

Average Variance, Average Correlation and Currency Returns*

Gino Cenedese

Bank of England

Lucio Sarno

Cass Business School and CEPR

Ilias Tsiakas

University of Guelph

First version: August 2011 - Revised: February 2012

Abstract

This paper provides an empirical investigation of the predictive ability of average variance and average correlation on the return to carry trades. Using quantile regressions, we find that higher average variance is significantly related to large future carry trade losses, whereas lower average correlation is significantly related to large gains. This is consistent with the carry trade unwinding in times of high volatility and the good performance of the carry trade when asset correlations are low. Finally, a new version of the carry trade that conditions on average variance and average correlation generates considerable performance gains net of transaction costs.

Keywords: Exchange Rates; Carry Trade; Average Variance; Average Correlation;

Quantile Regression.

JEL Classification: F31; G15; G17.

***Acknowledgements:** The authors are indebted for constructive comments to Bernard Dumas (editor), two anonymous referees, Pasquale Della Corte, Walter Distaso, Miguel Ferreira, Peter Gruber, Alex Maynard, Lukas Menkhoff, Francesco Nucci, Angelo Ranaldo, Pedro Santa-Clara, Maik Schmeling, Adrien Verdelhan and Paolo Volpin as well as to participants at the 2011 European Conference of the Financial Management Association, the 2011 ECB-Bank of Italy workshop on Financial Determinants of Exchange Rates, the 2011 Rimini Centre Workshop in International Financial Economics, the 2010 INFINITI Conference on International Finance at Trinity College Dublin, the 2010 Transatlantic Doctoral Conference at London Business School, and the 2010 Spring Doctoral Conference at the University of Cambridge. The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. *Corresponding author:* Lucio Sarno, Faculty of Finance, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK. Tel: +44 (0)20 7040 8772. Email: lucio.sarno@city.ac.uk. *Other authors' contact details:* Gino Cenedese, email: gino.cenedese@bankofengland.co.uk. Ilias Tsiakas, email: itsiakas@uoguelph.ca.

1 Introduction

The carry trade is a popular currency trading strategy that invests in high-interest currencies by borrowing in low-interest currencies. This strategy is designed to exploit deviations from uncovered interest parity (UIP). If UIP holds, the interest rate differential is on average offset by a commensurate depreciation of the investment currency and the expected carry trade return is zero. There is extensive empirical evidence dating back to Bilson (1981) and Fama (1984) that UIP is empirically rejected. In practice, it is often the case that high-interest rate currencies appreciate rather than depreciate.¹ As a result, over the last 35 years, the carry trade has delivered sizeable excess returns and a Sharpe ratio more than twice that of the US stock market (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011). It is no surprise, therefore, that the carry trade has attracted enormous attention among academics and practitioners.

An emerging literature argues that the high average return to the carry trade is no free lunch in the sense that high carry trade payoffs compensate investors for bearing risk. The risk measures used in this literature are specific to the foreign exchange (FX) market as traditional risk factors used to price stock returns fail to explain the returns to the carry trade (e.g., Burnside, 2012). In a cross-sectional study, Menkhoff, Sarno, Schmeling and Schrimpf (2012) find that the large average carry trade payoffs are compensation for exposure to global FX volatility risk. In times of high unexpected volatility, high-interest currencies deliver low returns, whereas low-interest currencies perform well. This suggests that investors should unwind their carry trade positions when future volatility risk increases. Christiansen, Ranaldo and Söderlind (2011) further show that the risk exposure of carry trade returns to stock and bond markets depends on the level of FX volatility. Lustig, Roussanov and Verdelhan (2011) identify a slope factor in the cross-section of FX portfolios based on the excess return to the carry trade itself constructed in similar fashion to the Fama and French (1993) “high-minus-low” factor. Burnside, Eichenbaum, Kleshchelski and Rebelo (2011) propose that the high carry trade payoffs reflect a peso problem, which is a low probability of large negative

¹ The empirical rejection of UIP leads to the well-known forward bias, which is the tendency of the forward exchange rate to be a biased predictor of the future spot exchange rate (e.g., Engel, 1996).

payoffs. Although they do not find evidence of peso events in their sample, they argue that investors still attach great importance to these events and require compensation for them. Brunnermeier, Nagel, and Pedersen (2009) suggest that carry trades are subject to crash risk that is exacerbated by the sudden unwinding of carry trade positions when speculators face funding liquidity constraints. Similar arguments based on crash risk and disaster premia are put forth by Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009) and Jurek (2009).

This paper investigates the intertemporal tradeoff between FX risk and the return to the carry trade. We contribute to the recent literature cited above by focusing on four distinct objectives. First, we set up a predictive framework, which differentiates this study from the majority of the recent literature that is primarily concerned with the cross-sectional pricing of FX portfolios. We are particularly interested in whether current market volatility can predict the future carry trade return. Second, we evaluate the predictive ability of FX risk on the full distribution of carry trade returns using quantile regressions, which are particularly suitable for this purpose. In other words, we relate changes in FX risk to large future gains and losses to the carry trade located in the tails of the return distribution. Predicting the full return distribution is useful for the portfolio choice of investors (e.g., Cenesizoglu and Timmermann, 2010), and can also shed light on whether we can predict currency crashes (Farhi *et al.*, 2009; and Jurek, 2009). Third, we define a set of FX risk measures that capture well the movements in aggregate FX volatility and correlation. These measures have recently been studied in the equities literature but are new to FX. Finally, we assess the economic gains of our analysis by designing a new version of the carry trade strategy that conditions on these FX risk measures.

The empirical analysis is organized as follows. The first step is to form a carry trade portfolio that is rebalanced monthly using up to 33 US dollar nominal exchange rates. Our initial measure of FX risk is the market variance defined as the variance of the returns to the FX market portfolio. We take a step further by decomposing the market variance in two components: the cross-sectional average variance and the cross-sectional average correlation, implementing the methodology applied by Pollet and Wilson (2010) to predict equity returns. Then, using quantile regressions, we assess the predictive ability of average variance and average correlation on the full distribution of carry trade returns. Quantile regressions provide

a natural way of assessing the effect of higher risk on different parts (quantiles) of the carry return distribution.² Finally, we design an augmented carry trade strategy that conditions on average variance and average correlation. This new version of the carry trade is implemented out of sample and accounts for transaction costs.

We find that the product of average variance and average correlation captures more than 90% of the time-variation in the FX market variance, suggesting that this decomposition works very well empirically. More importantly, the decomposition of market variance into average variance and average correlation is crucial for understanding the risk-return tradeoff in FX. Average variance has a significant negative effect on the left tail of future carry trade returns, whereas average correlation has a significant negative effect on the right tail. This implies that: (i) higher average variance is significantly related to large losses in the future returns to the carry trade, potentially leading investors to unwind their carry trade positions, and (ii) lower average correlation is significantly related to large future carry trade returns by enhancing the gains of diversification. Market variance is a weaker predictor than average variance and average correlation because, by aggregating information about the latter two risk measures into one risk measure, market variance is less informative than using average variance and average correlation separately. Finally, the augmented carry trade strategy that conditions on average variance and average correlation performs considerably better than the standard carry trade, even accounting for transaction costs. Taken together, these results imply the existence of a meaningful predictive relation between average variance, average correlation and carry trade returns: average variance and average correlation predict currency returns when it matters most, namely when returns are large (negative or positive), whereas the relation may be non-existent in normal times.

In addition, we find that average variance is a significant predictor of the left tail of the exchange rate component to the carry trade return. We then show that the predictive ability of average variance and average correlation is robust to the inclusion of additional predictive variables. It is also robust to changing the numeraire from the US dollar to a composite

² Cenesizoglu and Timmermann (2010) estimate quantile regressions and relate them to the intertemporal capital asset pricing model of Merton (1973, 1980). Their results show that predictive variables (such as average variance and average correlation) have their largest effect on the tails of the return distribution.

numeraire that is based on the US dollar, the euro, the UK pound and the Japanese yen. We further demonstrate that implied volatility indices, such as the VIX for the equities market and the VXY for the FX market, are insignificant predictors of future carry returns, and hence cannot replicate the predictive information in average variance and average correlation. Finally, the predictive quantile regression framework allows us to compute a robust measure of conditional skewness, which is predominantly positive at the beginning of the sample and predominantly negative at the end of the sample.

The remainder of the paper is organized as follows. In the next section we provide an overview of the literature on the intertemporal tradeoff between risk and return, and motivate the empirical predictions we examine in this paper. In Section 3 we describe the FX data set and define the measures for risk and return on the carry trade. Section 4 presents the predictive quantile regressions and discusses estimation issues. In Section 5, we report the empirical results, followed by robustness checks and further analysis in Section 6. Section 7 discusses the augmented carry trade strategies and, finally, Section 8 concludes.

2 The Intertemporal Tradeoff between FX Risk and Return: A Brief Review and Testable Implications

Since the contribution of Merton (1973, 1980), there is a class of asset pricing models suggesting that there is an intertemporal tradeoff between risk and return. These models hold for any risky asset in any market and hence can be applied not only to equities but also to the FX market. For the carry trade, the intertemporal risk-return tradeoff may be expressed as follows:

$$r_{C,t+1} = \mu + \kappa\sigma_t^2 + \varepsilon_{t+1}, \quad (1)$$

$$\sigma_t^2 = \varphi_0 + \varphi_1 AV_t + \varphi_2 AC_t, \quad (2)$$

where $r_{C,t+1}$ is the return to the carry trade portfolio from time t to $t+1$; σ_t^2 is the conditional variance of the returns to the FX market portfolio, also known as FX market variance; AV_t is

the equally weighted cross-sectional average of the variances of all exchange rate excess returns at each time t ; AC_t is the equally weighted cross-sectional average of the pairwise correlations of all exchange rate excess returns at each time t ; and ε_{t+1} is a normally distributed error term. These variables will be formally defined in the next section. It is important to note now, however, that the return to FX market portfolio is simply an equally weighted average of all exchange rate excess returns. It is straightforward to show that this is also the excess return to a naive international bond diversification strategy that invests in all available foreign bonds with equal weights by borrowing domestically. The recent literature on cross-sectional currency pricing typically uses the FX market portfolio as a standard risk factor (e.g., Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2012).³

Equation (1) is a general characterization of the theoretical prediction that there is a positive linear relation between the conditional market variance and future excess returns. The coefficient κ on the conditional market variance reflects the investors' risk aversion and hence is assumed to be positive: as risk increases, risk-averse investors require a higher risk premium and the expected return must rise. There is an extensive literature investigating the intertemporal risk-return tradeoff, mainly in equity markets, but the empirical evidence on the sign and statistical significance of the relation is inconclusive. Often the relation between risk and return is found insignificant, and sometimes even negative.⁴

Equation (2) shows that the conditional FX market variance can be decomposed into

³The FX market portfolio is defined in a different way to the stock market portfolio. In equilibrium, the stock market portfolio has to be held collectively by risk-averse investors, so the risk-return tradeoff depends on the risk aversion of the representative agent. However, carry trades are long-short zero-investment portfolios, and representative agent models are not well suited to analyze the determinants of risk premia for these portfolios. For example, models based on heterogeneous risk-averse agents or incomplete consumption risk sharing or speculation may explain what drives such risk premia (e.g., Chan and Kogan, 2002; Constantinides and Duffie, 1996; Constantinides, 2002; Sarkissian, 2003). In short, our analysis does not provide a direct empirical test of the ICAPM in FX markets but rather investigates an intertemporal risk-return relation that is motivated by the ICAPM by making an assumption about the composition of the FX market portfolio that is fairly standard in the literature.

⁴See, among others, French, Schwert and Stambaugh (1987), Chan, Karolyi and Stulz (1992), Glosten, Jagannathan and Runkle (1993), Goyal and Santa-Clara (2003), Brandt and Kang (2004), Ghysels, Santa-Clara and Valkanov (2005), and Bali (2008). There is also a well-established literature that relates exchange rate returns to volatility, with mixed success (e.g., Diebold and Nerlove, 1989; and Bekaert, 1995). More recently, Christiansen (2011) finds a positive contemporaneous risk-return tradeoff in exchange rates but no evidence of a predictive risk-return tradeoff. This literature differs from our study in that it focuses on individual exchange rates and uses conventional measures of individual exchange rate volatility. In general, these papers cannot detect a meaningful link between volatility and exchange rate movements, and our analysis provides evidence that this is partly due to the way risk is measured.

average variance and average correlation, as shown by Pollet and Wilson (2010) for equity returns. This decomposition is an aspect of our analysis that is critical for distinguishing the effect of systematic and idiosyncratic risk on future carry trade returns as well as for delivering robust and statistically significant results in predicting future carry trade returns. We begin with a brief discussion of the relation between average variance and future returns. Goyal and Santa-Clara (2003) show that although the equally weighted market variance only reflects systematic risk, average variance captures both systematic and idiosyncratic risk and, as a result, average variance is a powerful predictor of future equity returns. This implies that idiosyncratic risk matters for equity returns and our analysis investigates whether this is also the case for FX excess returns.⁵

We now turn to the relation between average correlation and future returns. Standard finance theory suggests that lower correlations generally lead to improved diversification and better portfolio performance. This is also the case for carry trade strategies, as shown by Burnside, Eichenbaum and Rebelo (2008), who demonstrate that the gains of diversifying the carry trade across many currencies are large, raising the Sharpe ratio by over 50 percent. More to the point, Pollet and Wilson (2010) show that the average correlation of stocks is significantly positively related to future stock returns. In the context of the ICAPM, they argue that since the return on aggregate wealth is not directly observable (e.g., Roll, 1977), changes in aggregate risk may reveal themselves through changes in the correlation between observable stock returns. Hence an increase in aggregate risk can be related to higher average correlation and to higher future stock returns. In short, the first testable hypothesis of our empirical analysis is:

H1: FX risk measures based on average variance and average correlation are more powerful predictors than market variance for future FX excess returns.

The intertemporal risk-return model of Equations (1)–(2) can be applied to the full conditional distribution of returns. As shown by Cenesizoglu and Timmermann (2010), the

⁵There is a number of theoretical models and economic arguments that justify why a market-wide measure of idiosyncratic risk may matter for returns. For example, variants of the CAPM where investors hold undiversified portfolios for either rational (tax or transactions costs) or irrational reasons predict an intertemporal relation between returns and idiosyncratic risk (see Goyal and Santa-Clara, 2003, for a discussion of this literature). This relation also obtains in a cross-sectional context in models with incomplete risk sharing (e.g., Sarkissian, 2003).

conditional quantile function implied by this model has the form:

$$Q_{r_{C,t+1}}(\tau | AV_t, AC_t) = \mu + \varphi_0 (\kappa + Q_\tau^N) + (\kappa + Q_\tau^N) \varphi_1 AV_t + (\kappa + Q_t^N) \varphi_2 AC_t \quad (3)$$

$$= \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t, \quad (4)$$

where Q_τ^N is the τ -quantile of the normal distribution, which has a large negative value deep in the left tail and a large positive value deep in the right tail. If, as suggested by the ICAPM, $\kappa > 0$ and also $\varphi_1, \varphi_2 > 0$, then we expect AV_t and AC_t to have a negative slope in the left tail and positive in the right tail. This provides an economic justification for the use of quantile regressions in our empirical analysis to separate the effect of AV_t and AC_t on different quantiles of carry trade returns. In what follows, we provide further economic arguments that justify the use of quantile regressions in the context of the ICAPM.

The negative relation between volatility and large carry trade losses (i.e., the left tail of returns) is also consistent with the theoretical model of Brunnermeier and Pedersen (2008) and its empirical application to currency crashes by Brunnermeier, Nagel and Pedersen (2009). The economic mechanism of Brunnermeier and Pedersen (2008) gives a prominent role to liquidity by arguing that when funding liquidity is tight, traders are reluctant to take on positions in high margin securities. This lowers market liquidity further, leading to higher volatility. In this model, market illiquidity and volatility can be amplified by two distinct liquidity spirals thus having an asymmetric effect on the left tail of asset returns. First, a margin spiral emerges if margins are increasing in market illiquidity. In this case, a funding shock to traders lowers market liquidity, leading to higher margins, which tightens traders' funding constraints even further and so on. Second, a loss spiral arises if traders hold a large initial position that is negatively correlated with customers' demand shock. In this case, a funding shock lowers market liquidity, leading to trader losses on their initial position, forcing traders to sell more, causing a further price drop and so on.

This mechanism is empirically applied on currency crashes by Brunnermeier, Nagel and Pedersen (2009), who find that currency crashes are often the result of endogenous unwinding of carry trade activity caused by liquidity spirals. In particular, they show that when volatility

increases, the carry trade tends to incur losses and the volume of currency trading tends to decrease. More importantly, negative shocks have a much larger effect on returns than positive shocks. During high volatility, shocks that lead to losses are amplified when traders hit funding constraints and unwind their positions, further depressing prices, thus increasing the funding problems and volatility. Conversely, shocks that lead to gains are not amplified. This asymmetric effect indicates that not only is volatility negatively related to carry trade returns but high volatility has a much stronger effect on the left tail of carry returns than in the right tail.⁶ This economic mechanism based on liquidity spirals relates volatility to future carry trade losses and is consistent with Equations (3) and (4) describing the quantile regressions implied by the ICAPM. Hence, the second testable hypothesis of our empirical analysis is:

H2: The predictive power of average variance and average correlation varies across quantiles of the distribution of FX excess returns, and is strongly negative in the lower quantiles.

2.1 Other Related Literature

This paper is related to Bali and Yilmaz (2011), who estimate two types of predictive regressions based on the ICAPM: first, of individual FX returns on individual variances for which they find a positive but statistically insignificant relation; and second, of individual FX returns on the covariance between individual exchange rates and the FX market variance for which they find a positive and statistically significant relation. Our analysis, however, substantially deviates from Bali and Yilmaz (2011) in a number of ways: (i) we focus on the carry trade portfolio, not on individual exchange rates; (ii) we analyze a larger number of currencies (33 versus 6 exchange rates) and a longer sample (34 years versus 7 years); (iii) we decompose the market variance into average variance and average correlation; (iv) we assess predictability across the full distribution of carry trade returns using quantile regressions; and (v) we design a new carry trade strategy that conditions on average variance and average correlation leading

⁶Note that this is also consistent with the view that macroeconomic fundamentals determine which currencies have high and low interest rates and the long-run value of exchange rates. In the short-run, however, dramatic exchange rate movements (i.e., currency crashes) and high volatility occasionally happen without fundamental news announcements due to the unwinding of carry trades when speculators near funding constraints.

to substantial gains over the standard carry trade.

The risk measures employed in our analysis have been the focus of recent intertemporal as well as cross-sectional studies of the equity market. The intertemporal role of average variance is examined by Goyal and Santa-Clara (2003), as discussed earlier, and also by Bali, Cakici, Yan and Zhang (2005). These studies show that average variance reflects both systematic and idiosyncratic risk and can be significantly positively related to future equity returns. In the cross-section of equity returns, the negative price of risk associated with market variance is examined by Ang, Hodrick, Xing and Zhang (2006, 2009). They find that stock portfolios with high sensitivities to innovations in aggregate volatility have low average returns. Similarly, Krishnan, Petkova and Ritchken (2009) find a negative price of risk for equity correlations. Finally, Chen and Petkova (2010) examine the cross-sectional role of average variance and average correlation. They find that for portfolios sorted by size and idiosyncratic volatility, average variance has a negative price of risk, whereas average correlation is not priced.

3 Measures of Return and Risk for the Carry Trade

This section describes the FX data set and defines our measures for: (i) the excess return to the carry trade for individual currencies, (ii) the excess return to the carry trade for a portfolio of currencies, and (iii) three measures of risk: market variance, average variance and average correlation.

3.1 FX Data

We use a cross-section of US dollar nominal spot and forward exchange rates by collecting data on 33 currencies relative to the US dollar: Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Euro area, Finland, France, Germany, Greece, Hong Kong, Hungary, India, Ireland, Italy, Japan, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan and United Kingdom. The sample period runs from January 1976 to February 2009. Note that the number of exchange rates for which there are available data varies over time;

at the beginning of the sample we have data for 15 exchange rates, whereas at the end we have data for 22. The data are collected by WM/Reuters and Barclays and are available on Thomson Financial Datastream. The exchange rates are listed in Table 1.⁷

3.2 The Carry Trade for Individual Currencies

An investor can implement a carry trade strategy for either individual currencies or, more commonly, a portfolio of currencies. In practice, the carry trade strategy for individual currencies can be implemented in one of two equivalent ways. First, the investor may buy a forward contract now for exchanging the domestic currency into foreign currency in the future. She may then convert the proceeds of the forward contract into the domestic currency at the future spot exchange rate. The excess return to this currency trading strategy for a one-period horizon is defined as:

$$r_{j,t+1} = s_{j,t+1} - f_{j,t}, \quad (5)$$

for $j = \{1, \dots, N_t\}$, where N_t is the number of exchange rates at time t , $s_{j,t+1}$ is the log of the nominal spot exchange rate defined as the domestic price of foreign currency j at time $t + 1$, and $f_{j,t}$ is the log of the one-period forward exchange rate j at time t , which is the rate agreed at time t for an exchange of currencies at $t + 1$. Note that an increase in $s_{j,t+1}$ implies a depreciation of the domestic currency, namely the US dollar.

Second, the investor may buy a foreign bond while at the same time selling a domestic bond. The foreign bond yields a riskless return in the foreign currency but a risky return in the domestic currency of the investor. Hence the investor who buys the foreign bond is exposed to FX risk. In this strategy, the investor will earn an excess return that is equal to:

$$r_{j,t+1} = i_{j,t}^* - i_t + s_{j,t+1} - s_{j,t}, \quad (6)$$

where $i_{j,t}^*$ and i_t are the one-period foreign and domestic nominal interest rates respectively.

⁷Note that our data includes no more than 33 currencies to avoid having exchange rate series with short samples and non-floating regimes.

The carry trade return in Equation (6) has two components: the interest rate differential $i_{j,t}^* - i_t$, which is known at time t , and the exchange rate return $s_{j,t+1} - s_{j,t}$, which is the rate of depreciation of the domestic currency and will be known at time $t + 1$.

The returns to the two strategies are exactly equal due to the covered interest parity (CIP) condition: $f_{j,t} - s_{j,t} = i_t - i_{j,t}^*$ that holds in the absence of riskless arbitrage. As a result, there is an equivalence between trading currencies through spot and forward contracts and trading international bonds.⁸ The return $r_{j,t+1}$ defined in Equations (5) and (6) is also known as the FX excess return.

If UIP holds, then the excess return in Equations (5) and (6) will on average be equal to zero, and hence the carry trade will be unprofitable. In other words, under UIP, the interest rate differential will on average be exactly offset by a commensurate depreciation of the investment currency. However, it is extensively documented that UIP is empirically rejected so that high-interest rate currencies tend to appreciate rather than depreciate (e.g., Bilson, 1981; Fama, 1984). The empirical rejection of UIP implies that the carry trade for either individual currencies or portfolios of currencies tends to be highly profitable (e.g., Della Corte, Sarno and Tsiakas, 2009; Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011).

3.3 The Carry Trade for a Portfolio of Currencies

There are many versions of the carry trade for a portfolio of currencies. In this paper, we implement one of the most popular versions. We form a portfolio by sorting at the beginning of each month all currencies according to the value of the forward premium $f_{j,t} - s_{j,t}$. If CIP holds, sorting currencies from low to high forward premium is equivalent to sorting from high to low interest rate differential. We then divide the total number of currencies available in that month in five portfolios (quintiles), as in Menkhoff, Sarno, Schmeling and Schrimpf (2012). Portfolio 1 is the portfolio with the highest interest rate currencies, whereas portfolio 5 has the lowest interest rate currencies. The monthly return to the carry trade portfolio is

⁸There is ample empirical evidence that CIP holds in practice for the data frequency examined in this paper. For recent evidence, see Akram, Rime and Sarno (2008). The only exception in our sample is the period following Lehman's bankruptcy, when the CIP violation persisted for a few months (e.g., Mancini-Griffoli and Ranaldo, 2011).

the excess return of going long on portfolio 1 and short on portfolio 5. In other words, the carry trade portfolio borrows in low-interest rate currencies and invests in high-interest rate currencies. We denote the monthly return to the carry trade portfolio from time t to $t + 1$ as

$$r_{C,t+1}.$$

3.4 FX Market Variance

Our first measure of risk is the FX market variance, which captures the aggregate variance in FX. Note that this measure of market variance focuses exclusively on the FX market, and hence it is not the same as the market variance used in equity studies (e.g., Pollet and Wilson, 2010). Specifically, FX market variance is the variance of the return to the FX market portfolio. We define the excess return to the FX market portfolio as the equally weighted average of the excess returns of all exchange rates:⁹

$$r_{M,t+1} = \frac{1}{N_t} \sum_{j=1}^{N_t} r_{j,t+1}. \quad (7)$$

This can be thought of as the excess return to a naive $1/N_t$ currency trading strategy, or an international bond diversification strategy that buys N_t foreign bonds by borrowing domestically.¹⁰

We estimate the monthly FX market variance (MV) using a realized measure based on daily excess returns:

$$MV_{t+1} = \sum_{d=1}^{D_t} r_{M,t+d/D_t}^2 + 2 \sum_{d=2}^{D_t} r_{M,t+d/D_t} r_{M,t+(d-1)/D_t}, \quad (8)$$

where D_t is the number of trading days in month t , typically $D_t = 21$. Following French, Schwert and Stambaugh (1987), Goyal and Santa-Clara (2003), and Bali, Cakici, Yan and

⁹We use equal weights as it would be difficult to determine time-varying “value” weights on the basis of monthly turnover for each currency over our long sample range. Menkhoff, Sarno, Schmeling and Schrimpf (2011) weigh the volatility contribution of different currencies by their share in international currency reserves in a given year and find no significant differences relative to equal weights.

¹⁰Note that the direction of trading does not affect the FX market variance. If instead the US investor decides to lend 1 US dollar by buying a domestic US bond and selling N_t foreign bonds with equal weights, the excess return to the portfolio would be $r_{M,t+1}^* = -r_{M,t+1}$. However, the market variance would remain unaffected: $V(r_{M,t+1}^*) = V(-r_{M,t+1}) = V(r_{M,t+1})$.

Zhang (2005), among others, this measure of market variance accounts for the autocorrelation in daily returns.¹¹

3.5 Average Variance and Average Correlation

Our second set of risk measures relies on the Pollet and Wilson (2010) decomposition of MV into the product of two terms, the cross-sectional average variance (AV) and the cross-sectional average correlation (AC), as follows:

$$MV_{t+1} = AV_{t+1} \times AC_{t+1}. \quad (9)$$

The decomposition would be exact if all exchange rates had equal individual variances, but is actually approximate given that exchange rates display unequal variances. Thus, the validity of the decomposition is very much an empirical matter. Pollet and Wilson (2010) use this decomposition for a large number of stocks and find that the approximation works very well. As we show later, this approximation works remarkably well also for exchange rates.

We can assess the empirical validity of the decomposition by estimating the following regression:

$$MV_{t+1} = \alpha + \beta (AV_{t+1} \times AC_{t+1}) + u_{t+1}, \quad (10)$$

where $E[u_{t+1} | AV_{t+1} \times AC_{t+1}] = 0$. The coefficient β may not be equal to one because exchange rates do not have the same individual variance and there may be measurement error in MV_{t+1} , AV_{t+1} and AC_{t+1} . However, the R^2 of this regression will give us a good indication of how well the decomposition works empirically.

¹¹This is similar to the heteroskedasticity and autocorrelation consistent (HAC) measure of Bandi and Perron (2008), which uses linearly decreasing Bartlett weights on the realized autocovariances. Our empirical results remain practically identical when using the HAC market variance, and hence we use the simpler specification of Equation (8) for the rest of the analysis.

We estimate AV and AC as follows:

$$AV_{t+1} = \frac{1}{N_t} \sum_{j=1}^{N_t} V_{j,t+1}, \quad (11)$$

$$AC_{t+1} = \frac{1}{N_t(N_t - 1)} \sum_{i=1}^{N_t} \sum_{j \neq i}^{N_t} C_{ij,t+1}, \quad (12)$$

where $V_{j,t+1}$ is the realized variance of the excess return to exchange rate j at time $t + 1$ computed as

$$V_{j,t+1} = \sum_{d=1}^{D_t} r_{j,t+d/D_t}^2 + 2 \sum_{d=2}^{D_t} r_{j,t+d/D_t} r_{j,t+(d-1)/D_t}, \quad (13)$$

and $C_{ij,t+1}$ is the realized correlation between the excess returns of exchange rates i and j at time $t + 1$ computed as

$$C_{ij,t+1} = \frac{V_{ij,t+1}}{\sqrt{V_{i,t+1}} \sqrt{V_{j,t+1}}}, \quad (14)$$

$$V_{ij,t+1} = \sum_{d=1}^{D_t} r_{i,t+d/D_t} r_{j,t+d/D_t} + 2 \sum_{d=2}^{D_t} r_{i,t+d/D_t} r_{j,t+(d-1)/D_t}. \quad (15)$$

Note that we do not demean returns in calculating variances. This allows us to avoid estimating mean returns and has very little impact on calculating variances (see, e.g., French, Schwert and Stambaugh, 1987).

3.6 Systematic and Idiosyncratic Risk

Define $V_{j,d}$ as the variance of the excess return to exchange rate j on day d . In this section only, for notational simplicity we suppress the monthly index t . Then, $V_{j,d}$ is a measure of total risk that contains both systematic and idiosyncratic components. Following Goyal and Santa-Clara (2003), we can decompose these two parts of total risk as follows. Suppose that the excess return $r_{j,d}$ is driven by a common factor μ_d and an idiosyncratic zero-mean shock $\varepsilon_{j,d}$ that is specific to exchange rate j . For simplicity, further assume that the factor loading for each exchange rate is equal to one, the common and idiosyncratic factors are uncorrelated, and ignore the serial correlation adjustment in Equation (13). Then, the data generating

process for daily returns is:

$$r_{j,d} = \mu_d + \varepsilon_{j,d}, \quad (16)$$

and the return to the FX market portfolio for day d in a given month t is:

$$r_{M,d} = \frac{1}{N_t} \sum_{j=1}^{N_t} r_{j,d} = \mu_d + \frac{1}{N_t} \sum_{j=1}^{N_t} \varepsilon_{j,d}, \quad (17)$$

where the second term becomes negligible for large N_t .

It is straightforward to show that in a given month t :

$$MV = \sum_{d=1}^{D_t} r_{M,d}^2 = \sum_{d=1}^{D_t} \left[\mu_d^2 + \frac{2}{N_t} \mu_d \sum_{j=1}^{N_t} \varepsilon_{j,d} + \left(\frac{1}{N_t} \sum_{j=1}^{N_t} \varepsilon_{j,d} \right)^2 \right], \quad (18)$$

$$AV = \sum_{d=1}^{D_t} \left[\frac{1}{N_t} \sum_{j=1}^{N_t} r_{j,d}^2 \right] = \sum_{d=1}^{D_t} \left[\mu_d^2 + \frac{2}{N_t} \mu_d \sum_{j=1}^{N_t} \varepsilon_{j,d} + \frac{1}{N_t} \sum_{j=1}^{N_t} \varepsilon_{j,d}^2 \right]. \quad (19)$$

The first two terms of MV and AV are identical and capture the systematic component of total risk as they depend on the common factor. The third term that depends exclusively on the idiosyncratic component is different for MV and AV. For a large cross-section of exchange rates, this term is negligible for MV, and hence MV does not reflect any idiosyncratic risk. For AV, however, the third term is not negligible and captures the idiosyncratic component of total risk.

To get a better idea of the relative size of the systematic and idiosyncratic components, consider the following example based on the descriptive statistics of Table 2. In annualized terms, the expected monthly variances are:

$$E[MV] \times 10^3 = 5 = \underbrace{5}_{\text{Systematic}} + \underbrace{0}_{\text{Idiosyncratic}}, \quad (20)$$

$$E[AV] \times 10^3 = 10 = \underbrace{5}_{\text{Systematic}} + \underbrace{5}_{\text{Idiosyncratic}}. \quad (21)$$

Therefore, half of the risk captured by AV in FX is systematic and the other half is idiosyncratic.

catic, whereas all of the risk reflected in MV is systematic.

Similarly, the standard deviations are:

$$STD [MV] \times 10^3 = 2, \quad (22)$$

$$STD [AV] \times 10^3 = 3. \quad (23)$$

As a result, the *t*-ratio of mean divided by standard deviation is 2.5 for MV and 3.3 for AV. In other words, AV is measured more precisely than MV, which can possibly make AV a better predictor of FX excess returns.

4 Predictive Regressions

Our empirical analysis begins with ordinary least squares (OLS) estimation of two predictive regressions for a one-month ahead horizon. The first predictive regression provides a simple way for assessing the intertemporal risk-return tradeoff in FX as follows:

$$r_{C,t+1} = \alpha + \beta MV_t + \varepsilon_{t+1}, \quad (24)$$

where $r_{C,t+1}$ is the return to the carry trade portfolio from time t to $t + 1$, and MV_t is the market variance from time $t - 1$ to t . This regression will capture whether, on average, the carry trade has low or negative returns in times of high market variance.

The second predictive regression assesses the risk-return tradeoff implied by the variance decomposition of Pollet and Wilson (2010):

$$r_{C,t+1} = \alpha + \beta_1 AV_t + \beta_2 AC_t + \varepsilon_{t+1}, \quad (25)$$

where AV_t and AC_t are the average variance and average correlation from time $t - 1$ to t . For notational simplicity, we use the same symbol α for the constants in the two regressions. The second regression separates the effect of AV and AC in order to determine whether the decomposition provides a more precise signal on future carry returns.

The simple OLS regressions focus on the effect of the risk measures on the conditional mean of future carry returns. We go further by also estimating two predictive quantile regressions, which are designed to capture the conditional effect of either MV or AV and AC on the full distribution of future carry trade returns. It is possible, for example, that average variance is a poor predictor of the conditional mean return but predicts well one or both tails of the return distribution. After all, higher variance implies a change in the tails of the distribution and here we investigate whether this is true in a predictive framework. Using quantile regressions provides a natural way of assessing the effect of higher risk on different parts of the distribution of future carry returns. It is also an effective way of dealing with outliers. For example, the median is a quantile of particular importance that allows for direct comparison to the OLS regression that focuses on the conditional mean. It is well known that outliers may have a much larger effect on the mean of a distribution than the median. Hence the quantile regressions can provide more robust results than OLS regressions even for the middle of the distribution. In our analysis, we focus on deciles of the distribution of future carry returns.

The first predictive quantile regression estimates the conditional quantile function:

$$Q_{r_{C,t+1}}(\tau | MV_t) = \alpha(\tau) + \beta(\tau) MV_t, \quad (26)$$

where τ is the quantile of the cumulative distribution function of one-month ahead carry returns.¹²

The second predictive quantile regression yields estimates of the conditional quantile function:

$$Q_{r_{C,t+1}}(\tau | AV_t, AC_t) = \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t. \quad (27)$$

The standard error of the quantile regression parameters is estimated using a moving block bootstrap (MBB) that provides inference robust to heteroskedasticity and autocorrelation of

¹²We obtain estimates of the quantile regression coefficients $\{\alpha(\tau), \beta(\tau)\}$ by solving the minimization problem $\{\alpha(\tau), \beta(\tau)\} = \arg \min_{\alpha, \beta} E[\rho_\tau(r_{C,t+1} - \alpha(\tau) - \beta(\tau) MV_t)]$, using the asymmetric loss function $\rho_\tau(r_{C,t+1}) = r_{C,t+1}(\tau - I(r_{C,t+1} < 0))$. We formulate the optimization problem as a linear program and solve it by implementing the interior point method of Portnoy and Koenker (1997). See also Koenker (2005, Chapter 6).

unknown form (Fitzenberger, 1997). Specifically, we employ a circular MBB of the residuals as in Politis and Romano (1992). The optimal block size is selected using the automatic procedure of Politis and White (2004) and Patton, Politis and White (2009). The bootstrap algorithm is detailed in the Appendix.

5 Empirical Results

5.1 Descriptive Statistics

Table 2 reports descriptive statistics on the following variables: (i) the return to the carry trade portfolio; (ii) the return to the exchange rate and interest rate components of the carry trade return; (iii) the excess return to the FX market portfolio; (iv) the FX market variance (MV); and (v) the FX average variance (AV) and average correlation (AC). Assuming no transaction costs, the carry trade delivers an annualized mean return of 8.6%, a standard deviation of 7.8% and a Sharpe ratio of 1.092.¹³ The carry trade return is primarily due to the interest rate differential across countries, which delivers an average return of 13.7%. The exchange rate depreciation component has a return of -5.1% , indicating that on average high-interest rate currencies do not depreciate enough to offset the interest rate differential. The carry trade return displays negative skewness of -0.967 and kurtosis of 6.043. These statistics confirm the good historical performance of the carry trade and are consistent with the literature (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011). Finally, the average market return is low at 1.0% per year, and its standard deviation is the same as that of the carry trade return at 7.8%.

Turning to the risk measures, the mean of MV is 0.005. The mean of AV is double that of MV at 0.010, and the mean of AC is 0.471. MV and AV exhibit high positive skewness and massive kurtosis. The time variation of AV and AC together with the cumulative carry trade return are displayed in Figure 1.

Panel B of Table 2 shows the cross-correlations. The correlation between the excess returns on the carry and the market portfolio is 9.1%. The three risk measures are highly positively

¹³We fully account for the effect of transaction costs in a later section.

correlated with each other but are negatively correlated with the carry and market returns. This is a first indication that there may be a negative risk-return relation in the FX market at the one-month horizon.

5.2 The Decomposition of Market Variance into Average Variance and Average Correlation

The three FX risk measures of MV, AV and AC are related by the approximate decomposition of Equation (9). We evaluate the empirical validity of the decomposition by presenting regression results in Table 3. The first regression is for MV on AV alone, which delivers a slope coefficient of 0.493 for AV and $\bar{R}^2 = 76.8\%$. The second regression is for MV on AC alone, which delivers a slope of 0.015 for AC and $\bar{R}^2 = 23.5\%$. The third regression is for MV on AV and AC (additively, not using their product), which raises \bar{R}^2 to 86.8%. Finally, the fourth regression is for MV on the product of AV and AC, which is consistent with the multiplicative nature of the decomposition, and delivers a slope coefficient of 0.939 and $\bar{R}^2 = 93.0\%$. In all cases, the coefficients are highly statistically significant. In conclusion, therefore, the MV decomposition into AV and AC captures almost all of the time variation in MV.

5.3 Predictive Regressions

We examine the intertemporal risk-return tradeoff for the carry trade by first discussing the results of OLS predictive regressions, reported in Table 4. The first regression is for the one-month ahead carry trade return on the lagged MV. The table shows that overall there is a significant negative relation. In other words, high market variance is related to low future carry trade returns. This clearly indicates a negative risk-return tradeoff for the carry trade and suggests that in times of high volatility the carry trade delivers low (or negative) returns. It is also consistent with the cross-sectional results of Menkhoff, Sarno, Schmeling and Schrimpf (2012), who find that there is a negative price of risk associated with high FX volatility.¹⁴

We refine this result by estimating a second regression for the one-step ahead carry trade

¹⁴It is important to emphasize, however, that our result is set up in a predictive framework, not in a cross-sectional contemporaneous framework.

return on AV and AC. We find that AV is also significantly negatively related to future carry trade returns. AC has a negative but insignificant relation. The \bar{R}^2 is 1.2% in the first regression and rises to 1.8% in the second regression. At first glance, therefore, there is at best a slight improvement in using the decomposition of MV into AV and AC in a predictive regression.

These results explore the risk-return tradeoff only for the mean of carry returns. It is possible, however, that high market volatility has a different impact on different quantiles of the carry return distribution. We explore this possibility by estimating predictive quantile regressions. We begin with Figure 2 which plots the parameter estimates of the predictive quantile regressions of the one-month ahead carry trade return on MV. These results are shown in more detail in Table 5. MV has a consistently negative relation to the future carry trade return but this relation is statistically significant only for a few parts of the distribution. The significant quantiles are all in the left tail: 0.05, 0.3, 0.4 and 0.5. Also note that the constant is highly significant, being negative below the 0.3 quantile and positive above it.

The results improve noticeably when we move to the second quantile regression of the future carry trade return on AV and AC. As shown in Figure 3 and Table 6, AV has a strong negative relation to the carry trade return, which is highly significant in all *left-tail* quantiles. The lower the quantile, the more negative the value of the coefficient. Above the median, the AV coefficient revolves around zero (positive or negative) and is not significant. Furthermore, it is interesting to note that AC has a negative and significant relation to the future carry trade return in the *right tail* of the distribution, and especially for quantiles 0.7 and higher.

The pseudo- \bar{R}^2 reported in Table 6 ranges from 0.1% for the quantile regressions describing the 0.6-quantile, to 4.1% for the quantile regression describing the extreme left tail of the return distribution.¹⁵ This result adds to the evidence that different parts of the return distribution present different degrees of predictability. In most cases, the pseudo- \bar{R}^2 is below 2%, in line with the modest predictability of FX excess returns typically found in the literature. However, we show in Section 6 how this modest statistical predictability leads to significant economic gains by designing trading strategies that condition on AV and AC.

¹⁵We compute the pseudo- \bar{R}^2 as in Koenker and Machado (1999).

These results lead to three important conclusions. First, it is very informative to look at the full distribution of carry trade returns to better assess the impact of high volatility and get a more precise signal. High MV has a significant negative impact only in the left tail.

Second, the decomposition of MV into AV and AC is helpful in understanding the risk-return tradeoff in FX. AV has a much stronger and significant negative impact than MV on the left tail of the carry trade return. As shown in Figure 3, this establishes clearly that high volatility in FX excess returns is strongly related to low future carry returns in the left tail. This is a new result that is consistent with the large negative returns to the carry trade in times of high volatility that typically lead investors to unwind their carry trade positions. It is also consistent with the empirical result in equity studies that idiosyncratic risk captured by (equally weighted) AV is significantly negatively related to the conditional mean of future returns (Goyal and Santa-Clara, 2003).

Third, AC is significantly negatively related to the right tail of future carry returns. This is also a new result. When FX return correlations are low, the carry trade is expected to perform well over the next period. The lower the correlations on average, the stronger the diversification effect arising from a given set of currencies, which tends to produce high carry trade returns. This is consistent with Burnside, Eichenbaum and Rebelo (2008), who show that diversification (i.e., trading a larger set of currencies) can substantially increase the Sharpe ratio of carry trade strategies. However, our result adds to previous empirical evidence in that we show that this diversification benefit tends to have an asymmetric effect on the return distribution: AC significantly affects the probability of large gains on carry trades, but its relation to large losses is insignificant. We do not have a theoretical explanation for this asymmetric effect, but believe that it is an intriguing result that warrants further research.

6 Robustness and Further Analysis

6.1 The Components of the Carry Trade

The carry trade return has two components: (i) the exchange rate component, which on average is slightly negative for our sample; and (ii) the interest rate component, which on

average is highly positive.¹⁶ Note that the exchange rate component is the uncertain part of the carry trade return as it is not known at the time that the carry trade portfolio is formed. In contrast, the interest rate component is known and actually taken into account when the carry trade portfolio is formed. Therefore, predicting the exchange rate component (i.e., whether high-interest currencies will depreciate and vice versa) effectively allows us to predict the carry trade return.

Figure 4 illustrates that when AV is high, the returns of the left-tail exchange rate component become lower. This negative relation is highly significant for the left tail of the exchange rate component up to the median. This result is consistent with high-interest currencies depreciating sharply (i.e., the forward bias diminishing) when AV is high and we are in the left tail of the distribution.¹⁷ More importantly, it also implies that to some extent exchange rates are predictable. In other words, this constitutes evidence against the well-known result that exchange rates are unpredictable. In short, our results establish that AV is a significant predictor of the large negative returns to the exchange rate component of the carry trade.

6.2 Additional Predictive Variables

As a robustness test, we use two additional predictive variables to determine whether they affect the significance of AV and AC. These are the average interest rate differential (AID) and the lagged carry return (LCR). AID is equal to the average interest rate differential of the quintile of currencies with the highest interest rates minus the average interest rate differential of the quintile of currencies with the lowest interest rates. All interest rates are known at time t for prediction of the carry trade return at time $t + 1$. LCR is simply the carry trade return lagged by one month.

The predictive quantile regression results are in Figure 5 and Table 7, and can be summarized in three findings: (i) AID has a significant positive effect on future carry trade returns in the middle of the distribution;¹⁸ (ii) LCR has a significant positive effect in the left tail;

¹⁶Recall the descriptive statistics in Table 2.

¹⁷This case is also consistent with the hypothesis of flight to quality, safety (e.g., Ranaldo and Söderlind, 2010) or liquidity that may explain why high-interest currencies depreciate and low-interest currencies appreciate in times of high volatility (Brunnermeier, Nagel and Pedersen, 2009).

¹⁸This is consistent with Lustig, Roussanov and Verdelhan (2010), who find that the average forward

and, more importantly, (iii) the effect of AV and AC remains qualitatively the same (although their significance diminishes slightly in the relevant parts of the distribution). The \bar{R}^2 now improves to 4.8% for the 0.05 quantile. Overall, the effect of AV and AC remains significantly negative in the left and right tails, respectively, even when we include other significant predictive variables.

6.3 The Numeraire Effect

A unique feature of the FX market is that investors trade currencies but all exchange rates are quoted relative to a numeraire. Consistent with the vast majority of the FX literature, we have used data on exchange rates relative to the US dollar. It is interesting, however, to check whether using a different numeraire would meaningfully affect the predictive ability of AV and AC. This is an important robustness check since it is straightforward to show analytically that the carry trade returns and risk measures are not invariant to the numeraire.¹⁹ In essence, the question we want to address is: given that changing the numeraire also changes the carry returns and the risk measures, does the relation between risk and return also change?

We answer this question by reporting predictive quantile regression results using a composite numeraire that weighs the carry trade return, AV and AC across four different currencies. The weights are based on the Special Drawing Rights (SDR) of the International Monetary Fund (IMF) and are as follows: 41.9% on the US dollar-denominated measures, 37.4% on the Euro-denominated measures, 11.3% on the UK pound-denominated measures, and 9.4% on the Japanese yen-denominated measures. The SDR is an international reserve asset created by the IMF in 1969 to supplement its member countries' official reserves that is based on a basket of these four key international currencies. The IMF (and other international organizations) also use SDRs as a unit of account and effectively that is what we do in this exercise.

The advantage of this approach is that: (i) we capture the numeraire effect in a single

discount is a good predictor of FX excess returns. In their study, the average forward discount is equal to the difference between the average interest of a basket of developed currencies and the US interest rate.

¹⁹For example, consider taking the point of view of a European investor and hence changing the numeraire currency from the US dollar to the euro. Then, all previous bilateral exchange rates become cross rates and N_t of the previous cross rates become bilateral. Furthermore, converting dollar excess returns into euro excess returns replaces the US bond as the domestic asset by the European bond.

regression as opposed to estimating multiple regressions for each individual numeraire; (ii) it is popular among practitioners who often measure FX returns using a composite numeraire across these four main currencies;²⁰ (iii) it provides the interpretation of generating a new weighted carry trade portfolio that is effectively a composite numeraire; and (iv) the weighted AV and weighted AC are straightforward to compute.²¹

The results shown in Figure 6 and Table 8 confirm that this exercise does not affect qualitatively our main result: the weighted AV still has a significant negative effect on the future weighted carry trade return in the lower tail, and the weighted AC still has a significant negative effect on the future weighted carry trade return in the upper tail. This is clear evidence that there is a strong statistical link between average variance, average correlation and future carry returns for certain parts of the distribution even when we consider a broad basket of numeraire currencies.

6.4 VIX, VXY and Carry Trade Returns

Our analysis quantifies FX risk using realized monthly measures of market variance, average variance and average correlation based on daily FX excess returns. An alternative way of measuring risk is to use implied volatility (IV) indices based on the IVs of traded options that can be thought of as the market's expectation of future realized volatility. As a further robustness check, we estimate predictive quantile regressions using two IV indices: the VIX index, which is based on the 1-month model-free IV of the S&P 500 equity index and is generally regarded as a measure of global risk appetite (e.g., Brunnermeier, Nagel and Pedersen, 2009); and the VXY index, which is based on the 3-month IV of at-the-money forward options of the G-7 currencies. The sample period for the VIX begins in January 1990 and for the VXY in January 1992, whereas for both it ends in February 2009.

²⁰Based on our experience, the typical weights adopted by practitioners in measuring returns relative to a composite numeraire are: 40% on the US dollar, 30% on the euro, 20% on the Japanese yen and 10% on the UK pound.

²¹The weighted carry trade return is computed as follows: $r_{C,t+1}^W = \sum_{p=1}^P w_p r_{p,C,t+1}$, where $p = 1, \dots, P = 4$ is the number of numeraires. The weighted average variance is: $AV_{t+1}^W = \sum_{p=1}^P w_p AV_{p,t+1} = \frac{1}{N_t} \sum_{j=1}^{N_t} \sum_{p=1}^P w_p V_{p,j,t+1}$ for j currencies. The weighted average correlation is: $AC_{t+1}^W = \sum_{p=1}^P w_p AC_{p,t+1} = \frac{1}{N_t(N_t-1)} \sum_{i=1}^{N_t} \sum_{j \neq i} \sum_{p=1}^P w_p C_{p,ij,t+1}$ for i, j currencies.

We begin with Table 9, which reports OLS results for simple contemporaneous regressions of each of the two IV indices on MV, AV and AC. These results will help us determine the extent to which the VIX and VXY are correlated with the FX risk measures we use. We find that the VIX is significantly positively related to AV and significantly negatively related to AC. Together AV and AC account for 37.6% of the variation of VIX. The VXY is also significantly related to AV but the relation to AC is low and insignificant. AV accounts for 54.5% of the variation in VXY.

The predictive quantile regression results for VIX and VXY, reported in Figure 7 and Table 10, suggest that neither the VIX nor the VXY are significantly related to future carry trade returns for any part of the distribution. Although the coefficients are predominantly negative, there is no evidence of statistical significance. Therefore, the predictive ability of AV and AC is not captured by the two IV indices and further justifies the choice of AV and AC as risk measures. The lack of predictive ability for the VIX in one-month ahead predictive regressions is consistent with the results of Brunnermeier, Nagel and Pedersen (2009), who find that the VIX has a strong contemporaneous impact on the carry return but is an insignificant predictor.

6.5 Conditional Skewness

The carry trade return is well known to exhibit negative skewness due to large negative outliers.²² This has led to an emerging literature that investigates whether the high average carry trade returns reflect a peso problem, which is the low probability of large negative outliers (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011); and are compensation for crash risk associated with the sudden unwinding of the carry trade (e.g., Brunnermeier, Nagel and Pedersen, 2009).

The predictive quantile regression approach has the advantage that it can be used to compute a measure of conditional skewness that is robust to outliers. As in Kim and White (2004), we use the Bowley (1920) coefficient of skewness that is based on the inter-quartile

²²Indeed, it is often said in industry speak that the carry trade payoffs “go up the stairs and down the elevator” or that the carry trade is like “picking up nickels in front of a steam roller.”

range. Using quantile regression (27), we estimate skewness conditionally period-by-period as follows:

$$SK_t = \frac{\widehat{Q}_{0.75,t} + \widehat{Q}_{0.25,t} - 2\widehat{Q}_{0.5,t}}{\widehat{Q}_{0.75,t} - \widehat{Q}_{0.25,t}}, \quad (28)$$

where $\widehat{Q}_{0.75,t}$ is the forecast of the third conditional quartile at time t for next period, $\widehat{Q}_{0.25,t}$ is the forecast of the first conditional quartile at time t for next period, and $\widehat{Q}_{0.5,t}$ is the forecast of the conditional median at time t for next period.

For any symmetric distribution, the Bowley coefficient is zero. This measure allows us to explore whether an increase in total risk (e.g., a rise in average variance) typically coincides with an increase in downside risk (e.g., lower conditional skewness). For example, when the lower tail conditional quantiles decline more than the upper tail conditional quantiles, this leads to negative conditional skewness and an increase in downside risk.

Figure 8 plots the conditional skewness and illustrates that it tends to be positive at the beginning of the sample and negative from the mid-nineties onwards. This is indicative of higher crash risk in the last part of the sample.

7 Augmented Carry Trade Strategies

We further evaluate the predictive ability of average variance and average correlation on future carry trade returns by assessing the economic gains of conditioning on average variance and average correlation out of sample for three augmented carry trade strategies. These strategies are then compared to the benchmark strategy, which is the standard carry trade. Our discussion begins with a description of the strategies, and then reports results with and without transaction costs.

7.1 The Strategies

The first augmented carry trade strategy conditions on AV only and implements the following rule at each time period t : for the carry trade returns that are lower than the τ -quantile of the distribution, if AV has *increased* from $t - 1$ to t , we *close* the carry trade positions and

thus receive an excess return of zero; otherwise we execute the standard carry trade. This strategy is designed to exploit the negative relation between current AV and the one-month ahead carry return. The focus of the AV strategy is the left-tail quantiles of the carry trade return distribution, where the negative effect of AV is the strongest.

The second strategy conditions on AC only and implements the following rule at each time period t : for the carry trade returns that are higher than the $1 - \tau$ quantile, if AC has *decreased* from $t - 1$ to t , we *double* the carry trade positions and thus receive twice the carry return; otherwise we execute the standard carry trade. This strategy is designed to exploit the negative relation between current AC and the one-month ahead carry return. However, the focus of the AC strategy is the right-tail quantiles, where the negative effect of AC is the strongest.

Third, the combined AV and AC strategy makes the following decision at each time period t : for the carry trade returns that are lower than the τ -quantile, if AV has *increased* from $t - 1$ to t , we *close* the carry trade positions and thus receive an excess return of zero; and for the carry returns that are higher than the $1 - \tau$ quantile, if AC has *decreased* from $t - 1$ to t , we *double* the carry trade positions and thus receive twice the carry trade return; otherwise we execute the standard carry trade. The combined AV and AC strategy focuses at the same time on both the low and the high quantiles of the carry trade return distribution. For example, this strategy is first applied to the 0.1 (for AV) and 0.9 (for AC) quantiles for the full sample, then to 0.2 and 0.8 quantiles and so on.

It is important to note that all three strategies are implemented out of sample. Specifically, all strategies move forward recursively starting 3 years after the beginning of the sample.²³ The strategies do not directly use the parameter estimates from the quantile regressions but simply try to exploit the negative relation between future carry returns and current AV and AC separately for low and high quantiles of the distribution.

The economic evaluation of the three strategies focuses on the Sharpe ratio. We also report the mean and standard deviation of the augmented carry trade returns. All these measures are reported in annualized units. We assess the practical applicability of the strategies by

²³Note that as a result of excluding the first three years of data, the statistics of the standard carry trade returns reported in Table 11 differ slightly from those in Table 2.

computing the turnover ratio as the percentage of the currencies that on average are traded every period. The turnover ratio provides us with a sense of how much more trading and rebalancing is required to implement an augmented strategy relative to the standard carry trade.

7.2 No Transaction Costs

Panel A of Table 11 reports the results for no transaction costs. The AV strategy performs very well for most quantiles and, as expected, does increasingly better as we move to the lower quantiles. For example, at the 0.1 quantile, the Sharpe ratio of the AV strategy is 1.314 compared to 1.070 for the standard carry trade. These large economic gains require only slightly higher trading as the turnover ratio rises from 17.1% in the benchmark case to 20.8%. By design, the turnover ratio remains reasonably low for the lowest quantiles where the conditioning on AV is implemented less often.

In contrast to the AV strategy that performs best in the lowest quantiles, the AC strategy performs well across all quantiles. It appears, therefore, that low average correlations are economically beneficial to the carry trade regardless of the quantile we focus on. For example, at the 0.1 quantile the Sharpe ratio of the AC strategy is 1.208, whereas at the 0.9 quantile it is 1.134 compared to 1.070 of the standard carry trade. These economic gains require only slightly higher trading for the high quantiles as the turnover ratio rises from 17.1% in the benchmark case to 21.4% at the 0.9 quantile. By design, for the AC strategy the turnover ratio remains reasonably low for the highest quantiles.

Finally, the combined AV and AC strategy performs better than the standard carry trade for all quantiles. For example, at the highest and lowest quantiles (0.1 for AV and 0.9 for AC) the Sharpe ratio is 1.224 and the turnover ratio is 0.257. In short, therefore, there are sizeable economic gains in implementing augmented carry trade strategies that condition on average variance and average correlation. In addition to showing tangible out-of-sample economic gains, these results also highlight the negative predictive relation of AV and AC to future carry trade returns across different quantiles of the distribution.

7.3 The Effect of Transaction Costs

A realistic assessment of the profitability of the carry trade strategies needs to account for transaction costs in trading spot and forward exchange rates. Every month that we form a new carry trade portfolio, we take a position in one forward and one spot contract for each currency that belongs to either portfolio 1 (highest interest rate currencies) or portfolio 5 (lowest interest rate currencies). At the end of the month, the contracts expire and new contracts are entered (on the same or different currencies). Define $c_{j,t}^S$ and $c_{j,t}^F$ as the one-way proportional transaction cost at time t for trading at the spot and forward exchange rate j , respectively. These values are equal to half of the spot and forward proportional bid-ask spread.²⁴ It is straightforward to show that the return to the carry trade for an individual currency j net of transaction costs is equal to:

$$r_{j,t+1}^{net} = s_{j,t+1} - f_{j,t} - c_{j,t+1}^S - c_{j,t}^F. \quad (29)$$

Our analysis implements the carry trade strategies using the transaction costs listed in Table 1. Due to data availability, we use the median transaction costs across time, which are different for each currency.²⁵ These are taken from Datastream and are computed using the longest bid and ask times series available for each exchange rate for the sample range of January 1976 to February 2009. The cross-currency average of the median one-way transaction costs is 6.10 basis points for the spot rates, and 9.42 basis points for the forward rates. Our transaction costs are consistent with the values discussed in Neely, Weller and Ulrich (2009).

The performance of the carry trade strategies with transaction costs is shown in Panel B of Table 11. As expected, when accounting for transaction costs the Sharpe ratios of all strategies are lower. For example, the standard carry trade has a Sharpe ratio of 1.070

²⁴The proportional bid-ask spread is the ratio of the bid-ask spread to the mid rate.

²⁵Using the median as opposed to the mean of transaction costs mitigates the effect of few large outliers in the time series of bid-ask spreads. The results remain largely unchanged when using the average proportional bid-ask spread. Note also that generally the effective spread is lower than the quoted spread, since trading will take place at the best price quoted at any point in time, suggesting that the worse quotes will not attract trades (e.g., Mayhew, 2002). Although some studies consider effective transaction costs in the range of 50% to 100% of the quoted spread (e.g., Goyal and Saretto, 2009), we use the full spread, which will likely underestimate the true returns.

before transaction costs and 0.741 after transaction costs. More importantly, we find that the augmented strategies still perform substantially better than the benchmark. In particular, the AV strategy still dominates the standard carry trade in the left tail. The AC as well as the combined AV and AC strategies outperform the benchmark in all quantiles. For example, the AV strategy can deliver a Sharpe ratio net of transaction costs as high as 0.983, the AC strategy as high as 0.971, and the the combined strategy as high as 0.959, compared to 0.741 for the benchmark. We conclude, therefore, that the improvement in the performance of the carry trade when conditioning on the movements of AV and AC is robust to transaction costs.

8 Conclusion

The carry trade is a currency investment strategy designed to exploit deviations from uncovered interest parity. Its profitability is based on the empirical observation that the interest rate differential across countries is not, on average, offset by a depreciation of the investment currency. Hence, investing in high-interest currencies by borrowing from low-interest currencies tends to deliver large positive excess returns.

This paper fills a gap in the literature by demonstrating empirically the existence of an intertemporal risk-return tradeoff between the return to the carry trade and risk in a predictive setting. We measure FX risk by the variance of the returns to the FX market portfolio. We then take a step further by decomposing the market variance into the cross-sectional average variance and the cross-sectional average correlation of exchange rate returns. Our empirical analysis is based on predictive quantile regressions, which provide a natural way of assessing the effect of higher risk on different quantiles of the return distribution.

Our main finding is that average variance has a significant negative effect on the left tail of the distribution of future carry trade returns, whereas average correlation has a significant negative effect on the right tail. We take advantage of this finding by forming a new version of the carry trade that conditions on average variance and average correlation, and show that this strategy performs considerably better than the standard carry trade. These results imply that to some extent exchange rates are predictable, especially when it matters most: when

the carry trade produces large gains or large losses. In other words, if the carry trade is about “going up the stairs and down the elevator,” then average variance and average correlation can tell us something valuable about when the elevator is likely to go up or down. In the end, by focusing on the tails of the return distribution of carry trades, we uncover a negative risk-return tradeoff in foreign exchange.

A Appendix: Notes on the bootstrap procedure

A.1 Bootstrap Standard Errors

We estimate the standard error of the parameters of the predictive quantile regressions using a moving block bootstrap (MBB), which provides inference that is robust to heteroskedasticity and autocorrelation of unknown form (Fitzenberger, 1997). Specifically, we employ a circular MBB of the model residuals as in Politis and Romano (1992). The optimal block size is selected using the automatic procedure of Politis and White (2004), as amended by Patton, Politis and White (2009). The bootstrap algorithm implements the following steps:

1. Estimate the coefficients of the τ -th conditional quantile function:

$$Q_{y_{t+1}}(\tau \mid X_t) = X_t' \beta(\tau),$$

for $t = 1, \dots, T$, where y_{t+1} is the dependent variable (e.g., the carry trade return), X_t is a $K \times 1$ matrix of regressors (e.g., a constant, AV and AC), and $\beta(\tau)$ is the $K \times 1$ vector of coefficients. We denote the estimates as $\hat{\beta}(\tau)$ and obtain the residuals associated to the τ -th quantile as $\hat{\varepsilon}_{t+1} = y_{t+1} - X_t' \beta(\tau)$.

2. “Wrap” the residuals $\{\hat{\varepsilon}_1, \dots, \hat{\varepsilon}_T\}$ around a circle, i.e., define the new series $\hat{e}_t = \hat{\varepsilon}_t$ for $t = 1, \dots, T$, and $\hat{e}_t = \hat{\varepsilon}_{t-T}$ for $t = T + 1, \dots, T + l - 1$, where l is the length of the block defined below. This “circular” structure, specified by Politis and Romano (1992), guarantees that the first and last few observations have the same probability of being selected as observations in the middle of the series.
3. Construct a bootstrap pseudo-series $\{e_1^*, \dots, e_T^*\}$ by resampling with replacement of overlapping blocks of size l . The block size l is computed using the automatic procedure of Politis and White (2004). In our application, the estimated optimal block length ranges from 1 to 4 monthly observations.
4. Form the dependent variable y_{t+1}^* using the bootstrap residual series $e^* \equiv \{e_1^*, \dots, e_T^*\}'$

and the estimates $\hat{\beta}(\tau)$ as follows:

$$y_{t+1}^* = X_t' \hat{\beta}(\tau) + e_{t+1}^*,$$

and then estimate the conditional quantile function $Q_{y_{t+1}^*}(\tau | X_t) = X_t' \beta(\tau)$, obtaining the bootstrap estimate $\hat{\beta}^*(\tau)$.

5. Repeat steps 3 and 4 for $B = 10,000$ times. Denoting $\hat{\beta}_j^*(\tau)$ as the estimate of the j^{th} bootstrap, for $j = 1, \dots, B$, estimate the variance-covariance matrix of $\hat{\beta}(\tau)$ as follows:

$$\widehat{\text{Var}}\left(\hat{\beta}(\tau)\right) = \frac{1}{B} \sum_{j=1}^B \left(\hat{\beta}_j^*(\tau) - \bar{\beta}_j^*(\tau)\right) \left(\hat{\beta}_j^*(\tau) - \bar{\beta}_j^*(\tau)\right)'$$

A.2 Bootstrap Hypothesis Testing

In order to test the null hypothesis of no predictability (i.e., $\beta(\tau) = 0$), we perform a *double bootstrap* to compute the bootstrap p -values (see, e.g., MacKinnon, 2007). The procedure is as follows:

- I. Obtain the estimate $\hat{\beta}(\tau)$.
- II. Obtain the estimate $\widehat{\text{Var}}\left(\hat{\beta}(\tau)\right)$ using the block bootstrap described above for $B_2 = 500$ bootstrap samples.
- III. Compute the t -statistic for each coefficient $\beta_i(\tau)$, $i = 1, \dots, K$ as follows:

$$\hat{t}_i = \frac{\hat{\beta}_i(\tau)}{\sqrt{\widehat{\text{Var}}\left(\hat{\beta}_i(\tau)\right)}}.$$

- IV. Generate $B_1 = 1,000$ bootstrap samples using the bootstrap DGP as in step 4 above, but this time imposing the null $\beta_i(\tau) = 0$. Use each of the new samples to calculate $\hat{\beta}_{i,j}^{**}(\tau)$, $j = 1, \dots, B_1$.

V. For each of the B_1 bootstrap samples, perform steps II and III above in order to generate B_1 bootstrap test statistics $t_{i,j}^{**}$.

VI. Calculate the bootstrap p -values for \hat{t}_i using

$$\hat{p}_i^*(\hat{t}_i) = \frac{1}{B_1} \sum_{j=1}^{B_1} I(|t_{i,j}^{**}| > |\hat{t}_i|),$$

where $I(\cdot)$ denotes the indicator function, which is equal to 1 when its argument is true and 0 otherwise.

Table 1. Exchange Rates

The table lists the 33 US dollar nominal exchange rates used to construct the FX market and carry trade portfolios. The start date and end date of the data sample is shown for each exchange rate. The transaction costs reported below are the median one-way proportional costs for the spot and forward exchange rates, defined as half of the bid-ask spread divided by the mid rate, and are reported in basis points. The transaction costs are computed using the longest bid and ask times series available for each exchange rate for the sample range of January 1976 to February 2009.

	<i>Exchange Rate</i>	<i>Start of Sample</i>	<i>End of Sample</i>	<i>Transaction Costs</i>	
				<i>Spot</i>	<i>Forward</i>
1	Australian Dollar	December 1984	February 2009	5.42	7.31
2	Austrian Schilling	January 1976	December 1998	7.64	11.08
3	Belgian Franc	January 1976	December 1998	7.92	12.67
4	Canadian Dollar	January 1976	February 2009	2.72	4.75
5	Czech Koruna	January 1997	February 2009	5.97	6.75
6	Danish Krone	January 1976	February 2009	4.16	7.03
7	Euro	January 1999	February 2009	2.64	2.77
8	Finnish Markka	January 1997	December 1998	6.52	5.28
9	French Franc	January 1976	December 1998	4.98	7.56
10	German Mark	January 1976	December 1998	8.09	15.55
11	Greek Drachma	January 1997	December 2000	4.07	5.45
12	Hong Kong Dollar	October 1983	February 2009	0.64	1.92
13	Hungarian Forint	October 1997	February 2009	5.05	8.36
14	Indian Rupee	October 1997	February 2009	3.13	4.39
15	Irish Punt	January 1976	December 1998	4.70	9.00
16	Italian Lira	January 1976	December 1998	3.16	8.54
17	Japanese Yen	June 1978	February 2009	10.32	9.76
18	Mexican Peso	January 1997	February 2009	4.35	5.25
19	Netherlands Guilder	January 1976	December 1998	11.71	17.21
20	New Zealand Dollar	December 1984	February 2009	8.14	10.44
21	Norwegian Krone	January 1976	February 2009	4.65	7.29
22	Philippine Peso	January 1997	February 2009	13.61	15.85
23	Polish Zloty	February 2002	February 2009	6.52	7.41
24	Portuguese Escudo	January 1976	December 1998	18.76	34.43
25	Saudi Arabian Riyal	January 1997	February 2009	0.53	0.93
26	Singaporean Dollar	December 1984	February 2009	3.14	6.94
27	South African Rand	October 1983	February 2009	11.46	20.12
28	South Korean Won	February 2002	February 2009	2.24	6.44
29	Spanish Peseta	January 1976	December 1998	6.98	12.08
30	Swedish Krona	January 1976	February 2009	4.85	7.76
31	Swiss Franc	January 1976	February 2009	11.61	18.60
32	Taiwanese Dollar	January 1997	February 2009	2.95	8.28
33	United Kingdom Pound	January 1976	February 2009	2.67	3.54

Table 2. Descriptive Statistics

The table reports descriptive statistics for the monthly excess returns of two FX portfolios: the carry trade and the market; for the two components of the carry trade: the exchange rate depreciation and the interest rate differential; and for three monthly risk measures: market variance, average variance and average correlation. The sample of 33 US dollar nominal exchange rates runs from January 1976 to February 2009. The return to the FX market portfolio is an equally weighted average of all exchange rate excess returns. The carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Market variance is the variance of the monthly returns to the FX market portfolio. Average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. The mean, standard deviation and the Sharpe ratio are annualized and assume no transaction costs. AR(1) is the first order autocorrelation.

Panel A: Summary Statistics						
	<i>Mean</i>	<i>St. Dev.</i>	<i>Sharpe Ratio</i>	<i>Skewness</i>	<i>Kurtosis</i>	<i>AR(1)</i>
<i>Portfolio Returns</i>						
Carry Trade	0.086	0.078	1.092	-0.967	6.043	0.132
Market	0.010	0.078	0.131	-0.120	3.195	0.097
<i>Carry Trade Components</i>						
Exchange Rate	-0.051	0.079		-1.133	6.232	0.140
Interest Rate	0.137	0.023		2.506	15.578	0.655
<i>Variances and Correlations</i>						
Market Variance	0.005	0.002		3.380	19.367	0.512
Average Variance	0.010	0.003		5.174	47.780	0.539
Average Correlation	0.477	0.182		0.028	2.241	0.796

Panel B: Cross-Correlations					
	Carry Return	Market Return	Market Variance	Average Variance	Average Correlation
<i>Portfolio Returns</i>					
Carry Trade Return	1.000				
Market Return	0.091	1.000			
<i>Variances and Correlations</i>					
Market Variance	-0.250	-0.145	1.000		
Average Variance	-0.371	-0.150	0.877	1.000	
Average Correlation	-0.048	-0.054	0.478	0.191	1.000

Table 3. Market Variance Decomposition

The table presents the ordinary least squares results for regressions on alternative decompositions of the FX market variance. The dependent variable is the market variance defined as the variance of the monthly returns to the FX market portfolio, which is an equally weighted average of the excess returns of 33 US dollar nominal spot exchange rates. Average variance is the equally weighted cross-sectional average of the variances of all exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise correlations of all exchange rate excess returns. All variables are contemporaneous and, with the exception of average correlation, they are annualized. Newey-West (1987) *t*-statistics with five lags are reported in parentheses. The sample period runs from January 1976 to February 2009.

Regressions for the Market Variance				
	(1)	(2)	(3)	(4)
Constant	0.000 (0.056)	-0.002 (-2.324)	-0.004 (-11.079)	0.000 (1.473)
Average Variance	0.493 (9.993)		0.456 (13.883)	
Average Correlation		0.015 (8.960)	0.010 (13.784)	
(Average Variance) × (Average Correlation)				0.939 (24.281)
\bar{R}^2 (%)	76.8	23.5	86.8	93.0

Table 4. OLS Predictive Regressions

The table presents the ordinary least squares results for two predictive regressions. The first regression is: $r_{C,t+1} = \alpha + \beta MV_t + \varepsilon_{t+1}$, where $r_{C,t+1}$ is the one-month ahead carry trade return and MV_t is the lagged market variance. The second regression is: $r_{C,t+1} = \alpha + \beta_1 AV_t + \beta_2 AC_t + \varepsilon_{t+1}$, where AV_t is the lagged average variance and AC_t is the lagged average correlation. The return to the carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Market variance is the variance of the return to the market portfolio constructed as the equally weighted average of the excess returns on 33 US dollar nominal exchange rates. Average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. With the exception of average correlation, all variables are annualized. Newey-West (1987) t -statistics with five lags are reported in parentheses. Bootstrap p -values generated using 10,000 bootstrap samples are in brackets. The sample period runs from January 1976 to February 2009.

Regressions for the Carry Trade Return		
	(1)	(2)
Constant	0.115 (5.939) [0.122]	0.156 (3.888) [0.064]
Market Variance	-5.999 (-2.006) [0.046]	
Average Variance		-3.781 (-1.972) [0.051]
Average Correlation		-0.070 (-0.936) [0.350]
\bar{R}^2 (%)	1.2	1.8

Table 5. Market Variance

The table presents the regression results for the conditional quantile function: $Q_{r_{C,t+1}}(\tau | MV_t) = \alpha(\tau) + \beta(\tau)MV_t$, where τ is a quantile of the one-month ahead carry trade return $r_{C,t+1}$ and MV_t is the lagged market variance. The return to the carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Market variance is the variance of the return to the market portfolio constructed as the equally weighted average of the excess return on 33 US dollar nominal exchange rates. All variables are annualized. Bootstrap t -statistics generated using 1,000 bootstrap samples are reported in parentheses. Bootstrap p -values using the double bootstrap are in brackets. The pseudo- \bar{R}^2 is computed as in Koenker and Machado (1999). The sample period runs from January 1976 to February 2009.

Predictive Quantile Regressions for the Carry Trade Return											
Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.024 (-6.227)	-0.019 (-4.742)	-0.005 (-2.100)	0.003 (1.424)	0.007 (4.711)	0.011 (7.297)	0.015 (8.886)	0.020 (10.660)	0.025 (17.106)	0.034 (15.022)	0.044 (15.569)
	[0.000]	[0.001]	[0.058]	[0.169]	[0.001]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Market Variance	-23.709 (-3.256)	-4.473 (-0.731)	-6.514 (-1.577)	-6.691 (-2.265)	-5.885 (-2.364)	-5.358 (-2.192)	-4.017 (-1.533)	-3.206 (-1.139)	-2.835 (-1.109)	-5.588 (-1.546)	-8.031 (-1.598)
	[0.006]	[0.495]	[0.120]	[0.042]	[0.035]	[0.039]	[0.147]	[0.296]	[0.259]	[0.132]	[0.137]
\bar{R}^2 (%)	0.9	0.4	0.4	0.8	1.1	0.5	-0.1	0.0	0.3	0.7	0.8

Table 6. Average Variance and Average Correlation

The table presents the regression results for the conditional quantile function: $Q_{r_{C,t+1}}(\tau | AV_t, AC_t) = \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t$, where τ is a quantile of the one-month ahead carry trade return $r_{C,t+1}$, AV_t is the lagged average variance and AC_t is the lagged average correlation. The return to the carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. With the exception of average correlation, all variables are annualized. Bootstrap t -statistics generated using 1,000 bootstrap samples are reported in parentheses. Bootstrap p -values using the double bootstrap are in brackets. The pseudo- \bar{R}^2 is computed as in Koenker and Machado (1999). The sample period runs from January 1976 to February 2009.

Predictive Quantile Regressions for the Carry Trade Return											
Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.017 (-2.246) [0.041]	-0.009 (-0.988) [0.406]	-0.005 (-1.125) [0.287]	0.004 (1.064) [0.317]	0.010 (3.349) [0.004]	0.014 (4.686) [0.000]	0.020 (6.116) [0.000]	0.025 (8.705) [0.000]	0.032 (10.367) [0.000]	0.040 (8.833) [0.000]	0.055 9.383 [0.000]
Average Variance	-12.727 (-3.242) [0.014]	-9.538 (-3.165) [0.008]	-6.558 (-2.865) [0.008]	-5.971 (-3.572) [0.001]	-4.250 (-3.087) [0.007]	-2.059 (-1.551) [0.127]	-2.616 (-1.713) [0.095]	0.223 (0.159) [0.862]	-0.960 (-0.603) [0.538]	0.194 (0.102) [0.913]	-3.603 (-1.338) [0.191]
Average Correlation	-0.010 (-0.662) [0.496]	-0.010 (-0.616) [0.613]	0.007 (0.726) [0.474]	0.001 (0.130) [0.894]	-0.003 (-0.603) [0.557]	-0.008 (-1.329) [0.164]	-0.009 (-1.362) [0.233]	-0.015 (-2.601) [0.016]	-0.015 (-2.310) [0.031]	-0.019 (-2.141) [0.051]	-0.023 (-2.084) [0.080]
\bar{R}^2 (%)	4.1	1.0	0.9	1.1	1.2	0.6	0.1	0.8	0.9	1.6	1.8

Table 7. Additional Predictive Variables

The table presents the regression results for the conditional quantile function: $Q_{r_{C,t+1}}(\tau | AV_t, AC_t, AID_t, r_{C,t}) = \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t + \beta_3(\tau) AID_t + \beta_4(\tau) r_{C,t}$, where τ is a quantile of the one-month ahead carry trade return $r_{C,t+1}$, AV_t is the lagged average variance, AC_t is the lagged average correlation, AID_t is the average interest rate differential and $r_{C,t}$ is the lagged carry trade return. The return to the carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. AID is the average interest rate differential of the quintile of currencies with the highest interest rates minus the average interest rate differential of the quintile of currencies with the lowest interest rates. With the exception of average correlation, all variables are annualized. Bootstrap t -statistics generated using 1,000 bootstrap samples are reported in parentheses. Bootstrap p -values using the double bootstrap are in brackets. The pseudo- \bar{R}^2 is computed as in Koenker and Machado (1999). The sample period runs from January 1976 to February 2009.

Predictive Quantile Regressions for the Carry Trade Return											
Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.013 (-1.521) [0.162]	-0.015 (-2.008) [0.100]	-0.009 (-1.562) [0.185]	-0.003 (-0.695) [0.554]	0.003 (1.075) [0.252]	0.007 (2.303) [0.036]	0.012 (3.176) [0.007]	0.019 (5.515) [0.000]	0.028 (7.979) [0.000]	0.037 (8.012) [0.000]	0.054 (7.882) [0.000]
Average Variance	-10.434 (-2.568) [0.028]	-5.689 (-1.752) [0.090]	-5.958 (-2.472) [0.033]	-5.457 (-3.018) [0.007]	-3.704 (-2.499) [0.022]	-4.086 (-2.819) [0.010]	-2.063 (-1.394) [0.204]	-2.720 (-1.808) [0.079]	-1.874 (-1.136) [0.248]	-0.735 (-0.334) [0.736]	-2.868 (-0.946) [0.382]
Average Correlation	-0.014 (-0.972) [0.336]	-0.012 (-0.961) [0.428]	-0.002 (-0.181) [0.866]	0.003 (0.380) [0.722]	-0.004 (-0.800) [0.401]	-0.003 (-0.489) [0.611]	-0.006 (-1.061) [0.327]	-0.010 (-1.727) [0.102]	-0.016 (-2.597) [0.028]	-0.019 (-2.396) [0.039]	-0.026 (-2.125) [0.081]
Average Interest Differential	-0.398 (-0.936) [0.331]	0.218 (0.583) [0.622]	0.519 (1.946) [0.083]	0.463 (2.372) [0.042]	0.499 (3.295) [0.000]	0.443 (2.978) [0.007]	0.440 (2.466) [0.023]	0.502 (3.032) [0.004]	0.403 (2.515) [0.027]	0.226 (0.956) [0.339]	0.040 (0.128) [0.905]
Lagged Carry Return	0.255 (1.892) [0.084]	0.296 (2.722) [0.021]	0.154 (2.002) [0.060]	0.101 (1.686) [0.128]	0.089 (1.953) [0.051]	0.067 (1.461) [0.153]	0.035 (0.663) [0.553]	0.015 (0.300) [0.749]	-0.043 (-0.810) [0.431]	0.013 (0.193) [0.853]	0.048 (0.476) [0.692]
\bar{R}^2 (%)	4.8	3.1	2.1	2.4	2.9	2.0	1.1	1.6	1.1	1.6	1.4

Table 8. Weighted Average Variance and Weighted Average Correlation

The table presents the regression results for the conditional quantile function: $Q_{r_{C,t+1}^W}(\tau | AV_t^W, AC_t^W) = \alpha(\tau) + \beta_1(\tau) AV_t^W + \beta_2(\tau) AC_t^W$, where τ is a quantile of the one-month ahead weighted carry trade return $r_{C,t+1}^W$, AV_t^W is the lagged weighted average variance and AC_t^W is the lagged weighted average correlation. All variables are weighted in order to account for the effect of numeraire in both the carry trade return and the risk measures. The weights are based on the Special Drawing Rights (SDR) of the International Monetary Fund and are as follows: 41.9% on the US dollar-denominated measures, 37.4% on the Euro-denominated measures, 11.3% on the UK pound-denominated measures and 9.4% on the Japanese yen-denominated measures. The weighted return to the carry trade portfolio uses the SDR weights across the four numeraires, and for a given numeraire is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. Weighted average variance is the weighted average across the four numeraires, where for a given numeraire average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. Weighted average correlation is the weighted average across the four numeraires, where for a given numeraire average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. With the exception of average correlation, all variables are annualized. Bootstrap t -statistics generated using 1,000 bootstrap samples are reported in parentheses. Bootstrap p -values using the double bootstrap are in brackets. The pseudo- \bar{R}^2 is computed as in Koenker and Machado (1999). The sample period runs from January 1976 to February 2009.

Predictive Quantile Regressions for the Weighted Carry Trade Return											
Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.024 (-1.781) [0.074]	-0.015 (-1.011) [0.398]	-0.009 (-1.151) [0.237]	0.003 (0.438) [0.650]	0.014 (2.742) [0.016]	0.018 (3.717) [0.001]	0.026 (4.449) [0.000]	0.037 (7.676) [0.000]	0.041 (7.786) [0.000]	0.053 (6.618) [0.000]	0.060 5.114 [0.000]
Weighted Ave. Variance	-13.230 (-3.216) [0.013]	-14.365 (-4.174) [0.007]	-6.793 (-2.856) [0.011]	-6.619 (-3.418) [0.006]	-4.924 (-2.985) [0.013]	-3.755 (-2.407) [0.028]	-2.663 (-1.718) [0.093]	-1.633 (-1.131) [0.227]	-1.916 (-1.189) [0.211]	-1.560 (-0.747) [0.420]	-1.273 (-0.430) [0.723]
Weighted Ave. Correlation	0.004 (0.143) [0.846]	0.011 (0.331) [0.786]	0.017 (0.923) [0.335]	0.004 (0.304) [0.735]	-0.012 (-1.033) [0.306]	-0.016 (-1.485) [0.143]	-0.024 (-1.834) [0.081]	-0.038 (-3.622) [0.001]	-0.038 (-3.210) [0.008]	-0.048 (-2.768) [0.008]	-0.045 (-1.741) [0.162]
\bar{R}^2 (%)	4.4	1.8	1.1	1.0	1.1	0.7	0.6	1.4	1.7	1.4	1.0

Table 9. VIX, VXY and FX Risk Measures

The table presents ordinary least squares results for regressions of the VIX and VXY indices on monthly FX risk measures. The VIX index is based on the 1-month model-free implied volatility of the S&P 500 equity index. The VXY index is based on the 3-month implied volatility of at-the-money-forward options on the G-7 currencies. Market variance is the variance of the return to the market portfolio constructed as the equally weighted average of the excess returns on 33 US dollar nominal exchange rates. Average variance is the equally weighted cross-sectional average of the variances of exchange rate excess returns. Average correlation is the equally weighted cross-sectional average of the pairwise correlations of exchange rate excess returns. With the exception of average correlation, all variables are annualized. Newey-West (1987) t -statistics with five lags are reported in parentheses. The sample period for the VIX begins in January 1990 and for the VXY in January 1992. The sample period for all indices ends in February 2009.

Panel A: Regressions for the VIX					
	(1)	(2)	(3)	(4)	(5)
Constant	0.014 (13.962)	0.013 (17.851)	0.021 (16.058)	0.019 (15.103)	0.015 (15.302)
Market Variance	5.695 (2.113)				
Average Variance		3.556 (6.559)		4.003 (7.729)	
Average Correlation			-0.012 (-3.862)	-0.016 (-5.113)	
(Average Variance) \times (Average Correlation)					5.496 (2.159)
\bar{R}^2 (%)	13.5	24.7	6.9	37.6	12.4
Panel B: Regressions for the VXY					
	0.008 (35.786)	0.007 (30.391)	0.008 (19.371)	0.008 (19.868)	0.008 (35.791)
Constant					
Market Variance	3.305 (10.540)				
Average Variance		1.462 (7.142)		1.473 (6.869)	
Average Correlation			0.001 (1.269)	-0.001 (-0.519)	
(Average Variance) \times (Average Correlation)					3.201 (9.628)
\bar{R}^2 (%)	53.0	54.5	0.5	54.4	49.6

Table 10. VIX and VXY

The table presents the regression results for two conditional quantile functions. The first one is: $Q_{r_{C,t+1}}(\tau | VIX_t) = \alpha(\tau) + \beta(\tau)VIX_t$, where τ is a quantile of the one-month ahead carry trade return $r_{C,t+1}$ and VIX_t is the lagged VIX index. The second one is: $Q_{r_{C,t+1}}(\tau | VXY_t) = \alpha(\tau) + \beta(\tau)VXY_t$, where VXY_t is the lagged VXY index. The return to the carry trade portfolio is constructed every month by going long on the quintile of currencies with the highest interest rates and going short on the bottom quintile. The VIX index is based on the 1-month model-free implied volatility of the S&P 500 equity index. The VXY index is based on the 3-month implied volatility of at-the-money-forward options on the G-7 currencies. All variables are annualized. Bootstrap t -statistics generated using 1,000 bootstrap samples are reported in parentheses. Bootstrap p -values using the double bootstrap are in brackets. The pseudo- \bar{R}^2 is computed as in Koenker and Machado (1999). The sample period runs from January 1976 to February 2009.

Panel A: The Carry Trade Return on VIX											
Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.018 (-1.382) [0.195]	-0.021 (-1.875) [0.097]	-0.004 (-0.507) [0.621]	0.004 (0.686) [0.464]	0.004 (0.929) [0.359]	0.009 (2.013) [0.090]	0.014 (3.670) [0.003]	0.016 (3.576) [0.006]	0.019 (4.077) [0.003]	0.025 (3.972) [0.001]	0.023 (3.521) [0.006]
VIX	-0.103 (-1.620) [0.140]	-0.027 (-0.511) [0.605]	-0.023 (-0.580) [0.584]	-0.018 (-0.605) [0.511]	0.006 (0.283) [0.745]	0.008 (0.379) [0.699]	-0.001 (-0.029) [0.957]	0.015 (0.731) [0.520]	0.023 (1.057) [0.341]	0.026 (0.870) [0.444]	0.070 (1.988) [0.065]
\bar{R}^2 (%)	0.8	-0.2	-0.2	-0.3	-0.4	-0.4	-0.4	-0.2	0.2	0.1	2.1
Panel B: The Carry Trade Return on VXY											
Constant	-0.051 (-1.646) [0.150]	-0.028 (-1.660) [0.098]	0.003 (0.241) [0.795]	0.009 (0.953) [0.356]	0.012 (1.347) [0.199]	0.023 (2.379) [0.028]	0.028 (3.239) [0.004]	0.039 (4.258) [0.001]	0.027 (3.093) [0.010]	0.041 (4.029) [0.003]	0.024 (1.414) [0.173]
VXY	0.183 (0.628) [0.591]	0.073 (0.447) [0.615]	-0.096 (-0.777) [0.439]	-0.102 (-1.139) [0.282]	-0.082 (-0.959) [0.355]	-0.123 (-1.354) [0.185]	-0.140 (-1.704) [0.105]	-0.190 (-2.209) [0.044]	0.005 (0.061) [0.916]	-0.081 (-0.828) [0.373]	0.160 (0.956) [0.352]
\bar{R}^2 (%)	-0.2	-0.2	0.1	0.1	0.4	1.0	1.0	0.4	-0.5	-0.3	1.2

Table 11. Out-of-Sample Augmented Carry Trade Strategies

The table presents the out-of-sample performance of augmented carry trade strategies that condition on the movement of average variance and/or average correlation with and without transaction costs. Panel A assumes no transaction costs, whereas Panel B implements the carry trade strategies using the transaction costs listed in Table 1. Average variance is the equally weighted cross-sectional average of the monthly variances of all exchange rate excess returns. The average variance strategy implements the following rule at each time t : for the carry trade returns that are lower than the τ -quantile of the distribution, if average variance has increased from $t-1$ to t , we close the carry trade positions and thus receive a return of zero at $t+1$; otherwise we execute the standard carry trade. Average correlation is the equally weighted cross-sectional average of the pairwise monthly correlations of all exchange rate excess returns. The average correlation strategy implements the following rule at each time t : for the carry trade returns that are higher than the $1-\tau$ quantile, if average correlation has decreased from $t-1$ to t , we double the carry trade positions and thus receive twice the carry return at $t+1$; otherwise we execute the standard carry trade. The combined average variance and average correlation strategy makes the following decision at each time t : for the carry trade returns that are lower than the τ -quantile, if average variance has increased from $t-1$ to t , we close the carry trade positions and thus receive a return of zero; and for the carry trade returns that are higher than the $1-\tau$ quantile, if average correlation has decreased from $t-1$ to t , we double the carry trade positions and thus receive twice the carry return; otherwise we execute the standard carry trade. The mean, standard deviation and Sharpe ratio are reported in annualized terms. The turnover ratio is equal to the percentage of the currencies that on average are traded every month. The sample period runs from January 1976 to February 2009. All strategies move forward recursively starting 3 years after the beginning of the sample so that the first observation is for January 1979.

Panel A: No Transaction Costs										
Quantile	<i>Carry Trade</i>	<i>Average Variance Strategy</i>								
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Mean	0.084	0.094	0.091	0.085	0.079	0.073	0.069	0.066	0.059	0.054
St. Dev.	0.079	0.071	0.069	0.066	0.065	0.064	0.061	0.060	0.057	0.056
Sharpe ratio	1.070	1.314	1.326	1.279	1.213	1.142	1.120	1.098	1.022	0.968
Turnover	0.171	0.208	0.264	0.310	0.346	0.387	0.421	0.444	0.485	0.514
<i>Average Correlation Strategy</i>										
Mean	0.084	0.146	0.147	0.139	0.130	0.122	0.118	0.114	0.107	0.097
St. Dev.	0.079	0.121	0.115	0.111	0.107	0.104	0.100	0.096	0.091	0.085
Sharpe ratio	1.070	1.208	1.284	1.262	1.212	1.174	1.186	1.186	1.182	1.134
Turnover	0.171	0.557	0.511	0.454	0.412	0.370	0.340	0.294	0.252	0.214
<i>Combined Average Variance and Average Correlation Strategy</i>										
Mean	0.084					0.104	0.106	0.107	0.106	0.099
St. Dev.	0.079					0.095	0.092	0.088	0.084	0.081
Sharpe ratio	1.070					1.095	1.152	1.211	1.264	1.224
Turnover	0.171					0.589	0.522	0.440	0.352	0.257

(continued)

Panel B: With Transaction Costs

Quantile	<i>Carry Trade</i>	<i>Average Variance Strategy</i>								
		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Mean	0.058	0.070	0.068	0.062	0.058	0.055	0.053	0.050	0.043	0.040
St. Dev.	0.079	0.071	0.069	0.066	0.065	0.063	0.061	0.060	0.057	0.055
Sharpe ratio	0.741	0.977	0.983	0.941	0.892	0.866	0.869	0.836	0.761	0.729
Turnover	0.171	0.208	0.259	0.306	0.344	0.385	0.417	0.442	0.485	0.513
<i>Average Correlation Strategy</i>										
Mean	0.058	0.109	0.109	0.107	0.097	0.091	0.090	0.085	0.079	0.070
St. Dev.	0.079	0.121	0.115	0.111	0.107	0.103	0.099	0.095	0.090	0.085
Sharpe ratio	0.741	0.903	0.949	0.971	0.913	0.881	0.904	0.891	0.876	0.823
Turnover	0.171	0.557	0.513	0.458	0.411	0.375	0.341	0.296	0.252	0.215
<i>Combined Average Variance and Average Correlation Strategy</i>										
Mean	0.058					0.081	0.082	0.081	0.080	0.074
St. Dev.	0.079					0.094	0.091	0.088	0.084	0.081
Sharpe ratio	0.741					0.853	0.895	0.926	0.959	0.914
Turnover	0.171					0.594	0.521	0.438	0.347	0.259

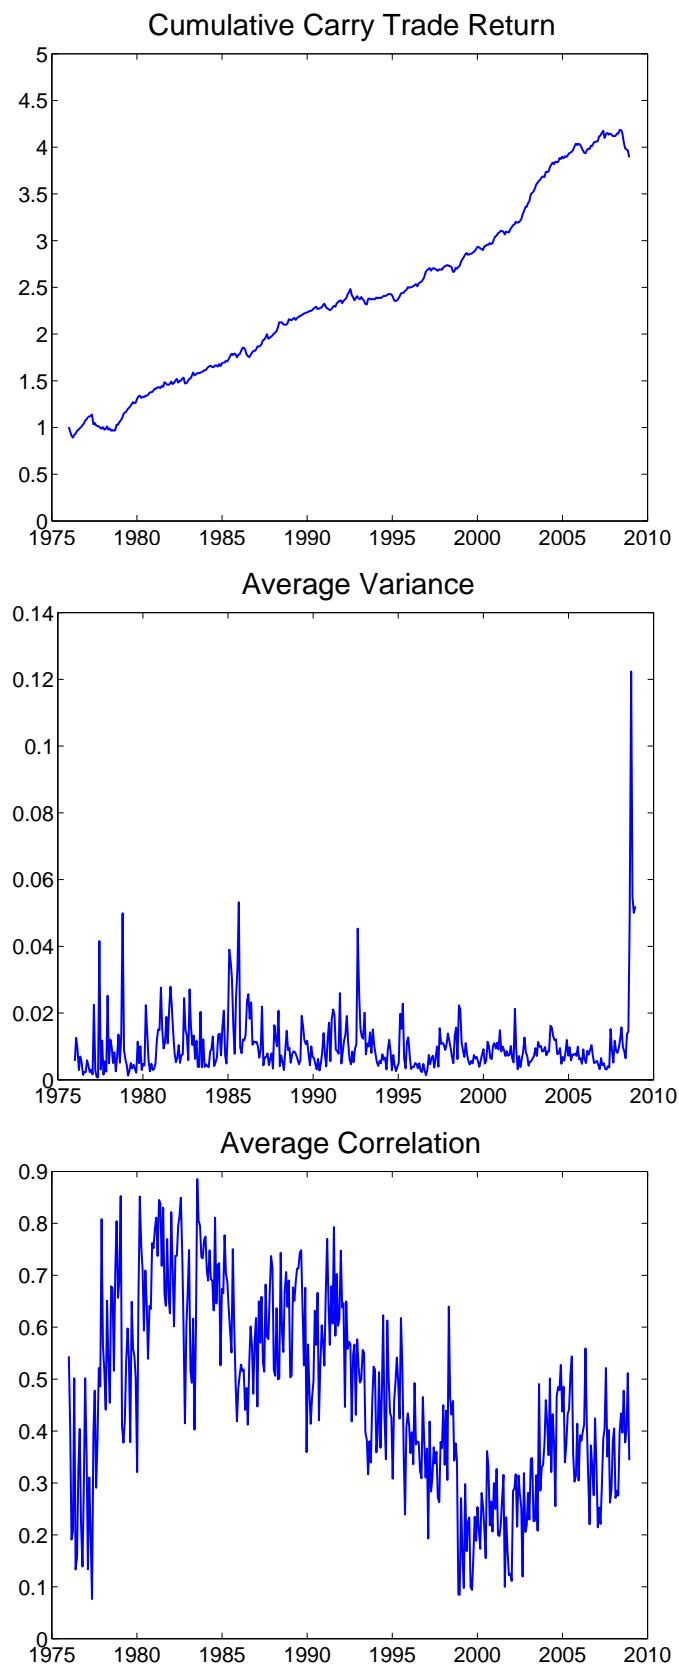


Figure 1. Carry Trade Return, Average Variance and Average Correlation

This figure displays the time series of the cumulative carry trade return, FX average variance and FX average correlation from January 1976 to February 2009.

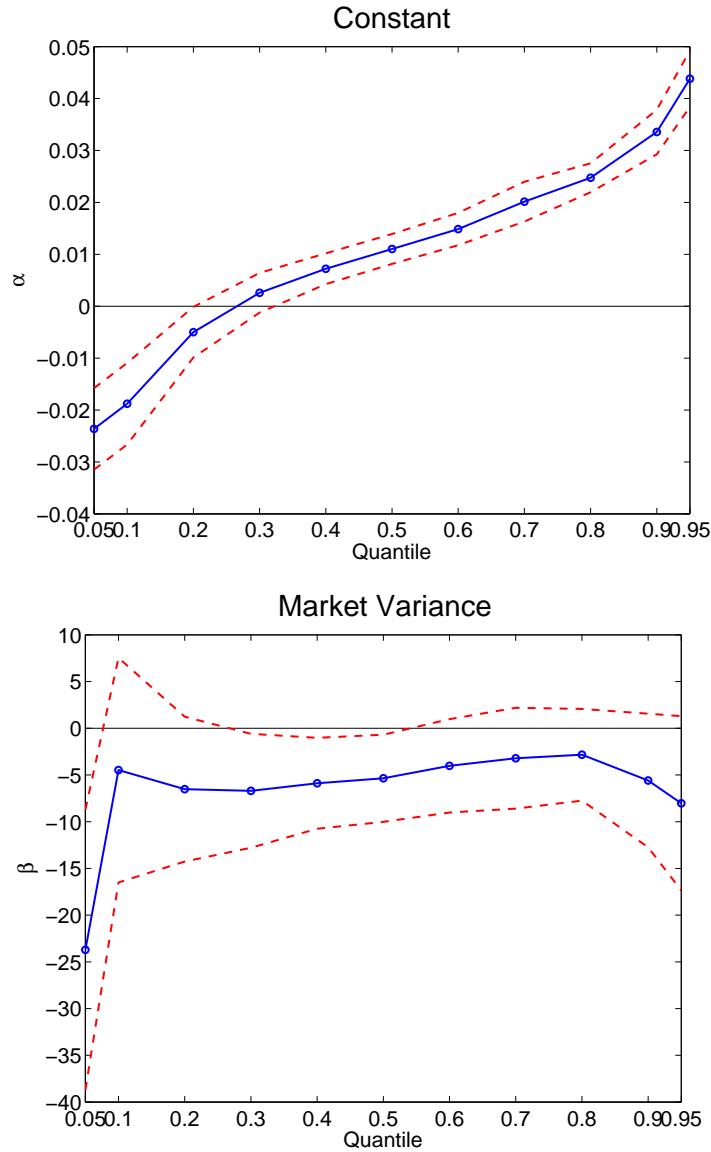


Figure 2. Market Variance

This figure shows the parameter estimates of the predictive quantile regression of the one-month-ahead carry trade return on the lagged market variance. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

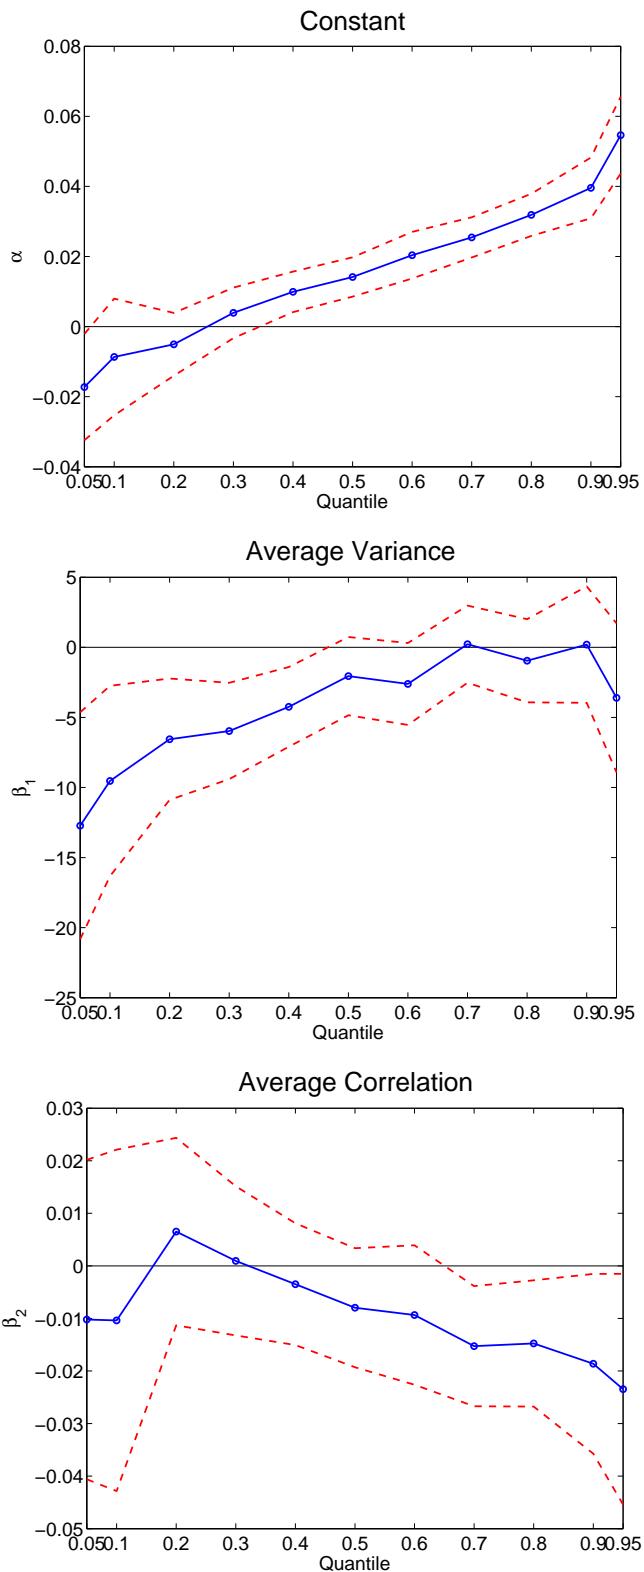


Figure 3. Average Variance and Average Correlation

This figure illustrates the parameter estimates of the predictive quantile regression of the one-month-ahead carry trade return on the lagged average variance and average correlation. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

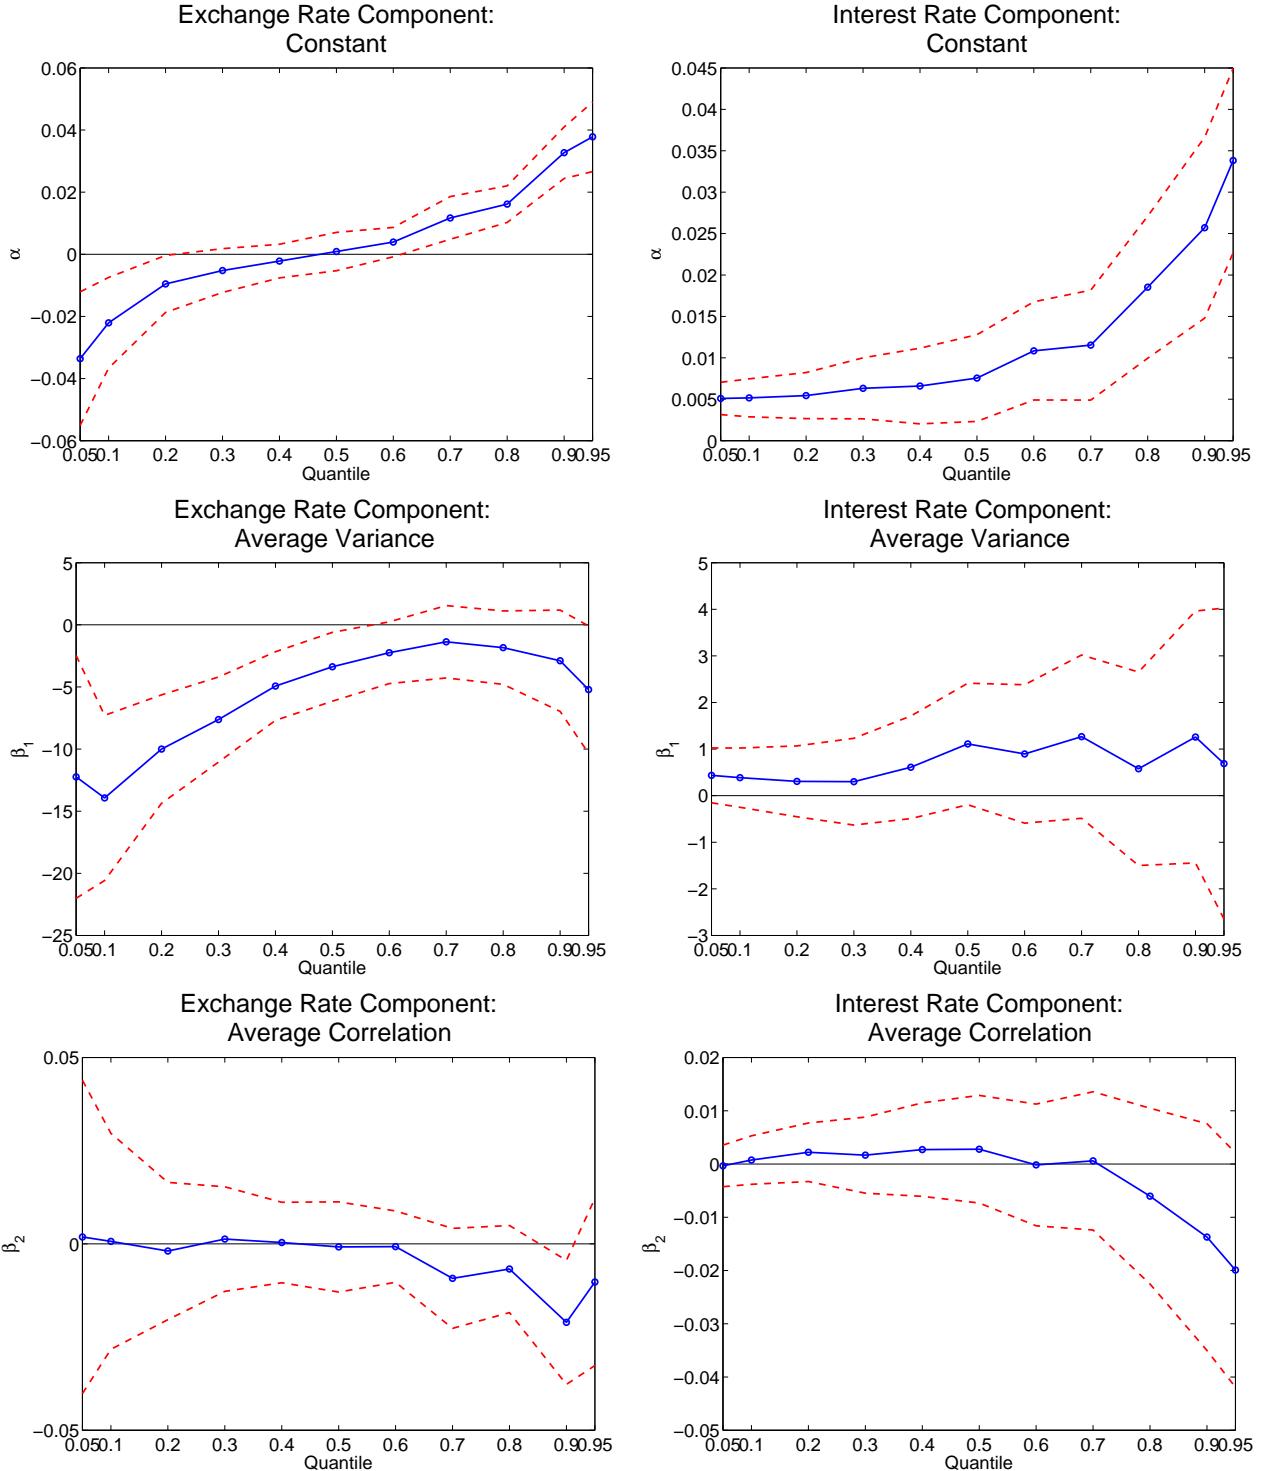


Figure 4. The Exchange Rate and Interest Rate Components of the Carry Trade

This figure exhibits the parameter estimates of two sets of predictive quantile regressions. The left panel shows the results for the one-month-ahead exchange rate component of the carry trade return on the lagged average variance and lagged average correlation. The right panel shows the results for the interest rate component of the carry trade return on the lagged average variance and average correlation. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

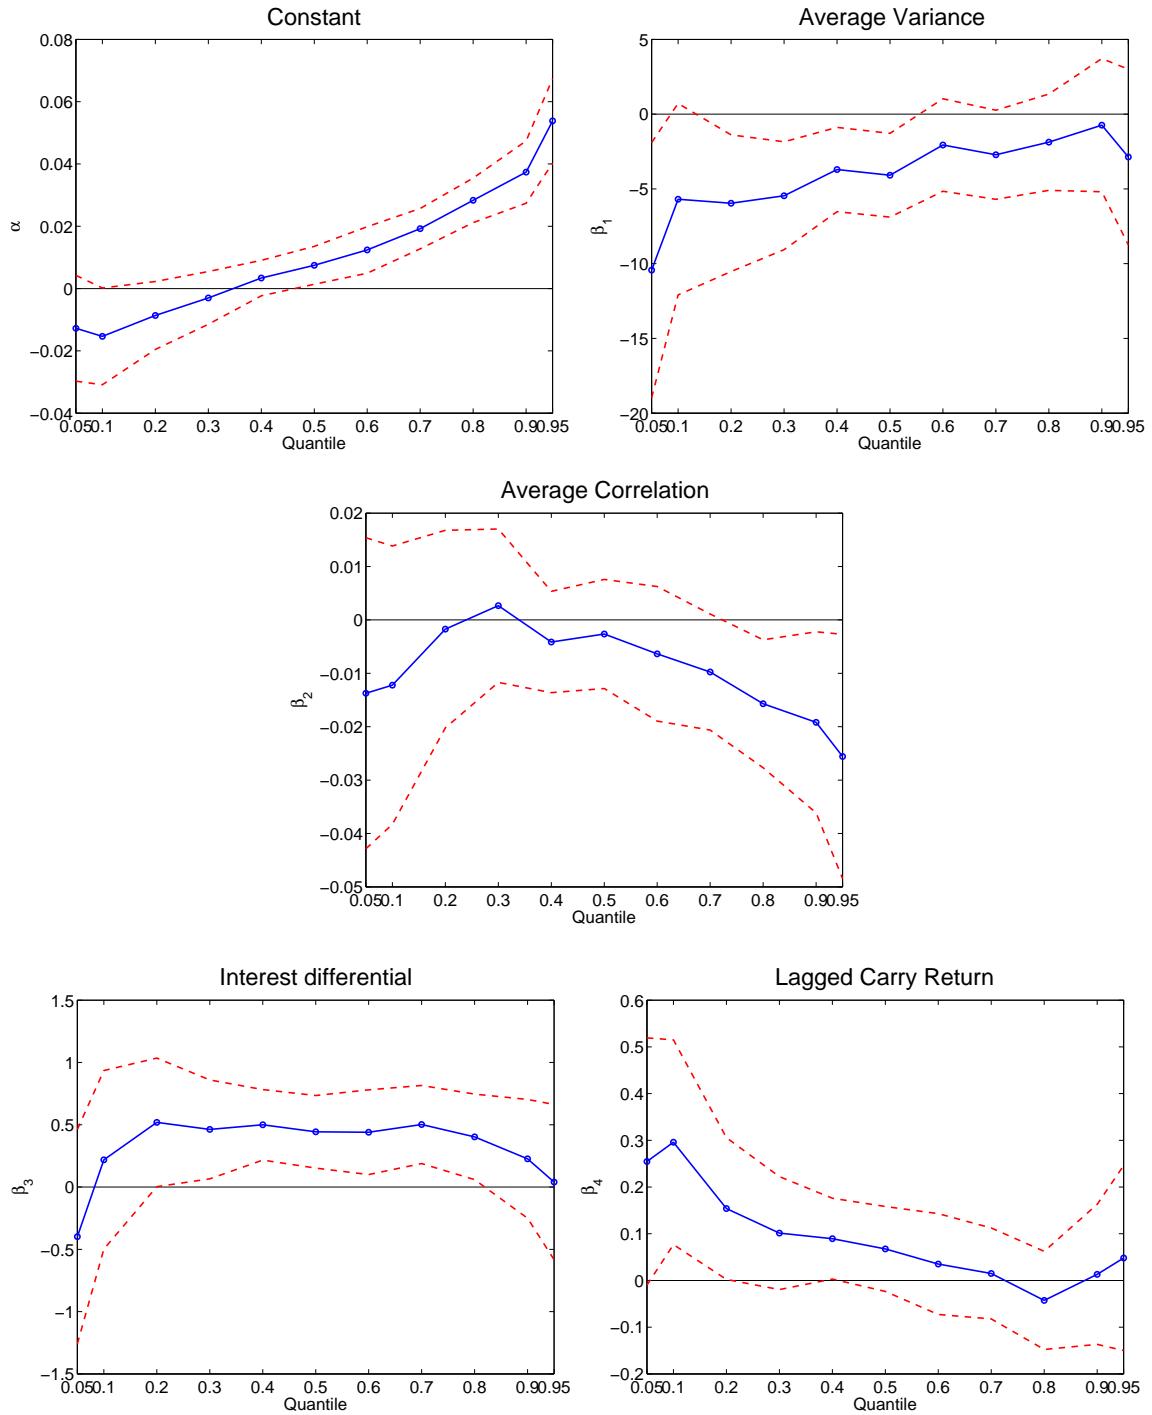


Figure 5. Additional Predictive Variables

This figure illustrates the parameter estimates of the predictive quantile regression of the one-month-ahead carry trade return on the lagged average variance, lagged average correlation and two additional predictive variables: the interest rate differential and the lagged carry trade return. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

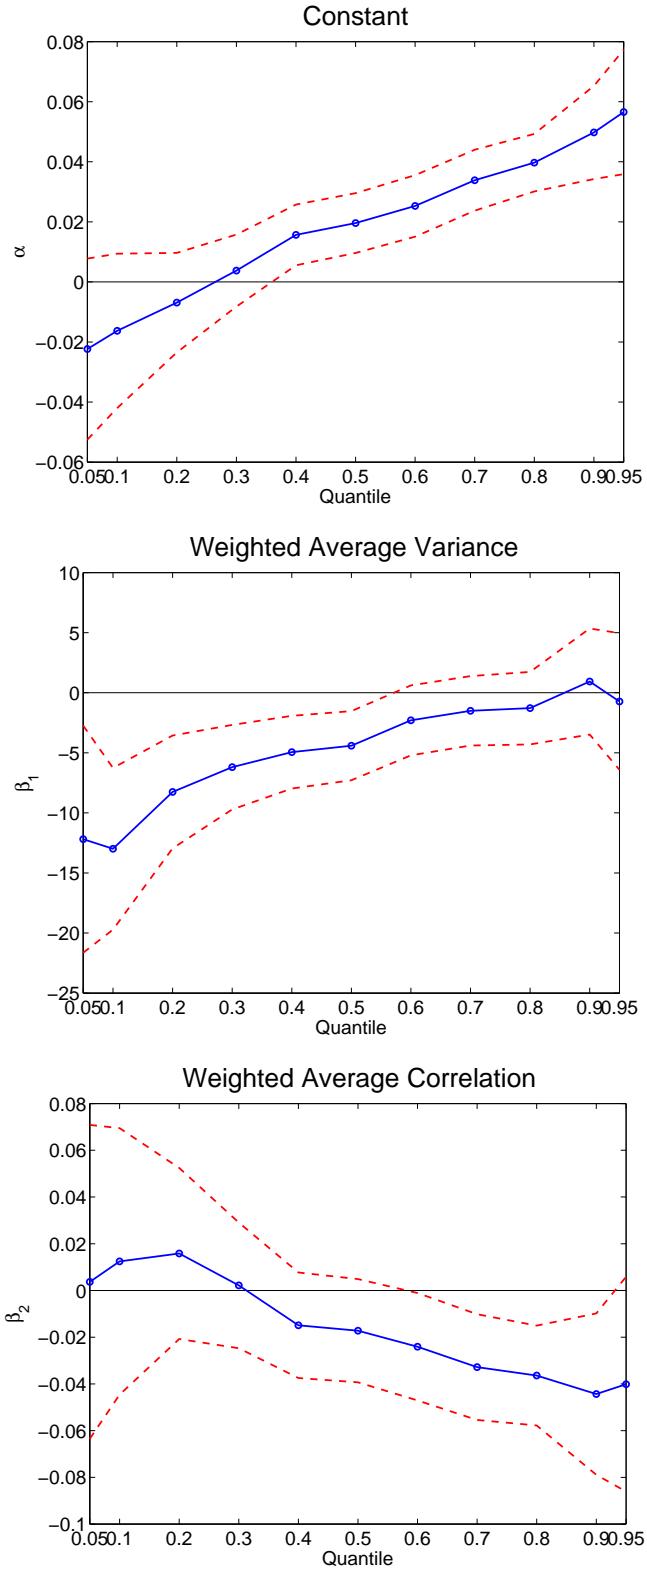


Figure 6. The Numeraire Effect

This figure displays the parameter estimates of the predictive quantile regression of the weighted one-month-ahead carry trade return on the lagged weighted average variance and lagged weighted average correlation. The weighted variables account for the numeraire effect using the IMF weights for Special Drawing Rights: 41.9% on the US dollar, 37.4% on the Euro, 11.3% on the UK pound and 9.4% on the Japanese yen. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

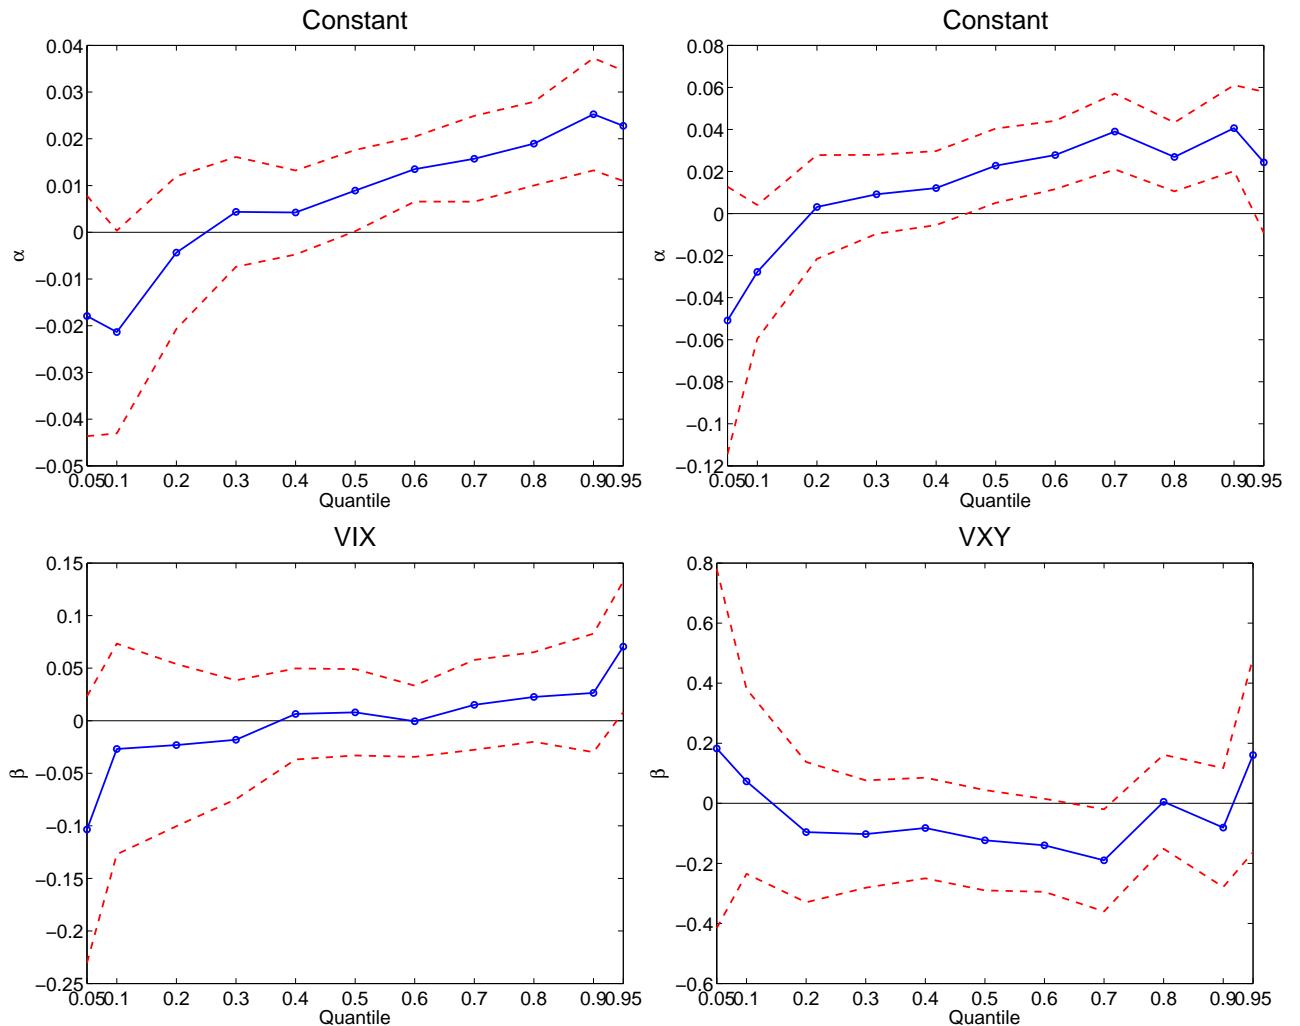


Figure 7. VIX and VXY

This figure exhibits the parameter estimates of two sets of predictive quantile regressions. The left panel shows the results for the one-month-ahead carry trade return on the lagged VIX index. The right panel shows the results for the one-month-ahead carry trade return on the lagged VXY index. The dashed lines indicate the 95% confidence interval based on bootstrap standard errors.

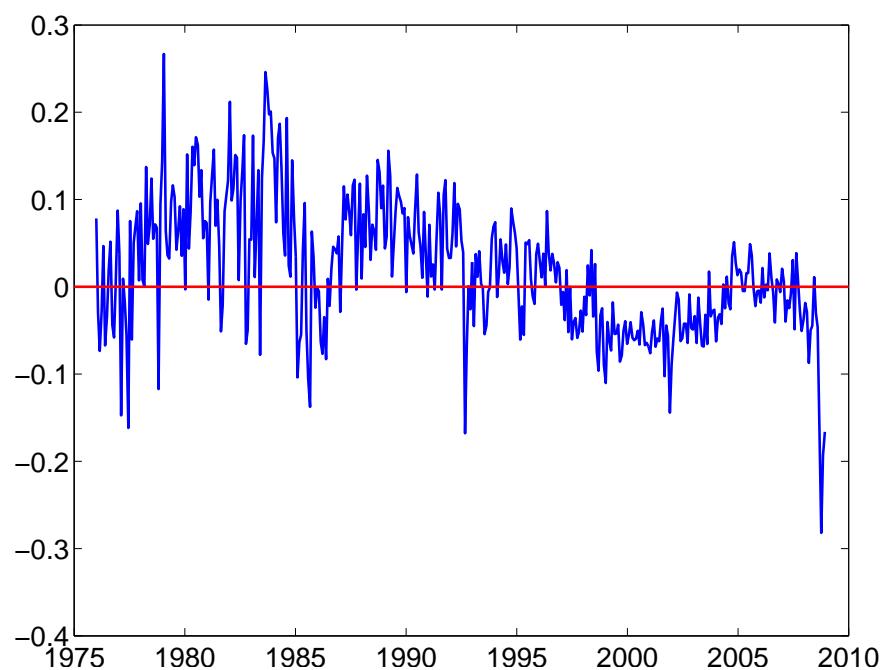


Figure 8. Conditional Skewness

This figure shows the time-variation of the robust conditional skewness measure of Kim and White (2004) based on the predictive quantile regressions. The measure uses the inter-quartile range ($a = 0.25$).

References

- Akram, F., D. Rime, and L. Sarno (2008). “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics* **76**, 237–253.
- Ang, A., R.J. Hodrick, Y. Xing, and X. Zhang (2006). “The Cross-Section of Volatility and Expected Returns,” *Journal of Finance* **61**, 259–299.
- Ang, A., R.J. Hodrick, Y. Xing, and X. Zhang (2009). “High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence,” *Journal of Financial Economics* **91**, 1–23.
- Bali, T.G. (2008). “The Intertemporal Relation between Expected Returns and Risk,” *Journal of Financial Economics* **87**, 101–131.
- Bali, T.G., N. Cakici, X. Yan, and Z. Zhang (2005). “Does Idiosyncratic Risk Really Matter?” *Journal of Finance* **60**, 905–929.
- Bali, T.G., and K. Yilmaz (2012). “The Intertemporal Relation between Expected Return and Risk on Currency,” *Handbook of Financial Econometrics and Statistics*, ed. C.-F. Lee. New York: Springer (forthcoming).
- Bandi, F.M., and B. Perron (2008). “Long-Run Risk-Return Trade-offs,” *Journal of Econometrics* **143**, 349–374.
- Bekaert, G., (1995). “The Time Variation of Expected Returns and Volatility in Foreign Exchange Markets,” *Journal of Business and Economic Statistics* **13**, 397–408.
- Bilson, J.F.O. (1981). “The ‘Speculative Efficiency’ Hypothesis,” *Journal of Business* **54**, 435–451.
- Bowley, A.L. (1920). *Elements of Statistics*. Scribner’s: New York.
- Brandt, M.W., and Q. Kang (2004). “On the Relationship between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach,” *Journal of Financial Economics* **72**, 217–257.

- Brunnermeier, M.K., S. Nagel, and L.H. Pedersen (2009). “Carry Trades and Currency Crashes,” *NBER Macroeconomics Annual 2008*, 313–347.
- Brunnermeier, M.K. and L.H. Pedersen (2009). “Market Liquidity and Funding Liquidity,” *Review of Financial Studies* **22**, 2201-2238.
- Burnside, C. (2012). “Carry Trades and Risk,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. London: Wiley (forthcoming).
- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). “Do Peso Problems Explain the Returns to the Carry Trade?” *Review of Financial Studies* **24**, 853–891.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2008). “Carry Trade: The Gains of Diversification,” *Journal of the European Economic Association* **6**, 581–588.
- Cenesizoglu, T., and A. Timmermann (2010). “Do Return Prediction Models Add Economic Value?” Unpublished working paper, HEC Montreal.
- Chan, Y.L., and L. Kogan (2002). “Catching up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices,” *Journal of Political Economy* **110**, 1255–1285.
- Chan, K.C., G.A. Karolyi, and R.M. Stulz (1992). “Global Financial Markets and the Risk Premium on U.S. Equity,” *Journal of Financial Economics* **32**, 137–167.
- Chen, Z., and R. Petkova (2010). “Does Idiosyncratic Volatility Proxy for Risk Exposure?” Unpublished working paper, Texas A&M University.
- Christiansen, C. (2011). “Intertemporal Risk-Return Trade-Off in Foreign Exchange Rates,” *Journal of International Financial Markets, Institutions and Money* (forthcoming).
- Christiansen, C., A. Ranaldo, and P. Söderlind (2011). “The Time-Varying Systematic Risk of Carry Trade Strategies,” *Journal of Financial and Quantitative Analysis* (forthcoming).
- Constantinides, G.M. (2002). “Rational Asset Prices,” *Journal of Finance* **57**, 1567–1591.
- Constantinides, G.M., and D. Duffie (1996). “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy* **104**, 219–240.

- Della Corte, P., L. Sarno, and I. Tsiakas (2009). “An Economic Evaluation of Empirical Exchange Rate Models,” *Review of Financial Studies* **22**, 3491–3530.
- Diebold, F.X., and M. Nerlove (1989). “The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model,” *Journal of Applied Econometrics* **4**, 1–21.
- Engel, C. (1996). “The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence,” *Journal of Empirical Finance* **3**, 123–192.
- Fama, E.F. (1984). “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* **14**, 319–338.
- Fama, E.F., and K. French (1993). “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics* **33**, 3–56.
- Farhi, E., S.P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009). “Crash Risk in Currency Markets,” Unpublished working paper, Harvard University.
- Fitzenberger, B., (1997). “The Moving Block Bootstrap and Robust Inference for Linear Least Squares and Quantile Regressions,” *Journal of Econometrics* **82**, 235–287.
- French, K.R., W. Schwert, and R.F. Stambaugh (1987). “Expected Stock Returns and Volatility,” *Journal of Financial Economics* **19**, 3–29.
- Ghysels, E., P. Santa-Clara, and R. Valkanov (2005). “There is a Risk-Return Trade-off After All,” *Journal of Financial Economics* **76**, 509–548.
- Glosten, L.R., R. Jagannathan, and D.E. Runkle (1993). “On the Relationship Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks,” *Journal of Finance* **48**, 1779–1801.
- Goyal A., and P. Santa-Clara (2003). “Idiosyncratic Risk Matters!” *Journal of Finance* **58**, 975–1007.
- Goyal, A., and A. Saretto (2009). “Cross-Section of Option Returns and Volatility,” *Journal of Financial Economics* **94**, 310–326.

- Jurek, J., (2009). “Crash-Neutral Currency Carry Trades,” Unpublished working paper, Princeton University.
- Kim, T-H., and H. White (2004). “On More Robust Estimation of Skewness and Kurtosis,” *Finance Research Letters* **1**, 56–73.
- Koenker, R. (2005). *Quantile Regression*. New York: Cambridge University Press.
- Koenker, R., and J. Machado (1999). “Goodness of Fit and Related Inference Processes for Quantile Regression,” *Journal of the American Statistical Association* **94**, 1296–1310.
- Krishnan, C.N.V., R. Petkova, and P. Ritchken (2009). “Correlation Risk,” *Journal of Empirical Finance* **16**, 353–367.
- Lustig, H., N. Roussanov, and A. Verdelhan (2010). “Countercyclical Currency Risk Premia,” Unpublished working paper, MIT.
- Lustig, H., N. Roussanov, and A. Verdelhan (2011). “Common Risk Factors in Currency Markets,” *Review of Financial Studies* (forthcoming).
- MacKinnon, J., (2007). “Bootstrap Hypothesis Testing,” QED Working Paper No. 1127.
- Mancini-Griffoli, T., and A. Ranaldo (2011). “Limits to Arbitrage during the Crisis: Funding Liquidity Constraints and Covered Interest Parity,” Swiss National Bank working paper 2010-14.
- Mayhew, S. (2002). “Competition, Market Structure, and Bid-Ask Spreads in Stock Option Markets,” *Journal of Finance* **57**, 931–958.
- Meese, R.A., and K. Rogoff (1983). “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics* **14**, 3–24.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012). “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance* (forthcoming).
- Merton, R.C., (1973). “An Intertemporal Capital Asset Pricing Model,” *Econometrica* **41**, 867–887.

- Merton, R.C., (1980). “On Estimating the Expected Return on the Market: An Exploratory Investigation,” *Journal of Financial Economics* **8**, 323–361.
- Neely, C.J., P.A. Weller, and J.M. Ulrich (2009). “The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market,” *Journal of Financial and Quantitative Analysis* **44**, 467–488.
- Newey, W.K., and K.D. West (1987). “A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica* **55**, 703–708.
- Patton, A., D. Politis, and H. White (2009). “Correction to ‘Automatic Block-Length Selection for the Dependent Bootstrap’,” *Econometric Reviews* **28**, 372–375.
- Politis, D., and J.P. Romano (1992). “A Circular Block-Resampling Procedure for Stationary Data,” in LePage, R., and L. Billard (eds.) *Exploring the Limits of the Bootstrap*. New York: Wiley.
- Politis, D., and H. White (2004). “Automatic Block-Length Selection for the Dependent Bootstrap,” *Econometric Reviews* **23**, 53–70.
- Pollet, J.M., and M. Wilson (2010). “Average Correlation and Stock Market Returns,” *Journal of Financial Economics* **96**, 364–380.
- Portnoy, S., and R. Koenker (1997). “The Gaussian Hare and the Laplacian Tortoise: Computability of Squared-Error Versus Absolute-Error Estimators, with Discussion,” *Statistical Science* **12**, 279–300.
- Ranaldo, A., and P. Söderlind (2010). “Safe Haven Currencies,” *Review of Finance* **14**, 385–407.
- Sarkissian, S. (2003). “Incomplete Consumption Risk Sharing and Currency Risk Premiums,” *Review of Financial Studies* **16**, 983–1005.