

# Variance risk on the FX market\*

Igor Pozdeev<sup>†</sup>

this version: August 2018

## Abstract

I use model-free implied variances and absence of triangular arbitrage on the FX market to recover the risk-neutralized conditional covariance matrices of currency returns and construct synthetic swaps on the variance of the carry trade portfolio and a number of currency indexes. Extending previous research about variance risk premium to the FX setting, I document both familiar and novel results. I show that implied variance of FX portfolios bears non-redundant predictive information for their future realized variance. In contrast to the stock market, the FX market exhibits significant heterogeneity in variance risk premium which ranges from significantly negative for USD and EUR to positive for the carry trade, AUD and CHF. Bridging the equity and FX variance risk premia, I document that the two co-move non-trivially, suggesting their common driver, but that at least a two-factor framework is necessary to explain observed cross-asset differences. Variance risk premium is found to be a powerful predictor of spot returns of currency portfolios including the carry trade. The term structure of forward variance is upward-sloping for most FX portfolios, especially at longer horizons; similar to results for the US stock market, the curves flatten out beyond maturity of 4-6 months, suggesting that the shocks to short-term uncertainty are priced much stronger.

---

\*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 (2017) for their helpful comments and insightful feedback.

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

# 1 Introduction

Stochastic variance is considered with little disagreement an inherent feature of asset dynamics on different markets and a major source of risk for investors – the FX market being no exception. The associated variance risk premium (*VRP*) – a quantity investors are ready to forgo on average to be reimbursed when variance is high – is a concept not only intuitive, but also empirically observed and recently embedded in influential asset pricing models. However, while *VRP* on the equity market has been studied both on the individual and aggregate level, namely for stocks and stock indexes (which are themselves traded assets and watched benchmarks), similar research for the FX market has been disconnected from popular currency portfolios such as the carry trade and US dollar index – despite their importance for academic asset pricing and investment management. One reason for this disparity has been reliance of the relevant econometric apparatus on availability of option-implied information about the asset of interest: obviously, while options on stock indexes are ubiquitous and extensively represented in financial datasets, options on the carry trade portfolio are hard to come by. I circumvent the problem by using a cross-section of options on all possible cross-rates of currencies that enter an FX portfolio, which makes construction of implied covariance matrices possible and quantifying *VRP* of any linear combination of currencies – straightforward. I use this result to study variance risk of the carry trade and various currency indexes: the former is arguably the most popular FX investment strategy, and the latter naturally represent the currency exposure of an internationally diversified investor. By doing so, I contribute to the growing body of asset pricing literature on variance risk in the FX market and the risk profile of currency trading strategies, as well as to the econometric literature on recovery of option-implied information and information content thereof.

Variance risk premium for holding an asset is most often studied through the relation between the risk-neutralized and objective expectations of its future realized return variance *RV*:

$$VRP(\tau) = E^{\mathbb{P}} [RV(\tau)] - E_t^{\mathbb{Q}} [RV(\tau)], \quad (1)$$

where  $\tau$  is the horizon that the variance is calculated over, and  $\mathbb{P}$  and  $\mathbb{Q}$  denote the objective and risk-neutralized measure respectively. The  $\mathbb{Q}$ -expectation of future  $RV$ , also called the (option-) implied variance  $IV$ , can be thought of as a fair value to swap the future realized variance for; under certain assumptions, it can be recovered from a cross-section of option prices on the stochastic payoff. The  $\mathbb{P}$ -expectation of future  $RV$  can be obtained from a plethora of models of various sophistication. When  $E^{\mathbb{P}}[RV]$  falls short of  $IV$ , a variance swap buyer would be losing money on average, suggesting that protection against rising variance is costly, or that variance risk premium is negative. Negative  $VRP$  like this 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)), commodities (Tee and Ting (2017)) and individual currency pairs (Della Corte et al. (2016)), meaning that receivers of returns from the aforementioned investments are eager to overpay for assets that provide protection against their rising variance. In this paper, I document evidence of *positive* variance risk premium of the (USD-neutral) carry trade strategy. The carry trade is a dynamic strategy that goes long currencies with high and short those with low interest rate difference to USD; for decades, it has enjoyed elevated profitability and attracted much attention of investors and academics alike. The profitability has been partially attributed to the overall riskiness of the strategy: its high volatility and GARCH-in-mean features have been previously proposed as a risk-based explanation for the 6% mean return and a Sharpe ratio of 0.6. For instance, Menkhoff et al. (2012) document the negative variance-return relation in the long leg of the carry portfolio and the positive one in the short leg, meaning that the portfolio performs relatively poorly when the aggregate FX market variance rises, and relatively well in the opposite case. My results thus challenge the variance risk-based explanation, as I present evidence that the cost of hedging this risk has been but too low. A positive  $VRP$  is not artifact of a peculiar FX option market structure: I show that currency indexes such as the USD index and EUR index do exhibit statistically significant negative premium that is comparable in magnitude to the premium for corresponding stock market indexes. That said, my findings add to the puzzle set forth 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.

I further show that variance risk for currency indexes co-moves non-trivially with that for the corresponding stock market indexes, the correlation ranging from 0.22 in the case of USD to 0.58 in the case of GBP. However, cross-asset differences in *VRP* estimates are not aligned, and at least a two-factor model is needed to explain them. Interestingly, while equity *VRP* estimates are strongly correlated with each other, the same is not true for FX *VRP* estimates, suggesting a richer within-class dependence structure of the FX variance risk.

Previous research about variance risk premium contains evidence of its significant predictive power for future returns. [Bollerslev et al. \(2009\)](#) and [Londono \(2015\)](#) document this for the stock market *VRP* and stock returns, and [Londono and Zhou \(2017\)](#) bring currency returns into the picture. I extend these results by showing that subsequent spot appreciation rates of currency portfolios, including but not limited to the carry trade, are predictable by the current FX variance risk premia estimates. When contrasted with one-factor predictive regressions featuring portfolio forward discounts, adding *VRP* as the second factor improves the  $R^2$  many-fold for most portfolios considered. I also find that implied variance that enter the calculation of *VRP* possesses non-redundant predictive power for the future realized variance. By running a horse race between several commonly used variance forecasting model, I discover that the currently observed *IV*, when taken as a pure variance forecast, tends to beat ARIMA(1,1,1), GARCH(1,1) and the commonly used martingale benchmark in terms of the two robust loss functions studied by [Patton \(2011\)](#), for most currency portfolios and horizons.

Finally, I extend the results of [Dew-Becker et al. \(2017\)](#) who find that of all the equity market uncertainty  $\tau$  periods into the future, only the shortest-term ( $< 4$  months) segment appears to be priced, as historical Sharpe ratios of rolling over more distant forward variance claims are significant before and quickly go to zero beyond this threshold. Hence, hedging news about future variance at longer horizons has been costless on average. The authors note that many financial and economic models, e.g. those of [Drechsler and Yaron \(2011\)](#) and [Bloom \(2009\)](#) imply that shocks to *expected* future variance are a priced source of risk, a fact mostly

confuted by the data. I show that for many currency indexes, especially for USD, EUR and the carry trade, the term structure of forward variance claims appears “too flat” beyond the same maturity of 4 months suggesting that FX market investors share the aversion to shocks about short-term and tolerance to shocks about long-term uncertainty.

My work belongs to and draws from the relatively young literature on the FX market variance risk. [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 how much the  $IV$  of each differs from the objective conditional variance. Over-weighting (under-weighting) currencies with high (low) difference, the authors construct a strategy that enhances the mean-variance investment opportunities on the FX market. [Ammann and Buesser \(2013\)](#) find further supporting evidence for negative variance risk premium in individual exchange rates. However, the variance risk of the carry trade and the dollar index has not received similar attention, although these portfolios are the two dominant priced factors in the cross-section of currency returns ([Lustig et al. \(2011\)](#), [Verdelhan \(2015\)](#)) and popular investment vehicles ([Curcuru et al. \(2011\)](#)). Moreover, making claims about the FX market  $VRP$  by studying individual exchange rates might miss the big picture because their covariance structure is not taken into account. I demonstrate my claim in Table 1 by replicating the strategy of [Della Corte et al. \(2016\)](#) and calculating the difference between its  $RV$  and  $IV$ . Additionally, using the implied and realized covariance matrices of currency returns, I calculate the maximum and minimum attainable difference between  $RV$  and  $IV$ , which is under some assumptions the *ex ante* maximum and minimum attainable variance risk premia. These turn out to be higher and lower respectively than that of the above strategy, such that repeating the sorting exercise of [Della Corte et al. \(2016\)](#) in the mini-universe of these three strategies would result in the long-short strategy rarely being actually selected.

[Table 1 about here.]

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 strate-

gies. To my knowledge, this is the first paper to do so. 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 USD: 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.<sup>1</sup> 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. Section 3 presents the findings. Section 4 concludes.

## 2 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 “returns” defined as  $R_x(t, t + 1) = S_x(t + 1)/S_x(t)$  are thus by default appreciation rates against USD. The USD index – a weighted

---

<sup>1</sup>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.

average of currency returns – is higher when the US dollar depreciates against a basket of currencies, but the same clearly applies to other currency indexes: the CHF index rises when the Swiss franc loses value to other currencies etc.

When explicitly written with a double subscript  $S_{xy}$ , exchange rate against currency  $y$  is meant rather than against USD, 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 one-period log-return of currency  $x$  is defined as:

$$r_x(t, t+1) = \Delta s_x(t+1) = s_x(t+1) - s_x(t). \quad (2)$$

## 2.2 Data

I use FX option data from Bloomberg. Here, 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 ( $\delta$ ) rather than against strike prices, forward delta<sup>2</sup> 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 usually multiplied by 100 for quotation purposes to become numbers such as 25, 15 etc.

---

<sup>2</sup>As provided by Bloomberg, delta is net of currency premium.

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<sup>3</sup>. 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 stamped with 17:00 New York time. The contracts are conveniently structured to be of constant maturity of 30, 60, 90, 120 days.

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 \(2014\)](#):

$$\sigma(C(\delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) + 0.5\sigma(BF(\delta)), \quad (3)$$

$$\sigma(C(1 - \delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) - 0.5\sigma(BF(\delta)) \quad (4)$$

where  $\sigma(Y(\delta))$  is the IV of contract of type  $Y \in \{ATM, BF, RR, C, P\}$  having delta  $\delta$ . 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 \(2009\)](#), (3)-(4) 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.

---

<sup>3</sup>Sometimes certain quotes would be missing on a given day; most often these are the less liquid 10- and 35-delta ones.



That said, it is possible to obtain  $2N + 1$  distinct  $(\delta, \sigma)$  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 [Bank of International Settlements \(2016\)](#). For instance, in the case of AUDCHF, the true AUD OIS rate will be taken and used to infer the CHF risk-free rate, although the CHF OIS rate is also available.

Local-economy stock indexes used for construction of equity variance risk premia estimates 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<sup>4</sup> the conditional variance of an arbitrage-free process to the conditional expectation of its quadratic variation.

---

<sup>4</sup>Under certain yet not implausible assumptions discussed therein.

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

$$r(t) = ds(t) = \lim_{\Delta t \rightarrow 0} s(t + \Delta t) - s(t). \quad (5)$$

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

$$[r, r]_{t, t+\tau} = \int_t^{t+\tau} (ds(k))^2, \quad (6)$$

and is closely related to the variance of the process:

$$\text{Var}_t(r(t, t + \tau)) = E_t([r, r]_{t, t+\tau}) = E_t \left( \int_t^{t+\tau} (ds(k))^2 \right), \quad (7)$$

The same Corollary suggests that a natural *ex post* estimator of the variance on the left-hand side of eq. (6) – the quantity the investor expects to receive – is the discretized version of the right-hand side<sup>5</sup>:

$$\text{Var}_{t+\tau}(r(t, t + \tau)) = \frac{254}{D_\tau} \sum_{d=1}^{D_\tau} r \left( t + \frac{d-1}{D_\tau}, t + \frac{d}{D_\tau} \right)^2 = RV(t, t + \tau), \quad (8)$$

where  $t + d/D_\tau$  denotes day  $d$  of time interval  $(t, t + \tau)$ ,  $D_\tau$  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. (8) 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 (8) have been favored in related econometric literature.

The fair price to swap the quantity in eq. (6) for (such that no money changes hands at time  $t$ )

---

<sup>5</sup>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. \(2016\)](#).

is by the standard argument its risk-neutralized expectation:

$$IV_t(\tau) = E_t^Q \left( \int_t^{t+\tau} (ds(k))^2 \right), \quad (9)$$

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 [Breedon and Litzenberger \(1978\)](#), equate<sup>6</sup> the conditional Q-expectation of the accumulated quadratic variation to the price of a continuous portfolio of options expiring in  $((t + \tau))$ , weighted by strikes:

$$E_t^Q \left( \int_t^{t+1} (ds(k))^2 \right) = 2 \int_0^\infty \frac{C_t(\tau, X) - (S(t) - X)^+}{X^2} dX, \quad (10)$$

where  $C_t(\tau, X)$  is the time- $t$  price of a call option on  $S$  with strike  $X$  and maturity of  $\tau$ , and  $(v)^+ = \max(v, 0)$ . The integral in eq. (10) can be evaluated numerically, but a careful inspection reveals two potential sources of errors in doing so: first, while the integration in eq. (10) 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<sup>7</sup>, 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.<sup>8</sup>

<sup>6</sup>Under assumptions that were relaxed by [Jiang and Tian \(2005\)](#) to include jump-diffusion processes.

<sup>7</sup>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.

<sup>8</sup>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.

[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:

$$vs(t, t + \tau) = RV(t, t + \tau) - IV_t(\tau), \quad (11)$$

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 in (11):

$$VRP_t(\tau) = E_t(vs(t, t + \tau)) = E_t(RV(t, t + \tau)) - IV_t(\tau), \quad (12)$$

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 hedging 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<sup>9</sup> dynamic trading strategy, rebalanced monthly and represented with a vector of weights  $(w_{1,t}, w_{2,t}, \dots, w_{N,t})'$  known at the end of the previous month. For instance, any position involving AUD, EUR and USD can be represented with a  $2 \times 1$  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

---

<sup>9</sup>On the ForEx, it would mean that having \$1 in the margin account it is only possible to open \$1 worth of positions.

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

$$v_{t+1} = (w_{1,t}, w_{2,t}, \dots, w_{n,t})', \quad (13)$$

$$E_t[v_{t+1}] = v_{t+1}, \quad (14)$$

$$\sum_{n=1}^N |w_{i,t}| + \left| \sum_{n=1}^N w_{i,t} \right| = 2, \quad (15)$$

where eq. (14) says that the composition of the portfolio at time  $t+1$  is known at time  $t$ . The constraint in (15) 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  $w_t = E_t[v_{t+1}] = (w_{1,t+1}, w_{2,t+1}, \dots, w_{N,t+1})'$ . 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:

$$f_{t+1} = w_t' r_{t+1}, \quad (16)$$

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

$$Var_t(f_{t+1}) = w_t' \Omega_t w_t, \quad (17)$$

where  $\Omega_t$  is the conditional covariance matrix of currency returns against USD.

## 2.5 Covariance matrices of currency returns

The cornerstone of this paper is the conditional covariance matrix of time- $(t + 1)$  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:

$$\Omega_t[x, y] = \begin{cases} \text{Var}_t(r_{x,t+1}), & x = y, \\ \text{Cov}_t(r_{x,t+1}, r_{y,t+1}), & x \neq y \end{cases}$$

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):

$$r_{xy} = r_x - r_y. \quad (18)$$

Applying the variance operator to both sides of eq. (18) results in:

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

which can be rearranged to isolate the covariance as follows:

$$\text{Cov}(r_x, r_y) = \frac{1}{2}(\text{Var}(r_x) + \text{Var}(r_y) - \text{Var}(r_{xy})). \quad (19)$$

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 Carry trade

Following [Lustig et al. \(2011\)](#), at time  $t$  I select two currencies with the highest, and two – with the lowest forward discount to the US dollar. Long positions are opened in the former, and short position – in the latter. A zero-cost strategy, it is often implemented on the subset of G10 currencies and has a historical Sharpe ratio of 0.6, an average return of 6% p.a., and billions of US dollars allocated to it, thus constituting an important asset class and asset pricing phenomenon. A common representation of the carry trade strategy is as follows (the counter currency is assumed to be the USD dollar):

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

### 2.6.2 Currency indexes

A return of a currency index is defined as the 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: for instance, for USD (the counter currency is assumed to be the USD dollar):

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

### 2.6.3 Extreme attainable variance risk premium strategy

At each time  $t$ , the maximum (minimum) attainable variance risk premium portfolio is portfolio  $w_t$  with two currencies in each the long and short leg, having the highest (lowest) possible

value of:

$$RV(t - \tau, t|w_t) - IV_t(\tau|w_t),$$

where  $RV(t - \tau, t)$  is the most recently realized variance of the portfolio return as the proxy for the expectation in (12), and  $IV_t(\tau)$  is the model-free implied variance of the return, both calculated from the option-implied covariance matrix of currency returns at time  $t$ .

## 3 Results

### 3.1 Realized and implied variance of currency portfolios

Figure 3 depicts the square root of the time-series averages of currency portfolios' realized (denoted as horizon-0) and implied variances. The carry trade portfolio is to little surprise the most volatile at all horizons, not only by the average realized, but also by the implied measure ( $14\% \pm 0.5\%$  p.a. depending on the horizon). The JPY index is the second most volatile, with 2.5% lower  $RV$  and 1% lower  $IV$  than that of the carry. Further down the variance level are the other typical carry portfolio constituents, namely NZD and AUD. The rest of the currency indexes have far more moderate variance estimates, falling short of the carry trade portfolio by 5-8%.

For many portfolios,  $IV$  tends to exceed  $RV$ , signaling a negative variance risk premium to be discussed in Section 3.2.

[Figure 3 about here.]

$IV$  is forward-looking by construction, but differs from the expected variance by the variance risk premium. Depending on whether the former virtue outweighs the latter flaw,  $IV$  can be used as a variance forecasting model. I compare it to several other popular models in application to variance of currency indexes and the carry trade portfolio. To evaluate the relative



forecasting performance, I use the squared error and QLIKE loss function (Patton (2011)):

$$\text{SE} : L(h, RV) = (RV - h)^2, \quad (20)$$

$$\text{QLIKE} : L(h, RV) = \log h + \frac{RV}{h}, \quad (21)$$

where  $RV$  is the realized variance and  $h$  is the variance forecast, assumed to be made for the same period that  $RV$  is sampled over. Different loss functions allow to obtain a broader picture of models' forecasting power, as the squared loss is symmetric and penalizes large forecast errors more heavily, whereas the QLIKE loss is asymmetric, less sensitive to large forecast errors and penalizes under-predictions more than over-predictions.

In addition to the pure  $IV$ , I use conditional variances from the ARIMA(1,1,1) and GARCH(1,1) models, as well as the lagged realized variance. ARIMA(1,1,1) is inspired by the ARIMA(0,1,3) of French et al. (1987) and found to have the best in-sample fit of all models of this class; it is calibrated to log-variances, and forecasts are made along the lines of the above paper, with Jensen's inequality accounted for. Lagged values of realized variance as a proxy for the expected variance have been used in numerous studies, e.g. in Moreira and Muir (2017), Della Corte et al. (2016) and Carr and Wu (2009). Favoring this model is tantamount to treating the variance process as a martingale. GARCH(1,1) of Engle (1982) and Bollerslev (1986) is a powerful model capturing essential facts about the variance dynamics. Economic logic suggests there is little room for asymmetric model specifications, because for currencies,  $r_{ij} = -r_{ji}$  and  $Var(r_{ij}) = Var(r_{ji})$ . Forecasts from all models are made in an out-of-sample fashion and refer to the portfolio composition looking forward (which is only relevant for the dynamically rebalanced carry trade): at time  $t - 1$ , portfolio weights for time  $t$  are obtained and the historical daily returns of the portfolio are calculated, then a model is calibrated to these returns, and finally  $\tau$ -period ahead forecasts are made. Details on ARIMA and GARCH models are given in Appendix A.

Figures 4A and 4B quantify the relative forecasting power of  $IV$  and the other models. In terms of the QLIKE loss,  $IV$  is found to outperform the other forecasts for most currency-horizons

with the horizon up to 6 months, for a total of 39 out of 45 such currency-horizons. In terms of the RMSE, *IV* is chosen as the preferred forecasting model for a total of 30 out of 45 such shorter currency-horizons. Curiously, at longer horizons (9-12 months) GARCH and lagged variance tend to do relatively better.

[Figure 4A about here.]

[Figure 4B about here.]

Treating the MFIV in the raw form as a variance forecasting model, I compare its absolute predictive power to those of the other models by estimating the following univariate regressions:

$$RV(t, t + \tau) = \alpha_i + \beta_i h_i(t, t + \tau) + \epsilon_{t+\tau}, \quad (22)$$

where  $h_i(t, t + \tau)$  is a forecast of the variance realized from  $t$  to  $t + \tau$  made from model  $i$  at time  $t$ . Additionally, as a check for efficiency, I estimate the multivariate regression with all considered forecasts on the right-hand side:

$$RV(t, t + \tau) = \alpha + \sum_i \beta_i h_i(t, t + \tau) + \epsilon_{t+\tau}. \quad (23)$$

Table 2 reports the estimates for the USD index and the carry trade portfolio for  $\tau = 1m$ . The MFIV enters univariate regressions with a beta closest to one and alpha closest to zero than any other model. Importantly, it turns out to be the only variance forecast with a significant beta in either joint regression, whereby the betas are not as far from one as those of the other forecasts (compared e.g. to the negative ARIMA betas). That is, the MFIV is an efficient variance forecast, in the sense that it subsumes the information available to the econometrician through time-series models.

[Table 2 about here.]

### 3.2 Time series of variance swap returns

On each trading day  $t$ , I compute the observed  $\tau$ -month implied variance  $IV_t(\tau)$  of a currency portfolio and the  $\tau$ -month-ahead realized variance  $RV(t, t + \tau)$  of the portfolio return. The difference between the two is the payoff of a variance swaps (VS) with maturity  $\tau$ , entered into on day  $t$  (see eq. 25). In addition to this payoff, and as a robustness check, I calculate the return on a fully collateralized swap position by normalizing the payoff by  $IV_t$ . For better readability, I multiply the payoff with 10000 and return with 100, and rescale both values to the monthly frequency:

$$\pi(vs; t, t + \tau) = \frac{1}{12\tau} \times (RV(t, t + \tau) - IV_t(\tau)) \times 10000, \quad (24)$$

$$r(vs; t, t + \tau) = \frac{1}{12\tau} \times \frac{RV(t, t + \tau) - IV_t(\tau)}{IV_t(\tau)} \times 100, \quad (25)$$

where  $\tau$  is in years.

Variance swap payoffs are closely related to variance risk premium, as a closer look at eq. (12) conveys. If expectations of future realized variance are unbiased and the realized variance process is stationary<sup>10</sup>,  $VRP$  would coincide with the time-series average  $r(vs; t, t + \tau)$ . With that in mind, the latter are sometimes called the *ex post* variance risk premium, and used to draw inference about the magnitude of  $VRP$ . Table 3 presents select descriptive statistics of the distribution of variance swap returns across horizons and portfolios, and Figure 5 visualizes the term structure of the averages values.

[Figure 5 about here.]

[Table 3 about here.]

The figure reveals, that unlike the significantly negative variance risk premium estimates for local stock markets documented elsewhere,  $VRP$  of currency indexes can appear not only sim-

---

<sup>10</sup>Which is admittedly a rather restrictive assumption.

ilarly significantly negative, e.g. for USD and EUR, but also moderately so, e.g. for CAD, GBP and NZD, as well as largely (yet not significantly) positive, e.g. for AUD and CHF. Interestingly, variance risk premium of the carry portfolio is also large and positive: the potential gain from hedging the strategy's variance is strictly greater than zero and ranges from \$23 per \$10000 notional for the 1-month horizon (or alternatively, 17% on the fully collateralized position) to \$3 for the 12-month horizon (3%). Although these estimates are not reliably different from zero, they dwarf those of other currency portfolios in magnitude, and their sheer non-negativity appears as a puzzle for a strategy as (superfluously) risky as the carry trade.

However, conclusions drawn from positive *average* variance swap returns are challenged to some extent by the returns' *median* values, as can be seen in Table 3. Distribution of variance returns is positively skewed and has a negative median irrespective of the sign of the mean. The carry trade portfolio is no exception, although it is still far from being the most expensive currency portfolio to variance-hedge. The median variance swap return of the carry trade ranges from  $-16\%$  at the shortest to  $-2\%$  at the longest horizon, with only CAD and GBP indexes having less negative estimates.

The strong disparity between the mean and median values of variance swap returns reflects the fact that occasional unexpected jumps in realized variance leading to large positive returns are not "matched" by an implied variance that would be sufficiently high on average. For several currency portfolios in 5, the magnitude of these jumps is large enough to turn the time series estimate of the risk premium positive, even though if not for them,  $IV$  would be found to exceed  $RV$ .

Finally, it is worth turning attention to the Sharpe ratios of variance swap returns. It has been previously documented for other asset classes ([van Binsbergen and Koijen \(2017\)](#), [Dew-Becker et al. \(2017\)](#)) that Sharpe ratios of returns decline in magnitude with investment horizon. Figure 6 show that the same is not true for the FX market: neither the cross-sectional average of (absolute) Sharpe ratios of variance swap returns for individual FX indexes, nor the Sharpe ratio of the carry trade variance swap returns is decreasing with the horizon. Table 3 further

conveys that this is happening because the standard deviation of returns tends to decrease in magnitude faster than the mean.

[Figure 6 about here.]

### 3.2.1 FX and stock market variance risk premium: a closer look

In order to draw a link between the stock and FX market variance risk premium, in Figure 7 I plot the average swap returns for stock market indexes denominated in the local currency against those for the respective currency index. There is little visible association between the  $VRP$  values across the two asset classes, which is also confirmed by the regression of average  $VRP$  estimates onto each other:

$$VRP_{fx} = 8.4069 + 0.2380 VRP_{eqt},$$

(8.64)            (0.30)

$$R_{adj}^2 = -0.16,$$

where standard errors of respective coefficient estimates are in parentheses. Hence, a one-factor model aspiring to explain the observed international differences in stock market  $VRP$  won't be able to explain the differences in FX  $VRP$ , and vice versa. A potentially successful model would need to feature at least two factors, as two distinct clusters of currencies form in Figure 7: CAD, AUD and CHF on the one hand and EUR, USD, GBP and JPY on the other, with a positive linear relation within each cluster.

[Figure 7 about here.]

To further relate variance risk premium of currency portfolios to that of stock indexes, I compute *daily* estimates of  $VRP$  by using GARCH(1,1) conditional variances for the  $RV$  in eq. (12). The previously deployed martingale proxy, although unbiased over the whole sample, makes

little sense on the day-by-day basis: by construction, it fails to react sufficiently quickly to return innovations and thus is a poor proxy for the  $\mathbb{P}$ -expected variance, a claim also supported by Figures 4A-4B. That said, Figure 8 shows the correlation structure of *VRP* estimates within each of the two asset classes as well as between them. Cell  $(i, j)$  in the part below the main diagonal shows the correlation between *VRP* estimates of currency portfolio  $i$  and  $j$ , whereas the part above – between *VRP* estimates of stock market index  $i$  and  $j$ . Stock market *VRP* estimates are in general more correlated with each other than currency portfolio *VRP* estimates. CHF is the least correlated with other currency indexes – all pairwise correlations are below 0.1 in magnitude – while EUR and CAD are the most correlated. The carry trade *VRP* is strongly and positively correlated with most currency indexes except for the USD. Cell  $(i, i)$  on the main diagonal (outlined with thick black borders) contains the correlation between *VRP* of stock market  $i$  and *VRP* of the corresponding currency index. Interestingly, although the correlations among currency index *VRP* estimates are far from large, each of the estimates is strongly correlated with *VRP* of “its” stock market index, the strongest association appearing for GBP and JPY (0.53).

[Figure 8 about here.]

### 3.3 Predicting returns of currency portfolios with *VRP*

Several papers document variance risk premium-induced predictability of returns in different asset classes. [Bollerslev et al. \(2009\)](#) show that S&P500 returns are predictable by the 1-month US stock market variance risk premium, whereby the strongest results are found for horizons of 2-5 months (that is, regression of returns over the next 2-5 months onto the currently observed *VRP* exhibits the largest  $R^2$ ). Evidence presented by [Londono \(2015\)](#) suggests that the US stock market *VRP* also predicts international stock returns, but at the same time, local stock market *VRP* has little to none predictive power for local stock returns. [Londono and Zhou \(2017\)](#) report that both the US stock market *VRP* and a cross-sectional average of international

stock market *VRP* have significant predictive power for the appreciation rates of currencies. In this section, I show that variance risk premium in currency index returns bears non-redundant predictive information for subsequent return dynamics, over and above a common benchmark such as the average forward discount in the portfolio.

To do so, I run two types of regression. First, as a benchmark, I regress  $\tau$ -month-ahead spot returns of a currency portfolio onto a constant and the portfolio forward discount<sup>11</sup>:

$$r(t, t + \tau) = \alpha + \beta d_t + \varepsilon(t, t + \tau), \quad (26)$$

where  $d_t$  is the log-forward discount of the basket of currencies. I then redo the exercise with the 1-month *VRP* added as the second regressor:

$$r(t, t + \tau) = \alpha + \beta d_t + VRP_t(1m) + \varepsilon(t, t + \tau), \quad (27)$$

whereby the IP-expectation of realized variance necessary to compute the *VRP* is obtained from the GARCH(1,1) model as in Section 3.1. I report the adjusted  $R^2$  values from the two regressions in Figures 9A and 9B respectively.

[Figure 9 about here.]

*VRP* raises predictability of spot returns a lot when compared to the benchmark forward discount-only specification. For EUR, GBP, JPY, NZD and USD,  $R^2$  increases for all horizons and maturities, sometimes as much as tenfold, once *VRP* is added as the second regressor. For AUD and CHF, this is only for the horizons of 1 to 4 months, and predictability gains are less pronounced. For CAD, no gain is detected except for the longest horizon. Measured by the  $R^2$ , predictability of the carry trade's spot return rises from 0.6 to 2.3 percent at the 1-month horizon, and from 2.0 to 6.9 percent at the 3-month horizon.

---

<sup>11</sup>Lustig et al. (2011) and Londono and Zhou (2017) document (limited) predictability of currency spot returns by forward discounts. Due to the covered interest parity, the forward discount is approximately equal to the average differential between the interest rate in the economy of the counter currency and that in the economy of the base currency.

### 3.4 Forward variance claims and news about future variance

Several 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 volatility (e.g. [Bloom \(2009\)](#), [Fernández-Villaverde et al. \(2011\)](#)) predict that investors are willing to hedge news about future uncertainty. To visualize the latter concept, imagine an investor in February 2016 when the date of the United Kingdom EU membership referendum date was set for June. The 4-month GBP variance swap rate gives her the price to hedge return variance of GBP accumulated from February to June, but she might instead wish to only hedge shocks to the expected June variance. To do that, she could buy the 4-month and sell the 3-month variance swap thus retaining sole exposure to the variance accumulated in June, as the return from such contract is simply:

$$\begin{aligned}
 r(vs; t, t + \tau) - r(vs; t, t + \tau - 1) &= \\
 RV(t, t + \tau) - IV_t(\tau) - RV(t, t + \tau - 1) + IV_t(\tau - 1) &= \\
 RV(\tau, \tau) + IV_t(\tau - 1) - IV_t(\tau), & \quad (28)
 \end{aligned}$$

where  $t$  indexes February, and  $\tau = 4$  such that  $t + \tau$  indexes June. Hence, by the risk-neutral pricing argument:

$$E_t^Q(RV(\tau, \tau)) = F_t(\tau) = IV_t(\tau) - IV_t(\tau - 1), \quad (29)$$

where  $F_t(\tau)$  is the forward price of the month- $(t + \tau)$  “variance strip”, or simply the forward variance price. It is the price to swap the variance accumulated in month  $t + \tau$  for. By rolling these strips over, the investor is effectively hedging shocks to uncertainty at the fixed horizon  $\tau$ ; to see why, assume that one month later the June forward variance price jumps up because the perceived uncertainty about the June referendum outcome increases – the investor sells her variance strip at a higher price and realizes a profit. Many model of the type mentioned above imply that in exchange for occasional profits in times of increased uncertainty about the future, the investor is ready to tolerate losing money on average. Surprisingly, [Dew-Becker et al. \(2017\)](#)



find the opposite is true on the stock market: investors (could) have been hedging news about future  $\tau$ -month S&P500 variance virtually for free for  $\tau > 3$ . On the other hand, hedging the shorter-term uncertainty as well as the realized variance itself (which can be thought of as the shortest uncertainty, with  $\tau = 1$ ) has been costly and resulting in high Sharpe ratios for the other side of the hedging transaction.

For each currency portfolio, I construct the time-series of forward variance prices as in eq. (29). To do so, I interpolate  $IV_t$  at the missing maturities (5, 7, 8, 10 and 11 months) using natural cubic splines calibrated to the cross-section at date  $t$ . By convention,  $F_t(0)$  denotes the *ex post* realized variance in month  $t$ <sup>12</sup>, and  $F_t(1)$  to denote the 1-month IV. Figure 10 shows the average term structure of forward variance, in the square root form.

[Figure 10 about here.]

Similar to the evidence from the stock market, the average slope of all curves is positive, although here it may be locally negative for some indexes, e.g. for AUD, CHF or the carry trade portfolio at the shortest segment. The steepest average slope (elevation of the 12-month IV over the 1-month IV) is observable for JPY (1.98% per year), and the steepest – for CAD (0.64%), meaning that hedging the distant variance of the JPY portfolio is relatively more costly than the more immediate variance, while the difference is minuscule for the CAD portfolio. The carry trade forward variance curve is downward sloping at the horizons of 1 and 2 months, and upward sloping elsewhere.

To estimate how costly it has been for an investor to hedge news about future uncertainty, I calculate returns of rolling over forward variance claims:

$$r_{t+1}(F(\tau)) = \frac{F_{t+1}(\tau - 1) - F_t(\tau)}{F_t(\tau)}, \quad (30)$$

where indexing refers to months<sup>13</sup>. Figure 11 shows Sharpe ratios of these returns.

<sup>12</sup>Since I deal with daily observations,  $F_t(0)$  is calculated over the 22 days following the observation day.

<sup>13</sup>Again, I construct daily overlapping returns with the convention that there are 22 days in a month.

[Figure 11 about here.]

For the USD and EUR indexes, the Sharpe ratio is strongly negative and statistically significant at the 1-month horizon ( $-1.12$  and  $-0.71$  respectively), a result mimicking that reported by [Dew-Becker et al. \(2017\)](#) for the US stock market. Beyond that, another economically significant spike is detected around the quarterly horizon, but the estimates cease to be significant altogether. With minor exceptions, the same pattern is evident for the other currency portfolios: if anything, only news about short-term uncertainty are priced.

For the carry trade portfolio, the familiar positive price of 1-month variance risk is surprisingly followed by a slightly positive price at the 2-month horizon, meaning that rolling over 2-month forward variance swaps and thus hedging news to uncertainty at that horizon has been costless on average as well. Similarly to the USD and EUR indexes, the Sharpe ratio gets negative around the quarterly horizon, then levels off in the vicinity of zero.

## 4 Conclusions

FX trading strategies, such as the carry trade, and currency portfolios, such as the dollar index, are at the heart of FX market investment management and academic research, yet have been eluding studies about 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 variance swaps for these strategies and study their variance risk.

I find evidence of significant negative variance risk premium for holding the diversified EUR- and USD-denominated currency portfolio, and statistically insignificant yet large positive risk premium for being invested in the carry trade. Given the high riskiness of the latter strategy, the finding appears as a puzzle. The reason for apparently positive premium estimates is a small number of variance shocks followed by huge positive variance swap payoffs; these shocks are not fully priced in.

The term structure of forward variance prices is upward-sloping for all currency portfolios and most horizons, but beyond the relatively short ones it levels off, such that – if anything – only the short-term variance risk bears a price.

Conditional FX variance risk premium estimates are weakly correlated across currency zones, but strongly so with the estimates for local stock markets. The joint distribution of unconditional premia however leaves little room for explaining cross-economy differences in a single-factor framework.

Two of my findings concern predictability: I show that implied variance of currency portfolios has non-redundant predictive power for the future realized variance, and that conditional variance risk premium estimates forecast future spot returns of currency portfolios, improving in-sample forecast accuracy over a standard benchmark.

## References

- Ammann, Manuel, and Ralf Buessler, 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.
- Bank of International Settlements, 2016, Triennial central bank survey of foreign exchange and derivative market activity in 2016, Technical report.
- 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.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–54.
- Bloom, N., 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307 – 327.
- 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.
- Curcuro, Stephanie, Clara Vega, and Jasper Hoek, 2011, Measuring carry trade activity, in Bank for International Settlements, ed., *Proceedings of the IFC Conference on*, volume 34, 436–453 (Bank for International Settlements).
- 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.
- Engle, Robert F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation, *Econometrica* 50, 987–1007.
- 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.
- French, Kenneth, G. Schwert, and Robert Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- 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.
- Londono, Juan M, 2015, The variance risk premium around the world .
- Londono, Juan M., and Hao Zhou, 2017, Variance risk premiums and the forward premium puzzle, *Journal of Financial Economics* 124, 415 – 440.
- Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

- 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, 2012, Carry trades and global foreign exchange volatility, *The Journal of Finance* 67, 681–718.
- Mirkov, Nikola, Igor Pozdeev, and Paul Soderlind, 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.
- 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.
- Patton, Andrew J., 2011, Volatility forecast comparison using imperfect volatility proxies, *Journal of Econometrics* 160, 246 – 256, Realized Volatility.
- 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.
- van Binsbergen, Jules H., and Ralph S.J. Koijen, 2017, The term structure of returns: Facts and theory, *Journal of Financial Economics* 124, 1 – 21.
- Verdelhan, Adrien, 2015, The share of systematic variation in bilateral exchange rates, *Journal of Finance* (forthcoming).
- 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).

# Appendix A Variance forecasting models

## A.1 ARIMA(1,1,1)

At time  $t$ ,  $N$ -day-ahead forecasts from this model are made by calculating non-overlapping  $N$ -day values of realized variance up until time  $t$ :

$$\{\sigma_k^2 = RV(t - kN + 1, t - (k - 1)N)\},$$
$$k = \dots, 2, 1$$

with  $RV$  defined as in eq. (8). The following model is calibrated to the log of this series:

$$(1 - \rho_1 L)(1 - L) \log \sigma_k^2 = \theta_0 + (1 - \theta_1 L)z_t,$$

where  $L$  is the lag operator. The  $N$ -day-ahead forecast of  $RV$  is in this representation the 1-period-ahead forecast of  $\log \sigma^2$ , corrected for the Jensen's term:

$$\sigma_{k+1}^2 = \exp\{\widehat{\log \sigma_{k+1}^2} + 0.5V(z_t)\},$$

where  $V(z_t)$  is the sample variance of model residuals. If the errors  $z_t$  are Normally distributed, the correction is exact.

## A.2 GARCH(1,1)

Returns are modeled as:

$$r_t = \sigma_t z_t,$$
$$z_t \sim \mathcal{N}(0, 1),$$
$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2,$$

where  $r$  is the return, and  $\sigma_t^2 = \text{Var}_t(r(t, t + 1))$  is its conditional variance. The  $N$ -step-ahead forecasts are computed as follows:

$$h(t, t + N) = \sigma^2 + (\alpha + \beta)^{N-1}(h(t, t + 1) - \sigma^2),$$

where  $\sigma^2 = \frac{\omega}{1-\alpha-\beta}$  is the unconditional variance of  $r$ , and  $h(t, t + 1)$  is the one-step-ahead forecast made at time  $t$ .



Table 1: Previously studied vs. extreme attainable currency risk premia.

	max	min	drs
mean	14.68 (1.78)	-4.64 (-0.70)	6.90 (0.93)
median	-10.69	-25.80	-16.07

This table presents select descriptive statistics of 1-month variance risk premium series: that of the strategy of [Della Corte et al. \(2016\)](#) (“drs”), as well as the maximum attainable (“max”) and the minimum attainable (“min”) variance risk premium. *VRP* calculated as the difference between the average realized variance of a strategy and its average *IV*, and reported in percent squared per year. The “drs” strategy is constructed by going long two currencies with the highest difference between *IV* and the previously realized variance, and short two with the lowest such difference; the maximum (minimum) attainable *VRP* strategy is constructed by selecting the portfolio with the highest (lowest) difference between *portfolio IV* and the previously realized *portfolio* variance. The *t*-statistics of the means calculated with the [Newey and West \(1987\)](#) adjustment and the automatic lag selection procedure of [Newey and West \(1994\)](#) are shown in parentheses. The sample period is from 06/2008 to 06/2017.

Table 2: Forecasting 1-month realized variance

	<i>Panel A: carry trade</i>					<i>Panel B: USD index</i>				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
$\alpha$	-27.77 (-0.66)	41.39 (1.02)	87.67 (2.56)	22.95 (0.54)	16.86 (0.51)	-1.00 (-0.06)	15.44 (1.74)	24.67 (3.14)	-0.12 (-0.01)	9.32 (1.85)
<i>arima</i>	1.38 (2.77)				-0.10 (-0.36)	0.89 (3.22)				-0.29 (-0.90)
<i>garch</i>		0.86 (1.94)			0.65 (1.21)		0.75 (3.79)			-0.28 (-1.17)
<i>lag</i>			0.57 (3.03)		-0.18 (-0.45)			0.62 (3.08)		0.17 (0.91)
<i>mfiv</i>				1.01 (1.29)	0.66 (2.27)				0.88 (5.02)	1.15 (4.07)
$R^2$	0.14	0.18	0.16	0.18	0.19	0.30	0.34	0.31	0.53	0.54


This table shows the output of regressing 1-month realized variance of the carry trade (Panel A) and USD index (Panel B) onto variance forecasts from different models. For each strategy, columns (1)-(4) keep the output of univariate predictive regressions with respective forecasts as the right-hand variable; column (5) keeps the output of the multivariate regression with all forecasts entering the right-hand side of the predictive regression jointly. The top row keeps regression constants; the bottom row – adjusted R-squared values. Respective  $t$ -statistics are shown in parentheses below coefficients. Variance is expressed in percent squared per year. The sample period is from 01/2006 to 12/2017.

Table 3: Variance swap returns: sample statistics

	horizon	1	2	3	4	6	9	12
aud	mean	14.03 (22.33)	12.86 (34.40)	11.76 (35.48)	8.62 (18.17)	7.03 (12.44)	3.83 (3.74)	3.71 (3.11)
	sharpe	0.10	0.13	0.15	0.15	0.20	0.20	0.27
	median	-17.95	-10.09	-7.37	-7.32	-3.79	-2.93	-2.14
cad	mean	-0.42 (5.96)	-0.45 (4.05)	-0.15 (4.19)	-0.00 (1.94)	0.18 (1.51)	0.11 (0.64)	0.06 (0.63)
	sharpe	-0.01	-0.02	-0.01	-0.00	0.02	0.02	0.01
	median	-9.89	-5.62	-4.50	-3.51	-2.42	-2.17	-1.44
chf	mean	30.25 (55.88)	16.26 (27.70)	11.78 (26.41)	8.83 (19.51)	7.25 (15.01)	4.88 (8.55)	3.28 (4.27)
	sharpe	0.08	0.11	0.13	0.14	0.18	0.22	0.26
	median	-20.81	-10.43	-7.43	-5.96	-4.17	-2.81	-2.05
eur	mean	-12.90 (8.14)	-7.11 (4.54)	-4.75 (3.73)	-3.81 (1.93)	-1.99 (1.18)	-1.21 (0.64)	-0.86 (0.45)
	sharpe	-0.20	-0.24	-0.23	-0.25	-0.18	-0.16	-0.15
	median	-28.39	-14.69	-10.25	-8.47	-5.99	-4.23	-3.35
gbp	mean	-0.36 (4.44)	-0.79 (3.21)	-0.60 (3.64)	-0.94 (1.69)	0.31 (2.02)	0.69 (1.43)	0.93 (1.52)
	sharpe	-0.01	-0.03	-0.03	-0.07	0.03	0.08	0.13
	median	-12.81	-7.00	-5.03	-4.21	-2.82	-2.67	-1.85
jpy	mean	2.09 (12.65)	4.06 (14.50)	4.02 (22.13)	3.68 (7.53)	1.85 (5.34)	1.65 (1.41)	0.78 (1.49)
	sharpe	0.02	0.07	0.09	0.10	0.09	0.12	0.09
	median	-23.97	-11.27	-9.05	-6.62	-3.60	-3.09	-2.76
nzd	mean	0.38 (7.38)	0.77 (6.84)	0.75 (5.69)	0.67 (2.29)	0.04 (1.47)	0.24 (0.80)	0.01 (0.62)
	sharpe	0.01	0.02	0.03	0.04	0.00	0.03	0.00
	median	-12.21	-6.39	-5.01	-3.26	-2.30	-1.95	-1.81
usd	mean	-12.91 (3.63)	-6.12 (2.79)	-3.93 (2.41)	-3.12 (1.78)	-1.98 (1.02)	-1.22 (0.66)	-0.86 (0.44)
	sharpe	-0.32	-0.27	-0.22	-0.21	-0.19	-0.17	-0.16
	median	-23.59	-11.59	-7.84	-6.48	-4.06	-2.55	-2.34
carry	mean	19.05 (22.94)	14.08 (21.30)	11.63 (24.35)	8.94 (10.89)	6.81 (7.31)	4.72 (5.00)	3.70 (3.60)
	sharpe	0.14	0.17	0.19	0.20	0.24	0.26	0.29
	median	-16.15	-10.13	-6.96	-5.55	-3.68	-1.30	-1.40

This figure shows average and median monthly values and Sharpe ratios of returns on synthetic swaps written on the variance of currency indexes and the carry trade portfolio. Column names convey the horizon that the columns keeps estimates for, in months. Standard errors of the mean, calculated with the [Newey and West \(1987\)](#) adjustment and the automatic lag selection of [Newey and West \(1994\)](#), are reported in parentheses below respective values. Definitions and sample period the same as in [Figure 5](#).

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. (10). 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.


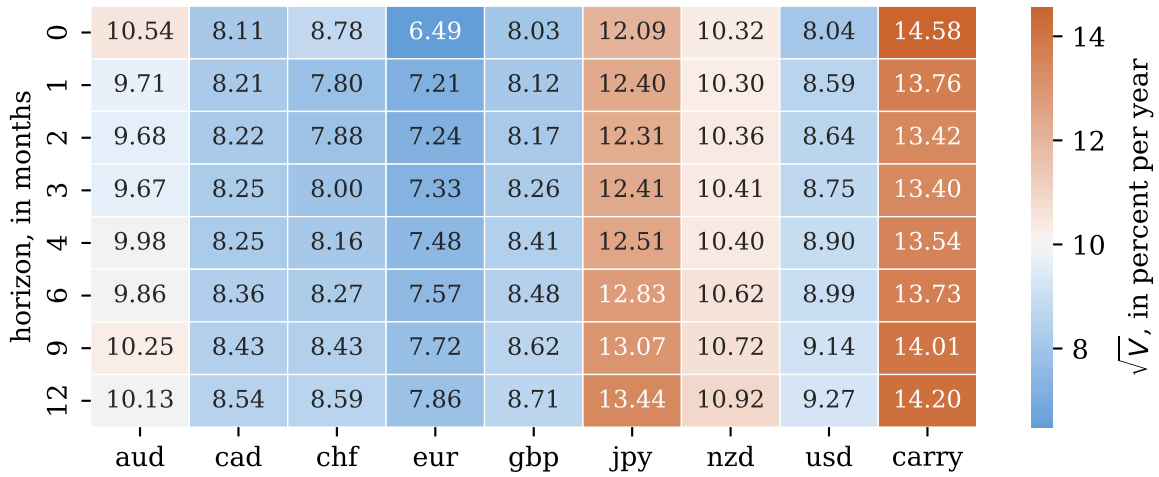



Figure 3: Realized and implied variance of currency portfolios



This figure shows the square root of the average implied variance of currency portfolios for horizons of 1 to 12 months, and of the average realized variance, put in the row corresponding to horizon 0. The former are calculated as in eq. (8), the latter – as in eq. (10) (by constructions, the average monthly realized variance is the same as quarterly, annual etc.). The observations of the implied variance are daily; the realized variance is calculated by averaging and rescaling squared daily returns over the whole sample. The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2006 to 12/2017.


Figure 4: Predictability of  $RV$  by  $IV$

(A):  $RMSE$



(continued on the next page)

(B): QLIKE




This figure shows RMSE (panel A) and average QLIKE loss (panel B) of variance forecasts. The  $y$ -axis keeps the forecasting horizons, in months; the  $x$ -axis keeps – four forecasting models for each considered currency portfolio. The forecasts are out-of-sample, obtained each day: from ARIMA(1,1,1) model calibrated to the series of  $\tau$ -horizon realized variances (“arima”); from GARCH(1,1) model by summing up consecutive  $n$ -day forecasts (“garch”) for  $n \in \tau$ ; as the  $\tau$ -horizon model-free implied variance (“iv”); by using the realized variance over the previous period  $\tau$  (“lagged”). RMSE inherits the units of variance, namely, squared percent per year, whereas QLIKE is unitless. In panel A, the tick mark on the  $y$ -axis corresponds to an annualized variance of 50% (). The lowest average loss value in each cell is highlighted in red. The sample period is from 01/2006 to 12/2017.




Figure 5: FX variance swaps

(A): Payoffs on 10000 local currency notional

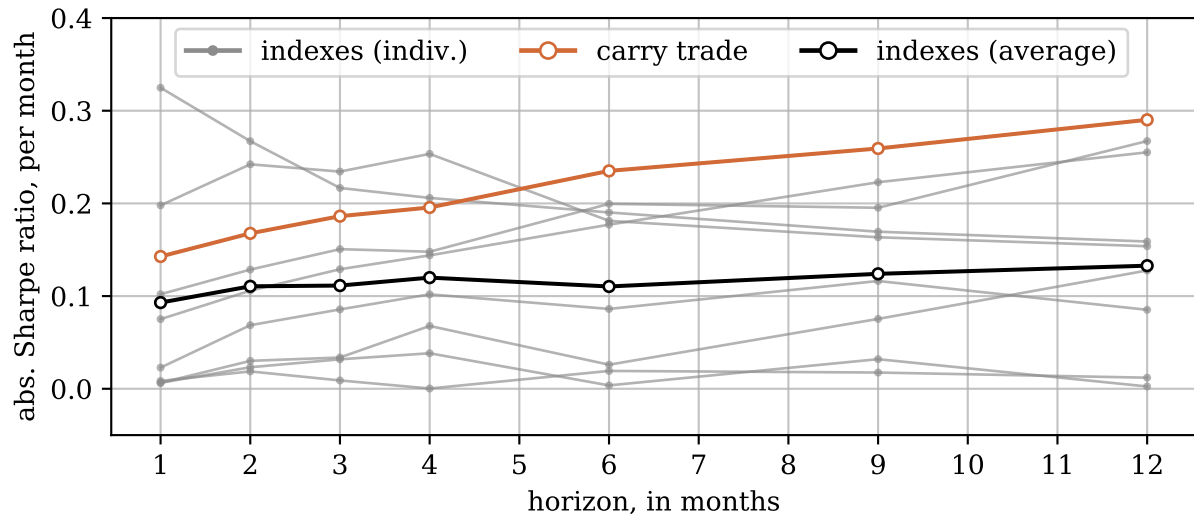


(B): Returns on fully collateralized positions




Panel A of this figure shows the average payoffs  $\times (RV(t, t + \tau) - IV_t(\tau))$ , and panel B – the average returns  $\times (RV(t, t + \tau) - IV_t(\tau)) / IV_t(\tau)$  of synthetic swaps on the variance of currency indexes and the carry trade portfolio. All values are per month. Estimates of the mean with the  $t$ -statistic above 2.0 are marked with an asterisk (\*), whereby the  $t$ -statistics are calculated with the [Newey and West \(1987\)](#) adjustment and the automatic lag selection of [Newey and West \(1994\)](#). The colorbar to the right maps heatmap colors to numeric values. The observations are daily and overlapping; the sample period is from 01/2006 to 12/2017.

Figure 6: Variance swap returns: Sharpe ratios



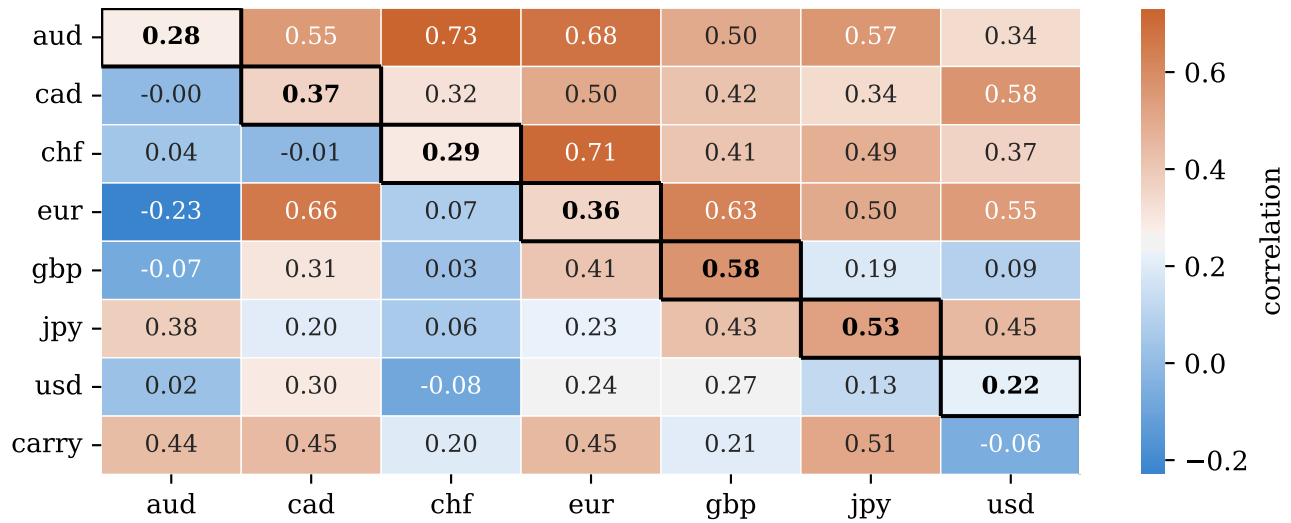
This figure depicts absolute Sharpe ratios of variance swap returns of currency indexes (in gray) and the carry trade portfolio (in red), previously presented in Table 3, as well as the cross-section average of the former. The returns are in percent per month, calculated as in eq. (25); the Sharpe ratios are per month. The sample period is from 01/2006 to 12/2017.

Figure 7: FX and stock market variance risk premium



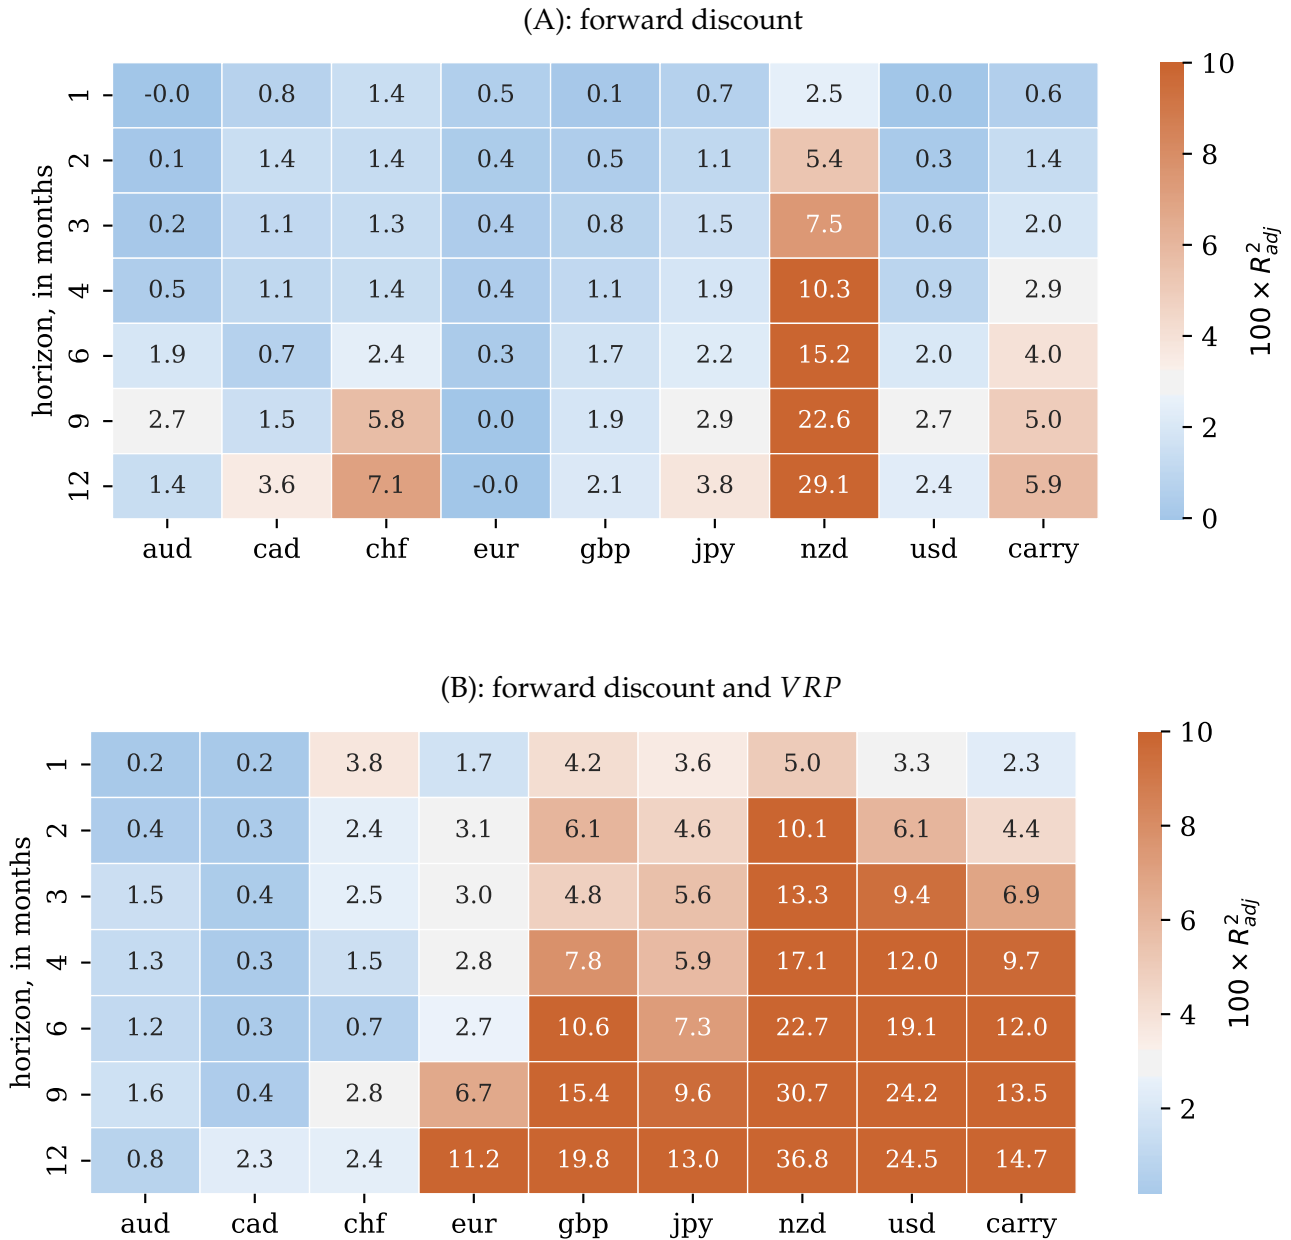
In this figure, the average 1-month variance swap return for stock market indexes denominated in the local currency (on the  $x$ -axis) is plotted against that of corresponding currency portfolios (on the  $y$ -axis). The solid line shows fitted values from the related regression. A variance swap return is defined as in eq. (25) using either VIX-like implied volatility index for stock market or MFIV for currency indexes, and expressed in percent per month. Stock and implied volatility indexes are as defined in Section 2.2. The sample period is from 01/2006 to 12/2017.

Figure 8: FX and stock market variance risk: correlation structure




Each cell in this heatmap shows correlation between 1-month variance risk premium estimates of two assets: for the main diagonal (cells with thick black borders), these are a local stock market index and the corresponding currency index; for the part below the diagonal, these are two currency indexes; for the part above the diagonal, these are two stock market indexes. Variance risk premium is defined with GARCH(1,1)-based variance forecasts on the left side of the difference in eq. (12) and either VIX-like implied volatility index for the stock market or MFIV for the currency portfolio. Stock and implied volatility indexes are as defined in Section 2.2. The carry trade portfolio does not have an associated stock market index; NZD is left out given no VIX data for the local stock market. The colorbar to the right maps colors in the figure to numeric values. The observations are daily and overlapping; the sample period is from 01/2006 to 12/2017.

Figure 9: Predicting spot returns of currency portfolios




This figure presents adjusted  $R^2$  values from regressions in eq. (26) (panel A) and (27) (panel B) for different currency portfolios (on the  $y$ -axis) and return horizons  $\tau$  (on the  $x$ -axis). The colorbar to the right maps heatmap colors to numeric values. The observations are daily and overlapping; the sample period is from 01/2006 to 12/2017.

Figure 10: Forward variance prices of FX portfolios: term structure



This figure shows the square root of the average forward variance price of currency portfolios for horizons along the  $x$ -axis, in percent per year. Each value is thus the one-to-one increasing transformation of the average price an investor pays at time  $t$  to receive the variance accumulated  $\tau$  months in the future. For maturities of 5, 7, 8, 10 and 11 months (shown as solid small dots), the  $IV$  used in calculation is fitted using natural cubic splines.  $F(0)$  is the average realized variance, and  $F(1) = IV(1)$ . The observations are daily and overlapping, one month is taken to be 22 days long. The sample period is from 01/2006 to 12/2017.

Figure 11: Rolling over forward variance claims: term structure of Sharpe ratios



This figure shows annualized Sharpe ratios of rolling over forward claims on variance of currency indexes and the carry trade portfolio, for horizons along the  $x$ -axis. Forward claims of maturity 1 correspond to the 1-month variance swaps. Values exceeding in magnitude  $2 \times$  own standard error are marked with an asterisk (\*). The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2006 to 12/2017.