The Cross-Section of Currency Volatility Premia*

Pasquale Della Corte　　Roman Kozhan　　Anthony Neuberger

March 31, 2020

Abstract

We identify a global risk factor in the cross-section of implied volatility returns in currency markets. A zero-cost strategy that buys forward volatility agreements with downward sloping implied volatility curves and sells those with upward slopes – a volatility carry strategy – generates significant excess returns. The covariation with volatility carry returns fully explains the cross-sectional variation of our slope-sorted portfolios. The lower the slope, the more the forward volatility agreement is exposed to volatility carry risk.

Keywords: Currency Volatility Risk Premia, Forward Volatility Agreement, Foreign Exchange Volatility, Term Structure.

JEL Classification: F31, F37, G01, G11, G12, G13, G15.

*We are grateful to Bill Schwert (editor), an anonymous referee, Andrea Appeddu, Pierre Collin-Dufresne, Riccardo Colacito, Federico Gavazzoni, Stefano Giglio, Kristjan Kasikov, Philippos Kassimatis, Robert Kosowski, Marie Lambert, Lucio Sarno, Tarun Ramadorai, Adrien Verdelhan, Christian Wagner as well as participants at the 2018 Vienna Symposium on Foreign Exchange Markets, 2017 Annual Meeting of the American Finance Association, 2016 Financial Econometrics and Empirical Asset Pricing Conference in Lancaster, 2016 Asset Pricing Workshop in York, 2016 INQUIRE UK Business School Seminar in London, 2016 Annual Hedge Fund Conference in London, Finance Forum 2016 in Madrid, Vienna Symposium on Foreign Exchange Markets 2018 and several research seminars for providing helpful comments and suggestions. Pasquale Della Corte is with Imperial College Business School, Imperial College London, and the Centre for Economic Policy Research (CEPR). Email: `P.DellaCorte@imperial.ac.uk`. Roman Kozhan is with Warwick Business School, University of Warwick. Email: `Roman.Kozhan@wbs.ac.uk`. Anthony Neuberger (corresponding author) is with Cass Business School, City, University of London. Email: `Anthony.Neuberger.1@city.ac.uk`.

1. Introduction

Recent research has documented the existence of pronounced volatility risk premia, especially at short horizons, and shows that investors care about uncertainty shocks.[1] The properties of these premia in the foreign exchange market, the world's largest financial market, remain underexplored. This paper documents that a carry trade in volatility, a long-short strategy that buys implied volatility at a discount and sells implied volatility at a premium, generates a significant Sharpe ratio. This finding resembles the well-known carry trade strategy whereby an investor is long currencies at discount and short currencies in the foreign exchange market (e.g., Lustig, Roussanov, and Verdelhan 2011; Menkhoff, Sarno, Schmeling, and Schrimpf 2012). The volatility carry returns, however, are virtually uncorrelated with the traditional carry trade and other popular currency strategies.

There are several features of the foreign exchange volatility market that make it of particular interest to financial economists. The over-the-counter currency options market is large and liquid with a daily average turnover equal to \$254 billion as of April 2016 and a notional amounts outstanding of \$11.7 trillion as of June 2016 (BIS 2016a, b). A wide range of strikes and maturities is traded, so volatility risk premia across different currency pairs and across maturities can be computed with precision. The risks being priced and traded in the FX market are macroeconomic in nature (e.g., Gabaix and Maggiori 2015; Zviadadze 2017; Colacito, Croce, Gavazzoni, and Ready 2018), so the market provides an excellent testbed for investigating the link between volatility risk premia in financial markets and real macroeconomic variables.

We conduct our analysis by examining the profitability of trading strategies using Forward Volatility Agreements (FVAs) – forward contracts that deliver the difference between the implied volatility of an exchange rate observed on the maturity date and the forward implied

[1] The literature on volatility risk premia in the equity, fixed income, and currency markets includes, among many others, Coval and Shumway (2001); Bakshi and Kapadia (2003); Bollerslev, Tauchen, and Zhou (2009); Broadie, Chernov, and Johannes (2009); Carr and Wu (2009); Christoffersen, Heston, and Jacobs (2009); Bakshi, Panayotov, and Skoulakis (2011); Della Corte, Sarno, and Tsiakas (2011); Ammann and Buesser (2013); Kozhan, Neuberger, and Schneider (2013); Della Corte, Ramadorai, and Sarno (2016); Londono and Zhou (2016).

volatility determined at the inception date. Della Corte, Sarno, and Tsiakas (2011) demonstrate the existence of volatility risk premia in the FX market through an examination of the average returns on short-term FVAs. In this paper, we extend the analysis by examining the cross-sectional differences in these returns and by extending the analysis to longer-term contracts. Following the pioneering work of Lustig and Verdelhan (2007), we identify a common risk factor in the data by building monthly portfolios of forward volatility agreements sorted by their implied volatility slopes for a broad range of maturity combinations. The first portfolio contains the highest (positive) slope currencies, while the last contains the lowest (negative). Similar to the work of Lustig, Roussanov, and Verdelhan (2011), we find that the first two principal components of the forward volatility returns account for most of the time-series variation. The first principal component is a level factor and it is virtually equal to the average excess return on all forward volatility returns (unconditional volatility premium). The second principal component is a slope factor and is highly correlated with the returns on the volatility carry - a zero-cost strategy that goes long in the last portfolio and short in the first portfolio. The covariation with the volatility carry risk factor fully explains the cross-sectional variation of our FVA portfolios. The R^2 ranges from 73% to 99%. The pricing errors of volatility excess returns are jointly insignificant for all maturity contracts ranging from 1 to 24 months. Our paper is the first to document the common factor in the currency volatility returns.

Focusing on the term structure of volatility risk premia allows us to investigate whether the volatility shocks that investors are exposed to are transitory or permanent in nature. There is clear evidence from the equity index market that spot and forward volatility markets behave rather differently. Dew-Becker, Giglio, Le, and Rodriguez (2016) show that while unconditional spot variance risk premia in the S&P500 index market are large, forward premia are insignificant at maturities in excess of a month or two. We find that much the same is true in the FX market, insofar as unconditional risk premia are concerned. However, volatility carry behaves in a strikingly different way. While the average return does decline steeply with maturity, the Sharpe ratio barely changes and remains statistically significant for up to two years. This suggests that the expected returns to volatility carry is related to permanent volatility shocks.

Our empirical evidence is robust to a number of additional exercises. First, we show that traditional currency factors (i.e., dollar, carry, global imbalance, global FX volatility, and FX liquidity), Fama-French global equity risk factors, and futures VIX returns cannot explain the cross-sectional variation of our implied volatility portfolios returns. Second, our volatility carry returns remain economically significant after accounting for the average bid-ask spreads of forward volatility agreements. Third, we find that different methodologies for the construction of forward implied volatility returns do not alter our key results. Finally, our results work equally well for a cross-section of 20 developed and emerging market countries and for a subset of 10 developed countries.

Our paper builds on the recent line of research that seeks to explain currency risk premia in a cross-sectional asset pricing setting.[2] Lustig, Roussanov, and Verdelhan (2011) find that the carry factor is a major source of risk in the cross-section of currency portfolios sorted by forward premia. Menkhoff, Sarno, Schmeling, and Schrimpf (2012) find that currency excess returns provide compensation for exposure to global FX volatility risk. Della Corte, Riddiough, and Sarno (2016) provide evidence that exposure to countries' external imbalances explains the cross-sectional variation of currency excess returns. Gabaix and Maggiori (2015) and Colacito, Croce, Gavazzoni, and Ready (2018) provide a theoretical basis for these empirical findings.

Our paper also contributes to the literature on the term structure of the volatility risk premium more generally. Dew-Becker, Giglio, Le, and Rodriguez (2016), Eraker and Wu (2016), and Johnson (2017) show that volatility risk premia in the equity market are the largest for short maturities and decrease at longer horizons. We also contribute to this research by showing that although unconditional FX volatility risk premium exhibit a similar pattern, the volatility carry premium remains both statistically and economically large al all horizons and is related

[2] The literature on carry trade is vast and includes, among many others, Brunnermeier, Nagel, and Pedersen (2009), Della Corte, Sarno, and Tsiakas (2009), Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011), Jurek (2014), Lustig, Roussanov, and Verdelhan (2014), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2015), Colacito, Croce, Gavazzoni, and Ready (2018), Bekaert and Panayotov (2016), Colacito, Croce, Gavazzoni, and Ready (2018), and Richmond (2019).

to permanent volatility shocks. Also, our exercise speaks to the related literature on the time-varying nature of exposure to volatility risk. The volatility risk premium varies with the level of volatility and market conditions (e.g., Bakshi and Kapadia 2003; Bakshi and Madan 2006; Todorov 2016; Aït-Sahalia, Karaman, and Mancini 2016; Barras and Malkhozov 2016). We show that exposure to the global risk factor that drives the local volatility risk premia co-varies with the slope of the implied volatility curve.

The rest of this paper is organized as follows. Section 2. sets the framework for our paper and describes the data set we use. Section 3. documents cross-sectional properties of the FX volatility risk premia and shows that a single factor, VCA, explains most of the cross-sectional variation in volatility excess returns. Section 4. contains a battery of robustness checks, and Section 5. concludes, identifying the implications of our findings for asset pricing theory. A separate Internet Appendix provides additional robustness tests and supporting analysis.

2. The Term Structure of Volatility Risk Premia

A natural way to trade the term structure of volatility risk premium in the FX market is through the use of forward volatility agreements (FVAs). They are over-the-counter (OTC) derivatives that allow traders to take positions on the future level of implied volatility. We show how to synthesize these agreements using quoted currency options and present empirical evidence on the behavior of volatility risk premia based on a large cross-section of currency pairs and different maturity combinations. This analysis motivates our key contribution reported in the following sections.

2.1. Forward Volatility Agreement

An FVA is a forward contract on the future implied volatility of a given exchange rate. The pay-off to the FVA is

$$\left(SVOL_{t+\tau_1}^{\tau_2} - FVOL_{t,\tau_1}^{\tau_2}\right) \times M. \tag{1}$$

The contract is written at time t, and matures at time $t + \tau_1$. The floating leg of the contract,

$SVOL_{t+\tau_1}^{\tau_2}$, is equal to the implied volatility at maturity for some specified horizon τ_2. The strike price or fixed leg of the contract, $FVOL_{t,\tau_1}^{\tau_2}$, is determined at time t. It is the forward implied volatility at time t for the period $(t + \tau_1, t + \tau)$, where $\tau = \tau_1 + \tau_2$. M denotes the notional dollar amount of the contract.

The time line is shown in Figure 1. It is worth noting that that the contract is a contract on implied volatility and not on realized volatility.

FIGURE 1 ABOUT HERE

We calculate the payoff to FVAs of different maturities over our sample period by computing the spot and forward implied volatilities from OTC currency options. Details are in the Appendix A.. We use the model-free approach of Britten-Jones and Neuberger (2000) to calculate implied variances. Following Della Corte, Ramadorai, and Sarno (2016), we calculate the implied volatility by taking the square root of the implied variance. This approach is subject to convexity bias. In our empirical analysis, we show that the impact of the convexity bias is negligible.

2.2. Currency Options Data

We collect daily OTC option implied volatilities on exchange rates vis-à-vis the US dollar from JP Morgan and Bloomberg. We use monthly data by sampling end-of-month implied volatilities from January 1996 to December 2015. Our sample includes 20 developed and emerging market countries: Australia, Brazil, Canada, Czech Republic, Denmark, Euro Area, Hungary, Japan, Mexico, New Zealand, Norway, Poland, Singapore, South Africa, South Korea, Sweden, Switzerland, Taiwan, Turkey, and United Kingdom. It starts with 9 currencies at the beginning of the sample in 1996 and ends with 20 currencies at the end of the sample in 2015.

OTC currency options are quoted in terms of Garman and Kohlhagen (1983) implied volatilities at fixed deltas (at-the-money, 10 delta call and put, and 25 delta call and put options) and fixed maturities. To convert deltas into strike prices and implied volatilities into option prices, we employ spot and forward exchange rates from Barclays and Reuters via Datastream, and interest rates from JP Morgan and Bloomberg.[3] This recovery exercise yields data on plain-vanilla European calls and puts for currency pairs vis-à-vis the US dollar for the following maturities: 1 month, 3 months, 6 months, 12 months and 24 months. We then construct spot and forward implied volatilities using the methodology presented above.

Although our main focus is on forward volatility risk premia, we will also analyze spot volatility risk premia extracted from volatility swaps. These premia are computed as the difference between the 1-month spot implied volatility at the end of month t and the realized volatility measured over the next following month. The realized volatility is computed using daily returns on forward exchange rates.

2.3. Testing for Volatility Risk Premia

Armed with spot and forward implied volatilities for different maturities and currencies, we can explore the term structure of volatility risk premia. For this exercise, we compute the volatility excess returns on an FVA over a month such that we work with non-overlapping monthly observations. Specifically, the excess return to an investor that holds a τ_1/τ FVA on a given currency pair between months t and $t+1$ is defined as

$$rx_{\ell,t+1} = \frac{FVOL^{\tau_2}_{t+1,\tau_1-1} - FVOL^{\tau_2}_{t,\tau_1}}{FVOL^{\tau_2}_{t,\tau_1-1}}, \tag{2}$$

where ℓ denotes a given maturity combination τ_1/τ (i.e., 1/3 months, 3/6 months, 6/12 months and 12/24 months). For example, we may consider the excess return from entering into a 3/6 months FVA today (i.e., you deliver the 3-month forward volatility in 3 month's time in exchange for the 3-month spot volatility observed in 3 month's time) and exiting in one month when it will be a 2/5 months FVA (i.e., you receive the 3-month forward volatility in 2 month's time against the 3-month spot volatility determined in 2 month's time).

By examining average excess returns for different maturities and currency pairs, we can readily establish whether there are significant volatility risk premia, and how such premia vary with

[3] We use money market rates and interest rate swap data from which we bootstrap zero-yield rates.

maturity. Table 1 documents the term structure of unconditional volatility risk premia across different maturities. Panel A aggregates excess returns using an equally-weighted scheme and reports an average excess returns of -2.90% (with a t-statistic of -3.07) for the 1/3 month maturity. Average excess returns, however, are small and statistically insignificant at longer horizons. The disappearance of the volatility risk premium at longer horizons is consistent with the findings of Dew-Becker, Giglio, Le, and Rodriguez (2016) for the equity market. Panel B employs GDP-weighted excess returns but results are largely comparable.

Table 1 about here

We now test if there exists a significant conditional volatility risk premium. Recent empirical evidence shows that this is so for the short-term end of the implied volatility curve in the FX market. Similar to the relationship between forward and future spot exchange rates, Della Corte, Sarno, and Tsiakas (2011) test the expectation hypothesis between forward and future spot implied volatilities using the analog of the Fama (1984) predictive regression. Using a cross-section of nine developed currency pairs, they find that the forward volatility premium is a biased predictor of the future spot implied volatility return.[4] We revisit and extend their analysis across different dimensions. In particular, we focus on a cross-section of 20 developed and emerging currency pairs between January 1996 and December 2015 and study the time-variation in volatility risk premia both at short and long horizons using different maturity combinations ℓ. We further enhance their testing framework by deriving a spot-forward implied volatility relationship for non-overlapping returns and define the forward volatility premium on date t for a given currency pair as $fvp_{\ell,t} = (FVOL^{\tau_2}_{t,\tau_1} - FVOL^{\tau_2}_{t,\tau_1-1})/FVOL^{\tau_2}_{t,\tau_1-1}$, where ℓ refers to a given maturity combination.[5] We provide more formal and detailed arguments

[4]For the equity market, Johnson (2017) shows that the shape of the VIX term structure is informative about time-varying variance risk premia, thus rejecting the expectations hypothesis.

[5]The justification for suspecting that conditioning on slope might capture time varying risk premia is simple. Suppose the spot volatility t is x_t, and the slope of volatility term structure (or forward premium) is s, then the profit from going long a one period volatility forward contract at time 0 and holding to maturity is $x_1 - x_0 - s$. The volatility risk premium $rx = m - s$ where m is the expected change in the spot volatility. In the absence of risk premia of course $s = m$. But in any model where there are time varying risk premia, then the risk premium must be negatively associated with slope unless the slope coefficient in a regression of the risk premium on the slope is $+1$ or more. This argument justifies the choice of the interest rate differential as a conditioning variable in the currency market, and is indeed the foundation of the carry trade more generally.

in the Internet Appendix B..

We build on this evidence and regress monthly volatility excess returns on the lagged forward volatility premia. We pool all currencies together and run the following panel regressions

$$rx_{i\ell,t+1} = \alpha + \gamma fvp_{i\ell,t} + fe_i + \varepsilon_{i\ell,t+1}, \tag{3}$$

where i denotes a given dollar-based currency pair and fe indicates the currency fixed effects. In the absence of risk premia, the current forward implied volatility is an unbiased predictor of the future implied volatility and the expected excess return is zero. We should then expect to see that α and γ are both equal to zero, and ε_{t+1} is serially uncorrelated.

Table 2 about here

We report the least-squares estimates of α and γ with t-statistics (reported in brackets) based on standard errors robust to heteroscedasticity, cross-sectional, and temporal dependence as in Driscoll and Kraay (1998) and Vogelsang (2012) in Panel A of Table 2. While the coefficient α is always statistically insignificant, the coefficient γ is always negative and statistically different from zero. The estimate of γ ranges between -0.65 (with a t-statistic of -4.73) for 1/3 month and -1.82 (with a t-statistic of -3.75) for 12/24 month. We further check whether our results are affected by convexity bias. Panel B of Table 2 shows that estimates of α and γ remain unchanged when implied volatilities are replaced with implied variances. Finally, we also run country-level predictive regressions but find consistent results across most of the currencies (see Table A1 in the Internet Appendix).

Taken together, this evidence strongly rejects the hypothesis that the forward implied volatility is an unbiased predictor of the future spot implied volatility. It shows there exists a time-varying volatility risk premium in the FX market at horizons up to 24 months. The strategy of selling implied forward volatility when it is at a premium to spot is generally profitable. This is similar to the carry trade strategy where an investor sells (buys) a currency at a forward premium (discount) in the forward market against the corresponding future spot exchange rate (e.g., Lustig, Roussanov, and Verdelhan 2011).

3. Cross-Section of Volatility Excess Returns

In this section, we study cross-sectional variation in volatility excess returns. The previous section shows that the volatility term structure is informative about future volatility excess returns. Motivated by this finding, we investigate whether information in the volatility term structure also predicts future FVA returns in the cross-section.

3.1. Implied Volatility Portfolios

Using forward volatility premia as a predictor is intuitively equivalent to extracting information from the slopes of the implied volatility term structures: selling (buying) an FVA with a positive (negative) forward volatility premium is tantamount to having a short (long) position on an FVA when the implied volatility curve is upward (downward) sloping. Guided by this intuition, we build portfolios of FVAs using the slopes of the implied volatility curves as the sorting variable.[6]

We measure the slope of the implied volatility curve for each currency on date t as

$$slope_t = \frac{SVOL_t^{24} - SVOL_t^{3}}{SVOL_t^{3}}. \tag{4}$$

At the end of period t, we allocate the FVAs to five baskets using the volatility slope of each currency pair observed on date t. We rank these portfolios from high to low slope such that Portfolio 1 contains the 20% of all FVAs with the highest slope and Portfolio 5 comprises the 20% of all FVAs with the lowest slope. We re-balance them monthly from January 1996 to December 2015 and compute the excess return for each basket as an equally weighted average of the volatility excess returns within that basket. This exercise is repeated for each maturity combination ℓ using a sample that includes up to 20 countries.

Similar to Lustig, Roussanov, and Verdelhan (2011), we also construct two additional portfolios: the level strategy, denoted LEV, which corresponds to a zero-cost strategy that equally

[6]Using the slope of the volatility term structure allows us to apply the same conditioning information across different maturities. The forward volatility premium is maturity specific and, hence, captures also information embedded in the curvature of the volatility curve. Results remain largely comparable when using forward volatility premia as conditioning variables.

invests in all implied volatility portfolios and the volatility carry strategy, denoted VCA, which is equivalent to a long-short strategy that buys Portfolio 5 and sells Portfolio 1. Table 3 presents the summary statistics for the five portfolios of FVAs. In brackets, we report t-stats based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection.

Table 3 about here

The average excess return increases monotonically from the first portfolio to the last portfolio for all maturity combinations. The average monthly excess return on Portfolio 1 (Portfolio 5) is about -4.66% (0.49%) in Panel A (1/3 months) and -0.40% (2.10%) in Panel D (12/24 months). There is no clear pattern for the standard deviation. Skewness is positive and higher for Portfolio 5 than Portfolio 1 in the shorter maturities but the ranking is reversed in the longer maturities. Also, there is evidence of positive return autocorrelation, especially for Portfolio 5.

We also report the summary statistics for the LEV and VCA portfolios. The average excess return of the LEV portfolio ranges from -2.39% (in Panel A) to 0.63% per month (in Panel D) but it is statistically significant only for 1/3 months. In contrast, the average excess return for the VCA strategy – long a portfolio of FVAs with the lowest volatility slope and short a portfolio of FVAs with the highest volatility slope – is always positive and highly statistically significant. The average excess return ranges between 5.15% (with a t-stat of 5.91) and 2.50% (with a t-stat of 5.67) per month for 1/3 months and 12/24 months, respectively. The corresponding annualized Sharpe ratios are also monotonically decreasing from 1.46 to 1.25.[7] The last row reports the frequency of portfolio switches ($freq$) computed as the ratio between the number of portfolio switches and the total number of returns at each date. It reveals substantial variation in the composition of the volatility portfolios.[8]

Overall, our descriptive statistics confirm that there exists substantial cross-sectional variation in excess returns of FVAs. Furthermore, we show that implied volatility slope has the

[7] The Internet Appendix Table A16 presents the summary statistics of portfolios sorted by forward volatility premia. We find qualitatively similar results.

[8] The Internet Appendix Table A2 presents the currency composition of these five slope-sorted portfolios.

ability to predict excess returns in the cross-section that are significant both statistically and economically.

Figure 2 about here

Figure 2 presents the one-year rolling Sharpe ratio for the VCA strategies (based on the slope-sorted portfolios) and their equally-weighted average. The strategies exhibit a clear counter-cyclical pattern producing higher risk-adjusted excess returns during the financial crisis and lower risk-adjusted excess returns otherwise. In particular, the Sharpe ratios are economically large during the financially troubled period of 1997-1999 which included the Asian financial crisis, the Russian sovereign default, and the collapse of the hedge fund LTCM. The Sharpe ratios of the VCA strategies are also high during the terrorist attacks on September 11, 2001, the wars in Afghanistan and Iraq, the recent global financial crisis that started with the collapse of Lehman Brothers in September 2008, and more recently during the European Sovereign crisis. Financial crises are generally characterized by a sudden increase in risk aversion and substantial exchange rate uncertainty which drive up the price of risk. Both factors are likely to be captured by the currency option implied volatilities (e.g., Marion 2010).

3.2. Common Variation in Volatility Excess Returns

A natural question to ask is whether volatility excess returns can be understood as compensation for risk, and if so, whether they respond to the same set of risk factors that price currency excess returns (e.g., Lustig, Roussanov, and Verdelhan 2011; Menkhoff, Sarno, Schmeling, and Schrimpf 2012). In this section, we study the (slope-sorted) implied volatility portfolios in a cross-sectional asset pricing framework and show empirically that they can be thought of as a reward for time-varying global risk. This is also where our analysis goes beyond Della Corte, Sarno, and Tsiakas (2011).

We start by examining whether average excess returns stemming from the cross-sectional predictability of implied volatility slopes reflect risk premia associated with exposure to a small set of risk factors. Similar to Lustig, Roussanov, and Verdelhan (2011), we employ principal component analysis on our implied volatility portfolios and find that up to 90% of the common variation in the excess returns of these portfolios can be explained by two factors.

TABLE 4 ABOUT HERE

Table 4 presents, for each maturity combination, the loadings of our volatility portfolios on the first two principal components as well as the fraction of the total variance (in bold) of portfolio returns associated with each principal component. Across the four maturities, the first two principal components explain between 88% and 90% of the common variation in portfolio returns.[9] The first principal component can be understood as a *level* factor as all portfolio load with similar coefficients on it, ranging between 0.52 on Portfolio 1 and 0.42 on Portfolio 5. The second principal component can be interpreted as a *slope* factor as loadings increase monotonically across portfolios, from −0.82 on Portfolio 1 to 0.49 on Portfolio 5.

Two candidate risk factors emerge from our principal component analysis. The first one can be approximated as the average excess return across all implied volatility portfolios (LEV) and can be seen as the average portfolio return of a US investor who buys all FVAs in the currency options market and represents the premium she is willing to pay to hedge her US volatility risk exposure. The second one can be approximated by the return difference between Portfolio 5 and Portfolio 1 (VCA) and can be interpreted as a zero-cost strategy that buys FVAs with the lowest implied volatility slopes and sells FVAs with the highest implied volatility slopes. The correlation of the first principal component with LEV is essentially one for all maturity combinations. The correlation of the second principal component with VCA is about 0.95 on average.[10]

[9] An alternative strategy would be to perform the principal component analysis on all the maturities simultaneously. The conclusions are very similar; the first two components correspond closely to the average of the maturity-specific component and capture 82% of the common variation in portfolio returns. Details of the analysis are contained in Table A3 of the Internet Appendix.

[10] The correlation of the LEV factor with the dollar factor of Lustig, Roussanov, and Verdelhan (2011) revolves around −0.45 whereas the correlation of the VCA factor with their carry factor is 0.01 on average and ranges from 0.13 for 1/3 months and −0.05 for 12/24 months.

3.3. Portfolio-level Asset Pricing Tests

We now turn to a more formal investigation of our portfolio excess returns using standard asset pricing methods as in Lustig, Roussanov, and Verdelhan (2011). In the absence of arbitrage opportunities, risk-adjusted excess returns have a price of zero and satisfy the Euler equation $E[m_t rx_t^j] = 0$, where rx_t^j is the excess returns of portfolio j at time t and $m_t = 1 - b' (f_t - \mu)$ is a stochastic discount factor (SDF) linear in the pricing factors f_t. The factor loadings are denoted by b and the factor means by μ. This specification also implies a beta pricing model $E[rx^j] = \lambda' \beta^j$ that depends on the factor risk prices λ and the risk quantities β^j, i.e., the regression coefficients of each portfolio's excess return rx_t^j on the risk factors f_t. The factor risk prices can be obtained as $\lambda = \Sigma_f b$, where $\Sigma_f = E\left[(f_t - \mu)(f_t - \mu)'\right]$ is the covariance matrix of the risk factors.

The factor loadings b are estimated via the Generalized Method of Moments (GMM) of Hansen (1982). To implement GMM, we use the pricing errors as a set of moments and a prespecified weighting matrix. Since the objective is to test whether the model can explain the cross-section of expected currency excess returns, we only rely on unconditional moments and do not employ instruments other than a constant and a vector of ones. The first-stage estimation (GMM_1) employs an identity weighting matrix. The second-stage estimation (GMM_2) uses an optimal weighting matrix based on a heteroskedasticity and autocorrelation consistent estimate of the long-run covariance matrix of the moment conditions. The model's performance is then evaluated using the cross-sectional R^2 and the HJ distance measure of Hansen and Jagannathan (1997), which quantifies the mean-squared distance between the SDF of a proposed model and the set of admissible SDFs.[11]

Motivated by the principal component analysis presented above, we study the risk exposure of our implied volatility returns using a two-factor SDF, i.e., with LEV and VCA as factors, and present asset pricing tests on the cross-sections of volatility portfolios as test assets in Table 5. We report estimates of the factor loading coefficients b and market prices of risk λ with the t-statistic in square brackets, the cross-sectional R^2, and the p-value of the HJ

[11] To test whether the HJ distance is statistically significant, we simulate p-values using a weighted sum of χ_1^2-distributed random variables (see Jagannathan and Wang 1996).

distance in parenthesis for all maturity combinations.

Table 5 about here

We find overall a positive and statistically significant price of VCA risk. In Panel A (the short term end of the implied volatility curve), the estimate of λ_{VCA} is about 4.75% per month (with a t-stat of 4.86) for the first-stage GMM. This implies that an asset with a beta of one earns a risk premium of 475 basis points per month. This estimate remains very similar in terms of magnitude and statistical significance when moving to the second-stage GMM or the FMB method. Since VCA is a tradable risk factor, its factor price of risk must equal its average excess return as the Euler equation applied to the risk factor itself would produce a coefficient β equal to one. This no-arbitrage condition is indeed satisfied in our exercise as the average monthly excess return on the VCA factor is 5.15%, slightly higher than the estimate of λ_{VCA}. A positive estimate of the VCA risk price indicates higher (lower) risk premia for implied volatility portfolios sorted on downward (upward) sloping implied volatility curves. We also uncover a strong cross-sectional fit in terms of R^2 and are unable to reject the null hypothesis that pricing errors are zero as measured by the HJ distance. Results for the other maturity combinations (see Panels B to D of Table 5) remain qualitatively very similar.

Table 5 also reports the price of LEV risk. Panel A displays a λ_{LEV} of -2.37% per month which compares well with the average return of -2.39% per month of the LEV portfolio. This factor is also statistically significant (with a t-stat of -2.20). b_{VCA} is positive and statistically significant (0.03 with a t-stat of 2.67) while b_{LEV} statistically insignificant (-0.01 with a t-stat of -1.31). We conclude that the LEV factor does not help explain variation in volatility excess returns given the presence of the VCA factor. Our finding remains qualitatively identical in Panels B to D of Table 5, thus confirming that the cross-section of the implied volatility portfolios can be priced just as well without the LEV factor as with it. While the level factor does not help explain the cross-sectional variation in expected returns, it is important for the level of average returns as it works as a constant that allows for common mispricing in the cross-sectional regression.

3.4. Country-level Asset Pricing Tests

Sorting asset returns into portfolios is widely used as it improves the estimates of the time-series slope coefficients. As shown by Lewellen, Nagel, and Shanken (2010), this practice produces a strong factor structure in test asset returns and can lead to misleading results, especially with a small cross-section of asset returns. We deal with this concern by also running asset pricing tests with country-level volatility excess returns as test assets, and LEV and VCA as risk factors. Table 6 reports Fama-MacBeth estimates of the market prices of risk λ (with t-statistics in brackets) and cross-sectional R^2 for all maturity combinations.[12]

Table 6 about here

Similar to Lustig, Roussanov, and Verdelhan (2011), we construct country-level excess returns between times t and $t + 1$ by going long (short) FVAs with implied volatility slopes lower (higher) than their median value at time t such that the strategy is dollar-neutral. Consistent with our previous findings, we uncover a positive and statistically significant factor price of volatility carry risk. The estimates of λ_{VCA} are economically large and statistically significant across all four maturity pairs (even after bootstrapping the standard errors) and the cross-sectional R^2 ranges between 48.9% to 76.0%. The estimates of λ_{LEV}, in contrast, are small and statistically insignificant, especially after bootstrapping the standard errors. In sum, we find that the volatility carry factor is the only source of priced risk in the cross-section of implied volatility returns for all four maturity combinations.

3.5. Time-series Exposure to VCA

If VCA is the only source of risk that matters in the cross-section, the volatility excess return should increase with its exposure to the VCA factor as measured by the factor betas. We estimate the exposure of each portfolio's excess return to the LEV and VCA factors

[12] As robustness, we also compute bootstrapped standard errors based on 10,000 replications but conclusions remain unchanged. We use the stationary bootstrap of Politis and Romano (1994) which resamples with replacement blocks of a random length of excess returns and pricing factors realizations from the original sample without imposing the model's restrictions. This procedure preserves both contemporaneous cross-correlations and serial correlations for excess returns and pricing factors.

by running the contemporaneous time-series regressions for each maturity combination. We present the least-squares estimates of constant α as well as the slope coefficients β_{LEV} and β_{VCA} in Table 7. In Panel A, we find that the first and the last portfolios have an estimates α of 0.81% per month, which is statistically significant at the 5% level. The estimates of α for the other portfolios are smaller and negative, and the null hypothesis that the alphas are jointly zero cannot be rejected at the 5% or 10% significance level since the p-value of the χ^2_α statistic is 0.21. The next column reports the beta estimates of the LEV factor which are all statistically significant and indistinguishable from one. This is expected as LEV is essentially the first principal component and does not explain any of the variations in average excess returns across portfolios.

Table 7 about here

The third column presents the beta estimates for the VCA factor which increase monotonically from −0.58 (with a t-stat of −13.37) for Portfolio 1 to 0.42 (with a t-stat of −9.76) for Portfolio 5. Moreover, the goodness of fit is very high since the R^2 is in the range between 86.0% and 93.7%. These results are broadly similar for the other maturity combinations.

3.6. Variation in VCA Loadings

We now study whether exposure to volatility carry also varies with the state variables driving the implied volatility slope by running panel regressions based on the following specification:

$$\begin{aligned} rx_{i\ell,t} \; &= \beta LEV_{i\ell,t} + \gamma VCA_{i\ell,t} + \phi X_{i,t-1} + \\ & \delta LEV_{i\ell,t} \times X_{i,t-1} + \lambda VCA_{i\ell,t} \times X_{i,t-1} + \alpha + fe + \varepsilon_{i\ell,t}, \end{aligned} \tag{5}$$

where $rx_{i\ell,t}$ is the monthly volatility excess return for currency i and maturity combination ℓ, $LEV_{i\ell}$ and $VCA_{i\ell}$ are the volatility level and volatility carry factors constructed as in Table 3 while excluding currency i, and X_i is the deviation of the implied volatility slope for currency i from the cross-country median value. In the spirit of Verdelhan (2018), we interact both $LEV_{i\ell,t}$ and $VCA_{i\ell,t}$ with $X_{i,t-1}$ to capture, respectively, time variation in the volatility level and volatility carry exposure with respect to unknown state variables. We absorb time-invariant unobserved characteristics with currency and maturity fixed effects, and unobserved variables that evolve over time but are constant across countries with time fixed effects. We refer to these fixed effects as fe.

Table 8 about here

Table 8 tabulates the parameter estimates with t-statistics reported in brackets. Specifications (1) and (4) have no fixed effects and are equivalent to pooled regressions. The estimates of β, i.e., the loadings on the volatility level factor, are always positive and statistically significant with values ranging between 0.96 and 0.69. These loadings capture the existence of a strong principal component that characterizes country-level volatility excess returns, meaning that an increase in LEV positively affects all countries. There exists some variation in the estimates of β, especially when time and currency fixed effects are turned on, suggesting that country-level volatility excess returns increase in different proportions in response to positive changes in the volatility level factor. The estimates of δ, the loadings on $LEV_{i\ell,t} \times X_{i,t-1}$, are always statistically insignificant and show that exposure to volatility level does not vary with the state variables underlying the implied volatility slopes.

Exposure to volatility carry, in contrast, varies with the implied volatility slopes since the estimates of λ, the loadings on $VCA_{i\ell,t} \times X_{i,t-1}$, are negative and statistically significant in all specifications, with values ranging between -0.99 and -1.03. This evidence suggests that exposure to volatility carry is lower when the implied volatility curve is steeper and vice versa. Putting it differently, for a given country, times of higher (lower) implied volatility slopes are also times of lower (higher) comovement with the volatility carry factor. We also report a Wald test for the null hypothesis that γ and λ on the unconditional and conditional volatility carry factors are jointly zero. The null hypothesis is always rejected at the 1% confidence level.

Table 9 about here

We also check whether our results are driven by a few extreme cases. We run country-level regressions with both unconditional and conditional exposure to volatility carry, while controlling for the volatility level, and report the results in Table 9. As in the previous table, the unconditional exposure to volatility level is always positive and statistically significant but the magnitude of β varies across countries: the lowest value is for Japan and the highest value is for Mexico. The unconditional exposure to volatility carry is captured by the coefficient γ, whose estimate is positive in 14 of the 20 countries, and significantly so (at the 10% level) in 8 of them. There is one country, Brazil, where the coefficient is significantly negative. The conditional exposure to volatility carry is reflected in the coefficient λ, whose estimate is negative and statistically significant (at the 10% level) in 11 out of 20 countries. The overall sensitivity of country-level excess returns to volatility carry depends on both unconditional and conditional exposure. We report via a Wald test that γ and λ are jointly statistically significant at the 10% confidence level in 17 out of 20 countries (the exceptions are Japan, Turkey, and South Africa). Overall, we confirm the findings reported in the previous exercise.

3.7. *Alternative Risk Factors*

We also check whether the volatility carry portfolios are explained by alternative risk factors such as currency factors, global equity factors, and US volatility factors. The set of currency factors includes the dollar (DOL), carry (CAR), global imbalance (IMB), FX global volatility (VOL) and FX global liquidity (LIQ) factor as in Lustig, Roussanov, and Verdelhan (2011), Menkhoff, Sarno, Schmeling, and Schrimpf (2012), and Della Corte, Riddiough, and Sarno (2016). We briefly describe how these traded factors are constructed in the Internet Appendix C.. As global equity risk factors, we use the Fama and French (2016) market (MKT), size (SMB), value (HML), profitability (RMW) and investment (CMA) factor taken from Kenneth French's website.[13] Finally, we use the VIX futures returns ranging from 1-month (R_1) to 6-month (R_6) to proxy the US equity volatility factors (as in Johnson 2017), which we collect from Travis Johnson's website. These data are only available between April 2004 and December 2015.

[13] We use ex-US equity factors as our test assets are dollar-neutral. We also use cum-US equity factors but results remain qualitatively identical.

Table 10 about here

We regress the excess return from our 20 implied volatility portfolios on the volatility level factor and on each group of alternative risk factors. We present the least-squares estimates of the alphas for each portfolio in Table 10. The alphas are statistically significant in most of the cases and the null hypothesis that the intercepts are jointly equal to zero is rejected at the 1% significance level. On the basis of this exercise, we conclude that the existing risk factors are unable to fully explain the variation in the excess returns of our implied volatility portfolios.[14]

3.8. Spot Volatility Risk Premia

So far, we have explored the behavior of volatility risk premia through the lens of FVAs. A natural question to ask, then, is how our analysis is related to volatility risk premia implied from other volatility derivatives such as the volatility swaps. Recall that while an FVA is a forward contract on future implied volatility, a volatility swap is a forward contract on future realized volatility. One may expect that such a relationship to be very close as buying a 3/6 months FVA is very similar to being long a 6-month volatility swap and short a 3-month volatility swap. There is a small convexity correction (since we are working with volatilities rather than variances) and a difference in the timing of the cash flows. Alternatively, buying a 1-month volatility swap can be seen as the limit case of a 0/1 month FVA strategy.

Motivated by these considerations, we also construct five baskets of 1-month volatility swaps sorted by the implied volatility slopes. Hence, we compute a strategy that equally invests in all portfolios as well as a long-short strategy that sells (buys) the first (last) portfolio. We refer to them as the 0/1 month LEV and VCA factors, respectively (see Table A9 in the Internet Appendix for the summary statistics of these portfolios). Figure 3 plots the average excess returns (Panel A) and the corresponding Sharpe ratios (Panel B) for both VCA and LEV across all maturity combinations.

[14]We report the full set of parameter estimates in the Internet Appendix Tables A11-A14. We show that the volatility level factor is always highly statistically significant whereas the explanatory power of all other factors is statistically insignificant with very few exceptions. Moreover, the R^2 (based on all factors) and the R^2_{LEV} (based on the level factor only) are almost identical.

Figure 3 about here

The behavior of LEV is strikingly similar to the behavior of the equity market risk premium shown in Dew-Becker, Giglio, Le, and Rodriguez (2016). There is a large and highly significant negative spot volatility risk premium that declines rapidly in magnitude with the horizon and is insignificantly different from zero at longer maturities. As they argue, this suggests that the spot volatility risk premium is related to transient volatility shocks. The behavior of VCA is quite different. While the magnitude declines with maturity, its volatility declines as well and the resulting Sharpe ratio is virtually independent of horizon. This suggests that the VCA risk premium is related to permanent volatility shocks.

4. Robustness Checks

In this section, we explore the robustness of our results with respect to data quality, methods of computing implied volatility, transaction costs, and selection of countries in our sample.

4.1. Traded vs. Quoted Implied Volatility

Our empirical analysis is based on implied volatility quotes from OTC currency options. We employ these data to synthesize both spot and forward model-free implied volatility. One may have some legitimate concerns about the indicative nature of these quotes, which we attempt to mitigate by studying the relationship between traded and quoted implied volatilities.

We collect transaction-level currency option contracts from the Depository Trust & Clearing Corporation (DTCC) between March 2013 and April 2019. We consider OTC options with a maturity between one month and two years for a sample of 20 developed and emerging currencies, thus matching our sample of quoted implied volatility. While currency options are generally quoted in terms of implied volatility and fixed deltas, transaction-level options are reported in terms of Garman and Kohlhagen (1983) option premia and strike prices. We filter out approximately 5% of the options, i.e., options with missing or zero premia, strikes, and notional amounts, options with negative maturities due to time-stamp errors, and in few cases options with unreasonable strike prices. This cleaning process leaves us with more than a million observations.

We then proceed to extract the implied volatility from traded options as follows. We first match option premia with spot exchange rates to the nearest second (the rates are from Thomson Reuters Tick History). We then employ a numerical procedure that uses the closed-form implied volatility solution of Brenner and Subrahmanyam (1988) as a starting value.[15] Regarding the implied volatility quotes, we first extend our sample up until April 2019 and then pair quoted with traded implied volatilities by currency, strikes, and maturities. We handle this exercise, in line with procedure presented in Section 2., by linearly interpolating implied variances at fixed deltas across maturities, converting then deltas into strikes using closed-form solutions, and finally using a cubic spline to interpolate between strikes with equal maturities.

Table 11 about here

Before running any regressions, we compare traded and quoted implied volatilities in terms of descriptive statistics. Panel A of Table 11 presents the descriptive statistics of traded implied volatilities for all countries in our sample. We work with daily observations, which we obtain by weighting intraday transactions by their corresponding notional value. Panel B reports the descriptive statistics of quoted implied volatilities aggregate daily via a simple average. We find that traded and quoted amounts are largely comparable both in terms of means and standard deviations. For example, the traded (quoted) implied volatility on USD/EUR has a mean of 8.64 (8.47) and a standard deviation of 1.85 (1.78). We also provide a visual inspection in Panel A of Figure 4, which plots the cross-country average implied volatility for traded and quotes options and shows that both quantities are very close to each other at the aggregate level. Panel B of Figure 4 plots the total volume of our traded options in billions of US dollars and the average is higher than $20 billion per day. Using the latest

[15] For this exercise, we rely on daily domestic and foreign interest rates as intraday money market rates are only available up to a maturity of three months. We also experiment with alternative closed-form solutions as starting values for our numerical exercise but results remain virtually unchanged.

BIS Triennial Survey as a reference (we do not observe the breakdown by currency pairs and maturity buckets), our transaction sample more than 11% of the global trading activity in the FX options market as of April 2019.

Figure 4 about here

We now run panel regressions with the following specification

$$IVOL^{T}_{ik\tau,t} = \alpha + \beta\, IVOL_{ik\tau,t} + \gamma' X_{i,t} + fe + \varepsilon_{ik\tau,t}, \tag{6}$$

where $IVOL^{T}_{ik\tau,t}$ is the traded implied volatility for currency i, strike k, and maturity τ on day t (we drop any intraday subscript to ease the notation), $IVOL_{ik\tau,t}$ is the corresponding quoted implied volatility, and $X_{i,t}$ denotes a set of currency-specific variables such option notional value, interest rate differential, foreign exchange liquidity, and spot exchange rate returns. Finally, we absorb time-invariant unobserved characteristics using currency and maturity fixed effects, and unobserved variables that evolve over time but are constant across countries with time fixed effects. We also add hour fixed effects to capture intraday seasonal patters. We refer to these fixed effects as fe. If the difference between trades and quotes is neither systematic nor driven by any omitted variable, then the estimate of the slope coefficient β should not be statistically different from one.

Table 12 about here

We present the estimates in Table 12 with different combinations of fixed effects. Since we are working with a large dimensional panel, we cluster standard errors in the currency and time dimension and then report the corresponding t-statistics in brackets. In all specifications, the estimate of β is not statistically different from one, the R^2 is well above 60%, and the $RMSE$ is about 2.5%. The figures suggest that there is no systematic difference between trades and quotes. The fitting, however, is not perfect and a possible source of noise could arise from the bid-ask spreads as well as the lack of intraday synchronization between quoted and traded volatilities. We do not directly observe the trading direction and are unable to identify bid and ask traded prices.

4.2. Traded vs. Synthetic Forward Volatility

In our analysis, we replicate forward implied volatilities by linearly interpolating spot implied volatilities. This process may introduce a measurement error that could potentially affect our results. To tackle this problem, we have manually collected the fixed leg of tradeable FVAs from the archive of a London based hedge fund. The sample ranges between October 2009 and January 2014, covers a cross-section of 9 developed and 6 emerging countries and contains bid and ask prices from different major dealer banks operating in London. We cannot use it to implement a trading strategy since we do not have regularly spaced prices on individual FVAs. We can however verify our method of synthesizing FVA prices by regressing the synthetic prices on actual prices:

$$FVOL^T_{ij\ell,t} = \alpha + \beta\, FVOL_{i\ell,t} + \gamma\, Spread_{ij\ell,t} + fe + \varepsilon_{ij\ell,t}, \tag{7}$$

where $FVOL^T_{ij\ell,t}$ is the tradeable forward volatility on day t for currency i, dealer j, and maturity combination ℓ (i.e., 1/2 months, 3/6 months, 6/12 months, and 12/24 months), $FVOL_{i\ell,t}$ is the synthetic forward volatility for currency i and maturity ℓ, $Spread_{ij\ell,t}$ is the bid-ask spread on tradeable forward volatility, and fe denotes dealer, currency, maturity, and time (monthly) fixed effects. We run the above specification separately using bid, ask, and mid prices for the dependent variable, while omitting any extra subscript to ease the notation.[16]

Table 13 about here

The results are set out in Table 13 with t-statistics reported in brackets. In Panel A, we employ mid prices for $FVOL^T_{ij\ell,t}$ and find that the estimate of β is generally not statistically different from one. Specification (3), for example, absorbs both currency and time fixed effects and reports an estimate for β of 0.99 (not statistically different from one), an R^2 larger than

[16] In this database, we observe 1/2 months as opposed to 1/3 months forward implied volatilities. We thus work with available maturities for both quoted and tradable volatilities in this exercise.

97%, and a small RMSE of about 0.37%. Panel B and Panel C repeat the exercise using bid and ask prices for $FVOL^{T}_{ij\ell,t}$, respectively, and show qualitatively similar results. Overall, this table suggests that our synthetic forward implied volatility is close to the fixed leg of a tradeable FVA.

4.3. Impact of Transaction Costs

We examine the effect of transaction costs on the profitability of the volatility carry strategies. Following Menkhoff, Sarno, Schmeling, and Schrimpf (2012), we model the impact of transaction costs by assuming that the trader bears half the cost of the bid-ask spread each time a position is opened or closed. Bid-ask spreads for OTC volatility derivatives such as FVAs are notoriously difficult to obtain, so we need to estimate them. The procedure we use to estimate the spread is as follows: we use the bid-ask spread quoted by Bloomberg for a delta neutral straddle in the same currency and for the same maturity as the FVA. The straddle spread is likely to reflect the cross-sectional and temporal variation in the spread on the corresponding FVA. But FVAs are likely to have higher spreads than straddles since they are less liquid. To correct for this bias we make use of data we have obtained from one of the largest dealer banks in the currency options market on the actual bid-ask spreads in their FVA quotes, which we have by maturity and currency over a short period of time. We use these actual FVA transaction costs to scale up the spreads on straddles.[17]

Following this procedure, the average bid-ask spread in our sample is about 100 *bps* for the shortest maturity and 80 *bps* for the longest maturity. While transaction costs significantly reduce the profitability of the volatility carry strategies, they remain significant both economically and statistically. The average Sharpe ratio of 1.32 recorded in Table 3 drops to 0.83 when charging the largest quoted spread. We report the full results in Table A4 in the Internet Appendix.

FIGURE 5 ABOUT HERE

[17] We thank Philippos Kassimatis from Maven Global for helpful discussions on how to proxy for time-varying bid-ask spreads on FVA contracts.

We can check the plausibility of the estimated VCA spreads against actual VCA bid and ask quotes that we obtained from the hedge fund, and which we introduced in section 4.2. Figure 5 summarizes the average bid-ask spreads on those contracts for different maturities and currencies. The average spread is 61 *bps*. It is lower for developed countries (58 *bps*) than for emerging currencies (103 *bps*). There is little variation by maturity, with the average spread ranging from 58 *bps* to 62 *bps* moving from the shortest to the longest maturity. These spreads fall by half when we take just the best bid and ask rather than the average across all dealers.[18] This evidence suggests that, if anything, we may have over-estimated the impact of transaction costs on the profitability of the VCA strategy.

4.4. Additional Exercises

We experiment with alternative methods to synthesize spot and forward implied volatility (see the Internet Appendix D.). Some countries in our sample may be affected by capital controls and their currency options might be rather illiquid. To avoid any danger of this biasing our results, we run the same tests on a subset of 10 developed currencies (see the Internet Appendix E.). In both cases the results are substantially unaltered.

5. Conclusions

We identify a common risk factor in the currency volatility returns by sorting currencies on the slope of their implied volatility term structure. A zero-cost portfolio strategy that buys forward volatility agreements in currencies with the lowest slopes (or forward volatility premia) and sells those with the highest slopes produces significant excess returns. A risk factor – volatility carry strategy – fully explains the cross-sectional variation of slope-sorted volatility excess returns. The lower the slope of the implied volatility curve, the more the forward volatility agreement return is exposed to this volatility carry premium. Unlike the unconditional volatility risk premium, which vanishes beyond 2-3 months, the volatility carry premium is manifest at maturities up to 24 months. The risk factors suggested in the literature – carry, global imbalance, global volatility, and liquidity – cannot explain the cross-sectional

[18] We tabulate these estimates (with t-statistics in brackets) in Table A15 in the Internet Appendix.

variation of the forward volatility agreement returns.

The existence of a substantial unconditional risk premium in the short term volatility market, and its absence from the forward volatility market (beyond 2-3 months) presents a challenge to standard asset pricing models. For while it is easy to see why agents might demand a premium for bearing volatility risk in general, it is harder to see why it is only volatility risk in the short term that is priced. Dew-Becker, Giglio, Le, and Rodriguez (2016) observe a similar phenomenon in the equity index market, and conclude that jump models are best able to explain these results; in effect, agents in the economy price sudden short-term bursts of volatility, but they do not price permanent shifts to the level of volatility. In the context of the FX market, a jump in the exchange rate between two economies is a jump in the ratio of their Stochastic Discount Factors (SDFs). The fact that the short-term risk premium is negative suggests that the jumps in the SDFs must be predominantly upward jumps.

Our finding that longer term volatility shocks are priced conditionally also has implications for the modelling of foreign exchange rates. The attenuation of the size of the risk premium with horizon, and the stability of the Sharpe ratio are both consistent with a story of volatility shocks being fairly persistent, but dying out over a period of multiple months. The linkage we find between the shape of the term structure of implied volatility and the risk premium may not be hard to explain. The term structure of implied volatility reflects both the term structure of expected volatility and the volatility risk premium. If the two are uncorrelated, then the slope of the term structure will be positively correlated with the risk premium, as indeed we find.

Our results do raise questions about why the sign of the long-term volatility risk premium varies both in the cross-section and in the time series, and points to the need for further research in this area. The fact that the forward volatility risk premium is similar when the two currencies are reversed implies that the covariance between the volatility swap rate and the two SDFs is similar. This in turn means that the conditional covariance between the volatility swap rate and the exchange rate (the ratio of the two SDFs) is small. This is consistent with the finding that traditional FX risks do not explain volatility risk premia.

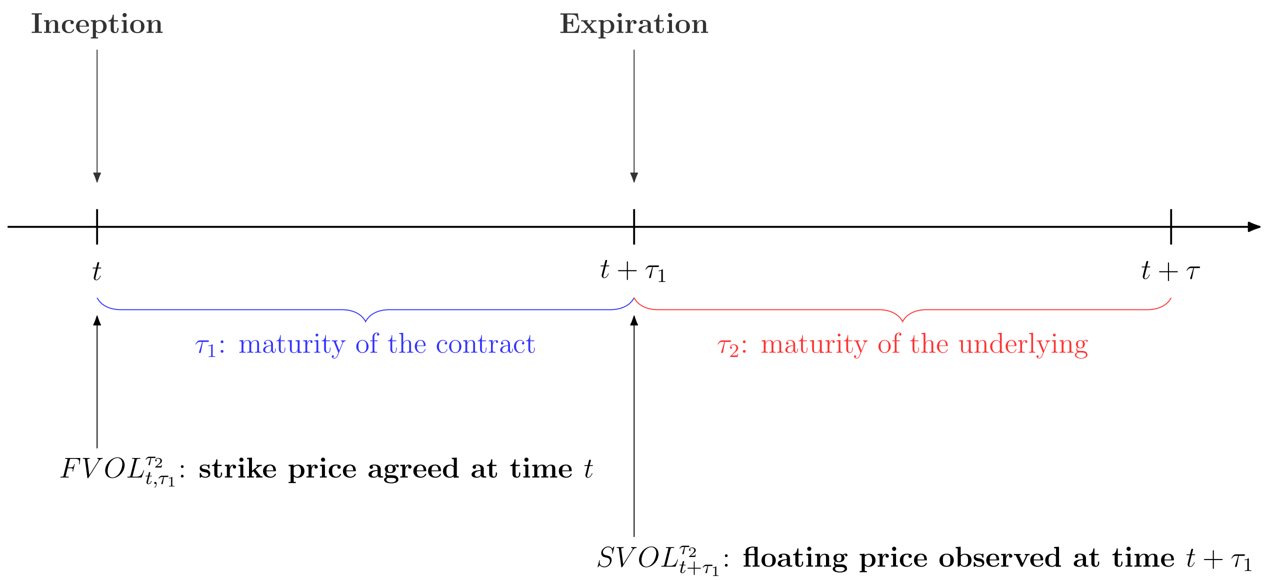

Figure 1. Forward Volatility Agreement

This figure describes a forward volatility agreement, i.e., a forward contract that exchanges the τ_2 period spot implied volatility ($SVOL_{t+\tau_1}^{\tau_2}$) observed at time $t + \tau_1$ against the forward implied volatility ($FVOL_{t,\tau_1}^{\tau_2}$) determined today but defined over the same future time interval. The buyer enters this contract at time t and receives on the maturity date $t + \tau_1$ a payoff equals to ($SVOL_{t+\tau_1}^{\tau_2} - FVOL_{t,\tau_1}^{\tau_2}$) for each dollar of notional amount. If $\tau_1 = \tau_2 = 3$ months, for example, $FVOL_{t,\tau_1}^{\tau_2}$ denotes the 3-month forward implied volatility in 3 month's time at time t and $SVOL_{t+\tau_1}^{\tau_2}$ is the 3-month spot volatility observed in 3 month's time.

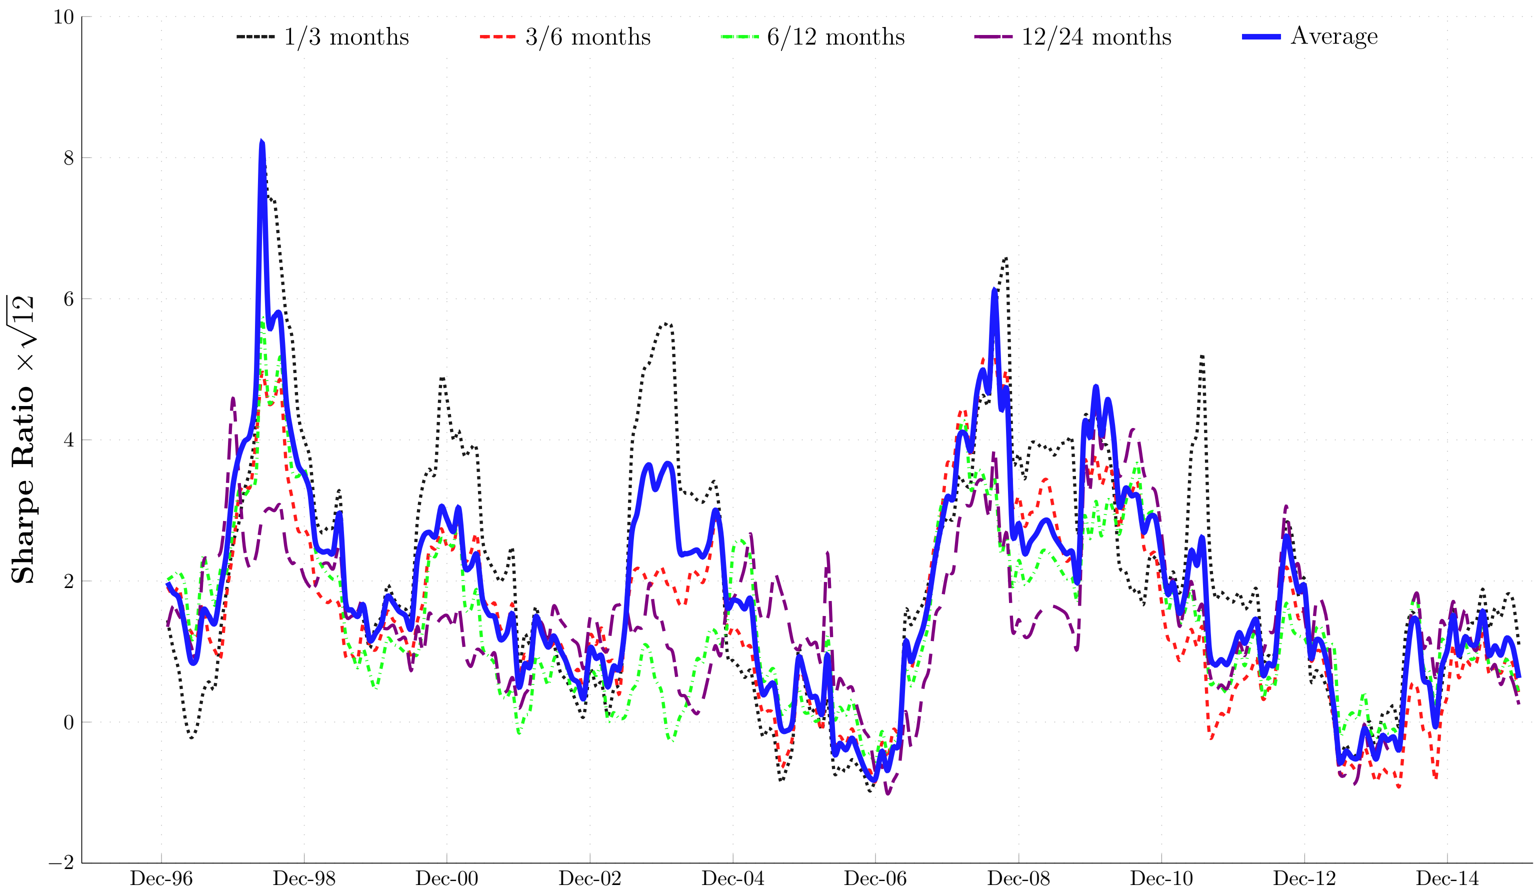

Figure 2. The Sharpe Ratios of Volatility Carry Strategies

This figure displays the annualized 1-year rolling Sharpe ratios of the volatility carry (VCA) strategies described in Table 3. Each strategy is constructed as a long-short strategy that buys a basket of forward volatility agreements with the lowest implied volatility slopes and sells a basket of forward volatility agreements with the highest implied volatility slopes. The implied volatilities are model-free as in Britten-Jones and Neuberger (2000) and Jiang and Tian (2005). The sorting variable for all the maturities is the ratio between the 24-month and 3-month implied volatility. *Average* denotes the rolling Sharpe ratio of an equally-weighted basket of volatility carry strategies. The volatility carry strategies are rebalanced monthly. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Figure A1 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

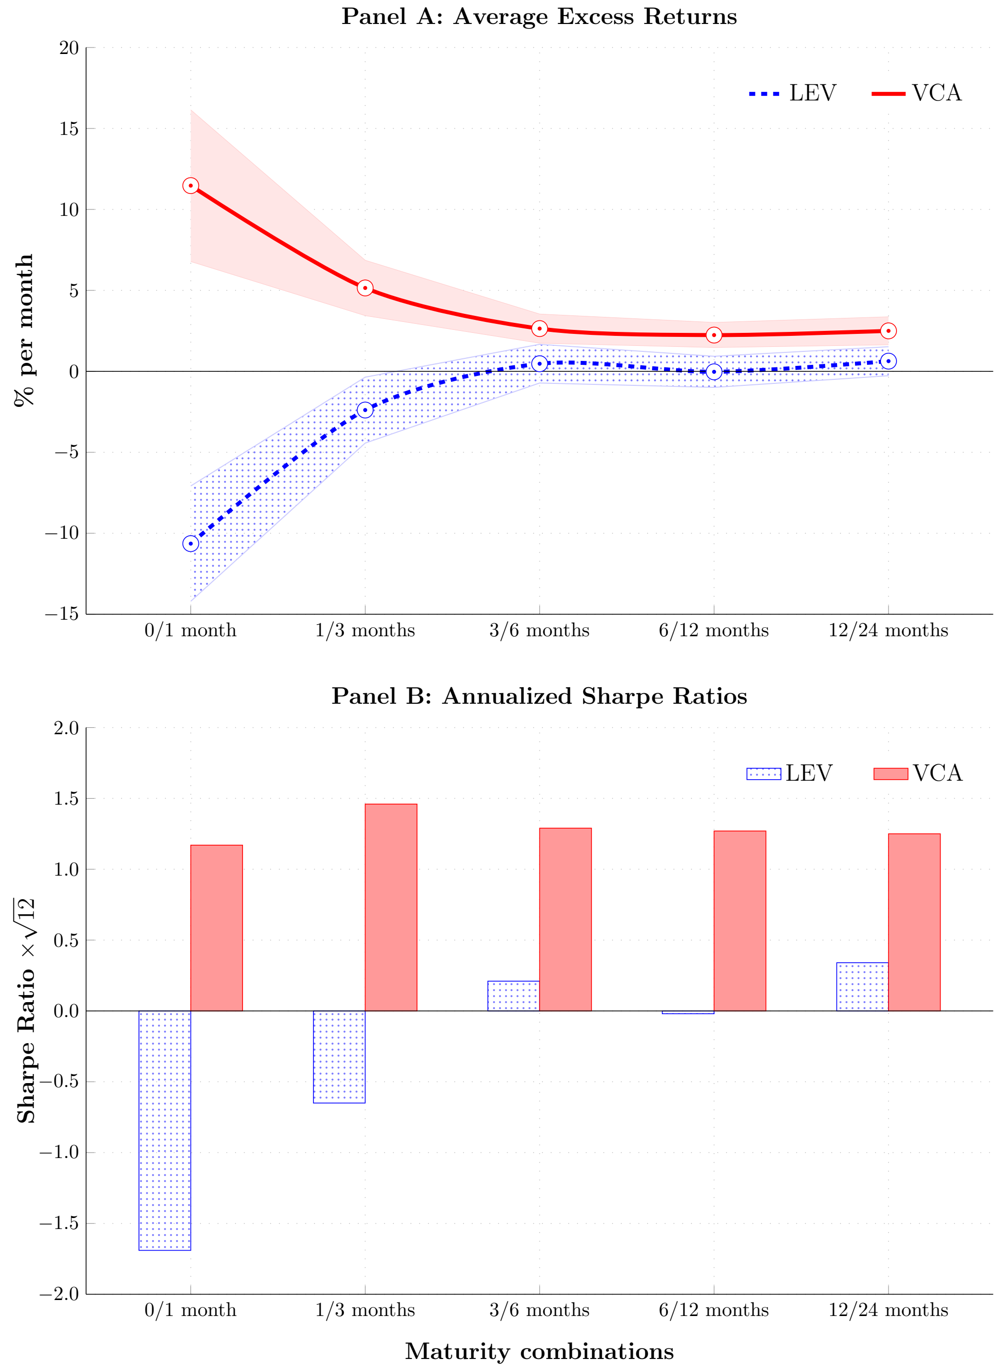

Figure 3. Performance of Volatility Carry Strategies

This figure reports average excess returns (Panel A) and annualized Sharpe Ratios (Panel B) of LEV and VCA, respectively. These are strategies based on slope-sorted portfolios of (i) volatility swaps for 0/1 month, and (ii) forward volatility agreements from 1/3 months to 12/24 months. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Figure A2 in the Internet Appendix displays results for a cross-section of 10 developed countries.

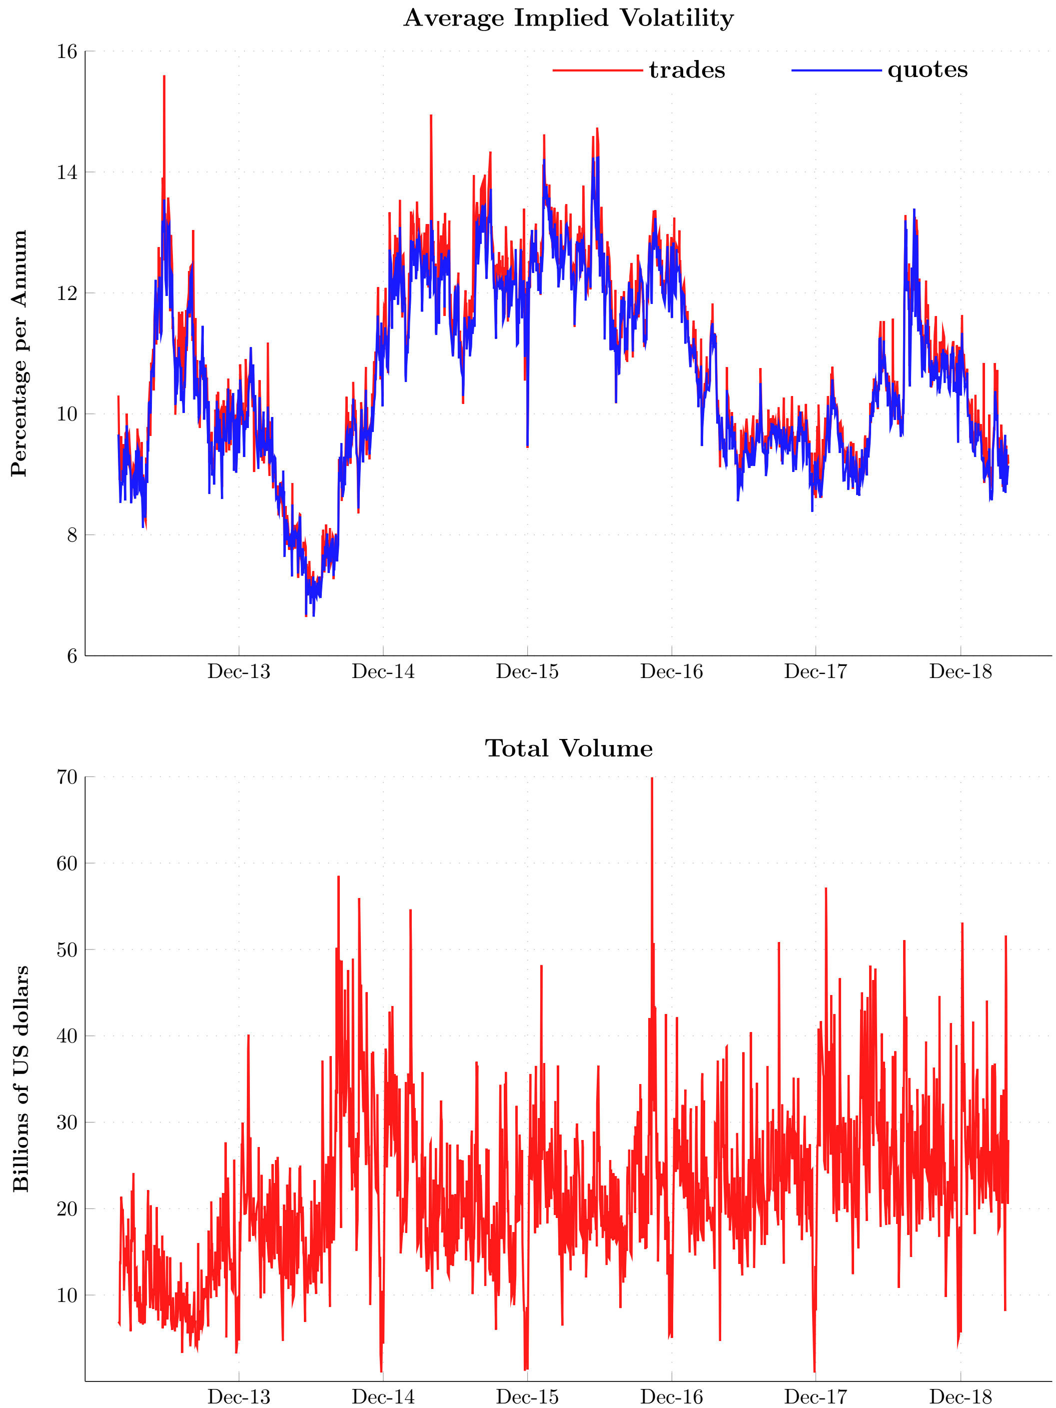

Figure 4. Traded vs. Quoted Implied Volatility

Panel A displays the implied volatility, averaged across countries and maturities, of over-the-counter traded and quoted currency options. Panel B presents the aggregate volume of traded options in billions of US dollars. The sample runs from March 2013 and April 2019 and includes a cross-section of 20 developed and emerging market countries. Traded options are from the Depository Trust & Clearing Corporation (DTCC) whereas quoted options are provided by Bloomberg and JP Morgan.

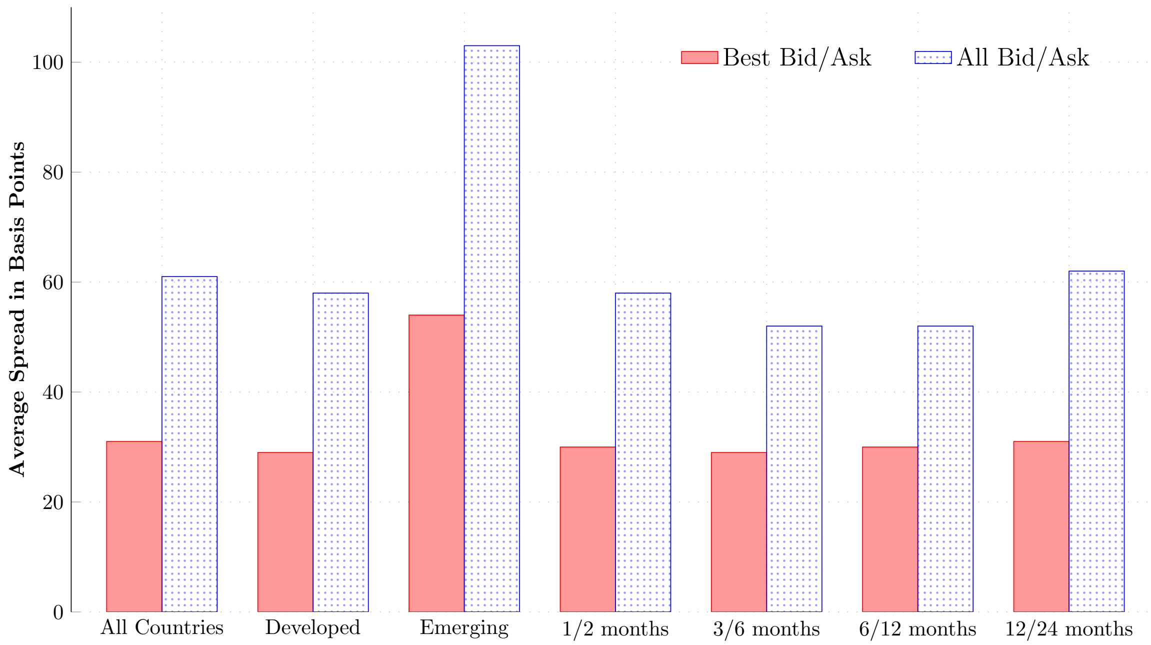

Figure 5. FVA Bid-Ask Spreads

This figure reports the average bid-ask spread on tradeable forward implied volatilities. The red bar denotes the average spread based on bid/ask prices from multiple dealer banks. The blue bar indicates the average spread based on the best bid/ask price. The sample ranges from October 2009 and January 2014 using a cross-section of 9 developed and 6 emerging countries. Bid and ask prices on forward implied volatilities have been manually collected from the archive of a London-based hedge fund.

Table 1. Descriptive Statistics: Volatility Excess Returns

This table presents descriptive statistics of equally-weighted (Panel A) and GDP-weighted (Panel B) volatility excess returns based on forward volatility agreements. Excess returns are computed using spot and forward model-free implied volatilities constructed as in Britten-Jones and Neuberger (2000) and Jiang and Tian (2005). The table also reports the Sharpe ratio (SR) and the first-order autocorrelation coefficient ac_1. t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. Excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A18 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	1/3 months	3/6 months	6/12 months	12/24 months
	Panel A: Equally-weighted			
mean	-2.90	0.30	-0.17	0.54
	[-3.07]	[0.56]	[-0.40]	[1.39]
sdev	11.80	7.14	5.74	5.59
skew	1.67	1.26	1.08	1.35
$SR \times \sqrt{12}$	-0.85	0.15	-0.10	0.34
ac_1	0.23	0.16	0.16	0.09
	Panel B: GDP-weighted			
mean	-2.82	0.23	-0.15	0.43
	[-3.15]	[0.45]	[-0.40]	[1.21]
sdev	11.41	7.10	5.77	5.46
skew	1.33	0.89	0.77	1.01
$SR \times \sqrt{12}$	-0.86	0.11	-0.09	0.27
ac_1	0.18	0.11	0.06	0.03

Table 2. Predictive Regressions: Volatility Excess Returns

This table presents panel regression estimates with currency fixed effects. In Panel A, the dependent variable is the volatility excess return whereas the explanatory variable is the lagged forward implied volatility premium, both computed using spot and forward model-free implied volatilities constructed as in Britten-Jones and Neuberger (2000) and Jiang and Tian (2005). In Panel B, implied volatilities are replaced with implied variances. The coefficient estimates α and γ should be equal to zero under the null hypothesis. t-statistics (reported in brackets) are based on standard errors robust to heteroscedasticity, cross-sectional, and temporal dependence as in Driscoll and Kraay (1998) and Vogelsang (2012). Excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A19 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	α		γ		$R^2(\%)$
		Panel A: Implied Volatilities			
1/3 months	-0.39	[-0.34]	-0.65	[-4.73]	7.7
3/6 months	0.38	[0.67]	-0.79	[-3.02]	2.6
6/12 months	0.58	[1.12]	-1.52	[-3.22]	1.9
12/24 months	0.29	[0.60]	-1.82	[-3.75]	2.0
		Panel B: Implied Variances			
1/3 months	1.47	[0.50]	-0.65	[-3.86]	5.8
3/6 months	1.68	[1.36]	-0.79	[-2.79]	2.2
6/12 months	1.82	[1.62]	-1.55	[-3.19]	1.8
12/24 months	1.19	[1.15]	-1.91	[-3.73]	1.9

Table 3. Descriptive Statistics: Slope-sorted Portfolios

This table reports descriptive statistics for five portfolios of forward volatility agreements sorted by their implied volatility slopes. Volatility excess returns are computed using spot and forward model-free implied volatilities constructed as in Britten-Jones and Neuberger (2000) and Jiang and Tian (2005). Slopes are computed using 3 months and 24 months model-free implied volatility. The first (last) portfolio P_1 (P_5) contains forward volatility agreements with the highest (lowest) implied volatility slopes. LEV denotes a strategy that equally invests in all five portfolios whereas VCA is a long-short strategy that buys P_5 and sells P_1. The table also reports the Sharpe ratio (SR), the first order autocorrelation coefficient ac_1, and the frequency of portfolio switches ($freq$). t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. The portfolios are rebalanced monthly and excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A20 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	P_1	P_2	P_3	P_4	P_5	LEV	VCA
			Panel A: 1/3 months				
$mean$	-4.66	-3.02	-2.35	-2.42	0.49	-2.39	5.15
	[-3.91]	[-2.82]	[-2.17]	[-2.61]	[0.38]	[-2.31]	[5.91]
$sdev$	16.33	14.08	13.41	12.13	14.16	12.72	12.25
$skew$	2.20	2.76	2.18	1.66	2.51	2.48	-1.34
$SR \times \sqrt{12}$	-0.99	-0.74	-0.61	-0.69	0.12	-0.65	1.46
ac_1	0.19	0.18	0.23	0.14	0.30	0.25	0.08
			Panel B: 3/6 months				
$mean$	-0.83	0.37	0.58	0.44	1.81	0.47	2.64
	[-1.31]	[0.50]	[0.93]	[0.86]	[2.58]	[0.78]	[5.75]
$sdev$	9.42	10.18	7.85	7.82	8.86	8.00	7.09
$skew$	1.64	5.22	1.63	1.30	2.34	2.73	-0.16
$SR \times \sqrt{12}$	-0.30	0.13	0.26	0.19	0.71	0.21	1.29
ac_1	0.08	0.17	0.21	0.02	0.18	0.18	0.00
			Panel C: 6/12 months				
$mean$	-1.13	-0.04	-0.08	-0.01	1.11	-0.03	2.24
	[-2.34]	[-0.06]	[-0.17]	[-0.03]	[1.92]	[-0.06]	[5.67]
$sdev$	7.28	8.47	6.51	6.47	7.49	6.49	6.12
$skew$	1.27	5.56	1.36	1.18	2.72	2.65	0.43
$SR \times \sqrt{12}$	-0.54	-0.02	-0.04	-0.01	0.52	-0.02	1.27
ac_1	0.08	0.19	0.18	0.00	0.20	0.18	0.01
			Panel D: 12/24 months				
$mean$	-0.40	0.38	0.37	0.68	2.10	0.63	2.50
	[-0.86]	[0.67]	[0.83]	[1.67]	[3.63]	[1.37]	[5.67]
$sdev$	7.01	8.05	6.39	6.42	8.14	6.31	6.95
$skew$	1.82	4.89	1.57	1.04	2.85	2.98	1.45
$SR \times \sqrt{12}$	-0.20	0.16	0.20	0.37	0.89	0.34	1.25
ac_1	0.08	0.12	0.12	-0.05	0.14	0.14	-0.04
$freq$	0.26	0.47	0.56	0.56	0.32		

Table 4. Principal Components: Slope-sorted Portfolios

This table presents the loadings for the first (PC_1) and second (PC_2) principal component of slope-sorted portfolios of forward volatility agreements presented in Table 3. In each panel, the last column reports the cumulative share of the total variance (CV) explained by the common factors. The portfolios are rebalanced monthly and excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A21 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	P_1	P_2	P_3	P_4	P_5	CV
Panel A: 1/3 months						
PC_1	0.52	0.46	0.43	0.40	0.42	0.82
PC_2	-0.82	0.10	0.20	0.20	0.49	0.90
Panel B: 3/6 months						
PC_1	0.46	0.53	0.40	0.40	0.44	0.81
PC_2	-0.79	0.01	0.08	0.18	0.58	0.89
Panel C: 6/12 months						
PC_1	0.43	0.54	0.41	0.40	0.45	0.80
PC_2	-0.76	-0.04	0.02	0.13	0.63	0.88
Panel D: 12/24 months						
PC_1	0.42	0.52	0.40	0.40	0.48	0.76
PC_2	-0.36	-0.26	-0.22	-0.10	0.86	0.88

Table 5. Asset Pricing Tests: Risk Price

This table presents cross-sectional asset pricing tests. Both test assets (slope-sorted portfolios of forward volatility agreements) and pricing factors (level and volatility carry strategies) are presented in Table 3. The table reports GMM (first and second-stage) estimates of the factor loadings b, the market price of risk λ, and the cross-sectional R^2. t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. HJ refers to the Hansen and Jagannathan (1997) distance (with simulated p-values in parentheses) for the null hypothesis that the pricing errors per unit of norm is equal to zero. The portfolios are rebalanced monthly and excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. A22 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	b_{LEV}	b_{VCA}	λ_{LEV}	λ_{VCA}	$R^2(\%)$	HJ
			Panel A: 1/3 months			
GMM_1	-0.01	0.03	-2.37	4.75	84.1	0.23
	[-1.31]	[2.67]	[-2.20]	[4.86]		(0.43)
GMM_2	-0.02	0.04	-2.30	4.86	73.0	
	[-1.92]	[4.33]	[-2.45]	[5.63]		
			-2.39	**5.15**		
			Panel B: 3/6 months			
GMM_1	0.01	0.05	0.47	2.59	96.8	0.11
	[1.31]	[5.19]	[0.76]	[5.45]		(0.66)
GMM_2	0.01	0.06	0.38	2.61	89.3	
	[1.07]	[6.16]	[0.75]	[5.78]		
			0.47	**2.64**		
			Panel C: 6/12 months			
GMM_1	0.00	0.06	-0.03	2.23	99.0	0.05
	[-0.41]	[5.61]	[-0.05]	[5.52]		(0.93)
GMM_2	0.00	0.07	-0.16	2.19	97.8	
	[-0.33]	[6.48]	[-0.39]	[5.88]		
			-0.03	**2.24**		
			Panel D: 12/24 months			
GMM_1	0.01	0.05	0.62	2.51	98.6	0.07
	[1.24]	[5.11]	[1.33]	[5.98]		(0.87)
GMM_2	0.01	0.05	0.55	2.40	97.8	
	[1.31]	[5.73]	[1.40]	[6.13]		
			0.63	**2.50**		

Table 6. Country-level Asset Pricing Tests

This table presents country-level cross-sectional tests. The test assets are implied volatility excess returns constructed by going long (short) forward volatility agreements with implied volatility slopes lower (higher) than the median implied volatility slope. The pricing factors are the volatility level (LEV) and the volatility carry (VCA) factors described in Table 3. The table reports Fama-MacBeth estimates of the factor price of risk λ and the cross-sectional R^2. t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. We bold λ when its statistical significance is at 5% (or lower) via 10,000 stationary bootstrap repetitions (e.g., Politis and Romano 1994). Excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A23 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	λ_{LEV}		λ_{VCA}		$R^2(\%)$
1/3 months	-2.94	[-2.05]	**9.13**	[4.29]	48.9
3/6 months	-0.46	[-0.62]	**3.96**	[3.72]	76.0
6/12 months	0.17	[0.29]	**1.93**	[2.25]	75.2
12/24 months	0.66	[1.10]	**2.19**	[2.82]	67.5

Table 7. Asset Pricing Tests: Factor Betas

The table reports least-squares estimates of time-series regressions. Both test assets (slope-sorted portfolios of forward volatility agreements) and pricing factors (level and volatility carry strategies) are presented in Table 3. t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. χ^2_α denotes the test statistics (with p-values in parentheses) for the null hypothesis that all intercepts α are jointly zero. The portfolios are rebalanced monthly and excess returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A24 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	α		β_{LEV}		β_{VCA}		$R^2(\%)$	χ^2_α
				Panel A: 1/3 months				
P_1	0.81	[2.11]	1.04	[29.68]	-0.58	[-13.37]	93.7	(0.21)
P_2	-0.65	[-1.79]	1.04	[30.98]	0.02	[0.64]	87.2	
P_3	-0.24	[-0.52]	0.98	[20.63]	0.05	[1.05]	86.0	
P_4	-0.72	[-2.07]	0.90	[25.49]	0.09	[2.15]	87.6	
P_5	0.81	[2.11]	1.04	[29.68]	0.42	[9.76]	91.6	
				Panel B: 3/6 months				
P_1	0.11	[0.63]	1.01	[48.02]	-0.54	[-14.00]	93.1	(0.76)
P_2	-0.17	[-0.80]	1.18	[9.41]	-0.01	[-0.14]	85.5	
P_3	0.15	[0.71]	0.91	[20.42]	0.00	[-0.03]	85.2	
P_4	-0.20	[-1.24]	0.91	[13.19]	0.08	[2.14]	85.4	
P_5	0.11	[0.63]	1.01	[48.02]	0.46	[12.16]	92.2	
				Panel C: 6/12 months				
P_1	0.07	[0.50]	0.99	[45.07]	-0.52	[-13.89]	92.4	(0.99)
P_2	0.02	[0.09]	1.20	[8.25]	-0.01	[-0.17]	84.2	
P_3	-0.06	[-0.36]	0.92	[15.17]	0.00	[0.01]	84.3	
P_4	-0.10	[-0.66]	0.90	[10.87]	0.05	[1.34]	82.0	
P_5	0.07	[0.50]	0.99	[45.07]	0.48	[12.71]	92.8	
				Panel D: 12/24 months				
P_1	-0.03	[-0.16]	0.99	[39.98]	-0.40	[-9.56]	89.5	(0.94)
P_2	-0.12	[-0.68]	1.18	[10.88]	-0.10	[-2.68]	85.6	
P_3	-0.01	[-0.04]	0.92	[18.20]	-0.08	[-1.96]	82.5	
P_4	0.18	[0.95]	0.91	[11.35]	-0.03	[-0.66]	78.5	
P_5	-0.03	[-0.16]	0.99	[39.98]	0.60	[14.44]	92.2	

Table 8. Exposure to Volatility Carry Risk

This table presents panel estimates from the following specification:

$$rx_{i\ell,t} = \beta LEV_{i\ell,t} + \gamma VCA_{i\ell,t} + \phi X_{i,t-1} + \delta LEV_{i\ell,t} \times X_{i,t-1} + \lambda VCA_{i\ell,t} \times X_{i,t-1} + \alpha + fe + \varepsilon_{i\ell,t},$$

where $rx_{i\ell}$ is the volatility excess return for currency i and maturity combination ℓ, $LEV_{i\ell}$ and $VCA_{i\ell}$ are the volatility level and volatility carry factors constructed as in Table 3 while excluding currency i, X_i is the implied volatility slope for currency i in deviation from the cross-country median value, and fe refers to currency, time (monthly) and maturity fixed effects. t-statistics (reported in brackets) are based on standard errors robust to heteroscedasticity, cross-sectional, and temporal dependence as in Driscoll and Kraay (1998) and Vogelsang (2012). $W_{\gamma\lambda}$ is the Wald test for the null hypothesis that γ and λ are jointly zero. The asterics ***, **, and * denote rejection of the null hypothesis at the 1%, 5%, and 10% confidence level, respectively. Returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg for a cross-section of 20 developed and emerging currencies. Table A25 in the Internet Appendix displays results for a cross-section of 10 developed currencies.

	(1)	(2)	(3)	(4)	(5)	(6)
$LEV_{i\ell,t}$	0.96	0.74	0.69	0.95	0.74	0.69
	[127.62]	[31.81]	[21.38]	[117.08]	[32.40]	[21.66]
$VCA_{i\ell,t}$	0.07	0.11	0.13	0.07	0.11	0.13
	[5.66]	[4.04]	[5.07]	[5.67]	[4.11]	[5.14]
$X_{i,t-1}$	−0.05	−0.07	−0.06	−0.05	−0.06	−0.06
	[−2.52]	[−2.94]	[−2.93]	[−2.47]	[−2.88]	[−2.88]
$LEV_{i\ell,t} \times X_{i,t-1}$				0.21	0.24	0.24
				[1.18]	[1.40]	[1.41]
$VCA_{i\ell,t} \times X_{i,t-1}$	−0.99	−1.02	−1.03	−0.99	−1.01	−1.02
	[−3.04]	[−3.02]	[−3.09]	[−3.15]	[−3.13]	[−3.21]
$W_{\gamma\lambda}$	***	***	***	***	***	***
$R^2(\%)$	58.7	58.8	59.0	58.7	58.9	59.0
# Observations	14,560	14,560	14,560	14,560	14,560	14,560
Time fe		Yes	Yes		Yes	Yes
Currency fe		Yes	Yes		Yes	Yes
Maturity fe			Yes			Yes

Table 9. Country-level Exposure to Volatiliy Carry Risk

This table presents country-level estimates from the following regression

$$rx_{i\ell,t} = \alpha + \beta LEV_{i\ell,t} + \gamma VCA_{i\ell,t} + \phi X_{i,t-1} + \lambda VCA_{i\ell,t} \times X_{i,t-1} + \varepsilon_{i\ell,t},$$

where $rx_{i\ell}$ is the volatility excess return for currency i and maturity combination ℓ, $LEV_{i\ell}$ and $VCA_{i\ell}$ are the volatility level and volatility carry factors constructed as in Table 3 while excluding currency i, and X_i is the implied volatility slope for currency i in deviation from the cross-country median value. t-statistics (reported in brackets) are based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. $W_{\gamma\lambda}$ is the Wald test for the null hypothesis that γ and λ are jointly zero. The superscripts ***, **, and * denote rejection of the null hypothesis at the 1%, 5%, and 10% confidence level, respectively. N denotes the number of observations. Returns are expressed in percentage per month. The sample runs from January 1996 to December 2015 and employs over-the-counter currency options from JP Morgan and Bloomberg.

	$LEV_{i\ell,t}$		$VCA_{i\ell,t}$		$X_{i,t-1}$		$VCA_{i\ell,t} \times X_{i,t-1}$		$R^2(\%)$	$W_{\gamma\lambda}$	N
AUD	1.03	[31.06]	0.08	[2.66]	-0.07	[-1.72]	-1.63	[-3.28]	75.4	***	956
BRL	1.05	[115.74]	-0.28	[-2.23]	-0.06	[-1.29]	0.98	[0.89]	54.7	*	476
CAD	0.87	[21.89]	0.05	[1.81]	-0.15	[-6.02]	0.07	[0.40]	59.2		956
CHF	0.83	[14.16]	0.00	[-0.06]	-0.01	[-0.29]	-3.00	[-4.40]	62.8	***	956
CZK	0.81	[23.85]	0.18	[4.93]	-0.13	[-2.01]	-0.54	[-0.38]	71.1	***	476
DKK	0.93	[35.92]	0.07	[1.92]	-0.07	[-2.49]	-3.68	[-8.38]	77.7	***	956
EUR	0.92	[29.57]	0.18	[7.92]	-0.03	[-1.43]	-2.94	[-9.16]	82.1	***	812
GBP	0.93	[33.09]	0.10	[3.82]	0.08	[1.67]	-1.79	[-3.43]	65.5	***	956
HUF	0.87	[22.89]	0.03	[0.81]	-0.05	[-0.80]	1.82	[2.01]	72.9	**	428
JPY	0.66	[13.62]	0.01	[0.21]	-0.09	[-3.98]	-0.16	[-0.41]	41.3		956
KRW	1.36	[23.06]	0.10	[0.78]	-0.14	[-2.68]	-1.91	[-2.30]	65.9	**	476
MXN	1.78	[6.83]	0.54	[2.09]	-0.02	[-0.21]	-5.67	[-3.27]	59.3	***	476
NOK	0.91	[28.83]	0.02	[0.58]	-0.08	[-1.68]	-1.43	[-1.83]	70.7	*	956
NZD	0.83	[30.24]	-0.05	[-1.48]	-0.02	[-0.36]	-3.87	[-6.60]	66.5	***	956
PLN	0.97	[14.17]	0.11	[2.10]	0.00	[-0.01]	-1.73	[-1.73]	76.4	***	476
SEK	0.93	[40.65]	0.01	[0.24]	-0.02	[-0.58]	-3.21	[-5.23]	77.9	***	956
SGD	1.09	[12.62]	-0.06	[-0.90]	-0.01	[-0.24]	-1.16	[-1.73]	76.2	***	476
TRY	0.90	[15.15]	0.21	[0.79]	0.06	[0.82]	-2.04	[-1.21]	42.9	*	476
TWD	0.88	[9.01]	-0.23	[-1.47]	-0.07	[-2.90]	0.92	[1.93]	36.4		716
ZAR	0.90	[13.02]	-0.08	[-1.00]	-0.03	[-1.07]	-0.74	[-1.19]	56.3	*	668

Table 10. Asset Pricing Tests with Alternative Risk Factors

This table presents the estimates of the alpha parameters from time-series asset pricing tests. The test assets (slope-sorted portfolios) are presented in Table 3. In addition to the level factor, the set of traded factors includes dollar, carry, global imbalance, foreign exchange volatility, and currency liquidity factors in Panel A; the Fama-French global equity factors, i.e., market excess return, size, value, profitability, and investment factors in Panel Ab; and the VIX futures returns ranging from 1-month to 6-month in Panel C. The superscripts ***, **, and * denote statistical significance at the 1%, 5%, and 10% confidence level, respectively, based on Newey and West (1987) standard errors with Andrews (1991) optimal lag selection. χ^2_α denotes the test statistics (with p-values in parentheses) for the null hypothesis that all intercepts α are jointly zero. Excess returns are expressed in percentage per month and range from January 1996 to December 2015 using a cross-section of 20 developed and emerging currencies. Tables A11–A14 in the Internet Appendix presents the estimation results for the rest of the parameters. Tables A26–A28 in the Internet Appendix displays similar results for a cross-section of 10 developed currencies.

	1/3 months	3/6 months	6/12 months	12/24 months	χ^2_α
		Panel A: Currency Risk Factors			
P_1	−3.76***	−0.48	−0.92***	−0.19	(<.01)
P_2	−2.38***	0.63**	0.17	0.66***	
P_3	−1.63***	0.92***	0.13	0.62***	
P_4	−1.96***	0.73***	0.24	0.90***	
P_5	0.73	1.94***	1.31***	2.33***	
		Panel B: Global Equity Risk Factors			
P_1	−4.06***	−0.40	−0.85***	−0.05	(<.01)
P_2	−2.46***	0.79***	0.26	0.62***	
P_3	−2.05***	0.76***	0.09	0.50**	
P_4	−2.06***	0.59***	0.19	0.79***	
P_5	1.26*	2.30***	1.46***	2.32***	
		Panel C: VIX Futures Returns			
P_1	−2.65**	0.23	−0.16*	0.25	(<.01)
P_2	−3.07***	−2.65**	0.17	−2.65**	
P_3	−2.06***	1.01***	0.30	−2.65**	
P_4	−2.33***	0.68***	0.16	0.72***	
P_5	0.01	1.90***	1.35***	1.99***	

Table 11. Descriptive Statistics: Traded vs. Quoted Implied Volatility

This table presents descriptive statistics of traded and quoted implied volatilities from over-the-counter currency options. In Panel A, for each currency pair, we first extract intraday implied volatilities from traded options and then aggregate them at the daily frequency using a volume-weighted average. Traded options are collected from the Depository Trust & Clearing Corporation (DTCC) and synchronized at the second level with spot exchange rates from Thomson Reuters Tick History. In Panel B, for each currency pair, we extract the corresponding implied volatilities matched by strikes and maturities from quotes and then aggregate them at the daily frequency via a simple average. Quotes on implied volatilities are provided by JP Morgan and Bloomberg. Implied volatilities are expressed in percentage per annum and their maturity ranges between one month and two years. Q_{95} and Q_5 denote the 95th and 5th percentile, respectively. The sample runs from March 2013 and April 2019 and includes a cross-section of 20 developed and emerging market currencies.

	mean	*sdev*	*skew*	Q_5	Q_{95}	*mean*	*sdev*	*skew*	Q_5	Q_{95}
	Panel A: Implied Volatility from Trades					Panel B: Implied Volatility from Quotes				
AUD	10.15	1.97	0.42	7.34	13.32	10.04	1.87	0.25	7.33	13.09
CAD	8.04	1.53	0.51	5.76	10.84	7.86	1.42	0.45	5.72	10.51
CHF	8.71	2.08	1.38	6.01	12.29	8.56	1.94	0.87	5.99	11.96
DKK	9.89	2.81	0.02	7.35	12.49	9.53	2.39	−0.10	6.86	11.76
EUR	8.64	1.85	0.60	6.03	12.15	8.47	1.78	0.51	5.92	11.85
GBP	9.12	2.38	1.51	5.92	13.25	8.96	2.24	1.44	5.90	13.01
JPY	9.77	2.14	0.44	6.63	13.30	9.57	2.16	0.36	6.38	13.22
NOK	10.33	2.43	1.35	7.43	14.69	10.16	2.19	0.89	7.42	14.50
NZD	10.95	2.25	0.77	8.00	14.47	10.81	2.09	0.28	7.92	14.21
SEK	10.04	2.14	3.00	7.56	13.95	9.78	1.73	0.97	7.58	13.36
BRL	15.67	3.75	0.82	10.66	22.98	15.28	3.49	0.62	10.35	21.85
CZK	10.17	2.14	0.28	6.51	14.37	10.24	1.93	0.04	7.22	13.74
HUF	10.80	2.31	0.49	7.65	14.86	10.62	2.35	0.64	7.48	14.89
KRW	9.70	1.85	0.04	6.61	12.63	9.63	1.77	0.06	6.72	12.46
MXN	12.38	2.60	0.11	7.47	16.48	12.29	2.45	0.20	7.79	16.20
PLN	11.13	2.24	0.80	7.85	14.74	10.95	1.90	0.20	7.96	13.93
SGD	5.89	1.60	1.22	3.80	8.61	5.73	1.40	0.43	3.77	8.18
TRY	14.55	4.92	1.70	8.78	23.72	14.36	4.87	1.98	9.07	24.04
TWD	4.77	1.19	0.46	3.12	6.61	4.31	1.14	1.05	3.04	5.76
ZAR	16.41	3.18	0.20	11.33	21.68	16.29	2.98	0.09	11.46	21.20

Table 12. Traded vs. Quoted Implied Volatility

This table presents panel regression estimates based on the following specification:

$$IVOL^T_{ik\tau,t} = \alpha + \beta\, IVOL_{ik\tau,t} + \gamma' X_{i,t} + fe + \varepsilon_{ik\tau,t},$$

where $IVOL^T_{ik\tau,t}$ is the traded implied volatility for currency i, strike k, and maturity τ on day t, $IVOL_{ik\tau,t}$ is the corresponding quoted implied volatility, $X_{i,t}$ denotes a set of currency-specific variables (option notional value, interest rate differential, foreign exchange liquidity, and spot exchange rate return), and fe refers to hour, currency, maturity, and time (monthly) fixed effects. t_β denotes t-statistics for the null hypothesis that β is equal to one. Implied volatilities are expressed in percentage per annum and their maturity ranges between one month and two years. t-statistics (reported in brackets) are based on standard errors by currency and time (month) dimension. The sample ranges from March 2013 and April 2019 using a cross-section of 20 developed and emerging currencies. Tables A30 in the Internet Appendix displays similar results for a cross-section of 10 developed currencies. Traded options are collected from the Depository Trust & Clearing Corporation (DTCC) whereas quoted implied volatilities are provided by JP Morgan and Bloomberg.

	(1)	(2)	(3)	(4)	(5)	(6)
IVOL	1.01	1.00	0.99	1.00	0.99	0.99
	[73.65]	[118.63]	[122.97]	[96.70]	[85.79]	[87.55]
t_β	[0.83]	[−0.59]	[−0.88]	[−0.40]	[−0.55]	[−0.76]
Option Notional				−0.76	−0.84	−0.90
				[−0.98]	[−1.13]	[−1.23]
Interest Rate Differential				0.02	0.01	0.01
				[3.46]	[0.45]	[0.51]
FX Liquidity				−0.61	0.22	0.22
				[−1.66]	[0.80]	[0.80]
FX Returns				−0.01	0.01	0.01
				[−1.39]	[0.79]	[0.77]
RMSE	2.53	2.52	2.52	2.53	2.52	2.52
$R^2(\%)$	66.5	66.6	66.7	66.6	66.6	66.7
# Observations	1,104,115	1,104,115	1,104,115	1,104,115	1,104,115	1,104,115
Hour fe		Yes	Yes		Yes	Yes
Time fe		Yes	Yes		Yes	Yes
Currency fe		Yes	Yes		Yes	Yes
Maturity fe			Yes			Yes

Table 13. Traded vs. Synthetic Forward Volatility

This table presents panel regression estimates based on the following specification:

$$FVOL^{T}_{ij\ell,t} = \alpha + \beta\, FVOL_{i\ell,t} + \gamma\, Spread_{ij\ell,t} + fe + \varepsilon_{ij\ell,t},$$

where $FVOL^{T}_{ij\ell,t}$ is the tradeable forward implied volatility on day t for currency i, dealer j, and maturity combination ℓ, $FVOL_{i\ell,t}$ is the synthetic forward implied volatility for currency i and maturity ℓ, $Spread_{ij\ell,t}$ is the bid-ask spread on the tradeable volatility, and fe refers to dealer, currency, maturity, and time (monthly) fixed effects. t_β denotes t-statistics for the null hypothesis that β is equal to one. t-statistics (reported in brackets) are based on standard errors robust to heteroscedasticity, cross-sectional, and temporal dependence as in Driscoll and Kraay (1998) and Vogelsang (2012). Volatilities are expressed in percentage per annum and the sample ranges from October 2009 to January 2014 for a cross-section of 15 developed and emerging currencies. Tables A31 in the Internet Appendix displays similar results for a cross-section of 9 developed currencies. Tradeable volatilities have been manually collected from the archive of a London based hedge fund. Mid synthetic volatilities are computed using data provided by JP Morgan and Bloomberg.

	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Mid Forward Implied Volatility			
$FVOL$	0.90	1.00	0.99	0.89	0.99	0.99
	[22.81]	[15.44]	[15.67]	[23.48]	[15.19]	[15.40]
t_β	[−2.64]	[−0.07]	[−0.13]	[−2.88]	[−0.09]	[−0.14]
Spread				−3.14	−0.39	−0.15
				[−1.06]	[−0.61]	[−0.17]
$RMSE$	0.63	0.37	0.36	0.63	0.37	0.36
R^2(%)	91.5	97.1	97.2	91.6	97.1	97.2
			Panel B: Bid Forward Implied Volatility			
$FVOL$	0.89	0.99	0.99	0.87	0.97	0.97
	[22.02]	[15.89]	[16.03]	[23.94]	[15.19]	[15.41]
t_β	[−2.84]	[−0.12]	[−0.19]	[−3.57]	[−0.42]	[−0.48]
Spread				−9.11	−6.74	−6.48
				[−3.71]	[−8.36]	[−6.33]
$RMSE$	0.65	0.39	0.37	0.61	0.36	0.36
R^2(%)	90.8	96.8	97.0	92.0	97.2	97.2
			Panel C: Ask Forward Implied Volatility			
$FVOL$	0.91	1.00	0.99	0.91	1.02	1.01
	[23.14]	[14.92]	[15.22]	[23.02]	[15.19]	[15.38]
t_β	[−2.38]	[−0.03]	[−0.08]	[−2.23]	[0.24]	[0.19]
Spread				2.83	5.96	6.19
				[0.81]	[8.66]	[6.56]
$RMSE$	0.65	0.40	0.39	0.65	0.38	0.37
R^2(%) 91.2	96.7	96.9	91.3	97.1	97.1	
# Observations	1,205	1,205	1,205	1,205	1,205	1,205
Currency fe		Yes	Yes		Yes	Yes
Time fe		Yes	Yes		Yes	Yes
Dealer fe			Yes			Yes
Maturity fe			Yes			Yes

References

Aït-Sahalia, Y., M. Karaman, and L. Mancini, 2016, "The Term Structure of Equity and Variance Risk Premia," Working Paper, Princeton University.

Ammann, M., and R. Buesser, 2013, "Variance Risk Premiums in Foreign Exchange Markets," *Journal of Empirical Finance*, 23, 16–32.

Andrews, D., 1991, "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817–58.

Bakshi, G., and N. Kapadia, 2003, "Delta-Hedged Gains and the Negative Market Volatility Premium," *Review of Financial Studies*, 16, 527–566.

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

Bakshi, G., G. Panayotov, and G. Skoulakis, 2011, "Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios," *Journal of Financial Economics*, 100, 475–495.

Barras, L., and A. Malkhozov, 2016, "Does Variance Risk Have Two Prices? Evidence from the Equity and Option Markets," *Journal of Financial Economics*, 1, 79–92.

Bekaert, G., and G. Panayotov, 2016, "Good Carry, Bad Carry," Working Paper, Columbia Business School.

BIS, 2016a, *OTC Derivatives Statistics at end-June 2016*. Bank for International Settlements, Basel.

——— , 2016b, *Triennial Central Bank Survey of Foreign Exchange and OTC Derivatives Markets Activity in 2016*. Bank for International Settlements, Basel.

Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.

Brenner, M., and M. Subrahmanyam, 1988, "A Simple Formula to Compute the Implied Standard Deviation," *Financial Analysts Journal*, 44, 80–83.

Britten-Jones, M., and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839–866.

Broadie, M., M. Chernov, and M. Johannes, 2009, "Understanding Index Options Returns," *Review of Financial Studies*, 22, 4493–4529.

Brunnermeier, M., S. Nagel, and L. Pedersen, 2009, "Carry Trades and Currency Crashes," *NBER Macroeconomics Annual*, 23, 313–347.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo, 2011, "Do Peso Problems Explain the Returns to the Carry Trade?," *Review of Financial Studies*, 24, 853–891.

Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22, 1311–1341.

Christoffersen, P., S. Heston, and K. Jacobs, 2009, "The Shape and Term structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work so Well," *Management Science*, 55, 1914–1932.

Colacito, R., M. Croce, F. Gavazzoni, and R. Ready, 2018, "Currency Risk Factors in a Recursive Multi-Country Economy," *Journal of Finance*, 73, 2719–2756.

Coval, J., and T. Shumway, 2001, "Expected Option Returns," *Journal of Finance*, 56, 983–1009.

Della Corte, P., T. Ramadorai, and L. Sarno, 2016, "Volatility Risk Premia and Exchange Rate Predictability," *Journal of Financial Economics*, 120, 21–40.

Della Corte, P., S. J. Riddiough, and L. Sarno, 2016, "Currency Risk Premia and Global Imbalances," *Review of Financial Studies*, 29, 2161–2193.

Della Corte, P., L. Sarno, and I. Tsiakas, 2009, "An Economic Evaluation of Empirical Exchange Rate Models," *Review of Financial Studies*, 22, 3491–3530.

Della Corte, P., L. Sarno, and I. Tsiakas, 2011, "Spot and Forward Volatility in Foreign Exchange," *Journal of Financial Economics*, 100, 496–513.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez, 2016, "The Price of Variance Risk," *Journal of Financial Economics*, 123, 225–250.

Driscoll, J. C., and A. C. Kraay, 1998, "Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data," *Review of Economics and Statistics*, 80, 549–560.

Eraker, B., and Y. Wu, 2016, "Explaining the Negative Returns to VIX Futures and ETNs: An Equilibrium Approach," *Journal of Financial Economics*, 125, 72–98.

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

Fama, E. F., and K. R. French, 2016, "International Tests of a Five-factor Asset Pricing Model," *Journal of Financial Economics*, 123, 441–463.

Farhi, E., S. P. Fraiberger, X. Gabaix, R. G. Ranciere, and A. Verdelhan, 2015, "Crash Risk in Currency Markets," Working Paper, Harvard University.

Gabaix, X., and M. Maggiori, 2015, "International Liquidity and Exchange Rate Dynamics," *Quarterly Journal of Economics*, 130, 1369 –1420.

Garman, M. B., and S. W. Kohlhagen, 1983, "Foreign Currency Option Values," *Journal of International Money and Finance*, 2, 231–237.

Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

Hansen, L. P., and R. Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.

Jagannathan, R., and Z. Wang, 1996, "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance*, 51, 3–53.

Jiang, G. J., and Y. S. Tian, 2005, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, 1305–1342.

Johnson, T., 2017, "Risk Premia and the VIX Term Structure," *Journal of Financial and Quantitative Analysis*, 52, 2461–2490.

Jurek, J. W., 2014, "Crash-neutral Currency Carry Trades," *Journal of Financial Economics*, 113, 325–347.

Kozhan, R., A. Neuberger, and P. Schneider, 2013, "The Skew Risk Premium in Index Option Prices," *Review of Financial Studies*, 26, 2174–2203.

Lewellen, J., S. Nagel, and J. Shanken, 2010, "A Skeptical Appraisal of Asset Pricing Tests," *Journal of Financial Economics*, 96, 175–194.

Londono, J. M., and H. Zhou, 2016, "Variance Risk Premiums and the Forward Premium Puzzle," *Journal of Financial Economics*, 124, 415–440.

Lustig, H., N. Roussanov, and A. Verdelhan, 2011, "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24, 3731–3777.

Lustig, H., N. L. Roussanov, and A. Verdelhan, 2014, "Countercyclical Currency Risk Premia," *Journal of Financial Economcis*, 111, 527–553.

Lustig, H., and A. Verdelhan, 2007, "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk," *American Economic Review*, 97, 89–117.

Marion, K., 2010, "Exchange Rates during Financial Crises," *BIS Quarterly Review*, March, 39–50.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf, 2012, "Carry Trades and Global Foreign Exchange Volatility," *Journal of Finance*, 67, 681–718.

Newey, W. K., and K. D. West, 1987, "A Simple Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Politis, D. N., and J. P. Romano, 1994, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89, 1301–1313.

Richmond, R. J., 2019, "Trade Network Centrality and Currency Risk Premia," *Journal of Finance*, 74, 1315–1361.

Todorov, V., 2016, "Variance Risk-premium Dynamics: The Role of Jumps," *Review of Financial Studies*, 23, 345–383.

Verdelhan, A., 2018, "The Share of Systematic Variation in Bilateral Exchange Rates," *Journal of Finance*, 73, 375–418.

Vogelsang, T. J., 2012, "Heteroskedasticity, Autocorrelation, and Spatial Correlation Robust Inference in Linear Panel models with Fixed-effects," *Journal of Econometrics*, 166, 303–319.

Zviadadze, I., 2017, "Term Structure of Consumption Risk Premia in the Cross Section of Currency Returns," *Journal of Finance*, 72, 1529–1566.