FEDERAL RESERVE BANK *of* ATLANTA WORKING PAPER SERIES

Foreign Exchange Predictability during the Financial Crisis: Implications for Carry Trade Profitability

Stanislav Anatolyev, Nikolay Gospodinov, Ibrahim Jamali, and Xiaochun Liu

Working Paper 2015-6
August 2015

Abstract: In this paper, we study the effectiveness of carry trade strategies during and after the financial crisis using a flexible approach to modeling currency returns. We decompose the currency returns into multiplicative sign and absolute return components, which exhibit much greater predictability than raw returns. We allow the two components to respond to currency-specific risk factors and use the joint conditional distribution of these components to obtain forecasts of future carry trade returns. Our results suggest that the decomposition model produces higher forecast and directional accuracy than any of the competing models. We show that the forecasting gains translate into economically and statistically significant (risk-adjusted) profitability when trading individual currencies or forming currency portfolios based on the predicted returns from the decomposition model.

JEL classification: F31, F37, C32, C53, G15

Key words: exchange rate forecasting, carry trade, positions of traders, return decomposition, copula, joint predictive distribution

---

Please address questions regarding content to Stanislav Anatolyev, New Economic School, 100A Novaya Street, Skolkovo, Moscow, 143026, Russia, sanatoly@nes.ru; Nikolay Gospodinov (corresponding author), Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street NE, Atlanta, GA 30309-4470, 404-498-7892, nikolay.gospodinov@atl.frb.org; Ibrahim Jamali, American University of Beirut, Olayan School of Business 449, P.O. Box 11-0236, Riad El-Solh, Beirut, Lebanon 1107 2020, ij08@aub.edu.lb; or Xiaochun Liu, University of Central Arkansas, University of Central Arkansas, EFIRM Department, 201 Donaghey Ave, COB 211, Conway, AR 72035, xliu@uca.edu.

1 Introduction

Modern international macroeconomic theory is founded on the belief that exchange rates are inherently predictable using economic fundamentals. Nonetheless, the empirical evidence is largely inconclusive or even completely unsupportive of this view. A large literature, starting with Meese and Rogoff (1983), has documented the empirical regularity that the random walk model of exchange rates is the best performing model in terms of out-of-sample forecasting. While a near-random-walk behavior in exchange rates is expected when the discount factor is near unity (Engel and West, 2005), the failure of the economic fundamentals and financial variables to exhibit any systematic predictive power is widely regarded as a major weakness of the modern international macroeconomics (Bacchetta and van Wincoop, 2006).[1] The absence of empirical evidence in support of exchange rate predictability, however, should not be construed as evidence of absence of predictability. In fact, exchange rate predictability may not have been detected due to possible hidden nonlinearities or slow-moving latent state variables whose effect passes undetected through the currency markets.[2]

The most significant departure from the lack of predictability of exchange rates has been documented in the carry trade literature. In a carry trade, an investor borrows in a low-interest currency and invests the borrowed funds in a high-yielding currency.[3] Under risk neutrality and uncovered interest rate parity (UIP), the carry trade should yield a zero average return. Despite the theoretical predictions, the carry trade has remained popular among investors and this led to widespread academic interest in the strategy's profitability.[4] The consensus emerging from the empirical research suggests that the carry trade has provided investors with statistically and economically significant positive returns over sustained periods. The documented profitability of the carry trade is consistent with the lack of empirical support for the UIP and with the voluminous

[1] Engel, Mark and West (2008) provide a more positive assessment of the predictive ability of monetary models for exchange rates.

[2] For example, Kilian and Taylor (2003) and Engel, Mark and West (2015) document some success at predicting exchange rates, especially at longer horizons, using nonlinear and factor models, respectively. Ferraro, Rogoff and Rossi (2015) also uncover a very short-term predictive relationship between commodity prices and exchange rates of commodity-dependent countries.

[3] In practice, a simple implementation of the carry trade would involve shifting the portfolio allocation from low-yielding currencies towards high-interest currencies. That is, an investor can perform the carry trade without directly borrowing or lending funds. As discussed subsequently, traders can also implement the carry trade using futures or forward contracts.

[4] For instance, by analyzing the Bank of International Settlements' triennial central bank survey of foreign exchange and derivatives market activity, Galati, Heath and McGuire (2007) note that foreign exchange turnover has increased the most for high-interest rate currencies.

literature on the forward premium puzzle (see Engel, 2015a, for a recent review of this literature).

A number of possible explanations have been advanced to account for the positive average returns of carry trade. In a classical asset pricing context, the positive average returns should reflect compensation for bearing a (possibly time-varying) risk premium and a number of recent contributions to the literature thoroughly examine the performance of common risk factors in currency pricing models (Lustig, Roussanov and Verdelhan, 2011; Burnside, Eichenbaum and Rebelo, 2011, Burnside, 2012). The findings emerging from these studies suggest that, with the exception of a global volatility risk factor, the TED spread (Bakshi and Panayotov, 2013), and term structure of interest rate variables (Ang and Chen, 2010; Lustig, Stathopoulos and Verdelhan, 2015), conventional equity and fixed-income market risk factors have demonstrated limited success in explaining the returns to the carry trade. In contrast, observable currency-specific risk factors, commonly used to assign individual currencies into portfolios, have shown more promise in predicting carry trade returns (Bakshi and Panayotov, 2013; Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b).

Another strand of the literature underscores the fact that positive carry trade returns are occasionally followed by large crash losses (Brunnermeier, Nagel and Pedersen, 2009; Berge, Jordà and Taylor, 2010; Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan, 2009; Jordà and Taylor, 2012) and entertains the importance of peso effects in driving the strategy's profitability (Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011; Burnside, 2012). The existence of crash returns and peso problems possibly reconciles the profitability of the carry trade with the predictions of UIP. A parallel literature investigates the importance of 'limits to arbitrage and hedging' in explaining the carry trade's profitability. In line with a number of studies invoking the limits to arbitrage and hedging in various asset markets,[5] the basic premise is to ascribe the existence of persistently positive average carry trade profits to liquidity frictions (Mancini, Ranaldo and Wrampelmeyer, 2013; among others) as well as to margin, short-selling or leverage constraints that prevent arbitrageurs and speculators from completely exploiting the carry trade's profitability (Brunnermeier, Nagel and Pedersen, 2009; Mancini-Griffoli and Ranaldo, 2011; among others).

In this paper, we adopt a statistical approach to uncovering and exploiting potential predictability in carry trade returns during and after the recent U.S. financial crisis. More specifically, we capitalize on the method proposed by Anatolyev and Gospodinov (2010) to decompose currency

[5] An excellent review of this literature is provided in Gromb and Vayanos (2010).

returns into two multiplicative components (sign and absolute returns) that individually exhibit much greater predictability than raw returns. We then model the joint conditional distribution of these components and use it to produce forecasts of future returns. We allow the two components to respond to currency-specific risk factors such as speculative pressure. This method of incorporating any implicit nonlinearities in a flexible, indirect fashion is motivated by: (i) the limited success of linear asset pricing models in explaining carry trade returns (Burnside, Eichenbaum and Rebelo, 2011), (ii) prior empirical evidence pointing to a statistically and economically significant element of nonlinear out-of-sample predictability in foreign exchange markets especially at long horizons (Kilian and Taylor, 2003; Engel, 2015b; Vlachev, 2015) and (iii) the profitability of trading based on the decomposition model of Anatolyev and Gospodinov (2010) for equity returns. By virtue of allowing for nonlinearities in currency returns, our paper also relates to an existing line of research that exploits non-linearities in the relationship between exchange rate returns and interest rate differentials (Engel, 2015b; Vlachev, 2015) or in carry trade returns (Daniel, Hodrick and Lu, 2014) for predictive purposes.

Several interesting results emerge from our analysis. First, the decomposition model exhibits substantial directional accuracy in predicting carry trade returns during the recent financial crisis. This is in sharp contrast with the pure carry trade strategies that recorded losses during the financial crisis (see also Lustig and Verdelhan, 2011). Second, the out-of-sample forecasting gains of the decomposition model (relative to the naïve historical mean and linear prediction models) translate into economically and statistically highly significant profitability. More specifically, trading individual currency forward contracts or forming portfolios based on the sign of the predicted return from the decomposition model generates larger (risk-adjusted) profits than any of the competing models.

We contribute to the existing literature on the carry trade along theoretical and empirical lines. From a modeling perspective, this paper offers a new approach to modeling and forecasting currency returns. We view the uncovered nonlinear predictability as a possible explanation of the limited success of linear asset pricing models in the context of currency markets. Note that the carry trade returns consist of two parts – future currency returns and interest rate differential – and while the pure carry trade strategies exploit only the differential in interest rates, both of which are near the zero lower bound over this period, we employ a model-based carry trade strategy that capitalizes on the predictability of currency returns. As we show in the paper, our decomposition model uncovers a large degree of predictability that generates highly profitable carry trade strategies. Finally, as a by-product of our analysis, we update the empirical performance of the commonly used carry trade strategies until the end of 2013. Overall, the post-2007 carry trade has been largely unsuccessful in replicating its profitability prior to the financial crisis.

The rest of the paper proceeds as follows. Section 2 outlines our asset pricing framework and motivates the use of nonlinear models in predicting carry trade returns. Section 3 provides a detailed discussion of the specific decomposition model that we employ. The data and variables employed in the empirical analysis are described in Section 4. Our empirical findings as well as the trading strategies we consider to assess the profitability of the decomposition model are discussed in Section 5. Section 6 offers some concluding remarks.

2 Pricing Kernel and Predictability of Currency Returns

To motivate our modeling and estimation framework, we adopt the approach by Constantinides (1992) and Backus, Foresi and Telmer (2001) to bond and currency pricing. We start with the fundamental pricing equation (Cochrane, 2001) where asset prices are obtained by 'discounting' future payoffs using a pricing kernel so that the expected present value of the payoff is equal to the current price. Let $\mathcal{F}_t$ denote the information set at time t. Associated with $\mathcal{F}$ is the space L^2 of all random variables with finite second moments that are in the information set $\mathcal{F}$. This space represents the collection of hypothetical claims that could be traded. Let R_t denote the domestic-currency denominated gross return on the asset at time t and $m_t \in L^2$ be an admissible pricing kernel that prices the asset correctly, i.e.,

$$E[R_{t+1}m_{t+1}|\mathcal{F}_t] = 1. \tag{1}$$

Following Backus, Foresi and Telmer (2001), the pricing kernel m^*_{t+1} associated with the foreign-currency denominated return $R^*_{t+1} = S_t R_{t+1}/S_{t+1}$, where S_t denotes the exchange rate between the domestic and foreign currency, satisfies the restriction

$$E[R^*_{t+1}m^*_{t+1}|\mathcal{F}_t] = E[R^*_{t+1}(S_{t+1}/S_t)m_{t+1}|\mathcal{F}_t] = 1, \tag{2}$$

where the second equality is obtained by substituting for $R_{t+1} = R^*_{t+1}(S_{t+1}/S_t)$ into (1).

Consider now the time $t+1$ payoff $F_t - S_{t+1}$, where F_t is the forward exchange rate, and since entering a forward contract has zero cost, we have the relationship

$$E[(F_t - S_{t+1})m_{t+1}|\mathcal{F}_t] = 0 \tag{3}$$

or equivalently

$$(F_t/S_t)E[m_{t+1}|\mathcal{F}_t] = E[m^*_{t+1}|\mathcal{F}_t] \tag{4}$$

by dividing (3) by S_t and invoking (2). Taking logarithms of (4) and rearranging, we have

$$s_t - f_t = \ln(E[m_{t+1}|\mathcal{F}_t]) - \ln(E[m^*_{t+1}|\mathcal{F}_t]),$$

where $s_t = \ln(S_t)$ and $f_t = \ln(F_t)$. Following Backus, Foresi and Telmer (2001), $\ln(E[m_{t+1}|\mathcal{F}_t])$ and $\ln(E[m^*_{t+1}|\mathcal{F}_t])$ can be expanded in terms of conditional cumulants which gives

$$s_t - f_t = \sum_{i=1}^{\infty} \frac{\kappa_{it+1} - \kappa^*_{it+1}}{i!}, \tag{5}$$

where κ^*_{it+1} and κ_{it+1} are the i-th cumulants (conditional on the information set $\mathcal{F}_t$) of m^*_{t+1} and m_{t+1}, respectively.

The main interest of this paper lies in forecasting carry trade (excess) returns for currency j defined as

$$er_{jt+1} = s_{jt+1} - f_{jt} = (s_{jt+1} - s_{jt}) + (s_{jt} - f_{jt}). \tag{6}$$

Note that, under the covered interest parity (CIP), the futures spread is equivalent to the interest rate differential $i_t - i^*_t$, where i_t and i^*_t denote the nominal interest rates in the domestic and foreign currency, respectively.[6] If both UIP and CIP hold, average excess carry trade returns are zero. Also note that the second component in the above equation is the expansion in terms of cumulants of m^*_{t+1} and m_{t+1} in (5). This suggests that the carry returns are a (possibly nonlinear) function of high-order moments of the state variables describing the dynamics of m^*_{t+1} and m_{t+1}. The signal contained in $(s_{jt} - f_{jt})$, however, is swamped by the much more volatile currency returns $(s_{jt+1} - s_{jt})$ (see, for example, Gospodinov, 2009) which presents challenges to identifying and extracting the predictable part of the carry trade returns.

[6] In fact, a common method of testing CIP is to employ the following regression:

$$f_t - s_t = \alpha + \beta(i_t - i^*_t) + error.$$

Under CIP and in the absence of transaction costs, the intercept and slope coefficients in the previous regression should be equal to zero and one, respectively.

There are several ways to proceed with developing a forecasting model of carry trade returns. The first approach recognizes that the pricing kernel is the intertemporal marginal rate of substitution and derives its induced time series properties but operationalizing this imposes strong assumptions on the economy and the preferences. Alternatively, one could follow Constantinides (1992) and Backus, Foresi and Telmer (2001) and parameterize and explore directly the time series properties of the pricing kernels m^*_{t+1} and m_{t+1}. This approach, however, also requires specifying a functional form for the pricing kernels and identifying the state variables driving their dynamic behavior. In this paper, we follow a flexible, albeit less structural, approach of modeling the conditional distribution of spot exchange rates. This method can capture inherent nonlinearities in the exchange rate dynamics that allow us to identify factors (affecting, for example, both m^*_{t+1} and m_{t+1}) which would pass through largely undetected in the traditional models (for a discussion of this, see Fisher and Gilles, 2000). This approach is also consistent with the recent findings by Bakshi and Panayotov (2013) of a predictability of carry trade payoffs by latent risk factors.

3 Decomposition Method

The decomposition method is based on exploiting predictability of multiplicative components in order to tease out nonlinear predictability from a series that is linearly unpredictable. Let us first illustrate this possibility with a simple example. Suppose that two processes, ε_t and x_t, are symmetrically distributed and serially and mutually independent at all leads and lags. Construct a_t to be a zero-mean AR(1) process

$$a_t = \rho a_{t-1} + \varepsilon_t.$$

Let a_t be observable at time t. This series is mean predictable from the past if $\rho \neq 0$, and the best predictor is ρa_{t-1}. Next, let x_t be observable at time t, and set the binary variable b_t to be the sign of the last period x_t:

$$b_t = \text{sign}(x_{t-1}).$$

This series is perfectly predictable from the past of x_t. Note that the series a_t and b_t have mean zero and are mutually independent at all lags and leads. Now let us construct the 'returns' series

$$r_t = a_t b_t$$

which has the properties that $E\left[r_t\right] = E\left[a_t\right] E\left[b_t\right] = 0$, and, for $j > 0$,

$$E\left[r_t r_{t-j}\right] = E\left[a_t b_t a_{t-j} b_{t-j}\right] = E\left[a_t a_{t-j}\right] E\left[b_t\right] E\left[b_{t-j}\right] = 0.$$

In other words, the 'returns' have mean zero and are serially uncorrelated, i.e., linearly unpredictable from their own past. Moreover, r_t is also linearly unpredictable from a_{t-j} for any $j > 0$, from b_{t-j} for any $j \geq 0$, and from x_{t-j} for any $j > 0$:

$$\begin{aligned} E\left[r_t a_{t-j}\right] &= E\left[a_t b_t a_{t-j}\right] = E\left[a_t a_{t-j}\right] E\left[b_t\right] = 0, \\ E\left[r_t b_{t-j}\right] &= E\left[a_t b_t b_{t-j}\right] = E\left[a_t\right] E\left[b_t b_{t-j}\right] = 0, \\ E\left[r_t x_{t-j}\right] &= E\left[a_t b_t x_{t-j}\right] = E\left[a_t\right] E\left[b_t x_{t-j}\right] = 0. \end{aligned}$$

This shows that the 'returns' r_{t+1} are linearly unpredictable from all observable histories. However, r_{t+1} is nonlinearly predictable from the observable past since

$$E\left[r_{t+1}|\mathcal{F}_t\right] = \rho a_t \operatorname{sign}(x_t).$$

The optimal nonlinear predictor is proportional to ρ, the degree of persistence in a_t, while the optimal linear predictor is zero irrespective of it. Note that a similar result would also hold if the best predictor of a_t was nonlinear in a_{t-1}. This example provides some intuition why the decomposition model is potentially able to detect certain forms of hidden predictability in the data; namely, when one (or both) of the multiplicative components exhibits persistence.

Now we describe the decomposition approach of Anatolyev and Gospodinov (2010) whose key insight is based on the return decomposition

$$r_t = |r_t| \operatorname{sign}(r_t) = |r_t|(2\mathbb{I}[r_t > 0] - 1),$$

where r_t are asset returns and $\mathbb{I}[\cdot]$ is the indicator function. The method then proceeds with the joint dynamic modeling of the two multiplicative components – 'volatility' $|r_t|$ and (a linear transformation of) 'direction' $\mathbb{I}[r_t > 0]$.

As in the above example, the driving force behind the predictive ability of the decomposition model is the predictability in the two components, $|r_t|$ and $\mathbb{I}[r_t > 0]$, that has been documented in previous studies. Consider first the model specification for absolute returns. Since $|r_t|$ is a positively valued variable, the dynamics of absolute returns is specified using the multiplicative error model (Engle, 2002)

$$|r_t| = \psi_t \eta_t,$$

where $\psi_t = E(|r_t||\mathcal{F}_{t-1})$ and η_t is a positive multiplicative error. This error has unit conditional mean and conditional distribution $\mathcal{D}$ which may be flexibly specified as a scaled Weibull distribution with a shape parameter ς. A convenient dynamic specification for the conditional expectation ψ_t is the logarithmic autoregressive model

$$\ln \psi_t = \omega_v + \beta_v \ln \psi_{t-1} + \gamma_v |r_{t-1}| + \rho_v \mathbb{I}\left[r_{t-1} > 0\right] + x'_{t-1}\delta_v. \tag{7}$$

This volatility model allows for persistence, regime-switching, a direct effect of last-period absolute return, and possible effects of other predictors x_{t-1}.

In the direction model, the conditional 'success probability' $p_t = \Pr\{r_t > 0|\mathcal{F}_{t-1}\}$ is parameterized as a dynamic logit model

$$p_t = \frac{\exp(\theta_t)}{1 + \exp(\theta_t)}$$

with

$$\theta_t = \omega_d + \phi_d \mathbb{I}\left[r_{t-1} > 0\right] + y'_{t-1}\delta_d, \tag{8}$$

allowing for regime-switching and possible effects of other predictors y_{t-1} that may be different from x_{t-1}. A direct persistence effect is not included because directional persistence is much lower than volatility persistence.

To describe the joint distribution of $R_t \equiv (|r_t|, \mathbb{I}\left[r_t > 0\right])'$, the copula approach is used. The conditional marginals of R_t are $(\mathcal{D}(\psi_t), \mathcal{B}(p_t))'$, where $\mathcal{B}(p_t)$ denotes the Bernoulli distribution with probability mass function $p_t^v (1 - p_t)^{1-v}$, $v \in \{0, 1\}$. Let $\mathcal{C}(w_1, w_2)$ denote a copula function on $[0, 1] \times [0, 1]$. Anatolyev and Gospodinov (2010) derive that, conditional on $\mathcal{F}_{t-1}$, the joint density/mass of R_t is given by

$$f_{R_t}(u, v) = f^{\mathcal{D}}(u|\psi_t)\varrho_t\left(F^{\mathcal{D}}(u|\psi_t)\right)^v \left(1 - \varrho_t\left(F^{\mathcal{D}}(u|\psi_t)\right)\right)^{1-v}, \tag{9}$$

where $f^{\mathcal{D}}(\cdot)$ and $F^{\mathcal{D}}(\cdot)$ are density and CDF of $\mathcal{D}$, and $\varrho_t(z) = 1 - \partial\mathcal{C}(z, 1 - p_t)/\partial w_1$.

Denote by α_t a time-varying copula parameter that captures the dependence between the two marginals. We consider the Frank copula[7] which has the form

$$\mathcal{C}(w_1, w_2) = -\frac{1}{\alpha} \ln\left(1 + \frac{[\exp(-\alpha w_1) - 1][\exp(-\alpha w_2) - 1]}{\exp(-\alpha) - 1}\right),$$

where $\alpha < 0$ ($\alpha > 0$) implies negative (positive) dependence while $\alpha = 0$ implies independence. For this copula, Anatolyev and Gospodinov (2010) deduce that

$$\varrho_t(z) = \left(1 - \frac{1 - \exp(-\alpha(1 - p_t))}{1 - \exp(\alpha p_t)} \exp(\alpha(1 - z))\right)^{-1}.$$

[7] The subsequent results are very similar with double Clayton copula and Farlie–Gumbel–Morgenstern copula and we omit the discussion of these two copula choices.

To allow for greater flexibility, we adopt a time-varying copula specification, where the dependence (copula) parameter is assumed to follow the dynamic process (see Manner and Reznikova, 2012)

$$\alpha_t = \lambda_0 + \lambda_1 \alpha_{t-1} + \lambda_2 |r_{t-1}| (1 - \mathbb{I}[r_{t-1} > 0])$$

with $|\lambda_1| < 1$ and $\alpha_t \in (-\infty, +\infty) \backslash 0$. In this specification, the forcing variable $|r_{t-1}|(1 - \mathbb{I}[r_{t-1} > 0])$ is equal to $|r_{t-1}|$ when r_{t-1} is negative and zero otherwise. As $\alpha_t \to 0$, the Frank copula approaches the independence copula and $\varrho_t(z) = p_t$.

The parameter vector $(\omega_v, \beta_v, \gamma_v, \rho_v, \delta_v', \varsigma, \omega_d, \phi_d, \delta_d', \lambda_0, \lambda_1, \lambda_2)'$ is estimated by maximum likelihood. From (9), the sample log-likelihood function to be maximized is given by

$$\sum_{t=1}^{T} \left\{ \mathbb{I}[r_t > 0] \ln \varrho_t \left(F^{\mathcal{D}}(|r_t| | \psi_t) \right) + (1 - \mathbb{I}[r_t > 0]) \ln \left(1 - \varrho_t \left(F^{\mathcal{D}}(|r_t| | \psi_t) \right) \right) + \ln f^{\mathcal{D}}(|r_t| | \psi_t) \right\}.$$

As our interest lies in the mean prediction of returns, we use the fact that

$$E(r_{t+1} | \mathcal{F}_t) = 2E(|r_{t+1}| \mathbb{I}[r_{t+1} > 0] | \mathcal{F}_t) - E(|r_{t+1}| | \mathcal{F}_t)$$

to construct the prediction of return at time $t+1$ as

$$\hat{r}_{t+1} = 2\hat{\xi}_{t+1} - \hat{\psi}_{t+1}, \tag{10}$$

where $\hat{\psi}_{t+1}$ and $\hat{\xi}_{t+1}$ are feasible analogs of ψ_{t+1} and ξ_{t+1}, and $\xi_{t+1} = E(|r_{t+1}| \mathbb{I}[r_{t+1} > 0] | \mathcal{F}_t)$ is the conditional expected cross-product of $|r_{t+1}|$ and $\mathbb{I}[r_{t+1} > 0]$. As shown in Anatolyev and Gospodinov (2010), ξ_{t+1} can be computed as

$$\xi_{t+1} = \int_0^1 Q^{\mathcal{D}}(z) \varrho_{t+1}(z) dz, \tag{11}$$

where $Q^{\mathcal{D}}(z)$ is a quantile function of the distribution $\mathcal{D}$. The feasible version $\hat{\xi}_{t+1}$ is obtained numerically (via the Gauss–Chebyshev quadrature) by evaluating the integral (11) using fitted $Q^{\mathcal{D}}(z)$ and fitted $\varrho_{t+1}(z)$.

This decomposition approach could be used directly to forecast the carry trade returns in (6). However, in our empirical analysis (Section 5), we choose to obtain the forecast of $s_{jt+1} - s_{jt}$ first and then construct the predicted carry trade returns by adding the (observed at time t) spread $s_{jt} - f_{jt}$. This modeling choice is motivated by two reasons. First, while $s_{jt+1} - s_{jt}$ is mean-zero and serially uncorrelated (see Table 1 below), the excess returns $s_{jt+1} - f_{jt}$ have a possibly non-zero mean and exhibit some serial correlation that need to be dealt with explicitly. Second, this allows us to relate our results to the large literature on (in-sample and out-of-sample) forecasting of currency returns and gain a better understanding of the source of the statistical and economic improvements of the carry trade strategies.

4 Data Description

Exchange rate data for the spot (S_t) and three-month US dollar forward (F_t) rates, expressed in US dollars per unit of the foreign currency, are obtained from Datastream. The currency returns for the j-th exchange rate are constructed as $r_{jt+1} = \ln(S_{jt+1}) - \ln(S_{jt})$ while the corresponding carry trade returns are constructed as $er_{jt+1} = \ln(S_{jt+1}) - \ln(F_{jt}) = r_{jt+1} - spread_{jt}$, where $spread_{jt} = \ln(F_{jt}) - \ln(S_{jt})$. We consider the exchange rates of four currencies – British pound (GBP), Canadian dollar (CAD), Japanese yen (JPY), and the Euro (EUR) – against the US dollar. This specific cross-section of currencies and three-month horizon for the forwards are dictated by the availability and characteristics of Commitment of Traders (COT) data described below.[8] Our empirical analysis is conducted at the weekly frequency. The sample for the first three exchange rates spans the period from the first week of November 1992 until the last week of December 2013 while the sample for the Euro is from the first week of January 1999 until the end of December 2013. The variables r_{jt+1}, er_{jt+1} and $spread_{jt}$ are appropriately transformed to represent weekly percentage returns.

In addition to $spread_{jt}$, we use speculative (or hedging) pressure as an external predictor. The speculative (hedging) pressure, constructed from data on commitment of traders, has been used extensively as a transaction-related predictor of future commodity price (de Roon, Nijman and Veld, 2000; Dewally, Ederington and Fernando, 2013; Szymanowska, de Roon, Nijman and van den Goorbergh, 2014) and currency (Klitgaard and Weir, 2004; Tornell and Yuan; 2012) movements. COT data for each currency is provided by the U.S. Commodity Futures Trading Commission (CFTC) and contains short and long futures positions of three categories of market participants: commercial, non-commercial and non-reportable. Commercial positions refer to the number of contracts held by commercial institutions (the CFTC classifies those as hedging), while the non-commercial category includes the number of contracts held by non-commercial users (classified by the CFTC

[8] For instance, the Australian and New Zealand dollars are likely to be "target currencies" in a carry trade due to the high interest rates in Australia and New Zealand. The incomplete COT data for these currencies prevents us from including them in our cross-section. We also do not include the Swiss Franc in our analysis due to its peg to the Euro during most of the post-crisis period.

as non-hedging and often interpreted as speculative; see Bessembinder, 1992). Notwithstanding some caveats related to trader classification, we follow the common practice in the literature by considering commercial users of futures contracts as hedgers and non-commercial users of foreign currency futures as speculators.[9]

Following Dewally, Ederington and Fernando (2013), we construct a measure of speculative pressure for week t and currency j as

$$sp_{jt} = \frac{\text{\# of long non-commercial positions } - \text{ \# of short non-commercial positions}}{\text{\# of long non-commercial positions } + \text{ \# of short non-commercial positions}}.$$

We also construct a measure of hedging pressure using data on commercial market participants (hedgers) instead of non-commercial users of futures contracts (speculators) as in the speculative pressure variable above. While over-the-counter (OTC) forward markets are considerably more liquid than currency futures markets, existing studies (Brunnermeir, Nagel and Pedersen, 2009; Galati, Heath and McGuire, 2007) underline the usefulness of the positions of non-commercial futures traders as a proxy of carry trade activity.[10] Therefore, the speculative and hedging pressure, which proved to be good predictors of spot exchange rate returns (Klitgaard and Weir, 2004; Tornell and Yuan; 2012), are also likely to be good predictors of carry trade returns.

The exchange rates are sampled as the daily close on Tuesday in order to match the weekly COT data which report the positions of traders every Tuesday over our sample period. Table 1 provides summary statistics of the exchange rate data and the currency-specific predictors. As documented elsewhere in the literature, the exchange rate returns appear to be serially uncorrelated while the spread is highly persistent with variability which is only a small fraction of the variability of the currency returns (Gospodinov, 2009). Due to the low variability of the spread, the carry trade returns, plotted in Figure 1, largely inherit the properties of the exchange rate returns although with slightly more pronounced non-zero mean and higher serial correlation which is not readily picked up by the first-order autocorrelation coefficient. Speculative pressure also exhibits high persistence.

[9] The literature indicates, for example, that commercial users of futures contracts tend to engage in speculative behavior (see, for example, the discussion in Galati, Heath and McGuire, 2007; among others). We should note that the exact classification of traders is not essential for the purposes of our empirical analysis.

[10] Galati, Heath and McGuire (2007) note that indicators of carry trade activity can be alternatively obtained from the Bank of International Settlements' (BIS) triennial survey of foreign exchange markets or from the BIS international banking statistics database. The use of measures of carry trade activity from these databases is not without its own shortcomings, however. According to Galati, Heath and McGuire (2007), "turnover data from OTC markets support the conclusions reached using futures data".

5 Empirical Results

5.1 Currency Return Predictability

This section presents in-sample estimation and out-of-sample forecasts for the exchange rate and carry trade returns.

5.1.1 In-Sample Estimation Results

The decomposition model is used to model the dynamic behavior of exchange rate returns r_t with marginals given by

$$\ln \psi_t = \omega_v + \beta_v \ln \psi_{t-1} + \gamma_v |r_{t-1}| + \rho_v \mathbb{I}\left[r_{t-1} > 0\right] + \delta_{1v} spread_{t-1} + \delta_{2v} sp_{t-1}$$

$$\theta_t = \omega_d + \phi_d \mathbb{I}\left[r_{t-1} > 0\right] + \delta_{1d} spread_{t-1} + \delta_{2d} sp_{t-1},$$

and a joint distribution based on the time-varying Frank copula

$$\mathcal{C}(|r_t|, \mathbb{I}\left[r_t > 0\right]) = -\frac{1}{\alpha_t} \ln \left(1 + \frac{[\exp(-\alpha_t |r_t|) - 1][\exp(-\alpha_t \mathbb{I}\left[r_t > 0\right]) - 1]}{\exp(-\alpha_t) - 1} \right),$$

where $\alpha_t = \lambda_0 + \lambda_1 \alpha_{t-1} + \lambda_2 |r_{t-1}|(1 - \mathbb{I}\left[r_{t-1} > 0\right])$. As specified above, the models for GBP, CAD, and JPY are estimated with weekly data from 10/13/1992 to 12/31/2013 (1109 observations) while the in-sample estimation for EUR uses weekly observations from 1/5/1999 to 12/31/2013 (782 observations).

Table 2 reports the estimated parameters and their corresponding standard errors. A few interesting observations emerge from these results. All of the exchange rates are characterized by strong persistence in volatility and dependence on the absolute value of lagged returns. The spread appears to be a useful predictor for both the volatility and the direction models. Speculative pressure[11] exhibits some predictive power in the direction model for GBP and EUR. Finally, the persistence parameter in the time-varying copula process is large and highly significant.

Even though Table 2 shows that some parameters in the decomposition model may be insignificant for one or more currencies, we do not undertake any model selection for individual currencies for the following reasons. First, some predictors may only exhibit occasional and short-lived predictive power (and hence are insignificant in the whole sample) but they might prove to be useful

[11] The specification with hedging pressure in addition to or in place of speculative pressure yields similar results and is not presented to conserve space. We also experimented with the S&P500 option-implied volatility index (VIX) as a global risk factor but did not find it to be a useful predictor in our model.

when needed the most (during the financial crisis or periods of liquidity shortage, for example). Second, keeping all predictors the same facilitates the comparison and the interpretation of the out-of-sample results across models and currencies. We acknowledge, however, that the forecasting performance can be further improved by a careful model selection based on information up to the time when the forecast is constructed.

5.1.2 Out-of-Sample Forecasting Results

This section reports results for one-step ahead out-of-sample forecasts for the realized carry trade returns, defined as

$$\hat{\mu}^i_{jt+1} = \text{sign}(\widehat{er}^i_{jt+1}) er_{jt+1}, \tag{12}$$

where $\widehat{er}^i_{jt+1} = \hat{r}^i_{jt+1} + (s_{jt} - f_{jt})$ and $\hat{r}^i_{jt+1}$ denotes the i-th model forecast of r_{jt+1} for currency j. The forecast period is from the second week of July 2007 until the last week of December 2013 and covers the recent financial crisis and the post-crisis periods. This period is selected for two reasons. First, it is driven by data availability, especially for the Euro whose positions of traders data start in the beginning of 1999. Second, one of the main objectives of this paper is to assess the recent carry trade profitability during and after the financial crisis. Thus, we set the length of the rolling window in the out-of-sample forecasts to be $R = 600$ for GBP, CAD, and JPY and $R = 442$ for EUR. The predictability of carry trade returns are evaluated over the following forecasting subsamples: (i) financial crisis period: 7/10/2007-12/28/2010; (ii) post-crisis period 1/4/2011-12/31/2013; and (iii) combined crisis and post-crisis period 7/10/2007-12/31/2013.[12]

We consider two competing models: a linear predictive model and the decomposition model. The linear model uses the same economic predictors as the decomposition model (i.e., *spread* and *sp*). The benchmark model is the random walk model. Table 3 presents the test proposed by Giacomini and White (GW, 2006) based on a quadratic function of the difference in realized returns $\hat{\mu}^i_{jt+1} - \tilde{\mu}_{jt+1}$, where $\tilde{\mu}_{jt+1}$ is the realized return of the benchmark (random walk) model. Positive values of the GW statistic indicate that the competing model outperforms the benchmark model, and the difference is significant if the test statistic is larger than the critical value from the standard normal distribution. While the decomposition model dominates RW and linear model for all currencies during the financial crisis and the combined period, the improvements of the

[12] While it would have been more appropriate to time the end of the financial crisis in the middle or in the second half of 2009 (according to NBER, the U.S. recession started in December 2007 and ended in June 2009), we wanted to split the number of observations more equally between the two subsamples.

decomposition model over RW are statistically significant at the 5% level only for CAD.

Next, we present evidence for the directional performance of the models which is exploited in the next section. In particular, we compute the statistics for the area under the receiver operating characteristic curve (AUC) and the corresponding returns-weighted AUC statistic (AUC*) suggested by Jordà and Taylor (2012).[13] Note that higher values for AUC and AUC* indicate a better directional forecast performance. As evident from Table 4, the directional forecasts of the decomposition model outperform those based on the linear and random walk models across all currencies. We next evaluate the economic significance of the accuracy of these directional forecasts in the context of trading strategies.

5.2 Trading Strategies

5.2.1 Active Trading Strategy for Individual Currencies

The trading strategy discussed in this section is based on the sign of predicted carry trade returns from one of the econometric models. The strategy, conducted from the perspective of a U.S. investor, consists of going long in the foreign currency when the sign of the predicted carry trade return from one of the econometric models is positive. Conversely, when the predicted sign of the carry trade return is negative, the trader shorts the foreign currency. As argued in Burnside, Eichenbaum and Rebelo (2011), the payoff from conducting the carry trade in the forward market is proportional, under CIP, to trading directly in the spot market (i.e., based on interest rate differentials). The payoff from this trading strategy for currency j and model i is given in (12). It is important to stress that the predicted returns $\hat{r}^i_{jt+1}$ are genuine out-of-sample forecasts (as in Section 5.1.2) that utilize information only up to time t.

Following the literature (Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2012b; Bakshi and Panayotov, 2013), we employ the bid-ask prices[14] to account for the transaction costs that investors would incur when implementing the trades. Let s^a_{jt} and s^b_{jt} denote the ask and bid (offered) prices, respectively, for currency j. The spot price, s_{jt}, is the midpoint of the bid and ask quotes. The trading strategy considered is as follows: the investor

[13] More specifically, let $d_{t+1} = \text{sign}(er_{t+1}) \in \{-1, 1\}$, $N_p = \sum_{d_{t+1}=1} 1$, $N_n = \sum_{d_{t+1}=-1} 1$, $\phi_k = \{\widehat{er}_{t+1} | d_{t+1} = 1\}$ for $k = 1, ..., N_p$, $\eta_l = \{\widehat{er}_{t+1} | d_{t+1} = -1\}$ for $l = 1, ..., N_n$, $B_{\max} = \sum_{d_{t+1}=1} er_{t+1}$ and $C_{\max} = \sum_{d_{t+1}=-1} |er_{t+1}|$. Then, AUC $= \frac{1}{N_p N_n} \sum_{k=1}^{Np} \sum_{l=1}^{N_n} \mathbb{I}[\phi_k > \eta_l]$ and AUC$^* = \sum_{k=1}^{Np} \frac{\phi_k}{B_{\max}} \sum_{l=1}^{N_n} \frac{\eta_l}{C_{\max}} \mathbb{I}[\phi_k > \eta_l]$. See Jordà and Taylor (2012) for more details.

[14] The bid and ask forward and spot data are collected by Thomson Reuters.

goes long in the foreign currency if the predicted carry trade return from one of the econometric models is positive and exceeds the transaction costs observed at time t; that is, if $\text{sign}(\widehat{er}^i_{jt+1}) > 0$ and $\widehat{er}^i_{jt+1} > s^a_{jt} - s_{jt}$. Conversely, the investor shorts the foreign currency if the predicted carry trade return from model i is negative and exceeds, in absolute value, the transaction costs observed at time t (if $\text{sign}(\widehat{er}^i_{jt+1}) < 0$ and $|\widehat{er}^i_{jt+1}| > s_{jt} - s^b_{jt}$). If the predicted carry trade return is lower than the transaction costs, no position is taken (i.e., the investor stays outside the market). The net-of-transaction costs returns of a long position are given by $\text{sign}(\widehat{er}^i_{jt+1})er_{jt+1} - (s^a_{jt+1} - s_{jt+1})$ while the net of transaction costs returns to a short position are $\text{sign}(\widehat{er}^i_{jt+1})er_{jt+1} - (s_{jt+1} - s^b_{jt+1})$.

We consider a benchmark strategy that uses the sign of the forward premium as an *ex ante* predictor of the sign of the carry trade returns. This strategy consists of selling currencies that are at a forward premium and buying currencies that are at a forward discount (Burnside, Eichenbaum and Rebelo, 2011). The predicted carry trade return of the benchmark strategy is $\widehat{er}^0_{jt+1} = -spread_{jt}$ which is equivalent to trading based on a random walk model for the spot exchange rate.[15]

The value of the initial investment is set equal to \$100. The value of the portfolio is re-calculated and re-invested every period. In addition to the models considered in the previous section, we also include strategies based on forecasts of currency returns obtained from their historical average up to time t. The results for the net-of-transaction-cost payoffs and the Sharpe ratios are presented in Tables 5 and 6, respectively.

Table 5 shows that the decomposition model outperforms the remaining models with some of the differences in payoffs (for CAD in the crisis period, for example) being quite large. Similar results hold for the risk-adjusted returns (Sharpe ratios) in Table 6.[16] We also considered another trading strategy in which the investor longs the foreign currency if $\text{sign}(\widehat{er}^i_{jt+1}) > 0$ and shorts the foreign currency if $\text{sign}(\widehat{er}^i_{jt+1}) < 0$. The investor thereby takes a position and incurs transaction costs every period. The results are very similar (with slightly lower payoffs in most cases) to the results reported in Table 5 for the trading strategy with the option to stay outside the market.

[15] Starting from $er_{jt+1} = s_{jt+1} - f_{jt} = (s_{jt+1} - s_{jt}) - (f_{jt} - s_{jt})$, a random walk (with no drift) for s_{jt+1} would imply that $\widehat{er}^0_{jt+1} = -spread_{jt}$.

[16] To evaluate the statistical significance of the computed Sharpe ratios, we constructed 90% bootstrap confidence intervals based on the moving block bootstrap with a blocksize equal to 8. Only the decomposition-based Sharpe ratios for the CAD during the financial crisis and the combined period are statistically larger than zero.

5.2.2 Currency Portfolios

In addition to investigating the profitability of trading individual currencies based on the sign of the econometric forecast, we examine the directional profitability of the econometric models from a portfolio perspective. More specifically, an equally-weighted portfolio of the four currencies (GBP, CAD, JPY, and EUR) is constructed by averaging the net-of-transaction costs realized returns from one of the econometric models across the four currencies. The excess return of the equally weighted portfolio therefore measures the returns accruing to an investor who trades the four currencies based on the predicted sign of the carry trade returns from one of the econometric models.

Following the literature (Brunnermeier, Nagel and Pedersen, 2009; Ang and Chen, 2010; Lustig, Roussanov and Verdelhan, 2011; and Accominotti and Chambers, 2014), we construct benchmark long-short carry trade portfolios on the basis of interest rate differentials (as proxied for, under CIP, by the $spread_{jt}$). The first portfolio, referred to as 'signed spread', consists of going long in a currency with a negative spread and going short in a currency whose spread is positive. The returns of this portfolio are equivalent to trading the four currencies based on a random walk model for each of the four spot exchange rates. In line with prior research, we also sort currencies based on the forward spread (premium). More specifically, we consider a 'one long/one short' portfolio of the four currencies in which the investor longs the currency with the largest negative spread and shorts the currency with the largest positive spread.[17] In all of the subsamples that we consider, JPY is typically the currency with the largest positive spread. This is consistent with this currency being a 'funding currency' as noted in previous research (Galati, Heath and McGuire, 2007). The 'target' or 'investment' currencies in the carry trade alternate between GBP, CAD and EUR.

In addition to portfolios formed on the basis of the spread, we employ the momentum portfolio of Menkhoff, Sarno, Schmeling and Schrimpf (2012a) as an additional benchmark. The momentum portfolio is constructed by sorting currencies on the basis of the lagged carry trade returns. The 'one long/one short' momentum portfolio consists of longing the currency with the largest carry trade return and shorting the currency with the lowest carry trade return at time t.[18] The investor

[17] We also experimented with a 'two long/two short' portfolio constructed by going long the two currencies with the largest negative spread and short the two currencies with the largest positive spread. The latter portfolio's profitability is lower than the 'one long/one short' portfolio (results are available from the authors). This decrease in profitability resonates with Bekaert and Panayotov (2015)'s empirical findings according to which the carry trade's profitability hinges on the currencies included in the portfolio. In fact, the authors distinguish between "good" and "bad" carry trades depending on the currencies comprising the portfolio.

[18] As noted in Li, Tsiakas and Wang (2015), the momentum portfolio allows investors to gain a long exposure in currencies that are on an upward trend and a short exposure in currencies that are on a downward trend.

realizes the returns at time $t+1$. All the portfolios (momentum and benchmark) are rebalanced weekly.

When constructing the long-short portfolios, the net-of-transaction costs returns to a currency in which the investor goes long are computed as $er_{jt+1} - (s^a_{jt+1} - s_{jt+1}) - (f_{jt} - f^b_{jt})$ whereas the net-of-transaction costs returns for a currency in which the investor goes short are $er_{jt+1} - (s_{jt+1} - s^b_{jt+1}) - (f^a_{jt} - f_{jt})$. The results (dollar values and Sharpe ratios) for the portfolio strategies for the financial crisis, post-crisis and combined sample periods are reported in Table 7. For comparability with the existing literature, Table 7 also reports the skewness and the kurtosis of each portfolio's realized returns.

Overall, the portfolio formed on the basis of the decomposition model generates higher profits and risk-adjusted returns than the other portfolios during the financial crisis and the combined period.[19] The decomposition model also exhibits lower kurtosis than the linear model, signed spread and momentum portfolios in the financial crisis, post-crisis and combined sample periods. In addition, the decomposition model's skewness is positive in the post-crisis period. These latter observations suggest that the decomposition model might be less prone to crash risk than these portfolios. Only in the post-crisis period, the risk-adjusted returns of the benchmark carry trade portfolio ('signed spread') are larger than those of the decomposition model. Furthermore, the risk-adjusted returns of the decomposition model portfolio also appear to be larger than most of the individual currencies. This is consistent with portfolio volatility being lower, due to better diversification, than individual currency return volatility.[20] We view our results as suggesting that a portfolio strategy might be less adversely affected than individual currencies by the unwinding of the carry trade.

6 Concluding Remarks

This paper examines the profitability of carry trade returns using a flexible approach that decomposes currency returns into multiplicative sign and absolute return components which exhibit

[19] The Sharpe ratio for the decomposition method during the financial crisis is statistically larger than zero using moving block bootstrap with a blocksize equal to 8. The Sharpe ratios of the other trading strategies are insignificant.

[20] As pointed out by Brunnermeir, Nagel and Pedersen (2009), elevated levels of volatility can lead to the unwinding of carry trade positions, and, consequently, to currency crashes. In this context, it is interesting to note that Menkhoff, Sarno, Schmeling and Schrimpf (2012b) identify a global currency volatility risk factor and show, by sorting currencies into five portfolios based on volatility, that high interest rate currencies yield lower payoffs when global volatility is high.

much greater predictability than raw returns. We allow the two components to respond to currency-specific risk factors and use the joint conditional distribution of these components, modeled as a time-varying copula, to produce forecasts of future returns.

Our out-of-sample forecasting results suggest that the decomposition model exhibits substantial predictability for Canadian dollar and Japanese yen returns during the recent financial crisis. We show that the out-of-sample forecasting gains of the decomposition model translate into economically and statistically highly significant profitability: trading individual currencies or forming portfolios based on the predicted carry trade returns from the decomposition model generates larger risk-adjusted profits than any of the competing models. Our empirical analysis also sheds light on the sources of improvement over the subdued profitability of pure carry trade strategies during and after the financial crisis. It would be interesting, as an avenue for future research, to examine if the carry trade's profitability of the decomposition model extends to other asset markets (such as commodity and bond markets) as in Koijen, Moskowitz, Pedersen and Vrugt (2013).

References

[1] Accominotti, O., and D. Chambers, 2014, Out-of-Sample Evidence on the Returns to Currency Trading, Working Paper, London School of Economics.

[2] Anatolyev, S., and N. Gospodinov, 2010, Modeling Financial Return Dynamics via Decomposition, *Journal of Business and Economic Statistics* **28**, 232–245.

[3] Ang, A., and J. S. Chen, 2010, Yield Curve Predictors of Foreign Exchange Returns, Working Paper, Columbia University.

[4] Bacchetta, P., and E. van Wincoop, 2006, Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?, *American Economic Review* **96**, 552–576.

[5] Backus, D. K., S. Foresi, and C. I. Telmer, 2001, Affine Term Structure Models and the Forward Premium Anomaly, *Journal of Finance* **56**, 279–304.

[6] Bakshi, G., and G. Panayotov, 2013, Predictability of Currency Carry Trades and Asset Pricing Implications, *Journal of Financial Economics* **110**, 139–163.

[7] Berge, T. J., Ò. Jordà, and A. M. Taylor, 2010, Currency Carry Trades, NBER Working Paper No. 16491.

[8] Bekaert, G., and G. Panayotov, 2015, Good Carry, Bad Carry, Columbia Business School Research Paper No. 15-53, Columbia University.

[9] Bessembinder, H., 1992, Systematic Risk, Hedging Pressure, and Risk Premiums in Futures Markets, *Review of Financial Studies* **5**, 637–667.

[10] Brunnermeier, M. K., S. Nagel, and L. H. Pedersen, 2009, Carry Trades and Currency Crashes, *NBER Macroeconomics Annual* **23**, 313–347.

[11] Burnside, C., 2009, Comment on "Carry Trades and Currency Crashes," *NBER Macroeconomics Annual* **23**, 349–359.

[12] Burnside, C., 2012, Carry Trades and Risk, in *Handbook of Exchanges Rates* (J. March, I. W. March and L. Sarno, eds.), 283–312.

[13] Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo, 2011, Do Peso Problems Explain the Returns to the Carry Trade? *Review of Financial Studies* **24**, 853–891.

[14] Burnside, C., M. Eichenbaum, and S. Rebelo, 2011, Carry Trade and Momentum in Currency Markets, *Annual Review of Financial Economics* **3**, 511–535.

[15] Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, NJ.

[16] Constantinides, G. M., 1992, A Theory of the Nominal Term Structure of Interest Rates, *Review of Financial Studies* **5**, 531–552.

[17] Daniel, K., R. J. Hodrick, and Z. Lu, 2014, The Carry Trade: Risks and Drawdowns, NBER Working Paper No. 20433.

[18] Dewally, M., L. H. Ederington, and C. S. Fernando, 2013, Determinants of Trader Profits in Commodity Futures Markets, *Review of Financial Studies* **26**, 2648–2683.

[19] de Roon, F. A., T. E. Nijman, and C. Veld, 2000, Hedging Pressure Effects in Futures Markets, *Journal of Finance* **55**, 1437–1456.

[20] Engel, C., 2015a, Exchange Rates and Interest Parity, in *Handbook of International Economics* (E. Helpman, K. Rogoff and G. Gopinath, eds.), Vol. 4, 453–522.

[21] Engel, C., 2015b, Exchange Rates, Interest Rates, and the Risk Premium, Working Paper, University of Wisconsin.

[22] Engel, C., N. C. Mark, and K. D. West, 2008, Exchange Rate Models Are Not As Bad As You Think, *NBER Macroeconomics Annual* **22**, 381–441.

[23] Engel, C., N. C. Mark, and K. D. West, 2015, Factor Model Forecasts of Exchange Rates, *Econometric Reviews* **34**, 32–55.

[24] Engel, C., and K. D. West, 2005, Exchange Rates and Fundamentals, *Journal of Political Economy* **113**, 485–517.

[25] Engle, R. F., 2002, New Frontiers for ARCH Models, *Journal of Applied Econometrics* **17**, 425–446.

[26] Farhi, E., S. P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan, 2009, Crash Risk in Currency Markets, NBER Working Paper No. 15062.

[27] Ferraro, D., K. Rogoff, and B. Rossi, 2015, Can Oil Prices Forecast Exchange Rates?, *Journal of International Money and Finance*, forthcoming.

[28] Fisher, M., and C. Gilles, 2000, Modeling the State-Price Deflator and the Term Structure of Interest Rates, Unpublished manuscript.

[29] Galati, G., A. Heath, and P. McGuire, 2007, Evidence of Carry Trade Activity, *BIS Quarterly Review* (September), 27–41.

[30] Giacomini, R., and H. White, 2006, Tests of Conditional Predictive Ability, *Econometrica* **74**, 1545–1578.

[31] Gospodinov, N., 2009, A New Look at the Forward Premium Puzzle, *Journal of Financial Econometrics* **7**, 312–338.

[32] Gromb, D, and D. Vayanos, 2010, Limits of Arbitrage, *Annual Review of Financial Economics* **2**, 251–275.

[33] Jordà, Ò, and A. M. Taylor, 2012, The Carry Trade and Fundamentals: Nothing to Fear but FEER Itself, *Journal of International Economics* **88**, 74–90.

[34] Kilian, L., and M. P. Taylor, 2003, Why Is It So Difficult to Beat the Random Walk Forecast of Exchange Rates?, *Journal of International Economics* **60**, 85–107.

[35] Klitgaard, T., and L. Weir, 2004, Exchange Rate Changes and Net Positions of Speculators in the Futures Market, *FRBNY Economic Policy Review* (May), 17–28.

[36] Koijen, R., T. Moskowitz, L. H. Pedersen, and E. Vrugt, 2013, Carry, Fama-Miller Working Paper.

[37] Li, J, I. Tsiakas, and W. Wang, 2015, Predicting Exchange Rates Out of Sample: Can Economic Fundamentals Beat the Random Walk?, *Journal of Financial Econometrics* **13**, 293–341.

[38] Lustig, H., and A. Verdelhan, 2011, The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Reply, *American Economic Review* **101**, 3477–3500.

[39] Lustig, H., N. Roussanov, and A. Verdelhan, 2011, Common Risk Factors in Currency Markets, *Review of Financial Studies* **24**, 3731–3777.

[40] Lustig, H., A. Stathopoulos and A. Verdelhan, 2015, The Term Structure of Currency Carry Trade Risk Premia, Working Paper, University of California Los Angeles.

[41] Mancini, L., A. Ranaldo, and J. Wrampelmeyer, 2013, Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums, *Journal of Finance* **68**, 1805–1841.

[42] Mancini-Griffoli, T., and A. Ranaldo, 2011, Limits to Arbitrage During the Crisis: Funding Liquidity Constraints and Covered Interest Parity, Working Paper, Swiss National Bank.

[43] Manner, H., and O. Reznikova, 2012, A Survey on Time-Varying Copulas: Specification, SImulations, and Application, *Econometric Reviews* **31**, 654–687.

[44] Meese, R. A., and K. Rogoff, 1983, Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?, *Journal of International Economics* **14**, 3–24.

[45] Menkhoff, L., L. Sarno, M. Schmeling and A. Schrimpf, 2012a, Currency Momentum Strategies, *Journal of Financial Economics* **106**, 660–684.

[46] Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf, 2012b, Carry Trades and Global Foreign Exchange Volatility, *Journal of Finance* **67**, 681–718.

[47] Szymanowska, M., F. A. de Roon, T. E. Nijman, and R. van den Goorbergh, 2014, An Anatomy of Commodity Futures Risk Premia, *Journal of Finance* **69**, 453–482.

[48] Tornell A., and C. Yuan, 2012, Speculation and Hedging in the Currency Futures Markets: Are They Informative to the Spot Exchange Rates?, *Journal of Futures Markets* **32**, 122–151.

[49] Vlachev, R., 2015, Exchange Rates and UIP Violations at Short and Long Horizons, Working Paper, Duke University.

Table 1. Descriptive statistics.

	Mean	Med	Std	Min	Max	Skew	Kurt	AC(1)
GBP								
r	-0.003	0.012	1.278	-6.406	5.719	-0.269	4.695	0.005
$spread$	-0.019	-0.011	0.021	-0.103	0.020	-0.591	2.316	0.983
er	0.016	0.034	1.277	-6.348	5.727	-0.262	4.663	0.005
sp	-0.030	-0.048	0.570	-1.000	0.962	0.035	1.684	0.878
CAD								
r	0.014	-0.007	1.121	-5.072	10.00	0.471	10.26	-0.049
$spread$	-0.003	-0.004	0.021	-0.092	0.049	-0.098	3.541	0.983
er	0.017	-0.007	1.122	-5.099	9.988	0.465	10.21	-0.047
sp	0.109	0.154	0.559	-0.963	1.000	-0.105	1.671	0.927
JPY								
r	0.011	-0.053	1.485	-6.078	8.990	0.378	5.475	0.018
$spread$	0.054	0.047	0.041	-0.075	0.136	0.073	1.516	0.992
er	-0.043	-0.102	1.487	-6.178	8.883	0.366	5.411	0.020
sp	-0.186	-0.253	0.528	-1.000	0.874	0.178	1.689	0.928
EUR								
r	0.019	0.050	1.431	-4.235	8.794	0.122	4.557	0.021
$spread$	0.004	0.001	0.025	-0.038	0.056	0.318	2.055	0.992
er	0.014	0.048	1.432	-4.238	8.812	0.129	4.557	0.023
sp	0.165	0.189	0.458	-0.968	0.978	-0.379	2.145	0.954

Notes: The table present the mean, median (Med), standard deviation (Std), minimum value (Min), maximum value (Max), skewness (Skew), kurtosis (Kurt) and the first-order autocorrelation coefficient (AC(1)) of the currency returns (r), basis ($spread$), carry trade returns (er), and speculative pressure (sp) for the period 10/13/1992 to 12/31/2013 (for GBP, CAD, and JPY) and 1/5/1999 to 12/31/2013 (for EUR).

Table 2. Estimation results for currency returns from the decomposition method.

	GBP		CAD		JPY		EUR	
	estim.	s.e.	estim.	s.e.	estim.	s.e.	estim.	s.e.
volatility								
κ	1.212	0.029	1.247	0.030	1.197	0.028	1.300	0.037
ω_v	0.018	0.010	0.021	0.010	-0.009	0.019	0.039	0.011
β_v	0.933	0.020	0.934	0.016	0.843	0.061	0.927	0.018
γ_v	0.032	0.008	0.052	0.011	0.037	0.011	0.023	0.009
ρ_v	-0.022	0.019	-0.009	0.020	0.062	0.035	-0.047	0.021
δ_{1v}	-0.255	0.107	0.089	0.090	0.142	0.149	-0.288	0.110
δ_{2v}	-0.008	0.007	0.008	0.006	0.001	0.032	-0.007	0.007
direction								
ω_d	-0.140	0.081	-0.060	0.084	0.079	0.112	-0.097	0.106
ϕ_d	0.235	0.122	0.013	0.117	-0.045	0.122	0.236	0.147
δ_{1d}	-2.243	1.451	-6.465	2.119	-2.156	1.442	-4.921	2.518
δ_{2d}	-0.262	0.107	0.085	0.111	0.124	0.125	0.327	0.163
copula								
λ_0	0.142	0.121	0.318	0.293	0.214	0.139	-0.040	0.040
λ_1	0.723	0.228	0.659	0.251	0.773	0.154	0.954	0.039
λ_2	-0.270	0.212	-0.439	0.373	-0.257	0.175	0.078	0.070

Notes: The estimation period is 10/13/1992 – 12/31/2013 for GBP, CAD, and JPY, and 1/5/1999 – 12/31/2013 for EUR.

Table 3. Giacomini and White (2006) test for forecast accuracy of carry trade returns.

	GBP		CAD		JPY		EUR	
	Linear	Decom	Linear	Decom	Linear	Decom	Linear	Decom
(1)	0.185	1.171	1.285	3.167	1.599	1.657	0.211	0.963
(2)	-2.099	-0.388	0.132	-0.030	-1.080	-0.368	-0.459	-1.151
(3)	-0.638	0.832	1.286	3.001	0.296	1.097	-0.029	0.832

Notes: The table reports the Giacomini and White (2006) test of differences of carry trade returns $\hat{\mu}^i_{jt+1} - \tilde{\mu}_{jt+1}$ under quadratic loss between the competing model i and the random walk model. The competing models are the linear and decomposition (Decom) models. The benchmark model is the random walk model. (1), (2) and (3) denote the three sample periods: financial crisis (7/10/2007-12/28/2010), post-crisis (1/4/2011-12/31/2013) and combined crisis and post-crisis (7/10/2007-12/31/2013), respectively.

Table 4. Out-of-sample directional forecast performance of carry trade returns.

	RW	Linear	Decom
GBP			
AUC	0.448	0.469	0.545
AUC*	0.528	0.445	0.531
CAD			
AUC	0.474	0.475	0.523
AUC*	0.478	0.464	0.531
JPY			
AUC	0.503	0.523	0.551
AUC*	0.507	0.484	0.580
EUR			
AUC	0.473	0.522	0.545
AUC*	0.424	0.526	0.619

Notes: The table reports results for the area under the receiver operating characteristic curve (AUC) and the corresponding returns-weighted AUC statistic (AUC*) of Jordà and Taylor (2012) for the random walk (RW), linear model, and the decomposition (Decom) model.

Table 5. Dollar values of trading strategies for individual currencies ($100 initial investment).

	Benchmark	Linear	HA	Decom
	Financial Crisis (July 2007 – Dec. 2010)			
GBP	$77.37	$60.09	$69.61	$101.28
CAD	$111.68	$86.96	$100.84	$152.01
JPY	$74.82	$112.77	$95.68	$115.80
EUR	$99.25	$87.91	$91.93	$106.51
	Post-Crisis (Jan. 2011 – Dec. 2013)			
GBP	$94.93	$85.34	$96.56	$99.84
CAD	$96.77	$92.60	$93.09	$94.57
JPY	$104.81	$73.13	$74.24	$87.20
EUR	$102.11	$96.07	$91.58	$105.07
	Combined (July 2007 – Dec. 2013)			
GBP	$73.45	$51.29	$67.51	$101.12
CAD	$108.08	$80.53	$93.88	$143.77
JPY	$78.42	$82.47	$71.04	$100.98
EUR	$101.35	$84.46	$84.19	$111.92

Notes: The table presents the payoffs to a $100 initial investment based on the sign of the predicted carry trade return from one of the competing models during the financial crisis, post-crisis and the combined crisis and post-crisis sample periods. Benchmark refers to a strategy of taking a long or a short position depending on the sign of the spread, Linear refers to the linear model, HA refers to the historical average, and Decom refers to the decomposition model.

Table 6. Sharpe ratios of trading strategies for individual currencies.

	Benchmark	Linear	HA	Decom
	Financial Crisis (July 2007 – Dec. 2010)			
GBP	-0.95	-1.22	-1.18	0.09
CAD	0.72	-0.22	0.07	0.94
JPY	-1.02	0.38	-0.08	0.46
EUR	0.01	-0.24	-0.12	0.21
	Post-Crisis (Jan. 2011 – Dec. 2013)			
GBP	-0.94	-0.83	-0.17	0.02
CAD	-0.24	-0.33	-0.30	-0.23
JPY	0.45	-1.07	-0.95	-0.60
EUR	0.16	-0.09	-0.30	0.28
	Combined (July 2007 – Dec. 2013)			
GBP	-0.80	-1.05	-0.79	0.06
CAD	0.29	-0.24	-0.05	0.55
JPY	-0.56	-0.25	-0.49	0.06
EUR	0.06	-0.18	-0.18	0.23

Notes: The table presents the annualized Sharpe ratios of a trading strategy based on the sign of the predicted carry trade return from one of the competing models during the financial crisis, post-crisis and the combined crisis and post-crisis sample periods. Benchmark refers to a strategy of taking a long or a short position depending on the sign of the spread, Linear refers to the linear model, HA refers to the historical average, and Decom refers to the decomposition model.

Table 7. Dollar values and Sharpe ratios of trading portfolios of carry trade returns ($100 initial investment).

	financial crisis				post-crisis				combined			
	value	skew	kurt	SR	value	skew	kurt	SR	value	skew	kurt	SR
Linear	$86.36	-0.42	5.05	-0.61	$86.94	-0.25	3.26	-1.11	$75.08	-0.40	5.67	-0.77
HA	$89.80	-0.34	3.56	-0.44	$88.95	-0.06	2.95	-0.80	$79.88	-0.27	3.74	-0.57
Decom	$119.50	-0.35	3.91	0.81	$96.98	0.16	3.00	-0.27	$115.89	-0.19	4.77	0.45
	Benchmark strategies											
Spread	$90.14	-0.26	7.85	-0.89	$99.74	-0.19	4.61	-0.03	$89.91	-0.35	8.79	-0.57
1l/1s(s)	$73.05	-0.53	3.89	-1.11	$94.84	-0.05	4.57	-0.27	$69.28	-0.47	4.44	-0.77
1l/1s(m)	$109.77	-0.75	7.14	0.36	$91.35	-0.39	4.82	-0.49	$100.28	-0.64	7.41	0.04

Notes: The table reports the annualized Sharpe ratios (SR), skewness (skew), kurtosis (kurt) and the payoffs to a $100 initial investment in an equally-weighted portfolio formed on the basis of the predicted sign from one of the competing models. The signed spread (Spread) portfolio benchmark is an equally-weighted portfolio which is long currencies with a negative spread and short currencies with a positive spread. The 'one-long/one-short' benchmark (1l/1s(s)) is a long-short portfolio which is long the currency with the largest negative spread and short the currency with the largest positive spread. The 'one-long/one-short' momentum portfolio (1l/1s(m)) is long the currency with the highest carry trade return and short the currency with the lowest carry trade return. Linear refers to the linear model, HA refers to the historical average, and Decom refers to the decomposition model.

Figure 1: Carry Trade Returns.