01]	639 [<0.01]	0.01
<i>CAD</i>	0.017 ^b (0.009)	0.294 (0.234)	- 3.01 [<0.01]	583 [<0.01]	0.01
<i>CHF</i>	-0.003 (0.009)	0.622 ^c (0.179)	- 2.12 [0.034]	606 [<0.01]	0.05
<i>EUR</i>	0.001 (0.012)	0.426 (0.289)	- 1.99 [0.046]	1031 [<0.01]	0.02
<i>GBP</i>	0.002 (0.013)	0.426 ^a (0.248)	- 2.31 [0.021]	745 [<0.01]	0.02
<i>JPY</i>	0.016 (0.011)	0.554 ^c (0.195)	- 2.28 [0.022]	421 [<0.01]	0.04
<i>NOK</i>	0.007 (0.010)	0.550 ^c (0.178)	- 2.52 [0.012]	610 [<0.01]	0.04
<i>NZD</i>	0.020 ^b (0.009)	0.352 ^b (0.146)	- 4.43 [<0.01]	610 [<0.01]	0.02
<i>SEK</i>	0.006 (0.010)	0.668 ^c (0.208)	- 1.60 [0.110]	572 [<0.01]	0.05

Table 5. Predictive Regressions for Non-Overlapping Observations

The table presents the ordinary least squares estimates of the predictive regression (16) for non-overlapping model-free spot and forward implied volatility changes on nine US dollar exchange rates for 1-month maturity. t^β is the t -statistic for the null hypothesis $\beta = 1$. LB is the Ljung-Box (1978) statistic for the null hypothesis of no autocorrelation in the regression residuals between 1 and 12 observations. R^2 is the coefficient of determination. Asymptotic standard errors are reported in parentheses and p -values in brackets. The superscripts a , b , and c indicate statistical significance at the 10%, 5%, and 1% level, respectively. The sample period comprises monthly observations from January 1996 to September 2009, except for EUR that starts in January 1999. The data are obtained from JP Morgan.

	α	β	t^β	LB	R^2
<i>AUD</i>	0.014 (0.012)	0.090 (0.190)	- 3.40 [<0.01]	17.3 [0.139]	0.01
<i>CAD</i>	0.016 (0.010)	0.273 (0.212)	- 2.96 [<0.01]	10.8 [0.546]	0.01
<i>CHF</i>	-0.004 (0.011)	0.641 ^c (0.221)	- 1.63 [0.105]	13.2 [0.352]	0.04
<i>EUR</i>	-0.001 (0.012)	0.586 ^b (0.242)	- 1.71 [0.090]	34.4 [<0.01]	0.04
<i>GBP</i>	0.002 (0.014)	0.392 ^a (0.225)	- 2.71 [<0.01]	2.5 [0.998]	0.01
<i>JPY</i>	0.016 (0.014)	0.508 ^b (0.217)	- 2.27 [0.025]	15.9 [0.194]	0.03
<i>NOK</i>	0.007 (0.012)	0.507 ^b (0.218)	- 2.26 [0.025]	4.5 [0.973]	0.03
<i>NZD</i>	0.019 ^a (0.012)	0.380 ^b (0.180)	- 3.45 [<0.01]	9.7 [0.642]	0.02
<i>SEK</i>	0.006 (0.012)	0.655 ^c (0.228)	- 1.51 [0.132]	10.2 [0.598]	0.04

Table 6. Predictive Regressions for Implied Variances

The table presents the ordinary least squares estimates of the predictive regression (16) using model-free spot and forward implied volatility changes on nine US dollar exchange rates for 1-month maturity. t^β is the t -statistic for the null hypothesis $\beta = 1$. LB is the Ljung-Box (1978) statistic for the null hypothesis of no autocorrelation in the regression residuals between 21 and 252 trading days. R^2 is the coefficient of determination. Newey-West (1987) standard errors are reported in parentheses and p -values in brackets, which are computed using a lag equal to the number of overlapping periods plus 1. The superscripts a , b , and c indicate statistical significance at the 10%, 5%, and 1% level, respectively. The sample period comprises daily observations from January 1996 to September 2009, except for EUR that starts in January 1999. The data are obtained from JP Morgan.

	α	β	t^β	LB	R^2
<i>AUD</i>	0.055 ^b (0.022)	0.031 (0.296)	- 3.27 [<0.01]	388 [<0.01]	0.01
<i>CAD</i>	0.053 ^c (0.020)	0.226 (0.260)	- 2.98 [<0.01]	487 [<0.01]	0.01
<i>CHF</i>	0.012 (0.020)	0.603 ^c (0.188)	- 2.12 [0.034]	568 [<0.01]	0.04
<i>EUR</i>	0.019 (0.029)	0.385 (0.326)	- 1.89 [0.059]	886 [<0.01]	0.01
<i>GBP</i>	0.030 (0.032)	0.376 (0.282)	- 2.21 [0.027]	554 [<0.01]	0.01
<i>JPY</i>	0.064 ^b (0.026)	0.439 ^a (0.227)	- 2.47 [0.013]	405 [<0.01]	0.01
<i>NOK</i>	0.036 (0.026)	0.524 ^c (0.199)	- 2.40 [0.017]	435 [<0.01]	0.02
<i>NZD</i>	0.061 ^c (0.021)	0.308 ^a (0.170)	- 4.08 [<0.01]	433 [<0.01]	0.01
<i>SEK</i>	0.034 (0.025)	0.640 ^c (0.233)	- 1.54 [0.123]	436 [<0.01]	0.03

Table 7. The Economic Value of Volatility Speculation

The table shows the in-sample and out-of-sample economic value of volatility speculation. The *Carry Trade in Volatility Strategy* conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and nine 1-month forward volatility agreements. The *Carry Trade in Currency Strategy* conditions on the forward bias by building an efficient portfolio investing in a US bond and nine 1-month forward exchange rates. The *Combined Strategy* conditions on both the forward bias and the forward volatility bias. Each strategy maximizes expected returns subject to a target volatility $\sigma_p^* = 10\%$. The benchmark strategy is riskless investing implied by unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR , respectively. Θ is the Goetzmann *et al.* (2007) performance measure, which is expressed in annual basis points and is for $\gamma = 6$. The results are reported net of the effective bid-ask spread, which is assumed to be equal to 50%, 75% and 100% of the quoted spread. The quoted spread is set to be equal to 160 basis points for trading forward volatility agreements and 10 basis points for trading spot and forward exchange rates. The sample period comprises daily observations from January 1996 to September 2009. The out-of-sample period proceeds forward using a 3-year rolling window. The data are obtained from JP Morgan.

<i>Strategy</i>	<i>In-Sample</i>				<i>Out-of-Sample</i>			
	μ_p	σ_p	SR	Θ	μ_p	σ_p	SR	Θ
<i>Effective Spread = 50% Quoted Spread</i>								
<i>Carry Trade in Volatility</i>	23.3	12.9	1.52	1441	21.3	12.5	1.68	1632
<i>Carry Trade in Currency</i>	11.2	10.1	0.73	446	19.6	13.3	1.22	1087
<i>Combined</i>	25.3	12.4	1.74	1655	36.7	14.8	2.26	2678
<i>Effective Spread = 75% Quoted Spread</i>								
<i>Carry Trade in Volatility</i>	19.9	12.9	1.25	1103	19.6	12.5	1.30	1166
<i>Carry Trade in Currency</i>	10.5	10.1	0.68	379	18.7	13.3	1.15	999
<i>Combined</i>	21.8	12.3	1.46	1322	31.9	14.7	1.94	2211
<i>Effective Spread = 100% Quoted Spread</i>								
<i>Carry Trade in Volatility</i>	16.5	12.9	0.99	764	14.9	12.5	0.92	698
<i>Carry Trade in Currency</i>	9.86	10.1	0.60	312	17.8	13.3	1.09	911
<i>Combined</i>	18.3	12.3	1.18	977	27.1	14.7	1.61	1740

Table 8. The Economic Value of Volatility Speculation for the Random Walk

The table shows the in-sample and out-of-sample economic value of volatility speculation when α and β are set equal to zero in the predictive regression (16). This is equivalent to implementing the naïve random walk model for implied volatility changes. The Random Walk strategy conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and nine 1-month forward volatility agreements. The strategy maximizes expected returns subject to a target volatility $\sigma_p^* = 10\%$. The benchmark strategy is riskless investing implied by unbiasedness. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR , respectively. Θ is the Goetzmann *et al.* (2007) performance measure, which is expressed in annual basis points and is for $\gamma = 6$. The results are reported net of the effective bid-ask spread, which is assumed to be equal to 50%, 75% and 100% of the quoted spread. The quoted spread is set to be equal to 160 basis points for trading forward volatility agreements. The sample period comprises daily observations from January 1996 to September 2009. The out-of-sample period proceeds forward using a 3-year rolling window. The data are obtained from JP Morgan.

<i>In-Sample</i>				<i>Out-of-Sample</i>			
μ_p	σ_p	SR	Θ	μ_p	σ_p	SR	Θ
<i>Effective Spread = 50% Quoted Spread</i>							
23.1	13.0	1.49	1404	25.2	13.3	1.63	1612
<i>Effective Spread = 75% Quoted Spread</i>							
19.6	13.0	1.22	1062	20.6	13.3	1.30	1165
<i>Effective Spread = 100% Quoted Spread</i>							
16.2	13.0	0.95	718	16.0	13.2	0.96	715




Figure 1. Portfolio Weights for Carry Trade in Volatility

The figure displays the average portfolio weights and the 95% confidence interval (range) for the 1-month carry trade in volatility strategy (CTV). The strategy conditions on the forward volatility bias by building an efficient portfolio investing in a US bond and nine 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.

Carry Trade Strategies: Correlation




Figure 2. Out-of-Sample Correlation of Carry Trade Strategies

The figure displays the correlation between the daily portfolio returns of the carry trade in volatility strategy (CTV) and the carry trade in currency strategy (CTC). The portfolio returns are generated using a three-year out-of-sample estimation window. The strategies assume that the effective bid-ask spread is equal to 75% of the quoted spread, where the latter is set to 160 basis points for the CTV strategy and 10 basis points for the CTC strategy. The shaded area indicates that the sample correlation is statistically different from zero with a 95% confidence interval. Statistical significance is assessed using the Fisher transformation.

Carry Trade Strategies: Sharpe Ratios




Figure 3. Out-of-Sample Sharpe Ratio of Carry Trade Strategies

The figure displays the annualized Sharpe ratio for the carry trade in volatility strategy (CTV) (solid line), and the carry trade in currency strategy (CTC) (dashed line). The portfolio returns are generated using a three-year out-of-sample estimation window. The strategies assume that the effective bid-ask spread is equal to 75% of the quoted spread, where the latter is set to 160 basis points for the CTV strategy and 10 basis points for the CTC strategy.

Carry Trade Strategies: Sharpe Ratios




Figure 4. Out-of-Sample Sharpe Ratio of Carry Trade Strategies

The figure displays the annualized Sharpe ratio for two carry trade in volatility (CTV) strategies. The solid line indicates the CTV strategy based on scaled returns (CTV_{Scaled}), whereas the dotted line refers to the CTV strategy based on unscaled returns ($CTV_{Unscaled}$). The CTV_{Scaled} strategy is based on the predictive regression $((SV_{t+1} - SV_t)/SV_t) = \alpha + \beta((FV_t^1 - SV_t)/SV_t) + \varepsilon_{t+1}$, while the $CTV_{Unscaled}$ strategy is based on the predictive regression $SV_{t+1} - SV_t = \alpha + \beta(FV_t^1 - SV_t) + \varepsilon_{t+1}$. The portfolio returns are generated using a three-year out-of-sample estimation window. The strategies assume that the effective bid-ask spread is equal to 75% of the quoted spread, where the latter is set to 160 basis points.

References

- Ang, A., R.J. Hodrick, Y. Xing, and X. Zhang (2006). “The Cross-Section of Volatility and Expected Returns,” *Journal of Finance* **61**, 259–299.
- 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.
- 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.
- Bakshi, G., P. Carr, and L. Wu (2008). “Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies,” *Journal of Financial Economics* **87**, 132–156.
- Bakshi, G., and N. Kapadia (2003). “Delta-Hedged Gains and the Negative Market Volatility Risk Premium,” *Review of Financial Studies* **16**, 527–566.
- Bakshi, G., and D. Madan (2006). “A Theory of Volatility Spreads,” *Management Science* **52**, 1945–1956.
- Bank for International Settlements (2007). “Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in 2007.”
- Bansal, R. (1997). “An Exploration of the Forward Premium Puzzle in Currency Markets,” *Review of Financial Studies* **10**, 369–403.
- Battalio, R.H., B. Hatch, and R. Jennings (2004). “Toward a National Market System for U.S. Exchange-Listed Equity Options,” *Journal of Finance* **59**, 933–962.
- Bates, D.S. (1991). “The Crash of ’87: Was it Expected? The Evidence from Options Markets,” *Journal of Finance* **46**, 1009–1044.
- 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.
- Berger, D., A. Chaboud, E. Hjalmarsson, and E. Howorka (2009). “What Drives Volatility Persistence in the Foreign Exchange Market?” *Journal of Financial Economics* **94**, 192–213.
- Bilson, J.F.O. (1981). “The ‘Speculative Efficiency’ Hypothesis,” *Journal of Business* **54**, 435–451.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). “Expected Stock Returns and Variance Risk Premia,” *Review of Financial Studies* **22**, 4463–4492.
- Britten-Jones, M., and A. Neuberger (2000). “Option Prices, Implied Price Processes, and Stochastic Volatility,” *Journal of Finance* **55**, 839–866.
- Broadie, M., and A. Jain (2008). “Pricing and Hedging Volatility Derivatives,” *Journal of Derivatives* **15**, 7–24.
- Brockhaus, O., and D. Long (2000). “Volatility Swaps Made Simple,” *Risk* **19**, 92–95.

- Brunnermeier, M.K., S. Nagel, and L.H. Pedersen (2009). “Carry Trades and Currency Crashes” *NBER Macroeconomics Annual 2008*, 313–347.
- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2008). “Do Peso Problems Explain the Returns to the Carry Trade?” NBER Working Paper 14054.
- 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.
- Campa, J.M., P.H.K. Chang, and R.L. Reider (1998). “Implied Exchange Rate Distributions: Evidence from OTC Option Markets,” *Journal of International Money and Finance* **17**, 117–160.
- Carr, P., and R. Lee (2009). “Volatility Derivatives,” *Annual Review of Financial Economics* **1**, 319–339.
- 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* **22**, 1311–1341.
- Christensen, B.J., and N.R. Prabhala (1998). “The Relation between Implied and Realized Volatility,” *Journal of Financial Economics* **50**, 125–150.
- Christoffersen, P., S. Heston, and K. Jacobs (2010). “Option Anomalies and the Pricing Kernel,” Working Paper, McGill University.
- Coval, J.D., and T. Shumway (2001). “Expected Option Returns,” *Journal of Finance* **56**, 983–1009.
- De Fontnouvelle, P., R.P.H. Fische, and J.H. Harris (2003). “The Behavior of Bid–Ask Spreads and Volume in Options Markets during the Competition for Listings in 1999,” *Journal of Finance* **58**, 2437–2464.
- 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* **22**, 3491–3530.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999). “A Guide to Volatility and Variance Swaps,” *Journal of Derivatives* **6**, 9–32.
- Driessen, J., P. Maenhout, and G. Vilkov (2009). “The Price of Correlation Risk: Evidence from Equity Options,” *Journal of Finance* **64**, 1377–1406.
- Duarte, J., and C.S. Jones (2007). “The Price of Market Volatility Risk,” Working Paper, University of Southern California.
- Engel, C. (1996). “The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence,” *Journal of Empirical Finance* **3**, 123–192.
- Fama, E.F. (1984). “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* **14**, 319–338.
- Farhi, E., S.P. Fraiburger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009). “Crash Risk in Currency Markets,” NBER Working Paper 15062.

- Fleming, J., C. Kirby, and B. Ostdiek (2001). “The Economic Value of Volatility Timing,” *Journal of Finance* **56**, 329–352.
- Galati, G., and M. Melvin (2004). “Why Has FX Trading Surged? Explaining the 2004 Triennial Survey,” *BIS Quarterly Review*, December, 67–74.
- Gatheral, J. (2006). *The Volatility Surface: A Practitioner’s Guide*, New York: Wiley.
- 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.
- Gonçalves, S., and H. White (2005). “Bootstrap Standard Error Estimates for Linear Regression,” *Journal of the American Statistical Association* **100**, 970–979.
- Goyal, A., and A. Saretto (2009). “Cross-Section of Option Returns and Volatility,” *Journal of Financial Economics* **94**, 310–326.
- Han, Y. (2006). “Asset Allocation with a High Dimensional Latent Factor Stochastic Volatility Model,” *Review of Financial Studies* **19**, 237–271.
- Harvey, C.R., and R.E. Whaley (1991). “S&P 100 Index Option Volatility,” *Journal of Finance* **46**, 1551–1561.
- Harvey, C.R., and R.E. Whaley (1992). “Market Volatility Prediction and the Efficiency of the S&P 100 Index Option Market,” *Journal of Financial Economics* **31**, 43–73.
- Hasings, C., F. Mosteller, J.W. Tukey, and C.P. Winsor (1947). “Low Moments for Small Samples: a Comparative Study of Order Statistics,” *Annals of Mathematical Statistics* **18**, 413–426.
- Hodrick, R.J. (1987). *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, London: Harwood.
- Hodrick, R.J., and S. Srivastava (1984). “An Investigation of Risk and Return in Foreign Exchange,” *Journal of International Money and Finance* **3**, 5–29.
- Jiang, G.J., and Y.S. Tian (2005). “The Model-Free Implied Volatility and Its Information Content,” *Review of Financial Studies* **18**, 1305–1342.
- Jiang, G.J., and Y.S. Tian (2007). “Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index,” *Journal of Derivatives* **14**, 35–60.
- Jorion, P. (1995). “Predicting Volatility in the Foreign Exchange Market,” *Journal of Finance* **50**, 507–528.
- Kandel, S., and R.F. Stambaugh (1996). “On the Predictability of Stock Returns: An Asset-Allocation Perspective,” *Journal of Finance* **51**, 385–424.
- Kim, M., and M. Kim (2003). “Implied Volatility Dynamics in the Foreign Exchange Markets,” *Journal of International Money and Finance* **22**, 511–528.
- Knauf, S. (2003). “Making Money from FX Volatility,” *Quantitative Finance* **3**, 48–51.
- Krishnan, C.N.V., R. Petkova, and P. Ritchken (2009). “Correlation Risk,” *Journal of Empirical Finance* **16**, 353–367.
- Ljung, G.M., and G.E.P. Box (1978). “On a Measure of a Lack of Fit in Time Series Models,” *Biometrika* **65**, 297–303.

- Low, B.S., and S. Zhang (2005). “The Volatility Risk Premium Embedded in Currency Options,” *Journal of Financial and Quantitative Analysis* **40**, 803–832.
- Lustig, H., and A. Verdelhan (2007). “The Cross-Section of Currency Risk Premia and US Consumption Growth Risk,” *American Economic Review* **97**, 89–117.
- Malz, A.M. (1997). “Estimating the Probability Distribution of the Future Exchange Rate from Option Prices,” *Journal of Derivatives* **5**, 18–36.
- 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.
- Mayhew, S. (2002). “Competition, Market Structure, and Bid-Ask Spreads in Stock Option Markets,” *Journal of Finance* **57**, 931–958.
- 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.
- Neuberger, A. (1994). “The Log Contract: A New Instrument to Hedge Volatility,” *Journal of Portfolio Management* **20**, 74–80.
- Newey, W.K., and K.D. West (1987). “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica* **55**, 703–708.
- Pastor, L., and R.F. Stambaugh (2000). “Comparing Asset Pricing Models: An Investment Perspective,” *Journal of Financial Economics* **56**, 335–381.
- 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.
- Richardson, M., and T. Smith (1991). “Tests of Financial Models in the Presence of Overlapping Observations,” *Review of Financial Studies* **4**, 227–254.
- Sarno, L. (2005). “Towards a Solution to the Puzzles in Exchange Rate Economics: Where Do We Stand?” *Canadian Journal of Economics* **38**, 673–708.
- Siegel, J.J. (1972). “Risk, Interest Rates and the Forward Exchange,” *Quarterly Journal of Economics* **86**, 303–309.
- Stein, J. (1989). “Overreactions in the Options Market,” *Journal of Finance* **44**, 1011–1023.
- Verdelhan, A. (2009). “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance* **65**, 123–146.
- 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.
- Windcliff, H., P.A. Forsyth, and K.R. Vetzal (2006). “Pricing Methods and Hedging Strategies for Volatility Derivatives,” *Journal of Banking and Finance* **30**, 409–431.