Correlated Risk Factors in Currency Markets

Tommy von Brömsen*

Sveriges Riksbank

January 15, 2020

Abstract

An asset that loads on two risk factors earns benefits of diversification across the factors. A change in the correlation between the factors has effects on the diversification, which directly affects the volatility of returns and potentially – via volatility timing – future Sharpe ratios. I find strong evidence of these economic links in currency markets. First, I document that the correlation (CORR) between the *dollar* factor and the *carry* factor is highly time-varying, across almost the entire $[-1, 1]$ interval. Second, for high (low) interest rate currency returns, a positive CORR (above 0.25) is associated with low (high) volatility and a high (low) future Sharpe ratio. The reverse holds for negative CORR (below -0.25). These results extend to the standard, high-minus-low carry trade (HML). Comparing negative to positive CORR, the average next-month Sharpe ratio of the HML carry trade is 0.20 and 0.99, respectively. Third, I show that CORR is procyclical, positively related to the US interest rate and global interest rate differentials. CORR is also linked to the time-varying characteristic of the US dollar as a safe haven. Fourth, CORR contains information about the predictability of carry trade crashes.

Keywords: Carry trade; Currency risk factors; FX; Interest rates; Predictability; Time-varying correlation

JEL Classification: E43, F31, G15.

*I thank Charlotte Christiansen, Magnus Dahlquist, Petri Jylhä (discussant), Valeri Sokolovski and Irina Zviadadze as well as the seminar participants at the Stockholm School of Economics, the Nordic Finance Network in Lund 2018, and Sveriges Riksbank for helpful comments and suggestions. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of Sveriges Riksbank. Contact e-mail: *tommy.vonbromsen@riksbank.se*.

1 Introduction

In a factor model, consider an asset with positive loadings on two, uncorrelated risk factors. Since the factors are uncorrelated, the asset returns are benefiting from a diversification across the factors. If the two risk factors begin to correlate positively, how do we expect the asset returns to change? Assuming perfect correlation, the asset becomes exposed to only a single factor and, as a result, the benefits of diversification disappear. According to conventional finance theory this should be associated with more volatile returns. Moreover, the literature on volatility timing suggests that a rise in volatility predicts lower Sharpe ratios. For example, Moreira and Muir (2017) show - for a variety of asset classes - that a rise in volatility predicts higher volatility also in the near future but it has no predictability over average returns; hence, the initial rise in volatility predicts a lower Sharpe ratio. Taken together, an increase in the correlation between the two risk factors should be associated with a contemporaneous rise in volatility and a lower predicted Sharpe ratio. Importantly, if the loadings on the two risk factors were of different signs (one positive and one negative) the implications for the returns would be reversed, i.e. a higher correlation between the risk factors would enhance the diversification benefits and be associated with lower volatility and, therefore, a higher predicted Sharpe ratio.

In this paper, I investigate whether there is empirical support for this conjecture. To obtain a high statistical power it is suitable to study an asset class in which two risk factors can explain a relatively large fraction of the variation in returns. For this purpose, we turn to currency markets. Lustig, Roussanov, and Verdelhan (2011) identify two risk factors that can price baskets of currencies (sorted by their interest rates) against the US dollar. These factors are: (i) the *dollar* factor (DOL), which is the portfolio of holding the US dollar against a basket of other currencies; and (ii) the *carry* factor (HML), which is the dollar-neutral, high-minus-low carry trade portfolio that invests in high interest rate currencies funded by short positions in low interest rate currencies. As shown by Lustig et al. (2011), the two factors can jointly explain roughly 80% of the variation in monthly returns. Lustig et al. (2011), as well as the subsequent related literature, largely treat these factors as orthogonal, which is a desirable feature since it facilitates the interpretation and the relative contribution of the factors.[1] Moreover, the cross-sectional factor loadings are as follows. Since all currency positions are taken against the US dollar they have the same, *negative* loading on the dollar factor (DOL). As for the carry factor, high (low) interest rate currencies have a positive (negative) loading, while currencies with neither high nor low interest rates have a loading close to zero. Hence, an increase in the correlation between the risk factors would enhance (reduce) the benefits of diversification for high (low) interest rate currencies. In other words, according to the conjecture, a higher correlation between the risk factors should, on average, be associated with a lower (higher) volatility of high (low) interest rate currency returns and lead to a higher (lower) future Sharpe ratio.

I find strong empirical support for the conjecture along with several important implications. First, I document that while the unconditional correlation between DOL and HML is indeed - as suggested by Lustig et al. (2011) - close to zero (-0.18 in my sample), their conditional, time-varying, correlation (CORR) is fluctuating significantly across almost the entire $[-1, 1]$ interval. While the baseline case computes CORR as an exponentially moving average across a rolling window of 112 daily observations, the results are robust also when using a simple equal-weighted moving average.

Second, I find strong and significant correlations between CORR and the volatility of high interest rate currencies (-0.60) and low interest rate currencies (0.44), respectively.[2] Even more pronounced is the correlation between CORR and the difference between the volatilities of low and high interest rate currencies, $\sigma^L - \sigma^H$, (0.87). Thus, in line with the conjecture, a change in the correlation between the risk factors is indeed associated with a change in the volatility of the test asset.

[1] Lustig et al. (2011) derive the two factors by essentially mapping them to the first two principal components from the cross-section of returns. By construction, principal components are orthogonal; hence, the implicit assumption that the risk factors are uncorrelated.

[2] In this exercise CORR is computed as the correlation between daily carry trade returns and dollar excess returns over the past month. Similarly, the currency volatilities are computed using daily currency returns over the past month.

Third, CORR contains predictive information about Sharpe ratios in the cross-section of currency returns. The same predictive information is also contained in $\sigma^L - \sigma^H$, as indicated by the result above, but the main focus in this paper will be on CORR since it lends itself nicely to economic interpretations. Further, to illustrate the predictability it is useful to divide CORR into three segments; negative, low and positive correlation.[3] For high interest rate currencies, the average next-month Sharpe ratio when CORR is negative is 0.32, while it is 1.03 when CORR is positive. For low interest rate currencies the corresponding next-month Sharpe ratios are 0.28 and −0.19, respectively. Thus, predicted Sharpe ratios of high (low) interest rate currencies are positively (negatively) related to CORR. These results extend to the HML carry trade strategy and are, in fact, even stronger. When comparing periods of negative and positive CORR, the average next-month Sharpe ratio of the HML excess returns are 0.20 and 0.99, respectively.

The predictability results are verified also in a regression analysis. In addition to CORR, the interest rate differential of the HML portfolio, $(r^H - r^L)$, is included as a regressor.[4] It is reasonable to include $r^H - r^L$ since interest rate differentials in general are known to contain predictive information about future currency returns. When regressing next-month HML returns on CORR, the interest rate differential, and their interaction term, all three regressors come out highly statistically significant and the adjusted R^2 of the regression is 4.9%. These results are robust to controlling for a large number of known predictors of foreign exchange returns.

Fourth, and finally, I find evidence that carry trade crashes are predictable when CORR is negative. From daily data I identify the 100 worst crashes of the HML carry trade strategy.[5] I ask the following question: given a positive return today, what is the likelihood that a carry

[3] CORR is considered negative, low and positive in the intervals $[-1, -0.25]$, $[-0.25, 0.25]$ and $[0.25, 1]$, respectively. Across the full sample period, the distribution of CORR over these intervals is 37%, 35% and 28%, respectively.

[4] The HML interest rate differential, $r^H - r^L$, is the difference between the average high and low interest rate, i.e. r^H and r^L denote the average interest rates among the currencies that the HML portfolio is taking a long and short position in, respectively. $r^H - r^L$ may also be referred to as the *carry* of the HML portfolio.

[5] Following, Daniel et al. (2016) and Sokolovski (2017), I define a *drawdown* as the cumulative return during periods of consecutive, negative daily returns. I then consider the 100 worst drawdowns as crashes.

trade crash starts tomorrow? As a simple test, I regress a crash indicator (1 if crash starts tomorrow; 0 otherwise) on a small set of potential predictors.[6] During periods of positive CORR the results are insignificant with an adjusted R^2 of 0.0%. When CORR is negative, most of the estimated coefficients are significantly different from zero and the adjusted R^2 is 5.2%, which is high given that these data are on *daily* frequency. One potential explanation for the large excess carry trade returns that I find, on average, for months that follow after observing a positive CORR, is that investors are requiring a compensation for the unpredictability of crashes. However, I leave further interpretations of the crash predictability to future research.

Since CORR holds predictive information about returns – potentially capturing a time-varying risk premium – it becomes important to understand the dynamics of CORR. I present two distinct economic channels that build on the definition of CORR. First, I find that large global interest rate differentials, and a high level of the US interest rate relative to the rest of the world, are associated with a high CORR. Intuitively, when US interest rates are high, the US dollar is more likely to be traded as an *investment* currency, i.e. investing in the carry trade involves taking a long position in the dollar. Subsequently, when carry trade positions are unwound, the US dollar is being sold. Thus, high US interest rates imply a positive relationship between carry trade returns and the US dollar, i.e. a positive CORR. The predictive power of CORR is then consistent with the findings of Atkeson and Kehoe (2009), who link interest rates to risks. They find that changes in short-term interest rates signal changes in bond risk premia, more so than changes in expected inflation. In our case, a higher CORR is associated with higher interest rates, which means more aggregate risk, in a broad sense.

Moreover, since interest rates are procyclical there is also a positive relationship between CORR and the US business cycle. As a result, the time-varying risk premium is procyclical, i.e. the HML carry trade is predicted to deliver high returns in good economic times and low returns in bad economic times. Since HML investors are earning low returns in times when

[6] The included regressors are: (i) the S&P 500 index; (ii) the VIX index, which is the option-implied volatility index used to gauge for equity and volatility risk (Daniel et al., 2016); (iii) the TED-spread, defined as the difference between the 3-month USD LIBOR rate and the 3-month US T-bill rate, used to gauge for funding liquidity (Brunnermeier et al., 2009); (iv) the HML interest rate differential; and (v) CORR.

their marginal rate of substitution is high, they will – according to conventional finance theory – require a positive average return for holding the carry trade portfolio over long periods of time. This is consistent with most risk-based explanations for the average, excess returns of carry trades.[7] Hence, the procyclical risk premia that CORR captures adds to our understanding of these abnormal returns.

The second economic channel is that CORR may capture the time-varying characteristic of the US dollar as a safe haven. When the conditional correlation between HML and DOL is negative, a sudden unwinding of carry trades is associated with an appreciation of the US dollar. The literature shows mixed evidence as to whether the US dollar can be regarded as a safe haven or not.[8] The mixed evidence itself is consistent with the view that the safe haven characteristic of the US dollar is, indeed, time-varying.

This is not the first paper to study correlations between returns in international capital markets. Ang and Bekaert (2002) introduce a regime-switching model to capture the empirical finding that correlations between international equity returns tend to increase during periods of high market volatility. Pollet and Wilson (2010) shows that higher aggregate risk in equity markets is associated with higher correlation between stocks. Following Pollet and Wilson (2010), Cenedese, Sarno, and Tsiakas (2014) decompose foreign exchange market variance into average variance and average correlation. They find that average variance holds predictive power of the left tail of future carry trade returns, while for the average correlation they find no predictive power. Moreover, Mueller, Stathopoulos, and Vedolin (2017) study correlations between FX pairs and find evidence that correlation risk is priced in FX markets. Aloosh and Bekaert (2017) study the co-movement of the G10 currencies, and use a new clustering technique to find a novel risk factor. Berg and Mark (2018) finds that global macroeconomic uncertainty is priced in currency markets. In contrast to these papers, I study the correlation

[7] For example, Menkhoff et al. (2012) argue similarly as they economically motivate FX volatility as a risk factor that can price carry trade returns. In a related paper, Christiansen, Ranaldo, and Söderlind (2011) find that carry trade strategies become more exposed to stock market risk during times of high FX volatility.

[8] See, for example, Kaul and Sapp (2006), Ranaldo and Söderlind (2010) and Hossfeld and MacDonald (2015).

between risk factors.

Over the years, the research literature has discovered a large number of risk factors and asset characteristics that contain predictive information about returns, mainly in equity markets. However, there is a growing concern about the actual contribution of these pricing factors. If there are significant correlations between some of the pricing factors then the number of true sources of risk may be much smaller than the number of pricing factors. In recent years, several papers have been applying different statistical techniques in an effort to reduce the number of pricing factors; see, for example, Feng, Giglio, and Xiu (2019) and Freyberger, Neuhierl, and Weber (2018), and a summary of over 300 of these pricing factors by Harvey, Liu, and Zhu (2016). This paper adds to this literature by studying the cross-sectional effects of correlations between pricing factors and thereby further motivating the need for reducing the number of actual pricing factors.

This paper is also related to the literature on volatility timing, which has already been applied to momentum strategies in equity markets (Barroso and Santa-Clara, 2015) and currency markets (Dahlquist and Hasseltoft, 2017), as well as to carry trade strategies (Daniel, Hodrick, and Lu, 2016). More generally, Moreira and Muir (2017) construct managed portfolios that are timing the returns in a variety of asset classes by switching away from risky assets when the return volatility is high. Similar to this literature I find that high volatility is associated with low future Sharpe ratios. However, while Moreira and Muir (2017) attribute the low Sharpe ratios solely to the increase in volatility, I find that higher average returns are contributing just as much. Moreover, this literature struggles with the puzzle of understanding why it is optimal for investors to switch away from risk during times of high volatility, when conventional finance theory suggests that expected returns should be high during those times.[9] In this regard, this paper contributes by being able to study CORR, rather than $\sigma^L - \sigma^H$, in the search for economic intuition and understanding.

Another related strand of the literature links the special properties of the US dollar to risk

[9] Periods of high volatility are typically associated with high risk aversion which, in turn, should be reflected in asset prices as an increase in expected returns.

premia in foreign exchange markets. For example, Rey (2015) studies the global financial cycle and suggests that the United States has a central role in determining the flow of global capital. In addition to the paper by Lustig et al. (2014) mentioned above, Verdelhan (2018) builds on the work by Lustig et al. (2011) and studies the dollar risk factor in a time-series regression of bilateral exchange rates. While these papers treat the dollar factor as being orthogonal to the carry trade factor, I show that the time-varying correlation between the two factors contain predictive information about future carry trade returns. As such, the findings in this paper pose a challenge to the existing models to account for this time-varying risk premia.

The rest of this paper is organized as follows. The next section describes the data and how we construct carry trades and the conditional correlation, CORR. Section 3 discusses the role of diversification benefits and makes the connection between CORR and future carry trade returns. Section 4 studies the predictability in a regression analysis. In Section 5 I show that CORR contains information about the predictability of carry trade crashes. Finally, Section 6 concludes.

2 Empirical setting

In this section, I describe how the carry trade strategy is implemented and present the data. I also construct the correlation measure and study some of its fundamental properties.

2.1 The carry trade strategy

Before constructing the carry trade strategy we introduce some notation and describe currency excess returns in general. Let the home currency be the US dollar, and any position (long or short) in a foreign currency is taken against the US dollar. Specifically, let S_t^j denote the exchange rate of currency j, expressed as the amount of US dollars per unit of currency j. Thus, an increase in S_t^j is associated with an appreciation of currency j relative to the US dollar. Similarly, let F_t^j denote the forward exchange rate at time $t+1$. Finally, let r_t denote the interest rate in the United States. Next, consider the trade of taking a long position in the foreign currency, and a corresponding short position in the US dollar. One way of implementing this trade is to use forward contracts. Specifically, the trade translates into buying currency j forward and then selling it in the spot market in the next period. The excess return, measured in US dollar, from doing this trade is

$$rx_{t+1}^{j} = \frac{1+r_t}{F_t^j}\left(S_{t+1}^{j} - F_t^j\right), \tag{1}$$

where the term $\frac{1+r_t}{F_t^j}$ is a scaling factor which ensures that the size of the position, at time t, is normalized to one US dollar.[10] Note that the trade could also be implemented in the spot market, as opposed to using forward contracts. However, these two alternatives are equivalent under the assumption that covered interest parity (CIP) holds.[11]

A carry trade strategy invests in high interest rate currencies, funded by short positions in low interest rate currencies. Correspondingly - since we are using forward contracts - a carry trade takes long (short) forward positions in currencies that are trading at a forward discount (premium); $F_t < S_t$ ($F_t > S_t$). To facilitate the interpretation of the results, throughout the remainder of this paper I will refer to interest rates rather than forward discounts/premiums. Moreover, let K_H and K_L denote sets of high and low interest rate currencies, respectively. Assume also that these sets contain the same number of currencies; K. We then consider an equal-weighted, dollar-neutral, carry trade strategy. Following the notation by Lustig et al. (2011), we refer to this strategy as *HML*, for high-minus-low interest rate currencies. The

[10] Normalizing the size of the position to one US dollar will facilitate the building of equal-weighted currency portfolios.

[11] Akram, Rime, and Sarno (2008) shows that prior to the financial crisis in 2008, deviations from CIP were small and insignificant for data on monthly frequency and higher. However, during and after the financial crisis, deviations from CIP have become both sizable and persistent; see, for example, Levich (2017) for a recent survey of the literature that looks at CIP deviations. Importantly however, these deviations have no implications for the monthly carry trade excess returns considered in this paper since we are using actual forward and spot rates, i.e. we do not, for example, rely on CIP to compute foreign interest rates.

excess return is then given by

$$rx_{t+1}^{HML} = \frac{1}{2K}\left[\sum_{i\in K_H} rx_{t+1}^{i} - \sum_{j\in K_L} rx_{t+1}^{j}\right]. \tag{2}$$

The total number of positions that the strategy takes is $2K$, which implies that the weight on each position is $1/2K$. Note that equal weights, and the same number of long and short positions, imply that the strategy is dollar-neutral, i.e. the short dollar exposure from positions in high interest rate currencies is exactly offset by the long dollar exposure from positions in low interest rate currencies.

Some papers take into account bid/ask spreads, in order to provide more realistic results. However, Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011) show, using a similar set of currencies and a similar carry trade strategy, that while bid/ask spreads are important in terms of quantities of payoffs, they do not explain the profitability of their carry trade. Hence, I follow them in abstracting from bid/ask spreads, and use mid-quotes for both long and short positions.

2.2 The dollar factor

The dollar factor, DOL, is the excess return on a portfolio that invests in the US dollar against a broad basket of other currencies. From (1), we see that the excess return from investing in the US dollar against currency j is $-rx_{t+1}^{j}$. Given a basket of N foreign currencies, DOL is defined as

$$DOL_{t+1} = -\frac{1}{N}\sum_{j=1}^{N} rx_{t+1}^{j}. \tag{3}$$

2.3 Data

Although the general analysis is carried out at the monthly frequency, we need daily data to construct several important quantities along the way; for example, the correlation measure, CORR. Therefore, we start from daily data on spot and 1-month forward prices, denominated in US dollars, for the G10 currencies. In addition to the US dollar (USD), these currencies include the Australian dollar (AUD), the British pound (GBP), the Canadian dollar (CAD), the Euro (EUR)[12] , the Japanese yen (JPY), the New Zealand dollar (NZD), the Norwegian krone (NOK), the Swedish krona (SEK), and the Swiss franc (CHF). The data are obtained from Reuters and Barclays (via Datastream), and the sample period stretches from December 31, 1984 to July 31, 2017. These data are available for all currencies over the entire sample period. The US interest rate is the 1-month USD LIBOR, downloaded from Bloomberg.

2.4 Excess carry trade returns

The HML carry trade strategy is implemented at the monthly frequency. At the end of each month, the G10 currencies (excluding the US dollar) are sorted by interest rates. The strategy then takes a long forward position in the three currencies with the highest interest rates, and a short forward position in the three currencies with the lowest interest rates. All positions are taken with equal weights, and against the US dollar, which leaves the total carry trade portfolio with no direct exposure to movements in the US dollar. The portfolio is then rebalanced at the end of each month.

In order to compute the correlation measure, CORR, we also need returns at the daily frequency. We obtain these from the monthly returns, in a procedure that partly follows that by Daniel et al. (2016). The main idea is to mark the portfolio to market each day. More specifically, we define a cumulative value function that starts each month with the value of zero, and ends each month with the corresponding monthly return. On each day within the month, the value function grows by the daily changes in exchange rates as well as a daily, equally weighted, proportion of the interest rate earned over the given month. Finally, using the value function, we follow Daniel et al. (2016) to back out the daily returns. The details of this procedure are presented in Appendix A.

[12] For data on the Euro prior to its introduction in 1999, we use the German Mark.

	Monthly	Daily
Mean	2.28	2.18
	[0.67]	[0.65]
Std dev	4.58	4.64
Skewness	−0.76	−0.71
Sharpe Ratio	0.50	0.47
N obs	391	8357

Table 1: Summary statistics of the HML returns, at monthly and daily frequency. Returns and standard deviations are expressed as percentages, and in annualized form; monthly (daily) mean returns are multiplied by 12 (252), and standard deviations are multiplied by $\sqrt{12}$ ($\sqrt{252}$). Reported in brackets are standard errors of the means, computed using stationary bootstrap following Politis and Romano (1994) with the average block size according to Politis and White (2004). The sample period is 12/31/1984 - 7/31/2017.

This HML strategy is basic, but still sophisticated enough to capture the typical carry trade characteristics.[13] Table 1 presents the summary statistics for the HML returns, on both monthly and daily frequency. Returns are annualized, so monthly (daily) average returns are multiplied by 12 (252), and standard deviations are multiplied by $\sqrt{12}$ ($\sqrt{252}$). The Sharpe ratio is the annualized mean return divided by the annualized standard deviation. As mentioned above, the HML returns exhibit the usual carry trade characteristics. The mean return on the monthly data is 2.28%, which is statistically different from zero. The average profitability of the carry trade is also depicted in Figure 1, which shows the cumulative HML returns over the sample period. Moreover, the standard deviation of 4.58% implies a Sharpe ratio of 0.50, which is close to what most other papers find. We also see the typical, negative skewness; in our case −0.76 for monthly returns.

Finally, we compare the summary statistics of the monthly and daily returns. We immediately see that these statistics are similar in magnitude. Thus, the method for computing daily returns is preserving the carry trade statistics, which is reassuring and gives credence to the method.

[13] The strategy is basic in the sense that we are using equal weights, as well as restricting our attention to the ten most liquid currencies. In the industry, examples of two carry trade investment products that focus on the G10 currencies are the iPath Optimized Currency Carry ETN and the Exchange Traded Fund Powershares DB G10 Currency Harvest Fund. There is a large number of papers in the academic literature that consider equal-weighted strategies for the G10 currencies, e.g. Bakshi and Panayotov (2013) and Daniel et al. (2016).

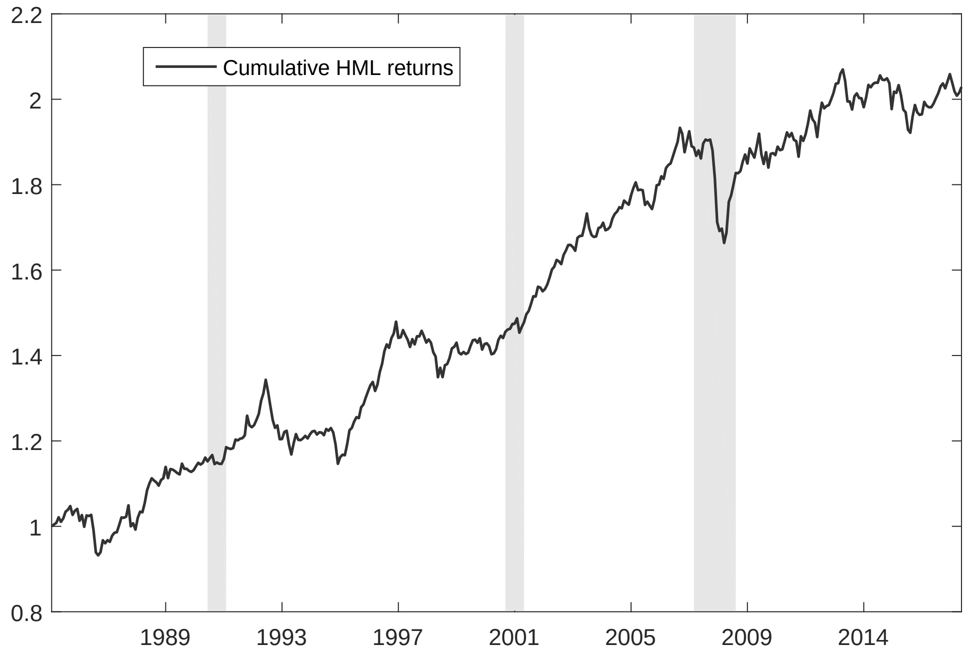

Figure 1: Cumulative returns on the HML portfolio, from investing 1 US dollar on January 1, 1985. Computed from monthly returns. Shaded areas are NBER recessions.

2.5 Constructing CORR

The correlation measure, CORR, is designed to capture the time-varying correlation between the carry trade returns (HML) and the excess returns from holding the US dollar (DOL). To compute the correlation we use a rolling window of historical, daily data points. We allow the weights of the observations within the rolling window to be exponentially distributed; giving higher weights to the more recent observations. This method for computing a conditional correlation requires the user to specify two parameters: the decay parameter, λ, which determines the relative weights; and the number of daily observations to include, T.

The same method for computing conditional correlations has been examined in detail by RiskMetrics™ (1996),[14] and they argue that it more accurately captures the true time-varying correlation than a simple equal-weighted average. Importantly, the main results of this paper

[14] RiskMetrics™ is a joint project between J.P. Morgan and Reuters. Their goal is to enhance the transparency on market risk, and make publicly available a set of techniques and data used to measure market risks. In 2010, RiskMetrics™ was acquired by MSCI. For more details, visit their website at https://www.msci.com/risk-performance.

remain strong and highly statistically significant also when CORR is computed using an equal-weighted rolling window. Moreover, in a study on foreign exchange spot data, RiskMetrics™ (1996) find that at the 0.1% tolerance level, the optimal parameters are 0.94 for the decay parameter, and 112 historical, daily data points. I follow RiskMetrics™ (1996) and set $\lambda = 0.94$, and $T = 112$ in my equations below.

CORR is computed as

$$CORR_t = \frac{\sigma_t^{HML,DOL}}{\sigma_t^{HML}\sigma_t^{DOL}},$$

where

$$\sigma_t^{HML,DOL} = (1-\lambda)\sum_{j=0}^{T-1}\lambda^j\left(rx_{t-j}^{HML} - \overline{rx}_t^{HML}\right)\left(DOL_{t-j} - \overline{DOL_t}\right),$$

$$\sigma_t^{HML} = \sqrt{(1-\lambda)\sum_{j=0}^{T-1}\lambda^j\left(rx_{t-j}^{HML} - \overline{rx}_t^{HML}\right)},$$

$$\sigma_t^{DOL} = \sqrt{(1-\lambda)\sum_{j=0}^{T-1}\lambda^j\left(DOL_{t-j} - \overline{DOL_t}\right)},$$

and where the averages are given by

$$\overline{DOL_t} = \frac{1}{T}\sum_{j=0}^{T-1} DOL_{t-j} \quad \text{and}$$

$$\overline{rx}_t^{HML} = \frac{1}{T}\sum_{j=0}^{T-1} rx_{t-j}^{HML}.$$

The resulting data series for CORR is at daily frequency, with 8246 observations. Since the main analysis in this paper uses monthly data, we construct the monthly series for CORR by using the end-of-month observations. The summary statistics for CORR is presented in Table 2. The mean is −0.06, which is indeed close to zero, while the *unconditional* correlation between HML and DOL is −0.18 (not reported in the table). Moreover, CORR varies strongly over the sample period. This is clear from the high standard deviation (0.44) as well as from

Summary statistics of CORR	
Mean	−0.06
Std dev	0.44
Min	−0.90
Max	0.92
N obs	386

Table 2: The table shows the summary statistics of CORR, defined as the conditional correlation between HML and DOL. The data are at monthly frequency, constructed from using end-of-month daily observations.

the minimum and maximum values; −0.90 and 0.92, respectively.

These characteristics of CORR are also visible in Figure 2, which plots CORR over the full sample period. From the figure we see that CORR is not only varying, but it is also fairly persistent. For example, in an AR(1) regression of CORR on the one-month lagged CORR, the estimated slope coefficient is 0.84. This persistence is of course partly mechanical, from the rolling window used in the construction of CORR. However, as a robustness check I restrict the rolling window to contain only the daily observations within each given month. That way, there is no mechanical correlation between the monthly observations of CORR. The corresponding estimated slope coefficient from the AR(1) regression is still high, at 0.66 and highly statistically significant. Most importantly, the main results of the paper are robust to using CORR constructed this way; with a non-overlapping rolling window.

3 Linking CORR to currency returns

In this section I construct currency portfolios to study the connection between CORR and the cross-section of currency returns. We will also study the dynamics of CORR to better understand some of its driving forces.

In order to study the cross-section of currency returns it is useful to consider portfolios of currencies since it reduces the noise in returns coming from currency-specific factors.[15]

[15] Lustig and Verdelhan (2007) were the first to study portfolios of currencies sorted by their interest rates.

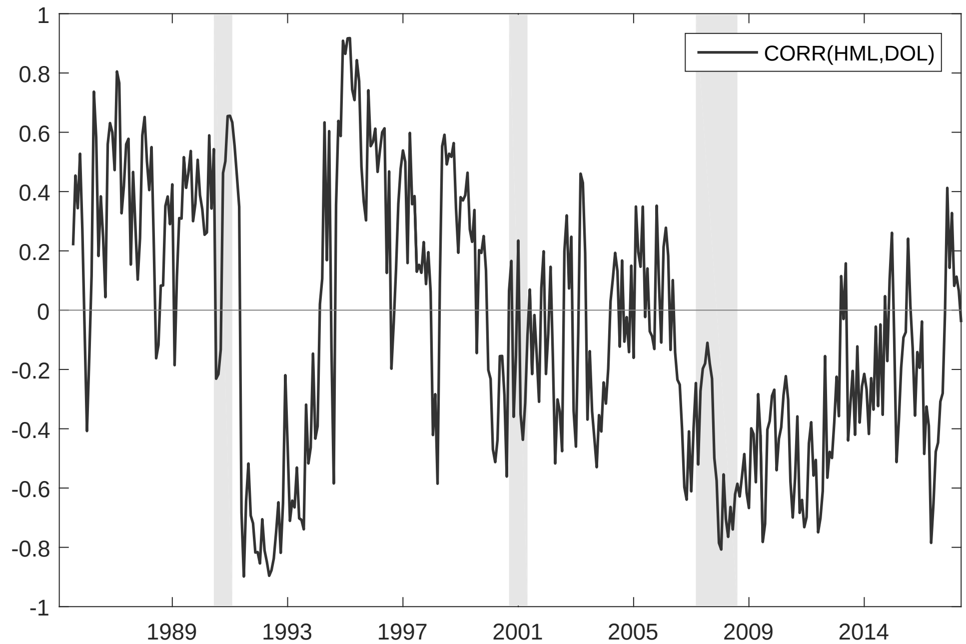

Figure 2: The figure plots CORR; the conditional (time-varying) correlation between HML and DOL. The data are at monthly frequency, constructed from using end-of-month daily observations. Each observation is computed using 112 historical daily observations; weighted exponentially giving higher weights to more recent observations. Shaded areas are NBER recessions.

Whereas Lustig et al. (2011) divide their developed currencies into five portfolios, I construct three portfolios, each containing three currencies. That is, at the end of each month, the nine non-USD G10 currencies are sorted by their interest rates. The currencies are then placed in three portfolios, denoted P1, P2 and P3, where portfolio P1 contains the currencies with the lowest interest rates, and P3 contains the currencies with the highest interest rates. Note that our HML carry trade portfolio is, by definition, equivalent to taking a long position in P3 funded by a short position in P1.

3.1 CORR and the cross-section of currency returns

Table 3 reports the summary statistics of the three portfolios when using the full sample. As expected, the annualized average returns of the portfolios are increasing with the interest rate; portfolios P1 and P3 have average returns of 0.6% and 5.1% respectively. The Sharpe ratios follow the same pattern, increasing from 0.06 for P1 to 0.49 for P3. Standard deviations, on the other hand, seem to show no clear pattern, which is consistent with the findings of Menkhoff et al. (2012). While unconditional standard deviations do not vary in the cross-section, conditional standard deviations can vary significantly, as we will see below.

One might be surprised to see a positive average return for the low-interest rate portfolio, P1. Lustig et al. (2011) report a negative average return on their portfolio containing the lowest interest rates. However, this can be explained by the negative excess return on the US dollar over our sample period. Since all portfolios (P1, P2 and P3) are short the US dollar, a depreciation of the US dollar elevates the returns of all portfolios. Specifically, the cumulative return from holding the US dollar against the other G10 currencies over our sample period is -63%. Most of these negative returns are earned during the first three years of the sample. Lustig et al. (2011) start their sample one year before our sample, and during that year, in 1984, the US dollar appreciated heavily.

Following Lustig et al. (2011), we assume a stochastic discount factor which is linear in the two factors; rx^{HML} and rx^{DOL}. Based on the implied beta-representation we can estimate the factor loadings from the following regression:

$$rx_t^P = \alpha + \beta_P^{DOL} \cdot rx_t^{DOL} + \beta_P^{HML} \cdot rx_t^{HML} + \epsilon_t, \tag{4}$$

where β_P^{DOL} and β_P^{HML} are the factor loadings on rx_t^{DOL} and rx_t^{HML}, respectively, and rx_t^P is the excess returns on a portfolio of currencies against the US dollar. For the three portfolios, we estimate the factor loadings and report them in Table 3.

As expected, all portfolios have a similar, negative loading on the dollar factor. The loading on the carry factor is negative for the low interest rate portfolio (P1) and positive for the high interest rate portfolio (P3). Portfolio P2 has no significant exposure to the carry factor. These findings are all consistent with those of Lustig et al. (2011).

At this point we may remind ourselves of the conjecture that motivated this paper. An increase in the correlation between the two risk factors should enhance the benefits of diversification for high interest rate currencies, which follows from the two factor loadings carrying *different* signs. This should be reflected in a contemporaneous decrease in volatility and, through the effects of volatility-timing, a higher Sharpe ratio in the following periods. The implications for low interest rate currencies are reversed since their factor loadings are of the same sign; both are negative.

Next, we test the conjecture by observing how the returns of the currency portfolios varies for different levels of CORR. It is natural to split the correlation interval into three parts, to capture the characteristics of negative, low, and positive correlation; these intervals are set to $[-1, -0.25]$, $(-0.25, 0.25)$ and $[0.25, 1]$, respectively. That way, CORR is well-represented across the sample period in all three intervals, as 37%, 35% and 28% of the observations are classified as negative, low and positive, respectively. Moreover, since focus is on the *predictive* power contained in CORR, the main analysis will consider CORR as being observed at the end of a given month, while currency returns are observed over the following month. The summary statistics for the next-month currency portfolio returns are presented in the last three columns of Table 3.

For high interest rate currencies, the next-month standard deviation of returns is, on average, 12.6% when CORR is negative. When CORR is positive the corresponding standard deviation is only 7.6%. For low interest rate currencies the relationship with CORR is reversed as the average next-month standard deviation of returns are 7.3% and 10.9% during negative and positive periods of CORR, respectively. Note that these are *next-month* standard deviations, while the corresponding, contemporaneous results are even stronger. In an additional test of the connection between CORR and standard deviations of cross-sectional returns, I compute the contemporaneous, unconditional correlation between CORR and high (low) interest rate currencies and obtain an estimate of -0.60 (0.44). When considering the difference between the standard deviations of low and high interest rate currencies, $\sigma^L - \sigma^H$, the correlation is even more pronounced, at 0.87. Taken together, this is evidence supporting

		Full sample	CORR from previous month		
			Negative $[-1, -0.25]$	Low $(-0.25, 0.25)$	Positive $[0.25, 1]$
	Mean	0.6	2.1	1.1	−2.1
P1	Std Dev	9.0	7.3	8.9	10.9
(low rates)	Skewness	0.21	0.14	0.51	0.14
	Sharpe Ratio	0.06	0.28	0.12	-0.19
	β^{DOL}	-1.04^{***}			
	β^{HML}	-1.02^{***}			
	Mean	2.3	1.9	2.4	3.0
P2	Std Dev	8.2	9.5	7.8	6.6
	Skewness	−0.36	−0.47	−0.08	−0.23
	Sharpe Ratio	0.28	0.20	0.30	0.45
	β^{DOL}	-0.93^{***}			
	β^{HML}	0.04			
	Mean	5.1	4.0	3.8	7.8
P3	Std Dev	10.4	12.6	9.6	7.6
(high rates)	Skewness	−0.32	−0.50	0.15	−0.17
	Sharpe Ratio	0.49	0.32	0.39	1.03
	β^{DOL}	-1.04^{***}			
	β^{HML}	0.98^{***}			
	Mean	2.3	1.0	1.3	5.0
HML	Std Dev	4.6	4.8	3.9	5.0
(P3-P1)	Skewness	−0.76	−0.47	−0.84	−1.15
	Sharpe Ratio	0.50	0.20	0.34	0.99
	Nr. of observations	386 (100%)	143 (37%)	136 (35%)	107 (28%)

Table 3: Summary statistics of next-month excess returns (after observing CORR). From the nine non-USD G10 currencies, the three high (low) interest rate currencies are placed in portfolio P3 (P1). The reported values for β^{DOL} and β^{HML} come from estimating the regression: $rx_t^P = \alpha + \beta^{DOL} \cdot rx_t^{DOL} + \beta^{HML} \cdot rx_t^{HML} + \epsilon_t$, where rx_t^P corresponds to the excess returns of portfolio $P1$, $P2$ and $P3$. The frequency of the data is monthly, and the sample period is December 31, 1984 to July 31, 2017. Returns and standard deviations are expressed as percentages, and in annualized form; mean returns are multiplied by 12 and standard deviations are multiplied by $\sqrt{12}$. Sharpe ratios are computed by dividing the annualized mean return by the annualized standard deviation. ***, ** and * denote significance at the 1%, 5% and 10% level, respectively.

our notion that a change in the correlation between the risk factors has effects on the diversification benefits across the factors and thereby also effects on the volatility of returns.

The pattern for next-month Sharpe ratios are also in line with the conjecture. For high (low) interest rate currency returns, the next-month Sharpe ratio is, on average, 0.32 (0.28) when CORR is negative, while during periods of positive CORR the corresponding Sharpe ratio is significantly larger (smaller); 1.03 (−0.19). Note that the differences in Sharpe ratios are not solely driven by the change in volatilities, as found in the cross section considered by Moreira and Muir (2017). The average returns tend to change in the direction that further amplifies the change in Sharpe ratios. For example, when comparing negative to positive periods of CORR, for high (low) interest rate currencies, the average next-month return is 4.0% (2.1%) and 7.8% (−2.1%), respectively.

The skewness of next-month returns also displays an interesting pattern, with implications for the US dollar. For all three portfolios in the cross-section, the highest skewness is found when CORR is low, i.e. when HML and DOL are roughly uncorrelated. During periods of a more distinct correlation - positive or negative - the skewness is significantly lower. A more negative skewness has been associated with higher crash risk (Brunnermeier et al., 2009), implying that the crash risk increases with a stronger correlation between HML and DOL. However, since the same pattern exists throughout the entire cross-section, it is most likely capturing a characteristic of the US dollar (recall that all portfolios are short the US dollar). Hence, the lower skewness when CORR is either negative or positive implies a higher risk of a sharp *appreciation* of the US dollar.

3.2 CORR and carry trade returns

Finally, we turn to the results on the carry trade portfolio, which are displayed in the bottom panel of Table 3. Focusing first on the volatility, we note that the cross-sectional differences in standard deviations cancels out for the carry trade returns. To see this, recall that the carry trade takes a long position in P3 and a short position in P1. Importantly, an increase in the volatility of either P3 or P1 will both result in a higher volatility of the carry trade returns. As a result, when comparing periods of negative and positive CORR, the higher volatility in P1 is offset by the lower volatility of P3. Further, the patterns for the next-month average returns and Sharpe ratios are clear. When CORR is negative, the average next-month carry trade returns are 1.0% and the Sharpe ratio is 0.20. During periods of positive CORR, the average next-month returns are 5.0% and the next-month Sharpe ratio is 0.99.

The skewness of the carry trade returns show a negative relationship with CORR. Comparing periods of negative and positive CORR, the average skewness of next-month carry trade returns are -0.47 and -1.15, respectively. Potentially, this indicates a higher crash risk when CORR is positive. These results are discussed further in Section 5 as we connect CORR to the predictability of carry trade crashes.

3.3 Understanding CORR

Since CORR contains predictive information about future currency returns it is important to understand the drivers behind CORR itself. From the previous section we know that it has an 87% correlation with the cross-sectional difference in volatilities, $\sigma^L - \sigma^H$. However, apart from that, it is less clear what characterizes periods of positive CORR compared to periods of negative CORR. The goal of this section is shed some light on these periods by studying common financial and macroeconomic variables.

Table 4 shows regression results and correlations using seven different variables. Far more variables have been tested, but these seven will suffice in painting the picture of the dynamics behind CORR. To explain the table, the column labeled $Corr(CORR, \cdot)$ shows the unconditional correlation between CORR and the variable in the first column. Column three and four - labeled β and R^2 - show the results from regressing CORR on the variable in the first column. Reported in the two columns are the estimated slope coefficients, β, and the adjusted R^2. The final column shows the results from regressing CORR on all variables in the first column, in a single regression. Moreover, the following analysis is divided into three parts: interest rates; volatility; and finally we take a broad look at liquidity, equity returns and macroeconomic indicators.

	$Corr(CORR,\cdot)$	β	R^2	Regression of CORR
T-bill 1-year	0.56	0.09***	31.2	0.06***
		(8.95)		(3.12)
$r^H - r^L$	0.43	74***	18.0	64**
		(6.66)		(2.50)
AFD	0.14	42	1.9	84**
		(1.19)		(2.42)
$FXvol$	−0.32	−92***	9.8	−33**
		(−5.39)		(−2.00)
TED	0.24	0.24*	5.4	−0.09
		(1.93)		(−1.12)
$S\&P500$	0.19	1.94***	3.4	1.50***
		(3.45)		(2.93)
IP	0.13	0.10**	1.5	0.02
		(2.12)		(0.66)
Intercept	-	-	-	−84**
	-	-	-	(−2.43)
N obs	-	-	-	378
$AdjR^2(\%)$	-	-	-	43.5

Table 4: The column labeled $Corr(CORR,\cdot)$ shows the unconditional correlation between CORR and the variable in the first column. Column 3 and 4, labeled β and R^2, show the results from regressing CORR on the variable in the first column. Reported in the columns are the estimated slope coefficients, β, and the adjusted R^2. The final column shows the results from regressing CORR on *all* variables in the first column, in a single regression. Parentheses are t-stats, computed from Newey and West (1987) standard errors, with optimal lags according to Andrews and Monahan (1992). All regressions and correlations are computed on contemporaneous variables, hence we have left out the time subscripts to simplify notation. ***, ** and * denote significance at the 1%, 5% and 10% level, respectively.

3.3.1 Interest rates

The strongest link to CORR is through the 1-year T-bill rate, which has an unconditional correlation with CORR of 0.56. In a regression of CORR on the 1-year T-bill rate, the estimated slope coefficient is highly significant and the resulting adjusted R^2 is 31.2%. Further, the average interest rate differential of the HML portfolio, $r^H - r^L$, is positively related to CORR, with a correlation of 0.43.

Another variable related to interest rates is the average forward discount (AFD), which is essentially the difference between the US interest rate and the average foreign interest rate. Lustig et al. (2014) show that the AFD contains predictive information about the US dollar. Also the AFD has a positive, although small, correlation with CORR; 0.14.

Taken together, these variables are capturing different aspects of interest rates. In the regression on the complete set of variables, all three variables come out significant, suggesting that each of them are providing important information for understanding the variation in CORR. The main takeaway is that a high US interest rate and large differences in global interest rates are associated with a high correlation between HML and DOL. One possible explanation is that when US interest rates are high, the US dollar is more likely to be traded as an investment currency. That is, when investors are taking long positions in the carry trade they tend to buy the US dollar. When carry trade positions are unwound, the US dollar is being sold. That way, a high US interest rates will lead to a positive relationship between carry trade returns (HML) and the US dollar (DOL).

3.3.2 Volatility

To proxy for the volatility in foreign exchange markets, I use the measure introduced by Menkhoff et al. (2012).[16] As reported in Table 4, its unconditional correlation with CORR is -0.32. However, this negative relationship is not equally pronounced across the $[-1, 1]$ correlation interval. Figure 3 shows a scatter plot of FX volatility against CORR. Clearly, much of the negative relationship comes from the 12 monthly observations with the highest FX volatility, which all coincided with strong negative CORR. Thus, we can conclude that spikes in FX volatility are associated with extreme negative correlation between HML and DOL. Finally, we note that the 11 months with the highest FX volatility occurred between August 2008 and October 2011; eight of those eleven months were in 2008-2009.

The extreme negative correlation during volatility spikes indicates that there is a flight-to-safety mechanism present. In times of high uncertainty, investors tend to unwind their risky positions and invest in safer assets (or, safe havens). If the US dollar is a safe haven, then higher-than-usual FX volatility may cause investors to unwind their carry trade positions and

[16] Another FX volatility proxy that I have tried is the one used to predict carry trade returns by Bakshi and Panayotov (2013). I have also used equity-based volatility proxies, such as the VIX; the option-implied volatility of the S&P500. However, none of those proxies carries significant information in explaining CORR beyond the information already captured in the Menkhoff et al. (2012) proxy.

invest in, for example, US Treasury bills. That behavior would impose a negative correlation between HML and DOL, which is what we see during extreme levels of high volatility. But the literature suggests that the US dollar may not always be regarded as a safe haven.

Ranaldo and Söderlind (2010) study several currencies over the period 1993-2008 to see if they qualify as safe havens. They find that only the Swiss Franc and the Japanese Yen qualify, while the US dollar does not. However, others have found that the US dollar indeed may serve as a safe haven. For example, Kaul and Sapp (2006) shows that the US dollar was used as a safe haven during the uncertainty surrounding the millennium shift. Moreover, Hossfeld and MacDonald (2015) find that after controlling for the flow of funds attributable to the unwinding of carry trades, the Swiss Franc and the US dollar qualify as safe havens; not the Japanese Yen.[17]

This mixed evidence is consistent with the view that the safe haven characteristic of the US dollar is time-varying. The conditional correlation between HML and DOL may well capture this behavior, i.e. when CORR is negative, the US dollar tends to serve as a safe haven. This is, for example, what we saw during the financial crisis of 2008-2009.

3.3.3 Liquidity, equity returns and macroeconomic indicators

The TED-spread is defined as the difference between the 3-month USD LIBOR rate and the 3-month US T-bill rate. It is a popular proxy for funding liquidity; see, for example, Brunnermeier et al. (2009). An increase in the TED-spread generally implies that banks are facing more liquidity problems. The correlation between TED and CORR is 0.24. However, the estimated slope coefficient for TED in the simultaneous regression on all variables is insignificantly different from zero. We reach the same conclusion when using the proxy for global equity market liquidity provided by Karnaukh et al. (2015).

Equity returns are also positively correlated with CORR. The slope coefficient for the S&P500 index comes out positive and significant, also in the regression on all variables, in-

[17] Hossfeld and MacDonald (2015) find that the appreciation of the Japanese Yen, in times of financial stress, can solely be attributable to the unwinding of carry trades, which typically involves purchasing the Yen.

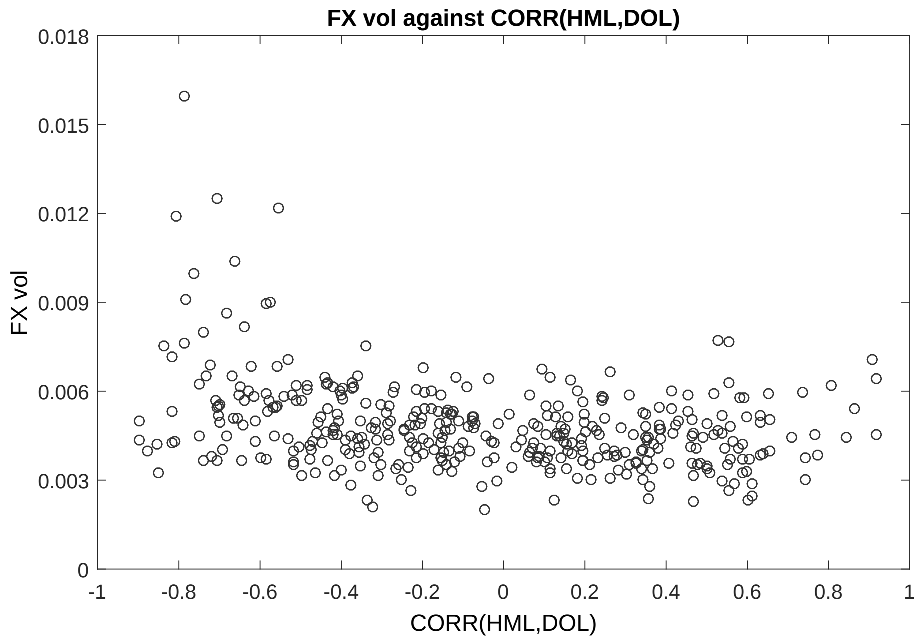

Figure 3: A scatter plot of FX volatility against CORR; the correlation between HML and DOL. The proxy for FX volatility is the one introduced by Menkhoff et al. (2012).

dicating that higher stock returns are associated with a higher correlation between HML and DOL. The same relationship holds when we look at *global* stock market returns; using the MSCI global stock market index.

Finally, as a macroeconomic indicator we use US industrial production (IP), since Ludvigson and Ng (2009), Duffee (2011) and Joslin et al. (2014) have shown that it contains information about bond risk premia which is not captured by interest rates. Given the strong connection that we have already found between interest rates and CORR, IP may hold additional information. Also, Lustig et al. (2014) document that US industrial production can be used to predict returns on baskets of currencies against the US dollar, which makes it a potential candidate for explaining CORR. However, the connection with CORR is only weakly positive, and its estimated slope coefficient in the simultaneous regression on all variables comes out insignificant. The same result is obtained when using data on US payrolls, which is another commonly used macroeconomic indicator.

To summarize, we have linked CORR to several variables; including interest rates, FX volatility and equity returns. Given these variables, we can explain almost 44% of the variation in CORR. We have also suggested that CORR may be linked to the time-varying characteristic of the US dollar as a safe haven. However, recall that CORR is directly linked to the cross-sectional dispersion in volatilities, $\sigma^L - \sigma^H$. On the other hand, by studying CORR - as we have done throughout this section - we can form a better economic intuition for the drivers of CORR.

4 Regression analysis

In this section I present a regression analysis to further study the links between CORR and next-month carry trade returns. The main regression takes the following form:

$$rx_{t+1}^{HML} = \beta_0 + \beta_1 \left(r_t^H - r_t^L\right) + \beta_2 CORR_t + \beta_3 \left(r_t^H - r_t^L\right) CORR_t + \epsilon_{t+1}, \quad (5)$$

where $r_t^H - r_t^L$ is the interest rate differential of the HML portfolio, i.e. the difference between the average interest rate of the three high-interest-rate currencies and the average interest rate of the three low-interest-rate currencies. That is, in equation (5) we regress next-month HML returns on the interest rate differential, CORR, and their interaction term. We proceed with a brief discussion on the interest rate differential before taking a closer look at the two-dimensional plane that CORR and $r^H - r^L$ span. The patterns observed there will then be reflected also in the regression results.

4.1 The interest rate differential

It is well established in the literature that the cross-sectional difference in interest rates hold information about short-run future currency returns - high interest rate currencies tend to depreciate less vis-à-vis low interest rate currencies than suggested by the uncovered interest parity (UIP).[18] This is the anomaly behind the abnormal returns on carry trade strategies. However, in addition to the cross-section, variation in interest rate differentials also contain important information across time; see, for example, Hassan and Mano (2014), who investigate several carry trade strategies while making a distinction between cross-sectional risk and time-varying risk. In a cross-time study of carry trades it is therefore natural to include the interest rate differential.

There is also economic intuition behind the importance of the interest differential across time. Specifically, it is reasonable to believe that during times when interest rate differentials are large, the carry trade is more attractive to investors - simply because they can lock in a larger interest rate differential. As a consequence, a larger change in the exchange rate is required to wipe out the profit earned from the interest rate differential. The literature offers many risk-based explanations for the excess returns on carry trades, and it is also reasonable to believe that those risks should be more pronounced during times of high interest rate differentials.[19] This paper is not the first to include the interest rate differential in a similar regression; see, for example, Verdelhan (2018) who includes the interest rate differential and interacts it with another regressor.

4.2 A simple look at the data

Together, CORR and the interest rate differential span a two-dimensional plane which contains information about future excess returns on our dollar-neutral carry trade strategy, HML. In order to visualize this in a graph, we discretize the two variables in the following way. CORR is grouped into six bins across the $[-1, 1]$ interval. The interest rate differential is

[18] The uncovered interest parity states that the exchange rate between two currencies is expected to change by the difference in their interest rates, such that the expected excess return from investing in the currency pair is zero. However, as Hansen and Hodrick (1980) and Fama (1984) (among others) have shown, the UIP fails to hold empirically as high interest rate currencies tend to depreciate less against low interest rate currencies than suggested by the UIP. High interest rate currencies actually tend to *appreciate* against low interest rate currencies. Bansal and Dahlquist (2000) show, however, that this result is less clear for emerging market currencies.

[19] See, for example, the first paper of this thesis, in which I find that the global interest rate differential, based on the G10 currencies, contains predictive information about the future excess returns of a basket of currencies against the US dollar.

grouped into either 'high' or 'low' relative to the median interest rate differential over the sample period. Figure 4 shows two panels, with bars that measure the average next-month HML return. The top (bottom) panel contains observations with high (low) interest rate differentials, at time t. Note that Figure 4 simply plots the data, without imposing any model or restrictions.[20]

In Figure 4, there are 12 bins in total. Over the sample period, we do not spend the same amount of time in each bin. Therefore, one concern may be that the bars in the figure are misleading in terms of displaying when the excess returns (gains and losses) are actually taking place. For example, the large average excess return when the interest rate differential is high and CORR is greater than 0.6 might be driven by only a few extreme observations. However, in a similar figure not included in the paper, I show that even after controlling for the time spent in each bin, the patterns that we see in Figure 4 remain just as clear.

Both panels in Figure 4 display distinct patterns. In the top panel - when the interest rate differential is high - there is a *positive* relationship between CORR and next-month carry trade. Whereas in the bottom panel - when the interest rate differential is low - the relationship between CORR and next-month carry trade returns is *negative*. From the figure we can also observe when large average gains and losses occur. Months with large excess returns on the HML strategy on average are preceded by months of either high interest rate differentials and strongly positive CORR (above 0.3), or low interest rate differentials and strongly negative CORR (below -0.6). On average, large losses on the HML strategy are preceded by months of high interest rate differentials and low CORR (below -0.3). This pattern can, at least partially, be explained by crashes in the carry trade connected to the US business cycle.

[20] If we only had one independent variable, the straight-forward way of visualizing the data would be to use a scatter plot. With two independent variables, however, it is useful to discretize the variables somehow, in order to produce an intuitive picture of the data.

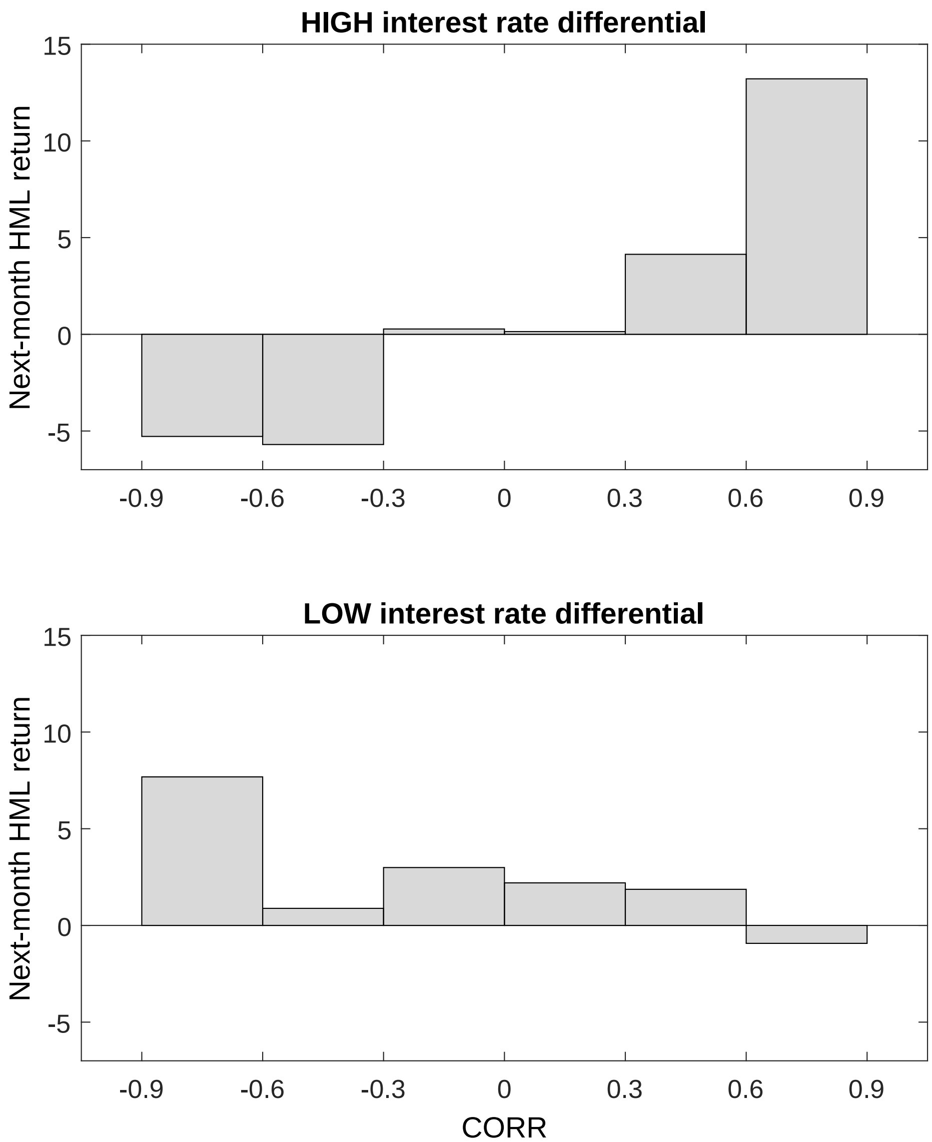

Figure 4: The distribution of next-month HML returns, across CORR and high/low HML interest rate differential. The 392 months in our sample are each placed into one of the 12 bins above. At the end of each month (at time t), we observe CORR (the conditional correlation between HML and DOL) and the interest rate differential (observed at time t, and applicable between t to $t+1$). The interest rate differential for a given month is then considered high (low) if it is higher (lower) than the median interest rate differential over the sample period. Each month is also placed in one of nine bins across the correlation interval $[-1, 1]$. The bars measure the average next-month excess return (annualized and measured in per cent) on the HML carry trade, earned between time t and $t+1$. The sample period is December 1984 to July 2017.

4.3 Connections to the US business cycle

There is a strong connection between CORR and the US business cycle. When the US economy is expanding, we typically see rising interest rates, as the Federal Reserve is hiking its policy rate to cool down the economy. When the levels of interest rates are rising we also tend to see rising interest rate differentials.[21] CORR is positively related to the US interest rate, suggesting that it tends to increase during economic expansions. Good economic times are typically also associated with credit growth and high risk appetite, and investors are attracted by risky investments such as the carry trade.[22] Thus, during economic expansions and towards the peak of the business cycle, we see high interest rate differentials and high CORR, as well as good carry trade returns. This is exactly what we see in the data, as depicted in the top panel of Figure 4.

The standard carry trade strategies, including the HML strategy, have been shown to crash occasionally as investors rapidly unwind their positions.[23] In general, these periods are associated with a sudden drop in risk appetite and investors therefore tend to shift from risky assets (e.g. carry trades) to safer assets (e.g. the US dollar). This shift results in depressed carry trade returns and, typically, an appreciating dollar, implying a negative correlation between HML and DOL, i.e. a negative CORR. Moreover, there are several potential reasons for the initial drop in risk appetite; for example, technical factors related to crowdedness (Sokolovski, 2017) and liquidity spirals (Brunnermeier et al., 2009). It could also be triggered by adverse macroeconomic shocks towards the end of the business cycle (Minsky, 1986; Cochrane, 2017), when interest rate levels, as well as interest rate differentials, tend to be high. In this case, heavy neg-

[21] For example, in our data, the unconditional correlation between the US 1-year T-bill rate and the HML interest rate differential is 0.68. Since the financial crisis of 2008 we have seen extremely low levels of interest rates globally, which may suggest that the high correlation is driven by the post-crisis part of our sample. However, even if we restrict our sample to 1985-2007, the unconditional correlation is 0.56.

[22] Collin-Dufresn, Goldstein, and Martin (2001) have shown that VIX serves as a good proxy for risk appetite, which is a link also made by Brunnermeier et al. (2009). In our data, the macroeconomic indicators (US payrolls and US industrial production) are negatively correlated with both VIX and the FX volatility proxy, suggesting that risk appetite is, in general, high during economic expansions and low during recessions.

[23] See, for example, Brunnermeier, Nagel, and Pedersen (2009), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2015), Farhi and Gabaix (2016), Sokolovski (2017) and Chernov, Graveline, and Zviadadze (2018).

ative returns related to deteriorating macroeconomic conditions are consistent with the pattern in Figure 4, as seen in the top panel with high interest rate differentials and negative CORR.

Shortly after these crashes, as risk sentiment tends to turn, the carry trade usually recovers by delivering positive excess returns.[24] At this point, it is likely that the Federal Reserve (the US central bank) has already reacted by lowering their policy interest rate, which typically spills over to lower global interest rate differentials.[25] Thus, this part of the business cycle is reflected in Figure 4, as the positive carry trade returns when CORR is negative and the interest rate differential is low.

4.4 Regression results

The general patterns in the data - as observed in Figure 4 - are captured also as we estimate our main regression specified in equation (5). The results are presented in Table 5, where regression specification (iii) corresponds to the main regression in equation (5). The estimated coefficients for the interest rate differential, CORR, as well as their interaction term come out highly statistically significant. The adjusted R^2 of the regression is 4.9%, which is high given the relatively short forecast horizon of one month.

How do we interpret the signs of the estimated coefficients in the main regression specification (iii), and can we connect them to the observed patterns in Figure 4? The estimated coefficients for both the interest rate differential and CORR are negative, while the interaction term has a positive estimate. The interaction term captures the joint movement of the two variables. That is, it takes a high absolute value when both the interest rate differential and CORR take high absolute values. For example, the positive coefficient estimate implies that when the interest rate differential is large and CORR is positive (HML and DOL are positively correlated), we predict high excess returns on the HML portfolio over the next month. This is

[24] For example, in connection to the great financial crisis the HML carry trade suffered large losses between August 2008 and January 2009, with a cumulative return of -13%. Over the five months that followed, the carry trade recovered and delivered a cumulative return of 10%; see Figure 1.

[25] Global interest rate are highly correlated with each other. The average unconditional correlation between the 1-month USD LIBOR and the corresponding interest rates for the other G10 economies is 0.74.

capturing the positive excess returns seen in the top panel of Figure 4.

Now, suppose that the interest rate differential is still large, but CORR is instead strongly negative. The interaction term then switches sign to negative, and the result is a prediction of negative excess returns on the HML portfolio over the next month. This is capturing the strongly negative next-month excess returns seen in the top panel of Figure 4. Note that as long as the interaction term is relatively large, it will dominate the single effects coming from the interest rate differential and CORR alone.

Next, suppose that the interest rate differential is low, i.e. $r^H - r^L$ is close to zero. Then, both the interaction term and the interest rate differential itself are close to zero, and only the coefficient on CORR matters. It has a negative sign, suggesting that a higher (lower) CORR predicts lower (higher) next-month excess returns. This is exactly what we see in the data, depicted in the lower panel of Figure 4.

Taken together, the high adjusted R^2 and the high statistical significance of the coefficient estimates are demonstrating the strong predictive information of carry trade returns contained in CORR and the interest rate differential.

The results are also economically significant. Let us see how the predicted next-month HML return changes as CORR changes from 0 to 0.5. The corresponding change in the interest rate differential is obtained from the fitted regression of the interest rate differential on CORR.[26] Expressed as effective annual rates, the interest rate differentials are 5.55% and 7.10% when CORR is 0 and 0.5, respectively.[27] Finally, the predicted next-month HML return changes from 0.63% to 4.34% as CORR changes from 0 to 0.5. That is, the predicted return changes by 3.71 percentage points, which almost corresponds to a one-standard-deviation change; recall from Table 1 that the standard deviation of HML returns is 4.58%.

One concern might be that these results are driven by a particular event in history, such as the financial crisis, or an extreme period, such as the late 1980s and early 1990s when

[26] The results from regressing the interest rate differential on CORR are as follows. The estimates for the intercept and the slope coefficient are 0.005 and 0.002, respectively, and both estimates are highly statistically significant. The adjusted R^2 is 18%.

[27] Note that the average interest rate differential over the sample period is 5.54%.

interest rates were high relative to the rest of the sample period. To address this I split the sample period in the middle to create two subsamples and run the main regression on each subsample. Although not reported in the paper, the main results are robust to this, as the coefficient estimates for both subsamples come out statistically significant and the adjusted R^2s remain large.

Another potential concern about the predictive power of CORR is that it may simply be picking up the effect of some other variable(s). To address this in the regression analysis I control for a set of other variables that could potentially be linked to future excess carry trade returns. Some of the results from those regressions are included in Table 5, in specifications (iv) - (ix). For each variable I also include its interaction term with the interest rate differential. The results clearly show that the predictive power of CORR and its interaction term with the interest rate differential remain highly statistically significant even after controlling for all other potential predictors. Nevertheless, from these regressions we make some useful observations about these other variables.

The proxy for FX volatility (Menkhoff et al., 2012) comes out significant when interacted with the interest rate differential. As expected, it carries a negative sign implying that higher FX volatility (coupled with large interest rate differentials) predicts lower excess carry trade returns. This is consistent with the previous literature; see, for example, Menkhoff et al. (2012) and Bakshi and Panayotov (2013).

Moreover, I also control for US payrolls, which is a common US macroeconomic indicator, available at monthly frequency. In regression specification (viii) its estimated coefficient is negative, which is counter-intuitive since we would expect good macroeconomic conditions to predict higher excess returns. However, if we exclude the interaction term between Payrolls and the interest rate differential (not reported in the table), the estimated coefficient on Payrolls itself becomes insignificant. On the other hand, in the final regression specification (ix), all control are included, and both variables with Payrolls come out significant. As with the case of CORR, the interaction term is positive, indicating that good economic conditions does indeed predict higher excess carry trade returns. Finally, we also note that when including all control variables, the coefficient for the TED-spread interacted with the interest rate differential becomes significant, with a negative sign. This is intuitive, since a higher TED-spread is interpreted as increased funding illiquidity, which should be associated with lower future excess returns.[28]

Not reported in the table are some other variables that I have controlled for, which also came out insignificant. For example, the foreign exchange volatility proxy introduced by Bakshi and Panayotov (2013), and the measure for global market illiquidity by Karnaukh, Ranaldo, and Söderlind (2015). Finally, the results are essentially unchanged if I substitute $\sigma^L - \sigma^H$ for CORR, i.e. the predictive information contained in CORR is reflected also in $\sigma^L - \sigma^H$, and the latter does not contain any new predictive information. In a horse race between the two (regressing next-month carry trade returns on the CORR, $\sigma^L - \sigma^H$, as well as their interaction terms with the interest rate differential and the interest rate differential itself), all variables come out insignificant, except for the interest rate differential, indicating that the regression model is unable to determine which variable is more significant.

[28] For a longer discussion on the connection between carry trades and the TED-spread see Brunnermeier, Nagel, and Pedersen (2009).

Dependent variable: rx_{t+1}^{HML}	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
$r^H - r^L$	0.01		−0.95***	−1.21	1.43	0.05	−1.00***	−1.85**	−0.63
	(0.04)		(−2.67)	(−0.95)	(1.08)	(0.07)	(−2.81)	(−2.13)	(−0.30)
$CORR$		0.003	−0.01***	−0.01***	−0.01***	−0.01***	−0.01***	−0.01*	−0.01*
		(1.47)	(−2.84)	(−2.80)	(−3.22)	(−2.65)	(−2.77)	(−1.78)	(−1.73)
$\left(r^H - r^L\right) CORR$			3.07***	3.24***	2.95***	3.47***	3.05***	2.33***	2.36***
			(4.15)	(4.27)	(4.57)	(3.77)	(4.19)	(3.31)	(2.64)
$T - bill$				0.00					−0.00
				(0.47)					(−0.60)
$\left(r^H - r^L\right) T - bill$				0.01					0.21
				(0.05)					(1.11)
$FXvol$					1.23				0.06
					(1.21)				(0.05)
$\left(r^H - r^L\right) FXvol$					−473*				−240
					(−1.74)				(−0.77)
TED						−0.00			0.00
						(−0.08)			(0.43)
$\left(r^H - r^L\right) TED$						−0.66			−1.34*
						(−1.14)			(−1.83)
$S\&P500$							−0.04		−0.04
							(−1.38)		(−1.20)
$\left(r^H - r^L\right) S\&P500$							4.87		2.79
							(1.11)		(0.44)
$Payrolls$								−0.03**	−0.04***
								(−2.12)	(−2.62)
$\left(r^H - r^L\right) Payrolls$								5.90	7.49**
								(1.58)	(2.09)
Intercept	0.002	0.002***	0.005***	0.005	−0.002	0.003	0.005***	0.008***	0.009
	(1.55)	(2.86)	(3.41)	(1.26)	(−0.33)	(0.98)	(3.66)	(2.97)	(1.16)
N obs	391	385	385	385	385	378	385	385	378
$AdjR^2(\%)$	−0.3	0.6	4.9	4.6	6.1	6.1	4.8	6.4	9.5

Table 5: Regressions of next-month carry trade returns on a set of independent variables. $r^H - r^L$ is the interest rate differential in the HML portfolio. CORR is the conditional correlation between HML and the US dollar. $T - bill$ is the yield of the 1-year US Treasury bill. $FXvol$ is the foreign exchange volatility proxy from Menkhoff et al. (2012). TED is the TED-spread; the difference between the 3-month USD LIBOR rate and the 3-month US T-bill rate. $Payrolls$ is Total Non-farm Payrolls of US employees, obtained from the FRED, S:t Louis Fed database. Parentheses are t-stats, computed from Newey and West (1987) standard errors, with optimal lags according to Andrews and Monahan (1992). ***, ** and * denote significance at the 1%, 5% and 10% level, respectively.

5 Predicting currency crashes

A common measure of crash risk is the skewness of returns; see, for example, Brunnermeier et al. (2009). Recall from Table 3 that the skewness of the carry trade returns are, on average, negative; −0.76. Interestingly, the skewness shows a negative relationship with CORR. When CORR is negative the average skewness is −0.47, compare to −1.15 during periods of positive CORR. But what evidence do we have that the crash risk is actually higher during periods of positive CORR? And if it is, could crash risk explain the higher returns during periods of positive CORR?

Negative skewness per se is not conclusive evidence of crash risk. For example, while Brunnermeier et al. (2009) suggests that there is a strong link between skewness and crash risk, and that excess carry trade returns are earned as a compensation for the negative skewness, Bekaert and Panayotov (2017) are able to create carry trade returns with abnormal excess returns without a significant, negative skewness. This is also the case for the dollar carry trade, introduced by Lustig et al. (2014), which has little skewness.

To further explore the properties of crashes in carry trade returns, we define *drawdowns* as the cumulative return during days of consecutive negative returns.[29] We then consider the 100 worst drawdowns as *crashes*. Figure 5 shows a bar plot of the distribution of crashes across CORR, where CORR is split into 10 bins. To capture the likelihood of observing a crash, the bars are scaled by the length of the sample period spent in each bin. We see that crashes are more likely to occur when the correlation is either strongly negative or strongly positive. In fact, the likelihood of a crash is greater during periods of negative CORR than positive CORR, which seems inconsistent with the more negative skewness.

Next, what if the average *size* of the crash is different across negative and positive CORR? The average cumulative returns of a crash during periods of negative and positive CORR are −2.1% and −1.9%, respectively. However, the differences between these averages are

[29] Drawdowns defined this way have also been studied by, for example, Daniel, Hodrick, and Lu (2016) and Sokolovski (2017).

statistically insignificant (p-value of 37%).

Further, the negative skewness may be a product of the higher average carry trade returns. We have seen in Table 3 that the average Sharpe ratio of the HML portfolio is 0.99 (0.20) in months following periods of positive (negative) CORR, while the skewness is -1.15 (-0.47). One might be tempted to conclude that the higher excess returns are earned as a compensation for the more negative skewness. However, the causality could also be reversed. Suppose that the higher excess returns are earned as a compensation for some other risk, such as the change in diversification benefits across the risk factors, as argued in this paper. Comparing periods of negative to positive CORR, the crashes are of the same size, as well as the standard deviations of returns (4.8% compared to 5.0%), only the average returns differ (1.0% compared to 5.0%). Thus, a crash observed during periods of positive CORR will be further to the left in the distribution of returns, than in the case of negative CORR. Hence, the more negative skewness of returns could be a result of the higher average returns.

As a final exercise, we investigate the predictability of crashes. Specifically, we ask the following question: given that we observe a positive carry trade return today, what is the likelihood that a crash will start tomorrow? The dependent variable in the following analysis will be an indicator variable, that takes the value 1 if there is a crash starting tomorrow, and 0 otherwise. As regressors, we consider some of the common financial variables that we have used earlier in this paper: the interest rate differential; CORR; returns on the S&P 500 equity index; the TED-spread; and the VIX.[30] I apply a standard OLS regression, which may be unsuitable when considering a probability as a dependent variable, since the fitted regression might predict a crash probability that is outside the $[0, 1]$ interval. However, in our case, we are simply looking for statistical significance between the variables, to provide us with a general sense of the predictability power. For a more economically robust model, one might consider, for example, a probit model.

[30] I intentionally keep the list of variables short in this test, mainly to save space and facilitate the interpretation of the results. Importantly, the main results are not changed even if I include all of the variables discussed earlier in the paper.

Figure 5: The distribution of carry trade crashes (of the HML portfolio) across CORR (the conditional correlation between HML and DOL). The distribution is scaled by the number of observations in each bin, to reflect the likelihood of a crash occurring.

The regression results are presented in Table 6. The sample is split into negative and positive values of CORR, as observed on day t; the day before the potential crash. Note that to achieve a relatively high power of the test, I here define negative and positive CORR to be greater than or less than zero, respectively.[31] The top (bottom) panel shows the regression results for the observations given a positive (negative) CORR. Without discussing the details of the significance of each regressor, we clearly see from the table that crashes are quite predictable when CORR is negative, while they are unpredictable when CORR is positive. In specification (vii), which includes all regressors, the adjusted R^2 when CORR is negative is 5.2%, while it is 0.0% when CORR is positive. Thus, we have uncovered yet another potential, non-competing, explanation for the high excess carry trade returns; during periods of positive CORR, investors may be compensated for the unpredictability of crashes.

[31] One potential problem with the power of these tests, is that the number of crashes (100) is small, compared to the sample size (4514). The number of crashes during positive (negative) CORR is 39 (61). However, the results are robust to instead considering the 1000 worst drawdowns as crashes.

POSITIVE CORR

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
$r^H - r^L$	2.62*		4.30				3.07
	(1.71)		(1.60)				(0.39)
$CORR$		−0.01	0.00				−0.05
		(−0.67)	(0.01)				(−0.70)
$(r^H - r^L)\,CORR$			−3.46				4.77
			(−0.74)				(0.26)
$S\&P500$				−0.46			−0.39
				(−1.20)			(−0.86)
TED					0.02		0.02
					(1.48)		(1.27)
VIX						0.00	−0.00
						(−0.14)	(−0.55)
Intercept	0.005	0.023**	0.003	0.482	−0.010	0.017	0.390
	(0.74)	(2.06)	(0.35)	(1.24)	(−0.55)	(1.05)	(0.83)
N obs	1982	1982	1982	1942	1847	1392	1358
$Adj\,R^2(\%)$	0.3	−0.0	0.3	0.0	0.2	−0.0	0.0

NEGATIVE CORR

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
$r^H - r^L$	2.19		2.36				−18.09**
	(1.14)		(0.58)				(−2.01)
$CORR$		−0.07**	−0.04				0.08
		(−2.39)	(−1.09)				(1.31)
$(r^H - r^L)\,CORR$			−6.24				−34.24*
			(−0.60)				(−1.91)
$S\&P500$				−0.77			−1.00*
				(−1.34)			(−1.85)
TED					0.06**		0.06*
					(2.12)		(1.96)
VIX						0.003**	0.002*
						(2.15)	(1.70)
Intercept	0.016*	−0.000	−0.007	0.792	−0.066*	−0.033	0.961*
	(1.90)	(−0.06)	(−0.47)	(1.37)	(−1.83)	(−1.63)	(1.83)
N obs	2532	2470	2470	2467	2337	2337	2279
$Adj\,R^2(\%)$	0.0	0.9	1.2	0.3	2.9	2.1	5.2

Table 6: In these regressions we study the predictability of carry trade crashes. The dependent variable is $\mathbb{1}^{HML}_{\{crash,t+1\}}$, which is 1 if there is a crash starting the next day, and 0 otherwise. The top (bottom) panel includes only observations when CORR is positive (negative). $r^H - r^L$ is the interest rate differential in the HML portfolio. CORR is the conditional correlation between HML and DOL. TED is the spread between the 3-month USD LIBOR rate and the 3-month US T-bill rate. VIX is the option-implied volatility of the S&P500 index. Parentheses are t-stats, computed from Newey and West (1987) standard errors, with optimal lags according to Andrews and Monahan (1992). ***, ** and * denote significance at the 1%, 5% and 10% level, respectively.

To sum up, the evidence for a higher crash risk during periods of positive CORR is mixed. We observe a more negative skewness, but we also note that the causality between skewness and higher excess returns could be reversed. When comparing crashes during periods of negative and positive CORR, we find the average size of the crashes are statistically not different. However, crashes during negative CORR are to a certain extent predictable, which might indicate that the higher excess returns during periods of positive CORR are earned as a compensation for the unpredictability of crashes.

6 Conclusions

This paper presents a conjecture suggesting that there is a link from the correlation between risk factors to future excess returns. Hence, the correlation supposedly contains predictive information about excess returns. More specifically, the conjecture states that when the correlation between two risk factors increases, the benefits of diversification for an asset having positive (or negative) loadings on *both* risk factors is reduced. This is a result of the two risk factors essentially merging into a single risk factor as they begin to correlate. The reduced benefits of diversification is directly associated with more volatile returns. Through the effects of volatility timing, the higher volatility should predict lower Sharpe ratios in the near future.

Resorting to currency markets, I find strong evidence in support of the conjecture. Lustig et al. (2011) have shown that two factors can price a large fraction of the returns in the cross-section of currencies (taken against the US dollar) sorted by interest rates. The two factors are the *dollar* factor and the *carry* factor. According to the conjecture, a higher correlation between the risk factors should, on average, be associated with a lower (higher) volatility of high-interest-rate (low-interest-rate) currency returns and lead to a higher (lower) future Sharpe ratio.

First, I document that the correlation (CORR) between the two risk factors is highly time-varying; across almost the entire $[-1,1]$ interval. Second, I find strong and significant correlations between CORR and the volatility of high interest rate currencies (-0.60) and low interest rate currencies (0.44), respectively. Even more pronounced is the correlation between

CORR and the difference between the volatilities of low and high interest rate currencies, $\sigma^L - \sigma^H$, (0.87).

Moreover, I compare periods of negative CORR (below -0.25, which corresponds to 37% of the sample period) to periods of positive CORR (above 0.25, which corresponds to 28% of the sample period). For high interest rate currencies, the average next-month Sharpe ratio when CORR is negative is 0.32, while it is 1.03 when CORR is positive. For low interest rate currencies, the corresponding Sharpe ratios are 0.28 and -0.19, respectively. This predictability extends to the standard, high-minus-low carry trade strategy (HML). Comparing periods of negative and positive CORR, the average next-month Sharpe ratio of the HML strategy is 0.20 and 0.99, respectively. These results are also found in a regression analysis, while being robust to controlling for a large set of known predictors of currency returns.

Finally, I find evidence that CORR contains information about the predictability of carry trade crashes. Specifically, these crashes are predictable - using a small set of common financial variables such as the VIX index and the TED-spread - during periods of negative CORR. When CORR is positive, however, there is no evidence of crash predictability.

To the best of my knowledge, this is the first paper to study the time-varying correlation between risk factors. The findings presented here calls for additional work in this area. In particular, I leave it to future research to investigate if the conjecture presented in this paper extends to other asset classes as well.

References

Akram, Q., D. Rime, and L. Sarno (2008): “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics*, 76, 237–253.

Aloosh, A. and G. Bekaert (2017): “Currency Factors,” Columbia Business School Research Paper 17-88.

Andrews, D. W. K. and J. C. Monahan (1992): “An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator,” *Econometrica*, 60, 953–966.

Ang, A. and G. Bekaert (2002): “International Asset Allocation With Regime Shifts,” *Review of Financial Studies*, 15, 1137–1187.

Atkeson, A. and P. J. Kehoe (2009): “On the Need for a New Approach to Analyzing Monetary Policy,” in *NBER Macroeconomics Annual 2008, Volume 23*, National Bureau of Economic Research, Inc, NBER Chapters, 389–425.

Bakshi, G. and G. Panayotov (2013): “Predictability of Currency Carry Trades and Asset Pricing Implications,” *Journal of Financial Economics*, 110, 139–163.

Bansal, R. and M. Dahlquist (2000): “The Forward Premium Puzzle: Different Tales From Developed and Emerging Economies,” *Journal of International Economics*, 51, 115 – 144.

Barroso, P. and P. Santa-Clara (2015): “Momentum has its Moments,” *Journal of Financial Economics*, 116, 111–120.

Bekaert, G. and G. Panayotov (2017): “Good Carry, Bad Carry,” Columbia Business School Research Paper 15-53.

Berg, K. A. and N. C. Mark (2018): “Global Macro Risks in Currency Excess Returns,” *Journal of Empirical Finance*, 45, 300–315.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," in *NBER Macroeconomics Annual 2008, Volume 23*, National Bureau of Economic Research, Inc, NBER Chapters, 313–347.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011): "Do Peso Problems Explain the Returns to the Carry Trade?" *The Review of Financial Studies*, 24, 853–891.

Cenedese, G., L. Sarno, and I. Tsiakas (2014): "Foreign Exchange Risk and the Predictability of Carry Trade Returns," *Journal of Banking & Finance*, 42, 302–313.

Chernov, M., J. Graveline, and I. Zviadadze (2018): "Crash Risk in Currency Returns," *Journal of Financial and Quantitative Analysis*, 53, 137–170.

Christiansen, C., A. Ranaldo, and P. Söderlind (2011): "The Time-Varying Systematic Risk of Carry Trade Strategies," *The Journal of Financial and Quantitative Analysis*, 46, 1107–1125.

Cochrane, J. H. (2017): "Macro-Finance," *Review of Finance*, 21, 945–985.

Collin-Dufresn, P., R. S. Goldstein, and J. S. Martin (2001): "The Determinants of Credit Spread Changes," *The Journal of Finance*, 56, 2177–2207.

Dahlquist, M. and H. Hasseltoft (2017): "Economic Momentum and Currency Returns," *Swedish House of Finance Research Paper*.

Daniel, K. D., R. J. Hodrick, and Z. Lu (2016): "The Carry Trade: Risks and Drawdowns," *Columbia Business School Research Paper*.

Duffee, G. R. (2011): "Information in (and not in) the Term Structure," *Review of Financial Studies*, 24, 2895–2934.

Fama, E. F. (1984): "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14, 319 – 338.

Farhi, E., S. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2015): "Crash Risk in Currency Markets," Working Paper 20948, Harvard University OpenScholar.

Farhi, E. and X. Gabaix (2016): "Rare Disasters and Exchange Rates," *The Quarterly Journal of Economics*, 131, 1–52.

Feng, G., S. Giglio, and D. Xiu (2019): "Taming the Factor Zoo: A Test of New Factors," NBER Working Papers 25481, National Bureau of Economic Research, Inc.

Freyberger, J., A. Neuhierl, and M. Weber (2018): "Dissecting Characteristics Nonparametrically," CESifo Working Paper Series 7187, CESifo Group Munich.

Hansen, L. and R. Hodrick (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88, 829–53.

Harvey, C., Y. Liu, and H. Zhu (2016): "Editor's Choice ... and the Cross-Section of Expected Returns," *Review of Financial Studies*, 29, 5–68.

Hassan, T. A. and R. C. Mano (2014): "Forward and Spot Exchange Rates in a Multi-currency World," NBER Working Papers 20294, National Bureau of Economic Research, Inc.

Hossfeld, O. and R. MacDonald (2015): "Carry Funding and Safe Haven Currencies: a Threshold Regression Approach," *Journal of International Money and Finance*, 59, 185 – 202.

Joslin, S., M. Priebsch, and K. J. Singleton (2014): "Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks," *The Journal of Finance*, 69, 1197–1233.

Karnaukh, N., A. Ranaldo, and P. Söderlind (2015): "Understanding FX Liquidity," *Review of Financial Studies*, 28, 3073–3108.

Kaul, A. and S. Sapp (2006): "Y2K Fears and Safe Haven Trading of the U.S. Dollar," *Journal of International Money and Finance*, 25, 760–779.

Levich, R. M. (2017): "CIP Then and Now: A Brief Survey of Measuring and Exploiting Deviations from Covered Interest Parity ," Prepared for the BIS Symposium: CIP - RIP?, Bank for International Settlements.

Ludvigson, S. and S. Ng (2009): "Macro Factors in Bond Risk Premia," *Review of Financial Studies*, 22, 5027–5067.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *The Review of Financial Studies*, 24, 3731.

——— (2014): "Countercyclical Currency Risk Premia," *Journal of Financial Economics*, 111, 527–553.

Lustig, H. and A. Verdelhan (2007): "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk," *American Economic Review*, 97, 89–117.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012): "Carry Trades and Global Foreign Exchange Volatility," *The Journal of Finance*, 67, 681–718.

Minsky, H. P. (1986): *Stabilizing an Unstable Economy*, A Twentieth Century Fund report, Yale University Press.

Moreira, A. and T. Muir (2017): "Volatility-Managed Portfolios," *The Journal of Finance*, 72, 1611–1644.

Mueller, P., A. Stathopoulos, and A. Vedolin (2017): "International Correlation Risk," *Journal of Financial Economics*, 126, 270 – 299.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Politis, D. N. and J. P. Romano (1994): "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89, 1303–1313.

Politis, D. N. 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.

Ranaldo, A. and P. Söderlind (2010): "Safe Haven Currencies," *Review of Finance*, 14, 385–407.

Rey, H. (2015): "Dilemma not Trilemma: The Global Financial Cycle and Monetary Policy Independence," Working Paper 21162, National Bureau of Economic Research.

RiskMetrics™ (1996): "RiskMetrics™ Technical Document," J.P. Morgan/Reuters, 4 ed.

Sokolovski, V. (2017): "Crowds, Crashes, and the Carry Trade," unpublished working paper.

Verdelhan, A. (2018): "The Share of Systematic Variation in Bilateral Exchange Rates," *The Journal of Finance*, 73, 375–418.

A Constructing daily HML returns

From monthly returns on the HML portfolio we compute *daily* returns. The idea is to mark the carry trade portfolio to market each day within the month. The forward contract is signed on day 0 (the last trading day of the previous month), with maturity on day T (the last day of the current month), and we are marking the portfolio to market on day $j \in (1, T-1)$. The cumulative market value of the portfolio is

$$V_j = \text{sign}(i^*_{0,T} - i_{0,T}) \frac{1 + i_{0,T}}{F_{0,T}} (S_j - F_{0,T}) - \left[i^*_{0,T} - i_{0,T} - \left((1 + \Delta_i)^j - 1 \right) \right], \quad (6)$$

where the daily interest rate differential, Δ_i, solves

$$(1 + \Delta_i)^T - 1 = i^*_{0,T} - i_{0,T}.$$

Note that for $j = 0$, equation (6) boils down to $V_0 = 0$. For $j = T$, we get the end-of-month value,

$$V_T = \text{sign}(i^*_{0,T} - i_{0,T}) \frac{1 + i_{0,T}}{F_{0,T}} (S_T - F_{0,T}). \quad (7)$$

The first term of equation (6) is the value of the contract as if it was the last day of the month, as in equation (7). On days in between, $j \in (1, T-1)$, the second term in equation (6) removes the interest rate differential earned between day j and T.

Finally, following Daniel et al. (2016), the excess daily return becomes

$$rx_j = \frac{V_j + (1 + i^*_{0,T})^{\frac{j}{T}}}{V_{j-1} + (1 + i^*_{0,T})^{\frac{j-1}{T}}} - (1 + i^*_{0,T})^{\frac{1}{T}}. \quad (8)$$