

Do Peso Problems Explain the Returns to the Carry Trade?*

Craig Burnside,[†] Martin Eichenbaum,[‡]
Isaac Kleshchelski,[§] and Sergio Rebelo[¶]

September 2008

Abstract

We study the properties of the carry trade, a currency speculation strategy in which an investor borrows low-interest-rate currencies and lends high-interest-rate currencies. This strategy generates payoffs which are on average large and uncorrelated with traditional risk factors. We investigate whether these payoffs reflect a peso problem. We argue that, with one proviso, they do. The proviso is that the defining characteristic of a peso event is a high value of the stochastic discount factor, not high carry-trade losses. We reach this conclusion by analyzing the payoffs to the hedged carry trade, in which an investor uses currency options to protect himself from the downside risk from large, adverse movements in exchange rates. We find that the same value of the stochastic discount factor that rationalizes the average payoffs to the carry trade also rationalizes the equity premium.

J.E.L. Classification: F31

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

*This paper is a substantially revised version of NBER Working Paper 12489 titled “The Returns to Currency Speculation.” We thank Ravi Jagannathan for numerous discussions and suggestions and the Chicago Mercantile Exchange for providing us with the currency-options data used in this paper. Burnside is grateful to the National Science Foundation for financial support (SES-0516697).

[†]Duke University and NBER.

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

[§]Washington University in Saint Louis.

[¶]Northwestern University, NBER, and CEPR.

1 Introduction

Currencies that are at a forward premium tend to depreciate. This ‘forward-premium puzzle’ represents an egregious deviation from uncovered interest parity (UIP). This paper studies the properties of a widely-used currency speculation strategy that exploits this anomaly. The strategy, known as the carry trade, involves selling currencies forward that are at a forward premium and buying currencies forward that are at a forward discount. Transaction costs aside, this strategy is equivalent to borrowing low-interest-rate currencies in order to lend high-interest-rate currencies, without hedging the associated currency risk. Consistent with results in the literature, we find that the carry-trade strategy applied to portfolios of currencies yields high average payoffs, as well as Sharpe ratios that are substantially higher than those associated with the U.S. stock market.

The most natural interpretation for the high average payoffs to the carry trade is that they compensate agents for bearing risk. However, we show that linear stochastic discount factors built from conventional measures of risk, such as consumption growth, the returns to the stock market, and the Fama-French (1993) factors, fail to explain the payoffs to the carry trade. This failure reflects the absence of a statistically significant correlation between the payoffs to the carry trade and traditional risk factors. Our results are consistent with previous work documenting that one can reject consumption-based asset-pricing models using data on forward exchange rates.¹ More generally, it has been difficult to use asset-pricing models such as the CAPM to rationalize the risk-premium movements required to account for the time-series properties of the forward premium.²

The most natural alternative explanation for the high average payoffs to the carry trade is that they reflect the presence of a peso problem. We use the term “peso problem” as defined by Cochrane (2001), i.e. “a generic term for the effects of small probabilities of large events on empirical work.” This definition of a peso problem is consistent with agents being risk averse and is equivalent to the “rare event” problem that has received substantial attention in the literature. In what follows we use the term “peso event” to refer to a rare event in which there are either large negative payoffs to the carry trade or unusually high values of the stochastic discount factor (SDF).

A number of authors have recently argued that the peso problem lies at the root of

¹See, for example, Bekaert and Hodrick (1992) and Backus, Foresi, and Telmer (2001).

²See, for example, Bekaert (1996) and De Santis and Gérard (1999).

the failure of UIP.³ Not surprisingly, peso problems can in principle also explain the positive average payoffs to the carry trade. To understand the basic argument, suppose that a foreign currency is at a forward premium, so that a carry-trade investor sells this currency forward. Assume that a substantial appreciation of the foreign currency occurs with small probability. The investor must be compensated for the negative payoff to the carry trade in this state of the world. The degree of compensation depends on the value of the SDF in the peso state and the magnitude of the negative payoff. Conceptually it is useful to distinguish between two extreme possibilities. The first possibility is that the salient feature of a peso state is large carry trade losses. The second possibility is that the salient feature of a peso state is a large value of the SDF. A key contribution of this paper is to assess the relative importance of these two possibilities. We do so by estimating the size of carry-trade losses and the level of the SDF in the peso state.

Our basic approach relies on analyzing the payoffs to a version of the carry-trade strategy that does not yield high negative payoffs in a peso state. This strategy works as follows. When an investor sells the foreign currency forward he simultaneously buys a call option on that currency. If the foreign currency appreciates beyond the strike price, the investor can buy the foreign currency at the strike price and deliver the currency in fulfillment of the forward contract. Similarly, when an investor buys the foreign currency forward, he can hedge the downside risk by buying a put option on the foreign currency. By construction, this “hedged carry trade” does not generate large negative payoffs in the peso state. To estimate the average payoff to the hedged carry trade we use data on currency options with a one-month maturity. At this stage of the analysis we wish to be eclectic about the size of the negative payoff in the peso state. So, our hedging strategy uses at-the-money options which pay off in all peso states, as well as in some non-peso states.

Using information on the average payoffs to the hedged and unhedged carry trade, we obtain estimates of the payoff to the unhedged carry trade and the SDF in the peso event. Our main findings can be summarized as follows. First, the average payoff to the hedged and unhedged carry trade are very similar (2.5 and 3.2 percent per annum, respectively). Second, the standard deviation of the payoffs to the hedged carry trade is actually substantially lower than those of the unhedged carry trade (3.6 percent versus 6.0 percent). Third, the payoff

³See Farhi and Gabaix (2008). In related work, Brunnermeier, Nagel, and Pedersen (2008) emphasize the importance of crash-related risk. Authors such as Rietz (1988), Barro (2006), and Gabaix (2007), argue that peso problems can explain other asset-pricing anomalies such as the equity premium.

to the unhedged carry trade in the peso state is only moderately negative. In particular, our point estimate of this payoff is only one standard deviation below the sample average of the payoff to the unhedged carry trade. Fourth, the SDF is over one-hundred times larger in the peso state than in the non-peso state.

The intuition for why the losses to the unhedged carry trade are small in the peso state is as follows. Any gains to the carry trade in the non-peso state must on average be compensated, on a risk-adjusted basis, by losses in the peso state. Since the average risk-adjusted gains to the hedged and unhedged carry trade in the non-peso state are similar, the risk-adjusted losses of these two strategies in the peso state must also be similar. Given that the value of the SDF in the peso state is the same for both strategies, the actual losses of the two strategies in the peso state must be similar. The options that we use in the hedged carry trade are always in the money in the peso state. So we know how much an agent loses in the peso state if he is pursuing the hedged carry trade. Since these losses turn out to be small, the losses to the unhedged carry trade in the peso state must also be small.

The rationale for why the SDF is much larger in the peso state than in the non-peso state is as follows. We just argued that the unhedged carry trade makes relatively small losses in the peso state. At the same time, the average, risk-adjusted payoff to the unhedged carry trade in the non-peso state is large. The only way to rationalize these observations is for the SDF to be very high in the peso state. So, even though the losses in the peso state are moderate, the investor attaches great importance to those losses.

A possible shortcoming of our methodology is that we can always produce values of the SDF and the carry-trade payoff in the peso state that rationalize the observed average payoffs to the carry trade. The skeptical reader may conclude that we have documented an interesting puzzle without providing a credible resolution of that puzzle. So, it is of interest to bring additional data to bear on the plausibility of our estimates. To this end, we consider two versions of an equity strategy that involves borrowing one dollar at the Treasury-bill rate and investing it in the stock market. In the first version the agent does not hedge against adverse movements in the stock market. In the second version the agent buys at-the-money put options which exactly compensate him for a fall in the stock market. We find that, in sharp contrast to the carry trade, the hedged stock market strategy yields large, negative average payoffs. Using the average payoffs to the two stock market strategies we generate an independent estimate of the value of the SDF in the peso state. Remarkably, the same

estimate of the peso state SDF that rationalizes the average payoffs to the carry trade also rationalizes the equity premium.

Our paper is organized as follows. In section 2 we describe the carry-trade strategy and discuss our method for estimating carry trade losses and the value of the SDF in the peso state. We describe our data in Section 3. In Section 4 we study the covariance between the payoffs to the carry trade and traditional risk factors, using both time series and panel data. In Section 5 we study the properties of the hedged carry trade. Together, Sections 4 and 5 provide the inputs for quantifying the ability of peso events to account for the observed payoffs to the carry trade. In Section 6 we report our results and generalize the analysis to multiple peso states. Section 7 concludes.

2 Peso problems and the carry trade

The failure of uncovered interest parity motivates a variety of speculation strategies. In this paper we focus on the carry trade, the strategy most widely used by practitioners (see Galati and Melvin (2004)). In this section we describe a procedure for analyzing peso-event explanations for carry-trade payoffs.

The carry trade consists of borrowing a low-interest-rate currency and lending a high-interest-rate currency. Abstracting from transactions costs, the payoff to the carry trade, denominated in dollars, is:⁴

$$y_t \left[(1 + r_t^*) \frac{S_{t+1}}{S_t} - (1 + r_t) \right]. \quad (1)$$

The variable S_t denotes the spot exchange rate expressed as dollars per foreign currency unit (FCU). The variables r_t and r_t^* represent the domestic and foreign interest rate, respectively. The amount of dollars borrowed, y_t , is given by:

$$y_t = \begin{cases} +1 & \text{if } r_t < r_t^*, \\ -1 & \text{if } r_t^* \leq r_t. \end{cases} \quad (2)$$

We normalize the amount of dollars we bet on this strategy (the absolute value of y_t) to one.

If S_{t+1} is a martingale:

$$E_t S_{t+1} = S_t, \quad (3)$$

⁴We study the impact of transactions costs in Section 4.

the expected payoff to the carry trade is positive and equal to the difference between the higher and the lower interest rates:

$$y_t (r_t^* - r_t) > 0.$$

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 ($F_t < S_t$). The value of x_t , the number of FCUs sold forward, is given by:

$$x_t = \begin{cases} +1/F_t & \text{if } F_t \geq S_t, \\ -1/F_t & \text{if } F_t < S_t. \end{cases} \quad (4)$$

This value of x_t is equivalent to buying/selling one dollar forward. The dollar-denominated payoff to this strategy at $t + 1$, denoted z_{t+1} , is

$$z_{t+1} = x_t (F_t - S_{t+1}). \quad (5)$$

Covered-interest-rate parity implies that:

$$(1 + r_t) = \frac{1}{S_t} (1 + r_t^*) F_t. \quad (6)$$

When equation (6) holds, the strategy defined by (4) yields positive payoffs if and only if the strategy defined by (2) has positive payoffs. This result holds because the two payoffs are proportional to each other. In this sense the strategies are equivalent. We focus our analysis on strategy (4) because of data considerations.

The impact of peso problems In this subsection we discuss our strategy to analyze peso-event explanations of carry-trade payoffs. Since the carry trade is a zero net-investment strategy, the payoff, z_t , must satisfy:

$$E_t (M_{t+1} z_{t+1}) = 0. \quad (7)$$

Here M_{t+1} denotes the SDF that prices payoffs denominated in dollars and E_t denotes the time- t conditional expectations operator. Taking unconditional expectation of equation (7) we obtain:

$$E (z_{t+1}) = -\frac{\text{cov} (M_{t+1}, z_{t+1})}{E (M_{t+1})}. \quad (8)$$

In light of equation (8) a natural explanation for the positive average payoffs to the carry trade is that these payoffs compensate agents for the negative covariance between M and z .

In our empirical work (see Section 4) we document that the covariance between the payoffs to the carry trade and a host of traditional risk factors is not statistically different from zero.⁵ This finding implies that traditional risk-based explanations are not a plausible rationale for the positive average payoffs to the carry trade.

An alternative explanation relies on the existence of peso events. To pursue this explanation we partition Ω , the set of possible states, s_t , into two sets. The first set, Ω^N , consists of those values of s_t corresponding to non-peso events. The second set, Ω^P , consists of those values of s_t corresponding to a peso event. For simplicity, we assume that for all $s_t \in \Omega^P$, $z(s_t) = z' < 0$ and $M(s_t) = M'$.

We denote by $G(s_{t+1})$ the unconditional distribution of s_{t+1} in non-peso states. For future reference we define $G(s_{t+1}|s_t)$ as the conditional distribution of s_{t+1} given s_t , where both s_{t+1} and s_t are in Ω^N . To simplify, we assume that the conditional and unconditional probability of the peso state is p . The unconditional version of equation (7) is:

$$(1 - p) \int M(s_{t+1})z(s_{t+1})dG(s_{t+1}) + pM'z' = 0. \quad (9)$$

Motivated by our empirical results, we assume that there are no peso events in our sample and that the covariance between M and z is zero in non-peso states.⁶ Equation (9) can then be re-written as:

$$(1 - p)E^G [M(s_{t+1})] E^G [z(s_{t+1})] + pM'z' = 0. \quad (10)$$

Here $E^G(\cdot)$ denotes the expectation over non-peso states, e.g. $E^G [z(s_{t+1})] = \int z(s_{t+1})dG(s_{t+1})$.

While there is no covariance between M and z in non-peso states, the unconditional covariance between M and z , $\text{cov}(M, z)$, can still be negative if $M' > E^G [M(s_{t+1})]$ or $z' < E^G [z(s_{t+1})]$. Under these circumstances the unconditional mean return over peso and non peso states, $(1 - p)E^G [z(s_{t+1})] + pz'$, can be positive. So, the existence of risk associated with peso events can rationalize positive payoffs to the carry trade, even in population.⁷ This

⁵See Villanueva (2007) for additional evidence on this point. Lustig and Verdelhan (2007) argue that aggregate consumption growth risk explains the cross-sectional variation in the excess returns to going long on currency portfolios that are sorted by their interest rate differential with respect to the U.S. Burnside (2007) challenges their results based on two findings. First, the time-series covariance between the excess returns to the Lustig-Verdelhan portfolios and standard risk factors, including aggregate consumption growth, is not significantly different from zero. Second, imposing the constraint that a zero beta asset has a zero excess return leads to a substantial deterioration in the ability of their model to explain the cross-sectional variation in excess returns to the portfolios.

⁶This assumption is consistent with Farhi and Gabaix's (2008) assumption that the stochastic discount factor for returns denominated in the world currency is constant in non-disaster states.

⁷See Farhi and Gabaix (2008) for a related discussion.

result is not useful in our context because we assume that there are no peso events in our sample. Explanations of the carry trade payoffs based on in-sample peso events run into the obvious problem that there is no covariance between standard risk factors and those payoffs.

Since z' is negative, equation (10) implies that the average return over non-peso states, $E^G[z(s_{t+1})]$, is positive. This observation captures the conventional view that a peso problem can rationalize positive average payoffs to the carry trade. The question we focus on is: can the existence of peso events provide a plausible rationale for our estimate of the average payoff to the carry trade in non-peso states? To study this question we develop a version of the carry-trade strategy that does not yield high negative payoffs when a peso event occurs. We call this strategy the “hedged carry trade.” We now describe this strategy in detail.

The hedged carry trade We begin by defining the notation we use to describe options contracts. A call option gives an agent the right, but not the obligation, to buy foreign currency with dollars at a strike price of K_t dollars per FCU. We denote the dollar price of this option by C_t . The payoff of the call option in dollars, net of the option price, is:

$$z_{t+1}^C = \max(0, S_{t+1} - K_t) - C_t(1 + r_t).$$

A put option gives an agent the right, but not the obligation, to sell foreign currency at a strike price of K_t dollars per FCU. We denote the dollar price of this option by P_t . The payoff of the put in dollars, net of the option price is:

$$z_{t+1}^P = \max(0, K_t - S_{t+1}) - P_t(1 + r_t).$$

Suppose that an agent sells one FCU forward. Then, the worst case scenario in the standard carry trade arises when there is a large appreciation of the foreign currency. In this state of the world the agent realizes large losses because he has to buy foreign currency at a high value of S_{t+1} to deliver on the forward contract. Suppose that an agent buys at time t a call option on the foreign currency with a strike price K_t . Then, whenever $S_{t+1} > K_t$, the agent buys FCUs at the price K_t . So, the minimum payoff of the hedged carry trade is:

$$(F_t - S_{t+1}) + (S_{t+1} - K_t) - C_t(1 + r_t) = F_t - K_t - C_t(1 + r_t). \quad (11)$$

Similarly, suppose that an agent buys one FCU forward. Then, the worst case scenario in the standard carry trade is a large depreciation of the foreign currency. In this state of the world the agent sells the foreign currency he receives from the forward contract at a low

value of S_{t+1} . Suppose that agents buy at time t a put option on the foreign currency with a strike price K_t . Then, whenever $S_{t+1} < K_t$, the agent sells FCUs at a price K_t . In this case the minimum payoff of the hedged carry trade is:

$$(S_{t+1} - F_t) + (K_t - S_{t+1}) - P_t(1 + r_t) = K_t - F_t - P_t(1 + r_t). \quad (12)$$

We define the hedged carry-trade strategy as:

- If $F_t \geq S_t$, sell $1/F_t$ FCUs forward and buy $1/F_t$ call options
- If $F_t < S_t$, buy $1/F_t$ FCUs forward and buy $1/F_t$ put options.

In order to normalize the size of the bet to one dollar, we set the amount of FCUs traded equal to $1/F_t$. The dollar payoff to this strategy is:

$$z_t^H = \begin{cases} z_{t+1} + z_{t+1}^C/F_t & \text{if } F_t \geq S_t, \\ z_{t+1} + z_{t+1}^P/F_t & \text{if } F_t < S_t, \end{cases} \quad (13)$$

where z_{t+1} is the carry-trade payoff defined in (5).

An alternative way to implement the hedged carry trade is to use options only, instead of using a combination of forwards and options. Under this alternative implementation we buy $1/F_t$ call options on the foreign currency when it is at a forward discount and $1/F_t$ put options on the foreign currency when it is at a forward premium. Using the put-call-forward parity condition,

$$(C_t - P_t)(1 + r_t) = F_t - K_t, \quad (14)$$

it is easy to show that this strategy for hedging the carry trade is equivalent to the one described above.⁸

The minimum payoff to the hedged carry trade, h_t , is negative. To see this we can use the put-call-forward parity condition, (14) and equations (11) and (12) to write the minimum payoffs as follows:

$$h_t = \begin{cases} -P_t(1 + r_t)/F_t & \text{if } F_t \geq S_t, \\ -C_t(1 + r_t)/F_t & \text{if } F_t < S_t. \end{cases} \quad (15)$$

Since option prices are positive, h_t is negative.

It is useful to summarize the realized payoffs to the hedged carry trade as follows:

$$z_{t+1}^H = \begin{cases} h_t & \text{if option is in the money,} \\ z_{t+1} - c_t(1 + r_t) & \text{if option is out of the money.} \end{cases}$$

⁸This equivalent requires that the strike price of the options be the same in the two strategies.

The variable c_t denotes the cost of the put or call option. Note that the option is in the money in the peso states as well as in some non-peso states.

Using options to assess the effect of peso problems Equation (7) implies that, conditional on being in a non-peso state at date t :

$$(1 - p) \int_{\Omega^N} [M(s_{t+1})z^H(s_{t+1})] dG(s_{t+1}|s_t) + ph(s_t)M' = 0. \quad (16)$$

Taking expectations over all non-peso states we obtain:

$$(1 - p)E^G [M(s_{t+1})z^H(s_{t+1})] + pE^G [h(s_t)] M' = 0. \quad (17)$$

In our empirical section we document that the covariance between $M(s_{t+1})$ and the payoff to the hedged carry trade, conditional on s_t being in the non-peso state, is not statistically different from zero. Using this fact we can re-write equation (17) as:

$$(1 - p)E^G [M(s_{t+1})] E^G [z^H(s_{t+1})] + pE^G [h(s_t)] M' = 0.$$

Using this equation to solve for $(1 - p)E^G [M(s_{t+1})]$ and replacing this term in equation (10), we obtain:

$$z' = E^G [h(s_t)] \frac{E^G [z(s_{t+1})]}{E^G [z^H(s_{t+1})]}. \quad (18)$$

We can estimate the variables on the right-hand side of equation (18) and compute an estimate of z' . In estimating z' we do not have to take a stand on the values of p , $E^G [M(s_{t+1})]$, or M' .

Given our estimate of z' and a value of p we can use equation (10) to estimate $M'/E^G [M(s_{t+1})]$,

$$\frac{M'}{E^G [M(s_{t+1})]} = \frac{(1 - p)E^G [z(s_{t+1})]}{p(-z')}. \quad (19)$$

There are two possible outcomes of these calculations. The first outcome is that the confidence interval for z' encompasses only very large negative values for z' . This outcome would support the conventional view that the peso event consists of a very large, negative payoff to the carry trade. The second outcome is that only relatively small, negative values of z' are consistent with equation (18). In this case a peso event can still explain the positive average payoff to the carry trade, but only if $M'/E^G [M(s_{t+1})]$ is large relative to $E^G [M(s_{t+1})]$. So, the carry trade makes relatively small losses in the peso event, but traders value those losses very highly.

A natural question is whether the implied value of $M'/E^G [M (s_{t+1})]$ is empirically plausible. To answer this question we consider an equity strategy whose payoff is also potentially affected by the peso event, s' . Using hedged and unhedged versions of this strategy we obtain an alternative estimate of $M'/E^G [M (s_{t+1})]$. We then assess whether this estimate of $M'/E^G [M (s_{t+1})]$ is consistent with the one implied by equation (19). The equity strategy involves borrowing one dollar at the Treasury-bill rate, r_t , and investing it in the S&P 100 index.⁹ We denote the ex-dividend price of the index and the associated dividend yield by V_t and d_t , respectively. The payoff to this strategy in non-peso states is given by:

$$x_{t+1} = V_{t+1}/V_t + d_t - (1 + r_t).$$

We denote by x' the payoff to this strategy in the peso state.

Now consider the following hedged version of the equity strategy: borrow at the Treasury-bill rate to invest in the S&P 100 index and buy at-the-money put options on the S&P 100 index. These put options exactly compensate an investor for a fall in the S&P 100. It follows that, any time the S&P 100 index falls, the payoff to the hedged stock strategy is the dividend yield on the stock index minus the dollar interest rate, and the price of the option ($c_t^x (1 + r_t)$). By assumption the stock index falls in the peso state as well as in some non-peso states. In these states the payoff to the hedged stock strategy is $d_t - c_t^x (1 + r_t)$.

In summary, the payoff to the hedged stock strategy net of the options cost is given by:

$$x_{t+1}^H = \begin{cases} x_{t+1} - c_t^x (1 + r_t) & \text{if } V_{t+1}/V_t - 1 \geq 0 \\ d_t - r_t - c_t^x (1 + r_t) & \text{if } V_{t+1}/V_t - 1 < 0 \end{cases}$$

The payoffs to the unhedged equity strategy must satisfy:

$$(1 - p) \int_{\Omega^N} M(s_{t+1})x(s_{t+1})dG(s_{t+1}|s_t) + pM'x' = 0, \quad (20)$$

Taking expectations with respect to non-peso states:

$$(1 - p)E^G [M(s_{t+1})x(s_{t+1})] + pM'x' = 0. \quad (21)$$

The payoffs to the hedged-equity strategy must satisfy:

$$(1 - p) \int_{\Omega^H} M(s_{t+1})x^H(s_{t+1})dG(s_{t+1}|s_t) + pM' \{d(s_t) - r(s_t) - c^x(s_t) [1 + r(s_t)]\} = 0. \quad (22)$$

Taking expectations with respect to non-peso states:

$$(1 - p)E^G [M(s_{t+1})x^H(s_{t+1})] + pM'E^G \{d(s_t) - r(s_t) - c^x(s_t) [1 + r(s_t)]\} = 0. \quad (23)$$

⁹The choice of this index is driven by data considerations.

We can use the payoffs from the hedged and unhedged stock strategy to generate estimates of $M'/E^G [M(s_{t+1})]$ and x' . We proceed as follows. We solve equation (21) for x' ,

$$x' = \frac{E^G [M(s_{t+1})x(s_{t+1})] E^G \{d(s_t) - r(s_t) - c^x(s_t) [1 + r(s_t)]\}}{E^G [M(s_{t+1})x^H(s_{t+1})]}. \quad (24)$$

Equation (23) implies:

$$M' = \frac{(1 - p)E^G [M(s_{t+1})x(s_{t+1})]}{p(-x')}. \quad (25)$$

We use the Fama-French (1993) model to compute a time series for $M(s_t)$ and estimate $E^G [M(s_{t+1})x(s_{t+1})]$ and $E^G [M(s_{t+1})x^H(s_{t+1})]$. Given a value of p we then estimate x' and $M'/E^G [M(s_{t+1})]$. The key test of the second interpretation of the peso event is whether the value of $M'/E^G [M(s_{t+1})]$ that emerges from this procedure is consistent with that implied by equation (19).

The next three sections of this paper provides the inputs necessary to implement the procedures just described. In Section 6 we report our results.

3 Data

In this section we describe our data sources for spot and forward exchange rates and interest rates. We also describe the options data that we use to analyze the importance of the peso problem.

Spot and forward exchange rates Our data set on spot and forward exchange rates, obtained from Datastream, covers 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 UK, and the U.S.

The data consist of daily observations for bid and ask spot exchange rates and one-month forward exchange rates. We convert daily data into non-overlapping monthly observations (see Appendix A for details).

Our data spans the period from January 1976 to January 2008. However, the sample period varies by currency (see Appendix A for details). Exchange rate quotes (bid, ask, and mid, defined as the average of bid and ask) against the British pound (GBP) are available beginning as early as 1976. Bid and ask exchange rate quotes against the U.S. dollar (USD)

are only available from January 1997 to January 2008. We obtain mid quotes over the longer sample against the dollar by multiplying GBP/FCU quotes by USD/GBP quotes.

Interbank interest rates and covered interest parity We also collected data on interest rates in the London interbank market from Datastream. These data are available for 17 countries/currencies: Australia, Belgium, Canada, Denmark, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, South Africa, Sweden, Switzerland, the UK, the U.S., and the Euro.

The data consist of daily observations for bid and ask eurocurrency interest rates. We convert daily data into non-overlapping monthly observations. Our data spans the period from January 1976 to January 2008, with the exact sample period varying by currency (see Appendix A for details).

To assess the quality of our data set we investigate whether covered-interest parity (CIP) holds taking bid-ask spreads into account. We find that deviations from CIP are small and rare. Details of our analysis are provided in Appendix B.

Option prices We use two options data sets. Our first data set is from the Chicago Mercantile Exchange (CME). These data consist of daily observations for the period from January 1987 to January 2008 on the prices of put and call options against the U.S. dollar for the Australian dollar, the Canadian dollar, the Euro, the Japanese yen, the Swiss franc, and the British pound. Appendix C specifies the exact period of availability for each currency.

Since we compute carry-trade payoffs at a monthly frequency, we use data on options that are one month from maturity (see Appendix C for details). We work exclusively with options expiring mid-month (on the Friday preceding the third Wednesday). We measure option prices using settlement prices for transactions that take place exactly 30 days prior to the option's expiration date. We measure the time- t forward, spot, and option strike and settlement prices on the same day, and measure the time $t + 1$ spot price on the option expiration date. To compute net payoffs we multiply option prices by the 30-day eurodollar interest rate obtained from the Federal Reserve Board. This 30-day interest rate is matched to the maturity of our options data set.

Our second options data is from J.P. Morgan. These data consist of daily observations on one-month at-the-money implied volatility quotes, and spot exchange rates for the following currencies: the Australian dollar, the Canadian dollar, the Danish krone, the Euro, the

Japanese yen, the Swiss franc, the British pound, the New Zealand dollar, the Norwegian krone, the Swedish krone, and the South African rand. Our sample period is from January 1996 to January 2008. We convert the implied volatility quotes to option prices using the Black-Scholes formula in combination with forward premia calculated using the data described in Appendix A. We use the same transactions dates as for the CME data. The implied volatilities in the two data sets are very similar.

Bid-ask spreads in exchange rates Table 1 displays median bid-ask spreads for spot and forward exchange rates measured in log percentage points ($100 \times \ln(\text{Ask}/\text{Bid})$). The left-hand panel reports spreads over the longest available sample for quotes against the British pound. The center panel reports spreads after the introduction of the Euro for quotes against the pound. The right-hand panel reports spreads over the longest available sample for quotes against the U.S. dollar.

Four observations emerge from Table 1. First, bid-ask spreads are wider in forward markets than in spot markets. Second, there is substantial heterogeneity across currencies in the magnitude of bid-ask spreads. Third, with the exception of South Africa, bid-ask spreads have declined for all currencies in the post-1999 period. This drop partly reflects the advent of screen-based electronic foreign-exchange dealing and brokerage systems, such as Reuters' Dealing 2000-2, launched in 1992, and the Electronic Broking Service launched in 1993.¹⁰ Fourth, over comparable sample periods, the bid-ask spreads for spot and forward exchange rates against the U.S. dollar are always lower than the analogous spreads against the British pound.

4 Payoffs to the carry trade

In this section we study the properties of the payoffs to the carry trade. First, we compute the mean and variance of the payoff to the carry trade with and without transactions costs. Second, we investigate whether there are large, negative payoffs in our sample, that could plausibly be referred to as peso events.¹¹ Third, we study the covariance between the payoffs to the carry trade and various risk factors using both time series and panel data.

¹⁰It took several years for these electronic trading systems to capture large transactions volumes. We break the sample in 1999, as opposed to in 1992 or 1993, to fully capture the impact of these trading platforms.

¹¹Recall that we are using Cochrane's (2001) definition of a peso event as a rare event that might or might not be realized in sample.

We consider two versions of the carry trade. In the ‘carry trade without transaction costs’ we assume that agents can buy and sell currency at the average of the bid and ask rates. We compute S_t as the average of the bid (S_t^b) and the ask (S_t^a) spot exchange rates,

$$S_t = (S_t^a + S_t^b) / 2,$$

and F_t as the average of the bid (F_t^b) and the ask (F_t^a) forward exchange rates,

$$F_t = (F_t^a + F_t^b) / 2.$$

The ask (bid) exchange rate is the rate at which a participant in the interdealer market can buy (sell) dollars from (to) a currency dealer.

In the ‘carry trade with transaction costs’ we take bid-ask spreads into account when deciding whether to buy or sell foreign currency forward and in calculating payoffs. In this case the number of FCUs sold forward, x_t , is given by:

$$x_t = \begin{cases} +1/F_t^b & \text{if } F_t^b/S_t^a > 1, \\ -1/F_t^a & \text{if } F_t^a/S_t^b < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (26)$$

The payoff to this strategy is:

$$z_{t+1} = \begin{cases} x_t (F_t^b - S_{t+1}^a) & \text{if } x_t > 0, \\ x_t (F_t^a - S_{t+1}^b) & \text{if } x_t < 0, \\ 0 & \text{if } x_t = 0. \end{cases} \quad (27)$$

4.1 Mean and variance of carry-trade payoffs

We consider the carry-trade strategy for individual currencies as well as for portfolios of currencies. For now we focus attention on the payoffs to an equally-weighted portfolio of carry-trade strategies.¹² This portfolio is constructed by betting $1/n_t$ of one unit of the home currency in each individual currency carry trade. Here n_t denotes the number of currencies in our sample at time t . In the remainder of the paper, unless otherwise noted, we use the term ‘‘carry-trade strategy’’ to refer to the equally-weighted carry trade. We report all statistics on an annualized basis. Table 2 reports the mean, standard deviation, and Sharpe ratio of the monthly payoffs to the carry trade, with and without transaction costs. We consider two alternative home currencies, the British pound and the U.S. dollar. Using the British pound as the home currency allows us to assess the importance of bid-ask spreads

¹²In Tables A2 and A3 of the Appendix we report results for individual currencies.

using a much longer time series than would be the case if we looked only at the U.S. dollar as the home currency.

Consider the results when the British pound is the home currency. Ignoring transaction costs, the Sharpe ratio of the equally-weighted carry-trade portfolio is roughly 0.811. Taking bid-ask spreads into account reduces the Sharpe ratio to 0.579. But the Sharpe ratio is statistically different from zero with and without transaction costs. Next, consider the results when the dollar is the home currency. Ignoring transaction costs, the Sharpe ratio of the equally-weighted carry-trade portfolio is roughly 1.061. Taking bid-ask spreads into account reduces the Sharpe ratio to 0.867. But, once again, the Sharpe ratio is statistically different from zero, both with and without transaction costs. The impact of transaction costs is smaller when the dollar is the base currency, because bid-ask spreads are lower for the dollar than for the pound (see Table 1).

The results in Table 2 may overstate the effect of transaction costs on the carry-trade payoff because there are alternative ways to execute the carry trade that can reduce these costs. We compute the payoffs to the carry trade executed through forward markets. However, when interest-rate differentials are persistent, it can be more cost efficient to execute the carry trade through money markets. To be concrete suppose that the Yen interest rate is lower than the dollar interest rate. We can implement the carry trade by borrowing Yen, converting the proceeds into dollars in the spot market and investing the dollars in the U.S. money market. This dollar investment and Yen loan are rolled over as long as interest rate differentials persist. When the strategy is initially implemented, the investor pays one bid-ask spread to convert the proceeds of the Yen loan into dollars. In the final phase of the strategy the investor pays a second bid-ask spread in the spot exchange market to convert dollar into Yen to pay back the initial Yen loan. In contrast, the strategy that underlies the payoffs in Table 2 incurs transaction costs associated with closing out the investor's position every month.

Taken together, our results indicate that, while transaction costs are quantitatively important, they do not explain the profitability of the carry trade. For the remainder of this paper we abstract from transaction costs and work with spot and forward rates that are the average of bid and ask rates.¹³ Given this decision we can work with the longer data set (from January 1976 to January 2008) using the U.S. dollar as the home currency.

¹³In an earlier version of this paper (Burnside, Eichenbaum, Kleshchelski, and Rebelo (2006)) we present a more comprehensive set of results for the carry trade payoffs taking bid-ask spreads into account.

Table 3 reports statistics for the payoffs to the equally-weighted carry trade and summary statistics for the individual-currency carry trades. The latter are computed by taking the average of the statistics for the carry trade applied to each of the 20 currencies in our sample. To put our results into perspective, we also report statistics for excess returns to the value-weighted U.S. stock market. Two results emerge from this table. First, there are large gains to diversification. The average Sharpe ratio across currencies is 0.479, while the Sharpe ratio for an equally-weighted portfolio of currencies is 0.972. This large rise in the Sharpe ratio is due to the fact that the standard deviation of the payoffs is much lower for the equally-weighted portfolio.¹⁴ Second, the Sharpe ratio of the carry trade is substantially larger than that of the U.S. stock market (0.972 versus 0.461). While the average excess return to the U.S. stock market is larger than the payoff to the carry trade (0.068 versus 0.050), the returns to the U.S. stock market are much more volatile than the payoffs to the carry trade (0.148 versus 0.051).

Figure 1 displays 12-month moving averages of the realized payoffs and Sharpe ratios associated with the carry trade. Negative payoffs are relatively rare and positive payoffs are not concentrated in a small number of periods. In addition, there is no pronounced time trend in either the payoffs or the Sharpe ratios.

4.2 Fat tails

So far we have emphasized the mean and variance of the payoffs to the carry trade. These statistics are sufficient to characterize the distribution of the payoffs only if this distribution is normal. We now analyze other properties of the payoff distribution. Figure 2 shows the sample distributions of the dollar payoffs to the carry trade and to the U.S. stock market.¹⁵ In addition we display a normal distribution with the same mean and variance as the empirical distribution of the payoffs. It is evident that the distributions of both payoffs are leptokurtic, exhibiting fat tails. This impression is confirmed by Table 3 which reports skewness and excess kurtosis statistics, as well as the results of the Jarque-Bera normality tests.¹⁶ While

¹⁴Since there are gains to combining currencies into portfolios, it is natural to construct portfolios that maximize the Sharpe ratio. See Burnside, Eichenbaum, Kleshchelski, and Rebelo (2006) for details on how to implement this strategy. For the sample considered in this paper the Sharpe ratios associated with the equally-weighted and optimally-weighted portfolios are very similar. For this reason we do not report results for the latter portfolio.

¹⁵Figure A1 in the Appendix shows the sample distributions of the dollar payoffs to the carry trade implemented for each of our 20 currencies.

¹⁶In Table A4 of the Appendix we report skewness, excess kurtosis, and the Jarque-Bera normality test for the dollar payoffs to the carry trade implemented for each of our 20 currencies.

both distributions have fat tails, the bad outcomes associated with the carry trade are small compared to those associated with the U.S. stock market (see Figure 2).

4.3 Risk factor analysis of carry-trade payoffs

In this subsection we show that the covariance of the payoffs to the carry trade and traditional risk factors is not statistically different from zero. We do so using both time-series and panel-data analysis. In what follows we consider real quarterly dollar-denominated payoffs, R_t^e , to our carry-trade strategies.¹⁷ These payoffs must satisfy:

$$E_t (R_{t+1}^e m_{t+1}) = 0. \quad (28)$$

where m_{t+1} is the SDF that prices real dollar-denominated payoffs. We consider linear SDFs of the form:

$$m_t = \xi [1 - (f_t - \mu)' b]. \quad (29)$$

Here ξ is a scalar, f_t is a vector of risk factors, $\mu = E(f_t)$, and b is a conformable vector. It follows from equation (28) and the law of iterated expectations that:

$$E (R_t^e m_t) = 0. \quad (30)$$

Equations (30) and (29) imply that:

$$E(R_t^e) = \beta \lambda$$

where

$$\begin{aligned} \beta &= \text{cov}(R_t^e, f_t') V_f^{-1}, \\ \lambda &= V_f b. \end{aligned} \quad (31)$$

Here V_f is the covariance matrix of the factors, β is a measure of the systematic risk associated with the payoffs, and λ is a vector of risk premia. Note that β is the population value of the regression coefficient of R_t^e on f_t . Our time-series analysis focuses on estimating the betas for different candidate risk factors. Our panel analysis provides complementary evidence on the importance of different risk factors by estimating alternative SDF models.

¹⁷In Appendix D we show how we convert monthly payoffs to real quarterly excess returns.

Time-series risk-factor analysis We consider the following risk factors: the excess returns to the value-weighted U.S. stock market, the Fama-French (1993) factors (the excess return to the value weighted U.S. stock market, the size premium (SMB), and the value premium (HML)), real U.S. per capita consumption growth (nondurables and services), the factors proposed by Yogo (2006) (the growth rate of per capita consumption of nondurables and services, the growth rate of the per capita service flow from the stock of consumer durables, and the return to the value-weighted U.S. stock market), luxury sales growth (obtained from Ait-Sahalia, Parker and Yogo (2004)), GDP growth, the Fed Funds Rate, the term premium (the yield spread between the 10 year Treasury bond and the three month Treasury bill), the liquidity premium (the spread between the three month Eurodollar rate and the three month Treasury bill), and two measures of volatility, the VIX and the VXO (the implied volatility of the S&P 500 and S&P 100 index options, respectively, calculated by the Chicago Board Options Exchange).

Table 4 reports the estimated regression coefficients associated with the different risk-factor candidates, along with the corresponding test statistics. Our key finding is that none of the risk factors covaries significantly with the payoffs to the carry trade. As Table 3 shows, the average payoff to the carry trade is statistically different from zero. Factors that have zero β s clearly cannot account for these payoffs.

Our procedure for assessing the importance of peso events assumes that the covariance between the payoffs to the carry trade and the SDF is zero in non-peso event states. Recall that there are no large, negative payoffs to the carry trade in our sample. The results of this subsection provide evidence for the zero-covariance assumption used in our procedure.

Panel risk-factor analysis We now discuss the results of estimating the parameters of SDF models built using the risk factors detailed in Table 4. In addition, we also use the Campbell-Cochrane (1990) SDF (see Appendix D for details on how we construct this SDF).¹⁸ We use the estimated SDF models to generate predicted average payoffs to the carry-trade strategy and the 25 Fama-French portfolios of U.S. stocks sorted on the basis of firm size and the ratio of book-to-market value. We then study how well the model explains the average payoff associated with the carry trade, as well as the cross-sectional variation of the different payoffs used in the estimation procedure.

¹⁸Verdelhan (2007) argues that open-economy models in which agents have Campbell-Cochrane (1999) preferences can generate non-trivial deviations from UIP.

We estimate b and μ by the generalized method of moments (GMM) using equation (30) and the moment condition $\mu = E(f_t)$. In practice, the variable R_t^e in (30) is a 26×1 vector of time- t payoffs to the carry-trade strategy and the 25 Fama-French portfolios. The first stage of the GMM procedure, which uses the identity matrix to weight the GMM errors, is equivalent to the Fama-MacBeth (1973) procedure. The second stage uses an optimal weighting matrix.¹⁹

It is evident from equations (29) and (30) that $\xi = E(m_t)$ is not identified. Fortunately, the point estimate of b and inference about the model's over-identifying restrictions are invariant to the value of ξ , so we set ξ to one for convenience. It follows from equations (29) and (30) that:

$$E(R_t^e) = -\frac{\text{cov}(R_t^e, m_t)}{E(m_t)} = E[R_t^e(f_t - \mu)' b]. \quad (32)$$

For each risk factor, or vector of factors, Table 5 reports the first and second-stage estimates of b , the R^2 , and the value of Hansen's (1982) J statistic used to test the over-identifying restrictions implied by equation (30).²⁰ The results fall into two categories, depending on whether the b parameters associated with a particular risk-factor model are estimated with any degree of precision. For the CAPM and the Fama-French model, the b parameters are precisely estimated and are statistically different from zero.²¹ But the over-identifying restrictions associated with these models are overwhelmingly rejected. Interestingly, the CAPM explains none of the cross-sectional variation in average payoffs. In contrast the Fama-French model explains a substantial component of the cross-sectional variation in expected payoffs.

The second category of results pertains to the remaining risk-factor models. For all these models, the b parameters associated with the corresponding risk factors are estimated with great imprecision. In no case can we reject the null hypothesis that the b parameters are equal to zero or that the model-implied excess return to the carry trade is equal to zero. Moreover, the R^2 statistics paint a dismal picture of the ability of these risk factors to explain the cross-sectional variation in expected payoffs. Indeed, most of the R^2 statistics are actually *negative*. However, because the b parameters are estimated with enormous

¹⁹Details of our GMM procedure are provided in Appendix E.

²⁰The R^2 measure is: $R^2 = 1 - \frac{(\bar{R}^e - \hat{R}^e)'(\bar{R}^e - \hat{R}^e)}{(\bar{R}^e - \bar{R}^e)'(\bar{R}^e - \bar{R}^e)}$. Here \hat{R}^e denotes the predicted mean payoff (the sample analogue of the right-hand side of equation (32)) evaluated at the point estimate of b . We denote by \bar{R}^e the actual mean payoff (the sample analogue of the left-hand side of equation (32)). Finally, we denote by \bar{R}^e the average across the elements of \bar{R}^e . This measure of R^2 is invariant to the value of ξ .

²¹An exception is the coefficient associated with the SMB factor in the Fama-French model.

imprecision, it is difficult to statistically rule out regions of the parameter space for which the model’s predictions for expected payoffs are consistent with the data. Since there is little information in the sample about the b parameters it is hard to statistically reject these factor models.²²

We now provide an alternative perspective on the performance of four SDF models that have received substantial attention in the literature. These models are: the CAPM model, the C-CAPM model, the Extended C-CAPM model, and the Fama-French model. Figure 3 plots the predictions of these models for $E(R_t^e)$ against the sample average of R_t^e . The circles pertain to the Fama-French portfolios, while the star pertains to the carry trade. It is clear that the first three models do a poor job of explaining the average payoffs to the Fama-French portfolios and the carry trade. Not surprisingly, the Fama-French model does a reasonably good job at pricing the payoffs to the Fama-French portfolios. However, the model greatly understates the average payoffs to the carry trade. The annualized excess return to the carry trade is 5.20 percent. The Fama-French model predicts that this return should equal -0.16 percent. The solid line through the star is a two-standard-error band for the difference between the data and model excess return, i.e. the pricing error. Clearly, we can reject the hypothesis that the model accounts for the average payoffs associated with the carry trade, i.e. from the perspective of the model the carry trade has a positive alpha.

Recall that our procedure for evaluating the importance of peso events assumes that the covariance between the payoffs to the carry trade and the SDF is zero in non-peso states.²³ Viewed overall, the results in this section provide strong support for this assumption.

5 Payoffs to the hedged carry trade

In this section we discuss the empirical properties of the hedged carry trade. As discussed in Section 3 our primary option data set is from the CME and covers six currencies and a shorter sample period (January 1987 to January 2008) than our data set on forward contracts. We compute the payoffs to the carry trade and hedged carry trade over the sample period and set of currencies for which options data are available.

²²We also estimated the parameters of these factor models using data beginning in 1948 for the Fama French portfolio returns. This extension has very little impact on the precision with which we estimate the b parameters.

²³We only consider linear stochastic discount factors. We do not rule out the possibility that some yet to be discovered non-linear stochastic discount factor models can simultaneously rationalize the cross-sectional variation in the carry-trade and Fama-French portfolios.

We implement the hedged carry trade using strike prices that are close to “at-the-money,” that is K_t is as close as possible to the current spot exchange rate, S_t . We choose these strike prices for two reasons. First, this choice ensures that, in a peso state, the options in the money. So, we do not have to take an a priori stand on the magnitude of z' . Second, options that are way out-of-the-money tend to be sparsely traded and relatively expensive.²⁴ To illustrate how trading volume varies with moneyness we use data from the CME that contains all transactions on currency puts and calls for a single day (November 14, 2007). This data set contains records for 260 million contract transactions. Figure 4 displays the volume of calls and puts of five currencies (the Canadian dollar, the Euro, the Japanese yen, the Swiss franc, and the British pound) against the U.S. dollar. In all cases the bulk of the transactions are concentrated on strike prices near the spot price. Interestingly, there is substantial skewness in the volume data. Most call options are traded at strike prices greater than or equal to the spot price. Similarly, most put options are traded at strike prices less than or equal to the spot price.

Table 6 reports the mean, standard deviation, and Sharpe ratio of the monthly payoffs to the carry trade, the hedged carry trade, and the U.S. stock market. Recall that we are abstracting from bid-ask spreads in calculating the payoffs to the hedged carry trade. In Section 4 we find that taking transaction costs into account reduces the average payoff to the unhedged carry trade executed with the U.S. dollar as the home currency by 9.0 percent. Using the data that underlies Figure 4 we compute average bid-ask spreads for puts and calls against the Canadian dollar, the Euro, the Japanese yen, and the Swiss franc. The average bid-ask spread in this data is 5.2 percent.²⁵ This estimate is slightly higher than the point estimate of 4.4 percent provided by Chong, Ding, and Tan (2003).²⁶ We use our estimate of the bid-ask spread to assess the impact of transaction costs on the average payoffs of the hedged carry trade. We find that the average payoff to the hedged carry trade declines by 12 percent as a result of transaction costs.²⁷ So, as with the unhedged carry trade, transaction

²⁴See Jurek (2008) for a detailed analysis of the impact of hedging using out-of-the-money options. Jurek finds that the payoffs to the carry trade hedged with these options is positive and highly statistically significant. See also Bhansali (2007) who considers hedging strategies in the course of investigating the relation between implied exchange-rate volatility and the payoffs to the carry trade.

²⁵The average bid-ask spreads for individual currencies are: Canadian dollar call 5.33 percent, put 4.39 percent, Euro call 4.26 percent, put 4.78 percent, Japanese yen call 5.26 percent, put 5.61 percent, Swiss franc call 5.33 percent, put 6.35 percent, and British pound call 4.29 percent, and put 4.57 percent.

²⁶Chong, Ding and Tan’s (2003) estimate is based on data from the Bloomberg Financial Database for the period from December 1995 through March 2000.

²⁷To assess the impact of transaction costs we increased the prices of the puts and calls used in our strategy

costs are significant for the hedged carry trade but do not eliminate the average payoff.

The average payoff to the hedged carry trade is lower than that of the carry trade (2.51 versus 3.32 percent). However, the average payoffs of the carry trade and the hedged carry trade are not statistically different from each other.

The first panel of Figure 5 displays a 12-month moving average of the realized payoffs for the hedged and unhedged carry-trade strategies. The second panel displays a 12-month moving average of the realized Sharpe ratios for both carry-trade strategies. The payoffs and Sharpe ratios of the two strategies are highly correlated. In this sense, the hedged and unhedged carry trade appear quite similar.

There is an important dimension along which the payoffs of the two carry-trade strategies are quite different. As Figure 6 shows, the distribution of payoffs to the unhedged carry trade has a substantial left tail. Hedging eliminates most of the left tail. This property reflects the fact that our version of the hedged carry trade uses options with strike prices that are close to at the money.

Based on the previous results we conclude that the profitability of the carry trade remains intact when we hedge away substantial losses. It is still possible, however, that hedging changes the nature of the payoffs so as to induce a correlation with traditional risk measures. We now investigate this possibility.

Recall from equation (31) that β is the population value of the regression coefficient of the carry-trade payoff on candidate risk factors. Table 7 reports our estimates of β for the hedged carry trade using the risk factors considered in Section 5. We find that, with the exception of GDP growth and the Fama-French HML factor, the estimated values of β are not significantly different from zero. So, these factors aside, we cannot reject the hypothesis that the payoffs to the hedged carry trade are not compensation for risk. Evidently, hedging away peso events does not change the payoffs in such a way that induces a statistically significant correlation between carry trade payoffs and risk factors. We return to the case of the Fama-French factors and GDP growth below.

We now turn to a panel risk-factor analysis of the hedged carry-trade payoffs. We estimate the parameters of the same SDF models considered in Section 5. Our estimation results are generated using a 26×1 vector of time- t payoffs to the hedged carry-trade strategy and the 25 Fama-French portfolios. We report our results in Table 8. The key finding is that the

by one half of the average bid-ask spread (2.6 percent).

results for the hedged carry trade are very similar in character to those reported in Table 5 for the unhedged carry trade over the longer sample period. These results can be summarized as follows. First, for the CAPM and the Fama-French model, the b parameters are precisely estimated and are statistically different from zero. The over-identifying restrictions associated with these models are overwhelmingly rejected. Second, the b parameters associated with the other risk-factor models are estimated with great imprecision. Not surprisingly, in these cases we cannot reject the over-identifying restrictions associated with the model. For these models we cannot reject either the null hypothesis that the b parameters are equal to zero or the associated implication that the model-implied excess return to the carry trade is equal to zero. Third, the only model for which the cross-sectional R^2 s are not negative is the Fama-French model. Finally, the SDF model based on GDP growth does very poorly in the sense that the R^2 is very low and the overidentifying restrictions are rejected.

Figure 7 displays the predictions of the CAPM, the C-CAPM, the extended C-CAPM models, and the Fama-French model for $E^G [R^e(s_t)]$ against the sample average of R_t^e . The first three models cannot account for the average payoffs to either the hedged carry trade or the Fama-French portfolios. The Fama-French model does a reasonable job of explaining the average payoffs to the Fama-French portfolios, but fails to explain the average payoffs to the hedged carry trade. From the perspective of this model the hedged carry trade has a positive alpha that is statistically significant.

6 Assessing the importance of peso events

In this section we implement the strategy for assessing the importance of peso events discussed in Section 2. This section is organized as follows. In subsection 6.1 we report estimates of z' and $M'/E^G [M(s_{t+1})]$ based on the average payoffs to the unhedged and hedged carry trade. These estimates are computed using our benchmark CME data set. In subsection 6.2 we incorporate stock returns into our empirical analysis. We assess the robustness of our results in subsection 6.3 using data from J.P. Morgan. Finally, in subsection 6.4 we extend our analysis to allow for multiple peso events. Up to this point we reported all statistics on an annualized basis. In this section we report monthly statistics so that our calculations are easier to follow.

6.1 Benchmark estimates

Recall that our strategy for estimating z' is based on equation (18), which we repeat below for convenience:

$$z' = E^G [h(s_t)] \frac{E^G [z(s_{t+1})]}{E^G [z^H(s_{t+1})]}. \quad (33)$$

The empirical analysis summarized in the previous two sections provide us with the inputs necessary to estimate z' . Our estimates of these inputs are summarized in Table 9.

Our estimate of the minimum net payoff to the hedged carry trade, $E^G [h(s_t)]$, is equal to -0.012 . We estimate $E^G [z(s_t)]$ and $E^G [z^H(s_t)]$ by their sample averages, 0.0027 and 0.0021 , respectively.

Substituting these estimates into equation (33) we obtain a point estimate of z' equal to -0.0154 . The corresponding standard error is 0.0033 . The implied two-standard-error band for z' is $(-0.0218, -0.0090)$. Our point estimate for z' is only one standard deviation below the estimated value of $E^G (z(s_t))$. Even the lower bound of the confidence interval for z' is only 1.4 standard deviations away from our estimate of $E^G (z(s_t))$. In our view these results do not support interpreting the peso event as a large negative payoff to the carry trade.

Given an estimate of z' we can estimate $M'/E^G [M(s_{t+1})]$ using equation (19), repeated here for convenience:

$$\frac{M'}{E^G [M(s_{t+1})]} = \frac{(1-p)E^G [z(s_{t+1})]}{p(-z')}. \quad (34)$$

Barro (2006) estimates a value of p equal to 0.017 . Motivated by Barro's estimate, we use a value of $p = 0.0014$. This value implies that with probability 0.983 no peso event occurs over a 12-month period. Using this value of p in equation (34) yields an estimate of $M'/E^G [M(s_{t+1})]$ equal to 121.7 with a standard error of 38 . This result supports the view that in a peso event the carry trade makes relatively small losses but traders value those losses very highly.

6.2 Incorporating stock market data into our analysis

In section 2 we develop estimators of $M'/E^G [M(s_{t+1})]$ and x' , the payoff to the stock market strategy in a peso state. Our estimators are based on the average payoffs in non-peso states to a hedged and unhedged stock market investment strategy. We repeat the two key equations

underlying these estimators for convenience:

$$x' = \frac{E^G [M(s_{t+1})x(s_{t+1})] E^G \{d(s_t) - r(s_t) - c^x(s_t) [1 + r(s_t)]\}}{E^G [M(s_{t+1})x^H(s_{t+1})]}. \quad (35)$$

$$M' = \frac{(1 - p)E^G [M(s_{t+1})x(s_{t+1})]}{p(-x')}. \quad (36)$$

We use estimates of the Fama-French (1993) model fit to the 25 Fama-French portfolios over the period 1986–2007 to compute a time series for $M(s_t)$.²⁸ We then use sample averages of $M(s_{t+1})x(s_{t+1})$ and $M(s_{t+1})x^H(s_{t+1})$ to estimate $E^G [M(s_{t+1})x(s_{t+1})]$ and $E^G [M(s_{t+1})x^H(s_{t+1})]$, respectively.

Our results are summarized in Table 9. We begin by contrasting the effect of hedging in stock markets and in currency markets. Hedging substantially reduces the excess return from investing in the stock market. The annualized rate of return drops from 9.8 percent to -5.2 percent as we go from the unhedged to the hedged stock market strategy. In sharp contrast, the annualized payoff to the carry trade only drops from 3.2 percent to 2.5 percent as we go from the unhedged to the hedged carry trade.

Using the same value of p discussed above, we estimate x' to equal -0.105 (see Table 9). This value of x' is roughly seven times larger in absolute value than z' . Moreover, our estimate of x' is roughly two and a half standard deviations away from the mean payoff to the unhedged equity strategy. In contrast, our estimate of z' is only one standard deviation away from the mean payoff to the unhedged carry trade. By either metric the peso event has a much larger impact on stock market payoffs than on carry trade payoffs. Finally, our point estimate of $M'/E^G [M(s_{t+1})]$ based on stock returns is equal to 107.8. Recall that our estimate of $M'/E^G [M(s_{t+1})]$ based on carry-trade payoffs is 121.7. Obviously, these two estimates are very similar. So the same value of $M'/E^G [M(s_{t+1})]$ can account for the equity premium and the observed average payoffs to the carry trade.

Taken together, the results of this subsection provide corroborating evidence for the view that the hallmark of a peso event is a large rise in the value of the SDF. Sampling uncertainty aside, this large rise is associated with large, negative stock market payoffs and modest, negative carry-trade payoffs.

²⁸The options data we use to construct the hedged equity strategy are available over this same time period.

6.3 Robustness analysis: J.P. Morgan data

To assess the robustness of our inference we begin by redoing our analysis using the six-currency version of the J.P. Morgan data set.²⁹ Our results are reported in Table 9.³⁰ Our estimates of $E^G(h(s_t))$, $E^G(z(s_t))$ and $E^G(z^H(s_t))$ imply an estimate of z' equal to -0.0181 with a standard error of 0.0039. Table 9 also reports results based on the 11 currency version of the J.P. Morgan data set. These estimates imply an estimate for z' equal to -0.0162 with a standard error of 0.0060. So, for both J.P. Morgan data sets, our estimate of z' is close to the estimate that we obtained with the CME data set (-0.0154). Once again, even taking sampling uncertainty into account, it does not appear that a peso event can be plausibly viewed as a large negative payoff to the carry trade.

6.4 Robustness analysis: allowing for multiple peso states

Under the assumption that there is a single peso state we find that the payoff to the unhedged carry trade is only moderately negative (-0.0154). However, there is trade in options that protect investors against much larger movements in exchange rates than those implied by our estimate of z' . At first glance, the fact that these way-out-of-the-money options are traded is a challenge for our interpretation of a peso event. We now show that this observation is not a problem for our interpretation by modifying our analysis to incorporate multiple peso states.

Suppose that there are L peso states of the world, $z'_i < 0$, $i = 1, \dots, L$ and that the value of M' is the same in all peso states. As above we assume that the probability of a peso state, both conditional and unconditional, is p . We also assume that, conditional on being in a peso state, the probability of z'_i is q_i . Here $\sum_{i=1}^L q_i = 1$. In this setting there can be many options with different strike prices. The payoffs to the unhedged carry trade must satisfy:

$$(1 - p)E^G[M(s_{t+1})] E^G[z(s_{t+1})] + pM' \sum_{i=1}^L q_i z'_i = 0. \quad (37)$$

Consider now the hedged carry trade strategy, where the hedging relies on at-the-money

²⁹If we use data from the CME over the JPM data set time period (1996-2007), the relevant values of the parameters are $h = -0.0028$, $c(1 + r) = 0.0090$ and $E_F(z) = 0.0039$, $E_F(z^H) = 0.0031$. These estimates imply a value of z' equal to: -0.0148 .

³⁰These values of $c(1 + r)$ and h are higher than in the CME data because the options in the JPM data are at the money, while those in the CME data set are slightly out of money.

options. Since these options are in the money in all peso states, it follows that:

$$(1 - p)E^G [M(s_{t+1})] E^G [z^H(s_{t+1})] + pM'h = 0. \quad (38)$$

Combining equations (37) and (38) we obtain:

$$\sum_{i=1}^L q_i z'_i = h \frac{E^G [z(s_{t+1})]}{E^G [z^H(s_{t+1})]}. \quad (39)$$

Our estimate of the right-hand side of equation (39) is: -0.0154 . It follows that the expected value of z' across all peso states is equal to -0.0154 . So, while there can be some large negative values of z' , these values have to have low probabilities.

In sum, the presence of multiple peso states renders our analysis consistent with the existence of currency options that have a wide array of strike prices. But, a large value of $M'/E^G [M(s_{t+1})]$ is still necessary to account for the average payoffs to the unhedged and hedged carry trade.

7 Conclusion

Equally-weighted portfolios of carry-trade strategies generate large positive payoffs and a Sharpe ratio that is almost twice as large as the Sharpe ratio associated with the U.S. stock market. We find that these payoffs are not correlated with standard risk factors. Moreover, standard SDF models do not explain the cross-sectional variation in expected equity and carry-trade payoffs.

A natural explanation for the positive average payoffs to the carry trade is that they reflect a peso problem. To investigate this possibility we develop a version of the carry trade that uses currency options to protect the investor from the downside risk from large, adverse movements in exchange rates. By construction, this hedged carry trade strategy eliminates the large negative payoffs associated with peso events. We show that the payoffs to the hedged carry trade are very similar to those of the unhedged carry trade. We argue that this result implies that the defining characteristic of a peso state is a high value of the SDF, not large losses in the carry trade. We also find that the same value of the SDF which rationalizes the observed payoffs to the carry trade also accounts for the observed equity premium.

REFERENCES

- Aït-Sahalia, Yacine, Jonathan Parker and Motohiro Yogo (2004) “Luxury Goods and the Equity Premium,” *Journal of Finance* 59 (6, Dec.): 2959–3004.
- Backus, David K., Silverio Foresi and Chris I. Telmer (2001) “Affine Term Structure Models and the Forward Premium Anomaly,” *Journal of Finance* 56 (1, Feb.): 279–304
- Backus, David K., Allan W. Gregory and Chris I. Telmer (1993) “Accounting for Forward Rates in Markets for Foreign Currency,” *Journal of Finance* 48 (5, Dec.): 1887–908.
- Bhansali, Vineer (2007) “Volatility and the Carry Trade,” *The Journal of Fixed Income*, 72-84, Winter.
- Barro, Robert, (2006) “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics* 121: 823–66.
- Bekaert, Geert (1996) “The Time-Variation of Expected Returns and Volatility in Foreign-Exchange Markets: A General Equilibrium Perspective,” *Review of Financial Studies* 9: 427–70.
- Bekaert, Geert and Robert J. Hodrick (1992) “Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets,” *Journal of Finance* 47 (2, Jun.): 467–509.
- Bilson, John F. O. (1981) “The ‘Speculative Efficiency’ Hypothesis,” *Journal of Business* 54 (3, Jul.): 435–51.
- Brunnermeier, Markus K., Stefan Nagel, Lasse H. Pedersen (2008) “Carry Trades and Currency Crashes,” forthcoming, *NBER Macroeconomics Annual 2008*.
- Burnside, Craig (2007) “The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk: A Comment,” National Bureau of Economic Research Working Paper No. 13129.
- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo (2006) “The Returns to Currency Speculation,” National Bureau of Economic Research Working Paper No. 12489.
- Campbell, John and John Cochrane, (1990) “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107: 205–51.
- Chong, Beng-Soon, David K. Ding and Kok-Hui Tan (2003) “Maturity Effect on Bid-Ask Spreads of OTC Currency Options,” *Review of Quantitative Finance and Accounting*, 21: 5–15.
- De Santis, Giorgio, and Bruno Gérard (1998) “How Big is the Premium for Currency Risk?” *Journal of Financial Economics* 49: 375–412.

- Engel, Charles (1996) “The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence,” *Journal of Empirical Finance* 3 (2, June), 123–92.
- Fama, Eugene F. (1984) “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* 14 (3, Nov.), 319–38.
- Fama, Eugene and Kenneth R. French (1993) “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics* 33 (1, Feb.), 3–56.
- Fama, Eugene and James D. MacBeth “Risk Return and Equilibrium: Empirical Tests,” *Journal of Financial Political Economy*, 71: 607-636, 1973.
- Farhi, Emmanuel and Xavier Gabaix (2008) “Rare Disasters and Exchange Rates,” National Bureau of Economic Research Working Paper No. 13805.
- Gabaix, Xavier (2007) “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance,” National Bureau of Economic Research Working Paper No. 13724.
- Galati, Gabriele and Michael Melvin (2004) “Why has FX Trading Surged? Explaining the 2004 Triennial Survey,” *BIS Quarterly Review* (Dec.), 67–74.
- Hansen, Lars Peter (1982) “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica* 50 (4, July), 1029–54.
- Julliard, Christian and Anisha Ghosh “Can Rare Events Explain the Equity Premium Puzzle?,” mimeo, London School of Economics, 2007.
- Jurek, Jakub W. “Crash-neutral Currency Carry Trades,” mimeo, Princeton University, 2008.
- Lustig, Hanno and Adrien Verdelhan (2007) “The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk,” *American Economic Review*, 97, 89–117.
- Rietz, Thomas A., (1988) “The Equity Risk Premium: A Solution,” *Journal of Monetary Economics* 22, 117–31.
- Verdelhan, Adrien (2007) “A Habit-Based Explanation of the Exchange Rate Risk Premium,” mimeo, Boston University.
- Villanueva, O. Miguel (2007) “Forecasting Currency Excess Returns: Can the Forward Bias Be Exploited?” *Journal of Financial and Quantitative Analysis* (42, Dec.), 963–90.
- Yogo, Motohiro (2006) “A Consumption-Based Explanation of Expected Stock Returns,” *Journal of Finance* 61 (2, April), 539–80.

TABLE 1
 MEDIAN BID-ASK SPREADS OF EXCHANGE RATES
 (percent)

	Quotes in FCU per GBP					Quotes in FCU per USD		
	Full Sample			1999:1-2007:1		Full Sample		
	Spot	1 Month	Sample	Spot	1 Month	Spot	1 Month	Sample
	Forward	Period		Forward		Forward	Period	
Austria	0.153	0.222	76:01-98:12			0.042	0.056	97:01-08:01
Belgium	0.158	0.253	76:01-98:12			0.111	0.118	97:01-98:12
Canada	0.054	0.095	76:01-08:01	0.070	0.076	0.043	0.047	97:01-08:01
Denmark	0.084	0.142	76:01-08:01	0.057	0.068	0.031	0.039	97:01-08:01
France	0.100	0.151	76:01-98:12			0.030	0.034	97:01-98:12
Germany	0.213	0.311	76:01-98:12			0.035	0.037	97:01-98:12
Ireland	0.094	0.180	79:04-98:12			0.141	0.150	97:01-98:12
Italy	0.063	0.171	76:01-98:12			0.062	0.068	97:01-98:12
Japan	0.193	0.240	76:01-08:01	0.055	0.063	0.040	0.043	97:01-08:01
Netherlands	0.234	0.344	76:01-98:12			0.032	0.038	97:01-98:12
Norway	0.093	0.147	76:01-08:01	0.099	0.107	0.072	0.079	97:01-08:01
Portugal	0.375	0.689	76:01-98:12			0.056	0.061	97:01-98:12
Spain	0.140	0.242	76:01-98:12			0.037	0.045	97:01-98:12
Sweden	0.097	0.157	76:01-08:01	0.086	0.097	0.067	0.073	97:01-08:01
Switzerland	0.239	0.389	76:01-08:01	0.083	0.088	0.059	0.063	97:01-08:01
USA/UK	0.054	0.072	76:01-08:01	0.025	0.027	0.026	0.028	97:01-08:01
Euro	0.054	0.056	99:01-08:01	0.054	0.056	0.030	0.032	99:01-08:01
Australia	0.090	0.095	97:01-08:01	0.084	0.089	0.065	0.068	97:01-08:01
New Zealand	0.114	0.125	97:01-08:01	0.100	0.108	0.084	0.092	97:01-08:01
South Africa	0.177	0.194	97:01-08:01	0.182	0.195	0.148	0.162	97:01-08:01

Note: Results are based on daily data, and are expressed in log percent.

TABLE 2

ANNUALIZED PAYOFFS OF THE CARRY-TRADE STRATEGIES

	No Transactions Costs			With Transactions Costs		
	Mean	Standard Deviation	Sharpe Ratio	Mean	Standard Deviation	Sharpe Ratio
British Pound is the Base Currency Jan-1976 to Jan-2008						
Equally-weighted carry trade	0.0321 (0.0079)	0.040 (0.002)	0.811 (0.205)	0.0307 (0.0110)	0.053 (0.003)	0.579 (0.212)
US Dollar is the Base Currency Jan-1997 to Jan-2008						
Equally-weighted carry trade	0.0477 (0.0160)	0.045 (0.003)	1.061 (0.346)	0.0443 (0.0176)	0.051 (0.003)	0.867 (0.343)

Note: Payoffs are measured either in British pounds, per pound bet, or in US dollars, per dollar bet. The carry-trade portfolio is formed as the equally-weighted average of up to 20 individual currency carry trades against either the British pound or the US dollar. The twenty currencies are indicated in Appendix Tables 2 and 3.

TABLE 3

ANNUALIZED PAYOFFS OF INVESTMENT STRATEGIES
 February 1976 to January 2008
 US Dollar is the Base Currency

	Mean	Standard Deviation	Sharpe Ratio	Skewness	Excess Kurtosis	Jarque-Bera Statistic
U.S. stock market	0.0682 (0.0250)	0.148 (0.009)	0.461 (0.181)	-0.758 (0.344)	2.65 (1.54)	149.3 (0.000)
Equally-weighted carry trade	0.0497 (0.0098)	0.051 (0.005)	0.972 (0.228)	-0.664 (0.606)	6.73 (2.30)	753.8 (0.000)
Average of individual-currency carry trade	0.0504	0.109	0.479	-0.259	1.03	31.2

Notes: Payoffs are measured in US dollars, per dollar bet. The payoff at time t to the US stock market is the value-weighted excess return on all US stocks reported in Kenneth French's database, divided by $1 + r_{t-1}$ (this normalizes the excess stock returns to the same size of bet as the carry-trade payoffs). The carry-trade portfolio is formed as the equally-weighted average of up to 20 individual currency carry trades against the US dollar. The individual currencies are indicated in Appendix Table 3. Standard errors are reported in parentheses, except for the Jarque-Bera statistic for which the p-value is reported in parentheses.

TABLE 4

FACTOR BETAS OF THE EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO EXCESS RETURN
1976Q2 to 2007Q4

Factors	Intercept	Beta(s)			R^2
CAPM	0.013 (0.003)	-0.017 (0.033)			0.002
Fama-French factors	0.013 (0.003)	0.016 (0.038)	-0.104 (0.066)	0.010 (0.053)	0.031
C-CAPM	0.015 (0.006)	-0.387 (0.931)			0.003
Extended C-CAPM	0.008 (0.009)	-0.691 (0.978)	0.817 (0.716)	-0.014 (0.035)	0.020
Luxury sales growth	0.013 (0.008)	-0.031 (0.050)			0.008
GDP growth	0.012 (0.003)	0.197 (0.347)			0.003
Fed Funds rate	0.011 (0.006)	0.035 (0.077)			0.002
Term premium	0.014 (0.004)	-0.052 (0.236)			0.001
Liquidity premium	0.014 (0.004)	-0.068 (0.348)			0.000
VIX volatility measure	0.009 (0.012)	0.014 (0.058)			0.001
VXO volatility measure	0.007 (0.009)	0.025 (0.038)			0.003

Notes: The table reports estimates of the equation $R_t^e = a + f_t' \beta + \epsilon_{t+1}$, where R_t^e is the quarterly real excess return of the equally-weighted 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 ($Mkt - Rf$), the Fama-French factors are the $Mkt - Rf$, SMB and HML factors (available from Kenneth French's database), the C-CAPM factor is real per capita consumption growth, the extended C-CAPM factors are real per capita consumption growth, real per capita durables growth, and the return on the value-weighted US stock market, the term premium is the 10 year T-bond rate minus the 3 month T-bill rate, and the liquidity premium is the 3 month eurodollar rate minus the 3 month T-bill rate. Details of the risk factors are provided in Appendix D. Heteroskedasticity-robust standard errors are in parentheses.

TABLE 5

GMM ESTIMATES OF LINEAR FACTOR MODELS

Test Assets are the Fama-French 25 Portfolios and the Equally-Weighted Carry-Trade Portfolio

	First Stage				Second Stage			
	μ	b	λ (%)	R^2	b	λ (%)	R^2	J
CAPM	0.0179 (0.0070)	3.59 (1.46)	2.22 (0.77)	-1.08	3.40 (1.28)	2.10 (0.71)	-1.11	59.52 (0.00)
Fama-French Factors								
<i>Mkt-Rf</i>	0.0179 (0.0070)	5.40 (1.98)	1.73 (0.71)	0.49	5.55 (1.85)	1.95 (0.72)	0.43	57.65 (0.00)
<i>SMB</i>	0.0077 (0.0046)	0.81 (2.00)	0.69 (0.46)		0.37 (1.81)	0.66 (0.46)		
<i>HML</i>	0.0110 (0.0066)	7.09 (2.18)	1.15 (0.56)		6.24 (1.91)	0.84 (0.53)		
C-CAPM	0.0048 (0.0005)	622.80 (680.45)	0.91 (1.06)	-2.78	105.00 (141.24)	0.15 (0.20)	-9.29	27.28 (0.34)
Extended C-CAPM								
<i>Consumption growth</i>	0.0048 (0.0005)	-183.54 (231.39)	-0.34 (0.35)	-0.98	-59.83 (86.36)	-0.10 (0.13)	-7.98	14.45 (0.91)
<i>Durables growth</i>	0.0102 (0.0019)	-137.14 (130.01)	-0.40 (0.32)		-28.31 (74.20)	-0.10 (0.16)		
<i>Market return</i>	0.0223 (0.0070)	3.88 (2.29)	2.20 (0.93)		0.92 (1.88)	0.47 (0.98)		
Luxury sales growth	0.0989 (0.0262)	15.70 (21.59)	14.84 (19.71)	-1.14	-1.01 (3.30)	-0.95 (3.18)	-13.97	16.38 (0.90)
GDP growth	0.0049 (0.0009)	-560.07 (755.43)	-3.16 (4.31)	-3.43	-6.44 (118.86)	-0.04 (0.67)	-11.99	9.66 (1.00)

Table 5 is continued on the next page

TABLE 5 (Continued)

GMM ESTIMATES OF LINEAR FACTOR MODELS

Test Assets are the Fama-French 25 Portfolios and the Equally-Weighted Carry-Trade Portfolio

	First Stage				Second Stage			
	μ	b	λ (%)	R^2	b	λ (%)	R^2	J
Fed Funds rate	0.0652 (0.0226)	-74.35 (155.17)	-9.26 (20.12)	-0.63	1.08 (13.01)	0.14 (1.63)	-12.53	0.98 (1.00)
Term premium	0.0166 (0.0035)	199.99 (193.95)	3.24 (3.10)	-0.25	15.85 (36.08)	0.26 (0.58)	-10.37	2.44 (1.00)
Liquidity premium	0.0087 (0.0033)	-386.25 (738.07)	-2.18 (4.41)	-0.28	12.15 (63.67)	0.07 (0.36)	-12.95	1.41 (1.00)
VIX volatility measure	0.1891 (0.0228)	-22.08 (28.57)	-7.34 (8.43)	-0.19	-1.40 (4.01)	-0.46 (1.32)	-10.25	11.87 (0.99)
VXO volatility measure	0.2034 (0.0209)	-12.37 (14.54)	-6.18 (6.37)	-0.33	-4.10 (3.07)	-2.05 (1.54)	-5.07	37.78 (0.05)
Campbell-Cochrane				-14.71				56.29 (0.00)

Notes: The table reports GMM estimates of the SDF $m_t = 1 - (f_t - \mu)'b$ using the moment conditions $E(R_t^e m_t) = 0$ and $E(f_t - \mu) = 0$, where R_t^e is a 26×1 vector containing the excess returns of the Fama-French 25 portfolios of US stocks sorted on size and the book-to-market value ratio as well as the quarterly real excess return of the equally-weighted carry-trade portfolio, and f_t is a scalar or vector of risk factors. The factors are described in more detail in the footnote to Table 4 and in Appendix D. The first stage of GMM is equivalent to the two-pass regression method of Fama and MacBeth (1973). The GMM procedure is described in more detail in Appendix E. Since $\hat{\mu}$ is the same for both GMM stages, the estimate is reported once. Estimates of the factor risk premia $\hat{\lambda} = \hat{V}_f \hat{b}$ are also reported (in percent), where \hat{V}_f is the sample covariance matrix of f_t . GMM-VARHAC standard errors are reported in parentheses for $\hat{\mu}$, \hat{b} and $\hat{\lambda}$. The table reports the R^2 measure of fit between the sample mean of R_t^e and the predicted mean returns, given by $d_T \hat{b}$, where $d_T = T^{-1} \sum_{t=1}^T R_t^e (f_t' - \hat{\mu})'$. Tests of the overidentifying restrictions are also reported. The test statistic, J , is asymptotically distributed as a χ_{26-k}^2 , where k is the number of risk factors. The p-value is in parentheses. The Campbell-Cochrane model is calibrated, as described in Appendix D, to match the mean equity premium and risk free rate in our sample period. Here we report a direct test of the moment condition $E(R_t^e m_t) = 0$ and the cross-sectional R^2 for the calibrated model. The sample period is 1976Q2 to 2007Q4.

TABLE 6

ANNUALIZED PAYOFFS OF INVESTMENT STRATEGIES
 February 1987 to January 2008
 US Dollar is the Base Currency

	Mean	Standard Deviation	Sharpe Ratio	Skewness	Excess Kurtosis	Jarque-Bera Statistic
U.S. stock market	0.0659 (0.0298)	0.147 (0.013)	0.450 (0.228)	-1.158 (0.435)	3.84 (2.22)	211.3 (0.000)
Equally-weighted carry trade	0.0322 (0.0123)	0.060 (0.004)	0.538 (0.218)	-0.672 (0.155)	1.09 (0.44)	31.5 (0.000)
Hedged, equally-weighted carry trade	0.0251 (0.0079)	0.036 (0.002)	0.707 (0.211)	0.751 (0.145)	0.43 (0.43)	25.6 (0.000)

Notes: Payoffs are measured in US dollars, per dollar bet. The payoff at time t to the US stock market is the value-weighted excess return on all US stocks reported in Kenneth French's database, divided by $1 + r_{t-1}$. The carry-trade portfolio is formed as the equally-weighted average of up to six individual currency carry trades against the US dollar. The individual currencies are the Australian dollar, the Canadian dollar, the Japanese yen, the Swiss franc, the British pound, and the euro. The hedged carry-trade portfolio combines the forward market positions with an options contract that insures against losses from the forward position (details are provided in the main text). Standard errors are in parentheses, except for the Jarque-Bera statistic for which the p-value is reported in parentheses.

TABLE 7

FACTOR BETAS OF THE HEDGED EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO EXCESS RETURN
1987Q2 to 2007Q4

Factor	Intercept	Beta(s)		R^2
CAPM	0.012 (0.002)	0.027 (0.030)		0.013
Fama-French factors	0.011 (0.002)	0.065 (0.036)	-0.014 (0.043) 0.088 (0.040)	0.078
C-CAPM	0.015 (0.003)	-0.552 (0.507)		0.009
Extended C-CAPM	0.017 (0.007)	-0.487 (0.569)	-0.191 (0.644) 0.023 (0.030)	0.021
Luxury sales growth	0.015 (0.004)	-0.031 (0.026)		0.024
GDP growth	0.017 (0.003)	-0.937 (0.358)		0.065
Fed Funds rate	0.021 (0.005)	-0.163 (0.100)		0.034
Term premium	0.008 (0.003)	0.268 (0.199)		0.028
Liquidity premium	0.020 (0.004)	-1.271 (0.539)		0.068
VIX volatility measure	0.012 (0.006)	0.012 (0.029)		0.001
VXO volatility measure	0.016 (0.006)	-0.017 (0.031)		0.004

Notes: The table reports estimates of the equation $R_t^e = a + f_t'\beta + \epsilon_{t+1}$, where R_t^e is the quarterly real excess return of the hedged equally-weighted carry-trade portfolio and f_t is a scalar or vector of risk factors (see the footnotes to Tables 4 and 6). The CAPM factor is the excess return on the value-weighted US stock market ($Mkt - Rf$), the Fama-French factors are the $Mkt - Rf$, SMB and HML factors (available from Kenneth French's database), the C-CAPM factor is real per capita consumption growth, the extended C-CAPM factors are real per capita consumption growth, real per capita durables growth, and the return on the value-weighted US stock market, the term premium is the 10 year T-bond rate minus the 3 month T-bill rate, and the liquidity premium is the 3 month eurodollar rate minus the 3 month T-bill rate. Details of the risk factors are provided in Appendix D. Heteroskedasticity-robust standard errors are in parentheses.

TABLE 8

GMM ESTIMATES OF LINEAR FACTOR MODELS

Test Assets are the Fama-French 25 Portfolios and the Hedged Equally-Weighted Carry-Trade Portfolio

	First Stage				Second Stage			
	μ	b	λ (%)	R^2	b	λ (%)	R^2	J
CAPM	0.0173 (0.0088)	3.07 (1.85)	1.91 (0.96)	-1.37	2.77 (1.63)	1.72 (0.88)	-1.45	56.81 (0.00)
Fama-French Factors								
<i>Mkt-Rf</i>	0.0173 (0.0088)	5.01 (2.45)	1.59 (0.89)	0.26	6.41 (2.21)	1.82 (0.88)	0.11	55.47 (0.00)
<i>SMB</i>	0.0026 (0.0060)	-0.81 (2.40)	0.31 (0.61)		-2.98 (2.14)	-0.14 (0.59)		
<i>HML</i>	0.0101 (0.0091)	5.88 (2.40)	1.10 (0.72)		6.92 (1.92)	1.30 (0.74)		
C-CAPM	0.0045 (0.0004)	677.51 (1118.70)	0.68 (1.13)	-7.11	184.31 (158.84)	0.19 (0.16)	-8.67	37.55 (0.05)
Extended C-CAPM								
<i>Consumption growth</i>	0.0045 (0.0004)	-12.58 (217.49)	-0.11 (0.25)	-1.05	18.76 (95.54)	0.03 (0.10)	-6.71	5.24 (1.00)
<i>Durables growth</i>	0.0103 (0.0025)	-242.84 (267.82)	-0.39 (0.42)		14.32 (76.41)	0.03 (0.12)		
<i>Market return</i>	0.0208 (0.0088)	1.59 (2.84)	2.10 (1.65)		0.73 (2.11)	0.40 (1.30)		
Luxury sales growth	0.0967 (0.0265)	17.66 (29.67)	16.48 (26.95)	-1.31	-0.41 (4.14)	-0.38 (3.89)	-10.39	13.35 (0.97)
GDP growth	0.0046 (0.0010)	-53.42 (138.83)	-0.14 (0.36)	-10.00	-8.20 (43.54)	-0.02 (0.11)	-10.03	54.18 (0.00)

Table 8 is continued on the next page.

TABLE 8 (Continued)

GMM ESTIMATES OF LINEAR FACTOR MODELS

Test Assets are the Fama-French 25 Portfolios and the Hedged Equally-Weighted Carry-Trade Portfolio

	First Stage				Second Stage			
	μ	b	λ (%)	R^2	b	λ (%)	R^2	J
Fed Funds rate	0.0486 (0.0082)	-67.34 (72.12)	-3.05 (3.10)	-2.96	-4.67 (12.53)	-0.21 (0.56)	-9.10	3.80 (1.00)
Term premium	0.0169 (0.0038)	141.89 (129.75)	1.93 (1.79)	-6.11	18.69 (23.41)	0.25 (0.32)	-9.08	4.90 (1.00)
Liquidity premium	0.0054 (0.0023)	-325.70 (489.33)	-0.48 (0.56)	-1.13	13.82 (56.30)	0.02 (0.09)	-10.82	6.93 (1.00)
VIX volatility measure	0.1891 (0.0228)	-22.08 (28.57)	-7.34 (8.43)	-0.31	-0.30 (3.94)	-0.10 (1.31)	-12.07	12.63 (0.98)
VXO volatility measure	0.2033 (0.0220)	-11.31 (13.35)	-5.91 (6.14)	-0.30	-3.68 (3.05)	-1.92 (1.60)	-4.74	38.38 (0.04)
Campbell-Cochrane				-11.74				49.65 (0.00)

Notes: The table reports GMM estimates of the SDF $m_t = 1 - (f_t - \mu)'b$ using the moment conditions $E(R_t^e m_t) = 0$ and $E(f_t - \mu) = 0$, where R_t^e is a 26×1 vector containing the excess returns of the Fama-French 25 portfolios of US stocks sorted on size and the book-to-market value ratio as well as the quarterly real excess return of the hedged equally-weighted carry-trade portfolio (see the note to Table 6), and f_t is a scalar or vector of risk factors. The factors are described in more detail in the footnote to Table 4 and in Appendix D. The first stage of GMM is equivalent to the two-pass regression method of Fama and MacBeth (1973). The GMM procedure is described in more detail in Appendix E. Since $\hat{\mu}$ is the same for both GMM stages, the estimate is reported once. Estimates of the factor risk premia $\hat{\lambda} = \hat{V}_f \hat{b}$ are also reported (in percent), where \hat{V}_f is the sample covariance matrix of f_t . GMM-VARHAC standard errors are reported in parentheses for $\hat{\mu}$, \hat{b} and $\hat{\lambda}$. The table reports the R^2 measure of fit between the sample mean of R_t^e and the predicted mean returns, given by $d_T \hat{b}$, where $d_T = T^{-1} \sum_{t=1}^T R_t^e (f_t - \hat{\mu})'$. Tests of the overidentifying restrictions are also reported. The test statistic, J , is asymptotically distributed as a χ_{26-k}^2 , where k is the number of risk factors. The p-value is in parentheses. The Campbell-Cochrane model is calibrated, as described in Appendix D, to match the mean equity premium and risk free rate in our sample period. Here we report a direct test of the moment condition $E(R_t^e m_t) = 0$ and the cross-sectional R^2 for the calibrated model. The sample period is 1987Q2 to 2007Q4.

TABLE 9

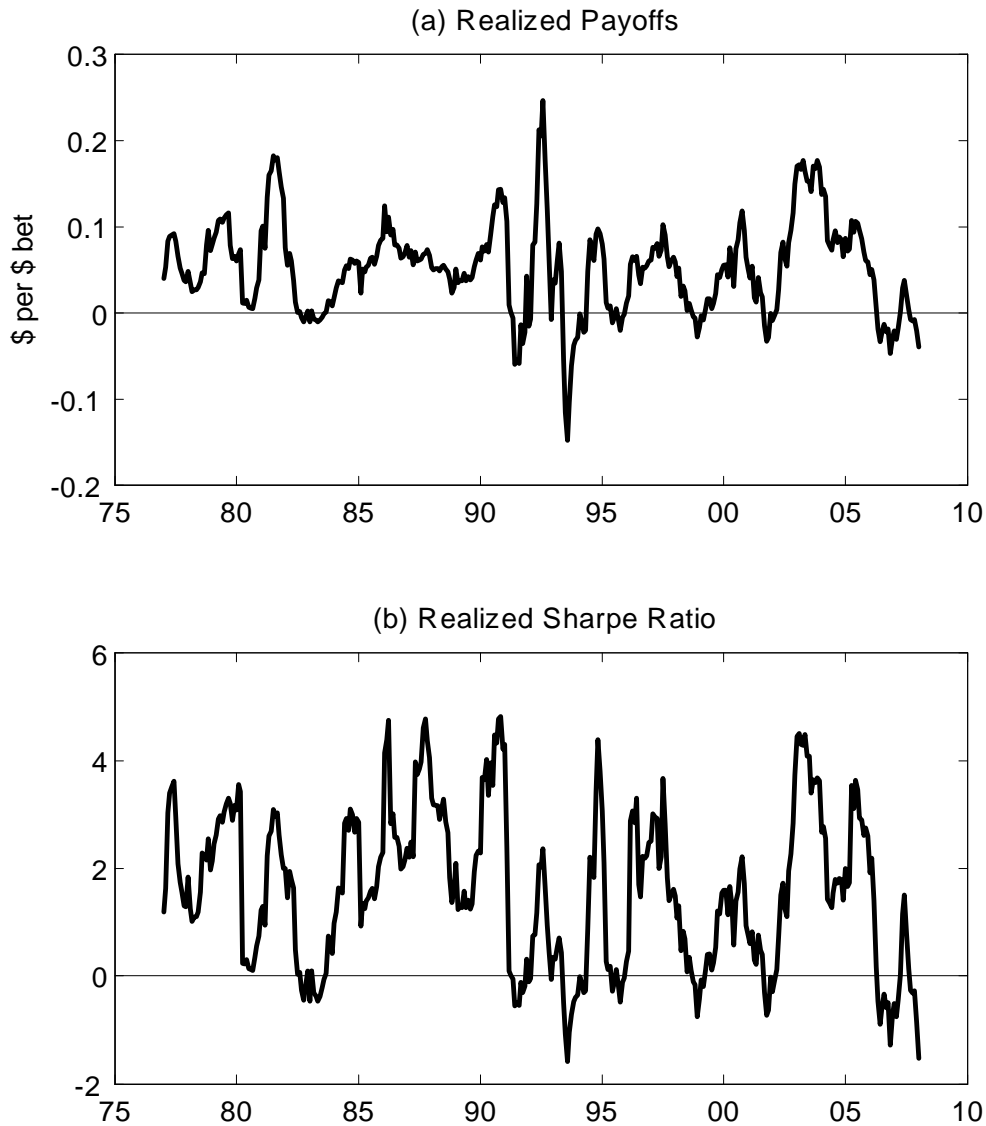
ESTIMATES OF MOMENTS USED AS INPUTS IN PESO EVENT CALCULATIONS

	CME Data	J.P. Morgan Data		Stock Market Data
	(1987–2007)	(1996–2007)		(1986–2007)
		6 currencies	11 currencies	
$E^G(z)$	0.0027 (0.0010)	0.0036 (0.0011)	0.0047 (0.0013)	
$E^G(z^H)$	0.0021 (0.0007)	0.0020 (0.0007)	0.0029 (0.0015)	
$E^G(x)$				0.0082 (0.0027)
$E^G(x^H)$				−0.0021 (0.0017)
$E^G(Mx)$				0.0194 (0.0027)
$E^G(Mx^H)$				0.0038 (0.0020)
$E^G(h)$	−0.0120 (0.0002)	−0.0100 (0.0003)	−0.0099 (0.0004)	
$E^G(d - r)$				−0.0023 (0.0003)
$E^G[c^x(1 + r)]$				0.0224 (0.0012)
z'	−0.0154 (0.0033)	−0.0181 (0.0039)	−0.0162 (0.0060)	
x'				−0.126 (0.052)
$M'/E^G(M)$	121.7 (38.0)			107.8 (56.3)

Notes: The variables z and z^H are, respectively, the payoffs to the carry trade, and the hedged carry trade in non-peso states. The variables x and x^H are, respectively, the excess returns to the equity strategy (defined as the S&P 100 index), and the hedged equity strategy (defined in the main text) in non-peso states. The variables z' and x' are, respectively, the payoffs to the carry trade and the equity strategy in the peso state. The variables M and M' are, respectively, the stochastic discount factor in non-peso and peso states. The variable h is the minimum payoff to the hedged carry trade strategy. The variable d is the dividend-yield of the S&P 100 index, r is the one-month eurodollar rate and c^x is the ex-ante cost of the option used to hedge the return of the index. The operator E^G is the unconditional expectations operator that applies to non-peso states of the world. The CME, J.P. Morgan, and stock market data are described in Appendix C.

FIGURE 1: ANNUALIZED REALIZED AVERAGE PAYOFFS AND SHARPE RATIOS OF THE EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO

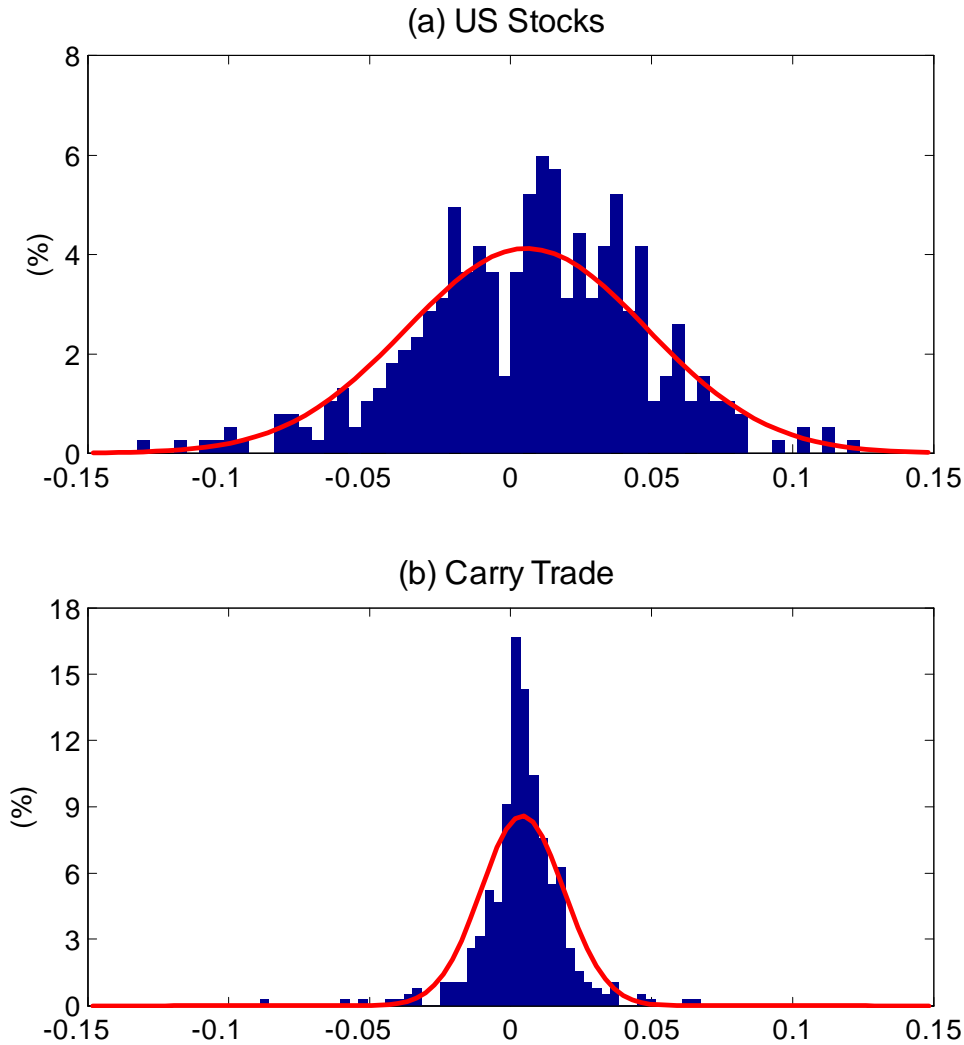
12-Month Rolling Window, February 1977–January 2008



Note: Plot (a) shows the annualized average payoff from month $t - 11$ to month t , in US dollars, per dollar bet in the carry trade. Plot (b) shows the ratio of the annualized average payoff, to the annualized standard deviation of the payoff, both being measured from month $t - 11$ to month t . The carry-trade portfolio is formed as the equally-weighted average of up to 20 individual currency carry trades against the US dollar.

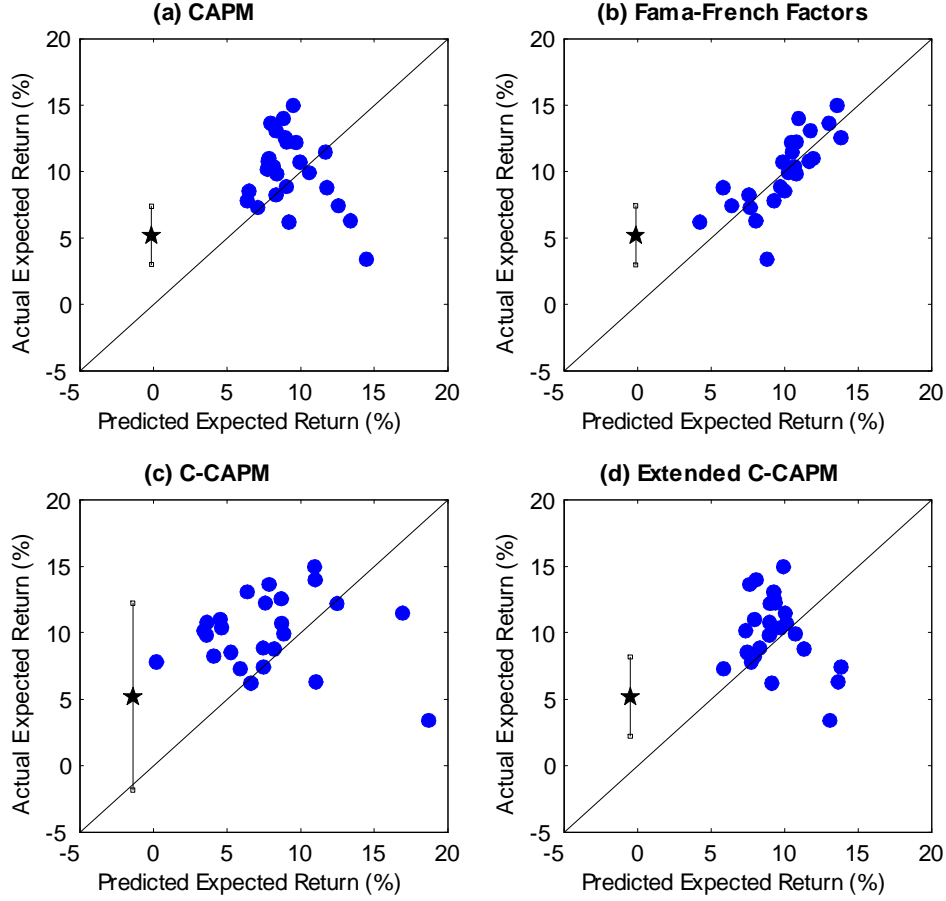
FIGURE 2: SAMPLING DISTRIBUTIONS OF THE EXCESS RETURNS OF THE EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO AND THE VALUE-WEIGHTED US STOCK MARKET

February 1976–January 2008



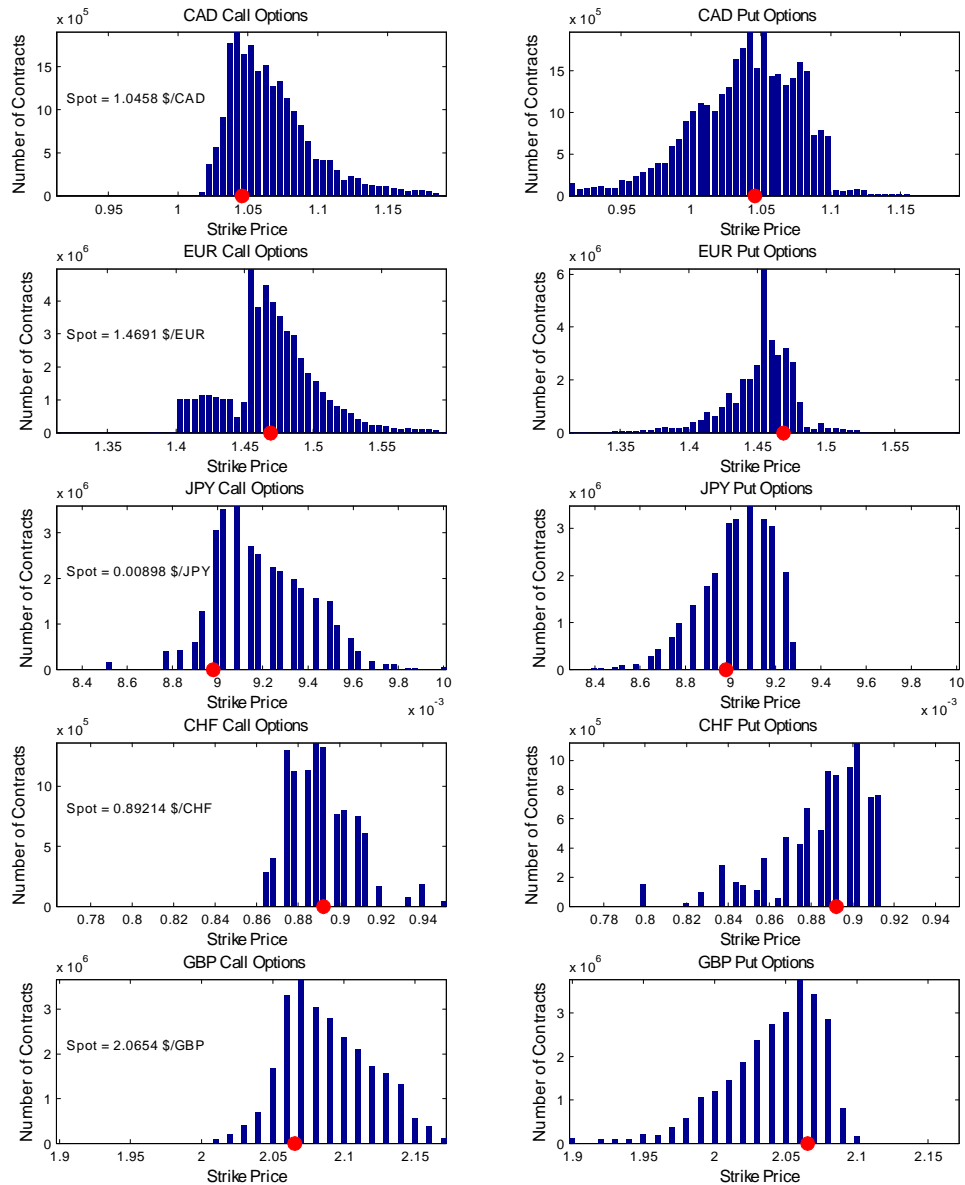
Note: In each plot the red line indicates the histogram implied by a normal distribution with the same mean and standard deviation as in the sampling distribution. The excess returns are computed at the monthly frequency. US stock excess returns are for the value-weighted US stock market from the Fama-French database. The carry-trade portfolio is formed as the equally-weighted average of up to 20 individual currency carry trades against the US dollar. Excess returns to the carry trade are payoffs scaled by $1 + r_t$.

FIGURE 3: CROSS-SECTIONAL FIT OF FACTOR MODELS ESTIMATED BY GMM
 Test Assets are the Fama-French 25 Portfolios & the Equally-Weighted Carry-Trade
 Portfolio



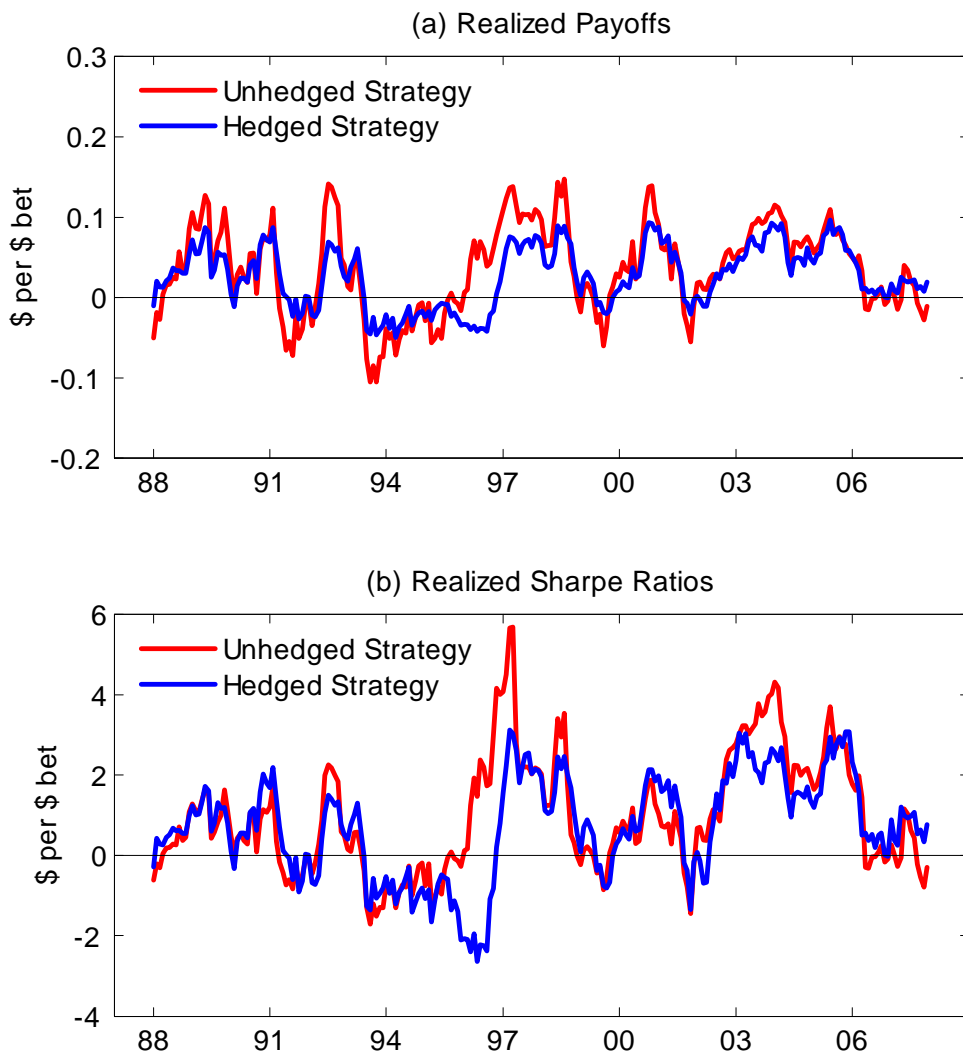
Note: In each case the parameters μ and b in the SDF $m_t = 1 - (f_t - \mu)' b$ are estimated by GMM using the method described in the text. The risk factors, f_t , are indicated by the title of each plot with details provided in the main text. The predicted expected return is $(1/T) \sum_{t=1}^T R_{it}^e (f_t - \hat{\mu})' \hat{b}$ for each portfolio's excess return, R_{it}^e . The actual expected return is $\bar{R}_i^e = (1/T) \sum_{t=1}^T R_{it}^e$. The blue dots correspond to Fama and French's 25 portfolios sorted on the basis of book-to-market value and firm size. The black star represents the carry-trade portfolio formed as the equally-weighted average of up to 20 individual currency carry trades against the US dollar. The black vertical line extending above and below the star is the actual expected return plus a two-standard error band for the pricing error of the carry-trade portfolio. When it does not cross the 45 degree line, the pricing error is statistically significant at the 5 percent level. Sample period is 1976Q2–2007Q4, and expected returns are annualized.

FIGURE 4: THE VOLUME OF CALLS AND PUTS AND MONEYNESS
November 14, 2007



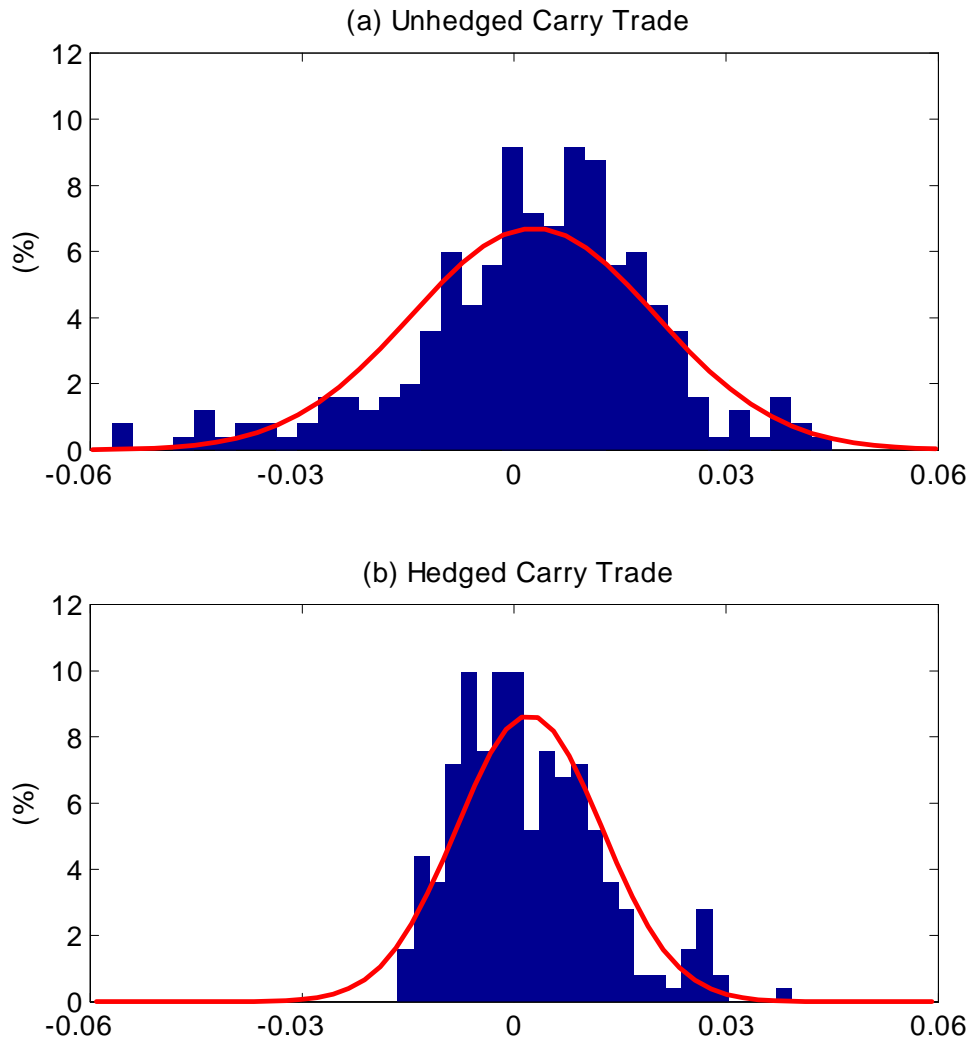
Note: Each plot indicates the number of contracts traded at different strike prices on Nov. 14 2007 for five currencies: the Canadian dollar (CAD), the Euro (EUR), the Japanese yen (JPY), the Swiss franc (CHF) and the British pound (GBP). The closing spot price of each currency is indicated by the red dot. In this plot currencies are quoted as USD/FCU. Source: the Chicago Mercantile Exchange.

FIGURE 5: ANNUALIZED REALIZED AVERAGE PAYOFFS AND SHARPE RATIOS OF THE EQUALLY-WEIGHTED HEDGED AND UNHEDGED CARRY-TRADE PORTFOLIOS
12-Month Rolling Window, February 1987–January 2008



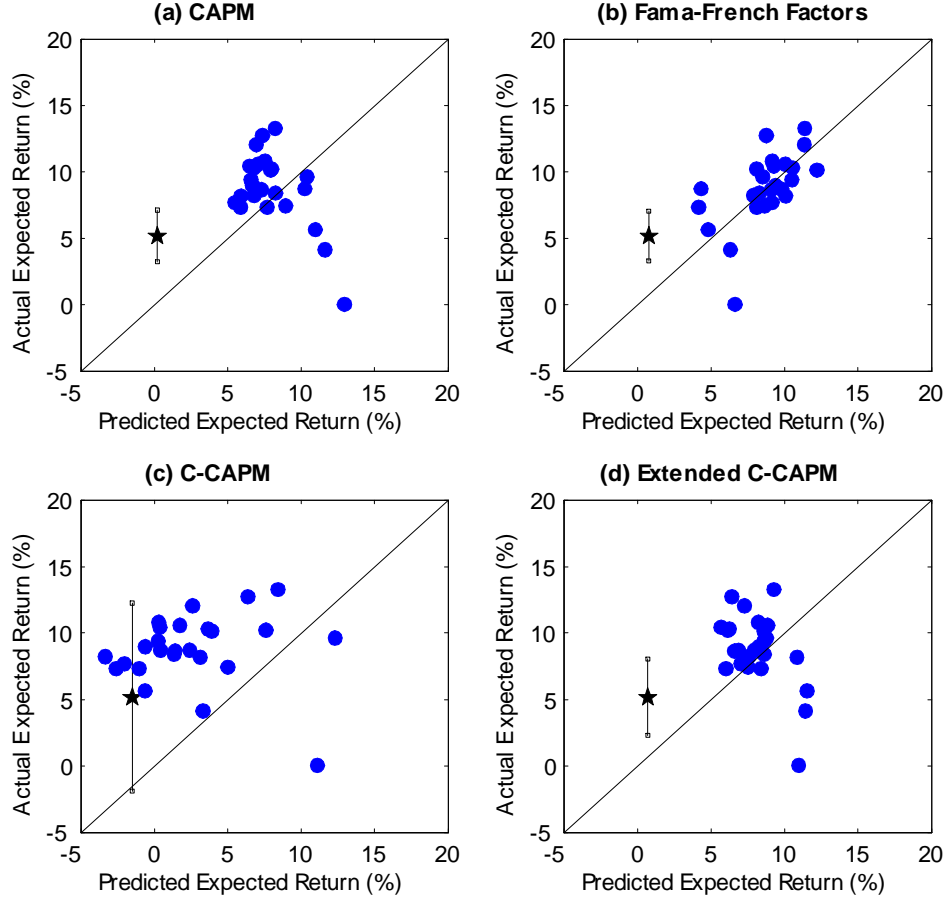
Note: Plot (a) shows the annualized average payoff from month $t - 11$ to month t , in US dollars, per dollar bet in the carry trade. Plot (b) shows the ratio of the annualized average payoff, to the annualized standard deviation of the payoff, both being measured from month $t - 11$ to month t . The unhedged portfolio is the equally-weighted carry-trade portfolio, described in the main text, formed by taking positions in the forward market currency-by-currency. The hedged position is formed by combining the forward position on each currency in the unhedged portfolio with a near-the-money option that insures against possible losses from the forward position. The carry-trade portfolios are formed as the equally-weighted averages of up to six individual currency carry trades against the US dollar.

FIGURE 6: SAMPLING DISTRIBUTIONS OF THE PAYOFFS OF THE EQUALLY-WEIGHTED
 CARRY-TRADE PORTFOLIOS
 February 1987–January 2008



Note: In each plot the red line indicates the histogram implied by a normal distribution with the same mean and standard deviation as in the sampling distribution. The excess returns are computed at the monthly frequency. The carry-trade portfolios are formed as the equally-weighted average of up to six individual currency carry trades against the US dollar. The unhedged portfolio is formed by taking positions in the forward market currency-by-currency. The hedged position is formed by combining the forward position on each currency in the unhedged portfolio with a near-the-money option that insures against possible losses from the forward position.

FIGURE 7: CROSS-SECTIONAL FIT OF FACTOR MODELS ESTIMATED BY GMM
 Test Assets are the Fama-French 25 Portfolios & the Equally-Weighted Hedged
 Carry-Trade Portfolio



Note: In each case the parameters μ and b in the SDF $m_t = 1 - (f_t - \mu)' b$ are estimated by GMM using the method described in the text. The predicted expected return is $(1/T) \sum_{t=1}^T R_{it}^e (f_t - \hat{\mu})' \hat{b}$ for each portfolio's excess return, R_{it}^e . The actual expected return is $\bar{R}_i^e = (1/T) \sum_{t=1}^T R_{it}^e$. The blue dots correspond to Fama and French's 25 portfolios sorted on the basis of book-to-market value and firm size. The black star represents the hedged carry-trade portfolio formed as the equally-weighted average of up to six individual currency carry trades against the US dollar. The hedged position is formed by combining the forward position on each currency in the unhedged portfolio with a near-the-money option that insures against possible losses from the forward position. The black vertical line extending above and below the star is the actual expected return plus a two-standard error band for the pricing error of the carry-trade portfolio. When it does not cross the 45 degree line the pricing error is statistically significant at the 5 percent level. Sample period is 1987Q2-2007Q4, and expected returns are annualized.

A: Spot and Forward Exchange Rate Data

Our foreign exchange rate data are obtained from Datastream. They are originally sourced by Datastream from the WM Company/Reuters. We use two data sets. The first data set consists of spot exchange rates and one month forward exchange rates for twenty currencies quoted against the British pound. This data set spans the period January 1976 to January 2008. The mnemonics for and availability of each currency are indicated in Table A5. With the exception of euro forward quotes, each exchange rate is quoted as foreign currency units (FCUs) per British pound (GBP). To obtain quotes in GBP/FCU we inverted the original quotes while swapping the bid and ask prices (except for the Euro forward quotes). The original data set includes observations on all weekdays. We sample the data on the last weekday of each month.

The second data set consists of spot exchange rates and one month forward exchange rates for twenty currencies quoted against the U.S. dollar. This data set spans the period December 1996 to January 2008. The mnemonics for and availability of each currency are indicated in Table A6. With the exception of the Irish punt, British pound, Euro (forwards only), Australian dollar, and New Zealand dollar, each exchange rate is quoted as foreign currency units (FCUs) per U.S. dollar (USD). To obtain USD/FCU quotes for the other currencies we inverted the original quotes while swapping the bid and ask prices. We also noticed a problem in the original Datastream data set: the bid and ask spot exchange rates for the Euro are reversed for all data available through 12/29/2006. We reversed the quotes to obtain the correct bid and ask rates. The original data set includes observations on all weekdays. We sample the data on the last weekday of each month.

When we ignore bid-ask spreads we obtain a data set running from January 1976 to January 2008 with all currencies quoted against the U.S. dollar. We convert pound quotes to dollar quotes by multiplying the GBP/FCU quotes by the USD/GBP quotes.

B: Interest Rate Data and CIP

Our eurocurrency interest rate data are obtained from Datastream. They are originally sourced by Datastream from the Financial Times and ICAP. The data set spans the period January 1976 to January 2008. The mnemonics for and availability of each interest rate is indicated in Table A7. The original data set includes observations on all weekdays. We

sample the data on the last weekday of each month.

To assess whether CIP holds it is critical to take bid-ask spreads into account. In this appendix the variables r_t^a and r_t^b denote the ask and bid interest rate in the domestic currency. The variables r_t^{*a} and r_t^{*b} denote the ask and bid foreign-currency interest rates.

In the presence of bid-ask spreads CIP is given by the following inequalities,

$$\pi_{CIP} = (1 + r_t^{*b}) \frac{F_t^b}{S_t^a} - (1 + r_t^a) \leq 0, \quad (40)$$

$$\pi_{CIP}^* = (1 + r_t^b) \frac{S_t^b}{F_t^a} - (1 + r_t^{*a}) \leq 0. \quad (41)$$

Equation (40) implies that there is a non-positive payoff (π_{CIP}) to the “borrowing domestic currency covered strategy.” This strategy consists of borrowing one unit of domestic currency, exchanging it for foreign currency at the spot rate, investing the proceeds at the foreign interest rate, and converting the payoff into domestic currency at the forward rate. Equation (41) implies that there is a non-positive payoff (π_{CIP}^*) to the “borrowing foreign currency covered strategy.” This strategy consists of borrowing one unit of foreign currency, exchanging the foreign currency into domestic currency at the spot rate, investing the proceeds at the domestic interest rate, and converting the payoff into foreign currency at the forward rate. Table A8 reports statistics for π_{CIP} and π_{CIP}^* for sixteen currencies.

Table A8 indicates that for all sixteen currencies, the median value for π_{CIP} and π_{CIP}^* is negative. Also the fraction of periods in which π_{CIP} and π_{CIP}^* are positive is small. Even in periods where the payoff is positive, the median payoff is very small.

Our finding that deviations from CIP are small and rare is consistent with the results in Taylor (1987) who uses data collected at 10-minute intervals for a three-day period, Taylor (1989) who uses daily data for selected historical periods of market turbulence, and Clinton (1988) who uses daily data from November 1985 to May 1986.

C: Options Data and Options-Based Strategies

Our first source of options data is the Chicago Mercantile Exchange (CME). We obtained daily quotes for put and call options for six currencies against the U.S. dollar. The currencies are available beginning on the following dates: Australian dollar (January 1994), Canadian dollar (August 1986), Euro (January 1999), Japanese yen (May 1986), Swiss franc (May 1985), British pound (January 1991). The data are available through the end of 2007. Due

to sparse coverage in the early part of the sample we begin our analysis no earlier than January 1987.

We use the following notation: the spot exchange rate (S), the one month forward exchange rate (F), the strike price on the closest to in-the-money call option on the dollar (K^C), the strike price on the closest to in-the-money put option on the dollar (K^P), the settlement price of the call option (C), the settlement price of the put option (P), and the one month eurodollar deposit rate, r . We obtained the eurodollar deposit rate from the Federal Reserve Board interest rate database (H.15). Since the CME data pertain to options on foreign currency, in what follows, the variables S , F , K^C and K^P are measured in USD/FCU, while the variables C and P are measured in USD per foreign currency unit transacted.

Since our analysis of the carry trade is done at the monthly frequency using one month forward exchange rates, we restrict attention to options that are one month from maturity. Since we work exclusively with options expiring mid month (on the Friday preceding the third Wednesday) we look for transactions taking place 30 days prior to expiration. To be concrete, take January 2007 as an example of an expiration date. The Friday preceding the third Wednesday is January 12th 2007. We therefore look for transactions involving options expiring on January 12th 2007 that took place on December 13th 2006 as these dates are 30 days apart. For the purpose of calculating payoffs we measure S_t , F_t , K_t^C , K_t^P , C_t , P_t and r_t on December 13th 2006. We measure S_{t+1} as the spot rate observed on January 12 2007.³¹

Our second source of options data is J.P. Morgan. We obtained daily one-month at-the-money implied volatility quotes, and spot exchange rates, for eleven currencies against the U.S. dollar. These data are available from January 1996 until January 2008 for the following currencies: Australian dollar, Canadian dollar, Danish krone, Euro (January 1998), Japanese yen, Swiss franc, British pound, New Zealand dollar, Norwegian krone, Swedish krone and South African rand. We convert the implied volatility quotes to option prices using the Black-Scholes formula in combination with forward premia calculated using the data described in Appendix A. We use the same transactions dates as for the CME data. The implied volatilities in the two data sets are very similar.

³¹Notice that this means one month's S_{t+1} is not necessarily the next month's S_t . For example, the February 2007 expiration date is February 16th 2007. So the transactions date we look for in January 2007 is January 17th not January 12th. In practice we ignore the fact that this timing creates some slightly overlapping months and some gaps, putting priority on matching maturities of forwards and options.

D: Details of the Risk-Factor Analysis

Defining Quarterly Real Returns The monthly payoffs to the carry trade, denoted generically here as z_t , were defined for trades where $1/F_t$ FCUs were either bought or sold forward. This is equivalent to selling or buying one dollar. It is useful, instead, to normalize the number of dollars sold or bought to $1 + r_{t-1}$, where r_{t-1} is the yield on a one-month Treasury bill at the time when the currency bet is made. That is, we define the monthly excess return

$$R_t^{e,m} = (1 + r_{t-1})z_t.$$

To see that $R_t^{e,m}$ can be interpreted as an excess return, consider the case where we buy foreign currency forward, so: $z_t = S_t/F_{t-1} - 1$. This value of z_t implies that $R_t^{e,m} = (1+r_{t-1})(S_t/F_{t-1} - 1)$. Assuming that CIP (equation (??)) holds, $R_t^{e,m} = (1+r_{t-1}^*)S_t/S_{t-1} - (1 + r_{t-1})$. So, when $(1 + r_{t-1})/F_{t-1}$ FCUs are bought forward $R_t^{e,m}$ is the equivalent to the excess return, in dollars, from taking a long position in foreign T-bills.

Let t index months, and let $s = t/3$ be the equivalent index for quarters. To convert the monthly excess return to a quarterly excess return we define:

$$R_s^{e,q} = \Pi_{j=0}^2(1 + r_{t-1-j} + R_{t-j}^{e,m}) - \Pi_{j=0}^2(1 + r_{t-1-j}).$$

This expression corresponds to the appropriate excess return because it implies that the agent continuously re-invests in the carry trade strategy. In month t he bets his accumulated funds from currency speculation times $1 + r_t$. To define the quarterly *real* excess return in quarter s , which we denote R_s^e , notice that this is simply:

$$R_s^e = \frac{R_s^{e,q}}{1 + \pi_s}$$

where π_s is the inflation rate between quarter $s - 1$ and quarter s .

To generate the returns we use the risk free rate data from Kenneth French's data library: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. These data correspond to the one-month Treasury bill rate from Ibbotson Associates (2006).

We convert nominal returns to real returns using the inflation rate corresponding to the deflator for consumption of nondurables and services found in the U.S. National Income and Product Accounts.

When we work with options data, the returns for the first quarter are the accumulated payoffs (as described above) realized mid-January, mid-February and mid-March. For the second, third and fourth quarters we use the analogous monthly payoffs.

Data Sources for Risk Factors and Other Variables The three Fama-French factors are from Kenneth French’s data library. The three factors are Mkt-Rf (the market premium, which we also use to define the CAPM factor), SMB (the size premium) and HML (the book to market premium). Each of these objects is an excess return. Nominal returns are converted to real returns as described above for our currency strategies.

Real per-capita consumption growth is from the U.S. National Income and Product Accounts which can be found at the website of the Bureau of Economic Analysis (BEA): www.bea.gov. We define real consumption growth as the weighted average of the growth rates of nondurables consumption and services consumption. The weights are the nominal shares of nondurables and services in their sum. We compute the growth rate of the population using the series provided by the BEA in the NIPA accounts. This series displays seasonal variation so we first pass it through the Census X12 filter available from the Bureau of Labor Statistics (www.bls.gov). The inflation series used in all our calculations is the weighted average of the inflation rates for nondurables and services with the weights defined as above.

The risk factors proposed by Yogo (2006) are the market return (Mkt-Rf plus the risk free rate), the real growth rate of per-capita consumption of nondurables and services, and the real growth rate of the per-capita service flow from the stock of consumer durables. To estimate the latter we proceeded as follows. Annual end-of-year real stocks of consumer durables are available from the U.S. National Income and Product Accounts, as are quarterly data on purchases of durables by consumers. Within each year we determine the depreciation rate that makes the quarterly purchases consistent with the annual stocks, and use this rate to interpolate quarterly stocks using the identity: $K_{t+1}^D = C_t^D + (1 - \delta^D)K_t^D$. Here K_t^D is the beginning of period t stock of consumer durables, C_t^D is purchases of durables, and δ^D is the depreciation rate. We assume that the service flow from durables is proportional to the stock of durables.

Real luxury retail sales growth is available from 1987Q1–2001Q4 and is obtained from Aït-Sahalia, Parker and Yogo (2004).

The quarterly index of industrial production is from the Federal Reserve Board of Governors (www.federalreserve.gov), Statistical Release Table G.17. We calculate the growth rate of this series.

The average monthly value of the Fed funds rate is from the Federal Reserve Board of

Governors (www.federalreserve.gov), Statistical Release Table H.15 (Selected Interest Rates), Effective Federal Funds Rate (mnemonic FEDFUNDS). We convert this to the quarterly frequency using the average of the three monthly values within each quarter.

The monetary policy shock is from Altig, et.al. (2004). Their estimates of the shock were updated through the end of 2007 by extending the data set. See Altig, et.al. (2004) for details of the underlying data.

Seasonally-adjusted monthly data on the stocks of M1, M2 and MZM are from the Federal Reserve Board of Governors (www.federalreserve.gov), Statistical Release Table H.6 (Money Stock Measures), (mnemonics M1SL, M2SL and MZMSL). We compute quarterly growth rates by taking the growth rate from the 3rd month of the previous quarter to the 3rd month of the current quarter.

The term premium is defined as the difference between the 10-year T-bond rate and the 3-month Treasury-bill rate. Data are from the Federal Reserve Board of Governors (www.federalreserve.gov), Statistical Release Table H.15 (Selected Interest Rates) for the 3-Month Treasury Bill Secondary Market Rate (mnemonic TB3MS) and the 10-Year Treasury Constant Maturity Rate (mnemonic GS10). We convert this to the quarterly frequency using the average of the three monthly values within each quarter.

The liquidity premium is defined as the difference between the 3-month eurodollar rate and the 3-month Treasury-bill rate. Data are from the Federal Reserve Board of Governors (www.federalreserve.gov), Statistical Release Table H.15 (Selected Interest Rates) for the 3-Month eurodollar rate (mnemonic EDM3). We convert this to the quarterly frequency using the average of the three monthly values within each quarter.

The VIX and VXO volatility measures were obtained at the daily frequency from Datastream (mnemonics CBOEVIX, available from February 1990, and CBOEVXO, available from February 1986). We convert these to the quarterly frequency by averaging across all daily observations within each quarter.

The Campbell-Cochrane SDF is constructed using the same consumption series for non-durables and services described above, and denoted here as C_t . The SDF is

$$m_t = \delta [S_t C_t / (S_{t-1} C_{t-1})]^{-\gamma}$$

where $s_t = \ln S_t$ is constructed recursively as follows:

$$\begin{aligned} s_t &= (1 - \phi)\bar{s} + \phi s_{t-1} + \lambda_t(\Delta \ln C_t - g) \\ \lambda_t &= \begin{cases} \sqrt{1 - 2(s_{t-1} - \bar{s})/e^{\bar{s}}} - 1 & \text{if } s_{t-1} < s_{\max} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

We calibrate the model parameters to the following values: $g = 0.0049$ (the average quarterly growth rate of real per capita consumption), $\sigma = 0.0052$ (the standard deviation of the quarterly growth rate of real per capita consumption), $\gamma = 2.88$, $\phi = 0.8766$, and $r_f = 0.0044$.

The remaining parameters are determined as

$$\begin{aligned} \bar{s} &= \ln[\sigma\sqrt{\gamma/(1-\phi)}] \\ s_{\max} &= \bar{s} + (1 - e^{2\bar{s}})/2 \\ \delta &= \exp(\gamma g - \gamma(1 - \phi)/2 - r_f). \end{aligned}$$

With these parameter values the model matches the average quarterly equity premium and real risk free rate in our sample, 1976Q2-2007Q4.

E: GMM Estimation

Generically we use GMM to estimate the linear factor model $m_t = 1 - (f_t - \mu)'b$ using the moment restrictions:

$$E(R_t^e m_t) = 0 \quad E(f_t) = \mu \quad (42)$$

where R_t^e is an $n \times 1$ vector of excess returns and f_t is a $k \times 1$ vector of risk factors. Define $u_{1t}(b, \mu) = R_t^e m_t = R_t^e [1 - (f_t - \mu)'b]$ and let $g_{1T}(b, \mu) = \frac{1}{T} \sum_{t=1}^T u_{1t} = \bar{R}^e - (D_T - \bar{R}^e \mu') b$ where $D_T = \frac{1}{T} \sum_{t=1}^T R_t^e f_t'$ and $\bar{R}^e = \frac{1}{T} \sum_{t=1}^T R_t^e$. Define $u_{2t}(\mu) = f_t - \mu$ and let $g_{2T}(\mu) = \frac{1}{T} \sum_{t=1}^T u_{2t} = \bar{f} - \mu$ and $\bar{f} = \frac{1}{T} \sum_{t=1}^T f_t$. Define $u_t = (u_{1t} \quad u_{2t})'$ and $g_T = (g_{1T} \quad g_{2T})'$. We consider GMM estimators that set $a_T g_T = 0$, where a_T is a $2k \times (n + k)$ matrix and takes the form

$$a_T = \begin{pmatrix} d_T' W_T & 0 \\ 0 & I_k \end{pmatrix}, \quad (43)$$

where $d_T = D_T - \bar{R}^e \bar{f}'$, and W_T is an $n \times n$ positive definite weighting matrix. It follows that the GMM estimators of b and μ are

$$\hat{b} = (d_T' W_T d_T)^{-1} d_T' W_T \bar{R}^e \quad (44)$$

$$\hat{\mu} = \bar{f}. \quad (45)$$

We consider two-stage GMM estimators. In the first stage $W_T = I_n$. In the second stage, $W_T = (P_T S_T P_T')^{-1}$ where $P_T = (I_n \quad \bar{R}^e \hat{b}')$ and S_T is a consistent estimator of $S_0 = \sum_{j=-\infty}^{+\infty} E(u_t u_{t-j}')$. Because u_{2t} may be serially correlated we use a VARHAC estimator, described in Burnside (2007), to compute S_T .

Let

$$\delta_T = \begin{pmatrix} -d_T & \bar{R}^e \hat{b}' \\ 0 & -I_k \end{pmatrix}. \quad (46)$$

A test of the pricing errors is based on

$$J = T g_T(\hat{b}, \hat{\mu}) (\hat{V}_g)^+ g_T(\hat{b}, \hat{\mu}), \quad (47)$$

where the + sign indicates the generalized inverse and

$$\hat{V}_g = A_T S_T A_T' \text{ with } A_T = I_{n+k} - \delta_T (a_T \delta_T)^{-1} a_T. \quad (48)$$

Equation (42) and the definition of m_t imply that

$$E(R_t^e) = E [R_t^e (f_t - \mu)'] b. \quad (49)$$

Corresponding to the right-hand side of (49) is a vector of predicted expected returns, $\hat{R}^e = d_T \hat{b}$. The cross-sectional R^2 measure is:

$$R^2 = 1 - \frac{(\bar{R}^e - d_T \hat{b})' (\bar{R}^e - d_T \hat{b})}{(\bar{R}^e - \tilde{R}^e)' (\bar{R}^e - \tilde{R}^e)}. \quad (50)$$

where $\tilde{R}^e = \frac{1}{n} \sum_{i=1}^n \bar{R}_i^e$ is the cross-sectional average of the mean returns in the data.

Equation (49) can be rewritten as

$$E(R_t^e) = E \underbrace{[R_t^e (f_t - \mu)']}_{\beta} \underbrace{V_f^{-1} V_f b}_{\lambda}. \quad (51)$$

The covariance matrix of f_t is estimated by GMM using the moment restriction

$$E [(f_t - \mu)(f_t - \mu)' - V_f] = 0.$$

An estimate of λ is given by $\hat{\lambda} = \hat{V}_f \hat{b}$ where \hat{V}_f is the sample covariance matrix of the factors. Standard errors for $\hat{\lambda}$ are obtained by the delta method using the joint distribution of \hat{b} , $\hat{\mu}$ and \hat{V}_f . The details are discussed in Burnside (2007).

APPENDIX REFERENCES

- Altig, David, Lawrence J. Christiano, Martin Eichenbaum and Jesper Linde (2004) “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” mimeo, Northwestern University.
- Burnside, Craig (2007) “Empirical Asset Pricing and Statistical Power in the Presence of Weak Risk Factors,” National Bureau of Economic Research Working Paper No. 13357.
- Clinton, Kevin (1988) “Transactions Costs and Covered Interest Arbitrage: Theory and Evidence,” *Journal of Political Economy* 96 (2, Apr.), 358–70.
- Cochrane, John H. *Asset Pricing*, Princeton University Press, 2001.
- Ibbotson and Associates (2006) *Stocks, Bonds, Bills and Inflation Yearbook*. Chicago, Ill.: Ibbotson Associates.
- Taylor, Mark P. (1987) “Covered Interest Parity: A High-Frequency, High-Quality Data Study,” *Economica* 54 (216, Nov.), 429–38.
- Taylor, Mark P. (1989) “Covered Interest Arbitrage and Market Turbulence,” *Economic Journal* 99 (396, June), 376–91.