Journal of Empirical Finance 42 (2017) 199–211

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

Foreign exchange predictability and the carry trade: A decomposition approach☆

Stanislav Anatolyev[a], Nikolay Gospodinov[b,*], Ibrahim Jamali[c], Xiaochun Liu[d]

[a] CERGE-EI, Czech Republic, and New Economic School, Russia
[b] Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, United States
[c] American University of Beirut, Lebanon
[d] University of Alabama, United States

ARTICLE INFO

JEL classification:
F31
F37
C32
C53
G15

Keywords:
Exchange rate forecasting
Carry trade
Return decomposition
Copula
Joint predictive distribution

ABSTRACT

In this paper, we decompose currency returns into multiplicative sign and absolute return components, which exhibit much greater predictability than raw returns, and use the joint conditional distribution of these components to obtain forecasts of future exchange rate returns. Our results suggest that the decomposition model produces higher forecast and directional accuracy than any of the competing models. We undertake trading exercises using carry trade returns and 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.

1. Introduction

Modern international macroeconomic theory is founded on the belief that exchange rate changes can be explained by economic fundamentals. Starting with Meese and Rogoff (1983), a large number of studies 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 and that economic fundamentals do not contain predictive power for exchange rate movements. 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 modern international macroeconomics (Bacchetta and van Wincoop, 2006).[1]

While this earlier literature supports the existence of an "exchange rate disconnect puzzle", the empirical evidence documenting some

☆ The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. The third author gratefully acknowledges financial support from the American University of Beirut's University Research Board.

* Corresponding author.
E-mail address: Nikolay.Gospodinov@atl.frb.org (N. Gospodinov).

[1] The disconnect between fundamentals and exchange rate changes is known as the "exchange rate disconnect puzzle"(see, for example, Della Corte and Tsiakas, 2012). Engel and West (2005) demonstrate analytically that exchange rates can exhibit near-random-walk behavior if fundamentals are integrated of order one and the factor used for discounting future fundamentals is near one. The authors also argue that, in such a setup, the near-random-walk behavior in exchange rates can still be consistent with a rational expectations present-value model and that their findings shed light on the reasons underlying the disconnect between exchange rates and fundamentals. Sarno and Sojli (2009) lend support to Engel and West (2005)'s analytical framework by providing empirical evidence that the discount factor is close to one.

http://dx.doi.org/10.1016/j.jempfin.2017.03.005
Received 8 July 2016; Received in revised form 10 March 2017; Accepted 28 March 2017
Available online 08 April 2017
0927-5398/

success at predicting exchange rate movements has steadily accumulated in subsequent contributions to the literature (Chinn and Meese, 1995; Clarida et al., 2003; Engel et al., 2007; Mark, 1995). More recently, additional empirical evidence suggesting short-horizon predictability in exchange rate returns emerged (Molodtsova and Papell; 2009; Ferraro et al., 2015).[2] However, the most robust evidence of predictive power of economic fundamentals for exchange rate returns is uncovered when the information from many predictors is aggregated. Li et al. (2015) provide evidence that exchange rate return forecasts generated from "kitchen sink" regressions with elastic-net shrinkage methods yield statistical and economic gains over the competing models. Similarly, Della Corte et al. (2009) show that combining the forecasts of different models generates economically significant improvements relative to the random walk.[3]

The findings emerging from recent contributions to the literature confirm that predictability is not confined to exchange rate returns but also extends to the cross-section and time series of carry trade returns. Even if exchange rates were to follow a random walk (and exchange rate returns are unpredictable), the returns to the carry trade, a popular trading strategy among investors, are predictable. In a carry trade, an investor borrows in a low-interest currency and invests the borrowed funds in a high-yielding currency. Note that the carry trade returns consist, as detailed next, of two parts- a future currency return and an interest rate differential. This implies that carry trade predictability can arise from predictability in either exchange rate returns or interest rate differentials. If exchange rates follow a random walk, carry trade returns would be equal to the interest rate differential between two currencies and are therefore predictable (Li et al., 2015).

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 and that carry trade returns are predictable. These latter two findings largely explain the strategy's popularity. Lustig et al. (2011) provide evidence that the interest rate differential (i.e. a "high-minus-low" interest rate factor) is an important cross-sectional predictor of carry trade returns while Menkhoff et al. (2012a) find that a global foreign exchange volatility factor is another cross-sectional predictor of carry trade returns.

A number of studies also uncover time series predictability in carry trade returns (Ang and Chen, 2010; Bakshi and Panayotov, 2013; Cenedese et al., 2014).[4] The predictability of the carry trade returns and the profitability of the carry trade are consistent with the lack of empirical support for Uncovered Interest Parity (UIP) and with the "forward premium puzzle".[5]

In this paper, we adopt a statistical approach to uncovering and exploiting potential predictability in future currency returns. More specifically, we capitalize on the method proposed by Anatolyev and Gospodinov (2010), extended to time-varying copulas, to decompose currency 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. This method of incorporating any implicit nonlinearities in a flexible, indirect fashion is motivated by the 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 (Cenedese et al., 2014; Clarida et al., 2003; Kilian and Taylor, 2003). The decomposition model thus exploits deviations from UIP to predict future currency returns which are the uncertain part of carry trade returns.[6] While we assess the statistical accuracy of the currency returns forecasts from the decomposition model, we examine the economic significance of our findings using carry trade returns in light of the popularity of the carry trade among currency investors.

By virtue of allowing for nonlinearities in currency returns, our paper relates to an existing line of research that exploits the existence of nonlinearities in foreign exchange rate returns for predictive purposes (Clarida et al., 2003; Kilian and Taylor, 2003). By way of conditioning on absolute returns for two quantiles, our approach also closely relates to that of Cenedese et al. (2014) who decompose market variance into average variance and average correlation and provide empirical evidence, using predictive quantile regressions, that average variance is a strong predictor of future carry trade returns.

Several interesting results emerge from our analysis. First, the decomposition model exhibits substantial directional accuracy in predicting future currency returns. Second, the out-of-sample forecasting gains of the decomposition model (relative to the random walk and linear prediction models) translate into economically and statistically highly significant profitability. More specifically,

[2] At first, the findings emerging from the literature pointed to long-horizon predictability in exchange rates (Chinn and Meese, 1995; Mark, 1995). Later studies also find short-horizon predictability in exchange rate movements using Taylor rule fundamentals (Molodtsova and Papell, 2009). Della Corte and Tsiakas (2012) articulately describe the state of the literature since Meese and Rogoff (1983) as coming "full circle"from finding no predictability in exchange rate movements to uncovering long-horizon predictability (Mark, 1995) then back to finding no predictability (Cheung et al., 2005) to eventually finding evidence of short-horizon predictability (Molodtsova and Papell, 2009).

[3] Della Corte and Tsiakas (2012) summarize the three approaches that have traditionally been used in the literature to uncover predictability in exchange rates returns. One group of studies assesses a (linear) model's statistical performance vis-à-vis the random walk and is therefore specific to the model in question and is concerned with the time series predictability in individual currency returns. Another strand of research evaluates the economic significance of several exchange rate forecasts within a portfolio allocation exercise while the third approach relies on forecast combination in order to account for model uncertainty. In this paper, we assess the statistical and economic significance of our forecasts and thereby employ the first and second approaches to assessing predictability.

[4] Bakshi and Panayotov (2013) identify the TED spread as an important predictor of carry trade returns while Ang and Chen (2010) find that term structure of interest rate variables are useful predictors of carry trade returns. The success of the TED spread and the currency volatility factors in predicting carry trade returns is consistent with the ability of these variables to proxy for crash (Brunnermeier et al., 2009; Berge et al., 2010; Farhi et al., 2009; Jordà and Taylor, 2012) and liquidity (Mancini et al., 2013; Mancini Griffoli and Ranaldo, 2011) risk.

[5] The "forward premium puzzle"is the empirical regularity that forward rates are biased predictors of future spot rates. For a recent review of the literature on the "forward premium puzzle", see Engel (2015). Under Covered Interest Parity (CIP), the "forward premium puzzle"is equivalent to an empirical violation of UIP (Cenedese et al., 2014; Della Corte and Tsiakas, 2012; Li et al., 2015). The empirical rejection of UIP, in turn, implies that exchange rates do not sufficiently adjust to eliminate the gains that investors can realize by borrowing in low-yielding currencies and investing in high-yielding currencies. The latter strategy is a simple implementation of the carry trade and violations of UIP therefore give rise to carry trade profitability (Della Corte and Tsiakas, 2012).

[6] As discussed next, carry trade returns are composed of two parts: the forward premium (or interest rate differential) which constitutes the fixed or perfectly predictable part of carry trade returns and the future currency return, which constitutes the uncertain part of the carry trade return. We use to decomposition model to generate forecasts of future currency returns.

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 foreign exchange predictability literature along the following lines. From a modeling perspective, this paper offers a new approach to modeling and forecasting currency returns. While pure carry trade strategies exploit only the differential in interest rates, both of which are near the zero lower bound over much of our out-of-sample forecasting period, we employ a model-based carry trade strategy that capitalizes on the predictability of future currency returns. As we show in the paper, our decomposition model uncovers a large degree of predictability that generates highly profitable trading strategies. On the empirical front, our analysis sheds light on the recent profitability of the carry trade and momentum portfolios as a by-product of the economic assessment of the decomposition model.

The rest of the paper proceeds as follows. Section 2 provides a motivating example for our decomposition approach as well as a detailed discussion of the specific decomposition model that we employ. The data and variables employed in the empirical analysis are described in Section 3. Our empirical findings as well as the trading strategies that we consider to assess the profitability of the decomposition model are discussed in Section 4. Section 5 offers some concluding remarks.

2. Decomposition method

2.1. Motivating example

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 $E[r_t] = E[a_t]E[b_t] = 0$, and, for $j > 0$,

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

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 $j > 0$:

$$E[r_t a_{t-j}] = E[a_t b_t a_{t-j}] = E[a_t a_{t-j}]E[b_t] = 0,$$
$$E[r_t b_{t-j}] = E[a_t b_t b_{t-j}] = E[a_t]E[b_t b_{t-j}] = 0,$$
$$E[r_t x_{t-j}] = E[a_t b_t x_{t-j}] = E[a_t]E[b_t x_{t-j}] = 0.$$

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

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

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.

2.2. General approach

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

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

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 $\psi_t = E(|r_t| \parallel \mathcal{F}_{t-1})$ 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

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

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

$$\theta_t = \omega_d + \phi_d \mathbb{I}[r_{t-1} > 0] + y_{t-1}{}'\delta_d, \tag{2}$$

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 $R_t \equiv (|r_t|, \mathbb{I}[r_t > 0])'$, 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 $C(w_1, w_2)$ 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

$$f_{R_t}(u, v) = f^{\mathcal{D}}(u|\psi_t)\varrho_t(F^{\mathcal{D}}(u|\psi_t))^{v}(1 - \varrho_t(F^{\mathcal{D}}(u|\psi_t)))^{1-v}, \tag{3}$$

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

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

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

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

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

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 $(\varsigma, \omega_v, \beta_v, \gamma_v, \rho_v, \delta_v{}', \omega_d, \phi_d, \delta_d{}', \lambda_0, \lambda_1, \lambda_2)'$ is estimated by maximum likelihood. From (3), the sample log-likelihood function to be maximized is given by

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

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

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

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

$$\hat{r}_{t+1} = 2\hat{\xi}_{t+1} - \hat{\psi}_{t+1}, \tag{4}$$

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

[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.

$$\xi_{t+1} = \int_0^1 Q^{\mathcal{D}}(z)\varrho_{t+1}(z)dz, \tag{5}$$

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 (5) using fitted $Q^{\mathcal{D}}(z)$ and fitted $\varrho_{t+1}(z)$.

3. Data description

Our empirical analysis is conducted with a cross-section comprising the ten major and most liquid (Della Corte and Tsiakas, 2012) currencies (G10) over the period January 1976 to December 2016. More specifically, monthly data on the spot (S_t) and one-month forward (F_t) rates, expressed in US dollars per unit of the foreign currency, for the British pound (GBP), Canadian dollar (CAD), Swiss franc (CHF), Euro (EUR), Japanese yen (JPY), Australian dollar (AUD), New Zealand dollar (NZD), Swedish krona (SEK) and Norwegian krone (NOK) are obtained from Barclays Bank through Datastream[8] for the period December 1984 to December 2016.[9]

One-month forward rates for the period January 1976 to November 1984 are implied using CIP. In order to apply CIP, we first obtain data on the one-month Euro deposit rates starting January 1976 for GBP, CAD, CHF, USD and in August 1978 for JPY from Datastream. The interest rate data are then combined with spot exchange rates obtained from the Federal Reserve's H.10 statistical release and the one-month forwards for these currencies are computed using CIP.[10]. Given that the one-month Euro deposit rate data for Japan are available only after August 1978, we follow Della Corte and Tsiakas (2012) by using spot and one-month forward data for JPY from Hai et al. (1997) to backfill our data to January 1976.[11] One month forward and spot data for the Euro are available starting only January 1999. Therefore, we extend the spot and forward rate data for the Euro back to January 1976 by splicing them with spot and one-month forward rate data on the Deutsche Mark (DEM) using the fixed conversion rate between the Euro and Deutsche Mark of January 1999.[12]

As a result of the unavailability of one-month Euro deposit rate data for Norway, Sweden, New Zealand and Australia prior to 1997, we again follow Della Corte and Tsiakas (2012) to extend the spot and forward rate data for these currencies back to 1976. For Norway and Sweden, we obtain spot and one-forward rate data, quoted in terms of British pound, from Datastream and convert these rates into USD using the GBP/USD exchange rate. We obtain money market rate data for New Zealand and Australia from Datastream and combine these interest rate data with spot exchange rate data in order to imply the NZD and AUD forward rates via CIP for the period January 1976 to November 1984.[13]

The currency returns for the j-th exchange rate are constructed as $r_{jt+1} = \ln(S_{jt+1}) - \ln(S_{jt})$. The main interest of this paper lies in forecasting carry trade (or excess currency) returns that are defined as $er_{jt+1} = \ln(S_{jt+1}) - \ln(F_{jt}) = s_{jt+1} - f_{jt} = (s_{jt+1} - s_{jt}) - (f_{jt} - s_{jt}) = r_{jt+1} - fp_{jt}$, where $fp_{jt} = \ln(F_{jt}) - \ln(S_{jt}) = f_{jt} - s_{jt}$. Under risk-neutrality and rational expectations, UIP implies that the expected carry trade return is equal to zero (Li et al., 2015).[14] A sizeable number of existing studies document the lack of empirical support for UIP and deviations from UIP imply that carry trade returns are different from zero.

Unlike UIP which is a predictive relationship, CIP is a contemporaneous arbitrage condition which stipulates that the forward premium is equal to the interest rate differential $f_{jt} - s_{jt} = i_t - i_t^*$, where i_t and i_t^* denote the nominal interest rates in the domestic and foreign currency, respectively.[15] In contrast to the weak empirical support for UIP, ample empirical evidence exists in support of CIP.[16]

To be exact, carry trade (or excess currency) returns are given by $er_{jt+1} = (s_{jt+1} - s_{jt}) - (i_t - i_t^*)$ as they measure the returns from investing in a foreign currency that is financed by borrowing in the domestic currency (Brunnermeier et al., 2009). Under CIP, the excess currency return can be equivalently written, in terms of the forward premium, as $er_{jt+1} = (s_{jt+1} - s_{jt}) - (f_{jt} - s_{jt})$.

Table 1 provides summary statistics of the exchange rate data. The currency returns for the nine exchange rates are plotted in Fig. 1. As documented elsewhere in the literature, the exchange rate returns appear to be serially uncorrelated while the forward

[8] The bid and ask data used in the profit computations of Section 4.2 are also obtained from Barclays Bank via Datastream.

[9] The cross-section of G10 currencies is used in a number of influential studies such as Della Corte and Tsiakas (2012), Li et al. (2015), Burnside et al. (2011a), Lustig et al. (2011), Daniel et al. (2017), and Bekaert and Panayotov (2016). The one-month forward and spot rates are end-of-month observations sampled from daily data.

[10] The pre-December 1984 spot exchange rates can be downloaded from: https://www.federalreserve.gov/ releases/h10/hist/default1989.htm

[11] Table 1 of Della Corte and Tsiakas (2012) provides an excellent overview of the data sources, including thorough information regarding series codes. This detailed data description provides guidance for assembling spot and one-month forward rate data for the G10 currencies starting in 1976.

[12] We thank an anonymous referee for suggesting that we follow the literature by merging data for EUR with data for DEM in order to extend the Euro spot and forward rates series backwards. Similar to the other currencies, we obtain the one-month forward rate data for DEM, used to splice the EUR series and extend it back to 1976, by applying CIP to spot and one-month Euro deposit rates for DEM.

[13] The source of the interest rate data for New Zealand and Australia is the International Financial Statistics database of the International Monetary Fund. The series codes can be found in Della Corte and Tsiakas (2012).

[14] More precisely, Li et al. (2015) note that the implications of UIP are threefold. First, UIP implies that the forward rate is an unbiased predictor of the future spot rate. The large literature on the "forward premium puzzle"provides strong empirical evidence against the unbiasedness of forward rates as predictors of future spot exchange rates. Second, as noted above, the expected currency excess (carry trade) returns are equal to zero under UIP. Hassan and Mano (2014) argue that while the forward premium puzzle implies non-zero currency excess returns, the two have to be treated as separate phenomena. Third, under CIP and UIP, the expected currency returns are equal to the forward premium.

[15] A common approach to testing CIP is to estimate 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 equal zero and one, respectively.

[16] See, for example, Akram et al. (2008) for evidence of empirical support for CIP. More recent research (Baba and Packer, 2009) suggests that short-lived deviations from CIP can exist during periods of financial turmoil (such as the period following the Lehman bankruptcy) due to insufficient liquidity (Mancini Griffoli and Ranaldo, 2011). Subsequent research (Borio et al., 2016; Du et al., 2016) also maintains that temporary deviations from CIP continued to occur after the financial crisis.

Table 1
Descriptive statistics.

currency	variable	mean	med	std	min	max	skew	kurt	AC(1)
GBP	r	-0.100	-0.077	2.974	-13.172	13.347	-0.257	4.789	0.073
	fp	-0.144	-0.087	0.213	-1.118	0.796	-0.633	5.312	0.827
	er	0.044	0.042	3.007	-12.655	13.782	-0.205	4.658	0.091
CAD	r	-0.056	-0.006	1.990	-12.452	8.984	-0.512	7.786	-0.051
	fp	-0.064	-0.050	0.140	-0.758	0.406	-0.541	4.726	0.675
	er	0.008	0.009	2.005	-12.623	8.976	-0.552	7.739	-0.043
CHF	r	0.192	0.114	3.445	-14.554	12.803	-0.016	4.097	-0.002
	fp	0.216	0.207	0.268	-0.885	1.052	0.303	3.969	0.841
	er	-0.024	-0.123	3.469	-15.331	12.598	-0.066	4.037	0.016
EUR	r	0.070	0.131	3.146	-11.139	10.006	-0.141	3.718	0.008
	fp	0.096	0.100	0.225	-0.546	0.868	-0.107	3.896	0.835
	er	-0.026	0.063	3.156	-10.950	9.543	-0.174	3.667	0.019
JPY	r	0.195	0.052	3.303	-10.363	15.980	0.350	4.349	0.052
	fp	0.223	0.186	0.268	-2.232	1.796	-0.803	18.405	0.515
	er	-0.027	-0.082	3.326	-11.070	15.552	0.293	4.223	0.072
AUD	r	-0.112	0.068	3.273	-19.553	8.930	-1.055	7.796	0.026
	fp	-0.208	-0.199	0.258	-1.035	0.712	-0.152	4.579	0.831
	er	0.097	0.278	3.291	-19.394	9.135	-0.994	7.620	0.037
NZD	r	-0.081	0.094	3.431	-24.850	12.263	-1.056	9.951	-0.011
	fp	-0.290	-0.230	0.365	-3.045	0.516	-2.624	14.927	0.819
	er	0.208	0.272	3.498	-24.981	12.469	-0.914	9.685	0.019
SEK	r	-0.148	-0.063	3.178	-16.959	8.777	-0.661	5.817	0.076
	fp	-0.142	-0.111	0.311	-2.845	1.117	-2.336	18.081	0.663
	er	-0.005	0.075	3.182	-16.563	8.824	-0.575	5.335	0.087
NOK	r	-0.090	0.062	3.055	-13.008	10.003	-0.310	4.314	0.018
	fp	-0.170	-0.137	0.321	-1.768	2.765	0.541	19.965	0.656
	er	0.080	0.201	3.057	-12.804	10.059	-0.274	4.141	0.038

Notes: The table presents 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), forward premium (fp) and carry trade returns (er) for the period January 1976 to December 2016. All series are in percent.

premium is highly persistent with variability which is only a small fraction of the variability of the currency returns (Gospodinov, 2009). The null hypothesis of a zero mean of exchange rate returns cannot be rejected at any commonly used significance level for all currencies. As a result, we use the decomposition approach described in the previous section to obtain the forecast of $s_{jt+1} - s_{jt}$ first and then construct the predicted carry trade returns by subtracting the (observed at time t) forward premium $f_{jt} - s_{jt}$. This also 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. Empirical results

4.1. Currency return predictability

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

4.1.1. In-sample estimation results

Let $\bar{r}_{t-1} = r_{t-1} - \frac{1}{t-1}\sum_{i=1}^{t-1} r_i$ denote the demeaned exchange rate returns. The predictors for the volatility model (1) and direction model (2) are $x_{t-1} = fp_{t-1}$ and $y_{t-1} = (fp_{t-1}, \bar{r}_{t-1}^2, \bar{r}_{t-1}^3)'$, respectively. The higher moments of returns are included in the direction model following Anatolyev and Gospodinov (2010). Thus, the decomposition model for the dynamic behavior of exchange rate returns r_t uses marginals that are given by

$$\ln \psi_t = \omega_v + \beta_v \ln \psi_{t-1} + \gamma_v |r_{t-1}| + \rho_v \mathbb{I}[r_{t-1} > 0] + \delta_v fp_{t-1}$$

$$\theta_t = \omega_d + \phi_d \mathbb{I}[r_{t-1} > 0] + \delta_{1d} fp_{t-1} + \delta_{2d} \bar{r}_{t-1}^2 + \delta_{3d} \bar{r}_{t-1}^3,$$

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

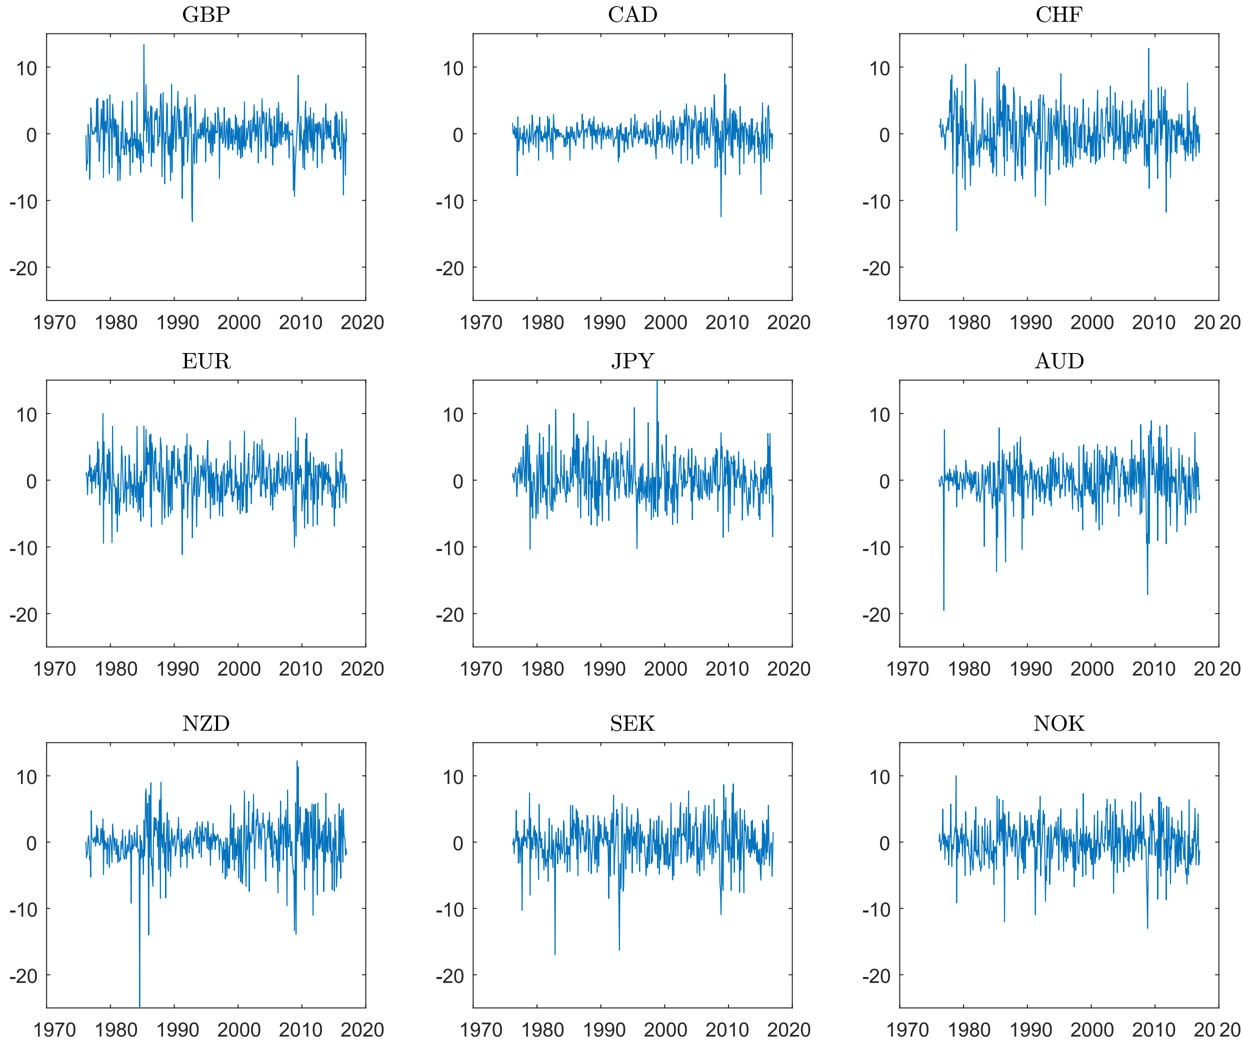

Fig. 1. Monthly returns for the spot exchange rates against USD (in percent).

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

where $\alpha_t = \lambda_0 + \lambda_1\alpha_{t-1} + \lambda_2|r_{t-1}|(1 - \mathbb{I}[r_{t-1} > 0])$. As specified above, the models for each of the nine exchange rates against the USD are estimated with monthly data for the period January 1976 – December 2016.

Table 2 reports the estimated parameters and their corresponding t -statistics. A few interesting observations emerge from these results. All of the exchange rates are characterized by strong persistence in volatility. The forward premium appears to be a useful predictor for both the volatility and the direction models.[17] Finally, the persistence parameter in the time-varying copula process is large and significant for GBP, CAD, CHF, JPY and NOK.

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 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.[18]

[17] In an earlier version of the paper, we also experimented with speculative and hedging pressure for each currency as well the S & P500 option-implied volatility index (VIX) as predictors.

[18] This point is accentuated in Sarno and Valente (2009) who conduct a thorough search for the best-performing model in terms of exchange rate return prediction. Their analysis also reveals frequent shifts in the relation between economic fundamentals and exchange rate movements which possibly explains the success of non-linear models of exchange rate return prediction.

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

	GBP	CAD	CHF	EUR	JPY	AUD	NZD	SEK	NOK
volatility									
ς	1.098	1.134	1.189	1.150	1.149	1.118	1.075	1.170	1.138
	(2.468)	(3.357)	(4.427)	(3.652)	(3.635)	(3.043)	(2.013)	(4.100)	(3.423)
ω_v	−0.530	−0.419	−0.828	−0.724	−0.722	−1.270	−0.367	−1.210	−2.584
	(−1.928)	(−3.277)	(−10.82)	(−1.218)	(−1.769)	(−3.439)	(−4.054)	(−2.373)	(−2.556)
β_v	0.869	0.916	0.800	0.815	0.824	0.674	0.910	0.684	0.347
	(12.54)	(35.09)	(25.06)	(5.293)	(7.962)	(7.393)	(42.02)	(5.283)	(1.317)
γ_v	1.971	4.146	3.012	2.011	1.359	4.461	2.642	3.003	4.116
	(2.305)	(3.534)	(3.156)	(1.025)	(0.733)	(2.582)	(4.291)	(1.590)	(1.322)
ρ_v	−0.047	0.008	0.025	−0.037	0.026	−0.210	−0.086	−0.122	0.028
	(−1.385)	(0.271)	(0.691)	(−0.814)	(0.538)	(−3.619)	(−3.029)	(−2.371)	(0.427)
δ_v	−4.959	6.027	2.753	2.051	10.535	−17.80	−1.970	−8.513	−4.215
	(−4.551)	(9.921)	(2.112)	(2.969)	(31.68)	(−15.50)	(−3.447)	(−3.101)	(−2.318)
direction									
ω_d	−0.220	−0.146	−0.089	−0.094	−0.004	−0.109	−0.283	−0.225	0.051
	(−1.491)	(−1.048)	(−0.528)	(−0.612)	(−0.020)	(−0.721)	(−1.892)	(−1.626)	(0.347)
ϕ_d	0.148	0.153	0.544	0.481	0.387	0.014	0.235	0.227	0.207
	(0.779)	(0.807)	(2.772)	(2.439)	(1.993)	(0.074)	(1.254)	(1.187)	(1.085)
δ_{1d}	−0.952	−1.599	−0.645	−0.143	−0.744	−0.851	−0.586	−0.581	0.122
	(−2.098)	(−2.396)	(−1.863)	(−2.108)	(−2.004)	(−2.342)	(−2.203)	(−2.804)	(1.790)
δ_{2d}	0.020	0.044	0.012	−0.076	0.034	−0.026	0.062	0.013	−0.075
	(1.322)	(1.007)	(1.187)	(−1.296)	(1.806)	(−1.410)	(1.107)	(1.044)	(−1.263)
δ_{3d}	0.026	0.013	−0.042	−0.028	−0.039	−0.028	−0.024	0.035	−0.013
	(1.256)	(1.314)	(−1.908)	(−1.324)	(−1.522)	(−1.611)	(−1.245)	(1.359)	(−1.712)
copula									
λ_0	0.222	−0.205	−0.018	−0.236	0.207	−0.348	−0.009	0.257	0.014
	(1.698)	(−2.018)	(−0.626)	(−0.858)	(1.543)	(−1.189)	(−0.048)	(0.851)	(0.307)
λ_1	0.623	0.750	0.968	0.109	0.818	0.516	0.478	0.100	0.980
	(2.634)	(2.209)	(17.29)	(0.247)	(4.746)	(1.081)	(1.868)	(0.199)	(20.04)
λ_2	−12.92	14.25	2.157	19.95	−11.25	8.512	−15.86	−31.03	−2.303
	(−7.191)	(7.060)	(0.662)	(5.795)	(−1.619)	(3.716)	(−2.595)	(−6.621)	(−1.389)

Notes: The table reports the maximum likelihood estimates and their *t*-statistics (in parentheses below the estimates) for the decomposition model with time-varying copula dependence. All *t*-statistics are for statistical significance except for the *t*-statistic for ς which is testing the null that $\varsigma = 1$.

4.1.2. Out-of-sample forecasting results

This section reports results for one-step ahead out-of-sample forecasts for currency returns. The forecast period is from January 2004 to December 2016. We set the length of the window in the out-of-sample forecasting exercise to be R=335, which is the length of the initial estimation sample, and consider two competing models: a linear predictive model and the decomposition model. The linear model uses the same vector of predictors, $(fp_{t-1}, \bar{r}^2_{t-1}, \bar{r}^3_{t-1})'$, as the decomposition model. The benchmark is a random walk (RW) model with drift.

Table 3 reports the out-of-sample pseudo-R^2 statistic of Campbell and Thompson (2008) for forecasting currency returns[19]

$$R^2_{oos} = 1 - \frac{\sum_{t=R}^{T-1} L(r_{jt+1} - \hat{r}_{jt+1})}{\sum_{t=R}^{T-1} L(r_{jt+1} - \tilde{r}_{jt+1})},$$

where $L(\cdot)$ is a loss function based on forecast errors, and $\hat{r}_{jt+1}$ and $\tilde{r}_{jt+1}$ denote the forecasts from one of the competing models and the benchmark (RW) model, respectively. More specifically, a positive R^2_{oos} statistic implies that the competing model outperforms the benchmark model in terms of out-of-sample prediction and vice versa. This paper considers loss functions based on both squared and absolute forecast errors.

The results suggest that the decomposition model dominates the random walk with drift and linear models for all currencies and loss functions. Under the quadratic (absolute) loss function, the largest forecast gains are obtained for EUR (AUD) followed by those for CHF and AUD (CHF and EUR). Albeit more moderate, the out-of-sample improvement for the other currencies is also impressive.

To determine the statistical significance of these findings, Table 3 also presents the test proposed by Giacomini and White (GW, 2006) based on the difference in currency j forecast losses $\triangle L_{jt+1} = L(r_{jt+1} - \hat{r}_{jt+1}) - L(r_{jt+1} - \tilde{r}_{jt+1})$, where $L(r_{jt+1} - \hat{r}_{jt+1})$ and $L(r_{jt+1} - \tilde{r}_{jt+1})$ are losses computed from the competing model and the RW model, respectively. As for R^2_{oos}, we use both quadratic and absolute error loss functions. Following Giacomini and White (GW, 2006), we use the conditional version of the test with a test function $(1, \triangle L_{jt})'$. The difference is statistically significant if the test statistic exceeds the critical value from a χ^2_2 distribution. A plus (minus) sign indicates that the corresponding model outperforms (underperforms) the benchmark statistically significantly at the

[19] While Campbell and Thompson (2008) and Welch and Goyal (2008) employ the statistic R^2_{oos} to assess the performance of competing models in predicting the equity risk premium, Della Corte and Tsiakas (2012) and Li et al. (2015) use R^2_{oos} to compare the out-of-sample predictive accuracy of competing models for foreign exchange rate returns.

Table 3
Statistics for out-of-sample forecast performance.

	R^2_{oos}				GW				AUC		
	quadratic		absolute		quadratic		absolute				
	linear	decom	linear	decom	linear	decom	linear	decom	RW	linear	decom
GBP	−2.63	5.91	−2.64	3.54	$7.67^{(-)}$	$4.62^{(+)(+)}$	$6.73^{(-)(-)}$	$5.62^{(+)}$	0.518	0.421	0.582
CAD	−17.80	4.24	−2.46	3.99	1.23	0.71	0.94	1.16	0.432	0.505	0.605
CHF	−5.14	7.41	−1.02	4.74	3.18	3.98	0.92	$6.83^{(+)}$	0.513	0.555	0.612
EUR	−2.06	8.39	−0.27	4.61	1.82	$5.12^{(+)}$	0.21	$5.85^{(+)}$	0.543	0.517	0.612
JPY	−1.65	0.75	−0.00	2.90	$4.80^{(-)}$	1.13	2.47	4.16	0.515	0.546	0.588
AUD	−8.74	6.81	−3.79	5.15	$4.92^{(-)}$	$5.55^{(+)}$	$4.81^{(-)}$	$10.18^{(+)}$	0.436	0.441	0.646
NZD	−10.13	0.70	−5.31	0.97	4.48	$6.18^{(+)}$	$5.60^{(-)}$	1.61	0.406	0.408	0.508
SEK	−4.49	4.96	−2.48	3.35	1.84	3.22	2.26	4.57	0.440	0.451	0.557
NOK	−4.10	4.46	−2.25	3.19	2.74	3.92	3.28	$7.78^{(+)}$	0.496	0.477	0.586

Notes: R^2_{oos} denotes the out-of-sample R^2, based on squared and absolute forecast errors, for the linear and decomposition (decom) models, using the random walk (RW) with drift as a benchmark model. GW denotes the Giacomini and White (2006) test of forecast performance under quadratic and absolute loss between the competing (linear or decomposition) model and the RW model. A plus (minus) sign for GW indicates that the model outperforms (underperforms) the benchmark statistically significantly at the 10% level. The statistics without any sign signify that the test does not reject the null of an equal predictive ability. AUC denotes the area under the receiver operating characteristic curve of Jordà and Taylor (2012) for RW, linear, and the decomposition models. The out-of-sample evaluation period is January 2004 – December 2016.

10% level. The results in Table 3 show that the decomposition model dominates the RW and linear models for all currencies. The improvements of the decomposition model over the RW benchmark are statistically significant for GBP, EUR, AUD and NZD under a quadratic loss function and for GBP, CHF, EUR, AUD and NOK under an absolute loss function.

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) suggested by Jordà and Taylor (2012).[20] Note that higher values for AUC indicate a better directional forecast performance. As evident from Table 3, the directional forecasts of the decomposition model outperform those based on the linear and RW with drift models across all currencies.

Since Meese and Rogoff (1973), researchers routinely compare models' statistical and directional forecast accuracy to that of the RW. Nonetheless, Della Corte and Tsiakas (2012) observe that even marginal statistical forecast accuracy gains for a model vis-à-vis the RW may translate into economic profitability. In addition, Melvin et al. (2013) persuasively argue that a model which correctly rank-orders currency return forecasts (relative to one other in the cross-section) is sufficient for successful currency investing. That is, from the perspective of investors, (losses) profits do not directly follow from a model's (in)ability to outperform the RW benchmark. Based on the prior observations, we next evaluate the economic significance of the accuracy of these directional forecasts in the context of trading strategies.

4.2. Trading strategies

4.2.1. 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. We conduct our trading exercises using the sign of the predicted carry trade returns, rather than the sign currency returns, so as to benchmark our results against the carry trade which is a popular strategy in active currency management (Li et al., 2015). 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 et al. (2011b), 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 realized return from trading based on model i is

$$\hat{\mu}^i_{jt+1} = \text{sign}(\widehat{er}^i_{jt+1}) er_{jt+1}, \tag{6}$$

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. It is important to stress that the predicted returns $\hat{r}^i_{jt+1}$ are genuine out-of-sample forecasts that utilize information only up to time t.

Following the literature (Lustig et al., 2011; Menkhoff et al., 2012a; Bakshi and Panayotov, 2013), we employ the bid-ask prices to account for the transaction costs that investors would incur when implementing the trades. Let ${s_{jt}}^a$ and ${s_{jt}}^b$ 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 goes long in the foreign currency if the predicted carry trade return from one of the econometric

[20] 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 $k = 1, \ldots, N_p$, and $\eta_l = \{\widehat{er}_{t+1} | d_{t+1} = -1\}$ for $l = 1, \ldots, N_n$. Then, AUC $= \frac{1}{N_p N_n} \sum_{k=1}^{N_p} \sum_{l=1}^{N_n} \mathbb{1}[\phi_k > \eta_l]$. See Jordà and Taylor (2012) for more details.

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 et al., 2011b). The predicted carry trade return of the benchmark strategy is $\widehat{er}^0_{jt+1} = -fp_{jt}$ which is equivalent to trading based on a random walk model for the spot exchange rate.[21]

The value of the initial investment is set equal to \$100. The value of the portfolio is re-calculated and re-invested every period. The results for the net-of-transaction-cost payoffs as well as the annualized average returns, volatilities and Sharpe ratios are presented in Table 4.

Table 4 shows that, with the exception of JPY and NZD (which are popular carry trade funding and investment currencies, respectively), the decomposition model outperforms the remaining models with some of the differences in payoffs (GBP, EUR and AUD) being large. Similar results hold for the risk-adjusted returns (Sharpe ratios) in Table 4.[22] 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 4 for the trading strategy with the option to stay outside the market.

4.2.2. *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 is constructed by averaging the net-of-transaction costs realized returns from one of the econometric models across all the currencies. The excess return of the equally weighted portfolio therefore measures the returns accruing to an investor who trades the currencies based on the predicted sign of the carry trade returns from one of the econometric models.

Following the literature (Accominotti and Chambers, 2014; Ang and Chen, 2010; Brunnermeier et al., 2009; Lustig et al., 2011; Li et al., 2015), we construct benchmark long-short carry trade portfolios on the basis of interest rate differentials (as proxied for, under CIP, by fp_{jt}). The first portfolio, referred to as 'signed forward premium' or 'signed fp', consists of going long a currency that is at a forward discount (i.e., a currency with a negative forward premium) and going short a currency that is at a forward premium (i.e., a currency whose forward premium is positive). The returns of this portfolio are equivalent to trading based on a random walk model for each of the exchange rates. In line with prior research, we also sort currencies based on the forward premium. More specifically, we consider three long-short 'forward premium' portfolios: a 'one long/one short' portfolio (1l/1s(fp)) in which the investor longs the currency with the largest negative forward premium and shorts the currency with the largest positive forward premium, a 'two long/two short' portfolio (2l/2s(fp)) in which the investor longs the two currencies with the largest negative forward premiums and shorts the two currencies with the largest positive forward premiums as well as a 'three long/three short' portfolio (3l/3s(fp)) in which the investor goes long the three currencies with the largest negative forward premiums and goes short the three currencies with the largest positive forward premiums.[23] For most of our out-of-sample period, JPY and CHF are the currencies with the largest positive forward premium.[24] This is consistent with these currencies being 'funding currencies' as noted in previous research (Brunnermeier et al., 2009; Galati et al., 2007). The 'target' or 'investment' currencies in the carry trade are AUD and NZD.

In addition to portfolios formed on the basis of the forward premium, we employ the momentum portfolios of Menkhoff et al. (2012b) constructed by sorting currencies on the basis of the lagged carry trade returns as additional benchmarks. Again, we consider three variants of the momentum portfolio: The 'one long/one short' momentum portfolio (1l/1s(m)) consists of longing the currency with the largest carry trade return and shorting the currency with the lowest carry trade return at time t, the 'two long/two short' portfolio (2l/2s(m)) consists of going long the two currencies with the largest carry trade returns and short the two currencies with the lowest carry trade returns at time t and the 'three long/three short' momentum (3l/3s(m)), used in Li et al. (2015), consists of going long the three currencies with the largest carry trade returns and short the three currencies with the lowest carry trade returns at time t.[25] The investor realizes the returns at time $t + 1$. All the portfolios (momentum and carry) are rebalanced monthly. While cognizant of the challenges inherent in identifying useful benchmarks for active currency investing discussed by Melvin et al. (2013), we elect to employ the carry trade and momentum portfolios as benchmarks due to their widespread use by academics (Li et al., 2015, among others) and investors alike.

[21] 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 $\widehat{er}^0_{jt+1} = -fp_{jt}$.

[22] 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 GBP, CHF and EUR are statistically larger than zero.

[23] While the 'three long/three short' portfolio is commonly used in existing research (Li et al., 2015), we consider three variants of the forward premium portfolios in light of Bekaert and Panayotov (2016)'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. The differences in the profitability of the three forward premium portfolios that we consider resonates with Bekaert and Panayotov (2016)'s findings.

[24] In 2015 and 2016, SEK and CHF become the two funding currencies.

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

Table 4
Dollar values ($100 initial investment), annualized average returns, standard deviations and Sharpe ratios of trading strategies for individual currencies.

	Payoff in $			Average return			Std. deviation			Sharpe ratio		
	bench	linear	decom	bench	linear	decom	bench	linear	decom	bench	linear	decom
GBP	75.67	91.47	211.85	-1.85	-0.30	6.18	7.55	8.76	8.99	-0.24	-0.03	0.68
CAD	72.11	75.45	111.92	-2.19	-1.71	1.33	7.97	9.51	9.71	-0.27	-0.18	0.13
CHF	92.65	84.11	114.07	0.10	-0.75	3.32	10.17	10.64	10.12	0.01	-0.07	0.32
EUR	115.78	93.89	205.40	1.43	-0.02	6.03	7.94	9.58	9.88	0.18	-0.00	0.61
JPY	110.62	84.83	105.20	1.05	-0.77	0.87	7.52	9.89	9.73	0.14	-0.07	0.08
AUD	115.00	56.23	209.41	1.93	-3.59	6.51	13.01	12.71	12.88	0.14	-0.28	0.50
NZD	125.69	86.15	105.03	2.71	-0.25	1.32	13.76	13.35	13.76	0.19	-0.01	0.09
SEK	86.45	53.96	149.22	-0.66	-4.12	3.67	9.58	11.11	10.93	-0.06	-0.37	0.33
NOK	55.79	65.21	128.77	-3.86	-2.67	2.52	11.00	11.05	10.82	-0.35	-0.24	0.23

Notes: The table presents the dollar payoffs to a $100 initial investment, annualized average returns in %, standard deviations and Sharpe ratios based on the sign of the predicted carry trade return from one of the competing models (benchmark (bench), linear and decomposition (decom)) during the out-of-sample evaluation period: January 2004 – December 2016. Benchmark refers to a strategy of taking a long or short position depending on the sign of the forward premium.

Table 5
Dollar values and descriptive statistics of trading portfolios of carry trade returns.

	value in $	mean	std	skew	kurt	SR
linear	79.95	-1.58	5.28	0.06	1.49	-0.29
decom	155.15	3.53	5.51	1.54	7.10	0.64
Benchmark strategies						
signed fp	96.73	-0.14	4.61	-0.52	1.43	-0.03
1l/1s(fp)	118.18	1.54	7.10	-1.39	4.79	0.21
2l/2s(fp)	128.03	2.06	5.65	-0.74	2.19	0.36
3l/3s(fp)	127.05	2.21	8.63	-0.87	3.07	0.25
1l/1s(m)	83.86	-1.12	6.86	0.66	3.90	-0.16
2l/2s(m)	97.46	-0.07	4.96	0.09	2.02	-0.01
3l/3s(m)	132.00	2.42	7.67	0.24	3.29	0.36

Notes: The table reports the payoffs to a $100 initial investment, the annualized average return in % (mean), standard deviation (std), skewness (skew), kurtosis (kurt) and Sharpe ratio (SR) in an equally-weighted portfolio formed on the basis of the predicted sign from the linear and decomposition (decom) models. The signed forward premium ('signed fp') portfolio benchmark is an equally-weighted portfolio which is long currencies with a negative forward premium and short currencies with a positive forward premium. The 'one-long/one-short' benchmark (1l/1s(fp)) is a long-short portfolio which is long the currency with the largest negative forward premium and short the currency with the largest positive forward premium. The 'two-long/two-short' benchmark (2l/2s(fp)) is a long-short portfolio which is long the two currencies with the largest negative forward premiums and short the two currencies with the largest positive forward premiums. The 'three-long/three-short' benchmark (3l/3s(fp)) is a long-short portfolio which is long the three currencies with the largest negative forward premiums and short the three currencies with the largest positive forward premiums. 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. The 'two-long/two-short' momentum portfolio (2l/2s(m)) is long the two currencies with the highest carry trade returns and short the two currencies with the lowest carry trade returns. The 'three-long/three-short' momentum portfolio (3l/3s(m)) is long the three currencies with the highest carry trade returns and short the three currencies with the lowest carry trade returns.

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 are reported in Table 5. For comparability with the existing literature, Table 5 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.[26] The decomposition model's portfolio returns also exhibit a lower volatility than the 'three long/three short' carry and momentum portfolios. In addition, the decomposition model's skewness is positive and larger than that of any of the benchmark portfolios. These latter observations suggest that the decomposition model's portfolio might be less prone to crash risk than these portfolios. The decomposition model's portfolio also outperforms any of the competing benchmarks on a risk-adjusted return basis (i.e. in terms of Sharpe ratio). Furthermore, with the exception of GBP, the risk-adjusted returns of the decomposition model's 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.[27] We view our results as suggesting that a portfolio strategy might be

[26] The Sharpe ratio for the decomposition method 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.

[27] As pointed out by Brunnermeier et al. (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 et al. (2012a) 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.

less adversely affected than individual currencies by the unwinding of the carry trade.

The statistical and trading success of the decomposition model appears to be stemming from two factors. First, the decomposition model models the entire conditional distribution of currency returns and incorporates, in a parsimonious and flexible manner, inherent nonlinearities in the exchange rate dynamics that may prove particularly important during periods of financial turbulence. Second and relatedly, by modeling absolute returns directly, which proxy for volatility, the decomposition model succeeds in avoiding large negative losses. In this respect, we interpret the success of the decomposition as being consistent with the findings of Cenedese et al. (2014) who provide, using predictive quantile regressions, compelling empirical evidence which suggests that higher variance relates to future carry losses. Furthermore, our decomposition of currency returns relates analytically to the approach of Cenedese et al. (2014). In fact, while we decompose currency returns into absolute and sign components, Cenedese et al. (2014) decompose market variance into average variance and average correlation and show that the predictive power for future carry trade returns arises from average variance. The authors also show that conditioning on market variance results in net-of-transaction costs trading gains. In a similar vein, our findings suggest that conditioning on absolute returns yields statistical and economic gains. Finally, the decomposition model delivers substantially improved directional forecasts which play a crucial role in the success of the model-based trading strategies.

5. 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 much greater predictability than raw returns. We 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 out-of-sample predictability. 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. We view the decomposition model's success as arising from its improved directional forecasting and its ability to avoid large losses by adequately capturing the tail of the distribution of future currency excess returns. 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 et al. (2017).

References

Accominotti, O., and Chambers, D. 2014. Out-of-Sample Evidence on the Returns to Currency Trading, Working Paper, London School of Economics.
Akram, Q.F., Rime, D., Sarno, L., 2008. Arbitrage in the foreign exchange market: turning on the microscope. J. Int. Econ. 76, 237–253.
Anatolyev, S., Gospodinov, N., 2010. Modeling financial return dynamics via decomposition. J. Bus. Econ. Stat. 28, 232–245.
Ang, A., and Chen, J.S. 2010. Yield Curve Predictors of Foreign Exchange Returns, Working Paper, Columbia University.
Baba, N., Packer, F., 2009. Interpreting deviations from covered interest parity during the financial market Turmoil of 2007–08. J. Bank. Financ. 33, 1953–1962.
Bacchetta, P., van Wincoop, E., 2006. Can information heterogeneity explain the exchange rate determination puzzle? Am. Econ. Rev. 96, 552–576.
Bakshi, G., Panayotov, G., 2013. Predictability of currency carry trades and asset pricing implications. J. Financ. Econ. 110, 139–163.
Berge, T. J., Ò. Jordà, and A. M. Taylor, 2010, Currency Carry Trades, NBER Working Paper No. 16491.
Bekaert, G., Panayotov, G., 2016. Good Carry, Bad Carry, Columbia Business School Research Paper No. 15-53. Columbia University.
Borio, C., McCauley, R., McGuire, P., Sushko, V. 2016. Covered Interest Parity Lost: Understanding the Cross-Currency Basis, *BIS Quarterly Review* (September), 45–64.
Brunnermeier, M.K., Nagel, S., Pedersen, L.H., 2009. Carry trades and currency crashes. NBER Macroecon. Annu. 23, 313–347.
Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2011. Do Peso problems explain the returns to the carry trade? Rev. Financ. Stud. 24, 853–891.
Burnside, C., Eichenbaum, M., Rebelo, S., 2011. Carry trade and momentum in currency markets. Annu. Rev. Financ. Econ. 3, 511–535.
Campbell, J.Y., Thompson, S.B., 2008. Predicting excess stock returns out of sample: can anything beat the historical average? Rev. Financ. Stud. 21, 1509–1531.
Cenedese, G., Sarno, L., Tsiakas, I., 2014. Foreign exchange risk and the predictability of carry trade returns. J. Bank. Financ. 42, 302–313.
Cheung, Y.-W., Chinn, M.D., Pascual, A.G., 2005. Empirical exchange rate models of the nineties: are any fit to survive? J. Int. Money Financ. 24, 1150–1175.
Chinn, M.D., Meese, R.A., 1995. Banking on currency forecasts: how predictable is change in money? J. Int. Econ. 38, 161–178.
Clarida, R.H., Sarno, L., Taylor, M.P., Valente, G., 2003. The out-of-sample success of term structure models as exchange rate predictors: a step beyond. J. Int. Econ. 60, 61–83.
Daniel, K., Hodrick, R.J., and Lu, Z. 2017. The Carry Trade: Risks and Drawdowns, Critical Finance Review, forthcoming.
Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. Rev. Financ. Stud. 22, 3491–3530.
Della Corte, P., Tsiakas, I. 2012. Statistical and Economic Methods for Evaluating Exchange Rate Predictability, In: Handbook of Exchange Rates (J. James, I. Marsh and L. Sarno, eds.), 221–264, John Wiley & Sons, Inc., Hoboken, NJ, USA.
Du, W., Tepper, A., Verdelhan, A. 2016. Deviations from Covered Interest Rate Parity, Working Paper, Federal Reserve Board.
Engel, C. 2015. Exchange Rates and Interest Parity. In: Handbook of International Economics (E. Helpman, K. Rogoff and G. Gopinath, eds.), Vol. 4, 453–522.
Engel, C., Mark, N.C., West, K.D., 2007. Exchange rate models are not as bad as you think. NBER Macroecon. Annu. 22, 381–441.
Engel, C., West, K.D., 2005. Exchange rates and fundamentals. J. Political Econ. 113, 485–517.
Engle, R.F., 2002. New frontiers for ARCH models. J. Appl. Econ. 17, 425–446.
Farhi, E., Fraiberger, S.P. Gabaix, X. Ranciere, R. Verdelhan, A. 2009, Crash Risk in Currency Markets, NBER Working Paper No. 15062.
Ferraro, D., Rogoff, K., Rossi, B., 2015. Can oil prices forecast exchange rates? J. Int. Money Financ. 54, 116–141.
Galati, G., Heath, A., and McGuire, P. 2007. Evidence of Carry Trade Activity, *BIS Quarterly Review* (September), 27–41.
Giacomini, R., White, H., 2006. Tests of conditional predictive ability. Econometrica 74, 1545–1578.
Gospodinov, N., 2009, A New Look at the Forward Premium Puzzle, Journal of Financial Econometrics 7, 312–338.
Hai, W., Mark, N.C., Wu, Y., 1997. Understanding spot and forward exchange rate regressions. J. Appl. Econ. 12, 715–734.
Hassan, T.A., and Mano, R.C. 2014. Forward and Spot Exchange Rates in a Multi-Currency World, NBER Working Paper No. 20294.
Jordà, Ò., Taylor, A.M., 2012. The carry trade and fundamentals: nothing to fear but FEER itself. J. Int. Econ. 88, 74–90.
Kilian, L., Taylor, M.P., 2003. Why is it so difficult to beat the random walk forecast of exchange rates? J. Int. Econ. 60, 85–107.
Koijen, R., Moskowitz, T., Pedersen, L.H., Vrugt, E. 2017. Carry, Journal of Financial Economics, forthcoming.

Li, J., Tsiakas, I., Wang, W., 2015. Predicting exchange rates out of sample: can economic fundamentals beat the random walk? J. Financ. Econ. 13, 293–341.
Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. Rev. Financ. Stud. 24, 3731–3777.
Mancini, L., Ranaldo, A., Wrampelmeyer, J., 2013. Liquidity in the foreign exchange market: measurement, commonality, and risk premiums. J. Financ. 68, 1805–1841.
Mancini Griffoli, T., Ranaldo, A. 2011, Limits to Arbitrage during the Crisis: Funding Liquidity Constraints and Covered Interest Parity, Working Paper, Swiss National Bank.
Manner, H., Reznikova, O., 2012. A survey on time-varying copulas: specification, simulations, and application. Econ. Rev. 31, 654–687.
Mark, N.C., 1995. Exchange rates and fundamentals: evidence on long-horizon predictability. Am. Econ. Rev. 85, 201–218.
Melvin, M., Prins, J., Shand, D., 2013. Forecasting Exchange Rates: An Investor Perspective. In: Handbook of Economic Forecasting 2B (G. Elliott and A. Timmermann, eds.), 721–750.
Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? J. Int. Econ. 14, 3–24.
Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012a. Carry trades and global foreign exchange volatility. J. Financ. 67, 681–718.
Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. J. Financ. Econ. 106, 660–684.
Molodtsova, T., Papell, D.H., 2009. Out-of-sample exchange rate predictability with taylor rule fundamentals. J. Int. Econ. 77, 167–180.
Sarno, L., Sojli, E., 2009. The feeble link between exchange rates and fundamentals: can we blame the discount factor? J. Money, Credit Bank. 41, 437–442.
Sarno, L., Valente, G., 2009. Exchange rates and fundamentals: footloose or evolving relationship? J. Eur. Econ. Assoc. 7, 786–830.
Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. Rev. Financ. Stud. 21, 1455–1508.