

The Returns to Currency Speculation in Emerging Markets

Craig Burnside*, Martin Eichenbaum[†] and Sergio Rebelo[‡]

February 2007

Abstract

The carry trade strategy involves selling forward currencies that are at a forward premium and buying forward currencies that are at a forward discount. We compare the payoffs to the carry trade applied to two different portfolios. The first portfolio consists exclusively of developed country currencies. The second portfolio includes the currencies of both developed countries and emerging markets. Our main empirical findings are as follows. First, including emerging market currencies in our portfolio substantially increases the Sharpe ratio associated with the carry trade. Second, bid-ask spreads are two to four times larger in emerging markets than in developed countries. Third and most dramatically, the payoffs to the carry trade for both portfolios are uncorrelated with returns to the U.S. stock market.

J.E.L. Classification: F31

Keywords: Uncovered interest parity, exchange rates, carry trade.

*Duke University and NBER.

[†]Northwestern University, NBER, and Federal Reserve Bank of Chicago.

[‡]Northwestern University, NBER, and CEPR.

Currencies that are at forward premium tend to depreciate. This empirical regularity underlies the carry trade, a currency speculation strategy that is widely used by practitioners. This strategy involves selling forward currencies that are at a forward premium and buying forward currencies that are at a forward discount. In this paper we compare the payoffs to the carry trade applied to two different portfolios. The first portfolio consists exclusively of developed country currencies. The second portfolio includes the currencies of both developed countries and emerging markets. Our main empirical findings are as follows. First, including emerging market currencies in our portfolio substantially increases the Sharpe ratio associated with the carry trade. Second, bid-ask spreads are two to four times larger in emerging markets than in developed countries. Carry trade strategies that ignore bid-ask spreads yield negative Sharpe ratios. Large positive Sharpe ratios emerge only when the trading strategy takes bid-ask spreads into account. Third, over our sample period the payoffs to the carry trade are essentially uncorrelated with returns to U.S. stock market. This result is consistent with Craig Burnside, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo (2006) who argue that it is difficult to rationalize the payoffs to the carry trade as compensating agents for bearing risk. Burnside et al. (2006) propose a market microstructure explanation of the large Sharpe ratios associated with the carry trade. We suspect that these explanations apply with even greater force to emerging market currencies.

Our data, which is from Datastream, covers the period from October 1997 to November 2006. Not all countries are included throughout the sample. Rather, countries are included as data on them becomes available. The original data are daily and represent quotes at 4 p.m. taken from the Reuters system. The rates are based on actual traded rates on the Reuters Dealing 2000-2 network along with other quoted rates contributed to Reuters by leading market participants (see Reuters, 2003). Each exchange rate is quoted as foreign currency units per U.S. dollar. We convert the daily data into weekly data by sampling the daily data on every Wednesday. We also construct a monthly data set by sampling the daily data on the 2nd day of each month. Our data includes both bid and ask exchange rates. The ask (bid) exchange rate is the rate at which a participant in the interdealer market can buy (sell) U.S. dollars from a currency dealer.

We have reliable data on both spot and forward rates quoted against the U.S. dollar for 63 countries. These are the countries included in the ‘large portfolio’ that we discuss below. We also consider a ‘small portfolio’ that includes ten developed country currencies

as in Burnside et al. (2006).¹

We obtained data on the weekly stock market premium ($Mkt-RF$) and the weekly Treasury Bill rate from Kenneth French's data library.²

The bid-ask spreads for both the spot and forward market are much larger for emerging markets than for developed economies. For developed countries, the median bid-ask spread in the spot market was between 0.039 and 0.051 percent depending on the sample period chosen, while the median spreads for 1-week and 1-month forwards were in the ranges 0.042–0.053 percent and 0.045–0.057 percent, respectively. We found the spreads to be between two and four times larger for emerging market currencies than for developed country currencies.

Let S_t^a and S_t^b denote the ask and bid spot exchange rates, respectively. Let F_t^a and F_t^b denote the ask and bid forward exchange rate, respectively, for forward contracts maturing at time $t+1$. The variable S_t denotes the average of S_t^a and S_t^b . Also, F_t denotes the average of F_t^a and F_t^b . All exchange rates are expressed as foreign currency units per U.S. dollar.

We consider two versions of the carry trade distinguished by how bid-ask spreads are treated. In both versions we normalize the size of the bet to one U.S. dollar. In the first version we sell x_t dollars forward according to the rule:

$$x_t = \begin{cases} +1 & \text{if } F_t \geq S_t, \\ -1 & \text{if } F_t < S_t, \end{cases} \quad (1)$$

We refer to this strategy as the 'naive carry trade'. This strategy is optimal if agents are risk neutral with respect to nominal payoffs, agents believe that $1/S_{t+1}$ is a martingale ($E_t(1/S_{t+1}) = 1/S_t$), and agents can trade at the average of bid and ask exchange rates.

In the second version we sell x_t dollars forward according to the rule:

$$x_t = \begin{cases} +1 & \text{if } F_t^b/S_t^a > 1, \\ -1 & \text{if } F_t^a/S_t^b < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We refer to this strategy as the 'transaction-cost-based carry trade'. This strategy is optimal if agents are risk neutral with respect to nominal payoffs and agents believe that $1/S_{t+1}^a$ and $1/S_{t+1}^b$ are martingales.

¹The currencies included in our large portfolio are as follows, with those also included in the small portfolio being indicated by (S): Argentina, Australia, Austria, Belgium (S), Brazil, Bulgaria, Canada (S), Chile, Colombia, Croatia, Cyprus, Czech Republic, Denmark, Egypt, Estonia, Euro (S), Finland, France (S), Germany (S), Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy (S), Japan (S), Kazakhstan, Kenya, Korea, Kuwait, Latvia, Lithuania, Malta, Mexico, Morocco, Netherlands (S), New Zealand, Norway, Pakistan, Philippines, Poland, Portugal, Qatar, Romania, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland (S), Taiwan, Thailand, Tunisia, Turkey, United Arab Emirates, UK (S), and Ukraine.

²http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

In both versions of the carry trade the realized payoffs are given by:

$$z_{t+1} = \begin{cases} x_t \left(F_t^b / S_{t+1}^a - 1 \right) & \text{if } x_t > 0, \\ x_t \left(F_t^a / S_{t+1}^b - 1 \right) & \text{if } x_t < 0, \\ 0 & \text{if } x_t = 0. \end{cases} \quad (3)$$

Recall that the sample period over which we have data varies by country. In the naive carry trade x_t is always different from zero. We consider two portfolios of naive carry-trade strategies. The first or ‘large’ portfolio gives equal weight at each point in time to all the currencies for which we have data. The second or ‘small’ portfolio is constructed in the same way but includes only the ten developed country currencies considered in Burnside et al. (2006).

We also consider two portfolios of the ‘transaction-cost-based carry trade’. For this strategy x_t can be equal to zero [see (2)], so the trader does not necessarily take positions in all available currencies. The first portfolio gives equal weight at each point in time to all the currencies for which we have data and for which x_t is different from zero. The second portfolio is constructed in the same way but includes only the ten developed country currencies considered in Burnside et al. (2006).

Table 1 reports statistics for the average, standard deviation, and Sharpe ratio of payoffs to our different portfolios. We begin by considering the payoffs to the naive carry trade computed under the assumption that it is possible to trade at the average of bid and ask rates, so that $z_{t+1} = x_t (F_t / S_{t+1} - 1)$. The resulting Sharpe ratio is very large: 0.27 on a weekly basis and 1.92 on an annualized basis. However, if we compute the payoffs to the carry trade according to (3), so that bid-ask spreads are taken into account, then the Sharpe ratio is actually *negative* (-0.05). We conclude that bid-ask spreads are very large in the sense that ignoring them leads to grossly misleading inference about the profitability of currency speculation strategies.

We now consider the payoffs to the transaction-cost-based carry trade. Table 1 indicates that the Sharpe ratio for this strategy is very high: 0.18 on a weekly basis and 1.32 on an annualized basis. So our simple modification of the naive carry trade is sufficient to transform a negative Sharpe ratio into a Sharpe ratio that is higher than that of the U.S. stock market.³ This large Sharpe ratio reflects, in part, that the trader takes transactions costs into account when deciding whether to take an active position in a given currency. There are 12 currencies

³The annualized monthly Sharpe ratio of the U.S. stock market is 0.23 for the period from October 1996 to August 2006 and 0.35 for the period from the beginning of Kenneth French’s data (January 1963) to August 2006.

for which the trader never takes a position and an additional 12 currencies for which the trader makes fewer than 10 trades over the whole sample. Moreover, the top ten (20) countries account for 60 (80) percent of all trades. We conclude that it is critical to take transactions costs into account in forming and evaluating currency speculation strategies.

Clearly a speculation strategy that involves fewer trades would mitigate the impact of transactions costs. The strategies discussed above involve settling existing positions and taking new positions on a weekly basis. To reduce transactions costs we could use forward contracts with longer maturities. To investigate the profitability of longer-horizon strategies we re-did the calculations underlying Table 1 using monthly forward contracts. We find that the Sharpe ratios for the weekly and monthly trading horizons are very similar. In principle it might be possible to reduce transactions costs and generate higher Sharpe ratios by extending the maturity of the forward contracts beyond one month. But, in practice, most forward contracts are available only for short horizons, typically less than a year. Because of data limitations we cannot investigate the profitability of carry-trade strategies with trading horizons exceeding one month.

One way to trade at horizons beyond those for which forward contracts are available is to borrow low-interest-rate currencies and lend high-interest-rate currencies. But this strategy can involve high transactions costs. Moreover, this strategy exposes the trader to default risk on both interest and principal invested in the high interest rate currency. In contrast, with the forward strategy the trader is exposed to default risk solely with respect to the forward contract payoff.

We now analyze the source of the high Sharpe ratio associated with the carry-trade strategy. The diversification effect that results from combining individual currencies into portfolios is the main reason why the Sharpe ratio is higher for the large portfolio than for the small portfolio. According to Table 1 the average payoff to the small and large portfolios are identical (0.0010). However, the standard deviation of the large portfolio payoff is roughly half the standard deviation of the small portfolio payoff (0.0100 versus 0.0054).

An important shortcoming of our results is that our data set happens to exclude major currency crises.⁴ To assess the potential effect of currency-crisis episodes on our results we proceed as follows. Suppose that a country incurs a currency crisis at time $t + 1$. We assume that $1/(N + 1)$ of the portfolio is invested at time t in the currency crisis country, where

⁴Recall that our sample period is October 1997 to August 2006, so it excludes the Asian currency crises that occurred in June and July 1997. The data set also excludes the Korean currency crisis of October 1997.

N is the number of countries in the previously defined ‘large portfolio’. We reweight the other currencies accordingly. In addition, we assume that the time $t + 1$ payoff to that portion of the portfolio is $S_t/S_{t+1} - 1$. This modification of the portfolio is conservative in the sense that it penalizes the portfolio by the full extent of the currency depreciation between t and $t + 1$, but ignores the interest differential earned during and before the crisis period. We also assume that the trader does not take any position on the currency until that point in time where our data set includes the currency in question. Using this procedure we incorporate five crises into our portfolio: Argentina (1/7/02), Brazil (1/19/99), Korea (12/10/97), Russia (9/7/98) and Turkey (2/23/01). In the case of Argentina we measure the devaluation in the week in which the fixed exchange rate regime was abandoned. For all other countries we use the largest weekly devaluation that occurred during the crisis period. Our procedure translates in to a one time weekly loss of -29 percent, -23 percent, -32 percent, -56 percent, and -29 percent for Argentina, Brazil, Korea, Russia, and Turkey, respectively. While these losses are very large, the diversified nature of the trader’s portfolio mitigates the impact of the crisis on the overall portfolio payoff.

According to Table 1, the currency crisis modification reduces the weekly Sharpe ratio, which falls from 0.183 to 0.097. However, this reduced Sharpe ratio is still roughly the same as that of the small portfolio. So, even under our conservative assumptions, an investor would do just as well with the crisis-modified portfolio as with the small portfolio.

To make concrete the properties of the different portfolio payoffs discussed above we proceed as follows. We use the realized payoffs to compute the cumulative realized return to committing one dollar in the beginning of the sample to a particular carry-trade portfolio and reinvesting the proceeds at each point in time. The agent starts with one U.S. dollar in his bank account and bets that dollar in the currency strategy. From that point forward the agent bets the balance of his bank account on the carry trade. The resulting payoffs are deposited or withdrawn from the agent’s account. Since the currency strategy is a zero-cost investment, the agent’s net balances stay in the bank and accumulate interest at the Treasury Bill rate. It turns out that the bank account balance never becomes negative in our sample.

The upper-left-hand panel of Figure 1 displays the cumulative realized returns to the large carry-trade portfolio. For comparison the upper-right-hand panel of Figure 1 displays these returns as well as the cumulative realized return associated with the U.S. stock market

and the Treasury-bill rate. Not surprisingly, the carry trade dominates a strategy of investing in Treasury bills. More interestingly, the total realized cumulative return to the carry-trade strategy is higher than that of the cumulative returns associated with the U.S. stock market. Moreover, the volatility of the large portfolio payoffs is much smaller than that of the U.S. stock market returns (0.0056 versus 0.0246). These observations explain why the Sharpe ratio associated with the large portfolio is so much larger than that of the stock market (0.183 versus 0.032 on a weekly basis).

The bottom left-hand panel displays the cumulative payoff to the large and small portfolios. The cumulative payoff of these two portfolios is similar but the small portfolio payoffs exhibit lower volatility. Finally, the bottom right-hand panel of Figure 1 displays the cumulative payoff to the large portfolio and the crises-adjusted portfolio discussed above. The main impact of the currency crisis correction is to reduce the payoff to the portfolio. Still, comparing the top and bottom right-hand panels of Figure 1 we see that the cumulative returns to the crisis-adjusted portfolio are higher than those of the U.S. stock market.

Why is the Sharpe ratio associated with the carry trade so high? Burnside et al. (2006) argue that it is difficult to interpret these payoffs as compensation for agents bearing risk, as measured by traditional risk factors such as real consumption growth. Because of data availability issues, the only conventional risk factor that we have available at the weekly frequency is the return to the U.S. stock market. We regress the payoffs to our portfolios on this variable and a constant. For the large portfolio the slope coefficient is 0.0194 with a standard error of 0.0087. For the small portfolio the slope coefficient is -0.0041 with a standard error of 0.0235. Finally, for the crisis-modified portfolio the slope coefficient is 0.0243 with a standard error of 0.0103. In all cases the slope coefficient estimates are small and, at best, marginally significant. Overall, these results are consistent with the key conclusion in Burnside et al. (2006): large Sharpe ratios are not a compensation for risk as traditionally measured.

So far we have emphasized the mean and variance of the payoffs to currency speculation. These statistics are sufficient to characterize the distribution of payoffs only under the assumption of normality. As Table 1 indicates, there is little evidence of skewness in either the small and large portfolio payoffs or in the U.S. stock market returns. While there is some evidence of kurtosis in the large portfolio, those payoffs do not appear to be more kurtotic than those associated with the U.S. stock market. We can easily reject the hypothesis that

the distribution is normal for the both the small and large portfolios of carry trade and the U.S. stock market. Not surprisingly there is more evidence of skewness and kurtosis in the crisis-modified carry-trade portfolio.

To assess the economic significance of these deviations from normality we confront a hypothetical trader with the possibility of investing in the U.S. stock market and wagering bets on the carry-trade. The trader's problem is given by,

$$\begin{aligned} \max_{\{C_t, X_{t+1}^s, X_{t+1}^c\}_{t=0}^{\infty}} U &= E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma} - 1}{1-\sigma} \right) \\ \text{s.t. } C_t &= Y_t + X_t^s(1 + r_t^s) + X_t^c r_t^c - X_{t+1}^s, \end{aligned} \quad (4)$$

Here C_t denotes consumption, Y_t is an exogenous income endowment normalized to one at time zero and assumed to grow at an annual rate of 1.9 percent (the growth rate of U.S. per capita GDP in our sample period), X_t^s and X_t^c are the end-of-period $t - 1$ investments in the U.S. stock market and the carry trade, respectively. The variables r_t^s and r_t^c are the time- t realized real return to the U.S. stock market, and the real payoff to the carry trade, respectively. We assume that r_t^c and r_t^s are generated by the joint empirical distribution of returns to the U.S. stock market and to the carry trade. We consider the large and small portfolio as well as the crisis-modified carry-trade portfolio.

It is useful to define the ratios $x_t^s = X_t^s/Y_t$ and $x_t^c = X_t^c/Y_t$. We assume that the trader chooses constant values $x_t^S = x^S$ and $x_t^C = x^C$ for all t . We choose a value of σ equal to five. For the large portfolio we find that $x^S = 0.26$, $x^C = 6.24$. For the small portfolio we find that $x^S = 0.37$, $x^C = 1.92$. For the crisis-adjusted portfolio we find that $x^S = 0.31$, $x^C = 2.54$. So, even though the distribution of payoffs to the crisis-adjusted portfolio has significantly fatter tails than those of a comparable normal distribution, the agent still wants to place very large bets on the carry trade.

It is always possible to rationalize the Sharpe ratios that we document by appealing to a peso problem, i.e. agents place positive weight on very large tail events that have not materialized in the sample. The problem with this explanation is that, by construction, there is no evidence to support it.

In summary, we show that there are large Sharpe ratios associated with the carry trade. The payoffs to the carry trade are uncorrelated with U.S. stock market returns. Our results raise an obvious question: if the large Sharpe ratios associated with the carry trade cannot be interpreted as compensating agents for risk, how can they persist as an equilibrium

phenomenon? Burnside et al. (2006) argue that transactions and microstructure frictions drive a wedge between average and marginal Sharpe ratios. By the latter we mean the Sharpe ratio associated with the last dollar bet on currency speculation strategies. Burnside et al. (2006) argue that for developed countries these frictions can explain large average Sharpe ratios and marginal Sharpe ratios that are close to zero. It remains an open question whether these frictions can explain the large Sharpe ratios associated with portfolios that include emerging market currencies.

REFERENCES


- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. 2006. “The Returns to Currency Speculation,” National Bureau of Economic Research Working Paper 12489.
- Reuters. 2003. “WM/Reuters Spot and Forward Rates: A Fixed Reference in a Changing World.”

TABLE 1—PAYOFFS TO PORTFOLIOS, 97:10—06:11

	Mean	Std. Dev.	Sharpe Ratio
Large carry-trade portfolio	0.0010 (0.0003)	0.0056 (0.0003)	0.1832 (0.0527)
Small carry-trade portfolio	0.0010 (0.0005)	0.0100 (0.0005)	0.1001 (0.0505)
Crisis-modified carry-trade	0.0006 (0.0003)	0.0066 (0.0005)	0.0971 (0.0544)
Value-weighted US stock market	0.0008 (0.0011)	0.0246 (0.0016)	0.0320 (0.0444)
	Skewness	Kurtosis	Jarque- Bera
Large carry-trade portfolio	-0.27 (0.21)	1.63 (0.36)	58 (0.00)
Small carry-trade portfolio	-0.55 (0.30)	2.11 (1.20)	111 (0.00)
Crisis-modified carry-trade	-1.71 (0.56)	9.54 (3.16)	2018 (0.00)
Value-weighted US stock market	-0.54 (0.34)	3.40 (1.40)	250 (0.00)

Notes: Portfolios are described in the main text. Payoffs are in dollars, per dollar traded. Standard errors in parentheses, p-values reported for the Jarque-Bera test.

FIGURE 1—THE CUMULATIVE RETURN TO INVESTING IN THE PORTFOLIOS



Note: The y-axis in each figure indicates the accumulated value (in U.S. dollars) of beginning with a balance of 1 U.S. dollar on 10/29/97 and rolling over the accumulated value of the investment weekly through 11/8/06. The portfolios are defined in the text.