Variance risk on the FX market*

Igor Pozdeev†

this version‡ November 2018

Abstract

I recover risk-neutralized covariance matrices of currency returns and combine them with ex post realized covariance matrices to study FX variance risk. I start by performing eigendecomposition of daily matrix differences and document three general findings: evidence for an overall negative FX variance risk premium; a special place for the US dollar index and Carry trade as the two portfolios with the most negative premium; existence of strategies with a significantly positive premium. On average, portfolios of negative spot return momentum and high recently realized variance are associated with a higher price of own variance risk. I find that the Carry trade variance risk dominates the US dollar variance risk as a priced factor, contributing to resolution of the differential pricing of "good and bad" carry portfolios. Finally, I show that hedging shocks to news about future FX variance is costly with a horizon-decaying pattern, similar to what is observed on the equity market. However, the costs do not become statistically or economically insignificant even at the longest horizons, suggesting a strong preference of FX market investors to hedge such shocks. My findings are thus relevant for FX risk management and theoretical asset pricing.

*I would like to thank (in alphabetical order) Dmitry Borisenko, Peter Carr, Robert Engle, Ralph Koijen, Chiara Legnazzi, Semyon Malamud, Paul Söderlind and Dare Wale, as well as participants of the SFI Research Days, Tandon School of Engineering brown bag seminar and University of St. Gallen seminar for their helpful comments and insightful feedback.

†Igor Pozdeev, NYU Stern School of Business and Swiss Institute of Banking and Finance, *E-mail*: igor.pozdeev@nyu.edu

‡Updated data and Jupyter notebooks with select replication code snippets can be found on my GitHub page: https://github.com/ipozdeev/variance-risk-paper-playground.

1 Introduction

Stochastic variance is considered with little disagreement an inherent feature of financial asset dynamics and a major source of risk for investors. The associated variance risk premium – the average payoff of an asset that is perfectly correlated with a variance process – is a phenomenon intuitive, backed by some of influential asset pricing models and empirically observed. Research on this type of risk was pioneered and has been shaped by its equity market strand, such that several important findings remain endemic to the latter. The FX market strand, which is where my paper belongs, has not only received less attention, but also failed to incorporate staple asset pricing factors and investment vehicles such as the Carry trade into the set of research objects, concentrating instead on USD currency pairs. That said, my work aims to align the two strands by replicating several results from the equity market in the FX market setting, and studying FX market variance risk from the perspective of currency portfolios.

Similar to Carr and Wu (2009), I perform my analysis through the lens of synthetic variance swaps. A variance swap is a contract that pays its buyer an amount equal to the realized variance of a return series in exchange for a predetermined swap rate; the variance risk premium is conveniently defined as the expectation of this net payoff. The notion of "synthetic" means the swaps I am talking about are not actually traded contracts as in Aït-Sahalia et al. (2018) and Dew-Becker et al. (2017), but rather no-arbitrage replicas thereof: under certain assumptions, the swap rate can be approximated by the model-free implied variance of the swap underlying, which in turn can be recovered from a cross-section of option prices. I follow Mueller et al. (2017) to recover such variances for all cross-rates of currencies in my sample and construct option-implied *covariance matrices* of currency returns, from which implied variances of generic portfolios naturally follow.

That said, I present my findings as follows. First, I show that the difference between ex post realized and option-implied covariance matrices is most of the time an indefinite, occasionally a negative definite and never a positive definite matrix. This means that in theory, strategies of both negative and positive variance swap payoffs are possible to construct, and hence that it is hard to make a definite statement about the sign of "the" FX market variance risk premium, although there is a stronger case for an overall negative one. Interestingly, the two eigenvectors corresponding to the most negative eigenvalues in the decomposition of the above matrix difference are aligned with the US dollar index and the Carry trade portfolio respectively, implying that these two portfolios are associated with the most negative variance risk premium of all and thus cementing the special place they have in the FX market fabric.

Second, I document that past spot return information influences conditional variance risk premium in the cross-section and over time. To do so, I take a set of portfolios and regress payoffs of swaps written on their variance onto a set of variables observed at swap inception. I find that portfolios with negative momentum exhibit a more negative price of variance risk, and hence are *ceteris paribus* relatively more expensive to insure against variance increase. Portfolios with a high recently realized variance of returns behave in a similar way, meaning that after periods of high realized variance, investors become more variance risk-averse.

Third, I turn attention to staple FX portfolios and show that protection against higher variance is priced for the Carry trade, Momentum, Value and the US dollar index – in the sense that a buyer of a swap written on the variance of any of them expects to lose money on average. These portfolios hold a special place in the FX market research. The former three have been discovered as profitable trading strategies likely exposed to non-diversifiable risks (Lustig et al. (2011), Asness et al. (2013), Menkhoff et al. (2011))[1]; the US dollar index has been shown as an important driver of the time-series dynamics of

[1]These portfolio are also on the list of tracked FX indexes published by Deutsche Bank.

currency returns (Verdelhan (2018)), and is clearly equivalent to the Dollar carry strategy (Lustig et al. (2014)) in terms of conditional variance of returns. The variance risk premium is negative and similar in magnitude across strategies; for the US dollar index it is comparable to the US equity market variance risk premium. Curiously, the Carry trade strategy, much more volatile than the other FX strategies and notorious for crashes and variance spikes, is far from being associated with the most negative or significant price of own variance risk, dominated in that respect by the US dollar index and the Value strategy in terms of statistical and economic significance respectively. That said, my findings add to the puzzle reported by Caballero and Doyle (2012) and Jurek (2014) who note that it has been strikingly cheap to hedge the Carry trade portfolio with FX options.

Fourth, motivated by my first set of results, I use the US dollar index and Carry trade variance swap payoff series as risk factors in the linear asset pricing framework to show that information in the former is subsumed by that in the latter. With the "good and bad" carry trade portfolios of Bekaert and Panayotov (2018) for test assets, I find only the Carry trade variance risk to be priced, the estimated price of risk being negative and close to the sample average value of the factor, whereas the US dollar index variance risk price is positive in some specifications and never significant.

Fifth and finally, I provide FX market-specific support for the equity market evidence that shocks to *news about future variance* are priced the weaker the more distant the future is, as if investors cared more about the next-period variance than about the variance expected farther ahead. In asset pricing models with preference for early resolution of uncertainty, and in macroeconomic models where shocks to variance expected tomorrow can lead to higher variance today, investors are willing to pay for protection against shocks to expected variance. Surprisingly, Dew-Becker et al. (2017) report that hedging shocks to expected long-term US stock market variance has been remarkably cheap, the Sharpe ratios of hedging strategies quickly going to zero and disappearing beyond the horizon of one quarter. To replicate their result in the FX market setting, I construct synthetic forward variance claims and calculate payoffs from rolling these claims over for various FX portfolios. I find that both payoffs and Sharpe ratios of rolling over forward claims on future quarterly variance tend to diminish with horizon. However, the decline is not as steep as on the equity market, and even for hedging shocks to the variance 4 quarters from now, the costs are statistically different from zero, and Sharpe ratios – comparable to those of the portfolio excess returns themselves. I also zoom in on the shortest segment of 1-4 months and discover a more or less flat term structure of shocks to future variance prices and respective Sharpe ratios. Overall, my findings point to a stronger desire of FX market investors to hedge shocks to long-term variance.

My work belongs to and draws from the relatively young literature on the FX market variance risk as a subset of the variance risk research in general. Negative price of variance risk has been previously documented for the US aggregate stock market (Carr and Wu (2009), Bollerslev et al. (2009)), non-US aggregate stock markets (Bollerslev et al. (2014)), US Treasury bonds (Choi et al. (2017)) and commodities (Tee and Ting (2017)). Della Corte et al. (2016) present a first evidence of the variance risk premium on the FX market by constructing a dollar-neutral portfolio of currencies sorted on a crude measure of conditional variance risk premium. Over-weighting (under-weighting) currencies with high (low) premium, the authors construct a strategy that enhances the mean-variance investment opportunities on the FX market. Ammann and Buesser (2013) find further support for negative variance risk premium in individual exchange rates. However, variance risk of FX portfolios such as the Carry trade and the US dollar index has not received similar attention.

Since variance of a portfolio is a quadratic form of the covariance matrix of portfolio constituents and portfolio weights, I first develop vector representation of currency trading strategies. Then, I recover risk-neutralized covariance matrices of currency returns by making use of the concept of model-free implied variance developed by Britten-Jones and Neuberger (2000) and Carr and Wu (2009) and the assumption of no triangular arbitrage. A similar exercise has been undertaken before. Walter and Lopez (2000) and Mueller et al. (2017) construct option-implied correlations between appreciation rates of currencies against the US dollar: the former paper criticizes the minuscule information content thereof, while the latter explores the properties of a trading strategy that is long (short) currencies with the highest (lowest) loading on the measure of FX correlation risk. My work differs from these studies both in research questions and methodology: namely, the covariances and correlations are never the center of my research, but rather a tool to calculate portfolio variances.[2] Relatedly, Jurek (2014) constructs the analogue of VIX for the carry trade portfolio, but only uses it to highlight the GARCH-in-mean effect in carry trade returns.

The rest of the paper is structured as follows. Section 2 describes the data I use and outlines the vector representation of currency trading strategies, discusses recovery of the option-implied covariances and construction of the variance risk premium estimates. Sections 3 and 4 present the findings. Section 5 concludes.

2 Data and methodology

2.1 Notation

By default, I treat currencies as assets from the point of view of an American investor, such that exchange rate S_x is the dollar "price" of currency x and could be referred to by quote XXXUSD if the three-letter ISO code of currency x is XXX. All τ-period "returns"

[2] One clear advantage of this is while the payoff of a variance swap can at least in theory be replicated with a portfolio of options (see Carr and Madan (1998)), no such replication is possible for individual correlations, which might impact their information content and can serve as the resolution of the critique of Walter and Lopez (2000). I thank Peter Carr for this observation.

defined as $R_x(t, t+\tau) = S_x(t+\tau)/S_x(t)$ are thus by default appreciation rates against USD.

When explicitly written with a double subscript S_{xy}, exchange rate against currency y is meant rather than against the US dollar, which can be thought of as "price" of currency x expressed in units of currency y, and "return" is then the appreciation rate of currency x against y.

In what follows, $s_x = \log S_x$, such that the log-return of currency x is defined as:

$$r_x(t, t+\tau) = \Delta s_x(t+\tau) = s_x(t+\tau) - s_x(t). \quad (1)$$

2.2 Data

I use FX option data from Bloomberg. In the next few paragraphs, I only present the basic facts about FX option data; a comprehensive introduction into the FX option market conventions is given in Wystup (2007) and Malz (2014).

Prices of FX options tend to be expressed in terms of the option Black-Scholes implied volatility (IV). This does not assume that the Black and Scholes (1973) model is understood to hold, but rather represents a one-to-one continuous mapping from the space of currency-denominated option prices to the space of unitless volatility, which allows for easier comparison between options of different strikes and maturities. Bloomberg provides implied volatility quotes against forward deltas (δ) rather than against strike prices, the forward delta[3] being the derivative of the Black-Scholes pricing function with respect to the forward rate of the underlying. Henceforth, delta is understood to be the forward delta. Just as implied volatility is a mapping from the option price, delta is a mapping from the strike price. As deltas are bounded between 0 and 1 in magnitude, they are

[3] As provided by Bloomberg, delta is net of currency premium.

usually multiplied by 100 for quotation purposes to become numbers such as 25, 15 etc.

The most liquid part of the FX option market is concentrated in at-the-money options (ATM) and option contracts, of which Bloomberg provides risk reversals (RR) and butterfly spreads (BF). The notion of at-the-money in Bloomberg is the so-called "delta parity", implying that a call option is at-the-money if it and an otherwise identical put option have the same absolute delta. The other two instruments are linear combinations of plain vanilla call and put options: risk reversals give the holder exposure to the skewness, and butterfly spreads – to the volatility of the underlying exchange rates. That said, for any given day, there are implied volatilities of 10-, 15-, 25- and 35-delta contracts of both types as well as of one ATM option provided by Bloomberg, a total of nine quotes[4]. These are indicative quotes, collected by the data vendor from its suppliers at a particular time of day (usually right after closing of exchanges in the ET time zone) and are snapped at 17:00 New York time. The contracts are conveniently structured to be of constant maturity of 1, 2, 3, 4, 6, 9, 12 months.

As the risk reversals and butterfly spreads are linear combinations of put (P) and call (C) options, it is possible to solve for two call option IVs given the IVs of both contracts and of the ATM. The only necessary assumption is that of no arbitrage (including the put-call parity relation). For example, as shown in Malz (2014):

$$\sigma(C(\delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) + 0.5\sigma(RR(\delta)), \tag{2}$$

$$\sigma(C(1-\delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) - 0.5\sigma(RR(\delta)) \tag{3}$$

where $\sigma(Y(\delta))$ is the IV of contract of type $Y \in \{ATM, BF, RR, C, P\}$ having delta δ. Omitted are the maturity of the contracts and the time subscripts required to be the same for all contracts in the above equations. As discussed in Reiswich and Wystup (2009), (2)-

[4]Sometimes certain quotes would be missing on a given day; most often these are the less liquid 10- and 35-delta ones.

(3) is only valid for small risk reversal implied volatilities. Still, it offers a handy, widely used relation, does not rely on a parametric form proposed in that paper and thus does not conflict with the cubic spline volatility smile interpolation I use later.

That said, it is possible to obtain $2N + 1$ distinct (δ, σ) points from N (RR, BF) quotes and one ATM quote, from where it is straightforward to get to price-strike pairs (C, K) as shown in Wystup (2007). The data other than option quotes needed for the calculations is also collected from Bloomberg. Specifically, for every currency pair I collect the spot exchange rate, and – for each considered maturity – the forward rate rate and two OIS rates as proxy for the risk-free rates. I substitute the risk-free rate of the less traded currency with a synthetic risk-free rate obtained from the covered interest parity (CIP), whereby the ranking is by turnover of OTC foreign exchange instruments reported in BIS (2016). For instance, in the case of AUDCHF, the true Australian dollar OIS rate will be taken and used to infer the Swiss franc risk-free rate, although the OIS rate for the latter currency is also available.

For construction of portfolio spot and excess returns, I use the same spot and forward rates from Bloomberg. The PPP data used for construction of the Value signal is from OECD.

My sample of currencies includes the Australian and Canadian dollar, Swiss franc, euro, British pound, Japanese yen, New Zealand and US dollar; the sample period is from 06/2009 to 12/2017.

Local-economy stock indexes used for construction of equity variance swap rates are obtained from websites of local stock exchanges: for Australia *S&P/ASX 300*, for Canada *S&P/TSX Composite*, for Switzerland *SMI*, for the Eurozone *Euro STOXX*, for the UK *FTSE 100*, for Japan *NIKKEI 225*, and for the US *S&P500*. Respective VIX-like indexes (usually referring to the same basket of stocks that the index is comprised by) are from Bloomberg.

As there is no VIX-like index for the New Zealand market, I exclude this country from calculations where both the stock and FX variance risk premium is considered.

2.3 Variance swap returns and variance risk premium

Imagine an investor at time t wishing to receive a payoff equal to the variance $Var_{t+\tau}(r(t, t+\tau))$ of a return over some time interval $(t, t+\tau)$. Two problems arise: first, since the variance is a latent characteristic of the return process, it is not observed and has to be estimated from data at time $t+\tau$ to be paid out; second, the fair price of this payoff at time t has to be determined.

As a solution to the first problem, Corollary 1 in Andersen et al. (2003) equates[5] the conditional variance of an arbitrage-free process to the conditional expectation of its quadratic variation. Assume that exchange rate S follows an arbitrage-free process $\{S(t)\}$, and let $s(t) = \log S(t)$ as before, such that the continuously compounded appreciation rate is

$$r(t) = ds(t) = \lim_{\Delta t \to 0} s(t+\Delta t) - s(t). \tag{4}$$

The quadratic variation accumulated from time t to time $t+\tau$ is defined as:

$$[r, r]_{t,t+\tau} = \int_t^{t+\tau} (ds(k))^2, \tag{5}$$

and is closely related to the variance of the process:

$$Var_t(r(t, t+\tau)) = E_t([r, r]_{t,t+\tau}) = E_t\left(\int_t^{t+\tau} (ds(k))^2\right), \tag{6}$$

The same Corollary suggests that a natural *ex post* estimator of the variance on the left-hand side of eq. (5) – the quantity the investor expects to receive – is the discretized

[5] Under certain yet not implausible assumptions discussed therein.

version of the right-hand side[6]:

$$Var_{t+\tau}(r(t,t+\tau)) = \frac{254}{D_\tau}\sum_{d=1}^{D_\tau} r\left(t+\frac{d-1}{D_\tau}, t+\frac{d}{D_\tau}\right)^2 = RV(t,t+\tau), \tag{7}$$

where $t + d/D_\tau$ denotes day d of time interval $(t, t+\tau)$, D_τ is the total number of days in that interval, and 254 is the annualisation factor approximately equal to the number of trading days in a year. This quantity is also called the realized variance, hence mnemonic RV. Eq. (7) assumes that daily returns have zero mean: although this seems restrictive and arguably more suitable at frequencies higher than daily, average currency spot returns are notoriously indistinguishable from zero (cf. Lustig et al. (2011)), such that slightly biased but lower-variance estimators of type (7) have been favored in related econometric literature.

The fair price to swap the quantity in eq. (5) for (such that no money changes hands at time t) is by the standard argument its risk-neutralized expectation:

$$IV_t(\tau) = E_t^Q\left(\int_t^{t+\tau}(ds(k))^2\right), \tag{8}$$

also called the variance swap rate. Mnemonic IV has to do with the fact that this rate is essentially the model-free implied variance of the log-price process $\{s(t)\}$. Britten-Jones and Neuberger (2000), building on the results of Carr and Madan (1998) and Breeden and Litzenberger (1978), equate[7] the conditional $\mathbb{Q}$-expectation of the accumulated quadratic variation to the price of a continuous portfolio of options expiring in $(t+\tau)$, weighted by strikes:

$$E_t^Q\left(\int_t^{t+1}(ds(k)^2\right) = 2\int_0^\infty \frac{C_t(\tau,X)-(S(t)-X)^+}{X^2}dX, \tag{9}$$

where $C_t(\tau, X)$ is the time-t price of a call option on S with strike X and maturity of τ,

[6]Estimators of this type are commonly used, see for instance Moreira and Muir (2017), Trolle and Schwartz (2010) and for daily observed FX returns Della Corte et al. (2016).

[7]Under assumptions that were relaxed by Jiang and Tian (2005) to include jump-diffusion processes.

and $(v)^+ = \max(v, 0)$. The integral in eq. (9) can be evaluated numerically, but a careful inspection reveals two potential sources of errors in doing so: first, while the integration in eq. (9) is from zero to infinity, options are traded over a much narrower range, and second, even in that range, strikes are far from being sampled continuously. Addressing these issues, I follow the literature and take the usual steps[8], which are graphically represented in Figure 1. First, as shown by the hollow dots, the observed option prices are transformed into the Black-Scholes implied volatilities to obtain a volatility smile. Then, as shown by the solid line, the smile is interpolated within the available strike range using a cubic spline. Third, as shown by the dashed line, the smile is extrapolated by keeping it constant at the level of the endpoints. Finally, the Black-Scholes volatilities are converted back into prices. For the estimations, I use a grid of 2000 points over the moneyness range between 2/3 and 5/3, and the Simpson's rule to perform the integration.[9]

[Figure 1 about here.]

That said, in form of a variance swap the investor purchases protection against rising variance, the time-$(t+\tau)$ return amounting to:

$$vs(t, t+\tau) = RV(t, t+\tau) - IV_t(\tau), \tag{10}$$

where $RV(t, t+\tau)$ is the realized variance over time interval $(t, t+\tau)$, and $IV_t(\tau)$ is the model-free implied variance observed from a cross-section of options at time t.

The variance risk premium is defined as the expected value of the variance swap return

[8] The same approach, if not without slight variations, has been used by Jiang and Tian (2005), Driessen et al. (2009), Buraschi et al. (2014), Della Corte et al. (2016) and many others.

[9] Another way to arrive at an estimate of a risk-neutral moment would be through calibration of a parametric density to the set of observed option prices, from which calculation of moments is straightforward (see Mirkov et al. (2016) for an example). I have ascertained that deviating from the model-free approach towards a parametric one does not lead to much different variance estimates.

in (10):

$$vrp_t(\tau) = E_t\left(vs(t, t+\tau)\right) = E_t\left(RV(t, t+\tau)\right) - IV_t(\tau), \tag{11}$$

Thus, it is the difference between the objective and risk-neutral expectation of the future realized variance of a stochastic return: if the difference is negative for some asset, investors are ready to pay for hendging the return variance of that asset.

2.4 Vector representation of currency trading strategies

Absence of triangular arbitrage allows to represent any currency position or trading strategy in a vector form, using only exchange rates against one common currency. A vector representation is necessary for computation of moments of trading strategy returns.

I define an N-currency portfolio as a zero-leverage[10] dynamic trading strategy, rebalanced monthly and represented with a vector of weights $(w_{1,t}, w_{2,t}, \ldots, w_{N,t})'$ known at the end of the previous month. For instance, any position involving the euro, the Australian and US dollar can be represented with a 2×1 vector $(w_{aud}, w_{eur})'$ of weights of AUDUSD and EURUSD. The strategy of having 40% of the portfolio in long AUDUSD and 60% in long EURUSD is defined as $(0.4, 0.6)'$, and the weights have to sum up to 1 in absolute value for the zero leverage constraint to bind. The strategy of having 40% in short AUD and 60% in short EUR is represented as $(-0.4, -0.6)'$, the weights again summing up to 1 in absolute value. Now, having 100% in long AUD and 100% in short EUR, represented as $(1.0, -1.0)'$ would also be a valid zero-leverage portfolio, as it is tantamount to be 100% long AUDEUR. In the latter case, the weights do not sum up to 1, but rather the short leg and the long leg do so separately. With that in mind, a currency portfolio must be a

[10] On the ForEx, it would mean that having \$1 in the margin account it is only possible to open \$1 worth of positions.

process $\{v_{t+1}\}$ of the form:

$$v_{t+1} = (w_{1,t}, w_{2,t}, \ldots, w_{n,t})', \tag{12}$$

$$E_t[v_{t+1}] = v_{t+1}, \tag{13}$$

$$\sum_{n=1}^{N} |w_{i,t}| + \left| \sum_{n=1}^{N} w_{i,t} \right| = 2, \tag{14}$$

where eq. (13) says that the composition of the portfolio at time *t+1* is known at time t. The constraint in (14) can be graphically summarized in Figure 2 for the example with AUDUSD and EURUSD: all combinations on the solid rhombus are valid (w_{aud}, w_{eur}) portfolios. Section 2.6 contains examples of popular trading strategies in vector representation.

[Figure 2 about here.]

To stress that the time-*(t+1)* portfolio composition is known at time t, I introduce $w_t = E_t[v_{t+1}] = (w_{1,t+1}, w_{2,t+1}, \ldots, w_{N,t+1})'$. Also, I denote the time-*(t+1)* return of any strategy as f_{t+1} to differentiate it from individual currency returns. This return is calculated in the usual way:

$$f_{t+1} = w_t' r_{t+1}, \tag{15}$$

where $r = (r_1, r_2, \ldots, r_n)'$ is the vector of individual currency returns. The conditional variance of the strategy return is:

$$Var_t(f_{t+1}) = w_t' \Omega_t w_t, \tag{16}$$

where Ω_t is the conditional covariance matrix of currency returns against the US dollar or any other currency.

2.5 Covariance matrices of currency returns

The cornerstone of this paper is the conditional covariance matrix of time-$(t+\tau)$ currency appreciation rates against a common counter currency, whereby the conditioning is w.r.t. the time-t information. Indexing the rows and columns of any such matrix with the base currencies, its (x, y) element reads:

$$\Omega_t[x,y] = \begin{cases} Var_t(r_x(t,t+\tau)), & x = y, \\ Cov_t\left(r_x(t,t+\tau), r_y(t,t+\tau)\right), & x \neq y \end{cases}$$

Absence of triangular arbitrage implies that the log-appreciation rate of currency x against y can be expressed in terms of their log-appreciation rates against a common currency (time subscripts can be dropped as long as returns are contemporaneous):

$$r_{xy} = r_x - r_y. \tag{17}$$

Applying the variance operator to both sides of eq. (17) results in:

$$Var(r_{xy}) = Var(r_x) + Var(r_y) - 2Cov(r_x, r_y),$$

which can be rearranged to isolate the covariance as follows:

$$Cov(r_x, r_y) = \frac{1}{2}(Var(r_x) + Var(r_y) - Var(r_{xy})). \tag{18}$$

The latter equation obviously holds irrespective of the time subscripts, of the set of information used for conditioning the moments, and of the measure under which the moments are taken. Thus, to obtain the covariance between returns of x and y one needs the variance of the appreciation rate of x against y as well as of each of them against a common currency.

2.6 Currency portfolios

2.6.1 US dollar index, *USD*

A return of a currency index is defined as an equally weighted average appreciation rate of foreign currencies against that particular currency: when the index goes up, the currency *depreciates* to the basket of other currencies. For each currency, the composition of this portfolio is the same after each rebalancing, e.g. for the US dollar:

aud	cad	chf	eur	gbp	jpy	nzd
1/7	1/7	1/7	1/7	1/7	1/7	1/7

2.6.2 Carry trade, *CAR*

As in Lustig et al. (2011), currencies are sorted by how much risk-free interest rates in respective economies exceed that in the US. Long positions are opened in the two currencies with the highest, and short position – in the two with the lowest such difference. Instead of taking any particular type of risk-free rate, I rely on the covered interest parity and proxy the above interest rate difference as the negative of the currency's forward discount, smoothing the series over 3 months to avoid FX market microstructure effects.[11] A common representation of the carry trade strategy is as follows:

aud	cad	chf	eur	gbp	jpy	nzd
0.5	0	-0.5	0	0	-0.5	0.5

[11]See Baba and Packer (2009) and Du et al. (2018) for a recent discussion on the parity deviations during the financial crisis.

2.6.3 Momentum, *MOM*

As in Asness et al. (2013), currencies are sorted by their cumulative appreciation against the US dollar over the past 12 months excluding the most recent month. Long positions are opened in the two currencies that have appreciated the most, and short positions – in those that have appreciated the least.

2.6.4 Value, *VAL*

As in Della Corte et al. (2016), currencies are sorted by their real exchange rate against the US dollar. Long positions are opened in the two currencies with the lowest, and short positions – in those with the highest real exchange rate. The real exchange rate is defined as $REER_t = PPP_t / S_t$, where PPP_t is the purchasing power parity rate from OECD data, and S_t is the nominal exchange rate.

2.6.5 Variance risk premium, *VRP*

As in Della Corte et al. (2016), currencies are sorted by a proxy of the time-t variance risk premium for holding that currency against USD. The premium is proxied as the difference between the realized in the previous year and the 1-year model-free implied variance of the currency appreciation rate against the US dollar. Long positions are opened in two currencies with the largest, and short – in those with the lowest such difference.

3 General results about FX variance risk

3.1 Extreme attainable variance risk premia

Access to both realized and implied covariance matrices makes it possible to determine which portfolios are associated with a consistently negative variance risk premium, and which – if any at all – with a positive one. The exercise can be restated as an eigenvalue problem.

As mentioned in Section 2.3 (specifically, eq. (11)), variance risk premium of a portfolio is equal to the expectation of the payoff of a swap written on the portfolio's variance. The payoff equation trivially involves quadratic forms defined by the objective and risk-neutralized covariance matrices:

$$vs(t, t+\tau) = w_i' \Omega^{\mathbb{P}}(t, t+\tau) w_i - w_i' \Omega_t^{\mathbb{Q}}(\tau) w_i \tag{19}$$

$$= w_i' \underbrace{\left(\Omega^{\mathbb{P}}(t, t+\tau) - \Omega_t^{\mathbb{Q}}(\tau)\right)}_{\Theta(t,t+\tau)} w_i, \tag{20}$$

where w_i is a vector of portfolio weights, and Θ, which I call *variance swap payoff structure matrix*, is introduced to denote the difference in question. Period-by-period, a condition necessary and sufficient to make *any* quadratic form defined by this matrix negative (positive), Θ must be negative (positive) definite, which it would if all of its eigenvalues are less (greater) than zero. Hence, I calculate the time series of eigenvalues of the realized $\Theta(t, t+\tau)$. I find that the proportion of days when only a negative variance swap payoff could have been realized, rises from 0.01 at the 1-month to 0.15 at the 4-month horizon, whereas the proportion of days with only a positive payoff being possible is exactly zero irrespective of the horizon. This speaks in favor of an overall negative FX variance risk premium, but is also clearly indicative of a theoretical possibility to construct portfolios with a positive one.

By construction, eigenvalues of Θ are equivalent to variance swap payoffs for portfolios formed with corresponding eigenvectors as weights. Since Θ is Hermitian, the lowest (highest) value is paired with the vector that minimizes (maximizes) the quadratic form in eq. (20) under the condition that vector elements sum up to one when squared.[12] For the 3-month horizon only, Figure 3 presents the time series average of the eigenvalues ranked from the lowest to highest, as well as the corresponding eigenvectors. There is strong evidence that a broad range of average payoffs have been possible, from very negative for portfolio 1 to positive for portfolio 7. Interestingly, portfolio 1 – associated with the lowest eigenvalues – is essentially the US Dollar index, as the weights of different currencies are all similar in magnitude and sign and identified precisely. Portfolio 2 is reminiscent of a typical carry trade, separating the Australian, New Zealand and Canadian dollar from the Japanese yen, Swiss franc and euro, which are the high and low interest rate currencies respectively. This further stresses the special role the US dollar index and the Carry portfolio play on the FX market. At the same time, portfolio 7 – the one associated with the highest eigenvalues – turns out to exhibit a consistently positive variance risk premium, suggesting a possibility of constructing such a .

[Figure 3 about here.]

Mostly identical results emerge when realized covariance entering the computation of variance swap payoff structure matrices is substituted with expected covariance, proxied by past realized covariance in the spirit of many previous variance risk premium studies.[13] The *ex ante* variance risk premium in this case is similarly negative (positive) for portfolios corresponding to the lowest (highest) eigenvalues, the time-series average of

[12] That said, the portfolios obtained from such eigendecomposition are leveraged.

[13] See, for example, Bollerslev et al. (2009), Della Corte et al. (2016), Zhou (2017) and others.

the first two eigenvectors again mimicking the US dollar index and a typical carry trade portfolio.

3.2 Explaining variance risk premium of FX strategies

Given the observed heterogeneity in variance risk premium estimates, a naturally arising question is whether the conditional variance risk premium is explainable by characteristics. This question translates to estimation of the coefficients β in the following panel relation:

$$w_i'\Theta(t, t+\tau)w_i = \beta_i' X_{i,t} + \epsilon_{i,t}, \tag{21}$$

where the left-hand side is the variance swap payoff of portfolio i reaped at time $t+\tau$, Θ and w_i are as defined in eq. (20), and X_i is a vector of portfolio characteristics. I choose X_i to comprise the portfolio average interest rate difference (weighted by w_i), momentum signal (return of portfolio i in the previous year without the most recent month), value signal (weighted average of the logarithms of individual value signals as in Section 2.6) and the recently realized variance (average squared return over the 5 days leading to t). Since there are infinitely many portfolios characterized by weights w_i to estimate eq. (21), I only concentrate on the subset of seven portfolios constructed from rescaled eigenvectors in Figure 3. The choice is admittedly arbitrary, but as I describe earlier, the portfolios therein include a version of the US dollar index, a static carry trade portfolio etc., thus offering a representative range of test assets. For this exercise, I use the 3-month variance risk premium estimates as the most statistically significant ones. As for the estimator, I choose panel OLS with time fixed effects and the Fama-MacBeth procedure. Additionally, I control for entity fixed effects in a separate specification. Within the panel OLS framework, time fixed effects are needed to pinpoint the sources of cross-sectional variation in risk premium estimates, and the entity fixed effects would allow to zoom in on the sources of time series variation over and above a common trend. The Fama-MacBeth estimator is used as a robustness check.

The results are presented in Table 1. Overall, it is the backward-looking spot return information that is a significant determinant of cross-sectional differences in FX variance risk premium estimates and a driver of conditional risk premium dynamics. Portfolios that have enjoyed positive spot return over the previous incomplete year tend to be cheaper to buy variance protection for, as evidenced by the positive and significant coefficient, robust across all estimators. Each extra percentage point of positive return translates *ceteris paribus* to 50-100%% of extra positive difference between the realized and implied variance. Equally significant and robust is the effect of higher recently realized variance: every extra squared percentage point thereof is associated with the variance risk premium becoming more negative by 1-2%%.

As to the other variables, the relation is either not significant, or not robust, or both. Portfolios with a relatively higher forward discount tend to exhibit slightly more positive variance risk premium, although the effect is reversed in the time dimension and barely detectable with the Fama-MacBeth estimator. Portfolios of currencies that are "cheap" in real terms are expected to have a more negative variance risk premium when judged by the panel OLS coefficients, but this is at stark odds with the Fama-MacBeth estimation results.

[Table 1 about here.]

4 Variance risk of popular FX trading strategies

4.1 Time series of FX strategies' variance swap payoffs

Next, I turn to documenting the variance risk premium estimates for a number of popular FX portfolios. As discussed in Section 2.3 (specifically, see eq. (11)), the unconditional premium can be expressed as the unconditional mean of variance swap payoffs, or as the difference between the average realized and implied variance.

Figure 4A shows (the square root of) the average realized and implied variance of portfolio returns for horizons of 1 to 4 months. Not surprisingly, the carry trade strategy is the most volatile of all, both in terms of the average realized and average implied variance, at all horizons, while the US dollar index is the least volatile. A restatement of the previously reported result from Section 3.1, the Carry trade is not the most expensive strategy to hedge the variance risk of, as Figure 4B suggests: its unconditional variance risk premium is either exceeded or matched in magnitude by that of the Value strategy, and both fall short of the US dollar index in terms of the Sharpe ratio of variance swap payoffs. This appears puzzling, given the reputation of the Carry trade as a strategy prone to crashes and volatility spikes much more than the other considered portfolios, and resembles the findings of Jurek (2014) who documents that hedging its downside with FX options increases the performance dramatically, as if the options on the carry-related cross-rates such as AUDJPY were "too cheap".

[Figure 4 about here.]

Table 2 suggests a possible venue for reconciliation of the puzzle: the time series average values of variance swap payoffs are pushed upwards by infrequent large positive observations. The *median* values on the other hand are uniformly twice as negative for the Carry as for the other strategies. Admittedly, these facts distort the analysis somewhat, but not dramatically: even at twice as higher a Sharpe ratio, the carry trade variance swap payoffs would only come close to resemble the US dollar index in terms of variance-riskiness.

[Table 2 about here.]

FX variance risk premium is dwarfed by that of local stock market returns across the economies in my sample, as Figure 5 indicates. Therein, I plot average local stock market variance swap *returns* and Sharpe ratios against the corresponding currency index values. A currency index is defined for different counter currencies similar to the US dollar index. A variance swap return is defined by scaling the variance swap payoff by the inverse of the swap rate, not dissimilar from return on a fully collateralized swap position:

$$vs(t, t+\tau) = \frac{RV(t, t+\tau) - IV_t(\tau)}{IV_t(\tau)} \times 100 \times \frac{12}{\tau}, \tag{22}$$

For this exercise, swap returns are preferred to payoffs because stock markets are on average much more volatile than the FX market, which would distort the cross-market comparison. At the 1-month horizon depicted in the Figure, FX variance swap returns and respective Sharpe ratios are twice as small as their equity market counterparts.

[Figure 5 about here.]

4.2 Asset pricing with variance risk

Given that the US dollar index and Carry trade strategy have been documented to play an important role in FX asset pricing (see e.g. Lustig et al. (2011) and Verdelhan (2018)), and taking into account my previously reported findings, namely coincidence of their composition with principal component loadings, as well as significance of respective variance risk premium estimates – it is of interest to test the corresponding sources of variance risk for cross-sectional asset pricing properties. To do so, I use a linear asset pricing framework and a range of "good and bad" carry trade portfolios constructed as in Bekaert and Panayotov (2018). Each such portfolio represents a carry trade set up on a narrower subsample of G10 currencies[14]. Bekaert and Panayotov (2018) show that these portfolios differ a lot in terms of profitability and distribution characteristics, and as such they provide a fruitful soil for testing asset pricing models.[15]

The model I estimate assumes that the stochastic discount factor of a representative investor linearly depends – among other things – on the US dollar and Carry trade variance swap payoffs, which leads to the following representation of the expected return of portfolio i:[16]

$$E[rx_i(t, t+\tau)] = \beta_i(\tau)'\lambda(\tau), \tag{23}$$

where rx_i is the excess return of the portfolio, λ is a vector of risk premia ("prices of risk"), and β_i is the vector of portfolio loadings on the risk factor ("quantities of risk"), which I allow to differ by horizon τ. Hence, the model implies that the cross-sectional differences in expected portfolio returns arise due to differences in their loadings on the risk factors, since the prices of risk are the same for each portfolio. For the risk factors, I take the US dollar index and Carry trade *returns* as in Lustig et al. (2011), together referred to as benchmark factors, and the two corresponding *variance swap payoffs*, together referred to as variance risk factors. I estimate the model in the two-step way as has become staple

[14]Here, in test assets construction I use the whole G10 sample, including the Danish and Norwegian krone and the Swedish krona for the portfolio construction. The US dollar index and Carry variance swap payoffs are constructed without these currencies, as before.

[15]Clearly, many of these portfolios have returns that are either not linearly independent of return of the other portfolios, or almost perfectly correlated with any other series, a nuisance exacerbated by the small sample size. After ensuring that the matrix of test asset returns has full column rank, I am left with 32 portfolios with the average pairwise return correlation of 0.6.

[16]For details, see Munk (2013)

in asset pricing with non-traded factors: the loadings are estimated in the first step, and then used to pin down the risk premia. Although the above four factors are theoretically traded (the benchmark ones clearly so, while the variance risk ones are such by analogy with the stock market variance swaps and a no-arbitrage condition), Table 2 shows that the time-series average payoffs as the estimates of the risk premium are rather imprecise, especially at the one-month horizon. I will demonstrate however that both the sign and the order of magnitude of the variance risk premium obtained in the second step are the same as those of the sample averages.

Table 3 shows the estimates of the four prices of risk from linear regressions of 1- to 4-month "good and bad" carry portfolios' returns on the risk factors. The first model is a standard specification with the returns of the broad US dollar index and the baseline Carry trade strategy on the right-hand side. A repetition of an old result, expected returns of the test assets are aligned with the exposures to the Carry, but not to the Dollar factor. The second model describes returns as a function of the two variance risk factors, namely the US dollar and Carry trade variance swap payoffs. There is clear evidence that only the Carry trade variance risk is priced: the corresponding lambdas are significant, and the price of risk estimates negative and of a similar magnitude as the sample values. The US dollar index variance risk exposures are not meaningfully aligned with expected returns, as the lambdas are economically minuscule and statistically indistinguishable from zero. In the third model, which embraces both the benchmark and variance risk factors, the price of the Carry trade variance risk becomes insignificant, dominated by the ever significant Carry trade return factor, but the numerical value is barely unchanged; the price of the US dollar index variance risk remains dwarfed by it. Noteworthy is the fact that the drop in significance is less pronounced for horizon of 3 months, where the variance risk premium is more precisely estimated in the first place.

[Table 3 about here.]

4.3 Hedging shocks to future variance

Previously, Figure 4B presented evidence of a negative significant variance risk premium in FX portfolio returns all the way up to the 4-month horizon, and suggested that it has been historically costly for investors to hedge variance of FX trading strategies. In this subsection, I answer a subtler question: do investors only care about hedging the shortest-term (next month or next quarter) variance, or do they also price shocks to news about future variance? In doing so, I closely follow the methodology and notation of Dew-Becker et al. (2017) who conduct a similar analysis for the US equity market. I deviate from their study in that I do not rely on interpolation of the term structure of forward variances, but rather report the results for monthly variance at the horizon of up to 4 month and for quarterly variance at longer horizons. Additionally – as everywhere throughout my work – I do not rescale variance swap payoffs to arrive at variance swap returns, but rather draw inference from time-series of the former.

To visualize the concept of shocks to future variance, consider a currency investor with a three-period investment horizon. There are two types of variance shocks that she will be exposed to in period 1: the first type are shocks to the variance in that period only, while the second type are shocks to the *expectation* of variance in periods 2 and 3. As noted by Dew-Becker et al. (2017), many established long-run risk models in finance (e.g. Bansal and Yaron (2004), Drechsler and Yaron (2011)) and recent works on the macroeconomic consequences of shocks to news about future uncertainty (e.g. Bloom (2009), Fernández-Villaverde et al. (2011)) predict that investors are willing to hedge such news.

To test this implication, I calculate effective costs of hedging shocks to future τ-period variance, $\tau = 2, 3, 4$ months or quarters, whereby the hedging is performed by rolling over forward variance claims. Given definition of variance in eqs. (6) and (8) and of the discretized estimator thereof in eq. (7), it is straightforward to express the variance realized solely in period τ as:

$$RV(\tau, \tau) = RV(t, t+\tau) - RV(t, t+\tau-1). \tag{24}$$

The price of this variance, or the period-τ forward variance claim price, follows trivially from the risk-neutral pricing argument:

$$F_t(\tau) = IV_t(t+\tau) - IV_t(t+\tau-1). \tag{25}$$

In other words, in period t it costs the investor $F_t(\tau)$ to buy forward the variance that will be realized in period $t+\tau$. Come period $t+1$, she can sell forward that variance (which would now be a contract expiring in $\tau - 1$ months), and realize a gain if there has been a positive shock to the expectation thereof. In order to maintain exposure to variance shocks at the τ-period horizon, she would buy forward the $t+\tau+1$ variance and then repeat the same steps, a succession familiar to futures market participants. A significantly large time-series average of payoffs from this succession would indicate that investors price shocks to variance at that particular horizon. By definition, period-1 forward variance is just the implied variance at that horizon, and the corresponding rollover payoff is simply the period-1 variance swap payoff.

Figure 6 shows the term structure of average forward monthly variance claims prices. For all portfolios, the curves are strictly upward sloping suggesting negative returns from the rollover strategy described above. Beyond the horizon of 1 month, the Carry trade exhibits a larger curve slope than the other strategies (the curves of those appear of about the same steepness), although at that horizon, the Value strategy and the Dollar index are the steepest, a repetition of the previously stated fact about the 1-month variance swap payoffs.

[Figure 6 about here.]

Figures 7A-7B depict the payoffs and respective Sharpe ratios of rolling over forward monthly variance claims at different horizons. All payoffs are negative, consistent with the upward-sloping term structure of forward variance prices; most are statistically significant, and Sharpe ratios are large for the US dollar index, Carry trade and Value strategy. More importantly, except for the US Dollar index, rolling over FX forward variance claims is as costly at the 4 months as it is at the 1 month horizon. Hence, it can be concluded that investors are willing to pay a non-negligible price to hedge shocks to news about variance at least within a popular rebalancing frequency[17].

[Figure 7 about here.]

Figures 8A-8B show the outcome of the same exercise conducted on quarterly forward variances at thresholds of 3, 6, 9 and 12 months: therein, $F(\tau)$ is the average payoff of rolling over forward claims on the variance realized over $(\tau - 3 < t < \tau)$, such that $F(3m)$ is again just the 3-month variance swap payoff. A somewhat different pattern emerges: shocks to the nearest and second nearest quarter variance are priced much stronger than shocks to more distant variance, as evidenced by declining payoffs and Sharpe ratios. For all portfolios, the fourth nearest quarter values are twice as low in magnitude as the second nearest quarter ones. However, most values are still statistically significant even at the longest horizons, and Sharpe ratios are as high as the those of the strategy excess returns.

[Figure 8 about here.]

Overall, there is evidence that investors care about shocks to news about future variance of FX portfolios, and despite the fact that these shocks are "cheaper" the farther the

[17] For instance, used in Deutsche Bank FX indexes.

hedged variance, investors are still ready to enter transactions of return profile comparable to trading strategies themselves to hedge event the shocks to news about the most distant variance away. This finding stands in contrast with what Dew-Becker et al. (2017) find for the US equity market: for a sample of US market variance swaps, the authors report that shocks to uncertainty at horizons larger than one quarter could have been hedged virtually for free.

5 Conclusions

FX portfolios, such as the US dollar index and Carry trade, are at the heart of FX market academic research and investment management, yet have been ignored by studies on FX variance risk premium. In this paper, I aim to redress the balance by using option-implied covariance matrices of currency returns and representing trading strategies in vector form to construct synthetic swaps on the variance of these strategies and study FX variance risk.

I find evidence of an overall negative FX variance risk premium, as the eigenvalues of the matrix difference between the implied and subsequently realized covariances are predominantly negative. The US dollar index and a typical carry trade strategy appear as the two portfolios with the most negative variance risk premium estimates of *all* possible portfolios, further stressing the own special place on the FX market. At the same time, strategies with a significantly positive ex post variance risk premium are achievable, up to the limits of predictability of the realized covariance.

Cross-sectional and time-series differences in conditional variance risk premium estimates can be explained by the past time-series information such as the spot return and realized variance: portfolios of negative momentum and high recently realized variance are on average the most expensive to insure against rising variance.

The Carry trade strategy, despite being the most volatile and crash-prone, does not stand out as the one with the highest negative variance risk price, superseded in that respect by the US Dollar index. Nevertheless, the Carry trade variance risk is an important factor for explaining the cross-section of FX returns, just as the Carry trade excess returns are a priced risk factor. The Carry trade variance risk is priced in the cross-section of "good and bad" carry portfolios, contributing to the resolution of their differential pricing, whereas the US dollar variance risk is not.

Shocks to news about future variance on the FX market are priced, the price being the smaller the more distant the expectation horizon. However, this decline is not as pronounced as for the equity market, where Sharpe ratios of strategies providing insurance against such shocks drop to zero beyond the horizon of 3 months. For FX portfolios, hedging shocks to news about future variance has been statistically and economically costly even for longer horizons of 6, 9 and 12 months.

References

Aït-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini, 2018, The term structure of variance swaps and risk premia, Research paper 18-37, Swiss Finance Institute.

Ammann, Manuel, and Ralf Buesser, 2013, Variance risk premiums in foreign exchange markets, *Journal of Empirical Finance* 23, 16 – 32.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–625.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *The Journal of Finance* 68, 929–985.

Baba, Naohiko, and Frank Packer, 2009, Interpreting deviations from covered interest parity during the financial market turmoil of 2007-08, *Journal of Banking & Finance* 33, 1953–1962.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

Bekaert, Geert, and George Panayotov, 2018, Good carry, bad carry, Research paper 15-53, Columbia Business School.

BIS, 2016, Triennial central bank survey of foreign exchange and derivative market activity in 2016, Technical report, Bank for International Settlements Editorial Group.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–54.

Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Bollerslev, Tim, James Marrone, Lai Xu, and Hao Zhou, 2014, Stock return predictability and variance risk premia: Statistical inference and international evidence, *Journal of Financial and Quantitative Analysis* 49, 633661.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *The Review of Financial Studies* 22, 4463–4492.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *The Journal of Business* 51, 621–651.

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *The Journal of Finance* 55, 839–866.

Buraschi, Andrea, Robert Kosowski, and Fabio Trojani, 2014, When there is no place to hide: Correlation risk and the cross-section of hedge fund returns, *The Review of Financial Studies* 27, 581–616.

Caballero, Ricardo J., and Joseph B. Doyle, 2012, Carry trade and systemic risk: Why are FX options so cheap?, Working Paper 18644, National Bureau of Economic Research.

Carr, Peter, and Dilip Madan, 1998, Towards a theory of volatility trading, in *Reprinted in Option Pricing, Interest Rates, and Risk Management, Musiella, Jouini, Cvitanic*, 417–427 (University Press).

Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *The Review of Financial Studies* 22, 1311–1341.

Choi, Hoyong, Philippe Mueller, and Andrea Vedolin, 2017, Bond variance risk premiums*, *Review of Finance* 21, 987–1022.

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno, 2016, Volatility risk premia and exchange rate predictability, *Journal of Financial Economics* 120, 21 – 40.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez, 2017, The price of variance risk, *Journal of Financial Economics* 123, 225–250.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *The Review of Financial Studies* 24, 1–45.

Driessen, Joost, Pascal J. Maenhout, and Grigory Vilkov, 2009, The price of correlation risk: Evidence from equity options, *The Journal of Finance* 64, 1377–1406.

Driscoll, John C., and Aart C. Kraay, 1998, Consistent covariance matrix estimation with spatially dependent panel data, *The Review of Economics and Statistics* 80, 549–560.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2018, Deviations from covered interest rate parity, *The Journal of Finance* 73, 915–957.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martin Uribe, 2011, Risk matters: The real effects of volatility shocks, *American Economic Review* 101, 2530–61.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

Jurek, Jakub W., 2014, Crash-neutral currency carry trades, *Journal of Financial Economics* 113, 325 – 347.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

Malz, Allan M, 2014, A simple and reliable way to compute option-based risk-neutral distributions, Federal Reserve Bank of New York Staff Report 677.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2011, Currency momentum strategies.

Mirkov, Nikola, Igor Pozdeev, and Paul Söderlind, 2016, Toward removal of the swiss franc cap: Market expectations and verbal interventions, Research Paper 2016/14.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *The Journal of Finance* 72, 1611–1644.

Mueller, Philippe, Andreas Stathopoulos, and Andrea Vedolin, 2017, International correlation risk, *Journal of Financial Economics* 126, 270 – 299.

Munk, Claus, 2013, *Financial Asset Pricing Theory*, first edition (Oxford University Press).

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Newey, Whitney K., and Kenneth D. West, 1994, Automatic lag selection in covariance matrix estimation, *The Review of Economic Studies* 61, 631–653.

Reiswich, Dimitri, and Uwe Wystup, 2009, FX volatility smile construction, CPQF Working Paper Series 20, Frankfurt School of Finance and Management, Centre for Practical Quantitative Finance (CPQF).

Tee, Chyng Wen, and Christopher Ting, 2017, Variance risk premiums of commodity ETFs, *Journal of Futures Markets* 37, 452–472.

Trolle, Anders B, and Eduardo S Schwartz, 2010, Variance risk premia in energy commodities, *The Journal of Derivatives* 17, 15–32.

Verdelhan, Adrien, 2018, The share of systematic variation in bilateral exchange rates, *The Journal of Finance* 73, 375–418.

Walter, Christian, and Jose Lopez, 2000, Is implied correlation worth calculating? evidence from foreign exchange options and historical data, Working Paper Series 2000-02, Federal Reserve Bank of San Francisco.

Wystup, U., 2007, *FX Options and Structured Products*, The Wiley Finance Series (Wiley).

Zhou, Hao, 2017, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, Review article for the annual review of financial economics.

Table 1: Explaining differences in variance risk premium estimates

	(1)	(2)	(3)
estimator	panel OLS	panel OLS	Fama-MacBeth
N obs.	17136	17136	17136
R^2	0.08	0.07	0.09
R^2 (within)	0.03	0.10	0.01
R^2 (between)	−0.39	0.31	0.97
R^2 (overall)	0.00	0.11	0.09
forward discount	3.44	−3.65	1.05
	(1.73)	(−0.73)	(0.47)
momentum signal	70.62	55.67	100.56
	(3.72)	(2.71)	(2.92)
value signal	−11.10	−41.69	44.63
	(−0.45)	(−1.46)	(3.10)
recent RV	−1.13	−1.12	−2.10
	(−3.28)	(−3.21)	(−2.21)
time f.e.	yes	yes	n/a
entity f.e.	no	yes	n/a

This table shows estimates of coefficients in eq. (21) which relates subsequent 3-month portfolio variance swap payoffs to portfolio characteristics observed at swap inception. Each column features a different estimator: in (1) panel OLS with time fixed effects is used, in (2) – panel OLS with time and entity fixed effects, and in (3) – Fama-MacBeth methodology. All variance-related variables are in percent-squared per year; returns are in percent per year, and value signal is dimensionless. In parentheses below coefficient estimates are respective t-statistics, calculated using Driscoll and Kraay (1998) standard errors. The sample period is from 01/2009 to 06/2018.

Table 2: Variance swap payoffs of for FX trading strategies: sample statistics

	usd	car	val	mom	vrp
1 month					
mean	−10.65	−14.19	−20.64	−12.25	−8.00
	(2.39)	(9.18)	(6.00)	(5.37)	(4.98)
sharpe	−1.16	−0.52	−1.20	−0.68	−0.43
std	31.79	95.13	59.37	62.07	65.16
median	−11.21	−18.82	−15.53	−11.82	−11.15
skewness	1.02	1.21	−2.14	1.05	3.11
2 months					
mean	−14.08	−23.89	−30.51	−18.75	−12.88
	(3.41)	(12.51)	(19.11)	(7.83)	(5.64)
sharpe	−1.24	−0.62	−0.63	−0.72	−0.53
std	27.70	94.35	119.13	63.58	59.13
median	−12.67	−24.04	−12.36	−12.84	−14.80
skewness	−0.23	0.42	−3.07	−1.33	1.93
3 months					
mean	−18.15	−33.05	−35.95	−22.69	−17.39
	(4.51)	(15.45)	(22.73)	(10.06)	(6.34)
sharpe	−1.26	−0.67	−0.61	−0.69	−0.63
std	28.83	98.05	118.72	65.56	54.95
median	−15.12	−28.03	−14.18	−15.63	−16.48
skewness	−0.49	−0.08	−2.94	−1.73	0.76
4 months					
mean	−21.99	−41.26	−40.80	−26.94	−22.87
	(5.51)	(17.65)	(25.48)	(11.99)	(7.06)
sharpe	−1.27	−0.72	−0.60	−0.69	−0.75
std	29.99	98.86	117.57	67.51	52.98
median	−18.05	−35.73	−17.43	−18.80	−18.70
skewness	−0.70	−0.38	−2.79	−1.81	−0.22

This table shows descriptive statistics of variance swap payoffs for a number of FX trading strategies and holding horizons τ (in months). A payoff is defined as in eq. (10), in percent-squared per year. Observations are daily and overlapping. Standard errors of the mean, calculated with the Newey and West (1987) adjustment with the number of lags set to $22 \times \tau$ are reported in parentheses below respective values. Sharpe ratios are annualized. Definitions of the strategies as in Section 2.6. The sample period is from 01/2009 to 06/2018.

Table 3: Risk premium estimates in the cross-section of "good and bad" carry trades

	(1)	(2)	(3)	(1)	(2)	(3)
		(A) 1 month			(B) 2 months	
rx_{CAR}	7.33		7.37	6.66		6.79
	(2.31)		(2.37)	(2.22)		(2.33)
rx_{USD}	−2.03		−1.87	2.12		−0.49
	(−0.45)		(−0.38)	(0.37)		(−0.09)
vsp_{CAR}		−38.74	−36.00		−45.30	−39.09
		(−2.09)	(−0.78)		(−2.39)	(−1.29)
vsp_{USD}		9.44	5.82		−0.50	0.29
		(0.92)	(0.67)		(−0.04)	(0.03)
		(C) 3 months			(D) 4 months	
rx_{CAR}	6.30		6.43	6.07		6.06
	(2.14)		(2.35)	(2.10)		(2.23)
rx_{USD}	1.51		−1.28	1.40		−0.14
	(0.25)		(−0.19)	(0.23)		(−0.02)
vsp_{CAR}		−53.17	−48.29		−53.87	−37.89
		(−2.27)	(−1.58)		(−2.07)	(−1.15)
vsp_{USD}		−7.97	−6.21		−12.22	−7.97
		(−0.46)	(−0.46)		(−0.65)	(−0.55)

Each panel of this table shows risk premium estimates and respective t-statistics from a set of linear asset pricing model. The models differ in terms of explanatory variables: in column (1), these are the US dollar index and Carry trade returns, in column (2) – the corresponding variance swap payoffs, and in column (3) – the four variables jointly. On the left-hand side are 1-, 2-, 3- and 4- month excess returns of the "good and bad" carry trade portfolios as in Bekaert and Panayotov (2018), the horizon increasing from panel A to D. Standard errors are calculated with the Newey and West (1987) adjustment (the number of lags set to $\times 22 \times$horizon). Returns are in percent p.a., variance swap payoffs are in currency per \$10000 notional. The sample period is from 01/2009 to 06/2018.

Figure 1: Preparing options data for integration

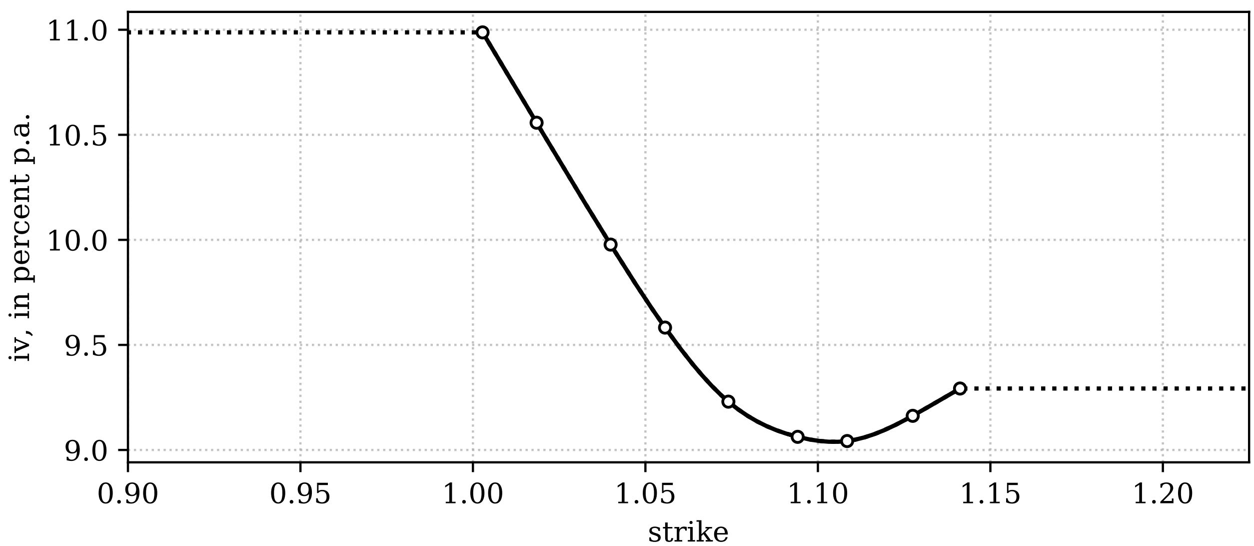

This figure shows the basic steps of inter- and extrapolating the observed option prices, before attempting the numerical integration of eq. (9). The hollow points correspond to the actually observed Black-Scholes implied volatilities of plain vanilla call options. The solid line depicts the cubic spline fitted to these volatilities to produce a smoothly interpolated smile. The smile is then extrapolated to the left and right of the observed strike range by keeping the respective endpoint volatilities constant, as shown by the dashed line. The example features call option prices extracted from the 10-, 15-, 25- and 35-delta risk reversals and butterfly spreads as well as the at-the-money option on EURUSD on 01/17/2017.

Figure 2: Portfolio constraints with AUDUSD and EURUSD.

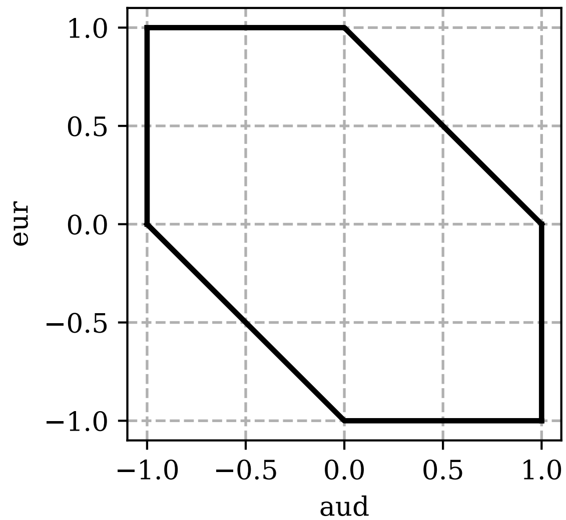

Figure 3: Eigendecomposition of variance swap payoff structure matrices

This figure shows the times series average of eigenvectors (left heatmap) and corresponding eigenvalues (right heatmap) calculated daily by performing eigendecomposition of the difference between two covariance matrices: the one realized over a 3-month period and the model-free option-implied one, observed at the start of the period (the observations of these matrix difference are thus daily and overlapping). The eigenvectors can be thought of as weights of USD exchange rates in portfolios, and eigenvalues – as portfolio variance swap payoffs. Values exceeding in magnitude 2× own standard error are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment with the automatic lag selection of Newey and West (1994). The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2009 to 06/2018.

Figure 4: Variance and variance risk of FX trading strategies

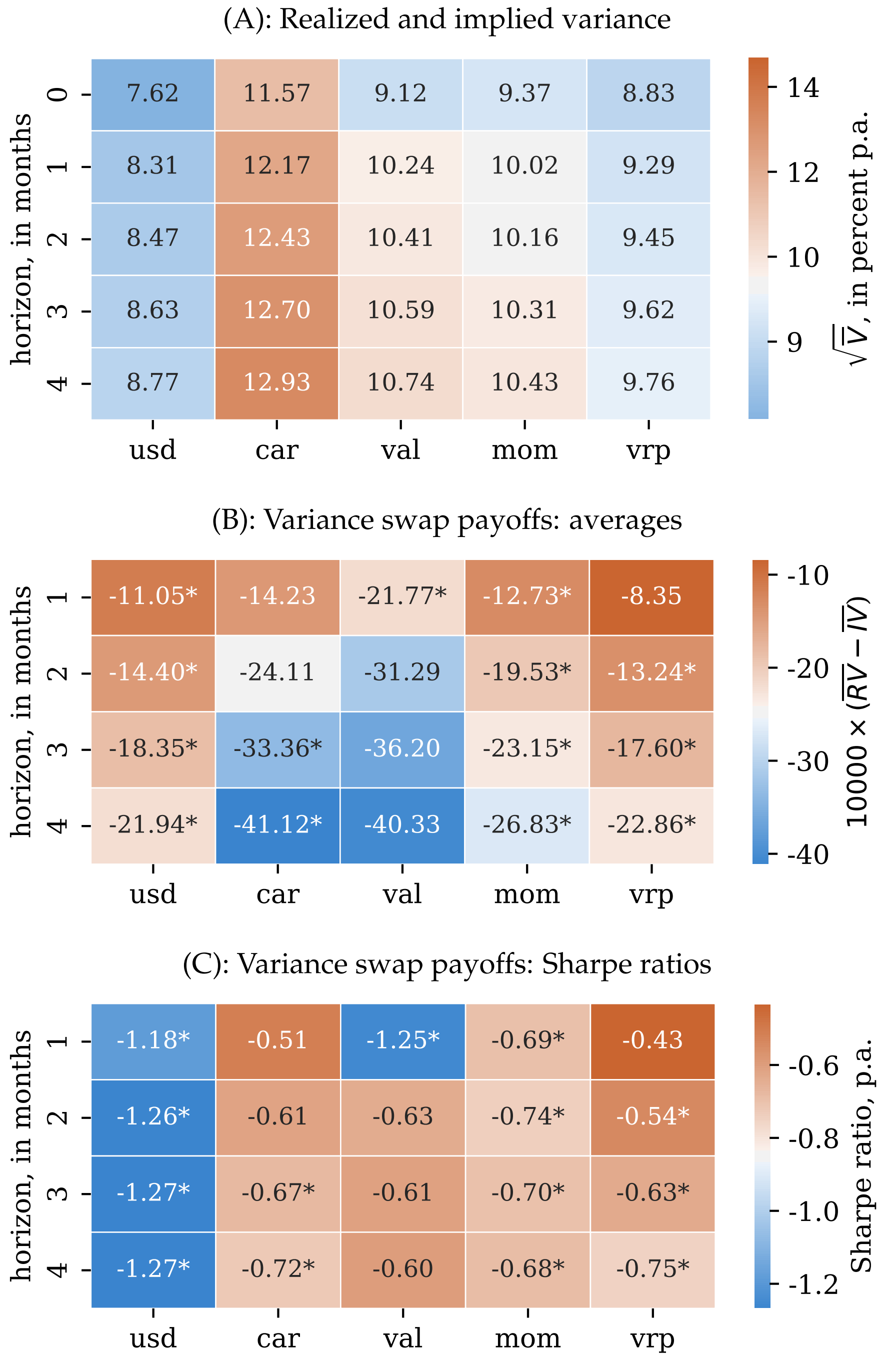

In this table, panel A shows the square root of the average realized (as horizon-0) and implied (for horizons of 1 to 4 months) variance of currency portfolios, in % p.a. The former are calculated as in eq. (7), the latter – as in eq. (9). Panel B shows the time-series averages of synthetic variance swap returns, calculated as in eq. (10), in percent squared p.a. Panel C shows the Sharpe ratios of these payoffs, annualized. In panels B and C the estimates with a t-statistic exceeding 2.0 in absolute value are marked with an asterisk. Observations are daily and overlapping; t-statistics are calculated with the Newey and West (1987) adjustment with the number of lags set to 22×horizon. The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2009 to 06/2018.

Figure 5: Variance swap returns of currency and local stock market indexes

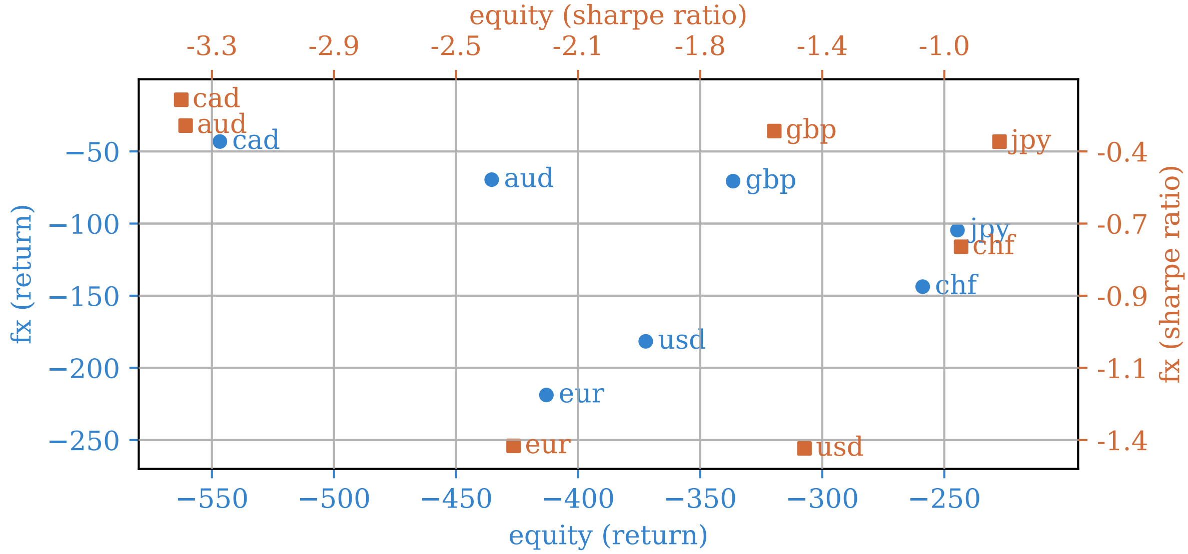

This figure depicts average 1-month variance swap returns (bottom x- and left y-axis, in blue) and respective Sharpe ratios (top x- and righ y-axis) for several local equity markets (on the x-axes) and corresponding currencies (on the y-axes). Returns are in percent, calculated as in eq. (22), whereby the FX variance swap rates are as in eq. (9), and the equity market swap rates are taken to be squared values of respective VIX indexes (details in Section 2.2. Both returns and Sharpe ratios are annualized. The sample period is from 01/2009 to 06/2018.

Figure 6: Forward variance claims prices: average values

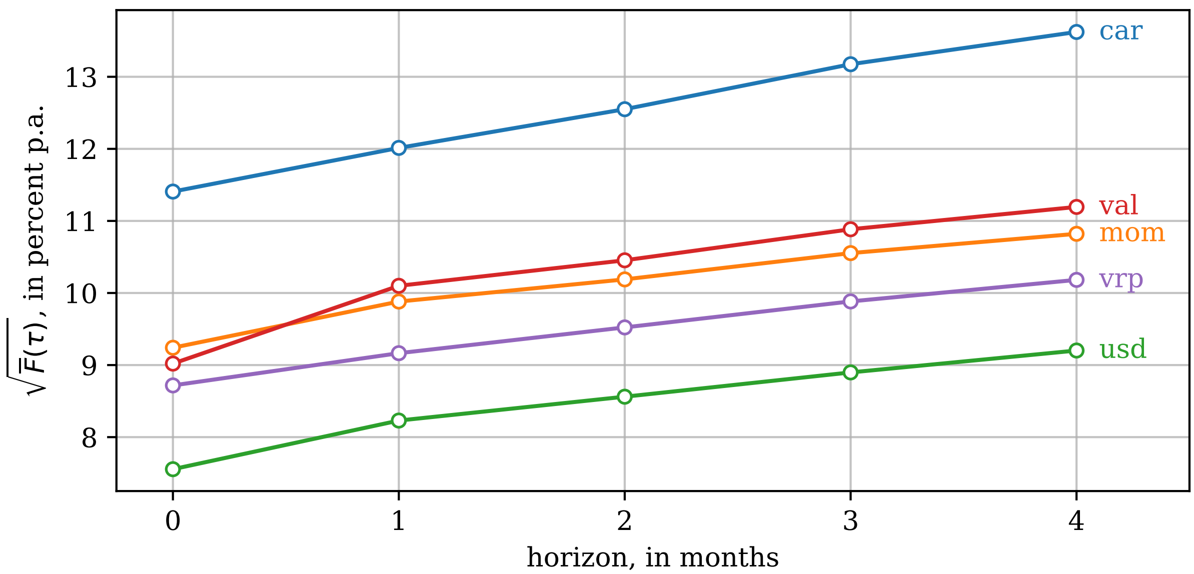

This figure shows the square root of the average τ-month forward variance prices for a number of currency portfolios (along the y-axis) and horizons (along the x-axis). Each value is thus a one-to-one increasing transformation of the average τ-month forward price of monthly variance. $F(0)$ is the average realized variance, and by definition, $F(1) = IV(1)$. Observations are daily and overlapping; one month is taken to be 22 days long. The sample period is from 01/2009 to 06/2018.

Figure 7: Rolling over forward variance claims: monthly variance

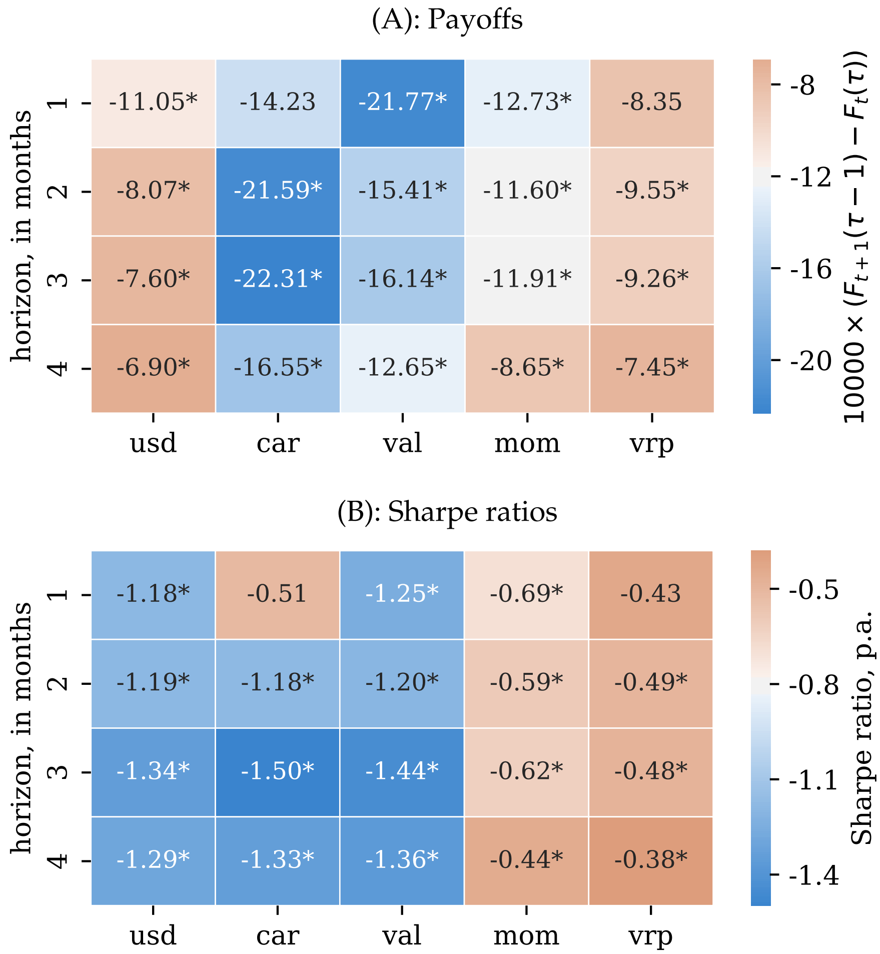

Panel A shows average monthly payoffs of rolling over forward monthly variance claims, and Panel B – respective Sharpe ratios. Horizon-1 claims are equivalent to 1-month variance swaps. Values exceeding in magnitude $2\times$ own standard error are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment (the number of lags set to $2 \times 22 \times$months). The colorbar to the right maps heatmap colors to numeric values. Observations are daily and overlapping; one month is taken to be 22 days long. The sample period is from 01/2009 to 06/2018.

Figure 8: Rolling over forward variance claims: quarterly variance

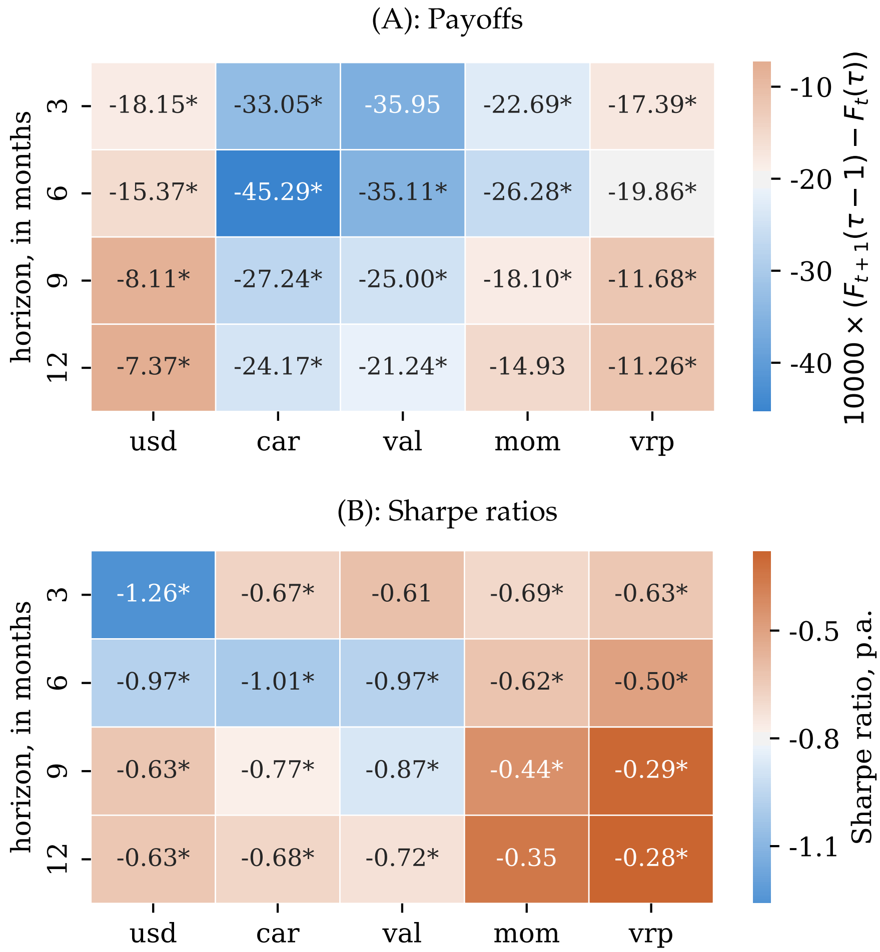

Panel A shows average quarterly payoffs of rolling over forward quarterly variance claims, and Panel B – respective Sharpe ratios. Horizon-3 claims are equivalent to 3-month variance swaps. Values exceeding in magnitude two own standard errors are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment (the number of lags set to $2 \times 22 \times$months). The colorbar to the right maps heatmap colors to numeric values. Observations are daily and overlapping; one quarter is taken to be 66 days long. The sample period is from 01/2009 to 06/2018.