CURRENCY CARRY TRADING WITH MGARCH-BASED CARRY-TO-RISK PORTFOLIO OPTIMIZATION*

Harald Schmidbauer†/ Angi Rösch‡/ Tolga Sezer§/ Vehbi Sinan Tunalıoğlu¶

(Last compiled: July 21, 2010)

Abstract

Currency carry trading is an investment strategy which borrows in a low-yield currency and/or invests in a high-yield currency, with the hope of profiting from the interest rate differential (the carry) between two money markets, as well as from currency movements. Technically, profit (or loss) will be generated by a difference between the exchange rate written in a forward contract, and the spot exchange rate on the day when the contract expires. The risk in this strategy stems from unforeseeable currency movements. Which currencies, and with which weights, should be included in a portfolio? One possible approach is to use carry-to-risk maximization. We extend the usual static approach by substituting an MGARCH-based conditional covariance matrix capable of forecasting portfolio risk for the next period. By comparing examples of time series of realized profits, our paper shows the benefit of forecasting the conditional covariance structure of currency movements when constructing a carry-trading portfolio.

Key words: Currency carry trading; portfolio optimization; carry-to-risk optimization; multivariate GARCH

1 Introduction

The basic idea of currency carry trade strategies is to borrow in lower yielding currencies and to invest in high yielding ones. These strategies are an outcome of the observation that the Uncovered Interest Rate Parity (UIP) may fail to hold. UIP is an arbitrage condition which states that there should be no profit opportunity from the differences in interest rates between two currencies. The present paper tests whether the economic benefit for carry trade investors can be enhanced by using advanced volatility forecasting tools.

Deviations from the UIP, also leading to what is termed "Forward Premium Puzzle", have been widely discussed in literature during the past 20 years. Surveys such as in Hodrick [13], Froot and Thaler [11] and Engel [9] are widely cited. In this branch of literature, the failure of the UIP to hold is attributed to the time-varying nature of risk premia, expectational errors

*This research project was presented on the 30th International Symposium on Forecasting ISF2010, held in San Diego, June 20-23. The paper is published in: The International Institute of Forecasters (ed.), *Proceedings of the 30th International Symposium on Forecasting ISF2010*, San Diego, June 20-23. ISSN: 1997-4124.

†Bilgi University, Istanbul, Turkey, & Ideal Analytix, Singapore; e-mail: harald@hs-stat.com

‡FOM University of Applied Sciences, Munich, Germany / Tai'an, China, & Ideal Analytix, Singapore; e-mail: angi@angi-stat.com

§Ideal Analytix, Singapore; e-mail: tolga.sezer@ideal.sg

¶Ideal Analytix, Singapore; e-mail: sinan.tunalioglu@ideal.sg

about the future spot exchange rate, and market microstructure issues such as liquidity and bid-ask spreads. However, a relatively small amount of literature has actually attempted to measure the economic benefit that can be achieved by an investor who tries to take advantage of this puzzle. This is astonishing from a practical point of view because investment practioners have been applying strategies, notably carry trading, to exploit this anomaly ever since its discovery by Bilson [3]. Evidence for the high volume of carry trading activity can be found in Galati et al. [12].

Furthermore, Backus et al. [1], Villanueva [16] and Wagner [17] have contributed to a new line of literature in this field. Villanueva [16] points out that “few studies examine forward bias trading profits even though significant returns define the economic significance of an anomaly”. In this line, Burnside et al. [4, 5, 6] have undertaken investigations aimed at tackling the literature shortcomings mentioned by Villanueva [16]. More specifically, they show how these carry trading strategies yield high Sharpe ratios on average. Using monthly data starting from 1976, they show that significant profits can be made using currency carry trading for individual currencies as well as for a portfolio formed in a classic Markowitz framework. However, they point out that these returns are not a compensation for risk. They argue that in the presence of microstructure frictions and limitations to speculation in foreign exchange, the marginal Sharpe ratios are zero even though average Sharpe ratios are positive.

Della Corte et al. [7, 8] point out that in the context of dynamic asset allocation strategies, there is no study investigating the economic value of the predictive ability of empirical exchange rate models, which condition on the forward premium while allowing for “volatility timing”. Further, they implicitly assume that the underlying returns are normally distributed. They allocate wealth according to different models such as a stochastic volatility model, a simple random walk, a monetary exchange rate model, and a GARCH-based model. To measure the economic benefit, the authors compare the level of average utility generated by each estimate of conditional mean and variance. Their findings lead them to conclude that there is significant economic benefit to an investor who exploits deviations from UIP by forecasting currency returns.

In Mayer [14], the returns to forward bias-trading in dynamic multi-currency strategies are investigated to assess the limits to speculation in foreign exchange markets, thereby challenging the findings by Burnside [4]. The results show that bias-trading strategies allow for economically significant excess returns and represent attractive diversification devices. Furthermore the author finds that carry trading portfolios containing emerging market currencies result in large diversification gains.

Galati et al. [12] point out that carry trading strategies are very widely used in the investment community and that a popular measure to gauge the attractiveness of the carry trade ex ante is to use the carry-to-risk ratio. The latter adjusts the interest rate differential w.r.t. the risk of future exchange rate movements, where this risk is proxied by the expected volatility (implied by foreign exchange options) of the relevant currency pair.

In the light of the existing literature, our goal is to extend extant approaches in this field while building on established properties of carry trading, and try to overcome certain shortcomings of currency carry trading, as practiced up to now:

1. Carry trade returns tend to fall in times of high volatility. Forecasting future exchange rate volatility by means of conditional volatility should enable us to better time when to invest and to unwind the carry trade (volatility timing effect).

2. The scope of the trade should be extended to form an optimal portfolio of currencies which maximizes the investor’s utility function by estimating the joint conditional risk of the portfolio.

3. A framework for the investment strategy should be formulated which allows the measurement of the ex ante attractiveness of the carry trading strategy in terms of the widely applied carry-to-risk ratio.

4. If carry trade returns tend to fall in times of high volatility, one should be able to short the high yielding currency, which is the currency prone to post negative returns.

This paper is organized as follows. Section 2 sets the stage for our investigation by introducing the relevant notions for the case of currency carry trading involving a single foreign currency. Portfolio construction and optimization, when investing in two foreign currencies, is the topic of Section 3. We obtain the dynamic risk forecast for the portfolio based on a bivariate GARCH model; details are explained in Section 4. Empirical results when investing in the Brazilian real and the Mexican peso, with the USD as base currency, are given and discussed in Section 5. Finally, Section 6 summarizes and draws some conclusions.

2 Currency Carry Trading: The Case of a Single Currency

In the following, we shall show the connection between different notions of *carry*, *profit*, volatility forecasts, and possible investment criteria. We also show the behaviour (i.e. realized profits and losses) of a monthly sequence of investments, based on the USD (US dollar) and the BRL (Brazilian real).

Our strategy

We take on the perspective of a USD investor, who contracts in a 1-month USD-BRL forward on Tuesday every four weeks. Our numerical example assumes that investments start in January 2005. The maturity date of each forward contract we enter is assumed to be the Tuesday four weeks later. On the date of maturity we exercise our contract, i.e. buy BRL at the contracted forward price, and immediately sell them at the observed spot price on this day. With this strategy, our realized (=observed) gross profit/loss at each date of maturity will be:

$$\frac{\text{contracted forward price}}{\text{observed spot price on the day of maturity}} \tag{1}$$

The sequence of realized gross profits/losses of this strategy, when performed from January 2005 onwards, is the red line in Figure 1.

Expectations at the beginning of the contract

According to the interest rate parity theorem (that is, according to UIP), the contracted forward price is supposed to be chosen so as to offset the interest rate differential between the USD and the BRL money market, which means that, on the date of maturity, spot prices are expected to equal the contracted forward prices. In this case, the profit/loss of our strategy will be zero, in other words, the *gross profit/loss at expected spot prices* will be one. The black line in Figure 1 refers to our expected gross profits/losses according to UIP.

Empirical studies, however, contradict this theory. They show that expected profits of carry trade investments are positive on the average. For example, in the case of constant spot prices (that is, the observed spot price on the day of maturity equals the spot price at the beginning of the contract), our gross profits/losses on the day of maturity will be

$$\frac{\text{contracted forward price}}{\text{spot price at the beginning of the contract}}, \tag{2}$$

which can be interpreted as the *gross profit/loss at constant spot prices*. This ratio is called the *carry of the contract*, in the sense of a contract parameter. Since its value is available at the

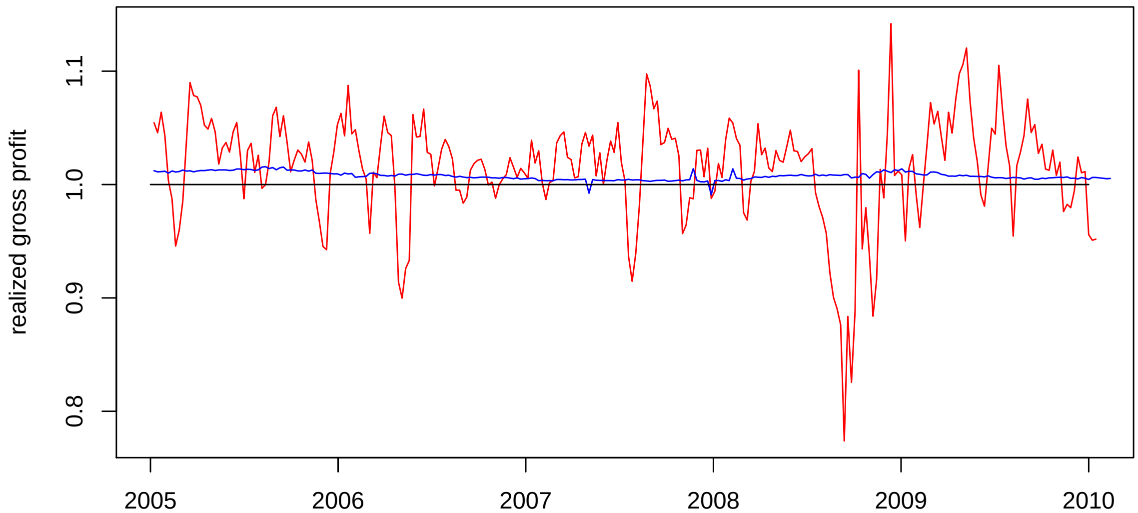

Figure 1: USD-BRL gross profit/loss sequences: realized (red line) and hypothetical (black and blue line)

date of decision whether to enter the contract or not, it may serve as an investment criterion, leading to a simple investment strategy: enter the contract if its carry is not smaller than one, otherwise don't. The sequence of the carries of the contract is the blue line in Figure 1. — This strategy clearly neglects the volatility of spot prices.

Our observation on the day of maturity

Our realized gross profits/losses (the red line in Figure 1) show much more volatility than in either of the two hypothetical scenarios at the beginning of the contract (black and blue line). Although our strategy provides an average gross profit significantly greater than one (1.014928, p-value: $2.343e-07$), it does not prevent us from entering the contract during the second half of 2008, where we encounter huge losses. Therefore, a simple decision strategy on the basis of the amount of the carry of the contract alone is insufficient.

Volatility forecast at the beginning of the contract

On each date of decision, we apply a GARCH(1,1) model to estimate the conditional volatility of the next week's return on spot prices.[1] From that we calculate a four-week-forecast of the return's volatility, matching the date of contract maturity. The sequence of volatility forecasts (assigned to the date of maturity, again in a weekly time setting) is shown in Figure 2.

Carry-to-risk forecast at the beginning of the contract

We combine carry of the contract and volatility forecasts to obtain a more sophisticated investment criterion, the so-called carry-to-risk forecast:

$$\frac{\text{carry of the contract}}{1 + \text{volatility forecast}/100} \tag{3}$$

The sequence of carry-to-risk indicators (assigned to the date of maturity, again in a weekly time setting) is shown in Figure 3. A possible investment criterion related to it is: Enter the contract if the carry-to-risk exceeds a given threshold, otherwise don't.

[1] All computations for this paper were carried out in R [15].

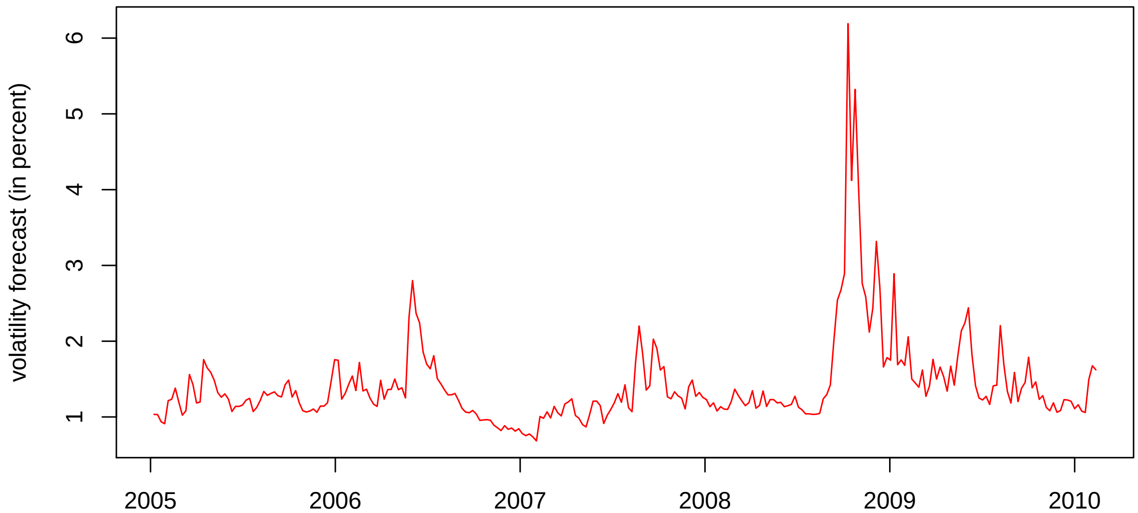

Figure 2: USD-BRL four-week volatility forecast

3 Portfolio Construction

This section deals with the problem of forming a currency carry trade portfolio when investing in two currencies. The performance of the portfolio is measured as the profit/loss generated after four weeks.

Let t denote our decision point of time. Then, the carry of the portfolio is the gross profit/loss of the portfolio we would expect at constant spot prices. It is composed of the single currency carries, here denoted in percent by

$$c_{i,t} = \left(\frac{f_{i,t}}{p_{i,t}} - 1 \right) \cdot 100\%, \tag{4}$$

where $i = 1, 2$ indicates the currency.

A portfolio is characterized by the weights ϕ_i invested in each currency. The carry-to-risk for such a portfolio can be computed as

$$\frac{\text{carry of the portfolio}}{1 + \text{portfolio volatility forecast}/100} = \frac{1 + \phi' \cdot \mathtt{c}_t/100}{1 + \sqrt{\phi' \cdot \hat{\mathtt{H}}_{t+4} \cdot \phi}/100}, \tag{5}$$

where $\phi' = (\phi_1, \phi_2)$, $\mathtt{c}_t = (c_{1,t}, c_{2,t})'$, and $\hat{H}_{t+4}$ designates the covariance matrix forecast of the return vector at time $t + 4$, i.e. four weeks after the initial investment has been made. The optimal portfolio is then determined by maximizing the carry-to-risk w.r.t. weights ϕ_1 and ϕ_2 under the constraints

$$|\phi_1| + |\phi_2| = 1, \quad -1 \leq \phi_1, \phi_2 \leq 1, \tag{6}$$

where a negative value of ϕ_1 or ϕ_2 indicates short selling. There are several approaches to obtain the covariance matrix forecast $\hat{H}_{t+4}$. Two obvious ways are:

- "static approach": Use rolling estimation (e.g. the 50 weeks preceding t) to obtain an "average" historical covariance pattern. The historical covariance matrix is taken as the forecast.
- "dynamic approach": Use a conditional covariance matrix, obtained from a multivariate GARCH model.

We shall compare these two approaches empirically in Section 5 below. The model to obtain the conditional covariance matrix forecast is described in the next section.

Figure 3: USD-BRL four-week carry-to-risk forecast

4 The Conditional Covariance Matrix Forecast

Let $r_t = (r_{1,t}, r_{2,t})$ be the vector of returns (exchange rate changes in percent) in week t. We assume that

$$r_t = \mu + \epsilon_t, \tag{7}$$

where μ is a constant vector and (ϵ_t) is modeled as a bivariate GARCH-BEKK process.[2] This mean specification is appropriate if no significant autocorrelation is found in the return series. The process (ϵ_t) is specified as

$$\epsilon_t = \mathtt{H}_t^{1/2} \cdot \nu_t, \tag{8}$$

where (ν_t) is a bivariate white noise process with $\text{cov}(\nu_t) = \mathtt{I}$ (the 2×2 unity matrix) and

$$\mathtt{H}_t \;=\; \mathtt{C}'\mathtt{C} \;+\; \mathtt{A}'\epsilon_{t-1}\epsilon_{t-1}'\mathtt{A} \;+\; \mathtt{B}'\mathtt{H}_{t-1}\mathtt{B} \tag{9}$$

with parameter matrices $\mathtt{C} = (c_{ij})$ $(c_{21} = 0)$, $\mathtt{A} = (a_{ij})$, $\mathtt{B} = (b_{ij})$. This is a symmetric model insofar as any ϵ_{t-1} will lead to the same $\mathtt{H}_t$ as $-\epsilon_{t-1}$.

The one-period ahead forecast of the covariance matrix, $\hat{\mathtt{H}}_{t+1}$, can be computed from equation (9) at time t. According to (5), a four-period ahead forecast for the covariance matrix is needed. This can be obtained recursively by substituting the respective conditional covariance matrix. The conditional expectations of the $\epsilon \cdot \epsilon'$ terms in (9) are estimated from averages over the four weeks preceding week t:

$$\hat{\mathtt{H}}_{t+i} \;=\; \mathtt{C}'\mathtt{C} \;+\; \mathtt{A}'\bar{\epsilon}_t\bar{\epsilon}_t'\mathtt{A} \;+\; \mathtt{B}'\hat{\mathtt{H}}_{t+i-1}\mathtt{B}, \quad i = 2, 3, 4. \tag{10}$$

5 Empirical Results — The Case of USD-BRL and USD-MXP

We assume a USD-based investor investing in a portfolio of 1-month-forward contracts in BRL (Brazilian real) and MXP (Mexican peso). Investments start on a Tuesday in January 2005.

[2]See Engle and Kroner [10], Bauwens and Rombouts [2].

Four weeks later, the contracts are exercised, and gross profits/losses can be observed in terms of the portfolio of realized gross profits/losses at observed spot prices on the day of maturity,

$$\frac{f_{i,t}}{p_{i,t+4}}, \tag{11}$$

where $i = 1, 2$ indicates the currency. On the same day, the portfolio is restructured with respect to the weights invested in each currency.

We compare two strategies of portfolio construction. Both strategies are based on maximizing carry-to-risk (5), but they use different input with respect to forecasting risk (via the covariance matrix) four periods ahead: In the static approach, the covariance matrix is estimated from historical data, starting 50 weeks previous to the date of the contract decision. Alternatively, we apply an MGARCH-BEKK model (9) to data from 2000 onwards to obtain a conditional covariance matrix.

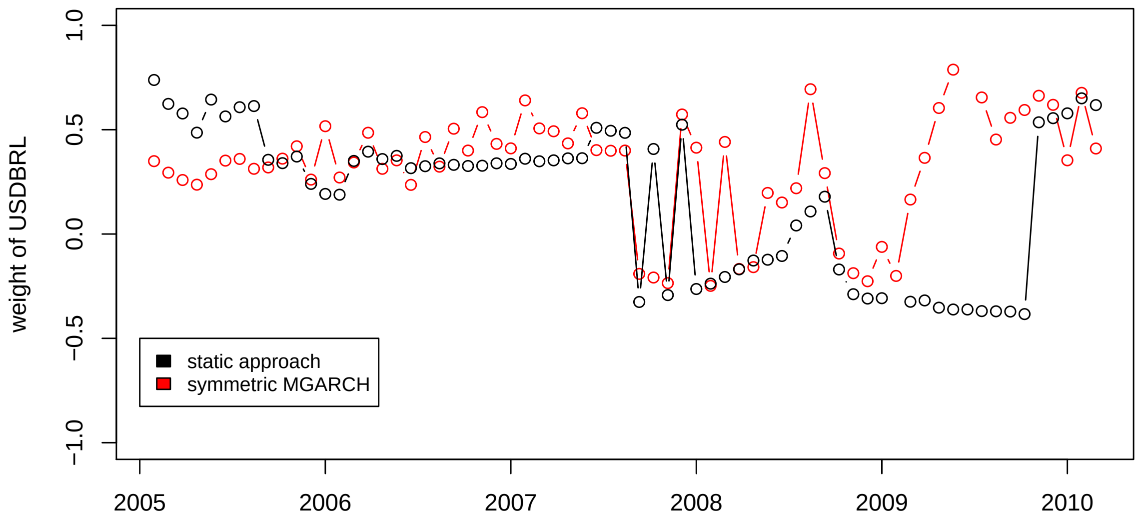

Figure 4: USD-BRL — sequences of optimal weights: static approach vs. MGARCH-BEKK

Figure 4 shows the way in which the portfolio is restructured from month to month, beginning in week 3 of January 2005. The plot shows only the weights of BRL, not those of MXP. Weights can be negative (short selling). A gap in the sequence of points indicates a month with no investment, since the portfolio carry was forecast negative for these months. The red dots refer to the MGARCH-BEKK approach. An investment criterion based on MGARCH obviously involves much more restructuring than a static approach based on an average covariance matrix (the black dots). This is no surprise, since conditional volatility and correlation change dynamically. In other words, the MGARCH-based portfolio will be more responsive to news.

The portfolio risk an investor encounters from month to month with either approach is shown in Figure 5 in a weekly time setting. The lines constitute the denominator of the carry-to-risk in equation (5). A comparison of this plot with Figure 2 reveals the decrease in risk when this portfolio is considered, rather than an investment in BRL alone. Again, the MGARCH reacts more promptly to new information, thus providing a more realistic evaluation of the risk involved than the static model.

The sequence of gross portfolio profits/losses, which could have been realized from January 2005 onwards (performed in a weekly time setting), as compared to one-month US treasury bonds, is shown in Figure 6. Each return is computed w.r.t. the investment made one month earlier, for each Tuesday. The greatest difference turns out to appear in 2009, in times of high volatility.

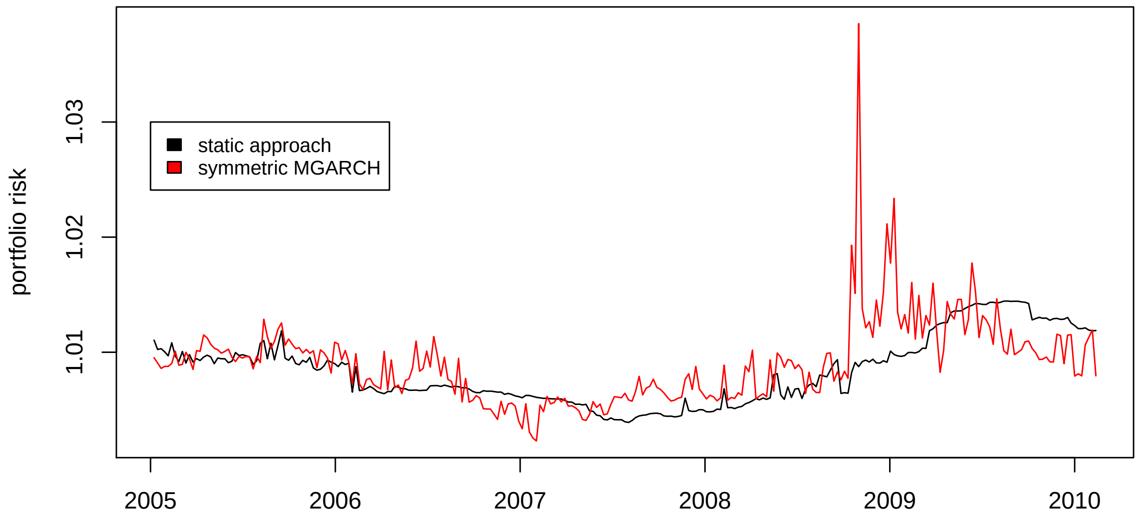

Figure 5: USD-BRL and USD-MXP — sequences of portfolio risks: static approach vs. MGARCH-BEKK

Return characteristics, as compared to a benchmark, can be described by Sharpe ratios. The Sharpe ratio can be defined as:

$$\frac{\mathrm{E}(R - R_f)}{\sqrt{\mathrm{var}(R - R_f)}},$$

where

$$\begin{aligned} R &= (\text{realized gross profit/loss} - 1) \cdot 100\%, \\ R_f &= \text{benchmark return} \end{aligned}$$

We use the Sharpe ratio empirically (ex post), not as a decision criterion. Values are shown in Table 1.

	no restriction w.r.t. trading	no trading in case of negative portfolio carry
symmetric MGARCH	0.14531	0.14652
static approach	0.00505	0.00074

Table 1: Sharpe ratios

The total profit accrued from month to month with either approach is shown in Figure 7. Here, our investment is beginning in week 3 of January 2005. A month with no investment results in a horizontal line for total accrued profit.

Figures 4, 5, 6 and 7 suggest that the effect of using a dynamic volatility model in portfolio optimization may be volatility-specific. To obtain more insight, weekly portfolio volatility, as predicted by the MGARCH, can be classified by magnitude: *very low*, *lower*, *higher*, and *very high*, each containing 25% of the observations. (The border between *higher* and *lower* defines the median.) Figure 8 shows the average surplus profit for the corresponding weeks when using the dynamic model vs. the static model, together with 95% confidence intervals for the expected surplus. It turns out that expected surplus is significantly positive when volatility is either very low or very high, and expected surplus is still positive in case of moderate volatility, even though the increase is not significant.

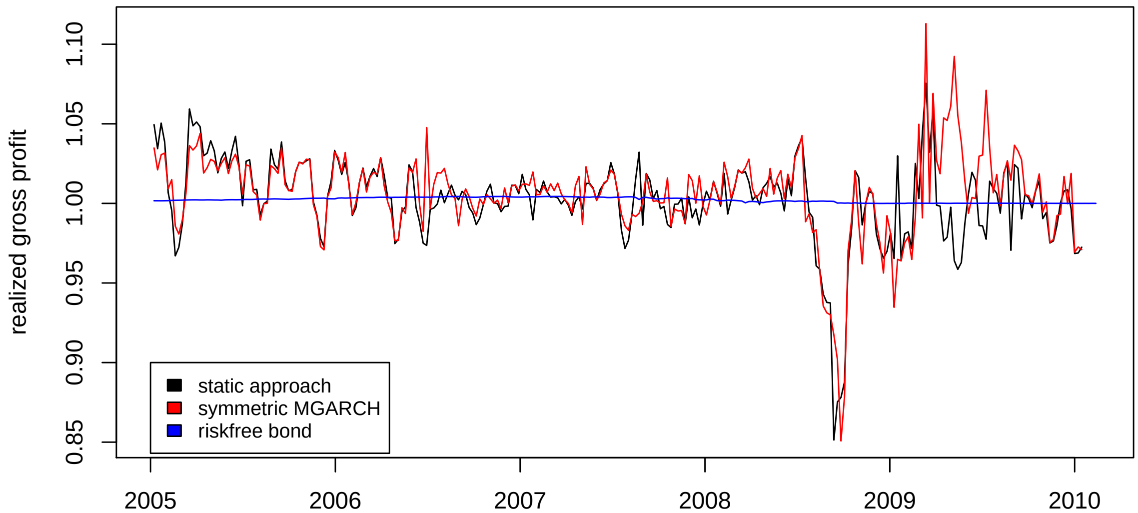

Figure 6: USD-BRL and USD-MXP — sequences of total gross profits/losses: static approach vs. MGARCH-BEKK

The better performance of the dynamic model in times of very high volatility is in accordance with item 1 (end of Section 1). A glance at Figure 4 provides empirical evidence for the conjecture in item 4 concerning a short position in times of high volatility.

6 Summary and Conclusions

The Uncovered Interest Rate Parity (UIP) condition states that there should be no arbitrage opportunity from the interest rate differential (the carry) between two currencies. From the observation that the UIP may fail to hold, a small branch of literature on currency carry trading strategies has emerged, aiming to measure the economic benefit to an investor who takes advantage of the failure of UIP.

The goal of our study is to try to overcome certain shortcomings of currency carry trading strategies as practised up to now. These shortcomings are related to the limited responsiveness to volatility movements of the usual static approach of volatility forecasting.

An MGARCH-BEKK model is used to obtain a conditional covariance matrix capable of forecasting portfolio risk for the next period. We compare this dynamic to the usual static approach by measuring realized profits/losses from the perspective of a USD-based investor who contracts in a portfolio of two currencies: the Brazilian real and the Mexican peso. The portfolio is restructured each month by means of carry-to-risk maximization. We found that this dynamic approach leads to a significant surplus profit in times of either very low or very high volatility. Conditionality, on which the dynamic model rests, is superior in capturing disruptions of continuity and triggering trading signals.

Further research needs to be undertaken in at least two directions. Firstly, from a practical point of view, it is desirable not to limit the portfolio to only two currencies. Secondly, from a more methodological point of view, an asymmetric MGARCH model may be superior in exploiting the information structure of past events and thus lead to more differentiated trading signals.

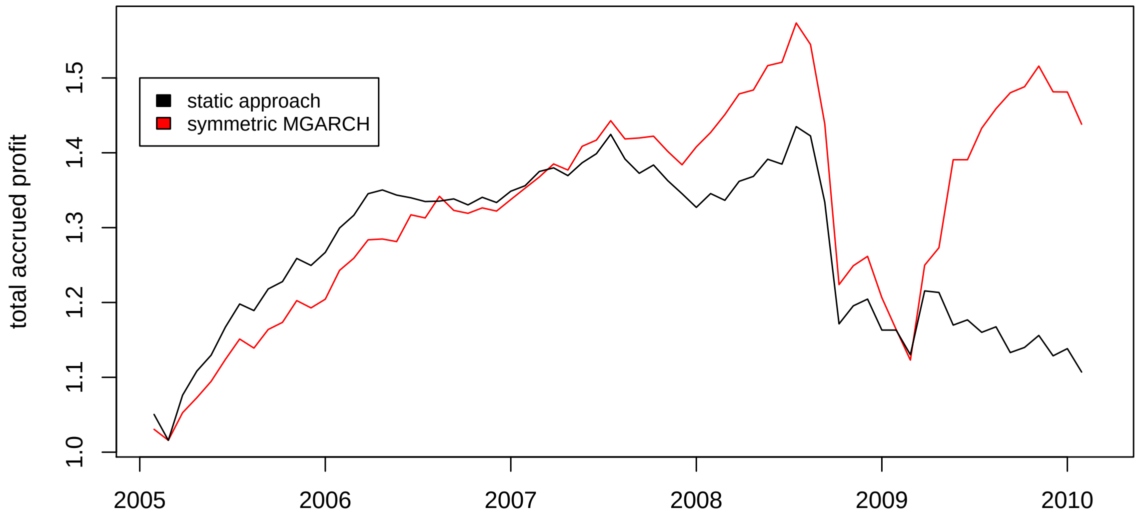

Figure 7: USD-BRL and USD-MXP — sequences of total accrued profit: static approach vs. MGARCH-BEKK

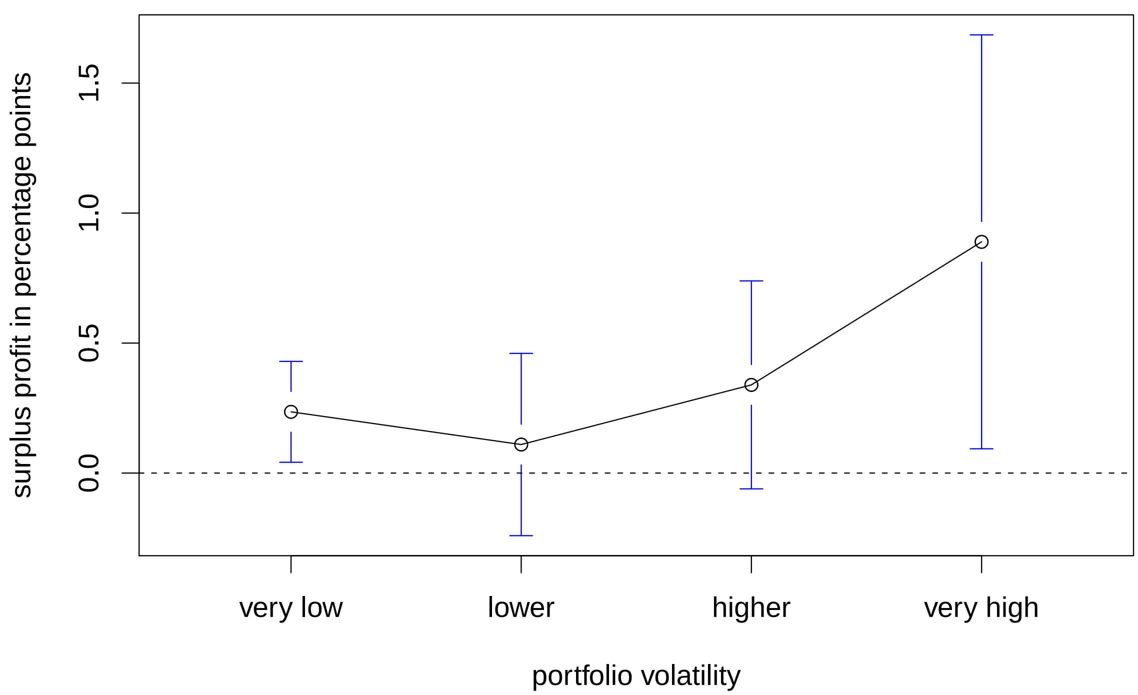

Figure 8: USD-BRL and USD-MXP — 95% confidence intervals for the surplus profit with MGARCH-BEKK volatility forecasts

References

[1] Backus, D.K., Gregory, A.W., & Telmer C.I. (1993): Accounting for forward rates in markets for foreign currency. *Journal of Finance* 48, 1887–1908.

[2] Bauwens, L., Laurent, S., & Rombouts, J.V.K. (2003): *Multivariate GARCH models: a survey*. CORE Discussion Paper 2003/31, Université Catholique de Louvain. Online available http://www.core.ucl.ac.be/ econometrics/Bauwens/Papers.

[3] Bilson, J.F.O. (1981): The "speculative efficiency" hypothesis. *Journal of Business* 54, 435–451.

[4] Burnside, C. (2008): Comment on M.K. Brunnermeier, S. Nagel, & L.H. Pedersen: *Carry Trades and Currency Crashes*. *NBER Macroeconomics Annual* 23, 349–360.

[5] Burnside, C., Eichenbaum, M., & Rebelo, S. (2006): Government finance in the wake of currency crises. *Journal of Monetary Economics* 53, 401–440.

[6] Burnside, C., Eichenbaum, M., & Rebelo, S. (2008): Carry trade: The gains of diversification. *Journal of the European Economic Association* 6, 581–588.

[7] Della Corte, P., Sarno, L., & Tsiakas, I. (2009): An economic evaluation of empirical exchange rate models. *Review of Financial Studies* 22, 3491–3530.

[8] Della Corte, P., Sarno, L., & Tsiakas, I. (2010): Spot and Forward Volatility in Foreign Exchange. CEPR Discussion Papers, no. 7893.

[9] Engel, C. (1996): The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of Empirical Finance* 3, 123–192.

[10] Engle, R.F., & Kroner, K. (1995): Multivariate simultaneous generalized ARCH. *Economic Theory* 11, 122–150.

[11] Froot, K.A., & Thaler, R.H. (1990): Anomalies: foreign exchange. *Journal of Economic Perspectives* 4, 179–192.

[12] Galati, G., Heath, A., & McGuire, P. (2007): Evidence of carry trade activity. *BIS Quarterly Review*, September 2007, 27–41.

[13] Hodrick, R.J. (1987): *The empirical evidence on the efficiency of forward and futures foreign exchange markets.* Harwood Academic Publishers, Chur, Switzerland.

[14] Mayer, G. (2010): Forward Bias Trading in Emerging Market. UniCredit & Universities Working Paper Series, no. 10.

[15] R Development Core team (2010): *R: A language and environment for statistical computing.* R foundation for statistical computing, Vienna, Austria. URL http://www.R-project.org.

[16] Villanueva, O.M. (2007): Forecasting currency excess returns: Can the forward bias be exploited? *Journal of Financial and Quantitative Analysis* 42, 963–990.

[17] Wagner, C. (2008): Risk-Premia, Carry-Trade Dynamics, and Speculative Efficiency of Currency Markets. Oesterreichische Nationalbank (Austrian Central Bank) Working Papers, no. 143.