Journal of Financial Economics 110 (2013) 139–163

Contents lists available at ScienceDirect

Journal of Financial Economics

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

Predictability of currency carry trades and asset pricing implications[☆]

Gurdip Bakshi[a,1], George Panayotov[b,*]

[a] Smith School of Business, University of Maryland, College Park, MD 20742, USA
[b] School of Business and Management, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

ARTICLE INFO

Article history:
Received 7 August 2011
Received in revised form
14 November 2011
Accepted 16 December 2011
Available online 30 April 2013

JEL classifications:
C23
C53
G11
G12
G13
C5

Keywords:
Currency carry trades
Predictability
Commodity returns
Currency volatility
Liquidity
Predictability-based decision rule
Currency-related risk factors

ABSTRACT

This paper studies the time series predictability of currency carry trades, constructed by selecting currencies to be bought or sold against the US dollar, based on forward discounts. Changes in a commodity index, currency volatility and, to a lesser extent, a measure of liquidity predict in-sample the payoffs of dynamically re-balanced carry trades, as evidenced by individual and joint *p*-values in monthly predictive regressions at horizons up to six months. Predictability is further supported through out-of-sample metrics, and a predictability-based decision rule produces sizable improvements in the Sharpe ratios and skewness profile of carry trade payoffs. Our evidence also indicates that predictability can be traced to the long legs of the carry trades and their currency components. We test the theoretical restrictions that an asset pricing model, with average currency returns and the mimicking portfolio for the innovations in currency volatility as risk factors, imposes on the coefficients in predictive regressions.

1. Introduction

Currency carry trade strategies involve borrowing in countries with low interest rates and investing in the currencies of countries that offer high interest rates. Such trading strategies have been pursued by practitioners, and their existence could be traced to exploitable disparities in global macroeconomic conditions.

> The economic malaise of the EU and Japan has central bankers and bondholders worried. But it has currency speculators and arbitragers licking their lips at the prospect of riding the euro and yen on their downward slide, while at the same time making fat profits on

☆ The feedback and advise of Robert Hodrick (the referee) and G. William Schwert are gratefully acknowledged. An earlier version of the paper was presented at Georgetown University, Hong Kong University, and Syracuse University. We welcome comments, including references to related papers we have inadvertently overlooked. We acknowledge helpful discussions with Turan Bali, Fousseni Chabi-Yo, Xiaohui Gao, Nikolay Gospodinov, Jeff Harris, Grace Hu,X Pete Kyle, Dilip Madan, Georgios Skoulakis, David Weinbaum, Fan Yang, and Yildiray Yildirim. Any remaining errors are our responsibility alone.

* Corresponding author. Tel.: +1 607 342 0970.
E-mail address: gbakshi@rhsmith.umd.edu (G. Bakshi).
[1] Tel.: +1 301 405 2261.

0304-405X/$ - see front matter
http://dx.doi.org/10.1016/j.jfineco.2013.04.010

> high-yield investments and rising currencies in emerging markets…. This sophisticated bit of financial wizardry is known as a carry trade, and it's a popular way to offset moves in equity and bond markets while also producing a handsome return. (*Time*, January 26, 2011)

Even if potentially lucrative, carry trades also present substantial downside risks, as witnessed, for example, by the strong appreciation of the yen in the wake of the Japanese earthquake in March 2011.

While advances have been made in understanding the role of peso effects in carry trade payoffs (e.g., Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011), and in identifying risk factors that determine average currency returns (e.g., Lustig, Roussanov, and Verdelhan, 2011; Ang and Chen, 2010; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012), questions that still warrant reconciliation are: Do carry trade payoffs exhibit time series predictability? Is such predictability statistically and economically significant? If so, can one link the identified predictors to some asset pricing model? Our innovation is the focus on time series predictability, as opposed to the cross section of currency portfolio returns, and we adopt as criteria the parameter significance in predictive regressions, out-of-sample predictability, and increase in Sharpe ratios and skewness of carry trade payoffs, achieved by exploiting a predictability-based decision rule.

We build time series of monthly carry trade payoffs, using spot and forward exchange rates of G-10 currencies and accounting for bid–ask spreads, and examine the profitability of carry trade strategies from two perspectives. First, we consider carry trades with fixed currency pairs, which borrow in a currency with traditionally low interest rate (e.g., Japanese yen), and then invest in a currency with high interest rate (e.g., New Zealand dollar) and, hence, are not necessarily anchored to funding in US dollars. Second, we follow the construction of investable carry trade indexes and dynamically rank-order currencies according to their interest rate differentials, as embedded in the forward discount, whereby the lowest- (highest-) yielding currencies are selected to be sold (bought).

Our approach elicits several findings about the payoffs of carry trade strategies. First, only the dynamically rebalanced carry trade strategies show evidence of profitability. In contrast to the fixed-pair strategies, with annualized Sharpe ratios not exceeding 0.10, certain dynamic carry trade strategies can generate Sharpe ratios as high as 0.50 in our sample and exhibit statistically significant average payoffs.

Next, we show that monthly payoffs of dynamic carry trades can be predicted in-sample, using as predictors changes in a commodity price index, changes in average currency volatility, and a variable that captures global liquidity. We establish both marginal and joint predictive power, and we show robustness using p-values obtained with the covariance estimators of Newey and West (1987) and Hodrick (1992), and via parametric bootstrap (e.g., Mark, 1995; Kilian, 1999; Amihud, Hurvich, and Wang, 2009). The coefficients in the predictive regressions are positive for commodity returns and negative for currency volatility. Moreover, decreasing liquidity is associated with lower future carry trade payoffs. A notable feature of our predictors is that they are not highly correlated among each other, which highlights the differences in their economic nature.

Our interest in a commodity-based predictor is motivated by the analysis in Chen and Rogoff (2003) and Chen, Rogoff, and Rossi (2010) and also by anecdotal evidence that commodity investing often coincides with an appetite for risk-taking. In considering average currency volatility as a predictor, we are guided by the perception, as in Bhansali (2008), that the profitability of carry trades appears to decline in volatile currency markets. In turn, the potential relevance of global liquidity is implied, among others, by the studies of Brunnermeier and Pedersen (2009) and Asness, Moskowitz, and Pedersen (2009).

The predictive ability is preserved in out-of-sample tests, as evidenced by the consistently positive values of the out-of-sample R^2 statistic of Campbell and Thompson (2008) and the low p-values associated with the MSPE-adjusted statistic of Clark and West (2007). In addition, combination forecasts, as in Stock and Watson (2004), yield p-values below 0.05 for the MSPE-adjusted statistic, strengthening the evidence for predictability. Furthermore, conditional carry trade strategies that exploit trading signals, generated using our predictors, enhance Sharpe ratios and mitigate negative skewness, relative to the unconditional carry trades. Complementing these results, the non-parametric market timing test of Henriksson and Merton (1981) indicates statistically significant timing ability based on our predictors.

Moreover, we investigate the robustness of our results using other predictors, in particular, the term structure variables found to be important in the context of Ang and Chen (2010), as well as the change in VIX, as in Brunnermeier, Nagel, and Pedersen (2009). The presence of these additional variables does not appear to diminish the predictive ability of our predictors.

We also find that focusing on distinct components of carry trade payoffs is informative about the nature of carry trades. In particular, we find that carry trade payoffs inherit their predictability from the long, and not the short, legs of the trades, but it is also seen that combining the short and long legs strengthens the evidence for predictability. Further, we observe that it is the currency component of payoffs that is captured by our predictors. We verify that these features appear to be specific to carry trade payoffs and do not pass on to the payoffs of the fixed currency pairs in our sample.

Building on the evidence, we perform GMM tests that suggest that the forecasting ability of our predictors is not inconsistent with a latent-variable model, as developed by Hansen and Hodrick (1983). We further examine whether the predictability of carry trade payoffs can be reconciled with an explicit asset pricing model and observable risk factors. Here we follow the approach developed in Kirby (1998) and assess the restrictions imposed by a stochastic discount factor model on the coefficients, obtained in predictive regressions. As prescribed in this approach, we set up a GMM system, allowing to statistically compare model-implied predictive coefficients with the unrestricted ones, and adopt a model with currency excess returns and innovations in currency volatility as risk factors (Lustig, Roussanov, and Verdelhan, 2011; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012). This exercise indicates a lack of consistency of the theoretical model with the documented predictability.

While we share with others the objective of studying carry trade strategies, our study differs in key respects. Departing from Jordá and Taylor (2009), we demonstrate from a different econometric perspective the predictive ability of certain observable macroeconomic variables, and we examine whether recent currency-related pricing models can reproduce the coefficients from predictive regressions. Whereas the evidence in Brunnermeier, Nagel, and Pedersen (2009) comes from dollar-based carry trade strategies, our objects of interest are dynamically re-balanced strategies. However, the studies of Ang and Chen (2010), Lustig, Roussanov, and Verdelhan (2011), and Menkhoff, Sarno, Schmeling, and Schrimpf (2012) put the thrust on the cross-sectional, rather than the time series return dimension, and on identifying common risk factors for currency returns. Furthermore, unlike Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011) or Jurek (2009), we do not focus on the role of the peso effect or negative skewness for explaining carry trade payoffs, but on predictability, and we contribute by establishing that conditional strategies using our predictors improve the skewness profile.

This paper is organized as follows. Section 2 describes the construction of the dynamic carry trade strategies while accounting for transaction costs. Section 3 examines the predictability of carry trades. Whether predictability can be traced to the short or long legs of carry trades or to the currency component of the trades is analyzed in Section 4. Here we also elaborate on the merits of our predictors. Section 5 provides statistical tests of a latent-variable model of carry trade payoffs and the testable restrictions that some factor models impose on predictive coefficients. Concluding remarks are in Section 6.

2. Description of currency carry trades

Throughout, we refer to currencies with low (high) interest rate as funding (investment) currencies. For the purpose of computing the payoff to a carry trade, one must choose a perspective, and we take it to be that of the US investor. In our setting, it means that the investment currencies are bought with US dollars, and the funding currencies are sold short against the US dollar, unless the US dollar itself qualifies to be an investment or funding currency.

In our effort to follow the implementation of the carry trade in an actual trading environment, we cast the trade in terms of forward currency contracts, as articulated previously, for example, in Brunnermeier, Nagel, and Pedersen (2009), Jurek (2009), Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011), and Lustig, Roussanov, and Verdelhan (2011). One advantage of using forward contracts is that they facilitate incorporating transaction costs (Bessembinder, 1994).

Most of our data comes from Barclays Capital (source: Datastream), which offers bid, ask, and mid-price quotes, all on a daily basis. Whereas quotes are available on forward contracts with maturities up to 12 months, we use only spot and one-month forward quotes and calculate one-month carry trade payoffs. This choice of frequency is dictated by considerations about sample length. Our sample starts in January 1985, ends in August 2011, and encompasses about 27 years of currency market history (320 observations).

Focusing on the most liquid currencies, we use data on spot and forward contracts of the G-10 currencies (excluding the euro, for its shorter history) quoted against the US dollar and omit currencies of emerging market countries. The currencies used are the Australian dollar, Canadian dollar, Swiss franc, British pound, Japanese yen, Norwegian krone, New Zealand dollar, and Swedish krona.[2] Several of the G-10 currencies in our sample are considered commodity currencies, and the relevance of this observation, following Chen and Rogoff (2003), will be discussed shortly.

We emphasize that a carry trade can be designed in various ways, distinguished, for example, by the number of currencies bought and sold, weight given to each currency, criteria for including a currency in the trade, re-balancing frequency, and treatment of leverage. In what follows, we keep to a design with equally weighted positions in one to four long and short G-10 currencies, with no leverage. The positions are re-balanced on a monthly basis, and currencies are included based only on their interest rate differentials as inferred through the forward discounts.

Specifically, for each month in the sample we construct payoffs, scaled to one dollar, to four trading strategies, in which the first strategy buys the highest-yielding among the G-10 currencies and sells the lowest-yielding counterpart, while the second, third, and fourth strategies buy the two, three and four highest-yielding currencies, respectively, and sell the two, three and four lowest-yielding ones.

To provide more detail, consider the trading strategy with one long and short pair of currencies, and denote by S_{t+1}^{lowest} (S_{t+1}^{highest}) the spot exchange rate per one US dollar at the end of month t+1 of the currency with the lowest (highest) interest rate at the end of month t and by $F_{t,t+1}^{\text{lowest}}$ ($F_{t,t+1}^{\text{highest}}$) the corresponding one-month forward exchange rate for delivery at the end of month t+1. Then the payoffs

[2] We consider a smaller set of currencies than in some other studies. Note, first, that the G-10 currencies subsume a predominant portion of investable carry trade indexes. Second, some emerging markets exhibit occasional violations of the covered parity condition in our sample. Our work is related to a strand of studies, including Backus, Foresi, and Telmer (2001), Lustig and Verdelhan (2007), Levich and Poti (2008), Darvas (2009), Beber, Buraschi, and Breedon (2010), Campbell, Medeiros, and Viceira (2010), Ranaldo and Soderlind (2010), Christiansen, Ranaldo, and Soderllind (2011), Burnside, Han, Hirshleifer, and Wang (2011), Farhi, Fraiberger, Gabaix, and Ranciere (2011), Jylha and Suominen (2011), Heyerdahl-Larsen (2011), and Yu (2011). We also added the euro to the set of G-10 currencies employed. The predictive power of our predictors was maintained, and these results are available upon request.

over month $t+1$ to the short and long component of the trade, respectively, and to the composite trade, all denominated in US dollars, are

$$z_{t+1}^{\text{short}_1} = 1 - \frac{F_{t,t+1}^{\text{lowest (ask)}}}{S_{t+1}^{\text{lowest (bid)}}}, \quad z_{t+1}^{\text{long}_1} = \frac{F_{t,t+1}^{\text{highest (bid)}}}{S_{t+1}^{\text{highest (ask)}}} - 1$$

$$\text{and} \quad z_{t+1}^{(1)} \equiv \frac{z_{t+1}^{\text{short}_1} + z_{t+1}^{\text{long}_1}}{2}. \tag{1}$$

The expression for $z_{t+1}^{(1)}$ could be seen as the payoff generated from betting half a dollar on each leg of the trade.

In general, when K currency pairs and, hence, K long and short positions are included in a carry trade, whereby K can be 1, 2, 3, or 4 in our context, we modify Eq. (1), and compute the payoff as

$$\overline{z}_{t+1}^{(K)} = \frac{1}{K} \sum_{k=1}^{K} z_{t+1}^{(k)}, \quad K = 1, \ldots, 4 \quad \text{with } z_{t+1}^{(k)} = \frac{z_{t+1}^{\text{short}_k} + z_{t+1}^{\text{long}_k}}{2}, \tag{2}$$

where $z_{t+1}^{\text{short}_k}$ ($z_{t+1}^{\text{long}_k}$) is the payoff over month $t+1$ to the short (long) position established at the end of month t in the k-th lowest-yielding (highest-yielding) currency.

When it happens that the US dollar is among, say, the lowest-yielding currencies included in a carry trade in a given month, then we do not short any other currency against it and adjust the weights accordingly, to maintain a zero-cost strategy. An analogous approach to constructing carry trades has been employed, among others, in Clarida, Davis, and Pedersen (2009), Melvin and Taylor (2009), and Ang and Chen (2010).[3] In the empirical work, we use two distinct vectors $\overline{\mathbf{z}}_{t+1} \equiv [\overline{z}_{t+1}^{(1)}\ \overline{z}_{t+1}^{(2)}\ \overline{z}_{t+1}^{(3)}\ \overline{z}_{t+1}^{(4)}]'$ and $\mathbf{z}_{t+1} \equiv [z_{t+1}^{(1)}\ z_{t+1}^{(2)}\ z_{t+1}^{(3)}\ z_{t+1}^{(4)}]'$.

We recognize that an alternative approach to studying carry trades can be followed, as, for example, in Brunnermeier, Nagel, and Pedersen (2009, Table 1), which entails individual currencies bought with US dollars. In contrast, our focus is on the predictability of dynamically re-balanced carry trades that are strictly based on forward discounts and can be constituted as zero-cost trading strategies and on statistical tests that assess whether predictability can be related to some asset pricing framework.

To illustrate the relative importance of the US dollar and other currencies in carry trades, Table A1 reports the number of occasions (end-of-month) when the respective currency has been the highest-yielding (lowest-yielding) or among the two, three, or four highest-yielding (lowest-yielding) currencies. Relative yields that are at the heart of this exercise are implied from the mid-quotes for spot and one-month forward contracts, i.e., from the forward discounts. Important to our approach, it is seen, for example, that the US dollar has been among the three lowest-yielding currencies about half the time, whereas the Japanese yen and the Swiss franc have been in this category almost every month in our sample. Providing an analogous distinction, the three highest-yielding currencies have consistently included the New Zealand and Australian dollars, while the US dollar appears in this group in 43 out of 320 months. The choice of currencies to be included in the long or short leg of a carry trade is not innocuous and could influence the profitability of carry trades and whether they display predictable variation.

3. Predictability of carry trade strategies

In this section we first provide evidence for the profitability of carry trades and demonstrate some specific features of carry trade payoffs that are not exhibited by the fixed currency pairs, which are most often involved in carry trades. Next, we turn to the predictability of carry trade payoffs, and our thrust is on investigating in-sample predictability, out-of-sample predictability, and whether a predictability-based decision rule can improve the Sharpe ratios and lower the negative skewness of carry trade payoffs, and then we focus on evaluating the relevance of combination forecasts. The importance of the considered issues stems from their relation to the theoretical puzzle that carry trade profitability has posed, whereby the predictability perspective could provide new insights.

3.1. *Only dynamically re-balanced carry trades appear to be profitable*

How profitable are carry trade strategies? Is there a risk-reward rationale for pursuing such strategies? Table 1 summarizes notable features of the payoffs to the four carry trade strategies and illustrates the profitability that underlies the existence of the carry trade as a popular trading strategy.

The annualized average payoffs range between 1.95% and 2.70%, with 95% confidence intervals in curly brackets. The confidence intervals are based on the stationary bootstrap (Politis and Romano, 1994), with optimal block size determined according to the algorithm of Politis and White (2004), which accounts for conditional heteroskedasticity (Goncalves and White, 2002), as well as non-normality. As seen from Table 1, the confidence intervals for the mean include zero for carry trade strategies 1 and 2, indicating statistical insignificance, while strategies 3 and 4 exhibit significantly positive average payoffs.

Another salient attribute of carry trade payoffs is their skewness, which is negative in all cases, ranging between −0.31 and −1.27. Importantly, the negative skewness gets less prominent for strategies involving larger number of currencies, in line with Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011, Table 2), who find negative, but not statistically significant, skewness for payoffs of carry trade portfolios. Fig. 1 illustrates this data feature, showing pronounced down-moves on several occasions for the $\overline{z}_t^{(1)}$ series, but visibly fewer and smaller down-moves for the $\overline{z}_t^{(4)}$ series.

While the average payoffs to carry trades do not appear to be impressive, their volatility is relatively low. The resulting annualized Sharpe ratios range between 0.25

[3] Weights in the carry trades could also be chosen based on optimization of certain risk-return characteristics, as is done, for example, in investable indexes such as Barclays Intelligent Carry Index or Credit Suisse's Rolling Optimized Carry Indices. As in our approach, JP Morgan's IncomeFX and IncomeEM and Deutsche Bank's Harvest investable indexes invoke equal weighting.

Table 1
Profitability of carry trades.
The payoffs over month t+1 of individual short and long positions within a carry trade are

$$z_{t+1}^{\text{short}_k} = 1 - \frac{F_{t,t+1}^{k\text{-th lowest (ask)}}}{S_{t+1}^{k\text{-th lowest (bid)}}} \quad \text{and} \quad z_{t+1}^{\text{long}_k} = \frac{F_{t,t+1}^{k\text{-th highest (bid)}}}{S_{t+1}^{k\text{-th highest (ask)}}} - 1,$$

where $S_t^{k\text{-th lowest}}$ ($S_t^{k\text{-th highest}}$) denote the spot exchange rate per one US dollar at the end of month t of the currency with the k-th lowest (k-th highest) interest rate (based on the forward discount), $F_{t,t+1}^{k\text{-th lowest}}$ ($F_{t,t+1}^{k\text{-th highest}}$) is the corresponding forward exchange rate for delivery at the end of month t+1. All payoffs are denominated in US dollars, and we consider the G-10 currencies (except the euro). The payoff over month t+1 of a carry trade with K long and short positions is, then

$$\overline{z}_{t+1}^{(K)} = \frac{1}{K}\sum_{k=1}^{K} z_{t+1}^{(k)}, \quad K = 1, \ldots, 4 \quad \text{with } z_{t+1}^{(k)} = \frac{z_{t+1}^{\text{short}_k} + z_{t+1}^{\text{long}_k}}{2}.$$

When the US dollar is among, say, the lowest-yielding currencies in a given month, no currency is shorted against it, and the weights are adjusted to be consistent with a zero-cost strategy. For each payoff series $\overline{z}_{t+1}^{(K)}$ indexed by the number K of long and short positions, we report the annualized mean, standard deviation and Sharpe ratio, monthly skewness and kurtosis, and first-order autocorrelation (ACF_1). We also report the 95% confidence intervals (denoted by C.I) for the mean payoff based on the stationary bootstrap (Politis and Romano, 1994), where the block size is based on the algorithm of Politis and White (2004). Similar statistics are also reported for six fixed currency pairs that are prominent in carry trade strategies (see Table A1), and the corresponding equally weighted portfolio. For example, JPY–AUD denotes a strategy where each month a US investor goes short the Japanese yen (the funding currency) and long the Australian dollar (the investment currency). The sample period is January 1985 to August 2011, with 320 observations.

Carry strategy	Mean (annual)	Stdev (annual)	Sharpe ratio (annual)	Stationary bootstrap 95% C.I for the mean	Skewness (monthly)	Kurtosis (monthly)	ACF_1
1	2.15	8.72	0.25	{−1.39 5.41}	−1.27	7.76	0.07
2	1.95	6.36	0.31	{−0.66 4.43}	−1.11	5.53	0.12
3	2.70	5.36	0.50	{0.44 4.76}	−0.79	4.66	0.11
4	2.21	4.48	0.49	{0.52 3.91}	−0.31	4.32	0.05
JPY–AUD	0.00	7.84	0.00	{−3.14 3.12}	−0.97	5.41	0.06
JPY–NZD	0.75	7.46	0.10	{−2.10 3.58}	−0.67	4.96	0.01
JPY–GBP	−0.13	6.30	−0.02	{−2.74 2.36}	−1.05	6.29	0.12
CHF-AUD	−0.57	7.20	−0.08	{−3.32 2.26}	−0.63	4.23	0.04
CHF-NZD	0.17	6.61	0.03	{−2.35 2.69}	−0.43	4.00	−0.01
CHF-GBP	−0.71	4.75	−0.15	{−2.49 1.09}	−1.13	8.95	0.05
Equally weighted	−0.08	5.16	−0.02	{−2.17 1.90}	−0.89	5.06	0.08

and 0.50, similar to those for the high-minus-low strategies with transaction costs, reported for developed countries in Lustig, Roussanov, and Verdelhan (2011, Table I) but lower than that for the equally weighted portfolios in Burnside, Eichenbaum, Kleshchelski, and Rebelo (2011, Table 2). Importantly, we observe that strategies that employ more currencies exhibit lower volatility, higher Sharpe ratios, and statistically significant average payoffs.

How do carry trades fare relative to other investment alternatives? Fig. 2 offers one perspective and plots the cumulative payoffs, generated by a $1 notional investment in two carry trade strategies in January 1985. These payoffs are compared with cumulative excess returns over the US risk-free rate $\Pi_{t=1}^{T}(1+er_t)$ from investments (all measured in US dollars) in (i) currencies (average return of the G-10 currencies held long against the US dollar), (ii) commodities (Raw Industrials subindex of the CRB commodity index), (iii) bonds (average total return of long-term bonds in our country sample), and (iv) equities (total return of the MSCI World equity index). The annualized Sharpe ratios for currencies, commodities, bonds, and equities are 0.35, −0.11, 0.72, and 0.33, respectively, providing additional indication for the desirability of carry trades within the universe of possible investments in our sample.

To further highlight the specifics of the payoffs to the dynamically re-balanced carry trades, Table 1 provides analogous statistics for strategies based on several fixed currency pairs. The pairs are chosen to involve the G-10 currencies that are most often among the lowest or highest-yielding in our sample, with the Japanese yen and Swiss franc being sold and the Australian dollar, New Zealand dollar, and British pound being bought against the US dollar. Comparing the payoffs to our carry trade strategies, the fixed pairs exhibit inferior average payoffs (at most 0.75%), insignificant in all cases; far lower annualized Sharpe ratios (at most 0.10), as also has been noted in Jordá and Taylor (2009, Table 2); and negative skewness, which is comparable to that of the dynamically re-balanced carry trades. In addition, the diversification achieved in an equally weighted portfolio of the six fixed currency pairs (as shown in the last line of Table 1) does not result in a markedly improved Sharpe ratio or less negative skewness. The fixed carry trade pairs are largely unprofitable.

One conclusion from our comparison with fixed currency pairs could be that the specific design of the carry trade can have a substantial impact on average payoffs but does not alter in an essential way the remaining statistical properties. It could also indicate that differences between results obtained in the literature could be largely due to the different ways in which the carry trade payoffs were constructed. Our approach strives to reflect the practical implementation of carry trades, in which the currency pairs to be included are chosen based on the ranking of interest rate differentials, as extracted from forward discounts.

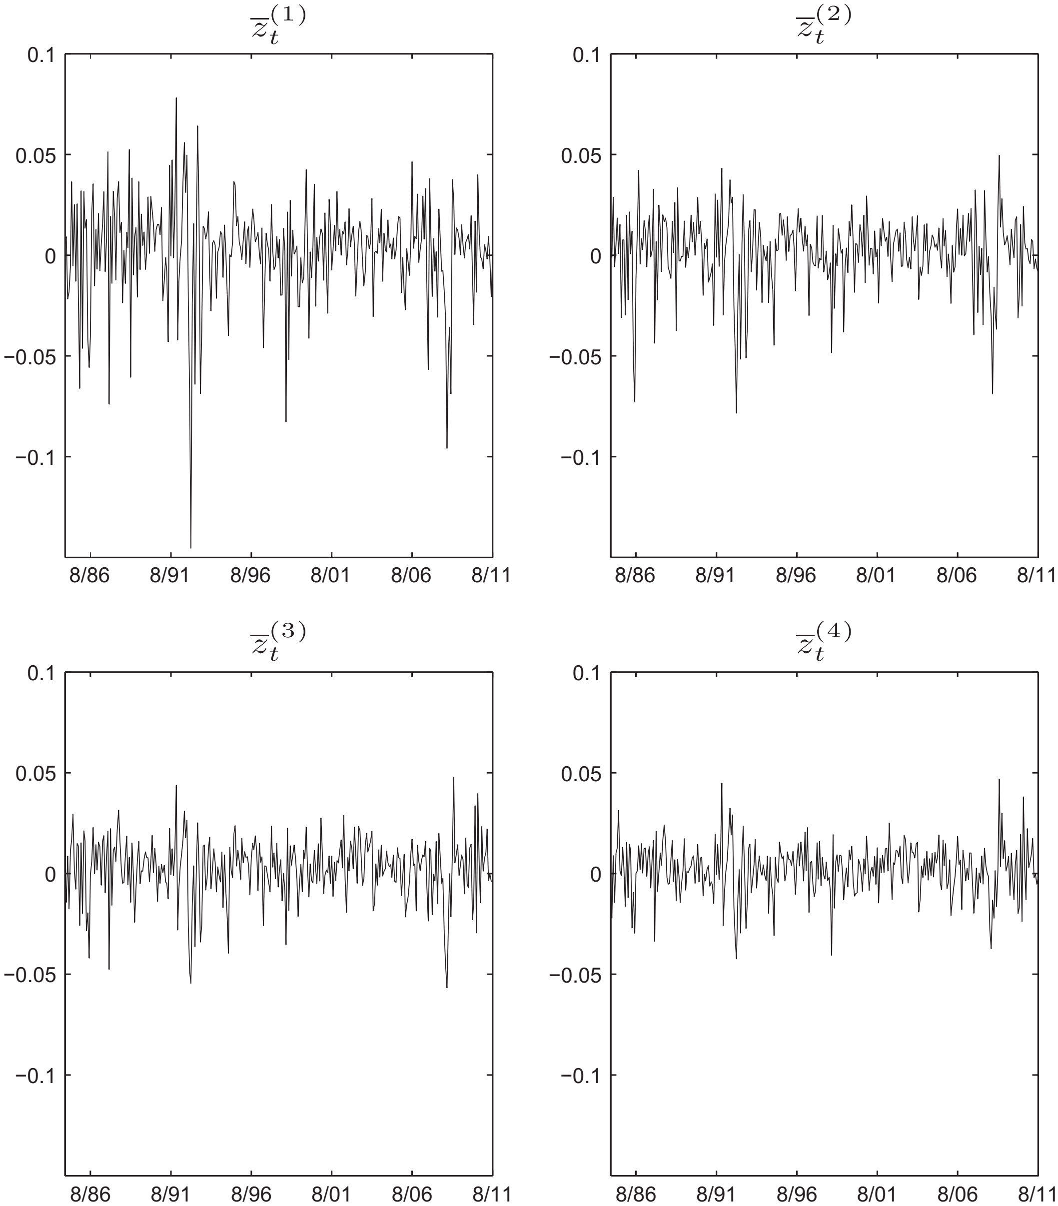

Fig. 1. Time variation in the payoff of carry trades. For each month over the sample period 01/1985 to 08/2011 (320 observations), we first compute the payoffs over month t+1 of individual short and long positions within a carry trade as

$$z_{t+1}^{\text{short}_k} = 1 - \frac{F_{t,t+1}^{k\text{-th lowest (ask)}}}{S_{t+1}^{k\text{-th lowest (bid)}}} \quad \text{and} \quad z_{t+1}^{\text{long}_k} = \frac{F_{t,t+1}^{k\text{-th highest (bid)}}}{S_{t+1}^{k\text{-th highest (ask)}}} - 1,$$

where $S_{t+1}^{k\text{-th lowest}}$ ($S_{t+1}^{k\text{-th highest}}$) denote the spot exchange rate per one US dollar at the end of month t+1 of the currency with the k-th lowest (k-th highest) interest rate at the end of month t, $F_{t,t+1}^{k\text{-th lowest}}$ ($F_{t,t+1}^{k\text{-th highest}}$) is the corresponding one-month forward exchange rate for delivery at the end of month t+1. All payoffs are denominated in US dollars, and we consider the G-10 currencies (except the euro). Then denoting $z_{t+1}^{(k)} = (z_{t+1}^{\text{short}_k} + z_{t+1}^{\text{long}_k})/2$, the payoff over month t+1 of a carry trade with K long and short positions is $\overline{z}_{t+1}^{(K)} = (1/K)\sum_{k=1}^{K} z_{t+1}^{(k)}$, for $K = 1, \ldots, 4$.

3.2. Motivation for the choice of predictors

We address the following questions: Are carry trade payoffs predictable? What macroeconomic determinants could be underlying such predictability? The focus on carry trade predictability distinguishes our study from work that considers the possible role of peso problems for understanding carry trade payoffs, the risk factors explaining average carry trade payoffs, or the return predictability of currencies held long against the US dollar.

To examine the questions, we focus on potential predictors with different economic nature, which could reveal different predictable components of carry trade payoffs. First, we consider predictors reflecting developments in the commodity sector, choose in particular the Raw Industrials subindex of the CRB Spot Commodity Index

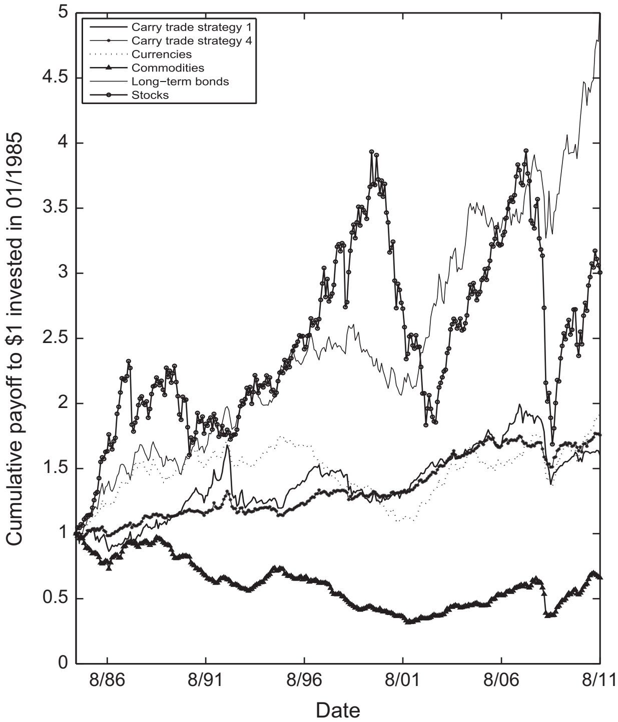

Fig. 2. Cumulative payoffs of carry trades versus alternative investments. Plotted is the cumulative payoff $\Pi_{t=1}^{T}(1+\text{er}_t)$, generated by investing $1 in January 1985 and reinvesting the proceeds on a monthly basis until 08/2011, where er_t is the monthly return of an investment in excess of the US risk-free rate. Because carry trades are constructed as zero-cost investments, there is no need to subtract the risk-free rate. Besides carry trade strategies 1 and 4, with payoffs $\overline{z}_{t+1}^{(1)}$ and $\overline{z}_{t+1}^{(4)}$, we consider investments (all measured in US dollars) in (i) currencies (average return of the G-10 currencies held long against the US dollar), (ii) commodities (Raw Industrials subindex of the CRB Spot Commodity Index), (iii) bonds (average total return of long-term bonds in our sample), and (iv) equities (total return of the MSCI World equity index). All data are from Datastream.

(available from Datastream), and denote it by CRB_t. We are motivated by the observation that two of the highest-yielding currencies (i.e., AUD and NZD) shown in Table A1, which are likely to be actively involved in carry trades, and possibly CAD as well, are among those typically considered commodity currencies (e.g., Chen and Rogoff, 2003). The particular predictor variable we use is the three-month log change, normalized to reflect monthly frequency,

$$\Delta\text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3}). \tag{3}$$

In adopting ΔCRB_t as a predictor of carry trade payoffs, we are also motivated by studies that investigate predictive relations between the exchange rates of commodity currencies and commodity prices (see, among others, Kearns, 2007; Groen and Pesenti, 2009; Chen, Rogoff, and Rossi, 2010).

Second, we explore a variable that could serve as a proxy for uncertainty in global currency markets and consider specifically the normalized log changes in average currency volatility:

$$\Delta\sigma_t^{\text{fx}} = \frac{1}{3}\log(\sigma_t^{\text{avg}}/\sigma_{t-3}^{\text{avg}}). \tag{4}$$

For each currency, we construct monthly volatility as the square root of the sum of squares of daily log changes in the exchange rate against the US dollar over a month (e.g., Andersen, Bollerslev, Diebold, and Labys, 2001), which is then averaged across the G-10 currencies in the sample and denoted by σ_t^{avg}. Bhansali (2008) argues for a link between currency uncertainty and carry trade payoffs, and Menkhoff, Sarno, Schmeling, and Schrimpf (2011, Tables II and III) show that innovations in currency volatility, calculated using absolute daily log returns, are important in cross-sectional tests of currency portfolios, sorted by the forward discount.

Third, we consider an aggregate liquidity variable, the global counterpart to the US TED spread that has been used as a predictor of carry trade payoffs in Brunnermeier, Nagel, and Pedersen (2009, Table 7). The departure in our approach is that we employ an average of the equivalents of the TED spread [Libor (London Interbank Offer Rate) minus the short rate proxy] across the G-10 currencies in our sample (except for NOK and NZD, for which the required data could not be obtained for the full sample), and denote this average as LIQ^{avg}. Motivated by Campbell (1991), the specific predictor variable employed is

$$\Delta\text{LIQ}_t = -\left(\text{LIQ}_t^{\text{avg}} - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}_{t-j}^{\text{avg}}\right). \tag{5}$$

Such an alternative to the US TED spread could represent global liquidity risk, and a version of it has been used, in a different context, by Asness, Moskowitz, and Pedersen (2009). For possible alternatives, see Adrian, Etula, and Shin (2010), who forecast log currency changes.

Table 2
In-sample predictability of carry trade payoffs with single predictors.

Reported are results from predictive regressions $\overline{z}^{(K)}_{t+1} = b_0 + b_x\, x_t + \epsilon^{(K)}_{t+1}$, where x_t is a single predictor. For $K = 1, \ldots, 4$, the payoff in month t+1 of the carry trade with K short and long positions is $\overline{z}^{(K)}_{t+1} = \frac{1}{K}\sum_{k=1}^{K} z^{(k)}_{t+1} = \frac{1}{K}\sum_{k=1}^{K}(z^{\text{short}_k}_{t+1} + r^{\text{long}_k}_{t+1})/2$, where $z^{\text{short}_k}_{t+1}$ ($z^{\text{long}_k}_{t+1}$) is the payoff of the short (long) position established at the end of month t in the k-th lowest-yielding (highest-yielding) currency. All payoffs are denominated in US dollars, and we consider the G-10 currencies (except the euro). The individual predictors x_t are defined as

$$\Delta \text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3}), \quad \Delta\sigma^{\text{fx}}_t = \frac{1}{3}\log(\sigma^{\text{avg}}_t/\sigma^{\text{avg}}_{t-3}), \quad \Delta \text{LIQ}_t = -(\text{LIQ}^{\text{avg}}_t - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}^{\text{avg}}_{t-j}).$$

CRB_t is the Raw Industrials subindex of the CRB Spot Commodity Index, σ^{avg}_t is the average currency volatility for month t across the G-10 currencies, where currency volatility is computed as the square root of the average squared daily log change over a month of a currency's spot exchange rate against the US dollar, and $\text{LIQ}^{\text{avg}}_t$ is the average TED spread (i.e., three-month Libor minus the three-month Treasury bill rate or its equivalent) for month t across the G-10 currencies in our sample (except for NOK and NZD, for which data could not be obtained for the full sample period). The estimates of the slope coefficients b_x are displayed along with the corresponding two-sided p-values NW[p], H[p], and B[p], respectively, based on (i) the heteroscedasticity and autocorrelation consistent (HAC) covariance matrix estimator from Newey and West (1987) (with automatically selected lag as in Newey and West (1994), and shown as 'NW lag"), (ii) the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability, and (iii) the parametric bootstrap, where the predictors are simulated under an ARMA–GARCH structure, chosen based on the BIC (see Panel B of Table A2). Adjusted R^2s are shown as $\overline{R}^2$. Regression intercepts are not reported to save on space.

Predictor	Carry strategy	b_x	NW[p]	H[p]	B[p]	$\overline{R}^2$ (%)	NW lag
Commodity returns, ΔCRB_t	1	0.24	0.00	0.02	0.02	3.9	1
	2	0.17	0.00	0.01	0.01	3.4	3
	3	0.12	0.01	0.02	0.01	2.4	4
	4	0.10	0.00	0.01	0.01	2.4	5
Currency volatility, $\Delta\sigma^{\text{fx}}_t$	1	−0.05	0.01	0.00	0.01	3.7	3
	2	−0.03	0.01	0.01	0.03	2.6	3
	3	−0.03	0.00	0.00	0.00	3.7	4
	4	−0.03	0.00	0.00	0.00	4.3	4
Liquidity, ΔLIQ_t	1	0.03	0.03	0.09	0.09	3.1	0
	2	0.02	0.02	0.07	0.09	2.6	4
	3	0.02	0.00	0.05	0.05	3.5	4
	4	0.02	0.00	0.04	0.03	2.5	3

Three additional points need to be made. First, our use of the volatility variable $\Delta\sigma^{\text{fx}}_t$ departs from the treatment in Lustig, Roussanov, and Verdelhan (2010, Section 4.5.3), who adopt average equity return volatility to explain average currency returns, and in Menkhoff, Sarno, Schmeling, and Schrimpf (2012), who adopt innovations in currency volatility as a pricing factor to study the cross section of currency portfolio returns. Second, we consider the merit of our predictors in Section 4.2, where we examine the global versions of two other predictors: term spread and change in interest rate of long-term bonds, which have been found to be important in a cross sectional setting by Ang and Chen (2010). Third, Table A2 elaborates on the statistical properties of our predictors, which generally can be represented by an ARMA–GARCH structure, and Appendix A provides discussion. Our predictors are intended to capture distinct aspects of the global macroeconomic environment, as seen, for example, in the moderate correlations between them.

3.3. *Carry trade payoffs display predictability in-sample*

Next, we examine predictive regressions of the type

$$\overline{z}^{(K)}_{t+1} = b_0 + \mathbf{x}'_t\mathbf{b}_x + \epsilon^{(K)}_{t+1} \quad \text{for } K = 1, \ldots, 4$$
$$\text{and} \quad \mathbf{x}_t = [\Delta\text{CRB}_t \;\; \Delta\sigma^{\text{fx}}_t \;\; \Delta\text{LIQ}_t]', \tag{6}$$

where $\overline{z}^{(K)}_{t+1}$ is the carry trade payoff for month t+1 as in Eq. (2) and $\mathbf{b} = [b_0 \mathbf{b}'_x]'$ is the vector of regression coefficients. We also consider univariate counterparts of (6), with individual components of the vector $\mathbf{x}_t$ as predictors. Whereas a predictor could show statistical significance in a univariate regression, one or more predictors in a multiple regression could exhibit diminished predictive ability. Examining the predictability of dynamically re-balanced carry trades distinguishes our exercise from the corresponding ones in Cutler, Poterba, and Summers (1991), Bekaert and Hodrick (1992), Mark (1995), Kilian (1999), Evans and Lyons (2002), Hau and Rey (2006), Corte, Sarno, and Tsiakas (2009), Lustig, Roussanov, and Verdelhan (2010), and Sarno, Schneider, and Wagner (2012).

Tables 2 and 3 present results with single and multiple predictors, respectively. We report the estimates of the slope coefficients $\mathbf{b}_x$ and the corresponding two-sided p-values NW[p], H[p], and B[p], based on the heteroscedasticity and autocorrelation consistent (HAC) covariance matrix estimator from Newey and West (1987), with automatically selected lag as in Newey and West (1994); the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability; and parametric bootstrap (as in Mark, 1995; Kilian, 1999; Amihud, Hurvich, and Wang, 2009, and also as discussed in Appendix B), respectively.

The simulation evidence in Ang and Bekaert (2007), Wei and Wright (2009), and Bakshi et al. (2011) indicates that the Hodrick (1992) 1B estimator is conservative and displays suitable size properties in the context of persistent predictors. With this in mind, our statistical inference relies predominantly on the Hodrick (1992) estimator,

Table 3

In-sample predictability of carry trades with multiple predictors.

Reported are results from predictive regressions $\overline{z}_{t+1}^{(K)} = b_0 + \mathbf{x}'_t\, \mathbf{b}_x + \epsilon_{t+1}^{(K)}$, where $\mathbf{x}_t$ is a vector of two or three predictors, and $\overline{z}_{t+1}^{(K)}$ is the payoff of a carry trade strategy with K short and long positions, for $K=1,\dots,4$. The predictors are $\Delta\text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3})$, $\Delta\sigma_t^{\text{fx}} = \frac{1}{3}\log(\sigma_t^{\text{avg}}/\sigma_{t-3}^{\text{avg}})$, and $\Delta\text{LIQ}_t = -(\text{LIQ}_t^{\text{avg}} - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}_{t-j}^{\text{avg}})$, where CRB_t is the Raw Industrials subindex of the CRB Spot Commodity Index, σ_t^{avg} is the average currency volatility for month t across the G-10 currencies, where currency volatility is computed as the square root of the average squared daily log change over a month of a currency's spot exchange rate against the US dollar, and $\text{LIQ}_t^{\text{avg}}$ is the average TED spread (i.e., three-month Libor minus the three-month Treasury bill rate or its equivalent) for month t across the G-10 currencies in our sample. The estimates of the slope coefficients $\mathbf{b}_x$ are displayed along with the corresponding two-sided p-values H[p] and B[p], respectively, based on (i) the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability, and (ii) the parametric bootstrap, where the predictors are simulated under an ARMA–GARCH structure, chosen based on the BIC (see Panel B of Table A2). The corresponding p-values for the test of the null hypothesis that the slope coefficients are jointly equal to zero are displayed as Joint H[p]. Adjusted R^2s are shown as $\overline{R}^2$. Regression intercepts are not reported to save on space.

	Panel A: All three predictors				Panel B: ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$				Panel C: ΔCRB_t and ΔLIQ_t				Panel D: $\Delta\sigma_t^{\text{fx}}$ and ΔLIQ_t			
	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
b_1 (ΔCRB_t)	0.18	0.13	0.09	0.08	0.20	0.14	0.09	0.08	0.22	0.15	0.11	0.09				
H[p]	0.06	0.04	0.05	0.04	0.04	0.02	0.05	0.04	0.03	0.02	0.02	0.01				
B[p]	0.06	0.04	0.09	0.07	0.04	0.02	0.06	0.05	0.04	0.02	0.03	0.02				
b_2 ($\Delta\sigma_t^{\text{fx}}$)	−0.04	−0.02	−0.02	−0.02	−0.04	−0.03	−0.03	−0.03					−0.05	−0.03	−0.03	−0.03
H[p]	0.03	0.05	0.03	0.01	0.01	0.02	0.01	0.00					0.01	0.02	0.01	0.00
B[p]	0.06	0.09	0.04	0.02	0.03	0.07	0.01	0.00					0.03	0.04	0.01	0.01
b_3 (ΔLIQ_t)	0.02	0.02	0.02	0.01					0.03	0.02	0.02	0.01	0.03	0.02	0.02	0.01
H[p]	0.27	0.23	0.15	0.20					0.16	0.14	0.08	0.08	0.19	0.16	0.12	0.14
B[p]	0.36	0.26	0.14	0.20					0.22	0.17	0.07	0.07	0.25	0.19	0.10	0.13
$\overline{R}^2$ (%)	7.6	5.9	6.8	6.6	6.1	4.8	4.9	5.5	6.1	5.1	5.3	4.6	5.6	4.2	5.8	5.5
Joint H[p]	0.02	0.01	0.02	0.00	0.01	0.01	0.01	0.00	0.04	0.02	0.02	0.01	0.01	0.02	0.01	0.00
Joint B[p]	0.04	0.02	0.02	0.00	0.02	0.01	0.01	0.00	0.04	0.02	0.04	0.01	0.03	0.03	0.01	0.00

together with the parametric bootstrap, which could reflect better the finite-sample properties of the estimators. Important for our task, we also test the null hypothesis that $\mathbf{b}_x = \mathbf{0}_{L\times 1}$, where L is the number of predictors, which amounts to assessing joint parameter significance. While Newey and West (1987) p-values are presented in Table 2 to provide perspective, they are omitted thereafter.

Our results convey several insights about the predictability of carry trade payoffs. To start with, each predictor is statistically significant in the univariate regressions in Table 2, where, out of 36 p-values, 32 do not exceed 0.05, and all are below 0.1.

In contrast, Table 3 shows that ΔLIQ_t generally loses statistical significance in the presence of either one or both of ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$, whereby 10 out of 12 H[p] or B[p] values for ΔLIQ_t exceed 0.1. At the same time, 23 (17) out of 24 H[p] (B[p]) values are below 0.05 for ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$. Thus, we find statistically strong predictive ability for ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$, while mixed evidence exists for predictability with ΔLIQ_t. However, we show in Table 5 that ΔLIQ_t can predict carry trade payoffs in a multivariate setting, but at horizons between four and six months.

In the univariate estimations, the adjusted R^2s range between 2.4% and 3.9% for ΔCRB_t, 2.6% and 4.3% for $\Delta\sigma_t^{\text{fx}}$, and 2.5% and 3.5% for ΔLIQ_t. These goodness-of-fit statistics appear comparable to those reported for currency returns, for example, Adrian, Etula, and Shin (2010, Tables 1 and 2), Lustig, Roussanov, and Verdelhan (2010, Table II), and Sarno, Schneider, and Wagner (2012, Table 8). When all three predictors are included together, we obtain adjusted R^2s in the range 5.9–7.6%.

Like Cochrane and Piazzesi (2005), we further measure the strength of predictability by the p-values H[p] and B[p] for the test of joint parameter significance, when either two or three predictors are included together. The key point is that all joint p-values are below 0.05, implying that carry trade payoffs have predictable components that reflect exposures to movements in commodity prices, currency volatility, and global liquidity.

Turning to the signs and magnitudes of the coefficient estimates, the slope estimates on ΔCRB_t are uniformly positive, across both the univariate and multivariate regressions, and vary between 0.1 and 0.24 in Table 2 and between 0.08 and 0.22 in Table 3. The positive slope estimates on ΔCRB_t would agree with the interpretation that commodity currencies are strengthened by past increases in commodity prices, similar to a result reported in Kearns (2007). Chen, Rogoff, and Rossi (2010) argue that the predictive relation is more likely from exchange rates to commodity prices, but they also point out that the dynamic relation between exchange rates and fundamentals is hard to unravel and is possibly time-varying (see also Mark, 2001). Our estimates, thus, present evidence about this relation, from the perspective of carry trade payoffs.

Other salient features of our predictive regressions are the negative (positive) slope estimates on $\Delta\sigma_t^{\text{fx}}$ (ΔLIQ_t). With slope estimates on $\Delta\sigma_t^{\text{fx}}$ between −0.02 and −0.05, a high level of $\Delta\sigma_t^{\text{fx}}$ predicts lower next-month carry trade payoffs. In turn, the documented positive slope estimates on ΔLIQ_t can be rationalized with the intuition that higher liquidity, reflected in lower spreads, results in larger amounts of funds flowing into carry trades, which shift relative exchange rates and generate higher carry trade payoffs. Such a liquidity channel has some support in our sample and appears to depart from the fundamentals-related interpretation, indicating possible multiple determinants of carry trade payoffs and their predictability.

The economic sensitivity of carry trade payoffs varies across the predictors but remains large in magnitude. For example, as can be inferred from Panel A of Table 3 and Panel A of Table A2, an increase of 1 standard deviation in the predictor can lead to an annualized change of 2.0%, −2.3%, and 1.6% in $\overline{z}^{(4)}_{t+1}$, when the predictor is, respectively, ΔCRB_t, $\Delta\sigma^{\text{fx}}_t$, and ΔLIQ_t. These changes are comparable to the average carry trade payoffs reported in Table 1, and we further elaborate on the economic significance of predictability in an out-of-sample context in Section 3.7 and Table 7.[4]

The message worth highlighting is that commodity risks and currency volatility risks, as reflected in ΔCRB_t and $\Delta\sigma^{\text{fx}}_t$, are largely orthogonal as predictors of subsequent carry trade payoffs and exhibit superior predictive ability, whereas ΔLIQ_t could have secondary importance at the monthly frequency.

Do carry trades differ, with respect to their predictability, from the fixed currency pairs that most often enter into the carry trades? This question is prompted by the contrast, discussed in Section 3.1, between the profitability of the carry trades and the component fixed currency pairs. Table A3 shows results from analogous predictive regressions, performed on individual pairs, and suggests that the fixed currency pairs do not appear to be predictable. We observe that six out of 18H[p] values are below 0.1, and none is below 0.05, which highlights further the importance of dynamic re-balancing in the design of carry trades. As seen, only the trade that involves borrowing in Japanese yen and investing in British pound displays some predictable pattern.

3.4. In-sample predictability is affirmed in a system of carry trade components

Following Ang and Bekaert (2007, Appendix B), we also test for the significance of a predictor in a system of predictive regressions. Here we stack up the four long and short components of carry trade payoffs, as shown in the second part of Eq. (2), and consider the specification

$$\mathbf{z}_{t+1} = \mathbf{A} + \mathbf{C}\mathbf{x}_t + \mathbf{e}_{t+1} \quad \text{recalling that } \mathbf{z}_{t+1} = [z^{(1)}_{t+1}\ \ z^{(2)}_{t+1}\ \ z^{(3)}_{t+1}\ \ z^{(4)}_{t+1}]', \tag{7}$$

where $\mathbf{A}$ is a 4×1 vector of regression intercepts and $\mathbf{C}$ is a 4×1, 4×2, or 4×3 matrix of slope coefficients. The coefficient estimates are obtained equation by equation via OLS, and we test whether $\mathbf{C}=\mathbf{0}$ using the covariance matrix estimator of Hodrick (1992).

With two (all three) predictors, the corresponding p-values, reported in Table 4, are based on the χ^2 statistic with 8 (12) degrees of freedom. The evidence from the system of equations broadly confirms the predictive power of our predictors, particularly ΔCRB_t and $\Delta\sigma^{\text{fx}}_t$. Moreover, adding ΔLIQ_t to the set of predictors worsens the joint p-values.

3.5. Predictors display differential predictive ability at longer horizons

In the spirit of Meese and Rogoff (1983), Cutler, Poterba, and Summers (1991), Mark (1995), Kilian and Taylor (2003), Ang and Chen (2010), and Lustig, Roussanov, and Verdelhan (2010), we ask whether our predictors maintain their predictive ability at longer horizons. For this purpose, we construct overlapping payoffs over two to six months and consider the regressions

$$\overline{z}^{(K)}_{t+j} = b_0 + \mathbf{x}'_t\mathbf{b}_x + \epsilon^{(K)}_{t+j}, \quad \text{for } K=1,\ldots,4 \quad \text{and} \quad j=2,\ldots,6, \tag{8}$$

where $\overline{z}^{(K)}_{t+j}$ is now the cumulated j-month payoff.

Table 5 shows that, as in the regressions with one-month forecasting horizon, the slope coefficients associated with $\Delta\sigma^{\text{fx}}_t$ are negative and decreasing in absolute value with the horizon, while those for ΔLIQ_t are positive and slightly increasing. The Hodrick (1992, Eq. (8)) p-values, which are adjusted for overlapping observations, reveal significance for $\Delta\sigma^{\text{fx}}_t$ for up to three-month forecast horizons, in contrast to those for ΔLIQ_t, which show significance only at horizons from four to six months.

The slope coefficients on ΔCRB_t are insignificant, indicating lack of predictability with ΔCRB_t for carry trade payoffs beyond the one-month horizon. However, in spite of the differential forecasting ability of the three predictors at various horizons, the joint p-values are mostly significant, with 14 (19) out of the 20 joint H[p]s reported in Table 5 being below 0.05 (0.1). The predictability with some of our variables is not confined to monthly carry trade payoffs.

3.6. Predictability is supported by out-of-sample statistics

The second set of statistical tests of predictability are based on the out-of-sample R^2 statistic, as suggested in Campbell and Thompson (2008), and the adjusted mean squared prediction error statistic (MSPE-adjusted) developed in Clark and West (2007).

Let $\widehat{\mu}_{t+1}$ be the prediction for month t+1 from a predictive regression of the form Eq. (6) and μ_{t+1} be the historical average payoff, both estimated using data up to and including month t. Then the out-of-sample R^2 statistic is $R^2_{OS} = 1 - \sum_{t=0}^{T-1}(\overline{z}^{(K)}_{t+1} - \widehat{\mu}_{t+1})^2 / \sum_{t=0}^{T-1}(\overline{z}^{(K)}_{t+1} - \mu_{t+1})^2$. A positive R^2_{OS} indicates that the mean-squared prediction error from the predictive regression model is lower than that from a prediction based on the historical average payoff.

The MSPE-adjusted statistic is obtained using $f_{t+1} = (\overline{z}^{(K)}_{t+1} - \mu_{t+1})^2 - [(\overline{z}^{(K)}_{t+1} - \widehat{\mu}_{t+1})^2 - (\mu_{t+1} - \widehat{\mu}_{t+1})^2]$ and equals the t-statistic from the regression of f_{t+1} on a constant, for which we report one-sided p-values. A p-value below 0.05 indicates a statistically significant improvement of the forecast when using the predictive regression model over the benchmark of using historical average payoff as predictor.

[4] We find further that monthly carry trade payoffs $\overline{z}^{(K)}_t$ display contemporaneous correlations between 0.09 and 0.17 with OECD industrial production growth. Carry trade payoffs, thus, appear to be mildly pro-cyclical, consistent with the directional effects of our predictors.

Table 4
Testing for joint predictability of carry trade returns based on a system of equations.

We consider a system of equations:

$$\mathbf{z}_{t+1} = \mathbf{A} + \mathbf{C}\mathbf{x}_t + \mathbf{e}_{t+1} \quad \text{with } \mathbf{z}_{t+1} = [z^{(1)}_{t+1} z^{(2)}_{t+1} z^{(3)}_{t+1} z^{(4)}_{t+1}]',$$

where $\mathbf{A}$ is a 4×1 vector of regression intercepts and $\mathbf{C}$ is a 4×1, 4×2, or 4×3 matrix of slope coefficients, and $\mathbf{e}_{t+1}$ is a 4×1 vector of regression residuals. The coefficient estimates are obtained equation by equation via OLS. Reported are the χ^2 test statistics, and the corresponding p-values for the Hodrick (1992) test that the slope coefficients are jointly equal to zero. The covariance matrix estimator is computed under the null hypothesis of no predictability (see also Ang and Bekaert, 2007, Appendix B). The degrees of freedom for the χ^2 statistic equal the number of exclusion restrictions, i.e., four, eight, or 12, with one, two, and all three predictors, respectively.

	Single predictor			Multiple predictors			
	ΔCRB_t	$\Delta\sigma^{\text{fx}}_t$	ΔLIQ_t	$\Delta\text{CRB}_t+\Delta\sigma^{\text{fx}}_t$	$\Delta\text{CRB}_t+\Delta\text{LIQ}_t$	$\Delta\sigma^{\text{fx}}_t+\Delta\text{LIQ}_t$	All three
χ^2	9.9	13.2	4.6	17.9	13.3	14.9	20.0
Joint H[p]	0.04	0.01	0.33	0.02	0.10	0.06	0.07

Table 5
In-sample predictability of carry trade payoffs at longer horizons.

Reported are results from predictive regressions $\overline{z}^{(K)}_{t+j} = b_0 + \mathbf{x}'_t\mathbf{b}_x + \epsilon^{(K)}_{t+j}$, for $j=2,\ldots,6$, where $\mathbf{x}_t$ is a vector of three predictors and $\overline{z}^{(K)}_{t+j}$ is the cumulative payoff over two to six months following month t of a carry trade strategy with K short and long positions, for $K=1,\ldots,4$. The predictors are $\Delta\text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3})$, $\Delta\sigma^{\text{fx}}_t = \frac{1}{3}\log(\sigma^{\text{avg}}_t/\sigma^{\text{avg}}_{t-3})$, and $\Delta\text{LIQ}_t = -(\text{LIQ}^{\text{avg}}_t - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}^{\text{avg}}_{t-j})$, where CRB_t is the Raw Industrials subindex of the CRB Spot Commodity Index, σ^{avg}_t is the average currency volatility for month t across the G-10 currencies, where currency volatility is computed as the square root of the average squared daily log change over a month of a currency's spot exchange rate against the US dollar, and $\text{LIQ}^{\text{avg}}_t$ is the average TED spread (i.e., three-month Libor minus the three-month Treasury bill rate or its equivalent) for month t across the G-10 currencies in our sample. The two-sided p-values for the slope coefficients are shown as H[p] and are based on the Hodrick (1992, Eq. (8)) covariance matrix estimator under the null of no predictability, which adjusts for overlapping observations. Next, the corresponding p-values for the test of the null hypothesis that the slope coefficients are jointly equal to zero are displayed as Joint H[p]. Regression intercepts are not reported to save on space.

Horizon	Carry strategy	b_1 (ΔCRB_t)	H[p]	b_2 ($\Delta\sigma^{\text{fx}}_t$)	H[p]	b_3 (ΔLIQ_t)	H[p]	$\overline{R}^2$ (%)	Joint H[p]
2-month	1	0.31	0.10	−0.09	0.01	0.02	0.36	11.5	0.02
	2	0.20	0.11	−0.06	0.01	0.03	0.15	10.5	0.00
	3	0.13	0.12	−0.04	0.01	0.02	0.13	8.4	0.01
	4	0.11	0.12	−0.04	0.00	0.02	0.19	10.0	0.00
3-month	1	0.37	0.17	−0.08	0.04	0.04	0.17	8.7	0.06
	2	0.21	0.23	−0.06	0.03	0.04	0.09	8.4	0.02
	3	0.13	0.31	−0.04	0.09	0.03	0.10	5.6	0.03
	4	0.09	0.39	−0.04	0.02	0.03	0.12	7.2	0.01
4-month	1	0.29	0.34	−0.07	0.13	0.06	0.05	5.9	0.08
	2	0.10	0.66	−0.06	0.06	0.05	0.04	5.9	0.02
	3	0.03	0.88	−0.04	0.13	0.04	0.05	3.7	0.04
	4	−0.01	0.97	−0.04	0.07	0.04	0.04	5.9	0.02
5-month	1	0.23	0.50	−0.05	0.21	0.06	0.04	3.4	0.10
	2	−0.03	0.92	−0.05	0.10	0.06	0.02	3.4	0.02
	3	−0.08	0.70	−0.03	0.23	0.04	0.03	1.8	0.05
	4	−0.12	0.52	−0.03	0.20	0.04	0.02	3.8	0.03
6-month	1	0.18	0.64	−0.03	0.44	0.07	0.02	2.6	0.10
	2	−0.14	0.62	−0.03	0.27	0.06	0.01	2.5	0.04
	3	−0.17	0.47	−0.02	0.45	0.04	0.04	1.1	0.13
	4	−0.22	0.31	−0.02	0.33	0.04	0.03	3.1	0.08

Both R^2_{OS} and the MSPE-adjusted statistic are computed relying on an expanding window with initial length of 180 months. Other choices of window length impart similar qualitative conclusions and are omitted. Panel A of Table 6 presents the out-of-sample statistics obtained with either two or three predictors.

When all three predictors are included in a regression, the out-of-sample R^2s range between 3.2% and 9.4%. Thus, the out-of-sample R^2s tend to be higher than those reported for the equity market, for instance, in Campbell and Thompson (2008) and Rapach, Strauss, and Zhou (2010). The uniformly positive out-of-sample R^2s strengthen the evidence for a predictable component in carry trade payoffs.

Next, the evidence from the MSPE-adjusted statistic across predictors appears to be aligned with that from the out-of-sample R^2s. When all three predictors are used, the MSPE-adjusted p-values are all below or equal to 0.03, while the p-values are somewhat higher for pairs including ΔLIQ_t. Overall, the out-of-sample statistics validate the predictability of carry trade payoffs.

Table 6
Predictability based out-of-sample R^2 and MSPE-adjusted p-values.

Reported are (i) out-of-sample R^2's (Campbell and Thompson, 2008) and (ii) one-sided p-values for the MSPE-adjusted statistic (Clark and West, 2007), obtained with two or three predictors and in combination forecasts (Stock and Watson, 2004). The predicted variables $\overline{z}^{(K)}_{t+1}$, for $K=1,\ldots,4$, are the payoffs of carry trade strategies with K short and long positions. The out-of-sample R^2 statistic is $R^2_{OS}=1-(\sum_{t=0}^{T-1}(\overline{z}^{(K)}_{t+1}-\widehat{\mu}_{t+1})^2/\sum_{t=0}^{T-1}(\overline{z}^{(K)}_{t+1}-\mu_{t+1})^2)$, and the MSPE-adjusted p-values are obtained by regressing $f_{t+1}=(\overline{z}^{(K)}_{t+1}-\mu_{t+1})^2-[(\overline{z}^{(K)}_{t+1}-\widehat{\mu}_{t+1})^2-(\mu_{t+1}-\widehat{\mu}_{t+1})^2]$ on a constant. An expanding window with initial length of 180 months is used to compute $\widehat{\mu}_{t+1}$ (the carry trade payoff for month t+1, predicted at the end of month t using one of our predictors), and μ_{t+1} (the average carry trade payoff observed at the end of month t). The combination forecasts use: (i) mean prediction, (ii) median prediction, and (iii) weighted prediction, where the weights depend on the performance of individual out-of-sample predictions over a hold-out period of q months. In particular, the weight ω^i_t assigned to the forecast obtained with predictor i at the end of month t is $\omega^i_t=1/\phi^i_t/\sum_{j=1}^{3}(1/\phi^j_t)$ for $\phi^i_t=\sum_{s=q}^{t-1}\theta^{t-1-s}(\overline{z}^{(K)}_{s+1}-\widehat{\mu}^i_{s+1})^2$, where $\widehat{\mu}^i_{s+1}$ is the prediction for month $s+1$ obtained with predictor i at the end of month s, $q=60$, and $\theta=0.9$ allows to attribute higher weights to more recent predictions.

	Panel A: Two or three predictors				Panel B: Combination forecasts		
Carry strategy	All three	ΔCRB_t $+\Delta\sigma^{fx}_t$	ΔCRB_t $+\Delta LIQ_t$	$\Delta\sigma^{fx}_t$ $+\Delta LIQ_t$	Mean	Median	Weighted
Out-of-sample R^2 (%)							
1	9.4	14.4	9.2	6.3	10.7	9.1	10.5
2	3.2	8.8	4.1	2.2	6.7	5.2	6.7
3	6.1	6.8	6.9	5.8	6.9	6.3	6.8
4	7.7	7.7	6.1	8.5	7.1	6.9	6.9
MSPE-adjusted one-sided p-values							
1	0.03	0.01	0.04	0.07	0.02	0.03	0.02
2	0.03	0.01	0.03	0.09	0.01	0.03	0.01
3	0.02	0.01	0.03	0.04	0.02	0.03	0.02
4	0.01	0.01	0.02	0.01	0.01	0.01	0.01

3.7. A predictability-based decision rule can sharpen Sharpe ratios and mitigate negative skewness

An additional perspective can be garnered by assessing the economic significance of the predictable variation in carry trade payoffs. The question to pose is: Can one improve the profitability or limit the downside of carry trade payoffs by taking into consideration their predictability? We adopt a decision rule to take a position in a carry trade (based on rank-ordering of forward discounts at the end of month t), if the predictive regression model equation (6) predicts a positive carry trade payoff for month t+1. If the model predicts a negative payoff, we refrain from entering into a carry trade.

To be consistent with the implementation in Section 3.6, we estimate the predictive regression model with an expanding window, taking an initial length of 180 months. We wish to compare the Sharpe ratios and skewness of the conditional payoffs thus obtained with their unconditional counterparts and gauge the potential economic impact of the predictability of carry trade payoffs. Further, we implement the non-parametric market-timing test of Henriksson and Merton (1981).

Panel A of Table 7 presents the Sharpe ratios and skewness of the payoffs, along with associated stationary bootstrap p-values in curly brackets. A p-value is computed as the proportion of 25,000 bootstrap samples, drawn under the null of no predictability, for which the Sharpe ratio (skewness) of a conditional strategy is lower (more negative) than that of the unconditional counterpart. Via these p-values, we can assess whether conditioning, using our predictors, can impact the Sharpe ratio or skewness of carry trade payoffs in a statistically significant manner.

Comparing Panel A of Table 7 with Table 1, predictability offers improvement in Sharpe ratios, more so for strategies 1 and 2. Equally relevant, the Sharpe ratios of some of the conditional strategies are comparable to the counterparts generated by the trading algorithm in Jordá and Taylor (2009, Tables 6 and 11).

Because many studies, as well as our Section 3.1, have highlighted the link between high negative skewness and average carry trade payoffs, we consider the interaction between predictability and skewness in the context of our decision rule. Relative to the skewness of carry trade payoffs in the out-of-sample period (−1.39, −0.89, −0.60, and −0.01, respectively, for strategies 1–4), the skewness of the conditional strategies uniformly increases, significantly so for strategies 1 and 2. Our results here are compatible with those in Jordá and Taylor (2009, Section 3, Table 11), who observe analogous reduction in the negative skewness of conditional payoffs.

Next, we examine the null hypothesis of no market timing ability based on our predictors and employ for this purpose the one-sided nonparametric market timing test of Henriksson and Merton (1981, Eq. (9)).[5] The test is applicable to observable forecasts of the sign of carry trade payoffs and employs certain forecast probabilities, conditional upon the realized outcomes, which are empirically captured by the proportions of time when our predictability-based decision rule suggests staying out of a carry trade (correctly or not). When all three predictors

[5] See also, among others, the related treatment in Breen, Glosten, and Jagannathan (1989), Pesaran and Timmermann (1994), and Christoffersen and Diebold (2006).

Table 7
Sharpe ratios and skewness of carry trade payoffs from a predictability-based decision rule.

Reported in this table are (i) Sharpe ratios and (ii) skewness of conditional carry trade payoffs, obtained with the three predictors (individually and together), and in combination forecasts (Stock and Watson, 2004). Conditional payoffs are based on a decision rule to take a position in a carry trade at the end of month t if a positive carry trade payoff is predicted for month $t+1$ and to do nothing if a negative payoff is predicted. The payoff predictions are obtained in predictive regressions with an expanding window with initial length of 180 months. In curly brackets we show p-values for the null hypothesis that the conditional strategy does not yield statistically higher Sharpe ratios or lower skewness than the unconditional strategy, which always takes a position in the carry trade, irrespective of the sign of the predicted payoff. A p-value is computed as the proportion of 25,000 stationary bootstrap samples, drawn under the null of no predictability, for which the Sharpe ratio (skewness) of a conditional strategy is lower (more negative) than that of the unconditional counterpart.

	Panel A: Two or three predictors				Panel B: Combination forecasts		
Carry strategy	All three	$\Delta CRB_t + \Delta\sigma_t^{fx}$	$\Delta CRB_t + \Delta LIQ_t$	$\Delta\sigma_t^{fx} + \Delta LIQ_t$	Mean	Median	Weighted
Sharpe ratios							
1	0.86	0.85	0.96	0.63	0.81	0.70	0.81
	{0.11}	{0.08}	{0.07}	{0.21}	{0.12}	{0.18}	{0.12}
2	0.86	0.95	0.93	0.64	0.70	0.54	0.70
	{0.12}	{0.09}	{0.09}	{0.25}	{0.21}	{0.36}	{0.21}
3	0.86	0.89	0.84	0.72	0.74	0.70	0.73
	{0.26}	{0.19}	{0.28}	{0.35}	{0.36}	{0.40}	{0.38}
4	1.02	0.81	1.13	0.79	0.76	0.77	0.76
	{0.13}	{0.32}	{0.08}	{0.35}	{0.39}	{0.37}	{0.39}
Skewness							
1	−0.36	0.16	−0.39	−0.72	−0.39	−0.42	−0.39
	{0.06}	{0.00}	{0.05}	{0.14}	{0.05}	{0.05}	{0.05}
2	−0.27	0.01	−0.22	−0.22	−0.29	−0.35	−0.29
	{0.11}	{0.03}	{0.09}	{0.10}	{0.12}	{0.15}	{0.12}
3	−0.22	−0.25	−0.23	−0.23	−0.27	−0.25	−0.25
	{0.23}	{0.23}	{0.24}	{0.20}	{0.24}	{0.23}	{0.23}
4	0.48	0.10	0.53	0.09	0.02	0.02	0.02
	{0.08}	{0.38}	{0.05}	{0.36}	{0.46}	{0.46}	{0.47}

Table 8
Henriksson and Merton market timing test.

Reported are results for the Henriksson and Merton (1981, Eq. (9)) one-sided nonparametric test of the null hypothesis of no market timing ability based on our predictors. The test is valid for observable forecasts of the sign of carry trade payoffs and relies on forecast probabilities, conditional upon the realized outcomes. The test reflects the proportion of observations when the predictive regressions correctly forecast a negative carry trade payoff, i.e., the forecast deviates from the notion that carry trades, obtained by sorting currencies based on the interest rate differential, are profitable. For different sets of predictors we report the respective proportion, whereby * (**) denotes that the test rejects the null at the 10% (5%) confidence level.

Carry strategy	All three	$\Delta CRB_t + \Delta\sigma_t^{fx}$	$\Delta CRB_t + \Delta LIQ_t$	$\Delta\sigma_t^{fx} + \Delta LIQ_t$
1	9.5*	8	8**	10.2
2	12.4**	10.2*	9.5**	10.9
3	9.5**	8**	5.1*	6.6
4	9.5**	7.3	7.3**	6.6

are employed, Table 8 shows that the null hypothesis can be rejected at the 5% confidence level for strategies 2, 3 and 4 and at the 10% level for strategy 1. Besides, the test rejects at the 5% (10%) level for strategies 1, 2 and 4 (3) when ΔCRB_t and ΔLIQ_t are used together and for strategy 4 (3) when ΔCRB_t and $\Delta\sigma_t^{fx}$ are used. These conclusions qualitatively match the ones reported for our predictability-based decision rule in Table 7. The similarity reflects the fact that the decision rule employs the sign of the forecasted carry trade payoff to construct trading strategies, in common with the market timing test considered.[6]

In conclusion, exploiting predictability generally improves the profile of carry trade payoffs when measured by the trade-off in mean and variance, and it also appears to help avoid undesirable payoff asymmetries. In addition, the market timing test provides some evidence for the economic value of carry trade predictability.

3.8. *Combination forecasts reinforce the evidence*

Guided by Stock and Watson (2004), our next tests involve combination forecasts, which are weighted averages of individual predictions and, arguably, perform better than the components.

Following convention, we employ three distinct weighting schemes, obtaining mean prediction, median prediction, and weighted prediction, in which the weights depend on the performance of individual predictions over

[6] Here a trade is initiated only if the predicted carry trade payoff is positive. We also implement a strategy in which we reverse the carry trade when the predictive regression implies a negative payoff, and our conclusions remain robust. These results are not reported to save on space but are available upon request.

Table 9
Predictability of the currency components and the short and long legs of carry trades.

Reported are results from the predictive regressions $\overline{z}_{t+1}^{(K)} = b_0 + \mathbf{x}'_t \mathbf{b}_x + \epsilon_{t+1}^{(K)}$, where $\mathbf{x}_t$ is a vector of three predictors: $\Delta \text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3})$, $\Delta\sigma_t^{\text{fx}} = \frac{1}{3}\log(\sigma_t^{\text{avg}}/\sigma_{t-3}^{\text{avg}})$, and $\Delta \text{LIQ}_t = -(\text{LIQ}_t^{\text{avg}} - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}_{t-j}^{\text{avg}})$. In Panel A, $\overline{z}_{t+1}^{(K)}$ for $K=1,\ldots,4$ is the payoff of a trading strategy with K short positions against the US dollar in the K lowest-yielding currencies. In Panel B, $\overline{z}_{t+1}^{(K)}$ for $K=1,\ldots,4$ is the payoff of a trading strategy with K long positions against the US dollar in the K highest-yielding currencies. In Panel C, $\overline{z}_{t+1}^{(K)}$ for $K=1,\ldots,4$ is the payoff of the carry trade strategy with short (long) positions in the K lowest-yielding (highest-yielding) currencies, from which the effect of interest rates has been eliminated, i.e., we have used (Eqs.) (1) and (2), where forward rates $F_{t,t+1}$ have been substituted by the spot rates S_t. The estimates of the slope coefficients $\mathbf{b}_x$ are displayed along with the corresponding two-sided p-values H[p] and B[p], respectively, based on (i) the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability and (ii) the parametric bootstrap, where the predictors are simulated under an ARMA–GARCH structure, chosen based on the BIC. Also shown are p-values for the test of the null hypothesis that the slope coefficients are jointly equal to zero. Adjusted R^2s are shown as $\overline{R}^2$. Regression intercepts are not reported to save on space.

	Panel A: Short leg of carry trade				Panel B: Long leg of carry trade				Panel C: Currency component of carry trade			
	1	2	3	4	1	2	3	4	1	2	3	4
b_1 (ΔCRB_t)	0.07	−0.03	−0.10	−0.11	0.30	0.28	0.22	0.23	0.19	0.13	0.09	0.08
H[p]	0.68	0.79	0.22	0.12	0.07	0.03	0.04	0.02	0.04	0.03	0.04	0.03
b_2 ($\Delta\sigma_t^{\text{fx}}$)	−0.02	0.01	0.01	0.00	−0.03	−0.04	−0.03	−0.03	−0.04	−0.02	−0.02	−0.02
H[p]	0.36	0.59	0.73	1.00	0.17	0.11	0.10	0.13	0.02	0.05	0.03	0.01
b_3 (ΔLIQ_t)	0.02	0.01	0.01	0.00	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.01
H[p]	0.27	0.46	0.66	0.95	0.53	0.51	0.34	0.34	0.26	0.22	0.14	0.18
$\overline{R}^2$ (%)	0.6	−0.7	−0.3	0.0	5.0	6.4	6.2	6.3	8.4	6.5	7.3	7.1
Joint H[p]	0.40	0.88	0.62	0.46	0.11	0.03	0.05	0.05	0.02	0.01	0.01	0.00

a holdout period that starts, say, q periods after the beginning of the sample. In this case, the weight ω_t^i assigned to individual prediction i at time t is

$$\omega_t^i = \frac{1/\phi_t^i}{\sum_{j=1}^{N}(1/\phi_t^j)} \quad \text{for } \phi_t^i = \sum_{s=q}^{t-1}\theta^{t-1-s}(\overline{z}_{s+1}^{(K)} - \widehat{\mu}_{s+1}^i)^2, \tag{9}$$

where $\widehat{\mu}_{s+1}^i$ is the i-th individual prediction for month $s+1$ and θ, if less than one, attributes higher weights to more recent predictions. We report results obtained with a hold-out period of 60 months, $q=120$, and $\theta=0.9$, following Rapach, Strauss, and Zhou (2010). We have also verified that other parameter sets yield similar results.

Panel B of Table 6 shows that the combination forecasts preserve the overall picture from Panel A of Table 6. For example, the highest out-of-sample R^2 of 10.7% is obtained for carry strategy 1 with the mean combination forecast, and typically the out-of-sample R^2s from the three forecasts are fairly close. Notably, the p-values for the MSPE-adjusted statistic in Panel B of Table 6 are all below 0.04, confirming the presence of predictable variation in carry trade payoffs.

Panel B of Table 7 reports the Sharpe ratios and skewness for the conditional carry trades. The combination forecasts bring improvement over the unconditional strategy, by delivering higher Sharpe ratios and less negative skewness, but the extent of the improvement is weaker statistically when measured by the stationary bootstrap p-values. In this case, we also find that the results from the combination forecasts, irrespective of the weighting scheme, generally do not dominate the counterparts with two or three predictors considered together.

4. Positioning our results in relation to the literature

The big picture that emerges is that, first, carry trade payoffs appear to be predictable at horizons up to six months, as corroborated by the joint p-values and the marginal significance of the predictors, whereby predictability appears to be stronger with ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$ and weaker with ΔLIQ_t. Second, the p-values for the MSPE-adjusted statistic and the out-of-sample R^2s agree in supporting out-of-sample predictability. Third, the results obtained with the predictability-based decision rule suggest that incorporating predictability can translate into both superior Sharpe ratios and less negative skewness of conditional carry trade payoffs. We examine market timing ability within the framework of Henriksson and Merton (1981). Finally, for most fixed currency pairs, little evidence exists for predictability, as shown in Table A3, implying a certain inherent distinction of the dynamically re-balanced carry trades.

Given the multifaceted evidence on predictability, this section addresses two key questions that could provide further clarity and help position our results with respect to the related literature. First, could predictability be traced to a specific component of carry trade payoffs? Identifying such a component could sharpen understanding of the findings and provide a possible basis for their interpretation. Second, do our predictors preserve their predictive ability in a wider context, when other predictors are present? In case they do, they could be pointing to predictive channels that could have been overlooked previously.

Table 10
Relation to other predictors.

The first three columns report results from the predictive regressions $\bar{z}_{t+1}^{(K)} = b_0 + b_x x_t + \epsilon_{t+1}^{(K)}$, where the single predictors x_t are described in Appendix Appendix C. The estimates of the slope coefficients b_x are displayed along with the corresponding two-sided p-values H[p], based on the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability and the adjusted R^2s, shown as $\overline{R}^2$. The fourth column presents Hodrick (1992) p-values H^m[p] for the respective predictor, obtained in multivariate predictive regressions that also include ΔCRB_t, $\Delta\sigma_t^{\text{fx}}$, and ΔLIQ_t as predictors. The next two columns report out-of-sample R^2's (Campbell and Thompson, 2008), and one-sided p-values for the MSPE-adjusted statistic (Clark and West, 2007), obtained with an expanding window with initial length of 180 months. The last four columns report Sharpe ratios and skewness of the conditional payoffs, obtained by applying a decision rule to take a position in a carry trade at the end of month t if a positive carry trade payoff is predicted for month $t+1$ and do nothing if a negative payoff is predicted. Shown also are stationary bootstrap p-values for the null hypothesis that the conditional strategy does not yield statistically higher Sharpe ratios or lower skewness than the unconditional strategy, which always takes a position in the carry trade, irrespective of the sign of the predicted payoff. Regression intercepts are not reported to save on space.

	In-sample regression				Out-of-sample statistics		Predictability-based decision rule			
Carry strategy	b_x	H[p]	$\overline{R}^2$ (%)	H^m[p]	$\overline{R}^2_{OS}$ (%)	MSPE-adj. p-value	Sharpe ratio	Bootstrap p-value	Skewness	Bootstrap p-value
Change in VIX, ΔVIX_t										
1	−0.03	0.16	0.9	0.47	2.5	0.02	0.51	0.30	−1.52	0.53
2	−0.02	0.33	0.2	0.29	0.1	0.34	0.45	0.51	−0.89	0.44
3	−0.01	0.32	0.2	0.25	0.2	0.33	0.66	0.42	−0.70	0.53
4	−0.02	0.17	0.6	0.49	1.2	0.09	0.66	0.45	0.07	0.46
Change in equity volatility, $\Delta\sigma_t^{\text{equity}}$										
1	−0.02	0.08	0.9	0.51	2.2	0.02	0.36	0.38	−1.48	0.53
2	−0.02	0.11	0.7	0.58	1.5	0.03	0.53	0.45	−1.00	0.49
3	−0.02	0.05	1.5	0.87	1.8	0.07	0.80	0.36	−0.28	0.13
4	−0.01	0.04	1.3	0.85	1.8	0.06	0.63	0.52	−0.17	0.46
Change in average interest rate of 10-year bonds, $\Delta\text{Long rate}_t$										
1	0.02	0.04	1.1	0.17	2.5	0.03	0.47	0.24	−1.21	0.32
2	0.01	0.04	1.0	0.19	1.9	0.04	0.57	0.40	−0.95	0.43
3	0.01	0.21	0.1	0.47	0.3	0.24	0.71	0.46	−0.65	0.47
4	0.01	0.14	0.4	0.33	0.7	0.13	0.64	0.51	−0.34	0.64
Average term premium, Term_t										
1	0.19	0.19	0.4	0.90	1.3	0.02	0.29	0.49	−1.39	0.43
2	0.15	0.15	0.5	0.74	1.4	0.04	0.64	0.29	−0.64	0.07
3	0.22	0.02	2.1	0.21	4.0	0.00	0.87	0.28	−0.35	0.14
4	0.19	0.01	2.5	0.13	4.2	0.00	0.79	0.28	−0.06	0.34

4.1. *Carry trade payoffs inherit their predictability from the long legs and currency component*

Returning to the construction of carry trade payoffs in Eq. (2), we first consider the long and short legs of each carry trade separately and investigate the predictability of their payoffs in isolation. Panel A of Table 9 shows predictability results for the short legs of the trades and offers sharp contrast to Panel B of the same table, which considers the results for the long legs. We maintain the predictive setting with all three predictors together.

Whereas all joint H[p] values are above 0.40 when predicting the short legs, it is revealing that three out of four joint H[p]s are below 0.05 when predicting the long legs of the carry trades. The disparity is also evident in the adjusted R^2, which range between 5.0% and 6.4% for the long-leg regressions, but never exceed 0.6% for the short-leg ones. In addition, while ΔLIQ_t is never statistically significant in Panels A or B, ΔCRB_t remains significant in most of the long-leg regressions. The main point to glean is that carry trade payoffs inherit their predictability from the long legs of the trades, but combining the short and long legs strengthens the evidence of predictability, as validated when comparing Tables 3 and 9.

Shifting the perspective, the interest rate differential is known at the beginning of a carry trade and, hence, we focus on the predictability of the underlying currency component. In the spirit of Eq. (1), we define the short and long currency component of the payoff of a carry trade with one short and long pair of currencies as $1-(S_t^{\text{lowest}}/S_{t+1}^{\text{lowest}})$ and $(S_t^{\text{highest}}/S_{t+1}^{\text{highest}})-1$, and we construct the composite currency component as the average. Our construction generalizes to carry trades with several currency pairs, following Eq. (2). We examine the question whether the predictability of carry trade payoffs emanates from the currency component.

Panel C of Table 9 shows that the currency components of carry trade payoffs appear to be predictable and elicits three observations. First, all joint Hodrick p-values are below 0.05. Second, ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$ are statistically significant in all cases. Third, the adjusted R^2 are between 6.5% and 8.4%, and they are comparable to those obtained for the carry trade payoffs. In essence, the predictability of the carry trades can be tied to the currency component.

We recognize that the above conclusion could be viewed as being at odds with studies that do not detect short-run currency predictability, notably, Meese and Rogoff (1983), Diebold and Nason (1990), Chinn and

Meese (1995), and Engel, Mark, and West (2008). This apparent discrepancy could be reconciled from two angles. First, we stress that it is the dynamically re-balanced currency component that is being predicted, and we have verified that the currency components of fixed-pair carry trade payoffs are generally not predictable with our set of predictors. Second, certain variables have achieved some success in forecasting currency changes, as articulated, for example, by Hau and Rey (2006), Della Corte, Sarno, and Tsiakas (2009), Adrian, Etula, and Shin (2010), Ang and Chen (2010), Lustig, Roussanov, and Verdelhan (2010), and Sarno, Schneider, and Wagner (2012). Therefore, the findings from the dynamically re-balanced carry trades complement the ongoing debate on currency predictability.

4.2. Further discussion of the merit of our predictors

Next, we examine variables that have been considered as predictors of currency risk premiums and changes in exchange rates against the US dollar, and we explore whether they also help to predict the payoffs of our dynamically re-balanced carry trades. Specifically, we consider the (i) change in VIX, denoted by ΔVIX_t (Brunnermeier, Nagel, and Pedersen, 2009), (ii) change in average equity return volatility, denoted by $\Delta\sigma_t^{\text{equity}}$, and employed in the cross-sectional tests of Lustig, Roussanov, and Verdelhan (2011), (iii) change in average interest rate on 10-year bonds across the countries in our sample, denoted by $\Delta\text{Long rate}_t$, and (iv) average term premium across the countries in our sample, denoted by Term_t. The impetus for considering $\Delta\text{Long rate}_t$ and Term_t comes from Ang and Chen (2010, Table 6), where the respective country-specific yield curve variables are shown to display explanatory power in panel estimations involving returns of currency portfolios. Appendix C provides more detail on the construction of these predictors.

To gauge the in-sample performance of these additional predictors, we examine the significance of the slope estimate in univariate regressions, as reflected in the H[p]'s, and marginal significance as measured by the respective predictor's p-value, denoted by $\text{H}^{\text{m}}[p]$, in multivariate regressions including ΔCRB_t, $\Delta\sigma_t^{\text{fx}}$, and ΔLIQ_t as well. We also examine out-of-sample statistics and the predictability-based decision rule, and Table 10 summarizes the results.

Turning to the findings, ΔVIX_t is never significant in-sample in univariate regressions, or in regressions with our three predictors, and exhibits high p-values in all cases for the MSPE-adjusted statistic, the incremental improvement in Sharpe ratios, and the reduction of the negative skewness offered by the conditional strategies.[7] Next, $\Delta\sigma_t^{\text{equity}}$ appears to show some promise in predicting in-sample carry trade payoffs and yields MSPE-adjusted p-values below 0.05, but it does not produce robust improvements in the Sharpe ratios. Further, we observe in-sample predictability with $\Delta\text{Long rate}_t$ (Term_t) for strategies 1 and 2 (3 and 4), which is also corroborated by the out-of-sample MSPE-adjusted statistic. However, the stand-alone predictive power of the four additional predictors diminishes when our three predictors are included, as reflected in the high p-values $\text{H}^{\text{m}}[p]$ (the minimum p-value is 0.13). Furthermore, the stationary bootstrap-based p-values for the predictability-based decision rule are all higher than 0.1, implying lack of statistical improvement in Sharpe ratios and skewness with the additional predictors.[8]

The main point of this section is that the predictability of carry trade payoffs with our predictors is likely linked to predictable variation in their currency components, as well as in their long components.

5. Link of predictability to risk factors

This section offers empirical tests that are aimed to relate the documented predictability to an asset pricing model. First, we test whether expected payoffs across the four carry trades move together, using a latent-variable model (as developed in Hansen and Hodrick, 1983; see also Campbell and Hamao, 1992; Ilmanen, 1995). Next, we employ the GMM-based methodology developed in Kirby (1998) to evaluate whether one can reconcile the predictability of carry trade payoffs with linear factor pricing models (Cochrane, 2005, Chapter 13).

In what follows, denote by E_t the expectation operator conditioned on the information set at time t.

5.1. Evidence for a time-varying risk premium interpretation of predictability

To facilitate an interpretation of predictability, we first consider a testing framework that examines whether carry trade payoffs could be consistent with a single latent-variable model. The model maintains that the expected excess return λ_t of the risk factor is linear in a set of instruments (constant plus predictors), which we take to be $\mathbf{x}_t^* \equiv [1\ \mathbf{x}'_t]' = [1 \Delta\text{CRB}_t \Delta\sigma_t^{\text{fx}} \Delta\text{LIQ}_t]'$, and that the conditional betas are constant. Hence

$$E_t(\mathbf{z}_{t+1}) = \boldsymbol{\beta}\lambda_t = \boldsymbol{\beta}\mathbf{c}'\mathbf{x}_t^* \quad \text{where } \boldsymbol{\beta} = [\beta_1\ \ \beta_2\ \ \beta_3\ \ \beta_4]' \quad \text{and}$$
$$\mathbf{c} = [c_0\ \ c_1\ \ c_2\ \ c_3]' \quad \text{are constant vectors.} \tag{10}$$

We normalize $\beta_1 = 1$ to allow identification. Within this framework, the relation between the predictors and the expected payoffs to different carry trades is synthesized by a latent variable.

[7] We explore two additional predictors: the level of the CRB spot index relative to an analogous futures index (equally weighted CCI, from Datastream) and VIX squared minus realized variance (Bollerslev, Tauchen, and Zhou, 2009, IV minus RVH). Based on Hodrick (1992) p-values, these variables were not found to have statistically significant predictive power for carry trades.

[8] We also investigate robustness with respect to three other predictive variables: (i) log change in crude oil (e.g., Kilian, 2009; Ferraro, Rogoff, and Rossi, 2011), (ii) average forward discount (e.g., Hansen and Hodrick, 1983; Bekaert and Hodrick, 1992; Lustig, Roussanov, and Verdelhan, 2010), and (iii) growth rate of OECD industrial production (see Appendix C for the details on variable construction). Our evidence indicates that each of these variables is insignificant in the presence of our predictors.

Table 11
GMM tests of a latent-variable characterization of carry trade payoffs.

Reported results are based on the GMM estimation of the single-latent-variable model Eqs. (11) and (12), whereby $E_t(\mathbf{z}_{t+1})=\boldsymbol{\beta}\mathbf{c}'\mathbf{x}_t^*$, with $\mathbf{x}_t^*=[1\mathbf{x}'_t]'=[1\ \Delta\text{CRB}_t\ \Delta\sigma_t^{\text{fx}}\ \Delta\text{LIQ}_t]'$. The table also presents results obtained with subsets of $\mathbf{x}_t^*$. Let $\mathbf{u}_{t+1}=\mathbf{z}_{t+1}-\boldsymbol{\beta}\mathbf{c}'\mathbf{x}_t^*$. Then under the null that the model is true, $g_T(\Theta)\equiv E(\mathbf{u}_{t+1}\otimes\mathbf{x}_t^*)=\mathbf{0}$. We minimize the quadratic form $g_T(\Theta)'\Omega_T g_T(\Theta)$, where Ω_T is a positive-definite weighting matrix and $\Theta=[\boldsymbol{\beta}'\mathbf{c}']'$ is the parameter vector. Shown in square brackets are the p-values for the coefficient estimates. The minimized value (multiplied by T), J_T, is χ^2-distributed with degrees of freedom (df) equal to the number of orthogonality conditions minus the number of parameters. We display the p-values corresponding to J_T in ⟨ ⟩.

	$\mathbf{z}_{t+1}=[z_{t+1}^{(1)}\ z_{t+1}^{(2)}\ z_{t+1}^{(3)}\ z_{t+1}^{(4)}]'$ subset of $\mathbf{x}_t^*$ employed				$\mathbf{z}_{t+1}=[z_{t+1}^{(1)}\ z_{t+1}^{(4)}]'$ subset of $\mathbf{x}_t^*$ employed			
	Constant ΔCRB_t $\Delta\sigma_t^{\text{fx}}$ –	Constant ΔCRB_t – ΔLIQ_t	Constant – $\Delta\sigma_t^{\text{fx}}$ ΔLIQ_t	Constant ΔCRB_t $\Delta\sigma_t^{\text{fx}}$ ΔLIQ_t	Constant ΔCRB_t $\Delta\sigma_t^{\text{fx}}$ –	Constant ΔCRB_t – ΔLIQ_t	Constant – $\Delta\sigma_t^{\text{fx}}$ ΔLIQ_t	Constant ΔCRB_t $\Delta\sigma_t^{\text{fx}}$ ΔLIQ_t
c_0 (Constant)	0.003	0.002	0.003	0.002	0.001	0.001	0.002	0.001
	[0.02]	[0.04]	[0.01]	[0.05]	[0.26]	[0.34]	[0.14]	[0.31]
c_1 (ΔCRB_t)	0.152	0.164		0.138	0.184	0.220		0.176
	[0.00]	[0.00]		[0.01]	[0.00]	[0.00]		[0.00]
c_2 ($\Delta\sigma_t^{\text{fx}}$)	−0.041		−0.043	−0.039	−0.045		−0.049	−0.041
	[0.00]		[0.00]	[0.00]	[0.00]		[0.00]	[0.00]
c_3 (ΔLIQ_t)		0.023	0.017	0.013		0.021	0.016	0.013
		[0.02]	[0.06]	[0.11]		[0.03]	[0.13]	[0.19]
β_2	0.349	0.414	0.354	0.334				
	[0.03]	[0.03]	[0.07]	[0.04]				
β_3	0.807	0.744	0.808	0.777				
	[0.00]	[0.00]	[0.00]	[0.00]				
β_4	0.339	0.256	0.366	0.419	0.306	0.241	0.379	0.362
	[0.09]	[0.26]	[0.10]	[0.03]	[0.13]	[0.29]	[0.10]	[0.06]
J_T-statistic	4.10	3.51	3.31	4.76	0.37	0.46	1.05	1.06
	⟨0.66⟩	⟨0.74⟩	⟨0.77⟩	⟨0.85⟩	⟨0.83⟩	⟨0.79⟩	⟨0.59⟩	⟨0.78⟩
df	6	6	6	9	2	2	2	3

The restrictions that the latent-variable model imposes can be tested by specifying

$$\begin{bmatrix} u_{t+1}^{(1)} \\ u_{t+1}^{(2)} \\ u_{t+1}^{(3)} \\ u_{t+1}^{(4)} \end{bmatrix} = \begin{bmatrix} z_{t+1}^{(1)} \\ z_{t+1}^{(2)} \\ z_{t+1}^{(3)} \\ z_{t+1}^{(4)} \end{bmatrix} - \begin{bmatrix} c_0 & c_1 & c_2 & c_3 \\ \beta_2 c_0 & \beta_2 c_1 & \beta_2 c_2 & \beta_2 c_3 \\ \beta_3 c_0 & \beta_3 c_1 & \beta_3 c_2 & \beta_3 c_3 \\ \beta_4 c_0 & \beta_4 c_1 & \beta_4 c_2 & \beta_4 c_3 \end{bmatrix} \begin{bmatrix} 1 \\ \Delta\text{CRB}_t \\ \Delta\sigma_t^{\text{fx}} \\ \Delta\text{LIQ}_t \end{bmatrix}. \tag{11}$$

Under the null that the model holds (Cochrane, 2005, Chapter 13)

$$g_T(\Theta)\equiv E(\mathbf{u}_{t+1}\otimes\mathbf{x}_t^*)=\mathbf{0} \quad \text{where } \mathbf{u}_{t+1}=[u_{t+1}^{(1)}\ u_{t+1}^{(2)}\ u_{t+1}^{(3)}\ u_{t+1}^{(4)}]'. \tag{12}$$

Applying GMM, we minimize $g_T(\Theta)'\Omega_T g_T(\Theta)$ to estimate the parameter vector $\Theta\equiv[\boldsymbol{\beta}'\mathbf{c}']'$, where Ω_T is a weighting matrix. The minimized value of the criterion function, multiplied by the number of observations T (J_T), is χ^2 distributed with degrees of freedom (df) equal to the number of orthogonality conditions minus the number of parameters in Θ.

Table 11 reports the estimation results for the system in Eq. (11), where p-values corresponding to the parameters are shown in square brackets. For robustness, we employ at most four instruments (including the constant), so as to keep the number of moment conditions to a manageable level, and we also consider a restricted payoff vector $\mathbf{z}_{t+1}=[z_{t+1}^{(1)}\ z_{t+1}^{(4)}]'$. Consistent with Tables 2 and 3, the coefficients on ΔCRB_t and ΔLIQ_t (c_1 and c_3) are positive, and the coefficient on $\Delta\sigma_t^{\text{fx}}$ (c_2) is negative. In addition, four out of the six p-values on c_3 are above 0.05, affirming the weaker predictive power of ΔLIQ_t for carry trade payoffs.

The coefficients β_2 and β_3 are statistically significant. Furthermore, in all cases the overidentifying restrictions imposed by the model are not rejected, as indicated by the p-values, reported below the J_T statistic. For example, the lowest p-value is 0.59, obtained with $\mathbf{z}_{t+1}=[z_{t+1}^{(1)}\ z_{t+1}^{(4)}]'$. In summary, the GMM tests suggest a time-varying risk premium interpretation of the documented predictability of carry trade payoffs.

5.2. Extant risk factors might not be able to reconcile the predictability of carry trade payoffs

This section considers observable risk factors, keeping in mind recent advances in the modeling of currency excess returns. The specific goal is to investigate whether the two currency-related risk factors, proposed in Lustig, Roussanov, and Verdelhan (2009) and Menkhoff, Sarno, Schmeling, and Schrimpf (2011), could provide a link

Table 12
Predictability implicit in an asset pricing model with currency returns and innovations in currency volatility as risk factors.

The GMM results reported in the first six columns of the table are based on the moment conditions from the system in Eq. (17) and help evaluate the restrictions that the model equation (13) imposes on the coefficients of the predictive regression $\overline{z}^{(K)}_{t+1} = \mathbf{x}^{*\prime}_t \mathbf{b} + \epsilon^{(K)}_{t+1}$, with $\mathbf{b} = [b_0\ \mathbf{b}'_x]'$ and $\mathbf{x}^*_t = [1\,\Delta\text{CRB}_t\ \Delta\sigma^{\text{fx}}_t\ \Delta\text{LIQ}_t]'$. Denote the vectors of unrestricted and restricted coefficients in the predictive regressions as $\mathbf{b}^u$ and $\mathbf{b}^r$, respectively, and define $\mathbf{b}^* = \mathbf{b}^u - \mathbf{b}^r$. The first four columns present the components of $\mathbf{b}^*$, with p-values in square brackets. The fifth column presents the Wald statistic W^a_T, distributed $\chi^2(3)$, that tests $\mathbf{b}^*_x = \mathbf{b}^u_x - \mathbf{b}^r_x = \mathbf{0}$, with p-values in ⟨⟩. The sixth column shows an analogous Wald statistic, denoted by W^b_T and distributed $\chi^2(4)$, that tests $\mathbf{b}^* = \mathbf{0}$ for each predictive regression. Next, we report the J_T statistic of model validity, distributed $\chi^2(4)$, and calculated by imposing the constraint $\mathbf{b}^u = \mathbf{b}^r$, which makes the system overidentified. The GMM tests follow the methodology in Kirby (1998) and are designed to assess whether the asset pricing model equation (13) is consistent with the predictability of carry trade payoffs with our three predictors. Two risk factors are used (as described in the text and as detailed in Appendix C): (i) average monthly excess return of the currencies in the sample, held long against the US dollar, $f_{t+1,1}$, and (ii) mimicking portfolio for currency volatility, $f_{t+1,2}$.

	Difference between unrestricted and restricted regression coefficients, $\mathbf{b}^*$				Wald tests		J-statistic
	Intercept	ΔCRB_t	$\Delta\sigma^{\text{fx}}_t$	ΔLIQ_t	W^a_T $\chi^2(3)$	W^b_T $\chi^2(4)$	J_T $\chi^2(4)$
$\overline{z}^{(1)}_{t+1}$	0.001	0.185	−0.004	0.021	20.76	27.46	16.21
	[0.35]	[0.01]	[0.02]	[0.13]	⟨0.00⟩	⟨0.00⟩	⟨0.00⟩
$\overline{z}^{(2)}_{t+1}$	0.001	0.129	−0.002	0.016	14.16	19.48	16.37
	[0.25]	[0.01]	[0.05]	[0.16]	⟨0.00⟩	⟨0.01⟩	⟨0.00⟩
$\overline{z}^{(3)}_{t+1}$	0.002	0.088	−0.002	0.016	17.82	28.46	19.90
	[0.01]	[0.03]	[0.02]	[0.06]	⟨0.00⟩	⟨0.00⟩	⟨0.00⟩
$\overline{z}^{(4)}_{t+1}$	0.002	0.075	−0.002	0.010	24.35	34.59	20.58
	[0.02]	[0.02]	[0.00]	[0.09]	⟨0.00⟩	⟨0.00⟩	⟨0.00⟩

between expected carry trade payoffs and predictability in the context of factor pricing models.

We first examine the restrictions on predictability implied by the following model (Cochrane, 2005, Chapter 8.3):

$$E_t(\overline{z}^{(K)}_{t+1} m_{t+1}) = 0 \quad \text{with } m_{t+1} = 1 - \mathbf{f}_{t+1}\mathbf{d}' \quad \text{and} \quad E(\mathbf{f}_{t+1}) = \mathbf{0}, \tag{13}$$

where m_{t+1} is a stochastic discount factor with $E(m_{t+1}) = 1$, $\mathbf{f}_{t+1} = [f_{t+1,1} f_{t+1,2}]'$ is a vector of factors, and $\mathbf{d}$ is a constant vector. In the empirical exercise, $f_{t+1,1}$ is the average monthly excess return of our eight currencies, held long against the US dollar (similar to Lustig, Roussanov, and Verdelhan, 2011, p. 3745; see Appendix C for details on the construction), and $f_{t+1,2}$ is the mimicking portfolio for the innovations in currency volatility [similar to Menkhoff, Sarno, Schmeling, and Schrimpf, 2012, Eq. (1)]. Both risk factors are in excess return over the US risk-free return, denominated in US dollars and demeaned, and have shown promise in describing the cross section of currency excess returns.

Can the pricing model equation (13) reproduce the predictability patterns established in our Table 3? To address this question, we appeal to the approach developed in Kirby (1998). First, let $\mathbf{b} = [b_0\ \mathbf{b}'_x]'$, and rewrite the regression equation (6) as

$$\overline{z}^{(K)}_{t+1} = \mathbf{x}^{*\prime}_t \mathbf{b} + \epsilon^{(K)}_{t+1},$$
where $\mathbf{x}^*_t \equiv [1\,\Delta\text{CRB}_t\ \Delta\sigma^{\text{fx}}_t\ \Delta\text{LIQ}_t]'$ and
$$\mathbf{b} = (E(\mathbf{x}^*_t \mathbf{x}^{*\prime}{}_t))^{-1} E(\overline{z}^{(K)}_{t+1}\mathbf{x}^*_t). \tag{14}$$

The testable restriction that the pricing model equation (13) imposes on the regression coefficients in the linear forecasting model is

$$\mathbf{b} = -(E(\mathbf{x}^*_t \mathbf{x}^{*\prime}{}_t))^{-1}\,\text{cov}(m_{t+1}, \overline{z}^{(K)}_{t+1}\mathbf{x}^*_t), \tag{15}$$

analogous to Kirby (1998, Eq. (13)). Briefly, the restriction equation (15) follows from multiplying both sides of $E_t(\overline{z}^{(K)}_{t+1} m_{t+1}) = 0$ in Eq. (13) by $\mathbf{x}^*_t$, then using the law of iterated expectations and substituting the resulting unconditional expectation $E(\overline{z}^{(K)}_{t+1}\mathbf{x}^*_t)$ in Eq. (14).

Proceeding with the econometric evaluation, let $\mathbf{b}^u = [b^u_0 (\mathbf{b}^u_x)']'$ and $\mathbf{b}^r = [b^r_0 (\mathbf{b}^r_x)']'$ be, respectively, the vector of unrestricted [as in Eq. (14)] and restricted [as in Eq. (15)] regression coefficients. Define

$$\mathbf{b}^* = \mathbf{b}^u - \mathbf{b}^r \quad \text{and} \quad \mathbf{b}^*_x = \mathbf{b}^u_x - \mathbf{b}^r_x. \tag{16}$$

Then whether the model in Eq. (13) is consistent with predictability can be evaluated in a GMM setting with 10 moment conditions:

$$E\begin{pmatrix} (1-\mathbf{f}_{t+1}\mathbf{d}')\otimes\mathbf{f}_{t+1} \\ (\overline{z}^{(K)}_{t+1} - \mathbf{x}^{*\prime}_t \mathbf{b}^u)\otimes\mathbf{x}^*_t \\ (\overline{z}^{(K)}_{t+1}\mathbf{f}_{t+1}\mathbf{d}' - \mathbf{x}^{*\prime}_t(\mathbf{b}^u - \mathbf{b}^*))\otimes\mathbf{x}^*_t \end{pmatrix} = 0, \tag{17}$$

similar to Eq. (26) in Kirby (1998), and given our normalization $E(m_{t+1}) = 1$.

The system in Eq. (17) treats $\mathbf{b}^*$ as a free parameter vector and is just identified. If the pricing model quantitatively matches the predictability relation, then the components of $\mathbf{b}^*$ should be statistically indistinguishable from zero. Such a GMM system allows for a comparison between unrestricted and restricted regression coefficients. We also estimate the system by imposing the constraint $\mathbf{b}^u = \mathbf{b}^r$, whereby the system becomes overidentified. Then we calculate Hansen's J_T-statistic for

Table A1
Funding versus investment currencies in carry trades.

At the end of each month over the sample period January 1985 to August 2011, we rank-order the G-10 currencies (excluding the euro) according to their interest rate differential against the US dollar, which is inferred from the forward discount and based on mid-quotes for spot and one-month forward exchange rates. The column labeled “Lowest” shows how many months (out of 320) the respective currency has been the lowest-yielding and, hence, the most attractive funding vehicle. The next three columns show likewise how many months has the respective currency been among the two, three, or four lowest-yielding ones. Analogous numbers in the last four columns show how often a currency has been the highest-yielding, or among the two, three or four highest-yielding ones. For example, on three (78) occasions the US dollar has been the highest-yielding currency (among the four highest-yielding ones) and, hence, has been an attractive investment vehicle, and on 31 (204) occasions it has been the lowest-yielding currency (among the four lowest-yielding ones).

	Lowest-yielding currencies (number of months)				*Highest-yielding currencies* (number of months)			
Currency	Lowest	Two lowest	Three lowest	Four lowest	Four highest	Three highest	Two highest	Highest
Australian dollar, AUD	0	3	10	30	276	228	162	54
Canadian dollar, CAD	1	3	47	167	59	22	5	3
Swiss franc, CHF	78	260	284	293	14	3	1	1
British pound, GBP	0	1	10	33	239	152	67	17
Japanese yen, JPY	206	308	314	316	4	2	2	1
Norwegian krone, NOK	1	1	37	105	188	162	89	46
New Zealand dollar, NZD	0	0	8	25	277	265	234	152
Swedish krona, SEK	3	11	73	107	145	83	68	43
US dollar, USD	31	53	177	204	78	43	12	3

overidentifying restrictions, and the corresponding p-value based on χ^2-statistic with four degrees of freedom.

The first four columns in Table 12 present the $\mathbf{b}^*$ coefficient estimates when all three predictors are employed, with p-values in square brackets. The predictability of carry trade payoffs with ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$ appears to be inconsistent with the model, with all corresponding p-values below 0.05. However, the equality of the unrestricted and restricted coefficients for ΔLIQ_t is not rejected, with all p-values above 0.06.

We also present the Wald statistics W_T^a, distributed $\chi^2(3)$, for the test of the null hypothesis that $\mathbf{b}_x^* = \mathbf{0}$ in the case with three predictors, with p-values in $\langle\ \rangle$. The W_T^a statistics are also not supportive of the pursued link between factor pricing and predictability and reject in all four cases, with zero p-values. These rejections testify to the high hurdle that predictability poses to asset pricing models.

Further evidence based on the Wald statistic, denoted by W_T^b, is presented in Table 12 and refers to joint tests of coefficient equality that include the intercepts as well, i.e., the null hypothesis is $\mathbf{b}^* = \mathbf{0}$. Going beyond the evaluation of a model strictly according to its ability to explain predictability, such tests could yield additional insight, based on the observation that a properly specified asset pricing model should also produce the correct intercept in the predictive regression and, hence, be able to explain average returns as well. The W_T^b statistic rejects the equality of the coefficients, with p-values equal to zero.

Table 12 also shows that the p-values for the J_T statistics are all equal to zero, indicating model inadequacy. While we recognize that the considered risk factors have proved versatile in capturing the return cross section of currency portfolios (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012), our estimations suggest that they could fail to internalize the extent of the predictability associated with ΔCRB_t and $\Delta\sigma_t^{\text{fx}}$.

Expanding on the above analysis, we follow Cochrane (1996, p. 583) and Lettau and Ludvigson (2001, p. 1245), and we modify the m_{t+1} specification in Eq. (13) to include scaled risk factors

$$m_{t+1} = 1 - f_{1,t+1}(d_{1,0} + d_{1,x}x_t) - f_{2,t+1}(d_{2,0} + d_{2,x}x_t), \quad (18)$$

$$= 1 - f_{1,t+1}d_{1,0} - f_{2,t+1}d_{2,0} - f_{1,t+1}x_t d_{1,x} - f_{2,t+1}x_t d_{2,x} \equiv 1 - \hat{\mathbf{f}}_{t+1}\hat{\mathbf{d}}', \quad (19)$$

where $E(\hat{\mathbf{f}}_{t+1}) = 0$ and x_t is one of our three predictors. Replacing $\mathbf{f}_{t+1}$ with $\hat{\mathbf{f}}_{t+1}$ and $\mathbf{d}'$ with $\hat{\mathbf{d}}'$ in Eq. (17), we estimate the GMM system with 12 moment conditions. We still obtain zero p-values for W_T^a, W_T^b, and J_T (not reported). Thus, the specifications in Eqs. (13) and (19) both yield model rejection. Our analysis suggests the need to refine the search for currency-related asset pricing models.

6. Conclusions and summary

This paper contributes by investigating the time series predictability of currency carry trades, whereas extant studies of these trades have mostly examined the cross-sectional aspects of their returns. We find predictable time variation in the payoffs of dynamically re-balanced carry trades, constructed by rank-ordering currencies according to their forward discounts. Specifically, we show that commodity returns, average currency volatility, and a measure of global liquidity predict currency carry trade payoffs. These predictors are not highly correlated among each other, indicating differences in their economic nature. Our evidence indicates that higher commodity returns predict higher carry trade payoffs, while higher average currency volatility or lower liquidity is accompanied by lower next-month carry trade payoffs.

We establish the predictive ability of our predictors from an in-sample perspective at horizons up to six months, whereby we find significant marginal and joint in-sample predictive power, evidenced by the p-values constructed from the covariance estimator of Hodrick

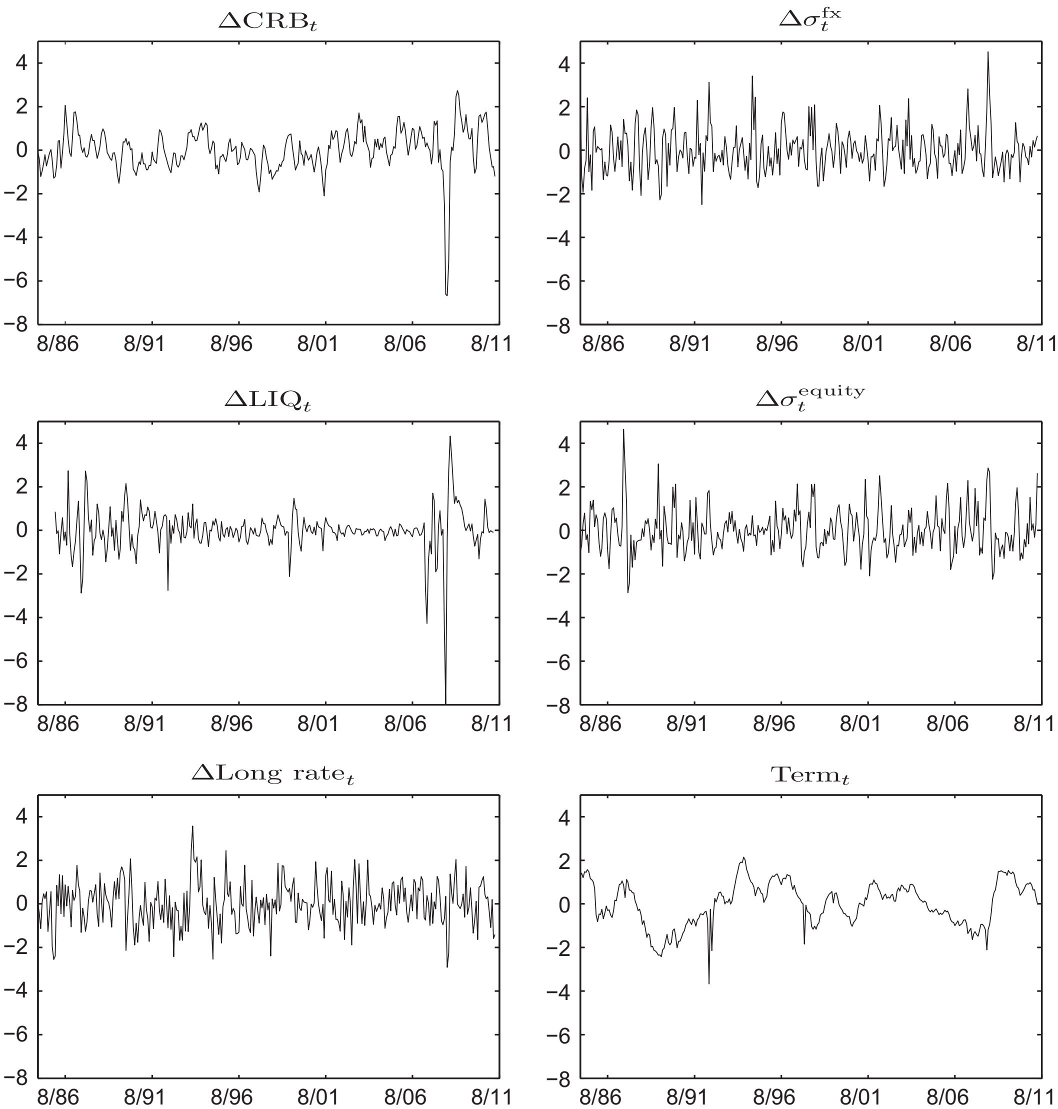

Fig. 3. Time variation in the predictors. Plotted are various predictors each standardized here to have *zero mean and unit variance*. Our three individual predictors are defined as

$$\Delta \text{CRB}_t = \frac{1}{3}\log(\text{CRB}_t/\text{CRB}_{t-3}), \quad \Delta\sigma_t^{\text{fx}} = \frac{1}{3}\log(\sigma_t^{\text{avg}}/\sigma_{t-3}^{\text{avg}}), \quad \text{and} \quad \Delta \text{LIQ}_t = -\left(\text{LIQ}_t^{\text{avg}} - \frac{1}{3}\sum_{j=1}^{3}\text{LIQ}_{t-j}^{\text{avg}}\right).$$

CRB_t is the Raw Industrials subindex of the CRB Spot Commodity Index, σ_t^{avg} is the average currency volatility for month t across the G-10 currencies, where volatility is computed as the square root of the average squared daily log change over a month of a currency's spot exchange rate against the US dollar, and $\text{LIQ}_t^{\text{avg}}$ is the average TED spread (i.e., 3-month Libor minus the 3-month Treasury bill rate or its equivalent) for month t across the G-10 currencies in our sample. The equity volatility predictor, $\Delta\sigma_t^{\text{equity}}$, is the average monthly equity volatility across the countries in the sample, whereby for each country, equity volatility for month t is computed as the square root of the average squared daily log change over the month of the country's equity index (total return index, in local currency). ΔLong rate_t is change in the interest rates of 10-year bonds, averaged across the countries in our sample, and Term $_t$ is the term premium averaged across the countries in our sample. The country term premium is the difference between the interest rate of a 10-year bond and one-month interest rate (Libor or its equivalent). Our data are from Datastream and Global Financial Data.

(1992) and via parametric bootstrap. Out-of-sample statistics and combination forecasts provide further evidence for predictability. We also find that conditional carry trade strategies that exploit trading signals generated using our predictors enhance Sharpe ratios and mitigate negative skewness, relative to the unconditional carry trades. Standard errors obtained by bootstrap reveal that these improvements are also statistically significant.

Probing deeper into these findings, we perform GMM tests that show consistency of the documented predictability with a latent-variable model. This evidence is supportive of a link between the time variation in risk premiums and the predictability of carry trade payoffs. Further differentiating our work from others, we examine whether a pricing model with observable risk factors, that we take to be average currency returns and innovations in

Table A2
Properties of the predictors.
Reported are the mean, standard deviation, skewness, kurtosis, and first-order autocorrelation (denoted by ACF_1). The three predictors are defined as

$$\Delta CRB_t = \frac{1}{3}\log(CRB_t/CRB_{t-3}), \Delta\sigma_t^{fx} = \frac{1}{3}\log(\sigma_t^{avg}/\sigma_{t-3}^{avg}) \quad \text{and} \quad \Delta LIQ_t = -\left(LIQ_t^{avg} - \frac{1}{3}\sum_{j=1}^{3} LIQ_{t-j}^{avg}\right).$$

CRB_t is the Raw Industrials subindex of the CRB Spot Commodity Index, σ_t^{fx} is the average currency volatility for month t across the G-10 currencies, where currency volatility is computed as the square root of the average squared daily log change over a month of a currency's spot exchange rate against the US dollar, and LIQ_t is the average TED spread (i.e., three-month Libor minus the three-month Treasury bill rate or its equivalent) for month t across the G-10 currencies in our sample (except for NOK and NZD, for which respective data could not be obtained for the full sample period). Also presented are contemporaneous correlations among the three predictors and with the carry trade payoffs. For each predictor, we estimate via maximum likelihood all models in the ARMA–GARCH(1,1) class, for AR and MA orders up to three, and choose the best model according to the BIC. These models are employed in our parametric bootstrap and to establish the stationarity of the predictors. Reported are the coefficient estimates for the best models, with asymptotic p-values in square brackets.

Panel A: Summary statistics and correlation matrix

	Mean	Stdev	Skewness	Kurtosis	ACF_1	Contemporaneous correlations $\Delta\sigma_t^{fx}$	ΔLIQ_t	$\bar{z}_t^{(1)}$	$\bar{z}_t^{(2)}$	$\bar{z}_t^{(3)}$	$\bar{z}_t^{(4)}$
ΔCRB_t	2.7E−03	0.021	−2.03	16.24	0.81	−0.24	0.19	0.24	0.25	0.24	0.22
$\Delta\sigma_t^{fx}$	−2.4E−05	0.096	0.76	3.59	0.30		−0.26	−0.31	−0.32	−0.35	−0.31
ΔLIQ_t	3.8E−03	0.136	−1.84	20.09	0.49			0.20	0.22	0.27	0.20

Panel B: ARMA–GARCH(1,1) models for the predictors

	Const.	AR(1)	AR(2)	AR(3)	MA(1)	MA(2)	MA(3)	Const.	ARCH	GARCH	BIC
ΔCRB_t	0.00	0.20			1.04	1.00		0.00	0.43		−6.6
	[0.00]	[0.00]			[0.00]	[0.00]		[0.00]	[0.00]		
$\Delta\sigma_t^{fx}$	0.00				0.36	0.31	−0.48	0.00	0.12	0.65	−2.2
	[0.82]				[0.00]	[0.00]	[0.00]	[0.15]	[0.10]	[0.00]	
ΔLIQ_t	−0.01	0.45	−0.13	−0.15				0.00	0.18	0.79	−1.9
	[0.20]	[0.00]	[0.08]	[0.01]				[0.00]	[0.00]	[0.00]	

currency volatility, can come close to mimicking the coefficients from our predictive regressions. Here we apply the econometric approach of Kirby (1998) and find that the adopted risk factors face difficulty replicating the pattern of predictability manifested in the data.

Given our results, one avenue of future research is to clarify on other macroeconomic predictors of carry trade payoffs. While an active body of literature links macroeconomic fundamentals to exchange rates, an analogous link to carry trade payoffs warrants further development. Some efforts could be directed toward an examination of pricing models in the context of carry trade predictability and possibly integrating the cross-sectional and time series aspects (e.g., Bali and Engle, 2010). Finally, more work could be done on examining whether information extracted from currency derivatives (straddles, risk reversals, or strangle margins), which have been found to be related to risk premiums in Bakshi, Carr, and Wu (2008), could help predict carry trade payoffs, when aggregated across countries, or can provide additional risk factors. Recent work on using factor models to predict currency returns (Engel, Mark, and West, 2009) testifies for the interest in this line of research.

Appendix A. Properties of predictors

This Appendix discusses some statistical features of our three predictors, including their persistence. Panel A of Table A2 displays the summary statistics and the correlations among the predictors. It is seen that the first-order autocorrelations are between 0.30 and 0.81, implying that statistical concerns regarding highly persistent predictors might not arise in our estimations.

In addition, the correlations between the three predictors are relatively low (at most 0.26 in absolute value) and could indicate that the predictors capture distinct aspects of the global macroeconomic environment. Fig. 3 delineates the time series variation in the predictors, each standardized to have zero mean and unit standard deviation for ease of comparison. The contemporaneous correlations between the predictors and carry trade payoffs $\bar{z}_t^{(K)}$ vary between −0.35 and 0.20.

Panel B of Table A2 shows the maximum likelihood estimates for ARMA–GARCH(1,1) models, fitted to the three predictor time series. The reported parameters in each case are for the model, selected according to the BIC among all ARMA–GARCH(1,1) models with AR and MA components up to order three, with asymptotic two-sided p-values displayed in square brackets.

The best model for ΔCRB_t, $\Delta\sigma_t^{fx}$, and ΔLIQ_t is ARMA (1,2)–GARCH(1,1), ARMA(0,3)–GARCH(1,1), and ARMA (3,0)–GARCH(1,1), respectively, all with significant AR and/or MA coefficients. Importantly, the individual and joint two-sided parametric bootstrap p-values reported in Tables 2 and 3 are based on these specifications for the predictors.

Table A3
Predictability for fixed currency pairs.

The first three columns report results from univariate predictive regressions, where the single predictor is ΔCRB_t, $\Delta\sigma_t^{fx}$, or ΔLIQ_t. All predicted payoffs are denominated in US dollars, and we consider six individual fixed currency pairs, whereby JPY and CHF are always shorted against the US dollar, and AUD, NZD and GBP are always bought against the US dollar. The estimates of the slope coefficients b_x are displayed along with the corresponding two-sided p-values H[p], based on the Hodrick (1992) 1B covariance matrix estimator under the null of no predictability, and the adjusted R^2s, shown as $\overline{R}^2$. The next two columns report out-of-sample R^2's (Campbell and Thompson, 2008), and one- sided p-values for the MSPE–adjusted statistic (Clark and West, 2007), obtained with an expanding window with initial length of 180 months. The last four columns report Sharpe ratios and skewness of the conditional payoffs, obtained by applying a decision rule to take a position in a fixed currency pair at the end of month t if a positive payoff for this pair is predicted for month $t+1$ and to do nothing if a negative payoff is predicted. Shown also are stationary bootstrap p-values for the null hypothesis that the conditional strategy does not yield statistically higher Sharpe ratios or lower skewness than the unconditional strategy, which always takes a position in the fixed pair, irrespective of the sign of the predicted payoff. Regression intercepts are not reported to save on space.

		In-sample regression			Out-of-sample statistics		Predictability-based decision rule			
	Carry strategy	b_x	H[p]	$\overline{R}^2$ (%)	$\overline{R}^2_{OS}$ (%)	MSPE-adj. p-value	Sharpe ratio	Bootstrap p-value	Skewness	Bootstrap p-value
ΔCRB_t	JPY–AUD	0.13	0.09	1.1	1.8	0.06	0.31	0.58	-0.66	0.09
	JPY–NZD	0.16	0.08	2.1	2.9	0.06	0.58	0.22	0.11	0.04
	JPY–GBP	0.15	0.06	2.8	5.2	0.07	0.28	0.13	–0.72	0.20
	CHF–AUD	0.05	0.42	–0.1	–2.8	0.83	–0.13	0.76	–2.05	0.95
	CHF–NZD	0.08	0.17	0.4	0.2	0.26	0.23	0.30	–0.45	0.57
	CHF–GBP	0.07	0.47	0.8	–4.4	0.77	–0.51	0.44	–0.96	0.00
$\Delta\sigma_t^{fx}$	JPY–AUD	–0.02	0.10	0.7	1.0	0.10	0.44	0.45	–0.52	0.09
	JPY–NZD	–0.02	0.13	0.6	1.3	0.03	0.32	0.54	–0.88	0.56
	JPY–GBP	–0.02	0.08	1.2	2.1	0.06	0.12	0.30	–0.42	0.08
	CHF–AUD	–0.01	0.68	–0.3	–0.3	0.75	–0.12	0.76	–1.91	0.72
	CHF–NZD	0.00	0.75	–0.3	–1.6	0.89	–0.26	0.86	–0.66	0.86
	CHF–GBP	0.00	0.62	–0.2	–0.8	0.87	–0.47	0.36	–4.51	0.64
ΔLIQ_t	JPY–AUD	0.02	0.15	1.9	1.3	0.19	0.24	0.64	–1.91	0.61
	JPY–NZD	0.03	0.14	2.5	2.3	0.19	0.34	0.55	–1.01	0.43
	JPY–GBP	0.03	0.06	3.9	6.5	0.06	0.29	0.09	–0.41	0.02
	CHF–AUD	0.01	0.53	–0.2	–2.2	0.86	–0.01	0.72	–1.01	0.57
	CHF–NZD	0.01	0.37	0.0	–1.0	0.77	0.02	0.72	0.01	0.28
	CHF–GBP	0.01	0.13	0.4	–0.1	0.25	–0.19	0.15	0.08	0.00

Appendix B. Description of the parametric bootstrap

This Appendix describes the parametric bootstrap procedure that we use, which is intended to safeguard inference in the predictive regressions against the possible impact of persistent predictors and to account for the finite-sample properties of the estimators.

Our implementation of the parametric bootstrap follows Mark (1995), Kilian (1999), and Amihud, Hurvich, and Wang (2009, Appendix C). To provide an illustration in the case of one predictor, let $z_{t+1} = b_0 + b_x x_t + \epsilon_{t+1}$. First, we estimate the parameters $\widehat{b}_0$ and $\widehat{b}_x$ via OLS, compute the Wald statistic $t_{\widehat{b}_x}$, and store the residuals ϵ_{t+1}. Next, given our estimated best ARMA–GARCH model for the predictor and the associated standardized residuals, we build iteratively bootstrapped predictor series x_t^n, for $n = 1, \ldots, 25{,}000$.

Then we build bootstrapped series of the predicted variable under the null $b_x = 0$, as $z_{t+1}^n = \widehat{b}_0 + \epsilon_{t+1}^n$, where ϵ_{t+1}^n are bootstrap samples of the regression residuals ϵ_{t+1}, which are drawn in a pairwise manner, to preserve the dependencies between regression residuals and predictors. Next, we estimate the predictive regression on each bootstrapped pair of predicted and predictor variable, $z_{t+1}^n = b_0^n + b_x^n x_t^n + u_{t+1}^n$, and store the Wald statistics $t_{\widehat{b}_x^n}$. Finally, we compute the proportion of $|t_{\widehat{b}_x^n}|$ s that exceed $|t_{\widehat{b}_x}|$, which yields the bootstrapped p-value for the $\widehat{b}_x$ estimate. An analogous procedure is followed with multiple predictors and to assess joint parameter significance.

Appendix C. Description of other predictor variables and risk factors

All predictors, except for the average forward discount and the term structure variables, are obtained as three-month log changes divided by three, similar to the way that ΔCRB_t and $\Delta\sigma_t^{fx}$ are defined.

- ΔVIX_t: change in the CBOE VIX volatility index, as in Brunnermeier, Nagel, and Pedersen (2009).
- $\Delta\sigma_t^{equity}$: change in average equity volatility, $\sigma_t^{equity,avg}$. For each country, equity volatility for month t is computed as the square root of the average squared daily log change over the month of the country's equity index (total return index, in local currency), and $\sigma_t^{equity,avg}$ is the average across these monthly equity volatilities for all countries in the sample (see, for example, Guo and Savickas, 2008 and Lustig, Roussanov, and Verdelhan (2011)).
- ΔLong rate$_t$: change in the interest rates of 10-year bonds, averaged across the countries in our sample (Ang and Chen, 2010, Table 1).

- Term_t: the term premium averaged across the countries in our sample, where the country term premium is the difference between the interest rate of a 10-year bond and one-month interest rate (Libor or its equivalent) (Ang and Chen, 2010, Table 1).
- ΔOIL_t: change in the price of oil (Brent Crude). We are guided by the observation that CAD and NOK feature among the carry trade currencies and that oil exports present a sizable share of exports for both Canada and Norway. It has been shown that countries often encounter exogenous terms-of-trade shocks due to oil price changes, possibly affecting their exchange rates (Backus and Crucini, 2000; Kilian, 2009; Ferraro, Rogoff, and Rossi, 2011).
- AFD_t: average forward discount. Forward discounts are extracted from spot and forward rates at the end of month t, with the US dollar as the reference currency, and then averaged across all countries in the sample, as in Lustig, Roussanov, and Verdelhan (2010) (see also Hansen and Hodrick, 1983; Bekaert and Hodrick, 1992).
- $\Delta\text{IP}_t^{\text{OECD}}$: change in industrial production, computed as change in total industrial production of the OECD countries, and meant to surrogate global economic growth.

The contemporaneous correlations among these variables, and with carry trade payoffs are shown next:

	ΔOIL_t	AFD_t	$\Delta\sigma_t^{\text{equity}}$	$\Delta\text{IP}_t^{\text{OECD}}$	$\Delta\text{Long rate}_t$	Term_t	$\overline{z}_t^{(1)}$	$\overline{z}_t^{(2)}$	$\overline{z}_t^{(3)}$	$\overline{z}_t^{(4)}$
ΔVIX_t	−0.05	−0.14	0.72	−0.03	−0.10	−0.13	−0.16	−0.22	−0.25	−0.23
ΔOIL_t		-0.07	-0.06	0.40	0.23	0.18	0.19	0.20	0.19	0.16
AFD_t			−0.07	−0.21	−0.12	−0.20	−0.02	−0.08	−0.02	0.01
$\Delta\sigma_t^{\text{equity}}$				0.03	−0.10	−0.16	−0.18	−0.22	−0.23	−0.19
$\Delta\text{IP}_t^{\text{OECD}}$					0.13	0.29	0.17	0.15	0.14	0.09
$\Delta\text{Long rate}_t$						0.20	0.04	0.13	0.12	0.12
Term_t							0.11	0.13	0.17	0.18

For example, ΔVIX_t and $\Delta\sigma_t^{\text{equity}}$ are negatively correlated with carry trade payoffs, ΔOIL_t and $\Delta\text{IP}_t^{\text{OECD}}$ are positively correlated, and AFD_t is practically unrelated. As could be expected, $\Delta\sigma_t^{\text{equity}}$ and ΔVIX_t are highly correlated (the correlation coefficient is 0.72). Our proxy for global economic growth, $\Delta\text{IP}_t^{\text{OECD}}$, is most correlated with ΔOIL_t.

While not reported, $\Delta\text{IP}_t^{\text{OECD}}$ displays varying correlations with our three predictors. Specifically, the correlation is 0.52 with ΔCRB_t, −0.08 with $\Delta\sigma_t^{\text{fx}}$, and −0.14 with ΔLIQ_t. Thus, our predictors can be viewed as capturing different aspects of the global economic environment.

Turning to the risk factors employed in Section 5.2, they are constructed as follows.

- $f_{t+1,1}$: (average monthly currency excess returns): In analogy to Lustig, Roussanov, and Verdelhan (2011, Section 2.1), we compute the currency return factor as the average across the excess returns of our eight currencies, held long against the US dollar;
- $f_{t+1,2}$ (mimicking portfolio for currency volatility): In analogy to Menkhoff, Sarno, Schmeling, and Schrimpf (2012, Eq. (8)), the currency volatility factor is computed as the fitted value from regressing innovations in currency volatility on the excess returns of a set of basis assets (Huberman, Kandel, and Stambaugh, 1987; Balduzzi and Robotti, 2008). The basis assets include the eight currencies used to construct the carry trades, together with the MSCI World equity index, the CRB Raw Industrials commodity index, and the long-term bonds (averaged across the countries in the sample). For robustness, we also use only the currencies as basis assets, but the results are similar and not reported in Table 12, to save space. In line with Menkhoff, Sarno, and Schmeling (2012, Eq. (4)), we build the currency volatility series by first averaging absolute daily log returns across currencies, and then averaging these daily values over each month.

The risk factors are denominated in US dollars, to conform with the US dollar perspective of the carry trade payoffs, and are expressed in excess return over the US risk-free return.

References

Adrian, T., Etula, E., Shin, H., 2010. Risk appetite and exchange rates. Unpublished working paper no. 361. Federal Reserve Bank of New York and Princeton University.

Amihud, Y., Hurvich, C., Wang, Y., 2009. Multiple-predictor regressions: hypothesis testing. Review of Financial Studies 22, 413–434.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 2001. The distribution of realized exchange rate volatility. Journal of the American Statistical Association 96, 42–55.

Ang, A., Bekaert, G., 2007. Stock return predictability: is it there? Review of Financial Studies 20, 651–707.

Ang, A., Chen, J., 2010. Yield curve predictors of foreign exchange returns. Unpublished working paper. Columbia University, New York, NY.

Asness, C., Moskowitz, T., Pedersen, L., 2009. Value and momentum everywhere. Unpublished working paper. University of Chicago and New York University.

Backus, D., Crucini, M., 2000. Oil prices and the terms of trade. Journal of International Economics 50, 185–213.

Backus, D., Foresi, S., Telmer, C., 2001. Affine term structure models and the forward premium anomaly. Journal of Finance 56, 279–304.

Bakshi, G., Carr, P., Wu, L., 2008. Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. Journal of Financial Economics 87, 132–156.

Bakshi, G., Panayotov, G., Skoulakis, G., 2011. Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios. Journal of Financial Economics 100, 475–495.
Balduzzi, P., Robotti, C., 2008. Mimicking portfolios, economic risk premia, and tests of multi-beta models. Journal of Business and Economic Statistics 26, 354–368.
Bali, T., Engle, R., 2010. The intertemporal capital asset pricing model with dynamic conditional correlations. Journal of Monetary Economics 57, 377–390.
Beber, A., Buraschi, A., Breedon, F., 2010. Differences in beliefs and currency risk premia. Journal of Financial Economics 98, 415–438.
Bekaert, G., Hodrick, R., 1992. Characterizing predictable components in excess returns on equity and foreign exchange markets. Journal of Finance 47, 467–509.
Bessembinder, H., 1994. Bid–ask spreads in the interbank foreign exchange markets. Journal of Financial Economics 35, 317–348.
Bhansali, V., 2008. Carry trade and currency volatility. Journal of Fixed Income 146, 329–341.
Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4493.
Breen, W., Glosten, L., Jagannathan, R., 1989. Economic significance of predictable variations in stock index returns. Journal of Finance 44, 1177–1189.
Brunnermeier, M., Nagel, S., Pedersen, L., 2009. Carry trades and currency crashes. NBER Macroeconomics Annual 23, 313–347.
Brunnermeier, M., Pedersen, L., 2009. Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.
Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S., 2011. Do peso problems explain the returns to the carry trade? Review of Financial Studies 24, 853–891.
Burnside, C., Han, B., Hirshleifer, D., Wang, T., 2011. Investor overconfidence and the forward premium puzzle. Review of Economic Studies 78, 523–558.
Campbell, J., 1991. A variance decomposition for stock returns. Economic Journal 101, 157–179.
Campbell, J., Hamao, Y., 1992. Predictable stock returns in the US and Japan: a study of long-term capital market integration. Journal of Finance 47, 43–69.
Campbell, J., Medeiros, K., Viceira, L., 2010. Global currency hedging. Journal of Finance 65, 87–121.
Campbell, J., Thompson, S., 2008. Predicting excess stock returns out of sample: can anything beat the historical average? Review of Financial Studies 21, 1509–1531.
Chen, Y., Rogoff, K., 2003. Commodity currencies. Journal of International Economics 60, 133–160.
Chen, Y., Rogoff, K., Rossi, B., 2010. Can exchange rates forecast commodity prices. Quarterly Journal of Economics 125, 1145–1194.
Chinn, M., Meese, R., 1995. Banking on currency forecasts: how predictable is change in money? Journal of International Economics 38, 161–178.
Christiansen, C., Ranaldo, A., Soderlind, P., 2011. The time-varying systematic risk of carry trade strategies. Journal of Financial and Quantitative Analysis 46, 1107–1125.
Christoffersen, P., Diebold, F., 2006. Financial asset returns, direction-of-change forecasting, and volatility dynamics. Management Science 52, 1273–1287.
Clarida, R., Davis, J., Pedersen, N., 2009. Currency carry trade regimes: beyond the Fama regression. Journal of International Money and Finance 28, 1375–1389.
Clark, T., West, K., 2007. Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics 138, 291–311.
Cochrane, J., 1996. A cross-sectional test of an investment-based asset pricing model. Journal of Political Economy 104, 572–621.
Cochrane, J., 2005. Asset Pricing. Princeton University Press, Princeton, NJ.
Cochrane, J., Piazzesi, M., 2005. Bond risk premia. American Economic Review 95, 138–160.
Cutler, D., Poterba, J., Summers, L., 1991. Speculative dynamics. Review of Economic Studies 58, 529–546.
Darvas, Z., 2009. Leveraged carry trade portfolios. Journal of Banking and Finance 33, 944–957.
Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. Review of Financial Studies 22, 3491–3530.
Diebold, F., Nason, J., 1990. Nonparametric exchange rate prediction? Journal of International Economics 28, 315–332.
Engel, C., Mark, N., West, K., 2008. Exchange rate models are not as bad as you think. NBER Macroeconomics Annual 2007, vol. 22, pp. 381–441.
Engel, C., Mark, N., West, K., 2009. Factor model forecasts of exchange rates. Unpublished working paper. University of Wisconsin and University of Notre Dame.
Evans, M., Lyons, R., 2002. Order flow and exchange rate dynamics. Journal of Political Economy 110, 170–180.
Farhi, E., Fraiberger, S., Gabaix, X., Ranciere, R., 2011. Crash risk in currency markets. Working paper. Harvard University.
Ferraro, D., Rogoff, K., Rossi, B., 2011. Can oil prices forecast exchange rates? Unpublished working paper. Duke University and Harvard University.
Goncalves, S., White, H., 2002. The bootstrap of the mean for dependent heterogeneous arrays. Econometric Theory 18, 1367–1384.
Groen, J., Pesenti, P., 2009. Commodity prices, commodity currencies, and global economic developments. Unpublished working paper no. 361. Federal Reserve Bank of New York.
Guo, H., Savickas, R., 2008. Average idiosyncratic volatility in G7 countries. Review of Financial Studies 21, 1259–1296.
Hansen, L., Hodrick, R., 1983. Risk aversion speculation in the forward foreign exchange market. In: Frenkel, J. (Ed.), Exchange Rates and International Macroeconomics, University of Chicago, Chicago.
Hau, H., Rey, H., 2006. Exchange rates, equity prices, and capital flows. Review of Financial Studies 19, 273–317.
Henriksson, R., Merton, R., 1981. On market timing and investment performance. II. Statistical procedures for evaluating forecasting skills. Journal of Business 54, 513–533.
Heyerdahl-Larsen, C., 2011. Asset prices and real exchange rates with deep habits. Unpublished working paper. London Business School.
Hodrick, R., 1992. Dividend yields and expected stock returns: alternative procedures for inference and measurement. Review of Financial Studies 5, 357–386.
Huberman, G., Kandel, S., Stambaugh, R., 1987. Mimicking portfolios and exact arbitrage pricing. Journal of Finance 42, 1–9.
Ilmanen, A., 1995. Time-varying expected returns in international bond markets. Journal of Finance 50, 481–506.
Jordá, O., Taylor, A., 2009. The carry trade and fundamentals: nothing to fear but FEER itself. Unpublished working paper. NBER # 15518.
Jurek, J., 2009. Crash-neutral currency carry trades. Unpublished working paper (SSRN # 1262934). Princeton University.
Jylha, P., Suominen, M., 2011. Speculative capital and currency carry trades. Journal of Financial Economics 99, 60–75.
Kearns, J., 2007. Commodity currencies: why are exchange rate futures biased if commodity futures are not? Economic Record 83, 60–73.
Kilian, L., 1999. Exchange rates and monetary fundamentals: what do we learn from long-horizon regressions? Journal of Applied Econometrics 14, 491–510.
Kilian, L., 2009. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. American Economic Review 99, 1053–1069.
Kilian, L., Taylor, M., 2003. Why is it so difficult to beat the random walk forecast of exchange rates? Journal of International Economics 60, 85–107.
Kirby, C., 1998. The restrictions on predictability implied by rational asset pricing models. Review of Financial Studies 11, 343–382.
Lettau, M., Ludvigson, S., 2001. Resurrecting the (C)CAPM: a cross-sectional test when risk premia are time-varying. Journal of Political Economy 109, 1238–1287.
Levich, R., Poti, V., 2008. Predictability and “good deals” in currency markets. Unpublished working paper. NBER # 14597.
Lustig, H., Roussanov, N., Verdelhan, A., 2010. Countercyclical currency risk premia. Unpublished working paper. NBER # 16427.
Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. Review of Financial Studies 24, 3731–3777.
Lustig, H., Verdelhan, A., 2007. The cross-section of foreign currency risk premia and consumption growth risk. American Economic Review 97, 89–117.
Mark, N., 1995. Exchange rates and fundamentals: evidence on long-horizon predictability. American Economic Review 85, 201–218.
Mark, N., 2001. International Macroeconomics and Finance: Theory and Econometric Methods. Blackwell, Oxford.
Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? Journal of International Economics 14, 3–24.
Melvin, M., Taylor, M., 2009. The crisis in the foreign exchange market. Journal of International Money and Finance 28, 1317–1330.
Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012. Carry trades and global foreign exchange volatility. Journal of Finance 67, 681–718.
Newey, W., West, K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Newey, W., West, K., 1994. Automatic lag selection in covariance matrix estimation. Review of Economic Studies 61, 631–653.

Pesaran, H., Timmermann, A., 1994. A generalization of the nonparametric Henriksson-Merton test of market timing. Economics Letters 44, 1–7.

Politis, D., Romano, J., 1994. The stationary bootstrap. Journal of the American Statistical Association 89, 1303–1313.

Politis, D., White, H., 2004. Automatic block-length selection for the dependent bootstrap. Econometric Reviews 23, 53–70.

Ranaldo, A., Soderlind, P., 2010. Safe haven currencies. Review of Finance 10, 395–407.

Rapach, D., Strauss, J., Zhou, G., 2010. Out-of-sample equity premium prediction: combination forecasts and links to the real economy. Review of Financial Studies 23, 821–862.

Sarno, L., Schneider, P., Wagner, C., 2012. Properties of foreign exchange risk premiums. Journal of Financial Economics 105, 279–310.

Stock, J., Watson, M., 2004. Combination forecasts of output growth in a seven-country data set. Journal of Forecasting 23, 405–430.

Wei, M., Wright, J., 2009. Confidence intervals for long-horizon predictive regressions via reverse regressions. Unpublished working paper. Federal Reserve Board and John Hopkins University.

Yu, J., 2011. A sentiment-based explanation of the forward premium puzzle. Unpublished working paper. University of Minnesota.