



City Research Online

City St George's, University of London

Citation: Della Corte, P., Ramadorai, T. & Sarno, L. (2016). Volatility Risk Premia and Exchange Rate Predictability. *Journal of Financial Economics*, 120(1), pp. 21-40. doi: 10.1016/j.jfineco.2016.02.015

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/13156/>

Link to published version: <https://doi.org/10.1016/j.jfineco.2016.02.015>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Volatility Risk Premia and Exchange Rate Predictability^{*†}

Pasquale Della Corte Tarun Ramadorai Lucio Sarno

First version: December 2013 - Revised: December 2014

***Acknowledgements:** We are grateful to Bill Schwert (editor), an anonymous referee, Adrian Buss, John Campbell, Bernard Dumas, Kenneth Froot, Federico Gavazzoni, Philippos Kassimatis, Lars Lochstoer, Elisa Luciano, Stefan Nagel, Andrea Vedolin, Adrien Verdelhan and participants at the Oxford-Man Institute 2013 Conference on Currency Trading and Risk Premia, the 2013 NBER Summer Institute in Asset Pricing, the 5th Global Alpha Forum in New York, the 2013 Annual Conference on Advances in the Analysis of Hedge Fund Strategies at Imperial College London, the 2013 Bloomberg FX Conference, the 2014 MAF Conference, the 2014 ISCEF Conference, and seminar participants at CUHK, EIEF, Erasmus, the HKMA, HKUST, INSEAD, LUMS, National Bank of Serbia, the NCU, NCHU, SAIF, SMU, TCD, University of Edinburgh, University of Kent, University of Manchester, University of Nottingham, and University of St. Gallen for helpful conversations and suggestions. We also thank JP Morgan and Aslan Uddin for the currency implied volatility data used in this study. Della Corte gratefully acknowledges the hospitality of the Hong Kong Institute for Monetary Research at the HKMA. Sarno acknowledges financial support from the ESRC (No. RES-062-23-2340) and the gracious hospitality of CERF at the University of Cambridge, and the Lee Kong Chian School of Business at SMU. All errors remain ours.

[†]Pasquale Della Corte is at Imperial College Business School, Imperial College London; email: p.dellacorte@imperial.ac.uk. Tarun Ramadorai is at the Saïd Business School, Oxford-Man Institute, University of Oxford and CEPR; email: tarun.ramadorai@sbs.ox.ac.uk. Lucio Sarno is at Cass Business School, City University London and CEPR; email: lucio.sarno@city.ac.uk.

Volatility Risk Premia and Exchange Rate Predictability

First version: December 2013 - Revised: February 2015

Abstract

We discover a new currency strategy with highly desirable return and diversification properties, which uses the predictive ability of currency volatility risk premia for currency returns. The volatility risk premium – the difference between expected realized volatility and model-free implied volatility – reflects the costs of insuring against currency volatility fluctuations, and the strategy sells high-insurance-cost currencies and buys low-insurance-cost currencies. A distinctive feature of the strategy's returns is that they are mainly generated by movements in spot exchange rates rather than interest rate differentials. We explore explanations for the profitability of the strategy, which cannot be understood using traditional risk factors.

Keywords: Exchange Rates; Volatility Risk Premium; Predictability, Efficient Currency Portfolios.

JEL Classification: F31; F37.

1 Introduction

For decades, finance practitioners and academics have struggled to understand currency fluctuations. The difficulty of explaining and forecasting nominal exchange rates was systematically documented by Meese and Rogoff (1983), and since then, it has continued to be difficult to find variables able to beat a random walk forecasting model for currencies (e.g., see Engel, Mark, and West, 2008). More recently, the literature on exchange rates has focused on a closely-related question, which is to document high returns to currency investment strategies such as carry and momentum.¹ Analogous to the difficulty of finding definitive answers about the source of currency fluctuations, there has been limited success in explaining the often high returns to these currency investment strategies in terms of compensation for systematic risk.

In this paper, we discover a new currency strategy with high risk-adjusted returns, excellent diversification benefits relative to the set of previously discovered currency strategies, and unusual properties that provide clues as to the underlying drivers of exchange rate movements. The key to this new strategy, which we dub *VRP*, is the significant predictive power of the currency volatility risk premium for changes in spot exchange rates.²

The desirability of the *VRP* strategy does not only derive from its profitability. The strategy is also a useful complement to other widely-studied currency strategies, as it has a low correlation with them. This unusually low correlation partly arises from the excellent performance of *VRP* during crises, and primarily from the fact that the excess returns of *VRP* are almost completely obtained through prediction of changes in exchange rates, rather than from interest rate differentials. This stands in sharp contrast with the performance of

¹See, for example, Lustig and Verdelhan (2007), Ang and Chen (2010), Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011), Lustig, Roussanov, and Verdelhan (2011), Menkhoff, Sarno, Schmeling, and Schrimpf (2012a,b) and Barroso and Santa Clara (2014), who all build currency portfolios to study return predictability and/or currency risk exposure.

²To be clear from the outset, our strategy does not trade volatility products. We simply use the currency volatility risk premium as conditioning information to sort currencies, build currency portfolios, and uncover predictability in currency excess returns and changes in spot exchange rates.

the carry strategy, which has primarily been driven by interest differentials rather than spot currency returns.³

The currency volatility risk premium is the difference between expected future realized currency volatility, and a model-free measure of implied volatility derived from currency options. A growing literature studies the variance or the volatility risk premium in different asset classes, including equity, bond, and foreign exchange (FX) markets.⁴ In general, this literature has shown that the volatility risk premium is on average negative: expected volatility is higher than historical realized volatility, and since volatility is persistent, expected volatility is also generally higher than future realized volatility.

Understood intuitively, the volatility risk premium represents compensation for providing volatility insurance, that is, the currency volatility risk premium can be interpreted as the cost of insurance against volatility fluctuations in the underlying currency. When it is high – realized volatility is higher than the option-implied volatility – insurance is relatively cheap, and vice versa.

We use the currency volatility risk premium to sort currencies into quintile portfolios at the end of each month. The *VRP* strategy buys currencies with relatively cheap volatility insurance, i.e., the highest volatility risk premium quintile, and sells short currencies with relatively expensive volatility insurance, i.e., the lowest volatility risk premium quintile. We track returns on this trading strategy over the subsequent period, meaning that these returns are purely out-of-sample, conditioning only on information available at the time of portfolio construction. We find that the performance of the strategy is remarkable, delivering per-

³We use interchangeably the terms spot currency returns and exchange rate returns to define the change in nominal exchange rates over time; similarly we use interchangeably the terms excess returns or portfolio returns to refer to the returns from implementing a long-short currency trading strategy that buys and sells currencies on the basis of some characteristic.

⁴See, for example, Carr and Wu (2009), Eraker (2008), Bollerslev, Tauchen, and Zhou (2009), Todorov (2010), Drechsler and Yaron (2010), Han and Zhou (2011), Mueller, Vedolin, and Yen (2011), Londono and Zhou (2012) and Buraschi, Trojani, and Vedolin (2014).

formance per unit of volatility that is better than or comparable to the highest of the set of widely-studied currency investment strategies that we consider.

Unusually for currency investment strategies, the performance of *VRP* stems virtually entirely from the predictability of spot exchange rates rather than from interest rate differentials. That is, currencies with relatively cheap volatility insurance tend to appreciate and those with relatively more expensive volatility insurance tend to depreciate over the subsequent month. The observed predictability of spot exchange rates associated with *VRP* is far stronger than that arising from carry (which is perhaps unsurprising given the well-documented fact that interest differentials are the proximate component of carry returns), and perhaps more importantly, stronger than that associated with currency momentum or any of the other well-known currency trading strategies that we consider. As mentioned earlier, this is part of the reason for the diversification benefits that the *VRP* strategy offers in a currency portfolio.

The contribution of our paper is purely empirical, and we do not have a formal theoretical model that links the volatility risk premium or its determinants to spot returns. However, we do provide empirical evidence on possible interpretations of our results. First, we consider the possibility that returns from the *VRP* strategy reflect compensation for risk, and test the pricing power of conventional risk factors for its returns using standard linear asset pricing models. We find no evidence that *VRP* returns can be explained by various sets of factors that have been used to explain time-series and cross-sectional variation in the returns to trading strategies more generally, and currency strategies more specifically.

We then extend our search for risk-compensation to check whether *VRP* returns capture fluctuations in aversion to global volatility risk. We check the relationship between *VRP* returns and global volatility risk in two ways: first, by using cross-sectional asset pricing tests of volatility risk premium-sorted portfolios on a global FX volatility risk factor, and second, by estimating time-varying loadings of currency returns on various proxies for global

volatility risk and building portfolios sorted on these estimated loadings. Neither of these tests produces evidence consistent with the proposed explanation. Indeed, the long-short strategy generated from estimated loadings on the global volatility risk factor produces substantially lower average returns than *VRP*; moreover, these returns are virtually uncorrelated with *VRP* returns. In sum, the data appear to reject an explanation based on fluctuations in aversion to global volatility risk and, more generally reject the hypothesis that *VRP* returns can be explained by exposure to common risk factors.

A second explanation that we consider relies on limits to arbitrage, and its effects on the interaction between hedgers and speculators in the currency market. There is a growing theoretical and empirical literature suggesting that such interactions are important in asset return determination (see, for example, Acharya, Lochstoer, and Ramadorai, 2013; Adrian, Etula, and Muir, 2013; and Gromb and Vayanos, 2010 for an excellent survey of the literature). Such an explanation for our results would rely on time-variation in the amount of arbitrage capital available to natural providers of currency volatility insurance (“speculators”), such as financial institutions or hedge funds. It would also require that risk-averse natural “hedgers” of currencies such as multinational firms are more (less) willing to hedge and hold currencies with relatively inexpensive (more expensive) volatility insurance. Such an explanation predicts price impact in the spot market in response to purchases or sales of currencies based on their relative cost of volatility insurance.

While we do not have a formal theoretical model of such a mechanism, we expect that when funding liquidity is lower (i.e., times of high capital constraints on speculators), and demand for volatility protection is higher (i.e., times of increased risk aversion of natural hedgers), we should detect increases in the spread in the cost of volatility insurance across currencies, as well as the spread in spot exchange rate returns across portfolios. We do find that increases in the TED spread – a commonly used proxy for funding liquidity (see, for example, Garleanu

and Pedersen, 2011) – are associated with higher *VRP* returns. Fluctuations in risk aversion, proxied by changes in the VIX, add significant additional explanatory power when interacted with the TED spread. We also measure capital flows to currency and global macro hedge funds, and find that when hedge fund flows are high, signifying increased funding and thus lower hedge fund capital constraints, the returns to *VRP* are lower and vice versa.⁵ In sum, there is some evidence consistent with limits to arbitrage in the currency market constituting part of the explanation for our results.

The predictive power of volatility risk premia for spot exchange rate returns is particularly interesting given the dismal performance of empirical exchange rate models in forecasting out-of-sample nominal exchange rate changes (see, for example, Meese and Rogoff, 1983, and Engel, Mark and West, 2008). Our analysis suggests that there is value in searching for predictive information in variables that are outside the conventional menu provided by international macro models of exchange rate determination. In particular, our results suggest that explaining the economic drivers of cross-sectional and time-series variation in currency volatility risk premia may help to shed light on the exchange rate determination puzzle.

The results in this paper also highlight intriguing similarities between the behavior of equity and currency options and their underlying asset markets. Several authors (see, for example, Goyal and Saretto, 2009, Bali and Hovakimian, 2009, and Buss and Vilkov, 2012) show that volatility risk premia have predictive power for the cross-section of stock returns. The similarity of the statistical relationships between equity options and underlying stocks, and currency options and underlying currencies suggests that there may be more general structural determinants of this relationship that are common across these markets. We leave

⁵Using CFTC data, we also find that commercial traders sell currencies which are more expensive to insure and buy currencies which are cheaper to insure, with financial traders trading in the opposite direction. This evidence also links our work to another stream of the exchange rate literature on forecasting currency returns using currency order flow. For example, Froot and Ramadorai (2005), Evans and Lyons (2005) and Rime, Sarno, and Sojli (2010) show that order flow has predictive power for exchange rate movements.

the exploration of these important issues to future work.

The paper is structured as follows. Section 2 defines the volatility risk premium and its measurement in currency markets. Section 3 describes our data and some descriptive statistics. Section 4 presents our main empirical results on the volatility risk premium-sorted strategy, while Section 5 investigates alternative mechanisms that could explain our findings. Section 6 concludes. A separate Internet Appendix provides robustness tests and additional supporting analyses.

2 Foreign exchange volatility risk premia

2.1 Volatility swap

A volatility swap is a forward contract on the volatility “realized” on the underlying asset over the life of the contract. The buyer of a volatility swap written at time t , and maturing at time $t + \tau$, receives the payoff (per unit of notional amount):

$$VP_{t,\tau} = (RV_{t,\tau} - SW_{t,\tau}) \quad (1)$$

where $RV_{t,\tau}$ is the realized volatility of the underlying, $SW_{t,\tau}$ is the volatility swap rate, and both $RV_{t,\tau}$ and $SW_{t,\tau}$ are defined over the life of the contract from time t to time $t + \tau$, and quoted in annual terms. However, while the realized volatility is determined at the maturity date $t + \tau$, the swap rate is agreed at the start date t .

The value of a volatility swap contract is obtained as the expected present value of the future payoff in a risk-neutral world. This implies, because $VP_{t,\tau}$ is expected to be 0 under the risk-neutral measure, that the volatility swap rate equals the risk-neutral expectation of the realized volatility over the life of the contract:

$$SW_{t,\tau} = E_t^{\mathbb{Q}} [RV_{t,\tau}] \quad (2)$$

where $E_t^{\mathbb{Q}}[\cdot]$ is the expectation under the risk-neutral measure $\mathbb{Q}$, $RV_{t,\tau} = \sqrt{\tau^{-1} \int_t^{t+\tau} \sigma_s^2 ds}$, and σ_s^2 denotes the (stochastic) volatility of the underlying asset.

2.2 Volatility swap rate

We synthesize the volatility swap rate using the model-free approach derived by Britten-Jones and Neuberger (2000), and further refined by Demeterfi, Derman, Kamal and Zou (1999), Jiang and Tian (2005), and Carr and Wu (2009).

Building on the pioneering work of Breeden and Litzenberger (1978), Britten-Jones and Neuberger (2000) derive the model-free implied volatility entirely from no-arbitrage conditions and without using any specific option pricing model. Specifically, they show that the risk-neutral expected integrated return variance between the current date and a future date is fully specified by the set of prices of call options expiring on the future date, provided that the price of the underlying evolves continuously with constant or stochastic volatility but without jumps.

Demeterfi, Derman, Kamal, and Zou (1999) show that the Britten-Jones and Neuberger (2000) solution is equivalent to a portfolio that combines a dynamically rebalanced long position in the underlying, and a static short position in a portfolio of options and a forward that together replicate the payoff of a “log contract.”⁶ The replicating portfolio strategy captures variance exactly, provided that the portfolio of options contains all strikes with the appropriate weights to match the log payoff. Jiang and Tian (2005) further demonstrate that the model-free implied variance is valid even when the underlying price exhibits jumps, thus relaxing the diffusion assumptions of Britten-Jones and Neuberger (2000).

The annualized risk-neutral expectation of the return variance between two dates t and $t + \tau$ can be formally computed by integrating option prices expiring on these dates over an

⁶The log contract is an option whose payoff is proportional to the log of the underlying at expiration (Neuberger, 1994).

infinite range of strike prices:

$$E_t^{\mathbb{Q}} [RV_{t,\tau}^2] = \kappa \left(\int_0^{F_{t,\tau}} \frac{1}{K^2} P_{t,\tau}(K) dK + \int_{F_{t,\tau}}^{\infty} \frac{1}{K^2} C_{t,\tau}(K) dK \right) \quad (3)$$

where $P_{t,\tau}(K)$ and $C_{t,\tau}(K)$ are the put and call prices at t with strike price K and maturity date $t + \tau$, $F_{t,\tau}$ is the forward price matching the maturity date of the options, S_t is the price of the underlying, $\kappa = 2 \exp(i_{t,\tau}\tau)$, and $i_{t,\tau}$ is the τ -period domestic riskless rate. The risk-neutral expectation of the return variance in Equation (3) delivers the strike price of a variance swap $E_t^{\mathbb{Q}} [RV_{t,\tau}^2]$, and is referred to as the model-free implied variance.

Even though variance emerges naturally from a portfolio of options, it is volatility that participants prefer to quote, as the payoff of a variance swap is convex in volatility and large swings in volatility, as we observed during the recent financial crisis, are more likely to cause large profits and losses to counterparties. Therefore, our empirical analysis focuses on volatility swaps, and we synthetically construct the strike price of this contract as

$$E_t^{\mathbb{Q}} [RV_{t,\tau}] = \sqrt{E_t^{\mathbb{Q}} [RV_{t,\tau}^2]} \quad (4)$$

and refer to it as model-free implied volatility.⁷

Computing model-free implied volatility requires the existence of a continuum in the cross-section of option prices at time t with maturity date τ . In the FX market, over-the-counter options are generally quoted in terms of Garman and Kohlhagen (1983) implied volatilities at fixed deltas. Liquidity is generally spread across five levels of deltas. From these quotes, we extract five strike prices corresponding to five plain vanilla options, and follow Jiang and Tian (2005), who present a simple method to implement the model-free approach when option prices are only available on a finite number of strikes.

Specifically, we use a cubic spline around these five implied volatility points. This interpolation method is standard in the literature (e.g., Bates, 1991; Campa, Chang, and Reider,

⁷See Della Corte, Sarno, and Tsiakas (2011) for a detailed discussion of convexity bias in this formula.

1998; Jiang and Tian, 2005; Della Corte, Sarno, and Tsiakas, 2011) and has the advantage that the implied volatility smile is smooth between the maximum and minimum available strikes. We then compute the option values using the Garman and Kohlhagen (1983) valuation formula,⁸ and use trapezoidal integration to solve the integral in Equation (3). This method introduces two types of approximation errors: (i) the truncation errors arising from observing a finite number, rather than an infinite set of strike prices, and (ii) a discretization error resulting from numerical integration. Jiang and Tian (2005), however, show that both errors are small, if not negligible, in most empirical settings.⁹

2.3 Volatility risk premium

In this paper we study the predictive information content in volatility risk premia for future exchange rate returns. To this end, we work with the ex-ante payoff or “expected volatility premium” to a volatility swap contract. The volatility risk premium can be thought of as the difference between the physical and the risk-neutral expectations of the future realized volatility.¹⁰ Formally, the τ -period volatility risk premium at time t is defined as

$$VRP_{t,\tau} = E_t^{\mathbb{P}}[RV_{t,\tau}] - E_t^{\mathbb{Q}}[RV_{t,\tau}] \quad (5)$$

where $E_t^{\mathbb{P}}[\cdot]$ is the conditional expectation operator at time t under the physical measure $\mathbb{P}$. Following Bollerslev, Tauchen, and Zhou (2009), we proxy $E_t^{\mathbb{P}}[RV_{t,\tau}]$ by simply using the lagged realized volatility, i.e., $E_t^{\mathbb{P}}[RV_{t,\tau}] = RV_{t-\tau,\tau} = \sqrt{\frac{252}{\tau} \sum_{i=0}^{\tau} r_{t-i}^2}$, where r_t is the daily log return on the underlying security. This approach is widely used for forecasting exercises – it makes $VRP_{t,\tau}$ directly observable at time t , requires no modeling assumptions,

⁸This valuation formula can be thought of as the Black and Scholes (1973) formula adjusted for having both domestic and foreign currency paying a continuous interest rate.

⁹In the Internet Appendix (Table A.10), we present results for different interpolation methods (Castagna and Mercurio, 2007) as well as a model-free approach that is robust to price jumps (Martin, 2012).

¹⁰Several papers define the volatility risk premium as difference between the risk-neutral and the physical expectation. Here we follow Carr and Wu (2009) and take the opposite definition as it naturally arises from the long-position in a volatility swap contract.

and is consistent with the stylized fact that realized volatility is a highly persistent process. Thus, at time t , we measure the volatility risk premium over the $[t, t + \tau]$ time interval as the difference between the ex-post realized volatility over the $[t - \tau, t]$ interval and the ex-ante risk-neutral expectation of the future realized volatility over the $[t, t + \tau]$ interval, i.e., $VRP_{t,\tau} = RV_{t-\tau,\tau} - E_t^{\mathbb{Q}}[RV_{t,\tau}]$.

For our purposes, we view currencies with high $VRP_{t,\tau}$ as those which are relatively cheap to insure at each point in time t , as their expected realized volatility under the physical measure (i.e., the variable against which agents hedge) is lower than the cost of purchasing option-based insurance – which is primarily driven by expected volatility under the risk-neutral measure. Conversely, we consider those currencies with relatively low $VRP_{t,\tau}$ as more expensive to insure at time t .

3 Data and currency portfolios

This section describes the data and the construction of the currency portfolios employed in our analysis. The data comprises spot and forward exchange rates, over-the-counter (OTC) currency options, hedge fund flows, and positions on currency futures and options.

3.1 *Exchange rate data*

We collect daily spot and one-month forward exchange rates (bid and ask prices) vis-à-vis the US dollar (USD) from Barclays and Reuters via Datastream. We use monthly data by sampling end-of-month exchange rates from January 1998 to December 2013. In our empirical exercise, we build currency portfolios using two sets of countries. The first sample comprises Australia, Canada, Denmark, Euro Area, Japan, New Zealand, Norway, Sweden, Switzerland, and the United Kingdom. These 10 countries have the most traded currencies and account for about 90% of the average daily turnover in FX markets according to the

Triennial Survey of the Bank for International Settlements (2013). We refer to this sample as the “Developed” countries sample. The second sample adds the most liquid emerging market currencies to the Developed country sample. Some currencies in this expanded “Developed and Emerging” countries sample may be subject to capital controls and, hence, not be tradable (in large amounts) in practice. To mitigate this concern, we follow Menkhoff *et al.* (2012b) and select the currencies for which the financial openness index of Chinn and Ito (2006) index – a measure of a country’s degree of capital account openness – is greater than or equal to zero. Ultimately, we only consider emerging market economies for which the capital account is sufficiently unrestricted so that trading in this currency can actually take place.¹¹ The final expanded sample includes: 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.

3.2 *Implied volatility data*

We calculate the volatility swap rate described in Section 2 using end-of-month implied volatility data on over-the-counter (OTC) currency options, obtained from JP Morgan. The OTC currency option market is characterized by specific trading conventions. While exchange traded options are quoted at fixed strike prices and have fixed calendar expiration dates, currency options are quoted at fixed deltas and have constant maturities. More importantly, while the former are quoted in terms of option premia, the latter are quoted in terms of Garman and Kohlhagen (1983) implied volatilities on baskets of plain vanilla options.

For a given maturity, quotes are typically available for five different combinations of plain-

¹¹More precisely, we start from 10 emerging market currencies and recursively apply the capital account openness index of Chinn and Ito (2006), available on Hiro Ito’s website. Data are available at the yearly frequency until 2011, and we construct monthly observations by forward filling, i.e., we keep end-of-period data constant until a new observation becomes available. Note that the Chinn-Ito index is not available for Taiwan. In this case, we rely on the capital account liberalization index of Kaminsky and Schmukler (2008), available on Graciela Kaminsky’s website.

vanilla options: at-the-money delta-neutral straddles, 10-delta and 25-delta risk-reversals, and 10-delta and 25-delta butterfly spreads. The delta-neutral straddle combines a call and a put option with the same delta but opposite sign such that the total delta is zero – this is the at-the-money (ATM) implied volatility quoted in the FX market. In a risk reversal, the trader buys an out-of-the money (OTM) call and sells an OTM put with symmetric deltas. The butterfly spread is constructed by buying a strangle and selling a straddle, and is equivalent to the difference between the average implied volatility of an OTM call and an OTM put, and the implied volatility of a straddle. From these data, one can recover the implied volatility smile ranging from a 10-delta put to a 10-delta call.¹² To convert deltas into strike prices, and implied volatilities into option prices, we employ domestic and foreign interest rates, obtained from Bloomberg.

This recovery exercise yields data on plain-vanilla European calls and puts for currency pairs vis-à-vis the US dollar, with maturity of one year. Practitioner accounts suggest that natural hedgers such as corporates prefer to hedge using intermediate-horizon derivative contracts rather than employing the more transactions-costs intensive strategy of rolling over short term positions in currency options. We therefore work with the one-year volatility swap in our empirical analysis.

3.3 *Hedge fund flows*

To construct a measure of new arbitrage capital available to hedge funds, we use data from a large cross-section of hedge funds and funds-of-funds from January 1998 to December 2013, which is consolidated from data in the HFR, CISDM, TASS, Morningstar, and Barclay-Hedge databases, and comprises roughly US\$ 1.5 trillion worth of assets under management (AUM) towards the end of the sample period. Ramadorai (2013), and Patton and Ramadorai (2013)

¹²In market jargon, a 10-delta call is a call whose delta is 0.10 whereas a 10-delta put is a put with a delta equal to -0.10 .

provide a detailed description of the process followed to consolidate these data.

We select a subset of 634 funds from these data, those self-reporting as currency funds or global macro funds, and construct the net flow of new assets to each fund as the change in the fund’s AUM across successive months, adjusted for the returns accrued by the fund over the month – this is tantamount to an assumption that flows arrive at the end of the month, following return accrual. We then normalize the figures by dividing them by the lagged AUM, and then value-weight them across funds to create a single aggregate time-series index of capital flows to currency and global macro funds.¹³

3.4 *Positions on currency futures*

We employ data from the Commitments of Traders report issued by the Commodity Futures Trading Commission (CFTC). The report aggregates the holdings of participants in the US currency futures and options markets (primarily based in Chicago and New York). It is typically released every Friday and reflects the commitments of traders for the prior Tuesday. The CFTC provides a breakdown of aggregate positions held by commercial traders and financial (or non-commercial) traders. The former are merchants, foreign brokers, clearing members or banks using the futures market primarily to hedge their business activities. The latter are hedge funds, financial institutions and individual investors, who are assumed to use the futures market for speculative purposes. We collect data from January 1998 to December 2013 on the Australian dollar, Brazilian real, British pound, Canadian dollar, Euro, Japanese yen, Mexican peso, New Zealand dollar, and Swiss franc relative to the USD dollar.

In our empirical analysis, we construct the net demand of currency options and futures - the difference between long and short positions scaled by the total open interest - for both

¹³We measure the net flow for each fund i as $Flow_t^i = AUM_t^i - AUM_{t-1}^i (1 + r_t^i)$, where AUM_t^i and r_t^i are assets under management and returns at time t , respectively. We then construct the AUM-weighted net flow scaled by the lagged AUM as $Flow_t = \sum_i w_{t-1} Flow_t^i$ where $w_t = (\sum_i AUM_t^i)^{-1}$. Finally, we winsorize $Flow_t$ at the 1 and 99 percentile points each month.

commercial and financial traders. We then examine how the buying and selling actions of different players in the futures and options market relate to the portfolio classifications associated with the *VRP* strategy.

3.5 *Other data*

We also collect monthly data on the VIX index, 3-month LIBOR and 3-month T-bill rate from Bloomberg, monthly data from the Federal Reserve Economic data website, and annual data for the purchasing power parity (PPP) spot rate from the OECD. The latter are published every March, and we retrieve monthly data by forward filling, i.e., we use the last available PPP rate until the next February.¹⁴

3.6 *Currency excess returns*

We define spot and forward exchange rates at time t as S_t and F_t , respectively. Exchange rates are defined as units of US dollars per unit of foreign currency such that an increase in S_t indicates an appreciation of the foreign currency. The excess return on buying a foreign currency in the forward market at time t and then selling it in the spot market at time $t + 1$ is computed as $RX_{t+1} = (S_{t+1} - F_t) / S_t$, which is equivalent to the spot exchange rate return minus the forward premium $RX_{t+1} = ((S_{t+1} - S_t) / S_t) - ((F_t - S_t) / S_t)$. According to the CIP condition, the forward premium approximately equals the interest rate differential $(F_t - S_t) / S_t \simeq i_t - i_t^*$, where i_t and i_t^* represent the domestic and foreign riskless rates respectively, over the maturity of the forward contract. Since CIP holds closely in the data at daily and lower frequencies (e.g., Akram, Rime, and Sarno, 2008), the currency excess return is approximately equal to an exchange rate component (i.e., the exchange rate change) minus an interest rate component (i.e., the interest rate differential): $RX_{t+1} \simeq ((S_{t+1} - S_t) / S_t) -$

¹⁴For Singapore and Taiwan, OECD's PPP spot data are not available and we use data from the Penn World Tables instead.

$(i_t - i_t^*)$.

We construct currency excess returns adjusted for transaction costs using bid-ask quotes on spot and forward rates as in Menkhoff, Sarno, Schmeling and Schrimpf (2012a). The total number of currencies in our portfolios changes over time, and we only include currencies for which we have bid and ask quotes on forward and spot exchange rates in the current and subsequent period.

3.7 Carry trade portfolios

At the end of each period t , we allocate currencies to five portfolios on the basis of their interest rate differential relative to the US, $(i_t^* - i_t)$ or forward premia since $-(F_t - S_t)/S_t = (i_t^* - i_t)$ via CIP. This exercise implies that Portfolio 1 comprises 20% of all currencies – those with the highest interest rate differentials (lowest forward premia), and Portfolio 5 also comprises 20% of all currencies – those with the lowest interest rate differentials (highest forward premia). We refer to the long-short portfolio formed by going long Portfolio 1 and short Portfolio 5 as *CAR*. We compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio, and individually track both the interest rate differential and the spot exchange rate component that make up these excess returns.

Lustig, Roussanov, and Verdelhan (2011) study these currency portfolio returns using their first two principal components. The first principal component implies an equally weighted strategy across all long portfolios, i.e., borrowing in the US money market and investing in foreign money markets. We refer to this zero-cost strategy as *DOL*. The second principal component is equivalent to a long position in Portfolio 1 (*investment currencies*) and a short position in Portfolio 5 (*funding currencies*), and corresponds to borrowing in the money markets of low yielding currencies and investing in the money markets of high yielding currencies. We refer to this long/short strategy as *CAR* in our tables – and we use both *DOL* and *CAR*

in risk-adjustment below.

3.8 *Momentum portfolios*

At the end of each period t , we form five portfolios based on exchange rate returns over the previous 3-months. We assign 20% of all currencies – those with the highest lagged exchange rate returns – to Portfolio 1, and 20% of all currencies – those with the lowest lagged exchange rate returns – to Portfolio 5. We then compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio. A strategy that is long in Portfolio 1 (winner currencies) and short in Portfolio 5 (loser currencies) is then denoted as *MOM*.¹⁵

3.9 *Value portfolios*

At the end of each period t , we form five portfolios based on the level of the real exchange rate.¹⁶ We assign 20% of all currencies – those with the lowest (highest) real exchange rates – to Portfolio 1 (Portfolio 5). We then compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio. A strategy that is long in Portfolio 1 (undervalued currencies) and short in Portfolio 5 (overvalued currencies) is then denoted as *VAL*.

3.10 *Risk reversal portfolios*

At the end of each period t , we form five portfolios based on out-of-the-money options. For each currency in each time period, we compute the risk reversal, which is the implied volatility

¹⁵Consistent with the results in Menkhoff, Sarno, Schmeling, and Schrimpf (2012b), sorting on lagged exchange rate returns or lagged currency excess returns to form momentum portfolios makes no qualitative difference to our results below. The same is true if we sort on returns with other formation periods in the range from 1 to 12 months.

¹⁶We compute the real exchange rate at the end of each month as $RER_t = S_t/PPP_t$, where S_t is the nominal exchange rate and PPP_t is the purchasing power parity rate computed using country CPI's.

of the 10-delta call less the implied volatility of the 10-delta put. We then assign 20% of all currencies with the lowest (highest) risk reversal to Portfolio 1 (Portfolio 5). We then compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio. A strategy that is long in Portfolio 1 (*high-skewness currencies*) and short in Portfolio 5 (*low-skewness currencies*) is then denoted as *RR*.

3.11 Volatility risk premia portfolios

At the end of each period t , we group currencies into five portfolios using the 1-year volatility risk premium constructed as described earlier. We allocate 20% of all currencies with the highest expected volatility premia, i.e., those which are cheapest to insure, to Portfolio 1, and 20% of all currencies with the lowest expected volatility premia, i.e., those which are expensive to insure, to Portfolio 5. We then compute the average excess return within each portfolio, and finally calculate the portfolio return from a strategy that is long in Portfolio 1 (*cheap volatility insurance*) and short in Portfolio 5 (*expensive volatility insurance*), and denote it *VRP*.

4 The *VRP* strategy: empirical evidence

4.1 Summary statistics and the returns to *VRP*

Table 1 presents summary statistics for the annualized average realized volatility $RV_{t,\tau}$, synthetic volatility swap rate $SW_{t,\tau} = E_t^{\mathbb{Q}}[RV_{t,\tau}]$, and volatility risk premium $VRP_{t,\tau} = RV_{t,\tau} - SW_{t,\tau}$ for the 1-year maturity ($\tau = 1$); in what follows, we drop the τ subscript, as it is always 1 year.

The table shows that, on average across developed currencies, RV_t equals 10.90%, with a standard deviation of 2.65%, and SW_t equals 11.68%, with a standard deviation of 2.71%. The average volatility risk premium VRP_t across these currencies, which is the difference of

these two variables, is equal to -0.78% , with a standard deviation of 1.64% . For the full sample of developed and emerging countries, RV_t and SW_t are slightly larger than for the sample of only developed currencies, as is the volatility risk premium, VRP_t , which equals -1.15 on average. We might expect to see this as the average price that hedgers have to pay to satisfy their demand for volatility insurance is larger when including emerging market currencies.¹⁷

Table 2 describes the returns (net of transactions costs) generated by our short expensive-to-insure, long cheap-to-insure currency strategy, reporting summary statistics for the five portfolios that are obtained when sorting on the volatility risk premium. In this table, P_L is the long portfolio that buys the top 20% of all currencies – those with the cheapest volatility insurance, P_2 buys the next 20% of all currencies, and so on until the fifth portfolio, P_S which is the portfolio that buys 20% of all currencies – those which are the most expensive to insure. VRP essentially buys P_L and sells P_S , with equal weights, so that $VRP = P_L - P_S$.

Table 2 reveals several facts about VRP . First, there is a general tendency of portfolio returns to decrease as we move from P_L towards P_S , although the decrease is not monotonic. The VRP average return is 4.95 (4.16) for the sample of Developed (Developed and Emerging) countries, and is statistically significantly different from zero (at the 5% level or better) for both excess returns and the FX return component. Second, the VRP return stems mainly from the long portfolio, P_L . Third, the return from P_L can almost completely be attributed to spot rate changes. Finally, the bottom panel of Table 2 shows the transition matrix between portfolios. This shows that there is currency rotation across quintile portfolios such that the steady-state transition probabilities are identical. Thus the performance of the strategy cannot simply be attributed to long-lived positions in particular currencies, a point we analyze in greater detail later in the paper.

¹⁷Table A.1 in the Internet Appendix reports summary statistics on the volatility risk premium for each currency.

The returns to VRP are very robust, based on a number of checks. First, we compute volatility risk premia using simple at-the-money implied volatility rather than the more complicated model-free implied volatility. We also implement the simple variance swap formula of Martin (2012), which allows for jumps. In both cases, results are virtually identical for developed countries, and improve for developed and emerging countries. We report these results in Internet Appendix Table A.10. Second, in our empirical work we also experiment with an AR(1) process for RV to form expectations of RV rather than simply using lagged RV over the previous 12 months to form these expectations. Again, we find that the results are virtually identical to those reported in Table 2. Third, in Internet Appendix Table A.7 we check whether a simple strategy based on sorting currencies by the difference between longer-term and short-term realized volatility effectively captures the returns from VRP . Using definitions of long-term ranging from six to 24 months and short-term from one to six months, we find that while there are a few high-return portfolios in the set, there is substantial variation in these returns across portfolios, leading to concerns of potential data-mining. Perhaps more importantly, these returns have low correlations with the returns of the VRP strategy, suggesting that implied volatility information from the options market is critical to the construction of the VRP strategy. Fourth, we show in Internet Appendix Table A.4 that the identities of the currencies most often found in the corner VRP portfolios are not easily recognizable from other currency strategies such as carry.

In the next section, we formalize this fourth exercise by explicitly comparing the returns of VRP to the conventional set of currency strategies considered in the literature thus far.

4.2 *Comparing VRP with other currency strategies*

In Table 3, we present the net returns to a number of long-short currency strategies computed using only time $t - 1$ information, to compare the predictability generated by strategies pre-

viously proposed in the literature with the new *VRP* strategy that we propose. We compare *CAR*, *MOM*, *VAL*, and *RR* with our *VRP* strategy. We report results for both subsamples (Developed, and Developed and Emerging) in our data.

Panel A of the table shows the results for the excess returns generated by these trading strategies. Consistent with a vast empirical literature (e.g., Lustig, Roussanov, and Verdelhan, 2011, Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011, and Menkhoff, Sarno, Schmeling, and Schrimpf, 2012a), *CAR* delivers a sizable average excess return, especially for the broader sample of countries analyzed. The Sharpe ratio of the carry trade is 0.38 for the sample of developed countries, and 0.53 for the full sample. *MOM* generates only small, yet positive, net excess returns, which is consistent with the recent evidence in Menkhoff, Sarno, Schmeling, and Schrimpf (2012b) that the performance of currency momentum has weakened substantially during the last decade. Both *VAL* and *RR* do quite well, with Sharpe ratios between 0.28 and 0.41 for *VAL*, and 0.41 and 0.42 for *RR*. However, with the exception of *CAR* for the Developed and Emerging sample and *RR* for the Developed sample, none of these common currency strategies generates average returns that are statistically significantly different from zero during the (admittedly short 16 year) period that we analyze, which includes the recent global financial crisis.

In contrast, the *VRP* strategy that we discover generates a Sharpe ratio of 0.61 and 0.51 for the two samples of countries considered, signifying that it outperforms all strategies for the Developed sample and is only slightly inferior to carry for the Developed and Emerging sample of countries. It is important to note that, for both samples, the *VRP* returns are clearly statistically significantly different from zero. Interestingly, the *VRP* strategy works better for the developed countries in our sample than for the whole sample of developed and emerging countries. One plausible explanation for this is that there is a greater prevalence of hedging using more sophisticated instruments such as currency options in developed markets

than in emerging markets.

While Panel A of the table suggests that the returns to the *VRP* strategy are comparable to or better than those of the other strategies that we provide as comparison, Panel B of the table introduces an important feature of the *VRP* strategy, namely that the major portion of these returns accrue as a result of spot rate predictability. This predictability is much larger than that associated with the competitor strategies over the sample period, generating an annualized mean spot exchange rate return of 5.45% for the developed countries, and 5.27% for the full cross-section of developed and emerging countries in our sample. In contrast, the exchange rate return from *CAR* is negative for both samples, and while other strategies have relatively better performance in predicting movements in the spot rate than *CAR*, the degree of predictability associated with any of these alternative strategies is also substantially smaller than *VRP*.

Several of the other moments presented in Panel B of Table 3 are also worth highlighting. First, the returns from *VRP* display desirable skewness properties, as its unconditional skewness is close to zero, and the maximum drawdown is far better (i.e., smaller in absolute size) than that of *CAR*. Finally, the table shows that the portfolio turnover of the *VRP* strategy (measured in terms of changes in the composition of the short and long legs of the *VRP* strategy, $Freq_S$ and $Freq_L$ in Table 3) is reasonable – lying in between the very low turnover of *CAR* and the high turnover of *MOM*.¹⁸

4.3 Combining *VRP* with other currency strategies

Panel C of Table 3 documents the correlation of the *VRP* strategy with the other strategies, and finds that the strategy tends to be mildly negatively correlated with *CAR* (with correlations of -0.08 and -0.06 for the two samples) and mildly positively correlated with *MOM* (with

¹⁸Table A.2 in the Internet Appendix reports the same information as Table 3 for gross, rather than net returns.

correlations of 0.11 and 0.15 for the two samples). The correlation with *VAL* for Developed countries is higher, but at 0.19 there is substantial orthogonal information in the strategy – indeed several of the other strategies are substantially more correlated with one another. Apart from showing that the strategy is distinct from those already studied in the literature, this also implies that combining *VRP* with *CAR*, *MOM*, *VAL*, and *RR* should yield sizable diversification benefits to an investor.¹⁹

Table 4 shows the subsample performance of the currency component of these strategies. Despite the inevitable attenuation of the sample period and the attendant difficulty of establishing statistical significance for each subperiod, the performance of *VRP* does seem substantially higher during crisis and NBER recession periods. However even outside of these recession periods, the return to *VRP* is still large and positive, and higher than that of all the competitor strategies. Even if *VRP* were to be used primarily as a hedge for a canonical currency strategy, it seems to exhibit desirable properties, delivering positive returns outside of crisis periods, and very high returns within crisis periods.²⁰

Taken together, the results from this section suggest that the *VRP* strategy has creditable excess returns overall, an important tendency to deliver returns during crisis periods that are far higher than the crashes commonly experienced with the carry trade, and far stronger predictive power for exchange rate returns, which is a unique feature in the space of alternative currency trading strategies. The importance of these features of the *VRP* strategy is twofold. First, a currency investor would likely gain substantial diversification benefits from adding

¹⁹It is also useful to note that the correlations for the excess returns from the strategies, presented in the table, are very close in magnitude to the correlations acquired from the exchange rate component of these returns – in other words, it is the currency component of the returns to this strategy that is the proximate source of the diversification benefits.

²⁰Figures A.1 and A.2 in the Internet Appendix plot the cumulative wealth generated by the strategies over the sample period, decomposing it into its two constituents: the exchange rate component (FX) and the interest rate differential component (yield). *VRP* returns are clearly distinct in that they are made up of a mildly negative yield component (for both samples of countries considered), and therefore the component due to spot return predictability is in fact larger than the full portfolio return, achieving Sharpe ratios above 0.50 in both samples of countries.

VRP to a currency portfolio to enhance risk-adjusted returns. Second, a spot currency trader interested in forecasting exchange rate fluctuations (as opposed to currency excess returns) might value the signals provided by *VRP*.

To better understand the value of the *VRP* strategy for a currency investor, we compute the optimal currency portfolio for an investor who uses all of the five strategies considered here: *CAR*, *VAL*, *RR*, *MOM*, and *VRP*. Specifically, consider a portfolio of N assets with covariance matrix Σ . The global minimum volatility portfolio is the portfolio with the lowest return volatility, and represents the solution to the following optimization problem: $\min w' \Sigma w$ subject to the constraint that the weights sum to unity $w' \iota = 1$, where w is the $N \times 1$ vector of portfolio weights on the risky assets, ι is a $N \times 1$ vector of ones, and Σ is the $N \times N$ covariance matrix of the asset returns. The weights of the global minimum volatility portfolios are given by $w = \frac{\Sigma^{-1} \iota}{\iota' \Sigma^{-1} \iota}$. We compute the optimal weights for both the Developed, and Developed and Emerging countries, and report the results graphically in Figure 1.

The results show that the optimal weight assigned to the *VRP* strategy is high, equal to 26% and 28% for the two sets of countries. The Sharpe ratio of the minimum volatility portfolio for the Developed sample, for instance, is quite impressive, at 0.69. However, this number drops to 0.60 if the investor is not given access to the *VRP* strategy, and only employs the other four currency strategies. Similarly for the Developed & Emerging sample, the Sharpe ratio equals 0.60 when the *VRP* strategy is included and drops to 0.50 when it is excluded from the menu of currency strategies. These findings confirm the value of *VRP* in a currency portfolio and its desirable correlation properties. However, it is important to take these particular results with caution given the short sample at our disposal for the estimation of the moments required in the optimization and the well-known sensitivity of mean-variance analysis to estimation errors.

Before turning to studying possible explanations of the performance of *VRP*, we check

whether such predictive power is purely cross-sectional. Specifically, one may be concerned whether – given the relatively short sample period of 16 years – the predictability recorded here stems from long-lived cross-sectional differences in the volatility risk premium, which happen to be related to cross-sectional differences in excess returns. To check whether this is the case, we construct a static VRP currency strategy, which we denote $\overline{VRP}$, which buys (sells) the currencies with the highest (lowest) average volatility risk premia over the sample period. This strategy requires no portfolio rebalancing, and its performance is informative of the extent to which the returns to the VRP strategy are due to unconditional differences in the volatility risk premium between currencies in the cross-section. However, this strategy does contain a look-ahead bias, since it assumes that an investor knows the unconditional mean of the VRP for each currency rather than having to learn it over time. As a result, the returns we compute here provide an upper bound of what a static strategy could achieve. We also compute analogous returns $\overline{CAR}$, $\overline{MOM}$, $\overline{VAL}$, and $\overline{RR}$. These returns can be thought of as the “static component” in the return decomposition proposed by Hassan and Mano (2013), which is designed to measure the relative importance of cross-sectional versus time-series predictability in FX strategies.

Table 5 shows the returns of these static strategies gross of transaction costs. Panel A presents the overall excess return and suggests that $\overline{VRP}$ performs well, with an average return of 3.51 (3.28) per annum for Developed (Developed and Emerging) Countries. However, Panel B of Table 5 shows that $\overline{VRP}$ returns are virtually entirely due to cross-sectional differences in the average interest rate differential, as there is basically no predictability in FX returns – this establishes that $\overline{VRP}$ is a distinct strategy from VRP , which derives virtually all of its performance from FX returns. Moreover, we cannot establish the statistical significance of $\overline{VRP}$ returns at conventional significance levels. Finally, Panel C of Table 5 shows that these static returns are highly correlated with one another, with $\overline{VRP}$ in particular displaying a

correlation of 0.73 with carry.²¹ Taken together, this table shows that time-series variation in currency volatility risk premia is important to explain the performance of *VRP*.

5 Understanding *VRP* returns

The empirical results reported earlier suggest that the currency volatility risk premium contains powerful predictive information for currency returns that is markedly different from the information contained in several common predictors studied in the literature. While the main contribution of our paper is empirical and we do not have a formal theoretical model that links the volatility risk premium (or its determinants) to spot currency returns, we examine possible mechanisms that may drive our results.

5.1 Risk premia

First, we consider the possibility that returns from the *VRP* strategy reflect compensation for risk. We begin by testing the pricing power of conventional risk factors for *VRP* returns, using standard linear asset pricing models, in both the cross-section and the time-series.

5.1.1 Time series tests

As a first step, Table 6 simply regresses the time-series of *VRP* returns on a number of risk factors proposed in the literature. First, Panel A confirms the results found in Tables 2 and 3, by using *DOL*, *CAR*, *MOM*, *VAL*, and *RR* as right-hand side variables, and shows that for both Developed and Developed and Emerging samples, there is substantial and statistically significant alpha relative to these factors. Panel B of the table uses the three Fama-French factors and adds equity market momentum, denoted *MOME*. Again, *VRP* has alpha relative to these factors which is very close to that in the prior panel. Finally,

²¹In Table A.3 in the the Internet Appendix, we examine the static, dynamic and dollar component of the *VRP* returns in a similar vein to Hassan and Mano (2014).

Panel C of Table 6 employs the Fung-Hsieh (2004) factor model, which has been used in numerous previous studies; see for example, Bollen and Whaley (2009), Ramadorai (2013), and Patton and Ramadorai (2013). The set of factors comprises the excess return on the S&P 500 index; a small minus big factor constructed as the difference between the Wilshire small and large capitalization stock indexes; excess returns on portfolios of lookback straddle options on currencies, commodities, and bonds, which are constructed to replicate the maximum possible return to trend-following strategies on their respective underlying assets; the yield spread of the US 10-year Treasury bond over the 3-month T-bill, adjusted for the duration of the 10-year bond; and the change in the credit spread of Moody's BAA bond over the 10-year Treasury bond, also appropriately adjusted for duration. Yet again, the table shows that the alpha of VRP is unaffected by the inclusion of these factors.

5.1.2 Cross-sectional tests

Our cross-sectional tests rely on a standard stochastic discount factor (SDF) approach (Cochrane, 2005), and we focus on a set of risk factors in our investigation that are motivated by the existing asset pricing literature on the returns to currency strategies. We begin by briefly reviewing the methods employed, and denote excess returns of portfolio i in period t by RX_t^i .

The usual no-arbitrage relation should apply meaning that risk-adjusted currency excess returns should have a zero price in expectation, satisfying the basic Euler equation:

$$\mathbb{E}[M_t RX_t^i] = 0, \tag{6}$$

with a linear SDF $M_t = 1 - b'(f_t - \mu)$, where f_t denotes a vector of risk factors, b is the vector of SDF parameters, and μ denotes factor means.

This specification implies a beta pricing model in which expected excess returns depend on factor risk prices λ , and risk quantities β_i , which are the regression betas of portfolio excess

returns on the risk factors for each portfolio i (see e.g., Cochrane, 2005):

$$\mathbb{E} [RX^i] = \lambda' \beta_i \quad (7)$$

The relationship between the factor risk prices in equation (7) and the SDF parameters in equation (6) is simply given by $\lambda = \Sigma_f b$, where Σ_f is the covariance matrix of the risk factors. Thus, factor risk prices can be easily obtained via the SDF approach, which we implement by estimating the parameters of equation (6) via the generalized method of moments (GMM) of Hansen (1982).²² We also present results from the more traditional two-stage procedure of Fama and MacBeth (1973) in our empirical implementation.

In our asset pricing tests we consider a two-factor linear model that comprises DOL and one additional risk factor, which is one of CAR and VOL_{FX} . DOL denotes the average return from borrowing in the US money market and equally investing in foreign money markets. CAR is the carry portfolio described earlier. VOL_{FX} is a global FX volatility risk factor constructed as the innovations to global FX volatility, i.e., the residuals from an autoregressive model applied to the average realized volatility of all currencies in our sample, as in Menkhoff, Sarno, Schmeling, and Schrimpf (2012a). In Internet Appendix Table A.8, we also consider innovations to global average percentage bid-ask spreads in the spot market (BAS_{FX}) and the option market (BAS_{IV}), which can be seen as global proxies for the FX spot market and the FX option market illiquidity, respectively.

In assessing our results, we are aware of the statistical problems plaguing standard asset pricing tests, recently emphasized by Lewellen, Nagel, and Shanken (2010). Asset pricing tests can often be highly misleading, in the sense that they can indicate strong but illusory explanatory power through high cross-sectional R^2 statistics, and small pricing errors, when in fact a risk factor has weak or no pricing power. Given the relatively small cross-section of

²²Estimation is based on a pre-specified weighting matrix and we focus on unconditional moments (i.e., we do not use instruments other than a constant vector of ones) since our interest lies in the performance of the model to explain the cross-section of expected currency excess returns (see Cochrane, 2005; Burnside, 2011).

currencies in our data, as well as the relatively short time span of our sample, these problems can be severe in our tests. As a result, when interpreting our results, we only consider the cross-sectional R^2 and Hansen-Jagannathan (HJ) tests on the pricing errors if we can confidently detect a statistically significant risk factor, i.e., if the estimates clearly point to a statistically significant market price of risk λ on a factor.

Table 7 reports GMM estimates of b , portfolio-specific β 's, and implied λ 's, as well as cross-sectional R^2 statistics and the HJ distance measure (Hansen and Jagannathan, 1997). In the table, standard errors are constructed as in Newey and West (1987) with optimal lag length selection according to Andrews (1991). Besides the GMM tests, we employ traditional Fama-MacBeth (FMB) two-pass OLS regressions (with Shanken (1992) corrected standard errors) to estimate portfolio betas and factor risk prices. Note that we do not include a constant in the second stage of the FMB regressions. Since DOL has virtually no cross-sectional relation to portfolio returns, it serves the same purpose as a constant that allows for common mispricing.

Panels A and B of Table 7 show clearly how none of the risk factors considered enters the SDF with a statistically significant risk price λ , and that this is the case for both the developed countries and the full sample. As expected, the FMB results in the table are qualitatively, and in most cases also quantitatively identical to the one-step GMM results. The bottom part of the panels show that there is little cross-sectional variation across the 5 portfolios sorted by the cost of currency insurance, which is what we confirm more formally in the asset pricing tests. While the HJ test delivers large p -values for the null of zero pricing errors in all cases, we attach no information to this result given the lack of clear statistical significance of the market price of risk.

5.1.3 *Aversion to volatility risk*

Next, we investigate the possibility that the currency-specific volatility risk premium captures fluctuations in aversion to volatility risk – i.e., a time-varying factor loading on the global volatility risk factor. We have already ascertained that a simple strategy allowing for static loadings on the Menkhoff, Sarno, Schmeling, and Schrimpf (2012a) strategy fails to explain the cross-section of *VRP* portfolio returns, but these tests do not account for the possibility that different currencies load differently on a global volatility shock at different points in time. There is also the possibility that market segmentation causes expected returns on different currencies to be determined independently – but this (remote) possibility is very difficult to evaluate, and if our strategy did indeed provide evidence of this, it would have far-reaching consequences.

To evaluate whether *VRP* returns can be explained by currencies exhibiting time-varying loadings on a global volatility shock, we estimate the loadings of currency returns on various proxies for global volatility risk, and build portfolios sorted on these estimated loadings. Specifically, we estimate the following rolling regression for each currency i :

$$RX_t^i = \alpha_i + \beta_i GVOL_t + \varepsilon_{it},$$

Here *GVOL* is a proxy for global volatility risk premia and we employ various measures, including the average volatility risk premium across our currencies (with equal weights); the first principal component of the currencies' volatility risk premia; and the equity volatility risk premium computed as the difference between the time- t one-month realized volatility on the S&P500 and the VIX index.

We estimate these regressions using rolling windows of 36 months. After obtaining estimates of the β_i coefficients, we sort currencies into five portfolios on the basis of these β_i estimates. Finally, we construct a long-short strategy which buys currencies with low betas

and sells currencies with high betas. In essence, this strategy exploits differences in exposure of individual currencies to global measures of volatility risk premia, which is a direct test of the above hypothesis.

The results using our three measures for *GVOL* are qualitatively identical and we report in Table 8 the results for *GVOL* set equal to the average volatility risk-premium across the currencies in our sample. Internet Appendix Tables A.5 and A.6 contain results for the other two measures. The table shows that the performance of this strategy is strictly inferior to the performance of the *VRP* strategy (in fact producing negative returns), and the correlation between the returns from the two strategies is close to zero. On the basis of this evidence, we conclude that there is no support for *VRP* returns being driven by aversion to global volatility risk in the data. Overall, the asset pricing tests reveal that it is not possible to understand the returns from the *VRP* strategy as compensation for global risk. Therefore, we turn to examining different explanations.

5.2 *Limits to arbitrage*

The second possible explanation that we consider is limits to arbitrage, in the spirit of Acharya, Lochstoer, and Ramadorai (2013). According to this explanation, the returns to *VRP* arise from the interaction between natural hedgers of FX risk, and currency market speculators. When the risk-bearing capacity of currency-market speculators is affected by shocks to the availability of arbitrage capital, this will make currency options across the board more expensive, with particular impacts on those currencies to which speculators have high exposure – for example, currency hedge funds may reduce their outstanding short put option positions in the currencies in which they trade (shorting put options is a favoured strategy of many hedge funds; see Fung and Hsieh, 1997, and Agarwal and Naik, 2004).

This will result in selling pressure on expensive-to-insure currencies as natural hedgers

such as corporations sell pre-existing currency holdings, abandon expensive currency hedges, and become more reluctant to denominate contracts in these currencies. Conversely, this mechanism results in relatively less pressure on cheap-to-insure currencies, for which natural hedgers are happy to hold higher inventories. This yields the positive long-short returns in the *VRP* portfolio. When capital constraints loosen, we should see the opposite behavior, i.e., a reversal in both the volatility risk premium and the spot currency position.

This explanation has several testable implications. First, for this mechanism to work demand pressure in the option market must have an impact on option prices, as demonstrated by Garleanu, Pedersen and Poteshman (2009) for stock options. Therefore, as a preliminary test, we run a similar regression to Garleanu, Pedersen, and Poteshman for FX markets, in an attempt to ascertain whether demand pressure in the FX derivatives used for hedging FX risk generates price impact which affects the volatility risk premium.

We estimate a panel regression (with fixed effects) of the volatility risk premium on a proxy for demand pressure in FX derivatives markets:

$$VRP_t^i = \alpha_i + \beta NDem_{t-lag}^i + u_{it}, \quad (8)$$

where VRP_t^i is the 1-year volatility risk premium for currency i (i.e., the difference between the realized volatility, RV_t and the synthetic volatility swap rate, SW_t),²³ and $NDem_t^i$ denotes the net demand of currency options and futures for end-users from the US Commodity Futures Trading Commission (CFTC). The net demand proxy is constructed as the difference between long and short positions scaled by the total open interest, and is available for two groups of end-users: commercial and financial.

For the left-hand side variable in these regressions, we employ several definitions of the volatility risk premium: the definition used in our core analysis, where RV is calculated using

²³Note that this is distinct from *VRP*, where the italics denote the returns to the trading strategy conditional on realizations of *VRP*, the level of the volatility risk premium.

daily exchange rate returns over the previous year and SW is computed as in Britten-Jones and Neuberger (2000) using 1-year currency option implied volatilities; in VRP_{si} , SW is computed using the simple variance swap method of Martin (2012); in VRP_{garch} , RV is the 1-year volatility forecast generated from the simple GARCH(1,1) applied to daily exchange rate returns; in VRP_{sv} , RV is the 1-year volatility forecast generated from a stochastic volatility model for daily exchange rate returns (e.g., Della Corte, Sarno, and Tsiakas, 2009; Sarno, Schneider and Wagner, 2012). Monthly CFTC data are collected on the last Tuesday of every month. All other variables are measured on the same day.

The regression results, reported in Table 9 for each of the two end-user groups, suggest that in a contemporaneous regression ($lag = 0$) the net demand proxy for commercial end-users always enters with a negative coefficient that is statistically significantly different from zero, regardless of the definition of the VRP on the left-hand-side. This is essentially the analogue of the result of Garleanu, Pedersen and Poteshman for the case of FX markets, and it implies that net demand for hedging in FX markets increases the cost of volatility insurance. It is also noticeable that this price impact is quite persistent in that the net demand proxy enters significantly also in a predictive regression ($lag = 1$ month). In contrast, the coefficient on financial end-users is positive and, in two regressions statistically significantly different from zero. Again this is consistent with the story of Garleanu, Pedersen and Poteshman, since financial users are providing volatility insurance to commercial customers, essentially acting as market makers.

Table 10 turns from the options market to currency excess returns, testing whether time-series variation in limits to arbitrage proxies predicts variation in VRP returns. The table shows results from predicting the exchange rate component of VRP ; the results for excess returns are, not surprisingly, qualitatively identical and quantitatively very similar. The first column in both panels shows the univariate regression of the exchange rate component of VRP

regressed on the lagged 12-month rolling average of the TED spread. The coefficient on this variable is positive and statistically significant for both sample of countries examined, which is consistent with the limits to arbitrage explanation – when funding liquidity is lower (i.e., times of high capital constraints on speculators), we find that the expected return from *VRP* increases. The second column shows that when the 12-month rolling average of changes in VIX (a proxy for increases in the risk aversion of market participants, yielding both greater limits to arbitrage and an increased desire to hedge) is positive, *VRP* returns increase (significantly for the full sample of countries), again consistent with the limits to arbitrage explanation.²⁴ Similarly, the third column shows that a general financial distress indicator (FSI, constructed by the Federal Reserve Bank of St. Louis) that captures the principal component of a variety of liquidity and volatility indicators is positive and, for the full sample of countries, statistically significant. The fourth column of the table interacts TED with changes in VIX, and finds strong statistically significant predictive power of this interaction for the FX returns on our strategy in both samples of countries, suggesting that when funding liquidity is constrained *and* risk aversion is high, *VRP* returns increase. The final column of the table adds in measures of capital flows into hedge funds. When aggregate capital flows into hedge funds are high, signifying that they experience fewer constraints on their ability to engage in arbitrage transactions, we find that returns for the *VRP* strategy are lower and vice versa, although the variable is only significant for the sample of developed countries.

The final five rows of Table 10 introduce several of the variables described above simultaneously to test their joint and separate explanatory power. We generally include TED, changes in VIX and the interaction separately to avoid potential collinearity in the regressions as these variables are highly correlated with one another – since they capture aggregate variation in funding liquidity and risk aversion, it is obvious that they contain a substantial

²⁴This is similar to the results in Nagel (2012), who shows that a strategy of liquidity provision in equity markets has returns which are highly correlated with VIX.

common component. Nonetheless, we find that all these variables retain their signs and are often statistically significant in these multivariate predictive regressions, offering some support to the limits to arbitrage explanation of our results. Table A.9 in the Internet Appendix reports results for the same regressions using raw measures of VIX, TED and FSI rather than rolling averages, and shows that the results are qualitatively identical.

Finally, we examine whether the observed buying and selling actions of different players in the currency market follow the pattern implied by the limits to arbitrage explanation, i.e., that currencies in the high volatility-insurance portfolio are sold and those in the low volatility insurance portfolio are bought by natural hedgers, with speculators taking the opposite position. We do so using the CFTC data on the position of commercial and financial traders in FX markets, essentially taking the currencies ranked by their volatility insurance costs, and documenting the traders' positions (cumulative net positions), rather than returns.²⁵ We view the CFTC position data as a proxy for cumulative order flow across different segments of FX market participants, given that there is evidence that the CFTC position data and currency order flow capture very similar information (e.g., Klitgaard and Weir, 2004).

The results of this exercise are reported in Figure 2, which plots the cumulative position in the currencies in the *VRP* portfolio for financial and commercial traders. We find that the position of commercial traders follows the pattern implied by the limits to arbitrage explanation – such traders sell expensive-to-insure currencies and buy cheaper-to-insure currencies. Financial traders display the opposite behavior, with a strongly negative position in the *VRP* portfolio, which is consistent with their acting as market-makers, providing liquidity to satisfy the buying (selling) demand for low (high)-insurance currencies.

Taken together, the results in this section lend support to a limits to arbitrage explanation

²⁵To allow for meaningful cross-currency comparisons, we need to ensure that net positions are comparable across currencies, as their absolute size differs across currencies. We therefore divide net positions by their standard deviation computed over a rolling window of 3 months.

for the spot predictability associated with VRP . While these findings are suggestive, they must be viewed in light of the fact that we do not provide a formal theoretical model to support this explanation, and there may of course be alternative explanations.

6 Conclusions

We show that the currency volatility risk premium has substantial predictive power for the cross-section of currency returns. Currencies with low implied volatility relative to historical realized volatility – those with relatively cheap volatility insurance – predictably appreciate, while currencies with relatively more expensive volatility insurance predictably depreciate. This predictive power is specifically related to future variation in spot exchange rate returns, and not to interest rate differentials. A portfolio of currencies (which we dub VRP) constructed by going long cheap volatility insurance currencies and short expensive volatility insurance currencies generates economically and statistically significant returns, which are largely uncorrelated with four widely-studied currency strategies.

While we do not have a formal theoretical model, we do provide empirical evidence pertaining to possible explanations for the performance of the strategy. We find that a comprehensive set of standard risk factors is unable to explain VRP returns, suggesting that these returns are not generated on account of compensation for systematic risk. We find some evidence in support of an explanation in which time-variation in limits to arbitrage causes volatility insurance costs to fluctuate across time and currencies, with consequences for the spot market as risk-averse currency hedgers become reluctant to take or hold positions in expensive-to-insure currencies.

Overall, the results in our paper provide new insights into the predictability of exchange rate returns, an area in which evidence has been difficult to obtain. We also introduce a new currency strategy with useful diversification properties into the rapidly-expanding research

on this topic. While our empirical results point to new, powerful predictive information for currency returns, our attempts to explain the drivers of this predictive power are limited by the absence of a formal theoretical model that links volatility risk premia and underlying asset returns. The development of such theory is an important avenue for future research.

References

- Acharya, V., Johnson, T., 2007. Insider trading in credit derivatives. *Journal of Financial Economics* 84, 110-141.
- Acharya, V., Lochstoer, L.A., Ramadorai, T., 2013. Limits to arbitrage and hedging: evidence from commodity markets. *Journal of Financial Economics* 109, 441-465.
- Adrian, T., Etula, E., Muir, T., 2013. Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, forthcoming.
- Agarwal, V., Naik, N.Y., 2004. Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17, 63–98.
- Akram, Q.F., Rime, R., Sarno, L., 2008. Arbitrage in the foreign exchange market: turning on the microscope. *Journal of International Economics* 76, 237–253.
- Andrews, D.W.K., 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59, 817–858.
- Ang, A., Chen, J., 2010. Yield curve predictors of foreign exchange returns. Unpublished working paper, Columbia Business School.
- Bali, T.G., Hovakimian, A., 2009. Volatility spreads and expected stock returns. *Management Science* 55, 1797–1812.
- Bank for International Settlements, 2013. Triennial central bank survey of foreign exchange and derivatives market in 2013. Bank for International Settlements Press, Basel.
- Barroso, P., Santa-Clara, P., 2014. Beyond the carry trade: optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, forthcoming.

- Bates, D.S., 1991. The crash of '87: was it expected? The evidence from options markets. *Journal of Finance* 46, 1009–1044.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- Bollen, N.P., Whaley, R.E., 2009. Hedge fund risk dynamics: implications for performance appraisal. *Journal of Finance* 64, 987–1037.
- Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463–4492.
- Breeden, D., and Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839–866.
- Buraschi, A., Trojani, F., Vedolin, A., 2014. When uncertainty blows in the orchard: comovement and equilibrium volatility risk premia. *Journal of Finance* 69, 101–137.
- Burnside, C., 2011. The cross section of foreign currency risk premia and consumption growth risk: comment. *American Economic Review* 101, 3456–76.
- Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2011. Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24, 853–891.
- Buss, A., Vilkov, G., 2012. Measuring equity risk with option-implied correlations. *Review of Financial Studies* 25, 3113–3140.
- Campa, J.M., Chang, P.H.K., Reider, R.L., 1998. Implied exchange rate distributions: evidence from OTC option markets. *Journal of International Money and Finance* 17, 117–160.

- Carr, P., Lee, R., 2009. Volatility derivatives. *Annual Review of Financial Economics* 1, 1–21.
- Carr, P., Wu, L., 2009. Variance risk premiums. *Review of Financial Studies* 22, 1311–1341.
- Castagna, A., Mercurio, F., 2007. The vanna-volga method for implied volatilities. *Risk* January, 106–111.
- Chinn, M.D., Ito, H., 2006. What matters for financial development? Capital controls, institutions, and interactions. *Journal of Development Economics* 81, 163–192.
- Cochrane, J.H., 2005. *Asset Pricing*. Princeton University Press, Princeton.
- Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. *Review of Financial Studies* 22, 3491–3530.
- Della Corte, P., Sarno, L., Tsiakas, I., 2011. Spot and forward volatility in foreign exchange. *Journal of Financial Economics* 100, 496–513.
- Demeterfi, K., Derman, E., Kamal, M., Zou, J., 1999. A guide to volatility and variance swaps. *Journal of Derivatives* 6, 9–32.
- Drechsler, I., Yaron, A., 2011. What’s vol got to do with it. *Review of Financial Studies* 24, 1–45.
- Easley, D., O’Hara, M., Srinivas, P.S., 1998. Option volume and stock prices: evidence on where informed traders trade. *Journal of Finance*, 53, 431–465.
- Engel, C., Mark, N.C., West, K.D., 2008. Exchange rate models are not as bad as you think. *NBER Macroeconomics Annual* 2007, 381–441.
- Eraker, B., 2010. The volatility premium. Unpublished working paper, Wisconsin School of Business.

- Evans, M.D.D., Lyons, R.K., 2005. Meese-Rogoff redux: micro-based exchange rate forecasting. *American Economic Review Papers and Proceedings* 95, 405–414.
- Fama, E.F., MacBeth, J.D., 1973. Risk, return, and equilibrium: empirical tests. *Journal of Political Economy* 81, 607–36.
- Froot, K.A., Ramadorai, T., 2005. Currency returns, intrinsic value, and institutional investor flows. *Journal of Finance* 60, 1535–1566.
- Fung, W., Hsieh, D.A., 1997. Empirical characteristics of dynamic trading strategies: the case of hedge funds. *Review of Financial Studies* 10, 275–302.
- Fung, W., Hsieh, D.A., 2004. Hedge fund benchmarks: a risk based approach. *Financial Analyst Journal* 60, 65–80.
- Gârleanu, N.B., Pedersen, L.H., 2011. Margin-based asset pricing and the law of one price. *Review of Financial Studies* 24, 1980–2022.
- Gârleanu, N.B., Pedersen, L.H., Poteshman, A., 2009. Demand-based option pricing. *Review of Financial Studies* 22, 4259–4299.
- Garman, M.B., Kohlhagen, S.W., 1983. Foreign currency option values. *Journal of International Money and Finance* 2, 231–237.
- Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. *Journal of Financial Economics* 94, 310–326.
- Gromb, D., Vayanos, D., 2010. Limits of arbitrage: the state of the theory. *Annual Review of Financial Economics* 2, 251–275.
- Han B., Zhou, Y., 2011. Variance risk premium and cross-section of stock returns. Unpublished working paper, University of Texas at Austin.

- Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50, 1029–54.
- Hansen, L.P., Jagannathan, R., 1997. Assessing specification errors in stochastic discount factor model. *Journal of Finance* 52, 557–590.
- Hassan, T., Mano, R., 2014. Forward and spot exchange rates in a multi-currency world. Unpublished working paper, University of Chicago.
- International Swaps and Derivatives Association, Emerging Markets Traders Association, Foreign Exchange Committee, 1998. The 1998 foreign exchange and currency option definitions. International Swaps and Derivatives Association, New York.
- Jiang, G.J., Tian, Y.S., 2005. The model-free implied volatility and its information content. *Review of Financial Studies* 18, 1305–1342.
- Kaminsky, G., Schmukler, S., 2008. Short-run pain, long-run gain: financial liberalization and stock market cycles. *Review of Finance* 12, 253–292.
- Klitgaard, T., Weir, L., 2004. Exchange rate changes and net positions of speculators in the futures market. *Federal Reserve Bank of New York Economic Policy Review* 10, 17–28.
- Lewellen, J., Nagel, S., Shanken, J., 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96, 175–194.
- Londono, J.M., Zhou, J.D., 2012. Variance risk premiums and the forward premium puzzle. Unpublished working paper, Federal Reserve Board.
- Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. *Review of Financial Studies* 24, 3731–3777.

- Lustig, H., Verdelhan, A., 2007. The cross section of foreign currency risk premia and US consumption growth risk. *American Economic Review* 97, 89–117.
- Martin, I., 2012. Simple variance swaps. Unpublished working paper, Stanford University.
- Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? *Journal of International Economics* 14, 3–24.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012a. Carry trades and global FX volatility. *Journal of Finance* 67, 681–718.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. *Journal of Financial Economics* 106, 620–684.
- Mueller, P., Yen, Y.M., Vedolin, A., 2011. Bond variance risk premiums. Unpublished working paper, London School of Economics.
- Nagel, S., 2012. Evaporating liquidity. *Review of Financial Studies* 25, 2005–2039.
- Neuberger, A., 1994. The log contract: a new instrument to hedge volatility. *Journal of Portfolio Management* 20, 74–80.
- Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Patton, A., Ramadorai, T., 2013. On the high-frequency dynamics of hedge fund risk exposures. *Journal of Finance* 68, 597–635.
- Ramadorai, T., 2013. Capacity constraints, investor information, and hedge fund returns. *Journal of Financial Economics* 107, 401–416.
- Rime, D., Sarno, L., Sojli, E., 2010. Exchange rates, order flow and macroeconomic information. *Journal of International Economics* 80, 72–88.

Sarno, L., Schneider, P., Wagner, C., 2012. Properties of foreign exchange risk premia. *Journal of Financial Economics* 105, 279–310.

Todorov, V., 2010. Variance risk premium dynamics: the role of jumps. *Review of Financial Studies* 23, 345–383.

Table 1

Volatility risk premia.

This table presents summary statistics for the 1-year volatility risk premium (VRP_t) defined as difference between the realized volatility (RV_t) and the synthetic volatility swap rate (SW_t). RV_t is calculated using daily exchange rate returns over the previous year. SW_t is computed as in Britten-Jones and Neuberger (2000) using 1-year currency option implied volatilities. Q_j refers to the j^{th} percentile. AC indicates the 1-year autocorrelation coefficient. VRP_t , RV_t , and SW_t are expressed in percent per annum, and averaged across two sets of currencies. The sample period comprises daily data from January 1998 to December 2013.

	VRP_t	RV_t	SW_t	VRP_t	RV_t	SW_t
	<i>Developed</i>			<i>Developed and emerging</i>		
<i>Mean</i>	-0.78	10.90	11.68	-1.15	10.96	12.11
<i>Sdev</i>	1.64	2.65	2.71	1.90	2.96	3.25
<i>Skew</i>	0.25	2.07	1.32	-0.36	2.23	1.80
<i>Kurt</i>	5.48	7.48	4.78	6.45	7.97	6.54
Q_5	-3.50	8.35	8.50	-4.33	8.41	9.04
Q_{95}	1.56	18.08	16.85	1.44	19.00	18.53
<i>AC</i>	-0.07	0.25	0.47	-0.06	0.22	0.43

Table 2

Volatility risk premia portfolios.

This table presents descriptive statistics of currency portfolios sorted on the 1-year volatility risk premia at time $t - 1$. The volatility risk premium is defined as difference between the realized volatility and the synthetic volatility swap rate both computed at time $t - 1$. The long (short) portfolio P_L (P_S) contains the top 20% of all currencies with the highest (lowest) volatility risk premium. VRP denotes a long-short strategy that buys P_L and sells P_S . The table also reports the first-order autocorrelation coefficient (AC), the annualized Sharpe ratio (SR), and the frequency of portfolio switches ($Freq$). *Panel A* displays the currency excess return whereas *Panel B* reports the exchange rate return component. *Panel C* presents the transition probability from portfolio i to portfolio j between time t and time $t + 1$. $\bar{\pi}$ indicates the steady state probability. The superscripts *, **, and *** indicate statistical significance for the mean at 10%, 5%, and 1%, respectively, based on Newey and West (1987) and Andrews (1991). Returns are expressed in percentage per annum and adjusted for transaction costs. The sample period comprises monthly data from January 1998 to December 2013.

Panel A: Currency excess returns												
	P_L	P_2	P_3	P_4	P_S	VRP	P_L	P_2	P_3	P_4	P_S	VRP
	<i>Developed</i>						<i>Developed and emerging</i>					
<i>Mean</i>	4.59*	3.12	1.10	3.27	-0.36	4.95**	4.71*	2.71	1.05	2.40	0.55	4.16**
<i>Sdev</i>	9.57	9.55	9.62	10.15	10.04	8.15	10.16	8.93	9.09	10.62	8.61	8.14
<i>Skew</i>	-0.20	-0.05	-0.08	-0.22	-0.22	-0.03	-0.23	-0.44	-0.32	-0.54	-0.18	0.01
<i>Kurt</i>	3.56	5.16	5.52	3.95	3.96	3.97	3.94	5.83	3.57	4.76	4.81	4.54
<i>SR</i>	0.48	0.33	0.11	0.32	-0.04	0.61	0.46	0.30	0.12	0.23	0.06	0.51
<i>AC</i>	0.04	-0.01	0.05	0.14	-0.01	0.05	0.04	0.08	0.11	0.07	0.03	-0.02
<i>Freq</i>	0.29	0.48	0.54	0.52	0.35	0.35	0.28	0.43	0.50	0.49	0.27	0.27
Panel B: Exchange rate returns												
<i>Mean</i>	4.60*	2.87	1.21	3.00	-0.85	5.45***	4.45*	2.30	1.04	1.72	-0.82	5.27**
<i>Sdev</i>	9.58	9.52	9.53	10.08	10.00	8.12	10.17	8.90	9.00	10.53	8.61	8.20
<i>Skew</i>	-0.25	-0.10	-0.10	-0.23	-0.24	-0.03	-0.29	-0.49	-0.35	-0.59	-0.27	0.09
<i>Kurt</i>	3.57	5.29	5.56	4.12	3.95	4.04	3.99	5.84	3.69	5.02	4.93	5.10
<i>SR</i>	0.48	0.30	0.13	0.30	-0.08	0.67	0.44	0.26	0.12	0.16	-0.10	0.64
<i>AC</i>	0.03	-0.02	0.04	0.13	-0.01	0.04	0.03	0.08	0.09	0.06	0.04	-0.02
<i>Freq</i>	0.29	0.48	0.54	0.52	0.35	0.35	0.28	0.43	0.50	0.49	0.27	0.27
Panel C: Transition matrix												
P_L	0.71	0.21	0.06	0.01	0.01		0.72	0.22	0.05	0.01	0.00	
P_2	0.20	0.53	0.18	0.06	0.02		0.18	0.58	0.17	0.05	0.02	
P_3	0.05	0.20	0.46	0.20	0.08		0.02	0.21	0.51	0.22	0.04	
P_4	0.01	0.04	0.22	0.48	0.24		0.01	0.07	0.22	0.51	0.19	
P_S	0.00	0.03	0.07	0.25	0.65		0.00	0.02	0.04	0.21	0.72	
$\bar{\pi}$	0.19	0.20	0.20	0.20	0.20		0.17	0.24	0.21	0.20	0.18	

Table 3

Currency strategies.

This table presents descriptive statistics of currency strategies formed using time $t - 1$ information. *CAR* is the carry trade strategy that buys (sells) the top 20% of all currencies with the highest (lowest) interest rate differential relative to the US dollar. Similarly, *MOM* is the momentum strategy that buys (sells) currencies with the highest (lowest) past 3-month exchange rate return, *VAL* is the value strategy that buys (sells) currencies with lowest (highest) real exchange rate, *RR* is the risk reversal strategy that buys (sells) currencies with the lowest (highest) 1-year 10-delta risk reversal, and *VRP* is the volatility risk premium strategy that buys (sells) currencies with the highest (lowest) 1-year volatility risk premium. The table also reports first order autocorrelation coefficient (*AC*), the annualized Sharpe ratio (*SR*), the Sortino ratio (*SO*), the percentage maximum drawdown (*MDD*), the frequency of portfolio switches for the long (*Freq_L*) and the short (*Freq_S*) position. *Panel A* displays the currency excess return whereas *Panel B* reports the exchange rate return component. *Panel C* presents the sample correlations of the currency excess returns. The superscripts *, **, and *** indicate statistical significance for the mean at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Returns are expressed in percentage per annum and adjusted for transaction costs. The sample period comprises monthly data from January 1998 to December 2013.

Panel A: Currency excess returns										
	<i>CAR</i>	<i>MOM</i>	<i>VAL</i>	<i>RR</i>	<i>VRP</i>	<i>CAR</i>	<i>MOM</i>	<i>VAL</i>	<i>RR</i>	<i>VRP</i>
	<i>Developed</i>					<i>Developed and emerging</i>				
<i>Mean</i>	4.10	0.92	3.66	5.10*	4.95**	4.90**	0.19	2.30	4.25	4.16**
<i>Sdev</i>	10.73	9.81	8.97	11.49	8.15	9.25	8.17	8.23	10.20	8.14
<i>Skew</i>	-0.71	0.26	-0.16	-0.47	-0.03	-0.65	0.07	-0.47	-0.52	0.01
<i>Kurt</i>	5.25	3.75	3.71	5.41	3.97	4.21	3.81	5.08	5.26	4.54
<i>SR</i>	0.38	0.09	0.41	0.44	0.61	0.53	0.02	0.28	0.42	0.51
<i>SO</i>	0.49	0.16	0.64	0.62	0.93	0.74	0.04	0.40	0.53	0.76
<i>MDD</i>	-37.8	-22.8	-15.1	-35.1	-17.0	-28.2	-18.8	-15.1	-31.5	-24.4
<i>AC</i>	0.08	-0.02	-0.02	0.08	0.05	0.05	-0.10	-0.09	0.09	-0.02
<i>Freq_L</i>	0.10	0.48	0.09	0.08	0.29	0.14	0.51	0.08	0.16	0.28
<i>Freq_S</i>	0.07	0.44	0.07	0.22	0.35	0.16	0.47	0.06	0.21	0.27
Panel B: Exchange rate returns										
<i>Mean</i>	-0.81	0.94	2.15	1.84	5.45***	-1.61	0.35	0.48	0.21	5.27**
<i>Sdev</i>	10.76	9.87	9.02	11.58	8.12	9.29	8.18	8.27	10.27	8.20
<i>Skew</i>	-0.72	0.33	-0.24	-0.50	-0.03	-0.72	0.09	-0.55	-0.55	0.09
<i>Kurt</i>	5.43	3.94	3.76	5.63	4.04	4.35	4.06	5.29	5.65	5.10
<i>SR</i>	-0.08	0.09	0.24	0.16	0.67	-0.17	0.04	0.06	0.02	0.64
<i>SO</i>	-0.10	0.17	0.36	0.22	1.01	-0.23	0.07	0.08	0.03	0.97
<i>MDD</i>	-43.3	-23.2	-22.5	-40.3	-14.5	-37.3	-18.2	-20.9	-38.0	-17.9
<i>AC</i>	0.09	-0.01	-0.02	0.09	0.04	0.08	-0.10	-0.08	0.11	-0.02
<i>Freq_L</i>	0.10	0.48	0.09	0.08	0.29	0.14	0.51	0.08	0.16	0.28
<i>Freq_S</i>	0.07	0.44	0.07	0.22	0.35	0.16	0.47	0.06	0.21	0.27
Panel C: Correlations										
<i>CAR</i>	1.00	-0.20	0.30	0.76	-0.08	1.00	-0.07	0.32	0.59	-0.06
<i>MOM</i>	-0.20	1.00	-0.20	-0.23	0.11	-0.07	1.00	-0.19	-0.17	0.15
<i>VAL</i>	0.30	-0.20	1.00	0.46	0.19	0.32	-0.19	1.00	0.57	-0.04
<i>RR</i>	0.76	-0.23	0.46	1.00	0.10	0.59	-0.17	0.57	1.00	0.09
<i>VRP</i>	-0.08	0.11	0.19	0.10	1.00	-0.06	0.15	-0.04	0.09	1.00

Table 4

Currency strategies: sub-samples.

This table presents descriptive statistics of the exchange return component to currency strategies formed using time $t - 1$ information. *CAR* is the carry trade strategy that buys (sells) the top 20% of all currencies with the highest (lowest) interest rate differential relative to the US dollar. Similarly, *MOM* is the momentum strategy that buys (sells) currencies with the highest (lowest) past 3-month exchange rate return, *VAL* is the value strategy that buys (sells) currencies with lowest (highest) real exchange rate, *RR* is the risk reversal strategy that buys (sells) currencies with the lowest (highest) 1-year 10-delta risk reversal, and *VRP* is the volatility risk premium strategy that buys (sells) currencies with the highest (lowest) 1-year volatility risk premium. The table also reports first order autocorrelation coefficient (*AC*) and the annualized Sharpe ratio (*SR*). The superscripts *, **, and *** indicate statistical significance for the mean at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Returns are expressed in percentage per annum and adjusted for transaction costs. The sample period comprises monthly data from March 2001 to November 2001, and from December 2007 to June 2009 in *Panel A*, from January 1998 to December 2006 in *Panel C*, and from January 2007 to December 2013 in *Panel D*.

Panel A: NBER recession periods										
	<i>CAR</i>	<i>MOM</i>	<i>VAL</i>	<i>RR</i>	<i>VRP</i>	<i>CAR</i>	<i>MOM</i>	<i>VAL</i>	<i>RR</i>	<i>VRP</i>
	<i>Developed</i>					<i>Developed and emerging</i>				
<i>Mean</i>	-9.63	11.01	4.58	-5.96	9.32	-11.72	7.83	-0.06	-8.78	10.65
<i>Sdev</i>	17.12	15.41	12.03	17.80	10.88	13.31	10.76	10.62	14.84	10.59
<i>Skew</i>	-0.44	0.28	-0.63	-0.79	-0.43	-0.79	0.62	-1.72	-1.14	0.63
<i>Kurt</i>	3.71	2.87	3.43	4.45	3.61	2.68	4.23	7.87	4.04	5.01
<i>SR</i>	-0.56	0.71	0.38	-0.33	0.86	-0.88	0.73	-0.01	-0.59	1.01
<i>AC</i>	0.35	0.12	-0.09	0.38	0.03	0.44	0.03	-0.24	0.38	0.05
Panel B: non-NBER recession periods										
<i>Mean</i>	0.70	-0.78	1.74	3.18	4.79**	0.12	-0.92	0.58	1.75	4.35**
<i>Sdev</i>	9.26	8.56	8.45	10.18	7.57	8.37	7.64	7.84	9.26	7.73
<i>Skew</i>	-0.54	-0.06	-0.08	0.06	0.09	-0.37	-0.28	-0.05	0.10	-0.22
<i>Kurt</i>	3.96	2.57	3.55	3.59	3.87	4.02	3.05	3.37	4.65	4.41
<i>SR</i>	0.08	-0.09	0.21	0.31	0.63	0.01	-0.12	0.07	0.19	0.56
<i>AC</i>	-0.08	-0.09	-0.01	-0.07	0.05	-0.11	-0.16	-0.04	-0.03	-0.06
Panel C: Pre-crisis period										
<i>Mean</i>	0.78	-0.11	1.76	3.41	4.54*	0.55	-0.64	0.30	2.58	5.28**
<i>Sdev</i>	8.22	7.96	9.94	9.96	7.45	8.33	7.37	8.81	9.61	7.90
<i>Skew</i>	-0.80	-0.04	-0.26	0.32	0.30	-0.79	-0.02	-0.08	0.33	0.27
<i>Kurt</i>	5.05	2.50	3.24	3.87	3.99	4.76	2.89	3.03	4.36	3.44
<i>SR</i>	0.09	-0.01	0.18	0.34	0.61	0.07	-0.09	0.03	0.27	0.67
<i>AC</i>	-0.09	-0.15	-0.03	0.00	-0.04	-0.10	-0.17	-0.03	0.04	-0.06
Panel D: Post-crisis period										
<i>Mean</i>	-2.85	2.28	2.65	-0.17	6.62*	-4.37	1.63	0.72	-2.83	5.26
<i>Sdev</i>	13.37	11.93	7.75	13.42	8.93	10.38	9.15	7.58	11.06	8.62
<i>Skew</i>	-0.55	0.39	-0.12	-0.84	-0.31	-0.59	0.13	-1.47	-1.25	-0.08
<i>Kurt</i>	4.27	3.58	4.65	5.53	3.95	3.79	4.40	10.27	6.00	6.53
<i>SR</i>	-0.21	0.19	0.34	-0.01	0.74	-0.42	0.18	0.10	-0.26	0.61
<i>AC</i>	0.18	0.07	0.01	0.15	0.11	0.22	-0.04	-0.17	0.16	0.02

Table 5

Static currency strategies.

This table presents descriptive statistics of static currency strategies. $\overline{CAR}$ is the carry trade strategy that buys (sells) the top 20% of all currencies with the highest (lowest) full sample average interest rate differential relative to the US dollar. Similarly, $\overline{MOM}$ is the momentum strategy that buys (sells) currencies with the highest (lowest) full sample average 3-month exchange rate return, $\overline{VAL}$ is the value strategy that buys (sells) currencies with lowest (highest) full sample average real exchange rate, $\overline{RR}$ is the risk reversal strategy that buys (sells) currencies with the lowest (highest) full sample average 1-year 10-delta risk reversal, and $\overline{VRP}$ is the volatility risk premium strategy that buys (sells) currencies with the highest (lowest) full sample average 1-year volatility risk premium. The table also reports first order autocorrelation coefficient (AC), the annualized Sharpe ratio (SR), the Sortino ratio (SO), the percentage maximum drawdown (MDD), the frequency of portfolio switches for the long ($Freq_L$) and the short ($Freq_S$) position. *Panel A* displays the currency excess return whereas *Panel B* reports the exchange rate return component. *Panel C* presents the sample correlations of the currency excess returns. The superscripts *, **, and *** indicate statistical significance for the mean at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Returns are expressed in percentage per annum. The sample period comprises monthly data from January 1998 to December 2013.

Panel A: Currency excess returns										
	$\overline{CAR}$	$\overline{MOM}$	$\overline{VAL}$	$\overline{RR}$	$\overline{VRP}$	$\overline{CAR}$	$\overline{MOM}$	$\overline{VAL}$	$\overline{RR}$	$\overline{VRP}$
	<i>Developed</i>					<i>Developed and emerging</i>				
<i>Mean</i>	4.62	2.58	-1.39	4.62	3.51	4.44	3.99**	-1.11	4.78*	3.28
<i>Sdev</i>	11.52	6.88	6.55	11.52	9.50	10.46	6.56	5.55	9.94	8.53
<i>Skew</i>	-0.66	0.51	-0.07	-0.66	-0.35	-0.63	0.31	0.62	-0.69	-0.16
<i>Kurt</i>	5.03	7.33	4.24	5.03	3.07	4.00	3.87	4.44	4.54	4.47
<i>SR</i>	0.40	0.38	-0.21	0.40	0.37	0.42	0.61	-0.20	0.48	0.38
<i>SO</i>	0.55	0.60	-0.31	0.55	0.55	0.61	1.12	-0.38	0.67	0.55
<i>MDD</i>	0.37	0.12	0.31	0.37	0.21	0.29	0.12	0.31	0.30	0.18
<i>AC</i>	0.07	-0.07	-0.11	0.07	0.07	0.10	-0.02	-0.10	0.11	-0.09
Panel B: Exchange rate returns										
<i>Mean</i>	-0.25	2.48	-1.51	-0.25	0.65	-1.80	4.81***	-0.88	-1.07	2.92
<i>Sdev</i>	11.58	6.91	6.55	11.58	9.56	10.49	6.59	5.54	10.00	8.60
<i>Skew</i>	-0.66	0.45	-0.05	-0.66	-0.33	-0.65	0.31	0.64	-0.72	-0.08
<i>Kurt</i>	5.12	7.27	4.27	5.12	3.03	4.04	3.88	4.62	4.59	4.66
<i>SR</i>	-0.02	0.36	-0.23	-0.02	0.07	-0.17	0.73	-0.16	-0.11	0.34
<i>SO</i>	-0.03	0.56	-0.34	-0.03	0.10	-0.24	1.35	-0.30	-0.15	0.49
<i>MDD</i>	0.43	0.12	0.31	0.43	0.27	0.41	0.11	0.23	0.36	0.20
<i>AC</i>	0.08	-0.07	-0.09	0.08	0.08	0.11	-0.02	-0.08	0.12	-0.08
Panel C: Correlations										
$\overline{CAR}$	1.00	0.17	-0.25	1.00	0.73	1.00	-0.18	-0.26	0.96	0.29
$\overline{MOM}$	0.17	1.00	-0.72	0.17	0.33	-0.18	1.00	-0.15	-0.07	0.20
$\overline{VAL}$	-0.25	-0.72	1.00	-0.25	-0.57	-0.26	-0.15	1.00	-0.27	-0.53
$\overline{RR}$	1.00	0.17	-0.25	1.00	0.73	0.96	-0.07	-0.27	1.00	0.27
$\overline{VRP}$	0.73	0.33	-0.57	0.73	1.00	0.29	0.20	-0.53	0.27	1.00

Table 6

Risk factors and volatility risk premium strategy: time series tests.

This table presents time-series regression estimates. The dependent variable is the volatility risk premium strategy (*VRP*) that buys (sells) currencies with the highest (lowest) 1-year volatility risk premium. *Panel A* uses the currency strategies described in Table 3 as explanatory variables. *Panel B* employs the Fama and French (1992) and the equity momentum factors whereas *Panel C* uses the Fung and Hsieh (2004) factors. The superscripts *, **, and *** indicate statistical significance at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Returns are annualized and adjusted for transaction costs (except the equity and the hedge fund factors). The sample period comprises monthly data from January 1998 to December 2013. Fama and French (1992) factors are from French's website whereas the Fung and Hsieh (2004) factors are from Hsieh's website.

Panel A: Currency factors									
	α	DOL	CAR	MOM	VAL	RR	R^2		
<i>Developed</i>	0.04**	−0.06	−0.24***	0.12	0.14	0.23***	0.08		
<i>Developed and emerging</i>	0.04**	0.23**	−0.22*	0.18**	0.02	0.14	0.08		
Panel B: Equity factors									
	α	R_m^e	SMB	HML	$MOME$		R^2		
<i>Developed</i>	0.06***	−0.06	−0.02	−0.07	−0.05*		0.01		
<i>Developed and emerging</i>	0.05***	−0.04	−0.08	−0.07*	−0.06*		0.02		
Panel C: Hedge fund factors									
	α	$Bond$ $Trend$	$Curr$ $Trend$	$Comm$ $Trend$	$Equity$ $Market$	$Size$ $Spread$	$Bond$ $Sarket$	$Credit$ $Spread$	R^2
<i>Developed</i>	0.05**	< .01	< .01	< .01	−0.04	−0.04	0.05	0.02	−0.03
<i>Developed and emerging</i>	0.04**	0.15	0.04	0.10	0.01	−0.10*	−0.20*	−0.17	0.02

Table 7

Asset pricing tests

This table reports asset pricing tests for a linear factor model that includes the dollar (DOL), the carry trade (CAR), and the foreign exchange global volatility (VOL_{FX}) factors. DOL is equivalent to a strategy that borrows in the US money market and equally invests in all foreign currencies, and serves as a constant in the cross-section. CAR is a long-short strategy that buys (sells) the top 20% of all currencies with the highest (lowest) interest rate differential relative to the US dollar. VOL_{FX} is computed as the innovations to a first order autoregressive process applied to the average foreign exchange rate volatility. The test assets are excess returns to five currency portfolios sorted on the 1-year volatility risk premium at time $t - 1$. *Panel A* reports GMM and Fama-MacBeth (FMB) estimates of the market price of risk λ , and the Hansen-Jagannathan distance HJ test for the null hypothesis that the pricing errors are jointly zero. *Panel B* reports least-squares estimates of time series regressions and the χ^2 test for the null that all intercepts are jointly zero. The superscripts *, **, and *** indicate statistical significance at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991) for GMM estimates, and Shanken (1992) for FMB estimates. Returns are annualized and adjusted for transaction costs. The sample period comprises monthly data from January 1998 to December 2013.

Panel A: Cross-section										
	λ_{DOL}	λ_{CAR}	R^2	HJ		λ_{DOL}	λ_{CAR}	R^2	HJ	
	<i>Developed</i>					<i>Developed and emerging</i>				
GMM_1	0.02	-0.10	0.19	0.18		0.02	-0.07	0.24	0.16	
GMM_2	0.02	-0.07	0.16			0.02	-0.04	0.19		
FMB	0.02	-0.10	0.19			0.02	-0.07	0.24		
	λ_{DOL}	λ_{VOL}	R^2	HJ		λ_{DOL}	λ_{VOL}	R^2	HJ	
GMM_1	0.02	0.07	0.24	0.17		0.02	0.01	0.14	0.16	
GMM_2	0.02	0.07	0.24			0.02	-0.02	0.14		
FMB	0.02	0.07	0.24			0.02	0.01	0.14		
Panel B: Time-series										
	α	β_{DOL}	β_{CAR}	R^2	χ^2	α	β_{DOL}	β_{CAR}	R^2	χ^2
P_L	0.02	0.94***	0.01	0.69	7.94	0.02**	1.06***	-0.02	0.75	7.77
P_2	0.01	0.99***	-0.01	0.75		0.00	0.95***	0.02	0.81	
P_3	-0.01	0.92***	0.08	0.72		-0.01	0.99***	-0.02	0.82	
P_4	0.01	1.14***	-0.14	0.82		0.00	1.19***	-0.09	0.84	
P_S	-0.03**	1.01***	0.05	0.76		-0.02*	0.82	0.11*	0.72	
	α	β_{DOL}	β_{VOL}	R^2	χ^2	α	β_{DOL}	β_{VOL}	R^2	χ^2
P_L	0.02*	0.95***	0.13	0.69	7.90	0.02*	1.05***	0.04	0.75	5.60
P_2	0.01	0.97***	-0.15	0.76		0.01	0.94***	-0.17	0.81	
P_3	-0.01	0.96***	0.09	0.71		-0.01	0.99***	0.04	0.82	
P_4	0.01	1.08***	0.07	0.80		0.00	1.17***	0.11	0.84	
P_S	-0.03***	1.02***	-0.15	0.75		-0.01	0.86***	-0.06	0.71	

Table 8

Beta-sorted portfolios: average volatility risk premia.

This table presents descriptive statistics of beta-sorted currency portfolios. Each beta is obtained by regressing individual currency excess returns on the average volatility risk premia using a 36-month moving window. The long (short) portfolio P_L (P_S) contains the top 20% of all currencies with the lowest (highest) beta. The table also reports the first order autocorrelation coefficient (AC), the annualized Sharpe ratio (SR), and the frequency of portfolio switches ($Freq$). *Panel A* displays the currency excess return whereas *Panel B* reports the exchange rate component. *Panel C* presents the pre- and post-formation β s, and the pre- and post-formation interest rate differential (if) relative to the US dollar. The superscripts *, **, and *** indicate statistical significance at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Returns are expressed in percentage per annum and adjusted for transaction costs. The sample runs from January 1998 to December 2013.

Panel A: Currency excess returns												
	P_L	P_2	P_3	P_4	P_S	P_L-P_S	P_L	P_2	P_3	P_4	P_S	P_L-P_S
	<i>Developed</i>						<i>Developed and emerging</i>					
<i>Mean</i>	3.84	2.20	2.87	3.12	8.85**	-5.01	3.41	3.11	3.70	2.87	7.84**	-4.44
<i>Sdev</i>	9.13	10.51	9.28	10.34	11.94	10.69	8.16	9.62	9.64	10.26	12.23	10.82
<i>Skew</i>	0.35	-0.11	-0.65	-0.27	-0.50	0.87	0.08	0.19	-0.50	-0.55	-0.84	1.06
<i>Kurt</i>	3.42	4.50	5.13	4.55	5.31	7.32	2.60	5.13	4.60	4.63	5.87	6.79
<i>SR</i>	0.42	0.21	0.31	0.30	0.74	-0.47	0.42	0.32	0.38	0.28	0.64	-0.41
<i>AC</i>	0.06	-0.02	0.17	-0.01	0.01	0.06	0.07	0.01	0.09	0.02	0.00	-0.03
<i>Freq</i>	0.16	0.25	0.28	0.27	0.11	0.11	0.16	0.22	0.26	0.22	0.12	0.12
Panel B: Exchange rate returns												
<i>Mean</i>	4.62*	2.35	2.51	2.39	6.19*	-1.57	4.35*	3.19	3.19	1.21	5.06	-0.71
<i>Sdev</i>	9.13	10.47	9.32	10.30	11.95	10.79	8.17	9.60	9.64	10.18	12.21	10.87
<i>Skew</i>	0.37	-0.12	-0.67	-0.28	-0.52	0.95	0.08	0.17	-0.51	-0.60	-0.90	1.16
<i>Kurt</i>	3.48	4.45	5.14	4.58	5.32	7.52	2.59	5.15	4.61	4.68	5.97	7.10
<i>SR</i>	0.51	0.22	0.27	0.23	0.52	-0.15	0.53	0.33	0.33	0.12	0.41	-0.07
<i>AC</i>	0.05	-0.03	0.17	-0.01	0.01	0.07	0.06	0.01	0.09	0.01	-0.01	-0.03
<i>Freq</i>	0.16	0.25	0.28	0.27	0.11	0.11	0.16	0.22	0.26	0.22	0.12	0.12
Panel C: Portfolio formation												
<i>pre-if</i>	-0.54	0.06	0.56	0.99	2.35		-0.71	0.14	0.72	1.92	2.37	
<i>post-if</i>	-0.54	0.05	0.54	0.97	2.29		-0.74	0.14	0.71	1.87	2.32	
<i>pre-β</i>	-0.49	-0.18	0.05	0.30	0.71		-0.47	-0.20	0.07	0.36	0.81	
<i>post-β</i>	-0.29**	-0.13	0.22**	0.04	0.16**		-0.21***	-0.18	0.06	0.05	0.12**	

Table 9

Net demand pressure and currency volatility risk premia.

This table presents fixed effects panel estimates of

$$VRP_t^i = \alpha_i + \beta NDem_{t-lag}^i + u_t^i$$

where VRP_{it} is the 1-year volatility risk premium for currency i whereas $NDem_t^i$ denotes the net demand of currency options and futures for two groups of end-users from the US Commodity Futures Trading Commission (CFTC). The net demand is constructed as difference between long and short positions scaled by the total open interest. VRP is defined as the difference between the realized volatility (RV_t) and the synthetic volatility swap rate (SW_t). RV is calculated using daily exchange rate returns over the previous year whereas SW is computed as in Britten-Jones and Neuberger (2000) using 1-year currency option implied volatilities. In VRP_{si} , SW is computed using the simple variance swap method of Martin (2012). In VRP_{garch} , RV is the 1-year volatility forecast generated from the simple garch(1,1). In VRP_{sv} , RV is the 1-year volatility forecast generated from a stochastic volatility model. The superscripts *, **, and *** indicate statistical significance at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). Monthly CFTC data are collected on the last Tuesday of every month. All other variables are measured on the same day. The sample runs from January 1998 to December 2013.

lag	VRP						VRP _{si}					
	α	β	R^2	α	β	R^2	α	β	R^2	α	β	R^2
	Commercial			Financial			Commercial			Financial		
0	-0.011***	-0.016**	0.025	-0.011***	0.021**	0.025	-0.010***	-0.012**	0.017	-0.009***	0.016*	0.016
1	-0.011***	-0.011**	0.012	-0.011***	0.013*	0.009	-0.009***	-0.008*	0.007	-0.009***	0.009	0.005
2	-0.010***	-0.001	< .001	-0.010***	-0.003	< .001	-0.009***	< .001	< .001	-0.009***	-0.004	0.001

lag	VRP _{garch}						VRP _{sv}					
	α	β	R^2	α	β	R^2	α	β	R^2	α	β	R^2
	Commercial			Financial			Commercial			Financial		
0	-0.007***	-0.003*	0.019	-0.007***	0.004	0.014	-0.007***	-0.003*	0.019	-0.007***	0.004	0.014
1	-0.007***	-0.003*	0.013	-0.007***	0.003	0.008	-0.007***	-0.003*	0.013	-0.007***	0.003	0.008
2	-0.007***	-0.002	0.004	-0.007***	0.001	0.001	-0.007***	-0.002	0.004	-0.007***	0.001	0.001

Table 10

Arbitrage risk proxies and VRP.

This table presents predictive regressions estimates. The dependent variable is the exchange rate return component of the *VRP* strategy at time t . This strategy is a long/short portfolio that buys (sells) the top 20% of all currencies with the highest (lowest) 1-year volatility risk premia at time $t - 1$. The set of predictors is measured at time $t - 1$, and includes the *TED* spread, the *VIX* index, the St. Louis Fed Financial Stress Index *FSI*, and the *Fund Flows* of currency and global macro funds constructed as the Asset under Management (AUM) weighted net flows scaled by the lagged AUM as in Patton and Ramadorai (2013). Δ denotes the first-difference operator and *TED*, ΔVIX , and ΔFSI are averaged on a 12-month rolling. The superscripts *, **, and *** indicate statistical significance at 10%, 5%, and 1% level, respectively, based on Newey and West (1987) and Andrews (1991). The exchange rate returns are annualized. The sample runs from January 1998 to December 2013.

α	<i>TED</i>	ΔVIX	ΔFSI	<i>TED</i> × ΔVIX	<i>Fund</i> <i>Flows</i>	R^2	α	<i>TED</i>	ΔVIX	ΔFSI	<i>TED</i> × ΔVIX	<i>Fund</i> <i>Flows</i>	R^2
<i>Developed</i>							<i>Developed and emerging</i>						
-0.01	0.12*					0.02	-0.02	0.14**					0.02
0.06***		0.03				0.01	0.06***		0.07**				0.03
0.06***			0.24			< .01	0.06***			0.50**			0.03
0.05***				0.07***		0.03	0.04***				0.10***		0.07
0.07***					-1.59**	0.02	0.06***					-1.31	0.01
0.02	0.09				-1.21	0.03	0.00	0.11*				-0.83	0.03
0.07***		0.02			-1.42**	0.02	0.06***		0.06***			-0.89	0.03
0.07***			0.16		-1.45*	0.02	0.06***			0.44**		-0.94	0.04
0.06***				0.06***	-1.20	0.04	0.05***				0.10***	-0.65	0.07
0.04	0.02	-0.06*		0.11**	-1.20	0.04	0.04	0.00	-0.04		0.14***	-0.71	0.06




Fig. 1. Global minimum volatility portfolios. The figure presents the global minimum volatility portfolio (MVP) and the efficient frontier (*solid line*) built using the currency strategies formed using $t - 1$ information. *CAR* is the carry strategy that buys (sells) the top 20% of all currencies with the highest (lowest) interest rate differential relative to the US dollar. Similarly, *MOM* is the momentum strategy that buys (sells) currencies with the highest (lowest) past 3-month exchange rate return, *VAL* is the value strategy that buys (sells) currencies with lowest (highest) real exchange rate, *RR* is the risk reversal strategy that buys (sells) currencies with the lowest (highest) 1-year 10-delta risk reversal, and *VRP* is the volatility risk premium strategy that buys (sells) currencies with the highest (lowest) 1-year volatility risk premium. The portfolio weights are reported in parentheses and computed as $w = (\Sigma^{-1}\iota)/(\iota'\Sigma^{-1}\iota)$ where Σ is the $N \times N$ covariance matrix of the strategies' returns, ι is a $N \times 1$ vector of ones, and N denotes the number of strategies. The dashed line denotes the efficient frontier that excludes the volatility risk premium (VRP) strategy. Excess returns are adjusted for transaction costs. The sample period comprises monthly data from January 1998 to December 2013.




Fig. 2. Net demand and volatility risk premium strategy. The figure presents the relation between the volatility risk premium (VRP) strategy and the net demand of currency options and futures from the Commodity Futures Trading Commission (CFTC). We sort currencies into four baskets using the volatility risk premia at time t , and then compute the average net demand of currency options and futures at time t . Finally, we cumulate the difference between the first (currencies with the cheapest volatility insurance) and the last (currencies with the most expensive volatility insurance) portfolio. The net demand is constructed as difference between long and short positions scaled by the total open interest for two groups of end-users. Commercial traders use the futures market primarily to hedge their business activities whereas financial (or non-commercial) traders use the futures market for speculative purposes. The data runs from January 1998 to December 2013 at weekly frequency (collected every Tuesday).