NBER WORKING PAPER SERIES

CARRY TRADES AND RISK

Craig Burnside

Working Paper 17278
http://www.nber.org/papers/w17278

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2011

I thank Martin Eichenbaum, Barry Rafferty, Sergio Rebelo, Nikolai Roussanov, Lucio Sarno and an anonymous referee for helpful comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Carry Trades and Risk
Craig Burnside
NBER Working Paper No. 17278
August 2011
JEL No. F31

ABSTRACT

Carry trades, in which an investor borrows a low interest rate currency and lends a high interest rate currency, have been profitable historically. The risk exposure of carry traders might explain their high returns, but conventional models of risk do not work because traditional risk factors, used to price the stock market, do not price currency returns. Less traditional factors that are more successful in explaining currency returns, are, however, unsuccessful in explaining the returns to the stock market. More exotic models of "crisis risk" are another possibility, but I show that any time-variation in the exposure of the carry trade to market risk has been insufficient, in sample, to explain the average returns earned by carry traders. Instead, peso events remain a candidate explanation of the returns to the carry trade.

Craig Burnside
Department of Economics
Duke University
213 Social Sciences Building
Durham, NC 27708-0097
and NBER
burnside@econ.duke.edu

I. Introduction

In a foreign exchange carry trade, an investor borrows funds in a low interest rate currency and lends those funds in a high interest rate currency. The uncovered interest parity (UIP) condition states that the interest rate differential between riskless assets denominated in foreign and domestic currency is equal to the rate at which the foreign currency is expected to depreciate against the domestic currency. If the UIP condition held, an investor engaged in the carry trade would, therefore, expect a zero net payoff. One motivation for investors to engage in the carry trade is, however, that UIP does not appear to hold in the data.[1] If anything, high interest rate currencies are more likely to appreciate than depreciate against low interest rate currencies. Consequently, in historical data, carry trades have earned positive average returns in excess of the interest differentials between the relevant currencies,

If investors expect to earn the interest differential, why do they limit their trading in foreign exchange? The most obvious explanation is that carry trades are risky, and that the average returns to carry trades reflect a risk premium. In this paper, I review the evidence for and against a variety of risk premium based explanations. I first explore traditional factor models, ones that have been used to explain the returns to stock market portfolios. Examples include the CAPM, the Fama-French three factor model and the consumption-CAPM. I find that these traditional models fail to explain the returns to the carry trade, defined either as an equally weighted portfolio of carry trades (as in Burnside et al. (2006, 2011)) or as high-minus-low portfolio of carry trades (as in Lustig, Roussanov and Verdelhan (2011)). Risk-based explanations of the returns to the carry trade rely on identifying risk factors that covary with the returns. Traditional factors are either uncorrelated with carry trade returns, i.e. they have zero betas, or the betas are much too small to rationalize the magnitude of the returns to the carry trade.

I also examine less traditional factor models. These models adopt risk factors constructed specifically to price currency returns. I begin by studying the returns to portfolios of currencies that have been sorted according to the size of their forward discount. This sorting approach to portfolio construction has a long tradition in the finance literature (see Fama and French (1993)), and was brought to the literature on currency returns by Lustig and Verdelhan (2007). In studying a similar set of currency portfolios I find that three factor models are quite successful in pricing the cross-section of returns. These models are based on

[1] See Hodrick (1987) and Engel (1996) for reviews of the large literature documenting the failure of UIP.

Lustig, Roussanov and Verdelhan's (2011) model, which uses a high-minus-low carry trade factor, Menkhoff et al.'s (2010) model, which uses a global currency volatility factor, and Rafferty's (2010) model, which uses a global currency skewness or "currency crash" factor. Although these models have some success in explaining currency returns, I find that they do not explain stock returns. Given, as I argued above, that models that do reasonably well in explaining stock returns do not explain currency returns, it appears that there is no unifying risk based explanation of returns in these two markets.

One plausible explanation for the fact that one set of factors works for currency returns, while another is more successful in explaining stock returns, is that there is some degree of market segmentation. I find this explanation unattractive for the following reason. Although segmentation between currency and stock markets is plausible, factors that price carry trade portfolios ought to have some success in pricing other currency portfolios, such as those based on momentum or value. Burnside, Eichenbaum and Rebelo (2011), and Sarno et al. (2011a) present evidence, however, that the same nontraditional factors that price carry trade portfolios are unable to price momentum portfolios defined using short-term historical returns. Sarno et al. (2011a) report more mixed evidence for momentum defined using longer term historical performance. Sarno et al. (2011b) find that individual currency characteristics appear to be important in explaining the returns to currency momentum. Alternatively, it could be argued that empirical exercises involving currency returns reveal one component of the global investor's SDF, while those involving stock returns reveal a different component of it. I find this explanation unsatisfying because it effectively renders untestable SDF-based explanations of asset return anomalies.

Finally, I provide evidence that time varying market risk is unlikely to explain the returns to the carry trade. During the recent financial crisis, carry trade returns and stock market returns became more highly correlated. This might suggest that covariance at times of market distress explains the returns to the carry trade (see, for example, Lustig, Roussanov and Verdelhan (2011)). While this is an interesting conjecture, as I show here the degree of covariance seen in the data is insufficient. An alternative explanation is the one pursued by Burnside et al. (2011), who argue that periods of extreme risk aversion that have not been observed in sample (peso events) can explain the returns to the carry trade and the stock market. I argue, below, that important challenges for future research on peso event based models is that they need to explain the empirical success of the nontraditional factor models described above, the time variation in "risk premia" needed to explain the UIP puzzle, and the cross-section of stock returns. Julliard and Ghosh (2010) suggest that explaining the cross-section of returns is difficult in a model with rare consumption disasters and constant relative risk aversion preferences because rare disasters tend to reduce the cross-sectional dispersion of the model-implied consumption betas.

In Section II I define the carry trade and measure the returns to two carry trade portfolios in historical data. In Section III I derive theoretical pricing equations for risk-based explanations and I outline the empirical methods used for assessing them. In Section IV I present empirical results. In Section V I discuss time varying risk, rare events, and peso problems. In Section VI I conclude.

II. The Carry Trade: Basic Facts

A. What is a Carry Trade?

In the carry trade, an investor borrows funds in a low-interest-rate currency and lends in a high-interest-rate currency. Here, I let the domestic currency be the U.S. dollar (USD), and denote the rate of interest on riskless USD denominated securities as i_t. I denote the interest rate on riskless foreign denominated securities as i_t^*. Abstracting from transactions costs, the payoff to borrowing one USD in order to lend the foreign currency is:

$$(1+i_t^*)\,\frac{S_{t+1}}{S_t} - (1+i_t)\,, \tag{1}$$

where S_t denotes the spot exchange rate expressed as USD per foreign currency unit (FCU). The payoff to the carry trade strategy is, therefore:

$$z_{t+1} = \operatorname{sign}(i_t^* - i_t)\left[(1+i_t^*)\,\frac{S_{t+1}}{S_t} - (1+i_t)\right]. \tag{2}$$

The carry-trade strategy can also be implemented by selling the foreign currency forward when it is at a forward premium ($F_t \geq S_t$) and buying the foreign currency forward when it is at a forward discount ($F_t < S_t$). If the number of FCUs transacted forward is normalized to be $(1+i_t)/F_t$, then the payoff to this version of the strategy, denoted z_{t+1}, is

$$z_{t+1} = \operatorname{sign}(F_t - S_t)\frac{1+i_t}{F_t}\,(F_t - S_{t+1})\,. \tag{3}$$

Covered interest rate parity (CIP) implies that:

$$\frac{1+i_t}{1+i_t^*} = \frac{F_t}{S_t}. \tag{4}$$

When CIP holds, the expressions for z_{t+1}in equations (3) and (2) are equal to each other. So the strategies are equivalent.

I measure payoffs to the carry trade using equation (3). My empirical analysis focuses on the carry trade implemented at a one month horizon, so I mainly work with monthly payoffs. In order to assess the importance of real risk factors that are measured at the quarterly frequency, I also compute quarterly real excess returns to the carry trade. Letting s be the time index for quarterly data, and t be the time index for monthly data, so that $s = t/3$, the quarterly excess return in quarter s is defined as:

$$z_s^q = R_t R_{t-1} R_{t-2} - R_t^f R_{t-1}^f R_{t-2}^f, \tag{5}$$

where

$$R_t = 1 + i_{t-1} + z_t \tag{6}$$

is the gross monthly rate of return to investing in the carry trade, and

$$R_t^f = 1 + i_{t-1} \tag{7}$$

is the gross monthly risk free return. The quarterly real excess return in quarter s is simply

$$z_s^{qr} = z_s^q / (1 + \pi_s) \tag{8}$$

where π_s is the growth rate of the deflator for the consumption of nondurables and services from the U.S. National Income and Product Accounts.

B. Measuring the Returns to the Carry Trade

To measure the returns to the carry trade, I consider trades conducted on a currency by currency basis against the U.S. dollar. I also consider portfolio based carry trade strategies. I implement the trades with historical data using the forward market strategy described above. My data set consists of spot and forward exchange rates from Reuters/WMR and Barclays, available on Datastream, for the euro and the currencies of 20 countries: Australia, Austria, Belgium, Canada, Denmark, France, Germany, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, South Africa, Spain, Sweden, Switzerland, the U.K., and the U.S. I use the risk free rate from Kenneth French's database as a measure of the U.S. interest rate.

The raw data are daily observations of spot and one-month forward exchange rates. I use end of month values of these data to create monthly observations. The data span the period January 1976 to October 2010, with the sample varying by currency. Reuters/WMR exchange rate quotes against the British pound (GBP) are available beginning in 1976. Reuters/WMR exchange rate quotes against the USD are available from January 1997 to October 2010. I construct USD quotes over the longer sample by multiplying GBP/FCU quotes by USD/GBP quotes. I augment the data set with USD quotes from Barclays for the Australian, New Zealand and South African currencies from 1983 through 1996. Details of the data set are provided in an online appendix available at http://web.duke.edu/~acb8/ctapp.pdf.

Table I provides summary information about the profitability of carry trades between the U.S. dollar and the other 20 currencies. On average the annual excess return to the individual currency strategies is 4.6% with a typical standard deviation of 11.3% (on an annual basis) and an annualized Sharpe ratio of 0.42. This compares favorably with the performance of the U.S. stock market over the same period, which had an average excess return of 6.3%, a standard deviation of 15.7% and a Sharpe ratio of 0.40. While the average carry trade was profitable, the performance of the individual carry trade varies across currencies, with trades against the Swiss franc earning a low 0.6% annual excess return, and trade against the Danish krone earning a high 9.3% annual excess return.

When carry trades are combined in portfolios, their performance is more impressive still. One strategy I consider is to combine all the individual currency positions in an equally weighted portfolio with the total value of the bet normalized to 1 USD at the time it is initiated. I refer to this strategy as the EW carry trade, and it is the same strategy studied by Burnside et al. (2011). As Table I indicates, over the historical sample, the EW strategy had an average annual payoff of 4.6%, with a standard deviation of 5.1% and a Sharpe ratio of 0.90. This shows that there are large diversification gains to combining carry trades in a portfolio.

A second strategy is constructed as follows. In each period, the available currencies in my sample, including the USD, are sorted into five bins according to their forward discount against the USD (of course, the USD's forward discount is 0). The first bin includes those currencies with the smallest forward discounts (the lowest interest rates), the second bin the next smallest, etc., with the fifth bin consisting of those currencies with the largest forward discounts (and, therefore, the highest interest rates). I then compute the payoff associated with borrowing one dollar in order to invest equally in the riskless securities of the constituent currencies of each bin.[2] This is equivalent to calculating the average value of

[2] This is the same procedure used by Lustig, Roussanov and Verdelhan (2010) and Menkhoff et al. (2010).

$(1 + i_t)(S_{t+1} - F_t)/F_t$ for the currencies with the bin. The USD is treated the same as other currencies, with the payoff being zero. This procedure produces five currency portfolios that I refer to as S1, S2, S3, S4 and S5. The second carry trade portfolio that I study involves investing 1 USD in S5 and -1 USD in S1. This is, effectively, equivalent to executing a carry trade in which the investor borrows the low interest rate currencies in S1 in order to lend funds in the high interest rate currencies in S5. Like Lustig, Roussanov and Verdelhan (2011), I refer to this portfolio as the HML carry trade portfolio.[3] As Table I indicates, over the historical sample, the HML carry trade strategy had an average annual excess return of 6.0%, with a standard deviation of 9.5% and a Sharpe ratio of 0.63. Although the HML portfolio had larger average returns than the EW portfolio over the historical sample, it is important to note that it is more highly leveraged than the EW portfolio since it involves a bet size of two dollars instead of one dollar.

Figure 1 shows 12-month moving averages of the carry trade portfolio and the U.S. stock market excess returns. Two features of the returns are worth noting. First, the EW and HML carry trade returns are positively correlated with each other (the correlation of the raw monthly returns is 0.51) but not perfectly so. Second, neither carry trade portfolio is strongly correlated with the US stock market, despite their common poor performance in the 2008-09 financial crisis.

Currency movements are often characterized as being highly skewed. For example, Brunnermeier, Nagel and Pedersen (2008) note the saying among traders that "exchange rates go up by the stairs and down by the elevator." While there is evidence that large forward discounts are associated with realized negative skewness of carry trade returns (and large premia with positive skewness), the amount of skewness exhibited by the EW and HML carry trade portfolios is less than that exhibited by the U.S. stock market, and for the EW portfolio it is not statistically significant (see Table I). Currency payoffs display excess kurtosis, with noticeable central peakedness, especially in the case of the EW portfolio.

[3]The HML portfolio is a close cousin of a market index, the Deutsche Bank G10 Currency Future Harvest (DBCFH). The DBCFH index takes positions in up to six currencies from a list of ten. The index is formed by taking equally-weighted long positions vis-à-vis the USD in the three currencies with the highest interest rates, and symmetric short positions vis-à-vis the USD in the three currencies with the lowest interest rates. The currency composition of the DBCFH portfolio is rebalanced quarterly, while the composition of my HML portfolio is rebalanced monthly.

III. Pricing the Returns to the Carry Trade

Risk-based explanations of the returns to the carry trade begin from the premise that there is an SDF that prices these returns. In particular, since the carry trade is a zero net-investment strategy, the payoff, z_t, must satisfy:

$$E_t\left(M_{t+1}z_{t+1}\right) = 0. \tag{9}$$

Here M_{t+1} denotes the SDF that prices payoffs denominated in dollars, while E_t is the mathematical expectations operator given information available at time t. Equation (9) implies that:

$$p_t \equiv E_t\left(z_{t+1}\right) = -\frac{\text{cov}_t\left(M_{t+1}, z_{t+1}\right)}{E_t\left(M_{t+1}\right)}. \tag{10}$$

The variable p_t is referred to as the conditional risk premium and corresponds to the conditional expectation of the payoff. As equation (10) suggests, one approach to learning about risk premia is to build a forecasting model for the payoffs to the carry trade. An approximation to the mathematical expectation in equation (10) is implicit in any forecasting model. Therefore, model forecasts correspond to estimates of the risk premium (Fama (1984)).

Consider an example of an individual currency carry trade in which the domestic interest rate exceeds the foreign interest rate, i.e., $i_t > i_t^*$, or, equivalently that the foreign currency is at a forward premium: $F_t > S_t$. Assume that the carry trader sells S_t^{-1} units, rather than $(1+i_t)/F_t$ units, of the foreign currency forward. In this case his payoff is:

$$z_{t+1} = \frac{F_t - S_{t+1}}{S_t}, \tag{11}$$

so that (10) becomes:

$$p_t = fp_t - E_t\delta_{t+1} \equiv \frac{\text{cov}_t\left(M_{t+1}, \delta_{t+1}\right)}{E_t\left(M_{t+1}\right)}, \tag{12}$$

where $\delta_{t+1} = \left(S_{t+1} - S_t\right)/S_t$ is the rate of appreciation of the foreign currency, and $fp_t = \left(F_t - S_t\right)/S_t$ is the forward premium.

Several features of equation (12) are worthy of note. First, to the extent that the exchange rate is well-approximated by a martingale, the risk premium to a carry trade is simply equal to the forward premium, i.e., if $E_tS_{t+1} = S_t$ then $p_t = fp_t$. Second, for many currency pairs δ_{t+1} and fp_t covary negatively in sample.[4] This implies, given equation (12), that for these currency pairs

$$\text{var}(p_t) \geq \text{var}\left(fp_t\right) + \text{var}(E_t\delta_{t+1}) \tag{13}$$

[4] For early surveys see Hodrick (1987) and Engel (1996). For recent evidence, see Burnside et al. (2006). Bekaert and Hodrick (1992) provide a broad set of evidence on the predictability of currency returns.

and

$$\text{cov}(p_t, E_t\delta_{t+1}) \leq -\text{var}(E_t\delta_{t+1}). \tag{14}$$

These inequalities, derived by Fama (1984) and discussed by Engel (1996), put restrictions on the time series properties of the risk premium that could clearly be tested for a particular model. Any good forecast based model, however, will satisfy (13) and (14) by construction. To see this, let $\hat{\delta}_{t+1}$ be the time series of one step ahead forecasts of δ_{t+1} produced by a forecasting model, and let the estimated risk premium be $\hat{p}_t = fp_t - \hat{\delta}_{t+1}$. As long as the forecasts have the property that $\text{cov}(\hat{\delta}_{t+1}, fp_t) \leq 0$ it follows that $\text{var}(\hat{p}_t) \geq \text{var}\,(fp_t) + \text{var}(\hat{\delta}_{t+1})$ and $\text{cov}(\hat{p}_t, \hat{\delta}_{t+1}) \leq -\text{var}(\hat{\delta}_{t+1})$. Finally, the challenge posed to economic researchers by (12) is that the risk premium is equal to the covariance term on the right hand side of the equation. A risk based explanation of the returns to the carry trade, therefore, relies on identifying an SDF that covaries with the rate of appreciation of the foreign currency. If in sample risk explains the returns to the carry trade, then this SDF should correspond to some observable time series. As I argue below, finding such an SDF remains an elusive goal of economic research.

My exploration of candidate SDFs focuses on the unconditional moment condition restriction corresponding to (9):

$$E\,(Mz) = 0. \tag{15}$$

I consider SDFs that are linear in vectors of risk factors:

$$M_t = \xi\left[1 - (f_t - \mu)' \, b\right]. \tag{16}$$

Here ξ is a scalar, f_t is a $k \times 1$ vector of risk factors, $\mu = E(f_t)$, and b is a $k \times 1$ vector of parameters.

Since the parameter ξ is not identified by (15) I set it equal to 1, so that $E(M) = 1$. Given this assumption, equation (15) implies that:

$$E\,(z) = -\text{cov}\,(M, z)\,. \tag{17}$$

Given the model for M given in equation (16), equation (17) can be rewritten as

$$E\,(z) = \text{cov}\,(z, f)\, b \tag{18}$$

or as

$$E\,(z) = \underbrace{\text{cov}\,(z, f)\, \Sigma_f^{-1}}_{\beta} \underbrace{\Sigma_f b}_{\lambda}, \tag{19}$$

where Σ_f is the covariance matrix of f_t. Equation (19) is the beta representation of the model. The betas, which are population coefficients in a regression of z_t on f_t, measure the risk exposure of the payoff, while λ is a $k \times 1$ vector of risk premia that is not specific to the payoff.

I assess risk based explanations of the returns to the carry trade in two ways. First, I ask whether there are risk factors for which the payoffs to the carry trade have statistically and economically significant betas. To answer this question, I run a simple time series regression of each portfolio's excess return on a vector of candidate risk factors:

$$z_{it} = a_i + f_t'\beta_i + \epsilon_{it}, \quad t = 1, \ldots, T, \text{ for each } i = 1, \ldots, n, \tag{20}$$

where T is the sample size, and n is the number of portfolios being studied.

Second, I ask whether these betas, combined with estimates of λ, can explain the returns to the carry trade according to (19). One way to answer this question is to run a cross-sectional regression of average portfolio excess returns on the estimated betas:

$$\bar{z}_i = \hat{\beta}_i'\lambda + \alpha_i, \qquad i = 1, \ldots, n, \tag{21}$$

where $\bar{z}_i = \frac{1}{T}\sum_{t=1}^{T} z_{it}$, $\hat{\beta}_i$ is the OLS estimate of β_i obtained above, and α_i is a pricing error. Let the OLS estimator of λ be $\hat{\lambda} = (\hat{\beta}'\hat{\beta})^{-1}\hat{\beta}'\bar{z}$, where $\bar{z}$ is an $n \times 1$ vector formed from the individual mean returns. Rather than actually running the cross-sectional regression I estimate the parameters of the SDF using the Generalized Method of Moments (GMM, Hansen 1982) and the moment restriction (15) along with $E(f) = \mu$. Equation (15) can be rewritten as

$$E\left(z\left[1 - (f - \mu)' b\right]\right) = 0, \tag{22}$$

where z is an $n \times 1$ vector of excess returns, and is equivalent to (18). The GMM estimators of μ and b at each step are $\hat{\mu} = \bar{f}$ and

$$\hat{b} = \left(d_T' W_T d_T\right)^{-1} d_T' W_T \bar{z}, \tag{23}$$

where d_T is the sample covariance matrix of z with f, and W_T is a weighting matrix. Estimates of λ are obtained from $\hat{b}$ as $\hat{\lambda} = \hat{\Sigma}_f \hat{b}$, where $\hat{\Sigma}_f$ is the sample covariance matrix of f. The model's *predicted mean returns* are $d_T\hat{b}$, and are estimates of the right hand side of (18). The *pricing errors* are the residuals, $\hat{\alpha} = \bar{z} - d_T\hat{b}$. In the first GMM step the weighting matrix is $W_T = I_n$, and the estimate of λ and the pricing errors are the same as the ones obtained from the cross-sectional regression described above. In subsequent GMM steps the weighting matrix is chosen optimally. A test of the pricing errors is obtained as $J = T\hat{\alpha}'V_T^{-1}\hat{\alpha}$, where V_T is a consistent estimate of the asymptotic covariance matrix of $\sqrt{T}\hat{\alpha}$ and the inverse is generalized. The test statistic is asymptotically distributed as a χ^2 with $n-k$ degrees of freedom. I mainly focus on results obtained by iterating on the GMM estimator to convergence. Burnside (2010a) shows that the first stage, second stage and iterated GMM estimators have similar size properties when calibrated linear factor models are used as the data generating process. However, the iterated estimator has much greater power to reject misspecified models. Burnside (2007) provides further details of the GMM procedure.

IV. Empirical Findings

In this section I use the empirical methods outlined in the previous section to determine whether there is a candidate SDF that can price the returns to the carry trade. I begin by studying risk factors that have traditionally been used to explain stock returns. These include the CAPM, the Fama-French three factor model, models in which industrial production growth and stock market volatility are used as risk factors, and consumption based models. I show that none of the SDFs specified in terms of these traditional risk factors explains the returns to the carry trade. I then turn to less traditional models in which the risk factors are, themselves, derived from currency returns. These models have varying degrees of success in pricing carry trade returns.

A. Traditional Risk Factors

A.1 Models for Monthly Returns

Table II summarizes the results of running the time series regressions described by equation (20) for monthly models. Consider, first the CAPM model, which uses the market excess return as a single factor. For the EW carry trade portfolio the beta is statistically insignificant. It is also economically small. To see this, consider that in order for the CAPM model to explain the return to holding the value weighted U.S. stock market, the annualized value of λ must be approximately equal to the average excess return of the stock market, since the beta of the market return is 1. Given that the beta of the EW carry trade portfolio is 0.028, the CAPM model therefore predicts that the average excess return of the EW carry trade should be 0.028 times the average return on the market, or 0.2%, as opposed to the 4.6% return indicated in Table I . For the HML carry trade portfolio the beta is larger (0.163), and statistically significant, but it is still economically small. The CAPM model predicts that the average excess return of the HML carry trade should be 0.163 times the average return on the market or 1.0%, as opposed to the 6.0% return indicated in Table I .

The second model is the three factor model proposed by Fama and French (1993), which uses the excess return to the value-weighted U.S. stock market (Mkt-Rf), the size premium (SMB), and the value premium (HML) as risk factors. Here, the beta associated with the Mkt-Rf factor is statistically significant for both carry trade portfolios but it remains economically small. The beta associated with the SMB factor is small and statistically insignificant for both carry trade portfolios. The beta associated with the HML factor is small and statistically insignificant for the EW carry trade portfolio, but is statistically significant and small for the HML carry trade portfolio.

The third model uses industrial production growth as a risk factor, while the fourth model uses industrial production growth in conjunction with the Mkt-Rf factor. The betas associated with industrial production growth are not statistically significant for either carry trade portfolio.

The fifth model uses realized stock return volatility (measured monthly using daily observations of Mkt-Rf) as a risk factor, while the sixth model uses stock volatility in conjunction with the Mkt-Rf factor. For the EW carry trade portfolio, the beta associated with stock market volatility is not statistically significant in either model. For the HML carry trade portfolio, the beta associated with stock market volatility is statistically significant, but not when Mkt-Rf is included in the regression.

Table III presents results from estimating each of these models using the iterated GMM estimator. The models are estimated using the EW and HML carry trade portfolios, as well as Fama and French's 25 portfolios sorted on the basis of book to market value and size. First, it is worth noting that in every case the pricing errors of the carry trade portfolios are statistically significant. None of the models explains the returns earned by these strategies. Second, all of the models are rejected, at the 5% level, by the test of the pricing errors. The only model with reasonably good fit is the Fama and French (1993) three factor model. For this model two of the slope coefficients (on Mkt-Rf and HML) are statistically significant, while the third coefficient (on SMB) is close to being significant. The R^2 measure of fit for the model is 0.38. But the model does a very poor job of explaining the returns to the carry trade portfolios.

A.2 Models for Quarterly Returns

Here I consider three risk factors: the growth rate of real consumption of nondurables and services, the growth rate of the service flow from the real stock of durables, and the market return. I consider each of these factors individually, but also use them together in a three factor model following Yogo (2006), who used the three factor model to study stock market returns. For the consumption variables I use both conventional timing (where consumption growth and the returns are measured in the same quarter) and Campbell's (2003) timing where consumption growth is measured in the quarter after the returns are realized.

Table IV summarizes the results of running the time series regressions of quarterly real excess returns on the risk factors described above. Consider, first the C-CAPM model, which uses real consumption growth (nondurables plus services) as a single factor. This model can be considered a linear approximation to a simple representative agent model in which households have constant relative risk aversion preferences with risk aversion parameter b. For both carry trade portfolios the beta is small and statistically insignificant. The betas are larger, but remain statistically insignificant, when Campbell's timing is used, and the beta for the HML carry trade has counterintuitive sign.

Another model that has received broad attention in the literature, is a three factor model, which I refer to as the Extended C-CAPM model, which includes consumption growth, durables growth and the market return as risk factors. This model is a linear approximation to a representative agent model in which households have recursive preferences over the two types of consumption good. Yogo (2006) uses this model to explain stock returns, while Lustig and Verdelhan (2007) use it to explain currency returns. Before turning to the full three factor model, I study models that use durables growth and the market return as single factors.

Consider, first, the model with the real growth rate of the stock of consumer durables (interpreted as the growth rate of the service flow from durables) as a single factor. For both carry trade portfolios the beta is statistically insignificant, and it has counterintuitive sign for the HML carry trade. The same is true for Campbell's timing, largely because the growth rate of the stock of durables is highly serially correlated.

Consider, next, the model that uses the market return, as opposed to the market excess return used in the CAPM model, as a single factor. As we saw for the CAPM using monthly data, the beta is small and statistically insignificant for the EW carry trade portfolio. For the HML carry trade portfolio the beta is statistically significant, but economically small.

Finally, I enter all three factors together in a single model. None of the factors is statistically significant for the EW carry trade portfolio, and neither consumption factor displays any significance for the HML carry trade portfolio. The market return has significance for the HML carry trade portfolio, but the coefficient is quantitatively small. This evidence, against a consumption-based explanation of the returns to the carry trade, is overwhelming.

Table V presents results from estimating the C-CAPM and extended C-CAPM models using iterated GMM. The models are estimated using the real excess returns of the EW and HML carry trade portfolios, as well as Fama and French's 25 portfolios sorted on the basis of book to market value and size. First, it is worth noting that none of the estimated b parameters are statistically significant at the 5% level, though the λ associated with the market return is statistically significant for the Extended C-CAPM model. Second, for all cases, the pricing errors of the carry trade portfolios are statistically significant, the models are rejected at the 5% level on the basis of the J-statistic, and the R^2 measures of fit are negative.

A.3 Discussion

The results in this section suggest that traditional risk factors cannot explain the returns to carry trade portfolios. At best, the models considered here explain very little of the average returns to the EW and HML carry trade portfolios, leaving unexplained economically large and statistically significant pricing errors. In every case the models can also be rejected based on statistical tests of the pricing errors.

The reader may be puzzled by the poor performance of the Extended C-CAPM model given its prior apparent success in explaining stock returns (Yogo (2006)). The bottom line is that the factors in the model simply do not produce the amount of spread in the betas required for the model to be a success. It is not easy to illustrate the problem for a three factor model, because there are betas in three dimensions and the partial explanatory power of each factor is what is relevant. To deal with this issue, I construct a calibrated SDF, $\hat{m}_t^Y = 1 - (f_t - \bar{f})'\hat{b}^Y$ where $\hat{b}^Y$ is Yogo's (2006) estimate of the b vector for the Extended C-CAPM model, f_t is the vector of relevant risk factors and $\bar{f}$ is their sample mean. I then compute betas for the two carry trade portfolios and the 25 Fama-French portfolios with respect to $\hat{m}_t^Y$. If the model explains the average returns on the 27 portfolios, then 27 estimated betas and the 27 average returns should line up (with negative slope) in a scatter plot. Instead, they are approximately uncorrelated in the cross-section (the correlation coefficient is -0.02). Also, only four out of the 27 betas are statistically significant at the 5 percent level. The model simply does not do a good job of explaining currency returns or stock returns over the period 1976 to 2010.

B. Factors Derived from Currency Returns

I turn, now, to less traditional factor models. Here the factors are, themselves, derived from currency returns. In creating factors in this way, the literature, beginning with Lustig, Roussanov and Verdelhan (2011), takes inspiration from the literature on stock returns, where it is common to choose risk factors that are, themselves, the returns to particular investment strategies. For example, having identified the size and value premia, Fama and French (1993) construct new risk factors which are the return differentials between, respectively, small and large firms (SMB), and high and low value firms (HML). Similarly, many researchers have used a momentum risk factor to explain stock returns, since Jegadeesh and Titman (1993) identified the momentum anomaly.

B.1 Currencies Sorted by Forward Discount

The Fama and French factors and the momentum factors are created by sorting firms on characteristics. Take, for example, Fama and French's SMB factor. To construct this factor, Fama and French (1993) sort firms by their market value. This sorting is done once per year at the end of June. Firms in the bottom third in terms of size are used to form a portfolio of small firms. Firms in the top third in terms of size are used to form a portfolio of large firms. The SMB factor is the return differential between the small firm portfolio and the large firm portfolio in each period.

In a similar way, Lustig and Verdelhan (2007) and Lustig, Roussanov and Verdelhan (2011) sort currencies into, respectively, eight and six portfolios according to their forward discount against the U.S. dollar. The sorting is done period by period. Each portfolio is equally weighted and represents the excess return to going long in the constituent currencies while going short in the U.S. dollar. In Section II.B I constructed the S1–S5 portfolios in an analogous way, and Menkhoff et al. (2010) follow a similar procedure. These five portfolios are the focus of my empirical work in this section.

Table VI shows the average returns of the five portfolios. Notice that they are monotonically increasing going from S1 to S5. This establishes that sorting currencies on the basis of the forward discount "works," in that it produces a set of portfolios with different expected returns, where the ordering of the expected returns aligns with an observed characteristic of the underlying assets. This result should not come as a surprise. Meese and Rogoff (1983) established that it is hard to produce a currency forecasting model that beats a random walk. If the change in the spot rate for each currency was exactly a random walk then the expected excess returns of S1 through S5 would be exactly equal to the forward discount period by period. In that circumstance, sorting done on the basis of the forward discount could not fail to provide portfolios with ordered average returns.

B.2 Factors Created from Sorted Portfolios

Lustig, Roussanov and Verdelhan (2011) take the sorting approach one step further by constructing two risk factors that they then use to price the cross-section of their six portfolios. The first risk factor, which they call the dollar risk factor, and denote RX, is simply the average excess return of the six portfolios. The second risk factor, which they denote $\mathrm{HML}_{\mathrm{FX}}$, is the return differential between the sixth portfolio (the largest forward discount) and the first portfolio (the smallest forward discount). In an analogous way, I construct two risk factors, one denoted DOL, which is simply the average excess return of the S1–S5 portfolios described above, the other being the excess return to the HML carry trade portfolio. Like Lustig, Roussanov and Verdelhan I use these two risk factors to study the cross-section of returns of the S1–S5 portfolios.

B.3 A Currency Volatility Factor

Menkhoff et. al (2010) use a factor analogous to DOL, and a factor that measures global currency volatility, to study portfolios similar to S1–S5. Their volatility factor is constructed on a monthly basis and is the average intramonth realized volatility of the daily log changes in the value of each currency (available in their sample) against the USD. In studying the importance of volatility as a pricing factor, they take inspiration from an earlier literature that suggests exposure to volatility helps explain stock returns (Ang et al. (2006), Da and Schaumburg (2009)). To re-examine their evidence, I construct a volatility factor, VOL, which is measured monthly, and is the average sample standard deviation of the daily log changes in the values of the currencies in my sample against the USD.

B.4 A Currency Skewness Factor

Rafferty (2010) uses a factor analogous to DOL, and a factor that measures global currency skewness, to study portfolios similar to S1–S5. He takes inspiration from the literature on currency crashes and liquidity, for example, Brunnermeier, Nagel and Pedersen (2009). They tell a story in which carry trades drive currency dynamics until liquidity dries up, traders draw back their positions in tandem, and the currencies which are the targets of their trades crash. Accordingly, one might expect that carry trades are risky because high interest rate currencies are exposed, in tandem, to these crashes. To measure coordinated "crashing" of target currencies, Rafferty constructs a global currency skewness factor. This factor sorts currencies into two groups, one with positive forward discounts (equivalently, positive interest differentials) and one with negative forward discounts. On a monthly basis he measures the realized skewness of the currencies in the first group, and the *negative* of the skewness of the currencies in the second group. The average, across available currencies, of these skewness statistics is his global currency skewness factor. I construct a similar factor using my data set and refer to it as SKW.

B.5 Betas of Currency Factors

Table VI summarizes the results of running time series regressions of the monthly excess returns to S1, S2, S3, S4, S5, the EW carry trade portfolio and the HML carry trade portfolio on three pairs of risk factors: DOL and the HML carry trade factor, DOL and VOL, and DOL and SKW. Consistent with the literature, the factors are entered in pairs in the time series regressions.

The DOL and HML Carry factors are highly correlated with the S1–S5 portfolio returns. The betas on the DOL factor are all close to 1 in value, and statistically significant. The betas of the HML Carry factor run from -0.48 for S1 and increase across the portfolios to a high of 0.52 for S5, although the betas for S2, S3 and S4 are all close to zero. The R^2 for the five regressions are all large as well. DOL and HML Carry also have positive and significant betas for the EW carry trade portfolio. Of course, the HML carry portfolio has a beta of 1 with respect to itself.

Should we be surprised by these findings? Not really. Recall, from Table VI, that sorting the portfolios on the basis of the forward discount produced a monotonic ordering of the expected returns. In this circumstance, the DOL and HML carry factors will, by construction, create a pattern in the betas similar to what we see in Table VI.[5] The reason is simple: DOL is the simple average of the returns to S1 through S5. And HML is the difference between the returns to S5 and S1. If the returns to S1 through S5 were mutually uncorrelated and had common variance the construction of the factors would then imply a beta of 1 for the DOL factor, and betas of -0.5, 0, 0, 0, and 0.5 for the HML factor. The observed pattern in the betas is very similar to this, with the difference being that the returns to S1 through S5 are not mutually uncorrelated and do not have exactly the same variance. In fact, S1 through S5 have an interesting factor structure. As Lustig, Roussanov and Verdelhan (2011) point out, the covariance structure of S1 through S5 implies that two important factors drive most of the time series variation in the five portfolio returns. Let the covariance matrix of the returns to S1–S5 be Σ_z, and diagonalize Σ_z as $\Sigma_z = P\Lambda P^{-1}$, where P is the matrix whose columns are the orthonormal eigenvectors of Σ_z and Λ has the corresponding eigenvalues of Σ_z on the diagonal and zeros elsewhere. Lustig, Roussanov and Verdelhan's point is that there are two large eigenvalues, with the others being much smaller. The eigenvectors corresponding to the two largest eigenvalues are close to being proportional to (1, 1, 1, 1, 1), and (-0.5, 0, 0, 0, 0.5). Therefore, if linear combinations of the returns are formed using these vectors, the two resulting "factors" are highly correlated with, respectively, DOL and HML.

Turning, now to the DOL and VOL factors, we see that using VOL as a factor, rather than HML Carry, has very little impact on the betas with respect to DOL. The betas with respect to VOL decrease monotonically as we go from S1 to S5 and are statistically significant for the extreme portfolios, being positive for S1 and negative for S5. The betas with respect to VOL are also negative and statistically significant for the EW Carry and HML Carry portfolios. These findings indicate that when global currency volatility increases, the returns to holding low interest rate currencies increase and the returns to holding high interest rate currencies decrease. That is, low interest rate currencies provide a hedge against increased volatility. The average value of the volatility factor in the sample is 0.6%, indicating that on an annualized basis volatility averages about $0.6 \times \sqrt{365} = 11.5\%$. The standard deviation of the volatility factor in the sample is 0.2%. The magnitude of the betas for S1 and S5 implies that if volatility went up by one standard deviation, the annualized return to S1 would be 5.5% higher than normal, while the annualized return to S5 would be 3.1% lower than normal.

[5] Burnside (2010b) goes through the details of this "by construction" result.

Similarly, when we consider the DOL and SKW factors, we see that using SKW as a factor has very little impact on the betas with respect to DOL. The betas with respect to SKW increase (almost) monotonically as we go from S1 to S5 and are statistically significant for all portfolios, except S3. The betas are negative for S1 and S2, and positive for S3 through S5, as well as for the EW Carry and HML Carry portfolios. These findings indicate that during episodes in which intramonth global currency skewness becomes more negative, the returns to holding low interest rate currencies increase and the returns to holding high interest rate currencies decrease. That is, low interest rate currencies provide a hedge against currency crashes, and high interest rate currencies are the most exposed to them. The average value of the skewness factor in the sample is -0.07, with a standard deviation of 0.26. The magnitude of the betas for S1 and S5 implies that if skewness became more negative by one standard deviation, the annualized return to S1 would be 4.1% higher than normal, while the annualized return to S5 would be 4.0% lower than normal.

B.6 Cross-Sectional Analysis of Currency Factors

Table VII presents estimates of the SDF for the three currency factor models, obtained using the first stage GMM estimator. This estimator is equivalent to running the cross-sectional regression, (21). The results in Table VII use only portfolios S1 through S5 in the cross-sectional analysis. Not surprisingly, given our discussion of the betas, each of the models appears to do quite well in explaining the cross-section of returns. To see why, recall that in each case the betas with respect to the DOL factor were similar across portfolios and close to 1. This means that the betas of the DOL factor act like a constant in the cross-sectional regression. The b and λ estimates associated with the DOL factor end up being statistically insignificant, in each case, because the betas associated with the other factor are centered near zero, and the cross-sectional average of the mean returns of S1–S5 is also, statistically, near zero. The b parameters and risk premia (λ) associated with the HML Carry, VOL and SKW factors are statistically significant at the 5% level (except the b associated with VOL, which is significant at the 10% level). Again, this is not too surprising, because we saw, in Table VI, that the betas with respect to these factors tend to increase or decrease monotonically across the five portfolios. Finally, for all models the cross-sectional R^2 statistic is large. The model that uses DOL and HML Carry ends up being rejected on the basis of the pricing errors at the 5% level. This is because the model does a relatively poor job of explaining the returns on the non-extreme and non-central portfolios. This is a typical, by-construction, result for factor models based on one-dimensional sorts, where the first factor is the average of the portfolio returns, and the second factor is the difference between the extremes (Burnside (2010b)).

Table VIII presents estimates of the SDF for the three currency factor models, obtained using the iterated GMM estimator. Qualitatively and quantitatively, most of the results in Table VII are robust to using the iterated estimator. Rafferty's (2010) skewness based model has the best fit, and has no individually significant pricing errors, as before. The R^2 associated with Menkhoff et al.'s (2010) model falls considerably but the model is still not rejected on the basis of the J-statistic despite there being three individually significant pricing errors. The poorer fit in terms of R^2 can be understood by the fact that the iterated GMM estimator attempts to shrink the pricing error associated with S2 but at the cost of increasing the pricing errors associated with S1 and S4. The GMM estimator tries to do this because these pricing errors are correlated with each other. Lustig, Roussanov and Verdelhan's HML Carry based model, despite having a high R^2, is statistically rejected, as before. With the iterated GMM estimator the poor fit of the non-central and non-extreme portfolios, S2 and S4 is highlighted by their statistical significance.

B.7 Discussion

Models with factors based on currency returns seem to do quite well at explaining the returns to sorted portfolios of currencies. In the case of the DOL-HML Carry model we have seen that this is not surprising, given that the sorting works (in that it produces an ordering of average returns), and given that both factors are constructed from the portfolios being priced. Indeed, it is somewhat dissatisfying to explain currency returns with the HML Carry portfolio when, previously, we were trying to explain the returns of the HML Carry portfolio. For the other currency based models there is no similar issue, but, at the same time, the factors used in these cases are only indirectly linked to theoretical models. In a sense, therefore, we are left with the unsettling question: Why do these models seem to work in pricing the cross-section of S1–S5?

If the estimated currency based models are really informative about the SDFs of investors, then these models should also price stock returns. To see whether they do I re-estimate the models of the previous section using the five sorted portfolios, as well as Fama and French's 25 portfolios sorted on the basis of size and value. When the stock market portfolios are added to the cross-section, the models fare quite poorly. As Table IX shows, the estimated parameters of the models do not change dramatically. However, the fit of the models and the results of the tests of the overidentifying restrictions suggest that the models simply cannot explain the cross-section of stock returns.

This finding suggests that the cross-sectional analysis does not identify investors' SDF, or, to put it differently, a simple risk-based story based on the moment condition (17) has not yet been identified. However, the currency based models are informative. They tell us that carry trade strategies (like EW Carry and HML Carry) do better when currencies are less volatile and daily returns are less skewed. They do worse when volatility increases, and skewness becomes more negative. This finding is informative about the forces that drive currency fluctuations even if a satisfactory SDF has yet to be identified.

There are two less pessimistic interpretations of my findings. The first, is that there is some degree of market segmentation, which would make one SDF applicable to stock returns, and another relevant to currency returns. Although this is a logical possibility, there is empirical evidence against it. The nontraditional risk factors, described above, that do reasonably well at explaining currency portfolios sorted on the basis of forward discount, ought to explain other currency portfolios, such as those based on momentum, if they are reflective of currency investors' SDF. Burnside, Eichenbaum and Rebelo (2011), and Sarno et al. (2011a) present evidence, however, that the same nontraditional factors that price carry trade portfolios are unable to price momentum portfolios defined using short-term historical returns (i.e., when momentum is defined in terms of the previous month's return). Sarno et al. (2011a) report more mixed evidence for momentum defined using longer term historical performance. Sarno et al. (2011b) find that individual currency characteristics appear to be important in explaining the returns to currency momentum.

The second interpretation of my findings is that empirical exercises involving currency returns identify one component of the global investor's SDF, while those involving stock returns identify another. According to this interpretation, in effect, the fully successful SDF is the sum of the individual SDFs identified by sorting stocks, currencies, and other assets, on the basis of each asset return anomaly. Again, this is a logically coherent explanation, but one I find rather unsatisfying in that it makes SDF-based explanations untestable.

V. Time Varying Risk and Rare Events

I concluded, in Section IV.A, that standard risk models do not explain the returns to the carry trade. There we saw that the beta of the HML carry trade portfolio with respect to the CAPM factor is statistically significant, but is much too small (0.163) to explain the risk premium of the carry trade. To explain the roughly 6% risk premium of the HML Carry portfolio, the beta would need to be about six times as large. Lustig, Roussanov and Verdelhan (2011) agree on this assessment, arguing that "the average beta of HML_{FX} with the US stock market return is too small to explain carry trade risk premia." However, they argue that the beta of the carry trade with respect to the stock market increases during times of financial market distress. Certainly, during the recent mortgage crisis my HML carry trade factor displayed more correlation than usual with the stock market. However, it seems unlikely that a simple conditional beta story can explain the returns to the carry trade. There are two reasons favoring this conclusion. First, as Figure 1 reveals, in historical data, there is no systematic relationship between distress in the stock market (measured by periods of sharp decline) and currency crashes (measured by period of big losses to the carry trade). Second, time variation in the carry trade's stock market beta, while significant, is quantitatively not large enough. To see this, consider Figure 2, which plots betas of the *daily* returns of the HML carry trade portfolio with respect to the market premium. The betas are computed with a 130 working day (6-month) backward-looking rolling window, but similar results emerge with different windows. Overall the betas, even at the extremes, are not that large. The figure also shows all dates at which the monthly return to the stock market was -10% or less. If these months initiated periods of stock market distress we might expect to see the rolling window betas increase in the 6-month windows inclusive of these events. In some cases, as in the recent crisis, this is what we observe. In other cases, such as the stock market crash at the end of the dot-com boom in 2000–01, there is no such increase in the beta.

Lustig, Roussanov and Verdelhan (2011) briefly explore a potentially related explanation of currency returns. They consider a two factor model in which one of the factors is DOL and the other factor is a measure of the change in global stock market volatility (the cross-country average of daily intramonth stock market volatility, measured using local currency MSCI indices). They find that the betas of their sorted currency portfolios with respect to stock market volatility decrease with the size of the interest differential. While the second factor is driven out by their HML carry factor in cross-sectional regressions, the relationship between stock market volatility and currency returns may shed light on a common economic explanation of the returns to the carry trade, especially because stock market volatility has been shown to have some power to explain the cross-section of stock returns (Ang et al., 2006).

An alternative explanation of the returns to the carry trade is that they reflect out of sample (or peso event) risk.[6] This explanation relies on the notion that (17) still holds, but that the observed historical data are not fully representative of the underlying population distribution of the payoffs and the SDF. Burnside et al. (2011) use currency options data to construct hedged and unhedged versions of the EW carry trade portfolio. By doing so, and by imposing that (17) holds after allowing for peso event risk, they are able to characterize the nature of the hidden peso events. They argue that peso events appear to be ones in which carry trades incur relatively modest losses. The defining characteristic of a peso event, instead, is the fact that the SDF increases sharply, indicating that investors fear disastrous outcomes.

What remains to be seen is whether the peso event based explanation of the returns to the carry trade can be connected to the results discussed above. Can peso risk induce a time-varying risk premium that explains the UIP puzzle? Can peso risk explain the correlation between volatility and skewness factors and carry trade payoffs? These are open questions for future research. A theme of this paper is that explanations of asset pricing puzzles that work across markets are not easily identifiable. So a peso risk story that works for currencies should also work for stock returns. Burnside et al. (2011) suggest that the same peso event that can rationalize carry trade returns can also rationalize the return on the overall stock market. Burnside, Eichenbaum and Rebelo (2011) suggest that the same peso event that can rationalize carry trade returns can also rationalize currency momentum returns. Julliard and Ghosh (2010), however, suggest that rare consumption disasters make the cross-section of stock returns harder to rationalize, because they reduce the spread of consumption betas in the cross-section. Also, if we observe sufficiently many extreme events in markets (e.g., the 2008 financial crisis, the 2011 European debt crises and the 2011 downgrade of U.S. Treasury debt by Standard and Poor's) the distinction between theories based on observed

[6] Krasker (1980), Lewis (1989) and Kaminsky (1993) explored the role of peso problems in explaining the behavior of foreign exchange markets. More recently Burnside et al. (2010) ask whether out of sample events can explain the returns to the carry trade. Farhi and Gabaix (2008) and Farhi et al. (2009) explore both in sample and out of sample rare events.

risk factors, especially those related to measures of volatility and skewness, and unobserved peso events may become less clear. In other words, the in-sample frequency of extreme events may end up being similar to their true frequency.

VI. Conclusion

Carry trades are, on average, profitable. As we have seen, conventional, stock market based, models of risk do not explain the returns to the carry trade. Less traditional factors, that are defined in terms of the currency fluctuations, are more successful in explaining currency returns, but do not, conversely, explain the returns to the stock market. This means that, at least for the moment, a unifying explanation of stock market and carry trade returns based on observed fluctuations in measures of risk remains elusive. An alternative explanation is that carry trade returns reflect investors' concerns about out of sample events. While this story has some appeal, it must, of course, grapple with the evidence that volatility and skewness (or crash risk) factors have explanatory power in sample.

REFERENCES

Ang, Andrew, Robert J. Hodrick, Yuhang Xing and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.

Bekaert, Geert and Robert J. Hodrick, 1992, Characterizing predictable components in excess returns on equity and foreign exchange markets, *Journal of Finance* 47, 467–509.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2009, Carry trades and currency crashes, *NBER Macroeconomics Annual* 23, 313–347.

Burnside, Craig, 2007, The cross-section of foreign currency risk premia and consumption growth risk: A comment, NBER Working paper 13129.

Burnside, Craig, 2010a, Identification and inference in linear stochastic discount factor models, NBER Working paper 16634.

Burnside, Craig, 2010b, Sorts and cross-sectional regressions, mimeo, Duke University.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2006, The returns to currency speculation, NBER Working paper 12489.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24, 853–891.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2011, Carry trade and momentum in currency markets, forthcoming, *Annual Review of Financial Economics*.

Campbell, John Y., 2003, Consumption-based asset Pricing, in George M. Constantinides, Milton Harris, and René Stulz, ed.: *Handbook of the Economics of Finance*, Vol. 1B (Elsevier).

Da, Zhi, and Ernst Schaumburg, 2009, The pricing of volatility risk across asset classes and the Fama-French factors, mimeo, University of Notre Dame.

Engel, Charles, 1996, The forward discount anomaly and the risk premium: A survey of recent evidence, *Journal of Empirical Finance* 3, 123–192.

Fama, Eugene F., 1984, Forward and spot exchange rates, *Journal of Monetary Economics* 14, 319–338.

Fama, Eugene and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Farhi, Emmanuel, Samuel P. Fraiberger, Xavier Gabaix, Romain Ranciere and Adrien Verdelhan, 2009, Crash risk in currency markets, NBER Working paper 15062.

Farhi, Emmanuel and Xavier Gabaix, 2008, Rare disasters and exchange rates, NBER Working paper 13805.

Hansen, Lars Peter, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.

Hodrick, Robert J., 1987, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets* (Harwood Academic Publishers, Chur, Switzerland).

Jegadeesh, Narasimhan and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Julliard, Christian, and Anisha Ghosh, 2010, Can rare events explain the equity premium puzzle? mimeo, Carnegie Mellon University.

Kaminsky, Graciela, 1993, Is there a peso problem? Evidence from the dollar/pound exchange rate, 1976-1987, *American Economic Review* 83, 450–472.

Krasker, William S., 1980, The 'peso problem' in testing the efficiency of forward exchange markets, *Journal of Monetary Economics* 6, 269–276.

Lewis, Karen K., 1989, Changing beliefs and systematic rational forecast errors with evidence from foreign exchange, *American Economic Review* 79, 621–636.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, SSRN paper 1139447. Forthcoming, *Review of Financial Studies.*

Lustig, Hanno and Adrien Verdelhan, 2007, The cross-section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Meese, Richard A. and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies : Do they fit out of sample? *Journal of International Economics* 14, 3-24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2010, Carry trades and global foreign exchange volatility, SSRN paper 1342968.

Rafferty, Barry, 2010, The returns to currency speculation and global currency realignment risk, manuscript, Duke University.

Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539–580.

TABLE I: **Annualized Excess Returns of Investment Strategies**

	Mean	Standard Deviation	Sharpe Ratio	Skewness	Excess Kurtosis
Average of individual-currency carry trades	0.0460	0.113	0.418	-0.233	1.61
Equally-weighted carry trade	0.0458	0.051	0.903	-0.531	4.23
	(0.0095)	(0.004)	(0.212)	(0.402)	(1.52)
HML carry trade	0.0597	0.095	0.626	-0.516	1.63
	(0.0164)	(0.006)	(0.192)	(0.206)	(0.77)
U.S. stock market	0.0634	0.157	0.403	-0.782	2.35
	(0.0277)	(0.010)	(0.189)	(0.278)	(1.14)

February 1976 to October 2010. Statistics are reported for annualized excess returns. The U.S. stock market return is the value-weighted excess return on all U.S. stocks reported in Kenneth French's database. The equally weighted carry trade portfolio is formed as the average of up to 20 individual currency carry trades against the U.S. dollar. The individual currencies are indicated in the text. The HML carry trade strategy is a portfolio that takes an equally weighted long position in the quintile of currencies with the largest forward discounts and an equally weighted short position in the quintile of currencies with the smallest forward discounts against the U.S. dollar. Heteroskedasticity consistent GMM standard errors are in parentheses. The mean excess return of the equally-weighted carry trade is not equal to the average mean excess return of the individual-currency carry trades because the sample periods for which the currencies are available varies.

TABLE II: **Monthly Factor Betas of the Carry Trade Portfolios**

	EW Carry Trade					HML Carry Trade				
Factors	Intercept	Beta(s)			R^2	Intercept	Beta(s)			R^2
CAPM	0.004	0.028			0.008	0.004	0.163*			0.072
	(0.001)	(0.017)				(0.001)	(0.038)			
Fama-French factors	0.004	0.042*	-0.034	0.037	0.020	0.003	0.185*	0.080	0.156*	0.100
(Mkt-Rf, SMB, HML)	(0.001)	(0.018)	(0.029)	(0.029)		(0.001)	(0.042)	(0.047)	(0.055)	
Industrial production	0.004	0.118			0.003	0.005	0.171			0.002
	(0.001)	(0.146)				(0.001)	(0.231)			
CAPM-IP	0.003	0.029	0.129		0.011	0.004	0.165*	0.232		0.075
(Mkt-Rf, I.P. growth)	(0.001)	(0.017)	(0.141)			(0.001)	(0.038)	(0.198)		
Realized stock volatility	0.005	-0.002			0.005	0.014	-0.010*			0.038
	(0.002)	(0.002)				(0.003)	(0.004)			
CAPM-Stock vol	0.005	0.023	-0.001		0.009	0.009	0.138*	-0.006		0.084
(Mkt-Rf, Stock vol)	(0.002)	(0.020)	(0.002)			(0.003)	(0.036)	(0.004)		

February 1976 to October 2010. The table reports estimates of the equation $z_t = a + f_t'\beta + \epsilon_{t+1}$, where z_t is the monthly excess return of a carry-trade portfolio and f_t is a scalar or vector of risk factors. The CAPM factor is the excess return on the value-weighted US stock market (Mkt-Rf), and the Fama-French factors are the Mkt-Rf, SMB and HML factors (available from Kenneth French's database). The industrial production factor is monthly industrial production growth. The stock volatility factor is realized daily volatility measured monthly. Heteroskedasticity consistent standard errors are in parentheses. Slope coefficients that are statistically significant at the 5 percent level are indicated by an asterisk (*).

TABLE III: **GMM Estimates of Monthly Linear Factor Models**

	b	λ (%)	R^2	J	Pricing Errors (α) EW Carry	Pricing Errors (α) HML Carry
CAPM	2.18	0.45*	-1.86	104	0.044*	0.051*
	(1.17)	(0.22)		(0.00)	(0.009)	(0.016)
Fama-French Factors						
Mkt-Rf	3.36*	0.49*	0.38	91.0	0.043*	0.040*
	(1.41)	(0.23)		(0.00)	(0.009)	(0.017)
SMB	3.15	0.24				
	(1.74)	(0.15)				
HML	6.52*	0.36*				
	(1.91)	(0.17)				
Industrial production	-5.40	-0.03	-9.59	108	0.046*	0.060*
	(28.1)	(0.14)		(0.00)	(0.009)	(0.016)
CAPM-IP						
Mkt-Rf	2.28	0.46*	-1.76	103	0.044*	0.050*
	(1.21)	(0.22)		(0.00)	(0.009)	(0.016)
I.P. growth	4.64	0.02				
	(29.6)	(0.14)				
Stock volatility	-0.30	-9.18	-3.49	85.77	0.044*	0.049*
	(0.25)	(6.40)		(0.00)	(0.009)	(0.017)
CAPM-Stock vol						
Mkt-Rf	-1.36	0.08	-6.70	80.5	0.044*	0.051*
	(1.70)	(0.23)		(0.00)	(0.010)	(0.018)
Stock volatility	-0.41	-11.1				
	(0.33)	(7.99)				

February 1976 to October 2010. Test assets are the Fama-French 25 portfolios, and the EW and HML carry trade portfolios. The CAPM factor is the excess return on the value-weighted US stock market (Mkt-Rf), and the Fama-French factors are the Mkt-Rf, SMB and HML factors (available from Kenneth French's database). The industrial production factor is monthly industrial production growth. The stock volatility factor is realized daily volatility measured monthly. The table reports iterated GMM estimates of the SDF parameter, b, and the factor risk premia, λ, reported in monthly percent. The R^2 is a measure of fit between the mean excess returns and the predicted mean returns. Test statistics, J, for the overidentifying restrictions are also reported. The annualized pricing errors of the carry-trade portfolios (α) are reported. Heteroskedasticity consistent standard errors are in parentheses, except for the J statistics, for which the p-value is in parentheses. An asterisk (*) indicates statistical significance at the 5 percent level.

TABLE IV: **Quarterly Factor Betas of the Carry Trade Portfolios**

	EW Carry Trade					HML Carry Trade				
Factors	Intercept	Beta(s)			R^2	Intercept	Beta(s)			R^2
C-CAPM (Δc)	0.012	-0.021			0.000	0.016	0.018			0.000
	(0.005)	(0.744)				(0.008)	(1.234)			
Durables (Δd)	0.007	0.524			0.009	0.018	-0.196			0.000
	(0.007)	(0.554)				(0.010)	(0.878)			
Market return (Mkt)	0.012	0.009			0.001	0.013	0.145*			0.061
	(0.002)	(0.031)				(0.011)	(0.068)			
Extended C-CAPM	0.007	-0.325	0.641	0.014	0.012	0.013	-0.437	0.130	0.148*	0.062
(Δc, Δc, Mkt)	(0.007)	(0.828)	(0.580)	(0.032)		(0.001)	(1.184)	(0.852)	(0.069)	
C-CAPM (Δc)	0.011	0.119			0.000	0.017	-0.384			0.001
(Campbell timing)	(0.004)	(0.605)				(0.007)	(1.170)			
Durables (Δd)	0.007	0.495			0.008	0.016	-0.044			0.000
(Campbell timing)	(0.007)	(0.576)				(0.011)	(1.010)			
Extended C-CAPM	0.007	-0.203	0.568	0.014	0.010	0.013	-1.646	0.604	0.172*	0.077
(Campbell timing)	(0.007)	(0.672)	(0.618)	(0.030)		(0.011)	(1.064)	(1.021)	(0.072)	

1976Q2–2009Q4. The table reports estimates of the equation $z_t^{qr} = a + f_t'\beta + \epsilon_{t+1}$, where z_t^{qr} is the quarterly real excess return of a carry-trade portfolio and f_t is a scalar or vector of risk factors. The C-CAPM factor is the log growth rate of real consumption of nondurables and services, the durables factor is the log growth rate of the service flow of durables assumed to be proportional to the real stock of consumer durables, the market return factor (Mkt) is from Kenneth French's database. "Campbell timing" is explained in the main text. Heteroskedasticity consistent standard errors are in parentheses. Slope coefficients that are statistically significant at the 5 percent level are indicated by an asterisk (*).

TABLE V: **GMM Estimates of Quarterly Linear Factor Models**

	b	λ (%)	R^2	J	Pricing Errors (α) EW Carry	HML Carry
C-CAPM	68.9	0.12	-7.25	66.2	0.047*	0.063*
	(50.2)	(0.09)		(0.00)	(0.012)	(0.021)
Extended C-CAPM						
Δc	3.30	0.00	-1.25	65.6	0.046*	0.054*
	(53.3)	(0.09)		(0.00)	(0.010)	(0.018)
Δd	5.41	0.01				
	(39.0)	(0.09)				
Mkt	2.26	1.57*				
	(1.21)	(0.73)				
Models with Campbell's timing for consumption						
C-CAPM	31.1	0.06	-7.27	68.5	0.047*	0.064*
	(41.4)	(0.07)		(0.00)	(0.011)	(0.019)
Extended C-CAPM						
Δc	-17.1	0.01	-4.20	43.4	0.045*	0.057*
	(51.9)	(0.08)		(0.01)	(0.011)	(0.018)
Δd	38.0	0.08				
	(42.9)	(0.11)				
Mkt	1.44	0.74				
	(1.33)	(0.74)				

1976Q2 to 2009Q4. Test assets are the Fama-French 25 portfolios, and the EW and HML carry trade portfolios. The C-CAPM factor is the log growth rate of real consumption of nondurables and services, the durables factor is the log growth rate of the service flow of durables assumed to be proportional to the real stock of consumer durables, the market return factor (Mkt) is from Kenneth French's database. The table reports iterated GMM estimates of the SDF parameter, b, and the factor risk premia, λ, reported in quarterly percent. The R^2 is a measure of fit between the mean excess returns and the predicted mean returns. Test statistics, J, for the overidentifying restrictions are also reported. The annualized pricing errors of the carry-trade portfolios (α) are reported. Heteroskedasticity consistent standard errors are in parentheses, except for the J statistics, for which the p-value is in parentheses. An asterisk (*) indicates statistical significance at the 5 percent level.

TABLE VI: **Factor Betas of the Sorted Currency Portfolios and the Carry Trade Portfolios**

		LRV (2010)			Menkhoff et. al. (2009)			Rafferty (2010)		
Portfolio	Mean Return	DOL	HML Carry	R^2	DOL	VOL	R^2	DOL	SKW	R^2
S1	-0.80%	0.959*	-0.478*	0.923	0.942*	0.0216*	0.736	0.948*	-0.0129*	0.723
	(1.86)	(0.020)	(0.020)		(0.039)	(0.0056)		(0.043)	(0.0030)	
S2	-0.28%	0.901*	-0.116*	0.857	0.893*	-0.0000	0.841	0.904*	-0.0066*	0.846
	(1.69)	(0.028)	(0.021)		(0.030)	(0.0046)		(0.028)	(0.0021)	
S3	2.80%	0.917*	-0.001	0.825	0.915*	-0.0030	0.826	0.910*	0.0044	0.827
	(1.80)	(0.027)	(0.026)		(0.027)	(0.0033)		(0.027)	(0.0024)	
S4	3.72%	0.959*	0.068*	0.868	0.961*	-0.0044	0.864	0.957*	0.0041*	0.865
	(1.76)	(0.024)	(0.022)		(0.025)	(0.0027)		(0.025)	(0.0021)	
S5	5.17%	0.959*	0.522*	0.933	0.986*	-0.0120*	0.717	0.973*	0.0126*	0.722
	(2.13)	(0.020)	(0.020)		(0.034)	(0.0052)		(0.036)	(0.0031)	
EW Carry	4.58%	0.184*	0.264*	0.371	0.196*	-0.0080*	0.140	0.179*	0.0140*	0.187
	(0.95)	(0.040)	(0.029)		(0.045)	(0.0039)		(0.043)	(0.0029)	
HML Carry	5.97%	0	1*	1	0.044	-0.0336*	0.071	0.026	0.0255*	0.061
	(1.64)				(0.062)	(0.0094)		(0.069)	(0.0052)	

February 1976 to October 2010. The table reports estimates of the equation $z_t = a + f_t'\beta + \epsilon_{t+1}$, where z_t is the monthly excess return of each of the portfolios indicated and f_t is a 2×1 vector of the indicated risk factors. The DOL factor is the average excess return to portfolios S1–S5. The HML carry portfolio is the excess return to being long portfolio P5 and short portfolio P1. The VOL factor is a measure of realized global currency volatility. The SKW factor is a measure of realized global currency skewness. Heteroskedasticity consistent standard errors are in parentheses. Slope coefficients that are statistically significant at the 5 percent level are indicated by an asterisk (*).

TABLE VII: **First Stage GMM Estimates of Linear Factor Models for Sorted Currency Portfolios**

	b	λ (%)	R^2	J	Pricing Errors (α) S1	S2	S3	S4	S5
Lustig, Roussanov and Verdelhan (2010) factors									
DOL	2.39	0.190	0.84	9.23	0.0010	-0.0158*	0.0071	0.0109	-0.0040
	(2.05)	(0.127)		(0.026)	(0.0060)	(0.0060)	(0.0068)	(0.0064)	(0.0046)
HML Carry	6.97*	0.539*							
	(2.08)	(0.133)							
Menkhoff et. al. (2009) factors									
DOL	0.96	0.198	0.73	4.31	0.0060	-0.0240*	0.0012	0.0070	0.0081
	(3.16)	(0.122)		(0.23)	(0.0043)	(0.0121)	(0.0092)	(0.0081)	(0.0098)
VOL	-3.07	-14.0*							
	(1.69)	(5.67)							
Rafferty (2010) factors									
DOL	-2.11	0.182	0.96	1.51	0.0033	-0.0061	-0.0029	0.0061	-0.0008
	(2.76)	(0.131)		(0.68)	(0.0054)	(0.0066)	(0.0091)	(0.0074)	(0.0071)
SKW	3.05*	20.7*							
	(0.95)	(6.19)							

February 1976 to October 2010. Test assets are the sorted currency portfolios S1–S5. The DOL factor is the average excess return to portfolios S1–S5. The HML carry portfolio is the excess return to being long portfolio P5 and short portfolio P1. The VOL factor is a measure of realized global currency volatility. The SKW factor is a measure of realized global currency skewness. The table reports first stage GMM estimates of the SDF parameter, b, and the factor risk premia, λ, reported in monthly percent. The R^2 is a measure of fit between the mean excess returns and the predicted mean returns. Test statistics, J, for the overidentifying restrictions are also reported. The annualized pricing errors of the S1–S5 portfolios (α) are reported. Heteroskedasticity consistent standard errors are in parentheses, except for the J statistics, for which the p-value is in parentheses. An asterisk (*) indicates statistical significance at the 5 percent level.

TABLE VIII: **Iterated GMM Estimates of Linear Factor Models for Sorted Currency Portfolios**

	b	λ	R^2	J	Pricing Errors (α)				
		(%)			S1	S2	S3	S4	S5
Lustig, Roussanov and Verdelhan (2010) factors									
DOL	2.35	0.184	0.83	9.37	-0.0012	-0.0159*	0.0077	0.0120*	-0.0001
	(2.03)	(0.127)		(0.025)	(0.0051)	(0.0060)	(0.0079)	(0.0059)	(0.0043)
HML Carry	6.30*	0.488*							
	(2.03)	(0.133)							
Menkhoff et. al. (2009) factors									
DOL	-0.68	0.053	0.184	5.78	0.0127*	-0.0085*	0.0185	0.0257*	0.0306
	(2.64)	(0.135)		(0.12)	(0.0062)	(0.0038)	(0.0139)	(0.0110)	(0.0176)
VOL	-2.26	-10.3*							
	(1.34)	(4.84)							
Rafferty (2010) factors									
DOL	-1.66	0.221	0.93	1.47	-0.0001	-0.0098	-0.0075	0.0012	-0.0064
	(2.72)	(0.136)		(0.69)	(0.0052)	(0.0091)	(0.0114)	(0.0069)	(0.0087)
SKW	3.14*	21.4*							
	(0.94)	(6.11)							

February 1976 to October 2010. Test assets are the sorted currency portfolios S1–S5. The DOL factor is the average excess return to portfolios S1–S5. The HML carry portfolio is the excess return to being long portfolio P5 and short portfolio P1. The VOL factor is a measure of realized global currency volatility. The SKW factor is a measure of realized global currency skewness. The table reports iterated GMM estimates of the SDF parameter, b, and the factor risk premia, λ, reported in monthly percent. The R^2 is a measure of fit between the mean excess returns and the predicted mean returns. Test statistics, J, for the overidentifying restrictions are also reported. The annualized pricing errors of the S1–S5 portfolios (α) are reported. Heteroskedasticity consistent standard errors are in parentheses, except for the J statistics, for which the p-value is in parentheses. An asterisk (*) indicates statistical significance at the 5 percent level.

TABLE IX: **GMM Estimates of Linear Factor Models for Sorted Currency and Stock Market Portfolios**

	b	λ (%)	R^2	J
Lustig, Roussanov and Verdelhan (2010) factors				
DOL	1.82	0.145	-1.92	97.0
	(1.97)	(0.128)		(0.000)
HML Carry	5.26*	0.407*		
	(1.94)	(0.134)		
Menkhoff et. al. (2009) factors				
DOL	0.30	0.053	-2.70	89.3
	(2.14)	(0.135)		(0.000)
VOL	-0.84	-3.86		
	(0.73)	(2.96)		
Rafferty (2010) factors				
DOL	-2.46	0.154	-2.42	56.8
	(2.44)	(0.138)		(0.001)
SKW	3.02*	20.4*		
	(0.86)	(4.26)		

February 1976 to October 2010. Test assets are the sorted currency portfolios S1–S5 and the Fama-French 25 portfolios sorted on size and value. The DOL factor is the average excess return to portfolios S1–S5. The HML carry portfolio is the excess return to being long portfolio P5 and short portfolio P1. The VOL factor is a measure of realized global currency volatility. The SKW factor is a measure of realized global currency skewness. The table reports iterated GMM estimates of the SDF parameter, b, and the factor risk premia, λ, reported in monthly percent. The R^2 is a measure of fit between the mean excess returns and the predicted mean returns. Test statistics, J, for the overidentifying restrictions are also reported. Heteroskedasticity consistent standard errors are in parentheses, except for the J statistics, for which the p-value is in parentheses. An asterisk (*) indicates statistical significance at the 5 percent level.

FIGURE 1

Annual Realized Excess Returns of the Carry-Trade and U.S. Stock Market

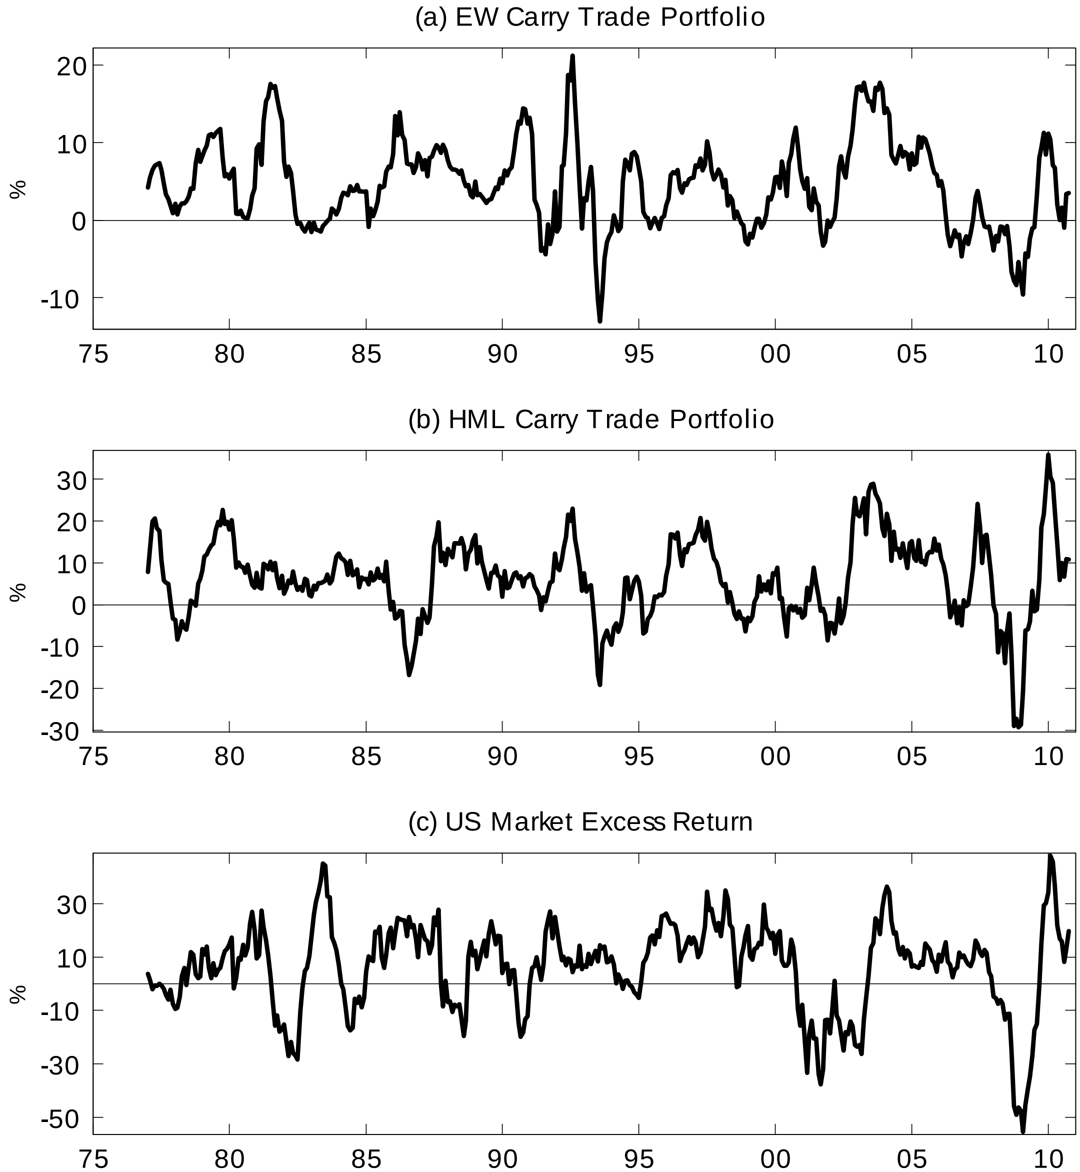

12-Month rolling window, February 1976–October 2010. The carry trade portfolios are described in detail the text. The EW carry trade is an equally weighted portfolio of carry trades in up to 20 currencies against the USD. The HML.carry trade portfolio is one in which the investor goes long in the highest interest rate currencies and short in the lowest interest rate currencies, defined in terms of sorted quintiles of up to 20 currencies. The US market excess return is the Mkt-Rf factor from Kenneth French's database.

FIGURE 2

The Time-Varying Market Beta of the HML Carry Trade Portfolio

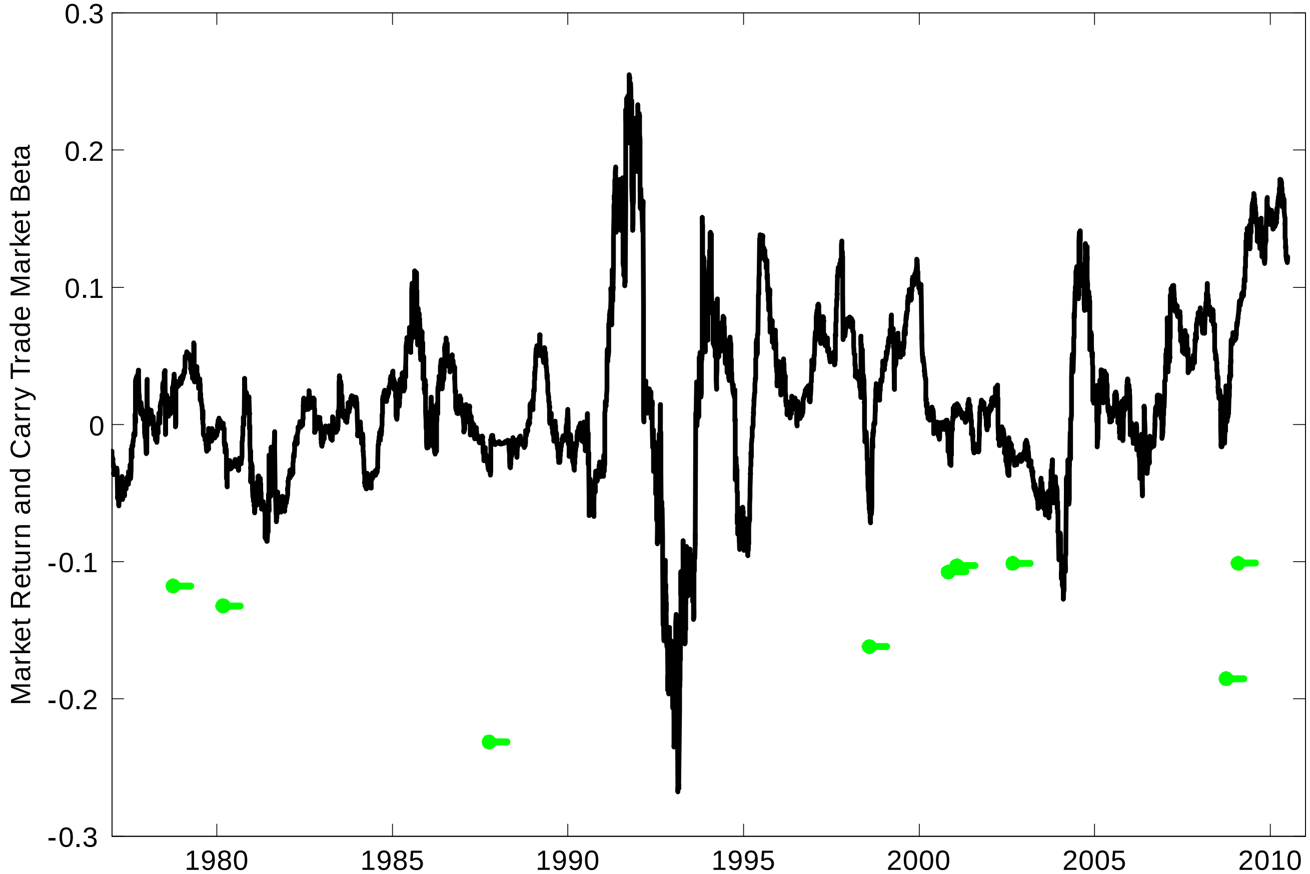

6-Month rolling window, January 1977–June 2010. The beta is computed by regressing daily returns to the HML carry trade portfolio on daily returns to the U.S. stock market, defined as the Mkt-Rf factor from Kenneth French's database. The green dots mark single months in which the market excess return was less than -10%. The green lines delineate a 6-month window after but inclusive of each of these dates.