

Foreign Exchange Risk and the Predictability of Carry Trade Returns*

Gino Cenedese
Bank of England

Lucio Sarno
Cass Business School and CEPR

Ilias Tsiakas
University of Guelph

February 2014

Abstract

This paper provides an empirical investigation of the time-series predictive ability of foreign exchange risk measures on the return to the carry trade, a popular investment strategy that borrows in low-interest currencies and lends in high-interest currencies. Using quantile regressions, we find that higher market variance is significantly related to large future carry trade losses, which is consistent with the unwinding of the carry trade in times of high volatility. The decomposition of market variance into average variance and average correlation shows that the predictive power of market variance is primarily due to average variance since average correlation is not significantly related to carry trade returns. Finally, a new version of the carry trade that conditions on market variance generates performance gains net of transaction costs.

Keywords: Exchange Rates; Carry Trade; Market Variance; Average Variance; Average Correlation; Quantile Regression.

JEL Classification: F31; G15; G17.

***Acknowledgements:** This paper is forthcoming in the *Journal of Banking and Finance*. The authors are indebted for constructive comments to Ike Mathur (editor), an anonymous referee, Pasquale Della Corte, Walter Distaso, Miguel Ferreira, Peter Gruber, Alex Maynard, Lukas Menkhoff, Francesco Nucci, Angelo Ranaldo, Pedro Santa-Clara, Maik Schmeling, Adrien Verdelhan, Paolo Volpin and Filip Žikeš 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; as well as seminar participants at Goldman Sachs, Carleton University, Leibniz University Hannover and McMaster University. 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 at the core of active currency management and 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 equal to 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. 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.¹

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 (2012a) find that the large average carry trade payoffs are compensation for exposure to global FX volatility risk. Christiansen, Rinaldo and Söderlind (2011) further show that the level of FX volatility also affects the risk exposure of carry trade returns to stock and bond markets. Mueller, Stathopoulos and Vedolin (2012) show that FX excess returns also carry a negative price of correlation risk. Lustig, Roussanov and Verdelhan (2011) identify a slope factor in the cross section of FX portfolios, constructed in similar fashion to the Fama and French (1993) “high-minus-low” factor. Burnside, Eichenbaum, Kleshchelski and Rebelo (2011) argue that the high carry trade payoffs reflect a peso problem, which is a low probability of large negative payoffs. Finally, 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

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; Sarno, 2005).

when speculators face funding liquidity constraints.²

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 the large future losses or gains of the carry trade located in the left or right tail of the return distribution respectively. Third, we define a set of FX risk measures that captures 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, fourth, we assess the economic gains of our analysis by designing a new version of the carry trade strategy that conditions on the FX risk measures.

The empirical analysis is organized as follows. The first step is to form two carry trade portfolios, which are rebalanced monthly: an advanced economy portfolio that includes ten major currencies relative to the US dollar for the sample period of January 1985 to April 2013; and a global portfolio of 22 currencies relative to the US dollar for the sample period of January 1998 to April 2013. Our main measure of FX risk is the market variance defined as the variance of the returns to an equally weighted portfolio of currencies (the “FX market portfolio” hereafter). The market variance is then decomposed in two components: the cross-sectional average variance and the cross-sectional average correlation of all bilateral exchange rates, implementing the methodology used by Pollet and Wilson (2010) to predict equity returns. Next, using quantile regressions, we assess the predictive ability of these risk measures 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. Finally, we design an augmented carry trade strategy that conditions on the risk measures and the return quantiles, which is implemented out of sample and accounts for transaction

² Similar arguments based on crash risk and disaster premia are put forth by Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2013) and Jurek (2014).

costs.

Our main finding is that FX market variance has a significant negative effect on the left tail of future carry trade returns. This implies a negative predictive relation between risk and realized returns in FX. It also indicates that higher market variance is significantly related to large losses to the carry trade, potentially leading investors to unwind their carry trade positions. Furthermore, more than 95% of the time-variation in the FX market variance can be captured by a decomposition into average variance and average correlation. The decomposition allows us to determine that the predictive power of market variance is primarily due to average variance: average variance also has a significant negative effect on the left tail of future carry trade returns, but average correlation does not contribute to the predictability of carry trade returns. Finally, an augmented carry trade strategy that conditions on market variance and the return quantile performs better than the standard carry trade, even when accounting for transaction costs.

Taken together, these results imply the existence of a meaningful predictive relation between market variance and carry trade returns, especially when returns are in the left tail of the distribution. In particular, our empirical analysis shows that information in both market variance and the return quantile is useful for predicting future carry trade returns. In this context, our main finding is that market variance predicts currency returns when it matters most, namely when returns have large negative values, whereas the relation is weaker in normal times. To be more precise, our trading strategy shows that when the carry trade displays a large loss, then market variance provides useful information about whether subsequent losses will occur.

The remainder of the paper is organized as follows. In the next section we discuss the theoretical foundations of the testable hypotheses we examine in this paper. In Section 3 we describe the FX data set and define the measures for risk and return to the carry trade. Section 4 presents the predictive quantile regressions. In Section 5, we report the empirical results, followed by a discussion of the augmented carry trade strategies in Section 6. Finally, Section 7 concludes.

2 Theoretical Motivation and Testable Implications

2.1 Market Variance and the ICAPM

Since the Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973, 1980), a class of asset pricing models has developed which suggests 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:

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

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

where $r_{C,t+1}$ is the return to the carry trade portfolio from time t to $t + 1$; MV_t is the conditional variance of the returns to the FX market portfolio at time t , termed the FX market variance; AV_t is the equally weighted cross-sectional average of the variances of all exchange rate excess returns at time t ; AC_t is the equally weighted cross-sectional average of the pairwise correlations of all exchange rate excess returns at time t ; and ε_{t+1} is a normally distributed error term at time $t+1$. These variables will be formally defined in the next section. It is important to note now, however, that the return to the FX market portfolio is simply an equally weighted average of all exchange rate excess returns. 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, 2012a,b).

Equation (1) is a general characterization of the theoretical prediction that there is a positive linear relation between market variance and future excess returns. The coefficient κ on market variance reflects 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.³

Equation (2) shows that market variance can be decomposed into average variance and average correlation (with $\varphi_1, \varphi_2 > 0$), as shown by Pollet and Wilson (2010) for equity returns. This decomposition is an aspect of our analysis that is critical for determining whether the potential predictive ability of market variance is due to movements in average variance or average correlation. In other words, the decomposition is used to clarify what is the source of the predictive information content of market variance. For example, Goyal and Santa-Clara (2003) show that the equally weighted market variance only reflects systematic risk, whereas average variance captures both systematic and idiosyncratic risk. In light of the above, the first testable hypothesis of the empirical analysis is as follows:

H1: Market variance is a predictor of future FX excess returns due to one or both of its two components: average variance and average correlation.

2.2 Quantile Regressions and the ICAPM

The intertemporal risk-return model of Equations (1)–(2) can be applied to the full conditional distribution of returns. The τ -th conditional quantile function for $r_{C,t+1}$ implied by Equations (1) and (2) is defined as:

$$\begin{aligned} Q_{r_{C,t+1}}(\tau | MV_t) &= \mu + (\kappa + Q_\tau^N) + (\kappa + Q_\tau^N) MV_t \\ &= \alpha(\tau) + \beta(\tau) MV_t, \end{aligned} \quad (3)$$

or, given the decomposition of MV into AV and AC, as:

$$\begin{aligned} Q_{r_{C,t+1}}(\tau | AV_t, AC_t) &= \mu + \varphi_0 (\kappa + Q_\tau^N) + (\kappa + Q_\tau^N) \varphi_1 AV_t + (\kappa + Q_\tau^N) \varphi_2 AC_t \\ &= \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t, \end{aligned} \quad (4)$$

where Q_τ^N is the τ -th quantile of the normal distribution, which has a large negative value deep

³See, among others, French, Schwert and Stambaugh (1987), Chan, Karolyi and Stulz (1992), Glosten, Jagannathan and Runkle (1993), Goyal and Santa-Clara (2003), Ghysels, Santa-Clara and Valkanov (2005), and Bali (2008).

in the left tail and a large positive value deep in the right tail. Given that $\kappa > 0$ and $\varphi_1, \varphi_2 > 0$, we expect the risk measures to have a negative slope in the left tail and a positive slope in the right tail. This provides an economic justification for the use of quantile regressions in our empirical analysis in order to separate the effect of the risk measures on different quantiles of carry trade returns. In what follows, we provide further economic arguments that justify the use of quantile regressions in our context.

2.3 Liquidity Spirals, Currency Crashes and the ICAPM

The negative relation between volatility and large carry trade losses (i.e., the left tail of returns) is consistent with the theoretical model of Brunnermeier and Pedersen (2008). The empirical application of this model to currency crashes is carried out 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. They show that when volatility increases, the carry trade tends to incur losses. 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 investors 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 on the right tail. This economic mechanism based on liquidity spirals leads to the second testable hypothesis of our empirical analysis:

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

3 Measures of Return and Risk for the Carry Trade

This section describes the FX data and defines our measures for risk and return to the carry trade.

3.1 FX Data

We use two separate data sets on US dollar nominal spot and forward exchange rates. The first data set focuses on advanced economies and contains US dollar exchange rates for 10 countries: Australia, Canada, Denmark, Euro area, Japan, New Zealand, Norway, Sweden, Switzerland and United Kingdom. The sample period for all exchange rates of the advanced economies runs from January 1985 to April 2013.⁴

The second data set contains 22 US dollar exchange rates, which include the 10 advanced economies plus the following 12 countries: Czech Republic, Hong Kong, Hungary, India, Kuwait, Mexico, Philippines, Saudi Arabia, Singapore, South Africa, Taiwan and Thailand. As our analysis uses these currencies to form portfolios, we refer to the first data set as the advanced economies portfolio and the second data set as the global portfolio. The sample period for the global portfolio runs from January 1998 to April 2013. This is the largest cross-section of currencies for this sample range that have available bid and ask quotes. Note that the composition of both the advanced economies portfolio and the global portfolio remains constant throughout their respective sample periods.

The data are collected by WM/Reuters and Barclays and are available on Thomson Financial Datastream. We have carefully checked and cleaned the data using the procedure of Darvas (2009). We control for typical data errors by replacing the present day's observation with the previous day's value if any of the following occurs: (i) bid and ask rates are equal; (ii) the spread of the forward exchange rate is less than the spread of the spot exchange rate; and (iii) the spot rate moves but the forward rate stays constant, and vice versa. Once the bid and ask rates have been cleaned, the mid rates used in our analysis are computed simply by taking the mid point of the bid-ask spread.

3.2 The Carry Trade for Individual Currencies

An investor can implement a carry trade strategy for either individual currencies or, more commonly in the practice of currency management, a portfolio of currencies. The carry trade

⁴ Before the introduction of the euro in January 1999, we use the US dollar-Deutsche mark exchange rate adjusted by the official conversion rate between the Deutsche mark and the euro.

strategy for individual currencies can be implemented by buying a forward contract now for exchanging the domestic currency into foreign currency in the future. The investor 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-month horizon is defined as: $r_{j,t+1} = s_{j,t+1} - f_{j,t}$, for $j = \{1, \dots, N\}$, where N is the number of exchange rates at month t , $s_{j,t+1}$ is the log of the nominal spot exchange rate defined as the domestic price of foreign currency j at month $t + 1$, and $f_{j,t}$ is the log of the one-month forward exchange rate j at month t , which is the rate agreed at month t for an exchange of currencies at $t + 1$. An increase in $s_{j,t+1}$ implies a depreciation of the domestic currency, namely the US dollar. Note that $r_{j,t+1}$ is also known as the FX excess return.⁵

If UIP holds, then the excess return will on average be equal to zero, and hence the carry trade will not be profitable. In other words, under UIP, the interest rate differential will on average be exactly offset by a commensurate depreciation of the investment currency. However, as mentioned earlier, 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 tends to be highly profitable (e.g., Darvas, 2009; Della Corte, Sarno and Tsiakas, 2009; Fong, 2010; 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 and, in this paper, we implement one of the most popular. 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 covered interest

⁵ Alternatively, the investor can implement the carry trade by buying a one-month foreign bond and, at the same time, selling a one-month domestic bond. This strategy delivers an excess return that is equal to: $r_{j,t+1} = s_{j,t+1} - s_{j,t} - (i_t - i_{j,t}^*)$, where i_t and $i_{j,t}^*$ are the one-month domestic and foreign nominal interest rates respectively. 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. Note that there is ample empirical evidence that CIP holds in practice for the data frequency examined in this paper (see, e.g., 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).

parity 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 (2012a). 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 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.

3.4 FX Market Variance

Our first measure of risk is the FX market variance, which captures the aggregate variance in FX, and is defined as 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: $r_{M,t+1} = \frac{1}{N} \sum_{j=1}^N r_{j,t+1}$. This can be thought of as the excess return to a naive $1/N$ currency trading strategy, or an international bond diversification strategy that buys N foreign bonds (i.e., all available foreign bonds in the opportunity set of the investor at time t) by borrowing domestically.

Our empirical analysis is based on a monthly measure of FX market variance (MV). A simple and popular way of estimating a realized measure of monthly MV is to use daily excess returns as follows:

$$MV_{t+1} = \sum_{d=1}^{D_t} r_{M,t+d/D_t}^2 + 2 \sum_{d=2}^{D_t} r_{M,t+d/D_t} r_{M,t+(d-1)/D_t}, \quad (5)$$

where D_t is the number of trading days in month t , and typically $D_t = 21$. Following French, Schwert and Stambaugh (1987), Goyal and Santa-Clara (2003), and Bali, Cakici, Yan and Zhang (2005), among others, this measure of market variance accounts for the autocorrelation in daily returns.⁶

⁶ An alternative way to construct a proxy for market variance, which we do not consider in this paper, is to use implied volatilities based on currency options data. See, e.g., Della Corte, Sarno and Tsiakas (2011).

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:

$$MV_{t+1} = AV_{t+1} \times AC_{t+1}. \quad (6)$$

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 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 estimate monthly measures of AV and AC as follows:

$$AV_{t+1} = \frac{1}{N} \sum_{j=1}^N V_{j,t+1}, \quad (7)$$

$$AC_{t+1} = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j \neq i}^N C_{ij,t+1}, \quad (8)$$

where $V_{j,t+1}$ is the realized monthly variance of the excess return to exchange rate j at month $t+1$, which is computed using daily excess returns as follows:

$$V_{j,t+1} = \sum_{d=1}^{D_t} r_{j,t+d/D_t}^2 + 2 \sum_{d=2}^{D_t} r_{j,t+d/D_t} r_{j,t+(d-1)/D_t}, \quad (9)$$

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

$$C_{ij,t+1} = \frac{V_{ij,t+1}}{\sqrt{V_{i,t+1}} \sqrt{V_{j,t+1}}}, \quad (10)$$

$$V_{ij,t+1} = \sum_{d=1}^{D_t} r_{i,t+d/D_t} r_{j,t+d/D_t} + 2 \sum_{d=2}^{D_t} r_{i,t+d/D_t} r_{j,t+(d-1)/D_t}. \quad (11)$$

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:

$$r_{C,t+1} = \alpha + \beta MV_t + \varepsilon_{t+1}, \quad (12)$$

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 returns following times of high market variance.

The second predictive regression assesses the risk-return tradeoff implied by the decomposition of market variance into average variance and average correlation:

$$r_{C,t+1} = \alpha + \beta_1 AV_t + \beta_2 AC_t + \varepsilon_{t+1}, \quad (13)$$

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 which component of MV can predict 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 effect of the risk measures on the full distribution of future carry returns. It is possible, for example, that market variance is a poor predictor of the conditional mean return but predicts well one or both tails of the return distribution. In our analysis, we focus on deciles of this distribution.

The first predictive quantile regression estimates:

$$Q_{r_{C,t+1}}(\tau | MV_t) = \alpha(\tau) + \beta(\tau) MV_t, \quad (14)$$

where $Q_{r_{C,t+1}}(\tau | \cdot)$ is the τ -th quantile function of the one-month ahead carry trade returns conditional on information available at month t . The second predictive quantile regression estimates:

$$Q_{r_{C,t+1}}(\tau | AV_t, AC_t) = \alpha(\tau) + \beta_1(\tau) AV_t + \beta_2(\tau) AC_t. \quad (15)$$

5 Empirical Results

5.1 Descriptive Statistics

Table 1 reports descriptive statistics on the FX risk and return measures for both the advanced economies and the global portfolio. Assuming no transaction costs, the carry trade delivers an annualized mean return of 6.2% (advanced economies) and 9.9% (global portfolio), the standard deviations are 11.1% and 9.0%, and the Sharpe ratios are 0.555 and 1.100. The carry trade return is primarily due to the interest rate differential across countries, which delivers an average return of 7.4% for the advanced economies and 11.9% for the global portfolio. The exchange rate depreciation component has a return of -1.2% and -2.0% respectively, indicating that on average exchange rates only partially offset the interest rate differential. Furthermore, the carry trade return displays negative skewness and higher-than-normal kurtosis. 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).

Turning to the risk measures, the mean of MV is about half the value of the AV, whereas the mean of AC is 0.461 and 0.317 for the two portfolios. MV and AV exhibit high positive skewness and massive kurtosis. The time variation of MV, AV and AC together with the cumulative carry trade return are displayed in Figure 1.

Panel B of Table 1 shows the cross-correlations between the variables. Note that these are contemporaneous cross-correlations and hence have no implications for prediction. We highlight that MV and AV are very highly correlated with each other but are negatively

correlated with the carry and market returns. It is also worth noting that the correlation between AV and AC is 21.3% and 26.8% for the two portfolios. This implies that FX variances and correlations are on average only moderately positively correlated.

Moreover, in the bottom row of Figure 1, we show the correlation between AV and AC over time using a five-year rolling window. The figure shows that this correlation fluctuates considerably over time, is often low and can even be negative. Interestingly, during the recent financial crisis only AV increases to unprecedented high levels, whereas AC is around normal values. This further justifies the decomposition of MV into AV and AC as the low correlation of AV and AC indicates that they may capture different types of risk.

5.2 The Decomposition of Market Variance

The three FX risk measures of MV, AV and AC are related by the approximate decomposition of Equation (6). We evaluate the empirical validity of the decomposition by presenting regression results in the Internet Appendix. In summary, we find that AV and AC together capture up to 96% of the time variation in MV, and the explanatory power of AV is substantially higher than AC.

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 2. This also tests the first hypothesis (*H1*) we set out in Section 2. In regressing the one-month ahead carry trade return on MV, the table shows that overall the relation is negative but not statistically significant. This is true for both advanced economies and the global portfolio. Therefore, just by using OLS regressions we cannot find evidence that high market variance is a significant predictor of future carry trade returns. This remains the case when we regress the one-month ahead carry trade return on AV and AC. Although AV and AC are also negatively related to the carry trade return, this result is not statistically significant. Therefore, in the context of OLS predictive regressions, the evidence does not support hypothesis *H1* as we find that the FX risk measures are not

significant predictors of future FX excess returns.

The OLS regressions explore the tradeoff between risk and mean 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, and hence test our second hypothesis ($H2$), by estimating predictive quantile regressions. We begin with Figure 2 which plots the parameter estimates $\beta(\tau)$ of the predictive quantile regressions of the one-month ahead carry trade return on MV. These results are shown in more detail in Table 3. The main finding here is that for both portfolios MV has a significant negative relation to the future carry trade return only for the quantiles in the left tail of the distribution. The lower the quantile, the more negative the value of the coefficient and the more statistically significant it is, all the way up to the median. In the right tail, the relation is primarily positive but insignificant. Also note that the constant is highly significant, being negative below the 0.3 quantile and positive above it.⁷

We refine these results by moving to the second quantile regression that conditions on AV and AC. We report these results in the Internet Appendix. Our main finding is that AV has similar predictive power to MV: AV has a significant negative relation to the carry trade return only in the left-tail quantiles. However, the relation of AC and carry trade returns is not statistically significant for any part of the distribution. In conclusion, therefore, the similar predictive power of MV and AV indicates that AV is the component of MV that allows us to predict the left tail of carry trade returns.

The results above 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. High MV has a significant negative impact only in the left tail. Second, the decomposition of MV into AV and AC is helpful in determining that the predictive ability of MV is primarily due to AV. These results establish a significant relation between FX volatility and large carry trade losses, which is consistent with investors unwinding their carry trade positions in times of high volatility. Third, AC is not significantly related to future carry returns for any part of the distribution. This further confirms that the predictive power of MV is due to AV and

⁷ As a robustness check, we show that the negative relation between MV and the left tail of future carry trade returns remains significant even when we change the numeraire currency from the US dollar to one of the following: the euro, the UK pound or the Japanese yen. See the Internet Appendix for the results.

not to AC.

In short, these results provide partial support for the second hypothesis ($H2$) since the predictive ability of MV and AV is not the same across all quantiles, and in fact is much stronger in the left tail of the distribution. However, in contrast to $H2$, we find that only MV and AV have a significant negative effect in the left tail, whereas AC does not have a significant relation with any part of the distribution of carry trade returns.

6 Augmented Carry Trade Strategies

We further evaluate the predictive ability of the risk measures on future carry trade returns by designing augmented carry trade strategies that condition on market variance. These strategies are then compared to the benchmark strategy, which is the standard carry trade.⁸

6.1 Trading Strategies

We first design an augmented carry trade strategy that only conditions on current MV. This strategy does not discriminate between different quantiles of the return distribution and hence is analogous to the OLS regressions reported earlier. We refer to this strategy as the 1.0* strategy as it involves the full distribution of returns and hence is distinct from strategies that condition on quantiles such as 0.1. In particular, the 1.0* strategy implements the following rule at each month t : if MV from $t - 1$ to t is “high” in the sense that it is higher than its median value up to that point, we close the carry trade positions and thus receive an excess return of zero at $t + 1$; otherwise we execute the standard carry trade.⁹

The second strategy is designed as described above but adds one further condition that relates to the quantiles of returns. In particular, the strategy now implements the following rule at each month t : only for the carry trade returns that are lower than the τ -quantile of the empirical distribution at month t , if MV from $t - 1$ to t is “high” in the sense that it is

⁸ In general, for the relation between statistical and economic measures of predictive ability see Cenesizoglu and Timmermann (2012).

⁹ We have also implemented this strategy by defining “high” MV as being in the top quartile or quintile of the MV distribution up to that point. The results are qualitatively similar to using the median as the threshold and are available upon request.

higher than its median value up to that point, we close the carry trade positions and thus receive an excess return of zero at $t + 1$; otherwise we execute the standard carry trade. We refer to the second strategy as the 0.1 strategy when we focus on the 0.1 quantile, or the 0.2 strategy for that quantile and so on.

The second augmented strategy is similar to the first one in that it exploits the negative relation between current MV and the one-month ahead carry return. However, whereas the first strategy only conditions on high MV, the second strategy also conditions on carry returns being in the left tail. Put differently, whereas the first strategy is consistent with the OLS regressions, the second one is guided by the quantile regressions by focusing on the left tail of returns.¹⁰

The third and final strategy ignores MV and instead focuses just on the quantile of returns. With this third strategy, we can check whether the possible improvement of the second strategy over the standard carry trade is due to negative returns being followed by negative returns or whether the additional MV condition of the second strategy matters and further improves performance. This strategy now implements the following rule at each month t : if the carry trade returns are lower than the τ -quantile of the empirical distribution at month t , we close the carry trade positions and thus receive an excess return of zero at $t + 1$; otherwise we execute the standard carry trade. As this strategy is primarily implemented for robustness, we focus on the 0.1 quantile and hence refer to it as the 0.1* strategy in order to distinguish it from the 0.1 strategy described above that also conditions on MV.

It is important to note that all strategies are implemented out of sample. Specifically, the strategies move forward recursively starting 3 years after the beginning of the sample so as to have enough observations to determine the quantiles of carry trade returns and the median of MV up to that point. The strategies do not directly use the parameter estimates from the regressions but simply try to exploit the negative relation between future carry returns and

¹⁰ Note that strictly speaking the quantile regressions capture the relation between the risk measures and particular quantiles of the one-step ahead carry trade returns. Since ex ante we do not know the future distribution of returns, the second trading strategy conditions on whether the current return is below a given quantile of the past distribution of returns. For example, it may be that the current quantile of returns is informative about future return quantiles. Moreover, given the nature of our testing hypotheses and the implementation of quantile regressions, it is natural to check whether being in the left tail of returns when volatility is high improves the performance of the trading strategy.

current MV for either the full distribution of returns or a particular quantile.

6.2 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 either new contracts are entered or the existing contracts are rolled over. The distinction between new and rolled-over positions is key to the analysis of carry trades because interest rate differentials tend to be persistent, leading to many rolled-over positions. In practice, the transaction cost for rolled-over positions is much lower than for new transactions. Following Darvas (2009), we account for the effect of transaction costs on the payoff to the carry trade, P_{t+1}^{net} , when trading each currency, as follows:

1. Buy foreign currency:

$$(a) \text{ New position: } P_{t+1}^{net} = A_t^{new} \cdot (S_{t+1}^b - F_t^a),$$

$$(b) \text{ Rolled-over position: } P_{t+1}^{net} = A_t^{rol} \cdot (S_{t+1}^b - F_t^a + (S_t^a - S_t^b)),$$

2. Sell foreign currency:

$$(a) \text{ New position: } P_{t+1}^{net} = A_t^{new} \cdot (-S_{t+1}^a + F_t^b),$$

$$(b) \text{ Rolled-over position: } P_{t+1}^{net} = A_t^{rol} \cdot (-S_{t+1}^a + F_t^b + (S_t^a - S_t^b)),$$

where A_t^{new} and A_t^{rol} are the notional amounts of the new and rolled-over positions respectively, the superscripts a and b denote ask and bid quotes, and the upper case S_t and F_t are the levels of spot and forward exchange rates.

To better understand the intuition of these calculations, we first define the spread of the swap points as the difference between the spread of the forward exchange rate minus the spread of the spot exchange rate. Then, the cost of a new transaction is equal to the full

spread of the spot rate plus half of the spread of the swap points. In contrast, the cost of a rolled-over transaction is just equal to half of the spread of the swap points. Therefore, rolling over a position has considerably lower transaction costs than opening a new position. For more details, including a particular example, see Darvas (2009).

6.3 The Performance of the Strategies

The performance of the strategies is reported in Table 4, where Panel A is for returns without transaction costs and Panel B for returns net of transaction costs. Our discussion below focuses on the results of Panel B that accounts for transaction costs. We begin with the 1.0* strategy, i.e., the MV strategy for the full distribution of returns, applied on the advanced economies portfolio. This augmented strategy has a lower return than the carry trade (3.76% vs. 5.34%), a lower volatility (6.38% vs. 10.28%) and a higher Sharpe ratio (0.59 vs. 0.52) that is marginally insignificant (p -value = 0.13). Therefore, the 1.0* strategy delivers an improved risk-return ratio over the carry trade but is insignificant.¹¹

Next, we continue looking at the augmented MV strategy for the advanced economies portfolio but we now add the second condition on the quantile of carry trade returns. Consider, for example, the 0.5 quantile. In this case, we shut down the carry trade every time that: (i) the current return is in the bottom half of the return distribution, and (ii) MV is high. Relative to the carry trade, this strategy delivers a slightly lower return (5.06% vs. 5.34%), a lower volatility (8.26% vs. 10.28%) and a higher Sharpe ratio (0.61 vs. 0.52) that is now significant (p -value = 0.07).

Now consider the 0.1 quantile. In this case, we shut down the carry trade every time that: (i) the current return is below the bottom decile of the return distribution, and (ii) MV is high. Relative to the carry trade, this strategy delivers a higher return (5.49% vs. 5.34%),

¹¹ We assess the statistical significance of the Sharpe ratios by testing whether the Sharpe ratio of an augmented strategy is significantly different to the Sharpe ratio of the standard carry trade. This is a one-sided test. The test assumes that returns are independent and identically distributed (iid), which is an assumption that was empirically validated using the BDS test of Brock, Dechert, Sheinkman and LeBaron (1996). The results of the BDS test are in the Internet Appendix. For iid returns, the Sharpe ratio estimator is asymptotically normally distributed with a standard error equal to: $SE(\widehat{SR}) = \sqrt{\frac{1}{T} \left(1 + \frac{1}{2}\widehat{SR}^2\right)}$, where \widehat{SR} is the estimated Sharpe ratio (see, e.g., Lo, 2002).

a lower volatility (9.15% vs. 10.28%) and a higher Sharpe ratio (0.60 vs. 0.52), although its statistical significance is borderline (p -value = 0.10).

Finally, consider the 0.1* strategy. In this case, we shut down the carry trade every time the current return is below the bottom decile of the return distribution irrespective of the value of MV. Relative to the carry trade, this strategy delivers a lower return (5.12% vs. 5.34%), a lower volatility (9.09% vs. 10.28%) and a higher Sharpe ratio (0.56 vs. 0.52) that is now insignificant (p -value = 0.25).¹²

When we move from the advanced economies to the global portfolio, the results are slightly different. The 1.0* strategy delivers a much lower return than the carry trade (2.69% vs. 8.10%), a much lower volatility (4.58% vs. 9.12%) and this time a lower Sharpe ratio (0.59 vs. 0.89). Hence for the full distribution of global portfolio returns, the 1.0* strategy yields a worse risk-return ratio than the carry trade. Nonetheless, this finding changes drastically when we also condition on the return quantile. For example, the 0.1 quantile MV strategy provides a higher return (8.58% vs. 8.10%), a lower volatility (7.53% vs. 9.12%) and a much higher Sharpe ratio (1.14 vs. 0.89) with a p -value of 0.01. Finally, the 0.1* strategy provides a lower return (7.65% vs. 8.10%), a lower volatility (7.66% vs. 9.12%) and a higher Sharpe ratio (1.00 vs. 0.89) with a p -value of 0.10.

To make our case more concrete consider Figure 3, which plots the cumulative returns of the 0.1 quantile MV strategy relative to the standard carry trade. As we can see in the figure, there is virtually no difference in the two strategies in the early part of the sample, when the standard carry trade tends to deliver an increasing cumulative return and the augmented strategy mirrors the standard strategy. However, the difference in the two strategies becomes more pronounced during the recent financial crisis when the augmented strategy performs better. Keep in mind that, in effect, this augmented strategy shuts down the carry trade for about 5% of the time: when the current return is below the bottom 0.1 quantile and MV is higher than its median. Although 95% of the time this augmented strategy is identical to the standard carry trade, it still manages to deliver a higher return, a much lower volatility

¹² Note that the 0.1 strategy that conditions on MV and the 0.1* strategy that does not, are similar in the following sense: most of the low returns in the bottom decile tend to be associated with high MV. This is consistent with the theory of Brunnermeier, Nagel and Pedersen (2009) on liquidity spirals and implies that the improvement of the 0.1 strategy over the 0.1* strategy is small but positive.

and significantly higher Sharpe ratio. In short, as this strategy focuses on the lowest return quantile, it establishes that not only is high FX volatility related to large carry trade losses, but also accounting for this relation can improve the performance of the carry trade. Overall, the results provide further evidence in support of hypothesis $H2$ to the extent it concerns MV.

These results suggest that, among the augmented strategies, the best way to improve the Sharpe ratio is to combine the condition of high MV with the condition of being deep in the left tail of current returns. This delivers a higher Sharpe ratio than the carry trade even though such an augmented strategy shuts down the carry trade only for a small percentage of the time. Furthermore, note that much of the improvement in the augmented strategies is due to a substantially lower volatility. Indeed, the volatility of the augmented strategies decreases monotonically as we move to a higher quantile since the higher the quantile, the more often we shut down the carry trade. As a result, the augmented strategies consistently deliver a higher Sharpe ratio than the standard carry trade. These Sharpe ratios are higher and more significant at the lower quantiles. To summarize, these results highlight the tangible out-of-sample economic gains of the augmented carry trade strategies.

7 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 in low-interest currencies tends to deliver large average returns. Not surprisingly, therefore, this investment strategy is at the heart of the currency management decision-making of key players in financial markets.

This paper fills a gap in the literature by demonstrating empirically the existence of a negative predictive relation between risk and the realized return to the carry trade. We measure FX risk by the market variance, which is 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 and out-of-sample trading simulations, which provide a natural way of assessing the effect of risk on different quantiles of the distribution of carry trade returns.

Our main finding is that market variance has a significant negative effect on the left tail of the distribution of future carry trade returns. This is due to the average variance component of market variance since average variance is also significantly related to large future carry trade losses. However, average correlation is the component of market variance that is not significantly related to any part of the carry return distribution. We take advantage of our findings by forming a new version of the carry trade that conditions on market variance and the return quantile, and show that this strategy performs better than the standard carry trade. To be more precise, our trading strategy shows that when the carry trade displays a large loss, then market variance provides useful information about whether subsequent losses will occur.

These results imply that to some extent exchange rates are predictable, especially when it matters the most: when the carry trade produces large losses. In other words, if the carry trade is about “going up the stairs and down the elevator,” then market variance can tell us something valuable about when the elevator is likely to go down. Market variance and the return quantile both contain useful information for predicting future carry trade returns. Combining these two conditions leads to augmented trading strategies that outperform the standard carry trade. In the end, by focusing on the left tail of the return distribution of carry trades, we uncover a significantly negative predictive relation between risk and realized returns in foreign exchange.

Table 1. Descriptive Statistics

The table reports descriptive statistics for the monthly excess returns of two currency portfolios: advanced economies and global portfolio. The advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013. 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 return to the FX market portfolio is an equally weighted average of all exchange rate excess returns. 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. The skewness, kurtosis and AR(1) are for monthly observations. AR(1) is the first order autocorrelation.

Panel A: Summary Statistics						
	<i>Advanced Economies</i>					
	<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	0.062	0.111	0.555	-0.999	5.498	0.037
Market	0.020	0.083	0.242	-0.313	3.875	0.056
<i>Carry Trade Components</i>						
Exchange Rate	-0.012	0.112		-1.076	5.662	0.061
Interest Rate	0.074	0.013		2.622	15.099	0.637
<i>Variances and Correlations</i>						
Market Variance	0.006	0.002		3.064	16.175	0.568
Average Variance	0.012	0.003		4.734	38.779	0.579
Average Correlation	0.461	0.151		0.019	2.473	0.569
<i>Global Portfolio</i>						
<i>Portfolio Returns</i>						
Carry Trade	0.099	0.090	1.100	-0.560	4.409	0.160
Market	0.009	0.066	0.134	-0.623	4.925	0.060
<i>Carry Trade Components</i>						
Exchange Rate	-0.020	0.086		-0.645	4.600	0.114
Interest Rate	0.119	0.024		1.784	5.113	0.777
<i>Variances and Correlations</i>						
Market Variance	0.004	0.001		3.639	20.776	0.700
Average Variance	0.010	0.003		5.201	41.390	0.599
Average Correlation	0.317	0.115		0.054	2.254	0.596

Panel B: Cross-Correlations

	Carry Return	Market Return	Market Variance	Average Variance	Average Correlation
<i>Advanced Economies</i>					
<i>Portfolio Returns</i>					
Carry Trade Return	1.000				
Market Return	0.154	1.000			
<i>Variances and Correlations</i>					
Market Variance	-0.219	-0.142	1.000		
Average Variance	-0.341	-0.197	0.934	1.000	
Average Correlation	0.083	0.062	0.451	0.213	1.000
<i>Global Portfolio</i>					
<i>Portfolio Returns</i>					
Carry Trade Return	1.000				
Market Return	0.457	1.000			
<i>Variances and Correlations</i>					
Market Variance	-0.403	-0.302	1.000		
Average Variance	-0.412	-0.323	0.945	1.000	
Average Correlation	-0.114	0.016	0.486	0.268	1.000

Table 2. OLS Predictive Regressions

The table presents the ordinary least squares results for two predictive regressions implemented on two currency portfolios: advanced economies and global portfolio. 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 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 average variance and AC_t is the 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 all exchange rate excess returns. 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 advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013.

Regressions for the Carry Trade Return				
	<i>Advanced Economies</i>		<i>Global Portfolio</i>	
	(1)	(2)	(1)	(2)
Constant	0.008 (2.650) [0.014]	0.006 (0.973) [0.350]	0.011 (3.790) [0.002]	0.016 (2.667) [0.018]
Market Variance	-5.273 (-0.861) [0.449]		-8.445 (-0.884) [0.485]	
Average Variance		-3.958 (-1.512) [0.220]		-1.282 (-0.373) [0.797]
Average Correlation		-0.008 (0.759) [0.462]		-0.021 (-1.333) [0.221]
\overline{R}^2 (%)	0.28	0.45	0.52	0.04

Table 3. 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 market variance. The results are for two currency portfolios: advanced economies and global portfolio. 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 all exchange rate excess returns. 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 advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013.

Predictive Quantile Regressions for the Carry Trade Return											
Quantile	<i>Advanced Economies</i>										
	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Constant	-0.036 (-3.464) [0.022]	-0.024 (-3.061) [0.013]	-0.009 (-2.670) [0.008]	-0.001 (-0.287) [0.731]	0.006 (2.111) [0.053]	0.012 (4.733) [0.000]	0.014 (6.004) [0.000]	0.022 (6.864) [0.000]	0.027 (10.606) [0.000]	0.038 (11.119) [0.000]	0.042 (8.797) [0.000]
Market Variance	-48.465 (-3.631) [0.005]	-17.672 (-1.726) [0.088]	-8.609 (-1.658) [0.094]	-9.277 (-2.310) [0.025]	-10.000 (-2.697) [0.011]	-6.571 (-1.821) [0.074]	-1.864 (-0.514) [0.569]	-0.197 (-0.046) [0.931]	5.293 (1.312) [0.186]	4.787 (0.934) [0.349]	17.325 (2.669) [0.027]
\bar{R}^2 (%)	3.6	2.6	1.2	0.9	0.2	-0.2	-0.2	-0.3	0.1	0.5	2.5
Quantile	<i>Global Portfolio</i>										
	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 (-2.722) [0.040]	-0.015 (-2.862) [0.033]	-0.007 (-2.073) [0.064]	0.001 (0.158) [0.843]	0.005 (1.478) [0.157]	0.012 (4.089) [0.000]	0.016 (5.137) [0.000]	0.023 (7.268) [0.000]	0.023 (11.611) [0.000]	0.032 (7.009) [0.000]	0.040 (8.785) [0.000]
Market Variance	-70.563 (-5.084) [0.002]	-29.850 (-2.521) [0.038]	-18.875 (-2.234) [0.043]	-17.947 (-2.263) [0.041]	-3.724 (-0.482) [0.669]	-6.415 (-0.875) [0.377]	3.171 (0.438) [0.654]	2.657 (0.386) [0.656]	21.160 (3.589) [0.003]	15.064 (1.619) [0.128]	16.812 (1.430) [0.166]
\bar{R}^2 (%)	9.2	3.3	2.1	0.5	-0.2	0.0	-0.5	-0.5	1.2	2.3	2.0

Table 4. Augmented Carry Trade Strategies

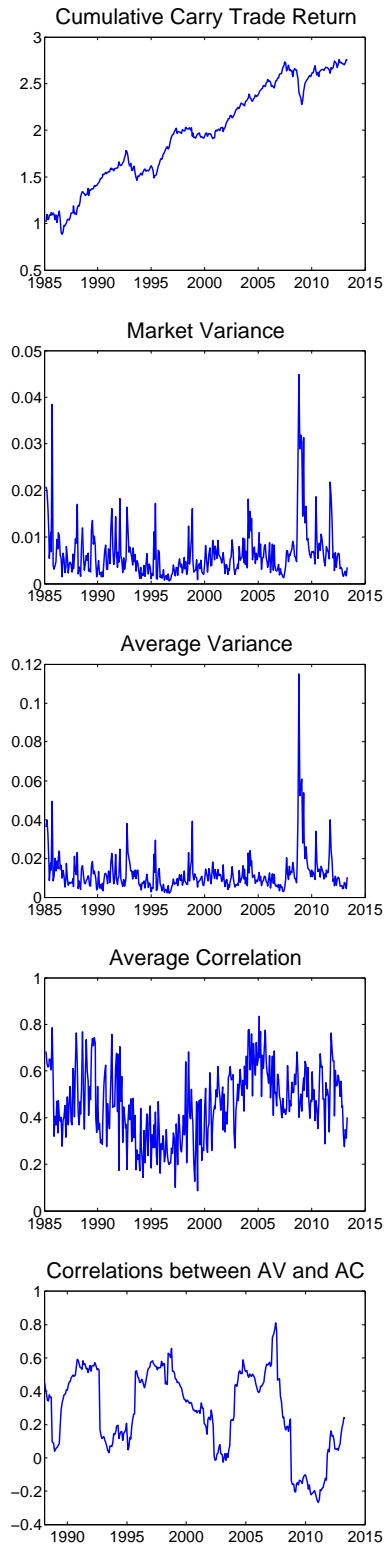
The table presents the out-of-sample performance of augmented carry trade strategies that condition on the movement of market variance. The results are for two currency portfolios: advanced economies and global portfolio. Panel A assumes no transaction costs, whereas Panel B implements the carry trade strategies with transaction costs. Market variance is the variance of the return to the market portfolio constructed as the equally weighted average all exchange rate excess returns. The quantile strategies implement the following rule at each time t : for the carry trade returns that are lower than the τ -quantile of the empirical distribution, if market variance from $t-1$ to t is higher than its median value, we close the carry trade positions and thus receive a return of zero at $t+1$; otherwise we execute the standard carry trade. The 0.1* strategy implements the following rule at each time t : if the carry trade return is lower than the 0.1-quantile of the empirical distribution, we close the carry trade positions and thus receive a return of zero at $t+1$; otherwise we execute the standard carry trade. The 1.0* strategy implements the following rule at each time t : if market variance from $t-1$ to t is higher than its median value, we close the carry trade positions and thus receive a return of zero at $t+1$; otherwise we execute the standard carry trade. The mean returns, volatilities and Sharpe ratios are reported in annualized terms. Volatility is the standard deviation of returns. Below each Sharpe ratio we report the p -value for a test of whether the Sharpe ratio of the augmented strategy is significantly different from the Sharpe ratio of the standard carry trade. The advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013. All strategies move forward recursively starting 3 years after the beginning of the sample.

Panel A: No Transaction Costs											
Carry Trade	Quantile										
	0.1*	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0*
<i>Advanced Economies</i>											
<i>Mean Returns (%)</i>											
6.27	6.30	6.49	6.45	5.89	5.98	6.04	5.94	5.13	4.91	5.23	4.53
<i>Volatilities (%)</i>											
10.25	8.99	9.08	8.80	8.69	8.43	8.22	8.16	7.97	7.41	6.75	6.39
<i>Sharpe Ratios</i>											
0.61	0.70 (0.08)	0.71 (0.06)	0.73 (0.03)	0.68 (0.15)	0.71 (0.06)	0.74 (0.03)	0.73 (0.04)	0.64 (0.31)	0.66 (0.21)	0.77 (0.01)	0.71 (0.06)
<i>Global Portfolio</i>											
<i>Mean Returns (%)</i>											
9.48	8.89	10.08	9.05	9.26	7.48	7.29	7.66	7.76	7.19	5.94	3.75
<i>Volatilities (%)</i>											
9.15	7.59	7.46	6.85	6.73	6.49	6.23	6.10	5.87	5.80	5.57	4.72
<i>Sharpe Ratios</i>											
1.04	1.17 (0.10)	1.35 (0.00)	1.32 (0.01)	1.38 (0.00)	1.15 (0.14)	1.17 (0.10)	1.26 (0.02)	1.32 (0.01)	1.24 (0.03)	1.06 (0.39)	0.80 (1.00)

Panel B: With Transaction Costs

Carry Trade	Quantile										
	0.1*	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0*
<i>Advanced Economies</i>											
<i>Mean Returns (%)</i>											
5.34	5.12	5.49	5.41	4.94	4.66	5.06	4.69	4.53	4.01	4.23	3.76
<i>Volatilities (%)</i>											
10.28	9.09	9.15	8.87	8.78	8.55	8.26	8.13	7.77	7.40	6.80	6.38
<i>Sharpe Ratios</i>											
0.52	0.56 (0.25)	0.60 (0.10)	0.61 (0.07)	0.56 (0.24)	0.54 (0.34)	0.61 (0.07)	0.58 (0.18)	0.58 (0.16)	0.54 (0.36)	0.62 (0.05)	0.59 (0.13)
<i>Global Portfolio</i>											
<i>Mean Returns (%)</i>											
8.10	7.65	8.58	7.65	7.46	5.50	5.47	5.65	6.17	6.00	4.67	2.69
<i>Volatilities (%)</i>											
9.12	7.66	7.53	6.97	6.77	6.51	6.21	6.11	5.92	5.69	5.44	4.58
<i>Sharpe Ratios</i>											
0.89	1.00 (0.10)	1.14 (0.01)	1.10 (0.02)	1.10 (0.02)	0.84 (0.68)	0.88 (0.53)	0.93 (0.36)	1.04 (0.07)	1.05 (0.05)	0.86 (0.62)	0.59 (1.00)

Advanced Economies



Global Portfolio




Figure 1. Carry Trade Return and FX Risk Measures

This figure displays the time series of the cumulative carry trade return, market variance, average variance, average correlation and the rolling correlation between average variance and average correlation using a 5-year rolling window. The left column is for a currency portfolio of advanced economies and the right column for a global portfolio. The advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013.

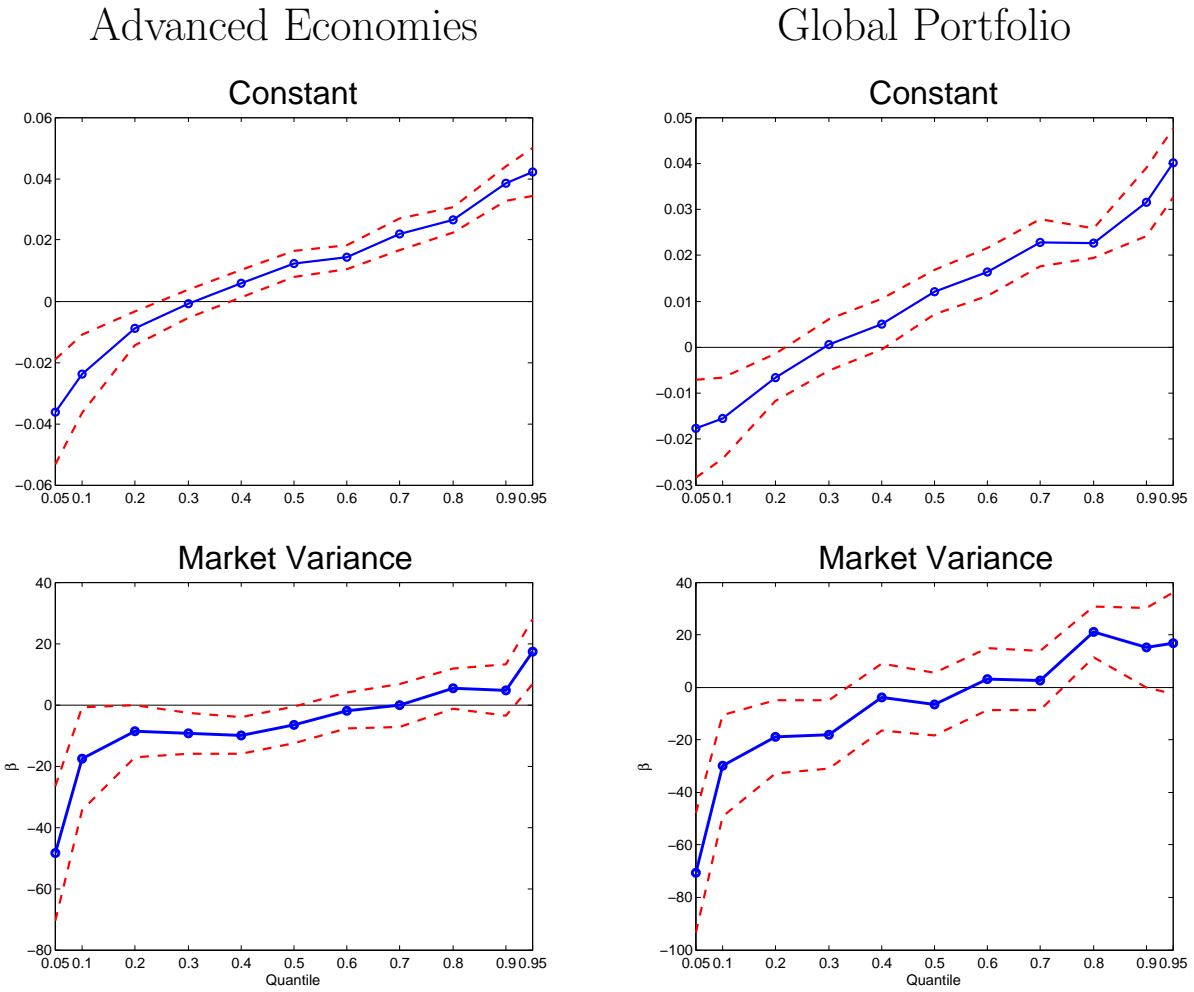


Figure 2. Market Variance

This figure shows the parameter estimates of the predictive quantile regression of the one-month ahead carry trade return on market variance. The left column is for a currency portfolio of advanced economies and the right column for a global portfolio. The advanced economies include 10 exchange rates relative to the US dollar for the sample period of January 1985 to April 2013. The global portfolio includes 22 exchange rates relative to the US dollar for the sample period of January 1998 to April 2013. The dashed lines indicate the 90% confidence interval based on bootstrap standard errors.




Figure 3. Cumulative Return of Augmented Carry Trade Strategies

This figure shows the cumulative return of augmented carry trade strategies in solid blue line versus the cumulative return of the standard carry trade in dashed red line. The augmented carry trade strategies are the market variance strategies for the 0.1 quantile, which implement the following rule at each month t : for the carry trade returns that are lower than the 0.1 quantile of the distribution, if market variance for the period from $t-1$ to t is higher than its median value, we close the carry trade positions and thus receive a return of zero at $t+1$; otherwise we execute the standard carry trade. The left column is for cumulative returns without transaction costs and the right column for cumulative returns net of transaction costs. All strategies are implemented out of sample.

References

- Akram, Q.F., D. Rime, and L. Sarno (2008). “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics* **76**, 237–253.
- 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.
- Bilson, J.F.O. (1981). “The ‘Speculative Efficiency’ Hypothesis,” *Journal of Business* **54**, 435–451.
- Brock, W., D. Dechert, J. Sheinkman, and B. LeBaron (1996). “A Test for Independence Based on the Correlation Dimension,” *Econometric Reviews* **15**, 197–235.
- 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*. Hoboken, NJ: Wiley.
- 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.
- Cenesizoglu, T., and A. Timmermann (2012). “Do Return Prediction Models Add Economic Value?” *Journal of Banking and Finance* **36**, 2974–2987.
- 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.
- Christiansen, C., A. Rinaldo, and P. Söderlind (2011). “The Time-Varying Systematic Risk of Carry Trade Strategies,” *Journal of Financial and Quantitative Analysis* **46**, 1107–1125.
- Darvas, Z. (2009). “Leveraged Carry Trade Portfolios,” *Journal of Banking and Finance* **33**, 944–957.
- Della Corte, P., L. Sarno, and I. Tsiakas (2009). “An Economic Evaluation of Empirical Exchange Rate Models,” *Review of Financial Studies* **22**, 3491–3530.
- Della Corte, P., L. Sarno, and I. Tsiakas (2011). “Spot and Forward Volatility in Foreign Exchange,” *Journal of Financial Economics* **100**, 496–513.
- 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. Fraiburger, X. Gabaix, R. Ranciere, and A. Verdelhan (2013). “Crash Risk in Currency Markets,” Unpublished Working Paper, Harvard University.
- Fong, W.M. (2010). “A Stochastic Dominance Analysis of Yen Carry Trades,” *Journal of Banking and Finance* **34**, 1237–1246.
- 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.
- Jurek, J., (2014). “Crash-Neutral Currency Carry Trades,” *Journal of Financial Economics* (forthcoming).
- 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.
- Lo, A.W. (2002). “The Statistics of Sharpe Ratios,” *Financial Analysts Journal* **58**, 36–52.
- Lustig, H., N. Roussanov, and A. Verdelhan (2011). “Common Risk Factors in Currency Markets,” *Review of Financial Studies* **24**, 3731–3777.
- 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.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012a). “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance* **67**, 681–718.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012b). “Currency Momentum Strategies,” *Journal of Financial Economics* **106**, 660–684.
- 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.
- Mueller, P., A. Stathopoulos, and A. Vedolin (2012). “International Correlation Risk” Unpublished Working Paper, London School of Economics.
- Newey, W.K., and K.D. West (1987). “A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica* **55**, 703–708.

Pollet, J.M., and M. Wilson (2010). "Average Correlation and Stock Market Returns," *Journal of Financial Economics* **96**, 364–380.

Sarno, L. (2005). "Towards a Solution to the Puzzles in Exchange Rate Economics: Where Do We Stand?" *Canadian Journal of Economics* **38**, 673–708.