

NBER WORKING PAPER SERIES

DO PESO PROBLEMS EXPLAIN THE RETURNS TO THE CARRY TRADE?

A. Craig Burnside
Martin S. Eichenbaum
Isaac Kleshchelski
Sergio Rebelo

Working Paper 14054
<http://www.nber.org/papers/w14054>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2008

This paper is a substantially revised version of NBER Working Paper 12489 titled "The Returns to Currency Speculation." We thank the editor, Geert Bekaert, two anonymous referees, John Cochrane, John Heaton, Jakub Jurek, and Ravi Jagannathan for very useful comments. We also thank 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). The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by A. Craig Burnside, Martin S. Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Do Peso Problems Explain the Returns to the Carry Trade?

A. Craig Burnside, Martin S. Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo

NBER Working Paper No. 14054

June 2008, Revised September 2010

JEL No. F31

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 argue that these payoffs reflect a peso problem. The underlying peso event features high values of the stochastic discount factor rather than very large negative payoffs.

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

Martin S. Eichenbaum
Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60208
and NBER
eich@northwestern.edu

Isaac Kleshchelski
Olin Business School
Washington University in St. Louis
1 Brookings Drive
St. Louis, MO 63130-4899
KLESHCHELSKI@WUSTL.EDU

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and NBER
s-rebelo@northwestern.edu

1 Introduction

The forward exchange rate is a biased forecaster of the future spot exchange rate. This fact is often referred to as the ‘forward-premium puzzle.’ We study the properties of a widely-used currency speculation strategy, known as the carry trade, that exploits this puzzle. This strategy 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 explanation for the high average payoff to the carry trade is that it compensates 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 unconditional 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 (e.g. Bekaert and Hodrick (1992) and Backus, Foresi, and Telmer (2001)). More generally, it has been difficult to use structural asset-pricing models to rationalize the risk-premium movements required to account for the time-series properties of the forward premium (see Bekaert (1996)).

The most natural alternative explanation for the high average payoff to the carry trade is that it reflects the presence of a peso problem. We use the term ‘peso problem’ to refer to the effects on inference caused by low-probability events that do not occur in sample.¹ Peso

¹In related work Farhi and Gabaix (2008) and Brunnermeier, Nagel, and Pedersen (2008) emphasize the importance of rare in-sample events. There are some moderately large negative payoffs to the unhedged carry trade in our sample. While these payoffs affect the profitability of the strategy, the average risk-adjusted

problems can in principle explain the positive average payoff 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. So, the average risk-adjusted payoff to the carry trade is positive in non-peso states.

In this paper we address the question of whether or not the large average payoff to the unhedged carry trade is compensation for peso-event risk. Our approach to answering this question 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’ is immune to large losses such as those potentially associated with a peso event.

We use data on currency options to estimate the average risk-adjusted payoff to the hedged carry trade. We find that this payoff is smaller than the corresponding payoff to the unhedged carry trade. This finding is consistent with the view that the average payoff to the unhedged carry trade reflects a peso problem. Given this result, an obvious question is: what is the nature of the peso event for which agents are being compensated? 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.² The second possibility is that the salient feature of a peso state is a large value of the stochastic discount factor (SDF). A key contribution of this paper is to assess the relative importance of these two possibilities. We find that a

payoff to the carry trade is positive.

²See Bates (1996a) for a related argument that exchange rate jumps can explain the volatility smile present in currency options.

peso event reflects high values of the SDF in the peso state rather than very large negative payoffs to the unhedged carry trade in that state.

The intuition for why the losses to the unhedged carry trade are small in the peso state is as follows. Any risk-adjusted payoffs associated with the carry trade in the non-peso states must on average be compensated, on a risk-adjusted basis, for losses in the peso state. According to our estimates, the average risk-adjusted payoffs of the hedged and unhedged carry trade in the non-peso states are not very different. Consequently, the risk-adjusted losses to these two strategies in the peso state cannot be very different. Since 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. By construction there is an upper bound to the losses of the hedged carry trade. This upper bound tell us 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 states 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 states 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 of the unhedged carry trade 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, we bring additional data to bear on the plausibility of our estimates. We consider two versions of an equity strategy that involve 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 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 risk-adjusted payoff to the two stock market strategies we generate an independent estimate of the value of the SDF in the peso state. Remarkably, this estimate is similar to the estimate of the peso state SDF that rationalizes the average risk-adjusted payoffs to the carry trade. We interpret these findings as supportive of the view that the positive average payoff to the unhedged carry trade reflects the presence of a peso problem.

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. In Section 6 we report our estimates of the payoffs to the carry trade and stock market strategies in the peso state, and our estimates of the SDF in the peso state. We also generalize the analysis to multiple peso states. Section 7 concludes.

2 Peso problems and the carry trade

The failure of uncovered interest rate parity (UIP) motivates a variety of speculative strategies. 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:

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

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. We normalize the amount of dollars bet on this strategy to one. The amount of dollars borrowed, y_t , is given by:

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

Consider the case in which S_t is a martingale:

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

where E_t denotes the time- t conditional expectations operator. This martingale property is not an implication of market efficiency, but it is a reasonable description of the data. Equation (3) implies that the expected payoff to the carry trade is positive and equal to the difference between the higher and the lower interest rate:

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

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 w_t , the number of FCUs sold forward, is given by:

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

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

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

Covered interest rate parity implies that:

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

When equation (6) holds, the payoffs to the strategies defined by equations (4) and (2) 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 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. Equation (7) implies that:

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

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 risk resulting from 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.³ 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 use the following notation. Let $\omega_t \in \Omega$ denote the state of the world at time t , let $z(\omega_t)$ denote the payoff to the carry trade in state ω_t , and $M(\omega_t)$ denote the value of the SDF in state ω_t . We partition Ω , the set of possible states, ω_t , into two sets. The first set, Ω^N , consists of those values of ω_t corresponding to non-peso events. The second set, Ω^P , consists of those values of ω_t corresponding to a peso event. For simplicity, we assume that for all $\omega_t \in \Omega^P$, $z(\omega_t) = z' < 0$ and $M(\omega_t) = M'$.

We denote by $\mathcal{F}^N(\omega_{t+1})$ the unconditional distribution of ω_{t+1} in non-peso states. For future reference we define $\mathcal{F}^N(\omega_{t+1}|\omega_t)$ as the conditional distribution of ω_{t+1} given ω_t , where both ω_{t+1} and ω_t are in Ω^N . To simplify, we assume that the conditional and unconditional

³See Villanueva (2007) for additional evidence on this point. Lustig and Verdelhan (2007) argue that a consumption-based SDF 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 consumption-based SDFs, 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.

probability of the peso state is p . The unconditional version of equation (7) is:

$$(1 - p)E^N(Mz) + pM'z' = 0, \quad (9)$$

where $E^N(\cdot)$ denotes the expectation over non-peso states,

$$E^N(Mz) = \int_{\Omega^N} M(\omega_{t+1})z(\omega_{t+1})d\mathcal{F}^N(\omega_{t+1}).$$

Since z' is negative, equation (9) implies that the average risk-adjusted payoff over non-peso states, $E^N(Mz)$, is positive. This observation captures the conventional view that a peso problem can rationalize positive average payoffs to the carry trade.

We focus on whether or not the existence of peso events provides 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 large losses when a peso event occurs. We call this strategy the ‘hedged carry trade.’ Below we describe this strategy in detail.

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

$$z_{t+1}^C = \max(0, S_{t+1} - K_t) - C_t(1 + r_t). \quad (10)$$

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 to the put option in dollars, net of the option price is:

$$z_{t+1}^P = \max(0, K_t - S_{t+1}) - P_t(1 + r_t). \quad (11)$$

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 the agent buys, at time t , a call option on the foreign currency with a strike price K_t . So, whenever $S_{t+1} > K_t$, the agent buys FCUs at the price K_t . In this case the minimum payoff to the hedged carry trade is:

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

Similarly, suppose that an agent buys one FCU forward. Then, the worst case scenario in the standard carry trade arises when there is a large depreciation of the foreign currency. In this state of the world the agent sells the foreign currency that he receives from the forward contract at a low value of S_{t+1} . Suppose that the agent buys, at time t , a put option on the foreign currency with a strike price K_t . So, whenever $S_{t+1} < K_t$, the agent sells FCUs at a price K_t . In this case the minimum payoff to the hedged carry trade is:

$$(S_{t+1} - F_t) + (K_t - S_{t+1}) - P_t(1 + r_t) = K_t - F_t - P_t(1 + r_t). \quad (13)$$

In order to normalize the size of the bet to one dollar, we set the amount of FCUs traded equal to $1/F_t$. 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.

The dollar payoff to this strategy is:

$$z_{t+1}^H = \begin{cases} z_{t+1} + z_{t+1}^C/F_t & \text{if } F_t \geq S_t, \\ z_{t+1} + z_{t+1}^P/F_t & \text{if } F_t < S_t, \end{cases} \quad (14)$$

where z_{t+1} , z_{t+1}^C , and z_{t+1}^P are the carry-trade payoffs defined in equations (5), (10) and (11), respectively.

An alternative way to implement the hedged carry trade is to use only options, 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,

$$(C_t - P_t)(1 + r_t) = F_t - K_t, \quad (15)$$

it is easy to show that this strategy for hedging the carry trade is equivalent to the one described above as long as the strike price of the options is the same in the two strategies.

The minimum payoff to the hedged carry trade, h_{t+1} , is negative. To see this property, use the put-call-forward parity condition, (15), and equations (12) and (13) to write the minimum payoffs as follows:

$$h_{t+1} = \begin{cases} -P_t(1 + r_t)/F_t & \text{if } F_t \geq S_t, \\ -C_t(1 + r_t)/F_t & \text{if } F_t < S_t. \end{cases} \quad (16)$$

Since option prices are positive, h_{t+1} is negative.

We summarize the realized payoffs to the hedged carry trade as follows:

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

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

Using options to assess the effect of peso problems Since the hedged carry trade is a zero net-investment strategy, the payoff, $z^H(\omega_{t+1})$, must satisfy:

$$(1 - p) \int_{\Omega^N} M(\omega_{t+1}) z^H(\omega_{t+1}) d\mathcal{F}^N(\omega_{t+1}|\omega_t) + pM'h(\omega_t) = 0.$$

Taking expectations over all non-peso states we obtain:

$$(1 - p)E^N(Mz^H) + pM'E^N(h) = 0. \quad (17)$$

Using equation (17) to solve for pM' and replacing this term in equation (9), we obtain:

$$z' = E^N(h) \frac{E^N(Mz)}{E^N(Mz^H)}. \quad (18)$$

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 or M' . Given our estimate of z' and a value of p we can use equation (9) to estimate M' ,

$$M' = \frac{(1-p)E^N(Mz)}{p(-z')}. \quad (19)$$

There are two possible outcomes of these calculations. The first possible outcome is that our estimate of z' is a large negative value, consistent with the conventional view about the payoffs to the carry trade in the peso state. The second possible outcome is that our estimate of z' is a small negative value. In this case a peso event can still explain the average payoff to the carry trade but only if M' is large. So, in this case, 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' is empirically plausible. To answer this question we consider an equity strategy whose payoff is also potentially affected by the peso event. Using hedged and unhedged versions of this strategy we obtain an alternative estimate of M' . We then assess whether this estimate of M' 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.⁴ 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 is given by:

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

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 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 equity strategy is the dividend yield of the 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

⁴The choice of this index is driven by data considerations.

these states the payoff to the hedged stock strategy is $d_t - r_t - c_t^x (1 + r_t)$. In summary, the payoff to the hedged equity strategy net of the options cost is given by:

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

The payoff to the unhedged equity strategy in the peso state is x' . The payoffs to the unhedged and hedged equity strategies must satisfy:

$$(1 - p)E^N(Mx) + pM'x' = 0. \quad (20)$$

$$(1 - p)E^N(Mx^H) + pM'E^N[d - r - c^x(1 + r)] = 0. \quad (21)$$

Combining these two equations we obtain:

$$x' = E^N[d - r - c^x(1 + r)] \frac{E^N(Mx)}{E^N(Mx^H)}. \quad (22)$$

$$M' = \frac{(1 - p)E^N(Mx)}{p(-x')}. \quad (23)$$

Given estimates of $E^N(Mx)$ and $E^N(Mx^H)$ we can use equations (22) and (23) to estimate M' and x' . Given a value of p we then estimate x' and M' . A test of the second interpretation of the peso event is whether the value of M' that emerges from this procedure is consistent with the value of M' implied by equation (19).

The procedure just described assumes that a peso-based explanation of the positive average payoff to the carry trade is required. After describing the data we motivate this assumption in Sections 4 and 5, at least in the context of linear asset pricing models. In Section 6 we report the results of implementing the procedure described in this section.

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 span the period January 1976 to July 2009. 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 July 2009. 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 January 1976 to July 2009, with the exact sample period varying by currency.

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. We provide details of our interest rate data and analysis of CIP in the online appendix (www.duke.edu/~acb8/bekr_app.pdf).

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 January 1987 to April 2009 on the prices of put and call options against U.S. dollar futures for the Australian dollar, the Canadian dollar, the euro, the Japanese yen, the Swiss franc, and the British pound. Appendix B specifies the exact period of availability for each currency. When a futures contract and an option contract expire on the same date, an option on currency futures is equivalent to an option on the spot exchange rate. This equivalence results from the fact that the price of a futures contract coincides with the spot exchange rate at maturity. In practice the expiration dates of the two contracts do not generally coincide in our data set. In Appendix B we describe how to modify our hedging strategy to take this fact into account. This modification involves adjusting both the number of options purchased and their strike price. Our modified hedging strategy is exactly equivalent to the hedging strategy described in Section 2 whenever interest rates are constant over the period between the expiration of the option and the expiration of the futures contract.

Since we compute carry-trade payoffs at the monthly frequency, we use data on options that are one month from maturity (see Appendix B for details). We work exclusively with options expiring near the beginning of each month (two Fridays prior to 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 $t + 1$ futures price on the option expiration date. To compute net payoffs we multiply option prices by the gross 30-day eurodollar interest rate obtained from the Federal Reserve Board. This 30-day interest rate is matched to the maturity of the options in our data set.

Our second options data set is from J.P. Morgan. These data consist of daily observations of one-month at-the-money-forward implied volatility quotes and forward 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, and the Swedish krone. Our sample period is February 1995 to July 2009. We convert the implied volatility quotes to option prices using the Black-Scholes formula. Given the structure of this data set, we select monthly trades that mature on a date near the end of the month and are initiated 30 days earlier. We select the last maturity date which is a business day and for which the initiation date is a business day.

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 the payoff distribution has fat tails. Third, we study the covariance between the payoffs to the carry trade and various risk factors using both time series and panel data.

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) foreign currency 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, w_t , is given by:

$$w_t = \begin{cases} +1/F_t^b & \text{if } F_t^b/S_t^a > 1, \\ -1/F_t^a & \text{if } F_t^a/S_t^b < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

The payoff to this strategy is:

$$z_{t+1} = \begin{cases} w_t (F_t^b - S_{t+1}^a) & \text{if } w_t > 0, \\ w_t (F_t^a - S_{t+1}^b) & \text{if } w_t < 0, \\ 0 & \text{if } w_t = 0. \end{cases} \quad (25)$$

4.1 Characteristics 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.⁵ This portfolio is constructed by betting $1/n_t$ U.S. dollars 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. Also, we report all statistics on an annualized basis. Table 1 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 using a much longer time series than would be the case if we used only 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 0.748. Taking bid-ask spreads into account reduces the Sharpe ratio to 0.507. 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 0.865. Taking bid-ask spreads into account reduces the Sharpe ratio to 0.694. 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.

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

⁵In the online appendix we report results for individual currencies.

paper we abstract from transaction costs and work with spot and forward rates that are the average of bid and ask rates.⁶ Given this decision we can work with the longer data set (from January 1976 to July 2009) using the U.S. dollar as the home currency.

Table 2 reports statistics for the payoffs to the equally-weighted carry trade and summary statistics for the individual-currency carry trades. We compute the latter 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 0.442, while the Sharpe ratio for an equally-weighted portfolio of currencies is 0.911. This large difference between the Sharpe ratios is due to the fact that the standard deviation of the payoffs is much lower for the equally-weighted portfolio. Second, the Sharpe ratio of the carry trade is substantially larger than that of the U.S. stock market (0.911 versus 0.373). While the average excess return to the U.S. stock market is larger than the payoff to the carry trade (0.058 versus 0.048), the returns to the U.S. stock market are much more volatile than the payoffs to the carry trade (0.156 versus 0.053).

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. Table 2 reports skewness and excess kurtosis statistics, as well as the results of the Jarque-Bera normality tests. It is evident that the distributions of both payoffs are leptokurtic, exhibiting fat tails.

4.2 The impact of the financial crisis on carry-trade payoffs

Panel A of Figure 1 shows the cumulative payoffs to investing one dollar in January 1976 in three different strategies. The first strategy involves investing in one-month Treasury bills. The second strategy involves investing in a value-weighted index of the universe of

⁶In 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.

U.S. stocks in the CRSP database. In both strategies the monthly proceeds are reinvested. The third strategy is the carry trade. Since this strategy involves zero net investment we compute the cumulative payoffs as follows. We initially deposit one U.S. dollar in a bank account that yields the same rate of return as the Treasury bill rate. In the beginning of every period we bet the balance of the bank account on the carry trade strategy. At the end of the period, payoffs to the carry trade are deposited into the bank account.

Three features of Figure 1 are worth noting. First, the cumulative payoffs to the carry trade and stock market strategies are very similar. Second, the payoffs to the carry trade are much less volatile than those of the U.S. stock market. These two features account for the Sharpe ratio of the carry-trade strategy being roughly 2.5 times higher than that of the U.S. stock market. Third, in the recent financial crisis the carry trade strategy loses money, but these losses are much smaller than those of the U.S. stock market. The U.S. stock market cumulative return peaked at \$44.32 in October 2007 and fell to a trough of \$21.47 in February 2009, a decline of 51.6 percent. The carry trade portfolio cumulative return peaked at \$31.22 in July 2008 and fell to a trough of \$27.87 in February 2009, a decline of 10.7 percent. Both the cumulative payoffs to the carry trade and the U.S. stock market strategies have partially recovered from their trough values.

The worst monthly payoffs (i.e. the largest drawdowns) to the carry trade from February 1976 to July 2009 are: -8.9 percent (March 1991), -5.8 percent (October 1992), and -5.1 percent (June 1993). The three worst monthly payoffs to the carry trade from July 2008 to July 2009 are: -4.2 percent (September 2008), -3.9 percent (August 2008), and -3.7 percent (January 2009).⁷ The three worst monthly payoffs to the stock market strategy in our sample are -23.0 (October 1987), -18.5 percent (October 2008), and -16.1 (August 1998). Note that the largest drawdowns of the carry trade strategy did not occur during the recent financial crisis. In contrast, one of the three worst payoffs to the stock market strategy did occur during the recent financial crisis.

⁷Because of data limitations we cannot compute the drawdowns on the carry trade strategy at a daily frequency.

It is worth emphasizing that, while there are some reasonably large negative payoffs to the carry trade in sample, the average payoff is still positive. We now turn to the question of whether the payoffs are correlated with traditional risk factors.

4.3 Risk factor analysis of carry-trade payoffs

In this subsection we show that the covariances between the payoffs to the carry trade and traditional risk factors are not statistically different from zero. We do so using both time-series and panel-data analysis. We consider data at both the monthly and quarterly frequencies. When data on the risk factors are available at the monthly frequency, we define a 26×1 vector R_t containing the time- t nominal payoff to the carry-trade strategy and the nominal excess returns of the 25 Fama-French (1993) portfolios of equities sorted by firm size and the ratio of book value to market value. When data on the risk factors are available at the quarterly frequency, we define a 26×1 vector R_t containing the time- t real quarterly payoff to the carry-trade strategy and the 25 Fama-French (1993) portfolios.⁸ These payoffs or excess returns must satisfy:

$$E_t(R_{t+1}m_{t+1}) = 0, \quad (26)$$

where, when the data are monthly, m_{t+1} is the SDF that prices nominal USD-denominated excess returns and, when the data are quarterly, m_{t+1} is the SDF that prices real USD-denominated excess returns. We consider linear SDFs of the form:

$$m_t = \xi [1 - (f_t - \mu)' b]. \quad (27)$$

Here ξ is a scalar, f_t is a vector of risk factors, $\mu = E(f_t)$, and b is a conformable vector. To simplify our analysis we abstract from non-negativity constraints on m_t (see Li, Xu and Zhang (2010) for a discussion of the potential importance of this issue).

It follows from equation (26) and the law of iterated expectations that:

$$E(R_t m_t) = 0. \quad (28)$$

⁸In Appendix C we show how we convert monthly payoffs to real quarterly excess returns.

Equations (28) and (27) imply that:

$$E(R_t) = \beta\lambda,$$

where

$$\begin{aligned}\beta &= \text{cov}(R_t, f_t')V_f^{-1}, \\ \lambda &= V_f b.\end{aligned}\tag{29}$$

Here V_f is the covariance matrix of the factors, β is a measure of the systematic risk of the payoffs, and λ is a vector of risk premia. Note that β is the population value of the regression coefficient of R_t on f_t . Our time-series analysis focuses on estimating the betas of the carry trade payoffs for different candidate risk factors. Our panel analysis provides complementary evidence on the importance of different risk factors by estimating alternative SDF models using the moment condition (28). One of these models is the Fama-French (1993) model that we later use to estimate M' .

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 Aït-Sahalia, Parker and Yogo (2004)), GDP growth, industrial production 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), the Pastor and Stambaugh (2003) liquidity measures, and three measures of volatility: the VIX, the VXO (the

implied volatility of the S&P 500 and S&P 100 index options, respectively, calculated by the Chicago Board Options Exchange) and the innovation to the VXO.

To conserve space this paper reports results for the first four risk models listed above. These results are representative of the results obtained using the broader list of factors (see the online appendix). Table 3 reports the estimated regression coefficients of the different risk-factors. Panels A and B of Table 3 report results for factors that are available at a monthly frequency and quarterly frequency, respectively. Our key finding is that none of the risk factors covaries significantly with the payoffs to the carry trade. Recall that the average payoff to the carry trade is statistically different from zero (see Table 2). Factors that have zero β s clearly cannot account for these positive average payoffs.

Panel risk-factor analysis We now discuss the results of estimating the parameters of SDF models built using the monthly and quarterly risk factors detailed in Table 3. In estimating the models we impose the null hypothesis that there are no peso events. We use the estimated SDF models to generate model-predicted expected payoffs to the carry-trade strategy and the 25 Fama-French portfolios.⁹ We then study how well the model explains the average payoff to the carry trade, as well as its overall ability to explain the cross-section of average excess returns used in the estimation procedure.

We estimate b and μ by the generalized method of moments (GMM) using equation (28) and the moment condition $\mu = E(f_t)$. The first stage of the GMM procedure, which uses the identity matrix to weight the GMM errors, is equivalent to the two-pass regression method commonly associated with Fama and MacBeth (1973). Because, at this stage, we are imposing the null of no peso events, all payoffs should reflect covariance with risk factors, so the first stage of the GMM procedure is equivalent to a cross-sectional regression with no constant. The second stage of the GMM procedure uses an optimal weighting matrix. We provide details of the GMM procedure in the online appendix.

⁹The online appendix also reports results for a quarterly version of the Campbell and Cochrane (1999) SDF. Verdelhan (2010) argues that open-economy models in which agents have Campbell-Cochrane (1999) preferences can generate non-trivial deviations from UIP.

It is evident from equations (27) and (28) that $\xi = E(m_t)$ is not identified. Fortunately, the point estimate of b and inference about the model’s overidentifying restrictions are invariant to the value of ξ , so we set ξ to one for convenience. It follows from equations (27) and (28) that:

$$E(R_t) = -\frac{\text{cov}(R_t, m_t)}{E(m_t)} = E[R_t(f_t - \mu)'b]. \quad (30)$$

Given an estimate of b , the predicted mean excess return is the sample analogue of the right-hand side of equation (30), which we denote by \hat{R} . The actual mean excess return is the sample analogue of the left-hand side of equation (30), which we denote by \bar{R} . We denote by \tilde{R} the average across the elements of \bar{R} . We evaluate different SDF models using the R^2 between the predicted and actual mean excess returns. This R^2 measure is invariant to the value of ξ and is given by:

$$R^2 = 1 - \frac{(\bar{R} - \hat{R})'(\bar{R} - \hat{R})}{(\bar{R} - \tilde{R})'(\bar{R} - \tilde{R})}.$$

For each monthly risk factor, or vector of factors, Table 4 reports estimates of b , the R^2 , and the value of Hansen’s (1982) J statistic used to test the overidentifying restrictions implied by equation (28). In addition, we report the alpha of the carry trade portfolio, i.e. the average payoff that is not priced by the risk factor.

The results fall into two categories, depending on whether the model is rejected based on the test of the overidentifying restrictions. For nine out of the 16 cases that we consider, the model is rejected, and the alpha of the carry trade is statistically significant. For the other seven cases the model is not rejected. However, in these cases the b parameters are estimated with enormous imprecision and the R^2 is negative.

To conserve space, here we report results for the four risk models for which we presented estimated betas. We present the additional results in the online appendix. From Table 4, Panel A we see that the CAPM and the Fama-French model are rejected at the one percent level. In addition, the alpha of the carry trade portfolio is statistically significant for these models.

We present results for the C-CAPM and the extended C-CAPM in Table 4, Panel B.

The b parameters of these models are estimated with enormous imprecision. Since there is little information in the sample about the b parameters it is hard to statistically reject these factor models. The R^2 for both of these models is negative.

We now provide an alternative perspective on the performance of the SDF models being considered. Figure 2 plots \hat{R} , the predictions of these models for $E(R_t)$ against \bar{R} , the sample average of R_t . The circles pertain to the Fama-French portfolios, and the star pertains to the carry trade. It is clear that the three CAPM models (panels a, c and d) do a poor job of explaining the excess returns to the Fama-French portfolios and the payoffs to the carry trade. Not surprisingly, the Fama-French model (panel b) does a reasonably good job of pricing the excess returns to the Fama-French portfolios. However, the model greatly understates the average payoffs to the carry trade. The annualized average payoff to the carry trade is 4.82 percent. The Fama-French model predicts that this average return should equal 0.19 percent. The solid line through the star is a two-standard-error band for the difference between the data and model average payoff, i.e. the pricing error. Clearly, we can reject the hypothesis that the model accounts for the average payoffs to the carry trade, i.e. from the perspective of the model the carry trade has a positive alpha.

In sum, this section provides evidence that it is difficult to explain the positive average payoff to the carry trade as compensation for exposure to conventional risk factors, at least in sample.

5 Payoffs to the hedged carry trade

In this section we analyze the empirical properties of the hedged carry trade. As discussed in Section 3, our primary options data set from the CME covers six currencies and a shorter sample period (February 1987 to April 2009) than our data set on forward contracts. For comparability we also compute the payoffs to the unhedged carry trade using the currencies and sample period covered by the CME data set.

We implement the hedged carry trade using strike prices that are close to ‘at the money’

(see Appendix B for details). We choose these strike prices because most of the options traded are actually close to being at the money (see the online appendix). Options that are way out of the money tend to be sparsely traded and relatively expensive.¹⁰

Table 5 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. The average payoff to the hedged carry trade is lower than that of the carry trade (1.58 versus 2.96 percent). Both the average payoff and the Sharpe ratio of the hedged carry trade are statistically different from zero.¹¹

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 annualized average payoff to the unhedged carry trade executed with the U.S. dollar as the home currency by 9 basis points on an annual basis. To assess the impact of transactions costs on the hedged carry trade we compute average bid-ask spreads for puts and calls against the Canadian dollar, the euro, the Japanese yen, and the Swiss franc. We base our estimates on 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. The average bid-ask spread in these data is 5.2 percent of the option price.¹² This estimate is slightly higher than the point estimate of 4.4 percent obtained by Chong, Ding, and Tan (2003) using the Bloomberg Financial Database for the period December 1995 to March 2000. To quantify the impact of transaction costs on the average payoffs of the hedged carry trade we increase the prices of the puts and calls used in our strategy by

¹⁰See 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.

¹¹The average payoffs and the Sharpe ratios for both carry-trade strategies are higher when we exclude the financial crisis period. Suppose that we compare a sample that stops on July 2008 to our sample period. The average payoff to the unhedged carry trade falls from 3.48 percent to 2.96 percent, while the Sharpe ratio falls from 0.586 to 0.476. The average payoff to the hedged carry trade falls from 1.80 percent to 1.58 percent, while the Sharpe ratio falls from 0.530 to 0.449.

¹²The average bid-ask spreads for individual currencies are: Canadian dollar call 5.3 percent, put 4.4 percent, Euro call 4.3 percent, put 4.8 percent, Japanese yen call 5.3 percent, put 5.6 percent, Swiss franc call 5.3 percent, put 6.4 percent, and British pound call 4.3 percent, and put 4.6 percent.

one half of the average bid-ask spread (2.6 percent). We find that the average payoff to the hedged carry trade declines from 0.0158 to 0.0121 as a result of transaction costs. This result suggests that transactions costs have a modest impact on the average payoffs of the hedged carry trade.

The first panel of Figure 3 displays a 12-month moving average of the realized payoffs for the hedged and unhedged carry-trade strategies. The second panel displays a time series of the realized Sharpe ratios over a 12-month moving window 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 payoffs appear quite similar. For both strategies 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.

There is an important dimension along which the payoffs of the two carry-trade strategies are quite different. As Figure 4 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.

Panel B of Figure 1 shows the cumulative payoffs to the unhedged carry trade, U.S. stocks and Treasury bills as well as the cumulative payoff to the hedged carry trade, beginning from a common initial date, December 1986. Two key features are worth noting. First, the cumulative payoff to the hedged carry trade is somewhat lower than that of the unhedged carry trade. This result reflects the cost of the options used in the hedged strategy and the fact that there are no large negative payoffs to the unhedged carry trade in sample. Second, the hedged carry trade payoffs are less volatile than those of the unhedged carry trade. This result reflects the fact that the hedging strategy truncates a subset of the negative payoffs to the unhedged carry trade that occur in the sample.

Time-series risk-factor analysis We now investigate whether the payoffs to the hedged carry trade are correlated with traditional risk measures. Table 6 reports our estimates of β , the regression coefficient of the carry-trade payoff on candidate risk factors, using the four risk models we have been considering. Note that two factors display a significant correlation with the payoffs to the hedged carry trade: the excess return to the value-weighted U.S. stock market and the value premium. This finding contrasts with our result that none of the risk factors are correlated with the payoffs to the unhedged carry trade.

Panel risk-factor analysis 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 4. Our estimation results are generated using a 26×1 vector of time- t excess returns to the hedged carry-trade strategy and the 25 Fama-French portfolios. We report our results in Table 7. As before, our results fall into two categories, depending on whether the model is rejected based on the test of the overidentifying restrictions. For the CAPM, the Fama-French model, and the C-CAPM the model is rejected at the one percent level. The extended C-CAPM model is not rejected because the b parameters are estimated with great imprecision.

Figure 5 displays the predictions of the CAPM, the C-CAPM, the extended C-CAPM models, and the Fama-French model for $E(R_t)$ against the sample average of R_t . The first three models do a very poor job at explaining the cross-sectional variation of excess returns. Indeed, the cross-sectional R^2 of these models is negative. The Fama-French model does a reasonable job of explaining the average excess returns to the Fama-French portfolios and the payoffs to the hedged carry trade.¹³ Recall that the Fama-French model does a very poor job at explaining the payoffs to the unhedged carry trade (see Figure 2).

In sum, this section, together with Section 4, indicates that it is difficult to explain the

¹³The use of options in the construction of the hedged carry-trade strategies can introduce non-linearities in the portfolio payoffs (see Glosten and Jagannathan (1994) and Broadie, Chernov and Johannes (2009)). Following Glosten and Jagannathan (1994), we re-do the cross-sectional analysis allowing for quadratic terms in the factors. Our results are unaffected by this extension.

payoffs to the unhedged and hedged carry trade as compensation for exposure to conventional measures of risk. We now turn to the question of whether the payoffs to the carry trade reflect a peso problem.

6 Characterizing the nature 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 $M'/E^N(M)$. We compute these estimates 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

Our estimates of z' and $M'/E^N(M)$ are based on equations (18) and (19) which require estimates of the risk-adjusted average payoffs to the unhedged and hedged carry trade. The estimates of these payoffs reported in Sections 4 and 5 are generated under the null hypothesis of no peso events. Since this null is inappropriate for the analysis in this section, we re-estimate these payoffs allowing for peso events. In practice this means allowing for a constant in our linear SDF model which is based on the Fama-French factors (see the online appendix).

The first column of Table 8 reports our benchmark estimates which are based on the CME data for the sample period 1987 to 2009. Note that the average payoffs to the unhedged and hedged carry trade are equal to 0.0025 and 0.0013, respectively. The mean minimum net payoff to the hedged carry trade, $E^N(h)$, is equal to -0.0105 .

We use our estimated SDF model to generate a time series for M_t . We then use this time

series to estimate $E^N(Mz)$ and $E^N(Mz^H)$. The resulting point estimate of z' is -0.0216 with a standard error of 0.0059 . The implied two-standard-error band for z' is $(-0.033, -0.010)$. The standard deviation of z in our sample is 0.018 . So, our point estimate for z' is roughly 1.3 standard deviations below the estimated value of $E^N(z)$. The lower bound of the two-standard-error band for z' (-0.033) is only two standard deviations below the average payoff to the unhedged carry trade. Thus, we find little evidence to support the view that z' is a very large negative value relative to the empirical distribution of payoffs to the carry trade. For robustness we calculated values of z' and M' under the assumption that the correlation between the SDF and the hedged and unhedged carry trade payoffs is zero in non-peso states. As reported in Table 8 these estimates are similar to the results reported above. Viewed overall, our results are consistent with Bates (1996b) who studies high-frequency data on options prices for the Deutsche Mark and the Yen against the U.S. dollar over the period 1984 to 1992. He argues that peso events, defined as a large, negative, value of z' , cannot account for the failure of UIP.

Given an estimate of z' and a value of p we can estimate $M'/E^N(M)$ using equation (19). Nakamura, Steinsson, Barro, and Ursúa (2010) provide a convenient benchmark value for p . Using consumption data covering 24 countries and more than 100 years, these authors estimate the annual probability of a disaster, thought of as a large drop in consumption, to be 0.017 . Nakamura et al. (2010) argue that the frequency of disasters in the U.S. is statistically consistent with their estimate of p . Motivated by their estimate we use a monthly value of p equal to 0.0014 . This value of p implies that a peso event occurs on average once every 700 months. Recall that our data cover a period of 268 months and we assume that there are no peso events in sample. Our benchmark value of p implies that the probability of not observing a peso event in a sample as long as ours is 68 percent. So, our assumption that a peso event did not occur in our sample is quite plausible from a purely statistical point of view.

In conjunction with equation (19), our benchmark value of p yields an estimate of

$M'/E^N(M)$ equal to 92.9, with a standard error 47.5. This result supports interpreting a peso event as primarily reflecting a high value of the SDF. That is, traders value losses very highly in a peso state.

There is clearly a great deal of uncertainty about the true value of p . But the basic result that a peso event is characterized by a high value of $M'/E^N(M)$ is robust for plausible values of p . Higher values of p imply lower values of $M'/E^N(M)$. Suppose, for example, that p is equal to $1/268$. In this case the probability of observing a sample of 268 months with no peso events is roughly 37 percent. The implied value of $M'/E^N(M)$ is roughly 35. Alternatively, suppose we choose $p = 0.0111$, so that the probability of not observing a peso event in our sample is only five percent. Even in this case the implied value of $M'/E^N(M)$ equals roughly 12.

6.2 Incorporating stock-market data into our analysis

In this subsection we use stock market data to generate an independent estimate of M' . 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. As Table 5 indicates, over the period February 1987–April 2009, the annualized rate of return to the unhedged stock-market strategy is 6.87 percent versus -4.79 for the hedged stock-market strategy. In sharp contrast, over the same sample period, the annualized payoff to the unhedged carry trade is 2.96 percent versus 1.58 percent for the hedged carry trade.

In section 2 we develop estimators of M' and x' , the payoff to the stock market strategy in a peso state. We base our estimators on equations (22) and (23). We use estimates of the Fama-French (1993) model fit to the 25 Fama-French portfolios over the period February 1986–July 2009 to compute a time series for M_t . The options data we use to construct the hedged equity strategy are available over this same time period. We use sample averages of $M_t x_t$ and $M_t x_t^H$ to estimate $E^N(Mx)$ and $E^N(Mx^H)$, respectively.¹⁴

¹⁴The results reported in the text are based on a linear SDF. We find that the same implied values of x' and $M'/E^N(M)$ cannot be rejected when we use an SDF that includes quadratic terms in some or all of the

Our estimate of x' is -0.188 . This value of x' is roughly ten times larger in absolute value than z' . By this metric the peso event has a larger impact on stock market payoffs than on carry trade payoffs. Also x' is more than four standard deviations away from our estimate of $E^N(x)$. Using our benchmark value of p , our point estimate of $M'/E^N(M)$ based on stock returns is equal to 67.6. Recall that our estimate of $M'/E^N(M)$ based on carry-trade returns is 92.9. These two estimates are not statistically significantly different from each other. In this sense, the same value of $M'/E^N(M)$ 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 relatively modest, negative carry-trade payoffs.

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. We report our results in Table 8. Our estimates of $E^N(h)$, $E^N(Mz)$ and $E^N(Mz^H)$ imply an estimate of z' equal to -0.0386 with a standard error of 0.0173. Table 8 also reports results based on the ten currency version of the J.P. Morgan data set. These estimates imply an estimate for z' equal to -0.0438 with a standard error of 0.0207. For both J.P. Morgan data sets, our estimate of z' is more negative than for the CME data set and represents a payoff roughly three standard deviations below the estimate of $E^N(z)$. Still, these payoffs are not particularly large, and our estimates of $M'/E^N(M)$ remain large.

6.4 Robustness analysis: allowing for multiple-peso states

Under the assumption that there is a single peso state we find that the value of z' is a relatively small negative number. Here we study the robustness of our results by allowing

Fama-French factors.

for a continuum of peso states. As above, we denote the set of peso states by Ω^P and assume that the probability of the collection of peso states is p , both conditionally and unconditionally. When the economy is in the peso state the value of ω_{t+1} is drawn from an i.i.d distribution. We denote by $\mathcal{F}^P(\omega_{t+1})$ the cumulative distribution of ω_{t+1} given that $\omega_{t+1} \in \Omega^P$. We let the functions $M'(\omega_{t+1})$ and $z'(\omega_{t+1})$ denote the values of the SDF and payoff to the carry trade for each $\omega_{t+1} \in \Omega^P$. The payoffs to the unhedged carry trade must satisfy:

$$(1-p)E^N(Mz) + p \int_{\Omega^P} M'(\omega_{t+1})z'(\omega_{t+1})d\mathcal{F}^P(\omega_{t+1}) = 0. \quad (31)$$

Consider now the hedged carry trade strategy, where the hedging relies on at-the-money options. Since these options are in the money in all peso states, it follows that:

$$(1-p)E^N(Mz^H) + pE^P(M')E^N(h) = 0 \quad (32)$$

where $E^P(M') = \int_{\Omega^P} M'(\omega_{t+1})d\mathcal{F}^P(\omega_{t+1})$.

Solving equation (32) for $(1-p)E^N(M)$ and substituting the result into equation (31) we obtain:

$$\int_{\Omega^P} M'(\omega_{t+1})z'(\omega_{t+1})d\mathcal{F}^P(\omega_{t+1}) = \frac{E^N(Mz)}{E^N(Mz^H)}E^P(M')E^N(h). \quad (33)$$

Letting $E^P(z') = \int_{\Omega^P} z'(\omega_{t+1})d\mathcal{F}^P(\omega_{t+1})$ we can rewrite equation (33) as:

$$E^P(M')E^P(z') + \text{cov}^P(M', z') = \frac{E^N(Mz)}{E^N(Mz^H)}E^P(M')E^N(h). \quad (34)$$

We assume that there is a tendency for worse peso states (large values of M') to be associated with worse payoffs (more negative values of z') so that $\text{cov}^P(M', z') \leq 0$. In this case equation (34) implies that:

$$E^P(z') \geq \frac{E^N(Mz)}{E^N(Mz^H)}E^N(h). \quad (35)$$

So, equation (35) implies that the expected value of z' across all peso states is greater than or equal to the estimate of z' implied by equation (18). While there can be some large negative values of z' , these values must have low probabilities. If we assume that

$\text{cov}^P(M', z') = 0$, then equation (35) implies that the average value of z' is equal to that implied by equation (18).

We now consider the implications of this extension for the average value of the SDF across peso states. Solving equation (32) for the average value of the SDF in the peso state we obtain:

$$E^P(M') = -\frac{1-p}{p} \frac{E^N(Mz^H)}{E^N(h)}. \quad (36)$$

Equation (19) is the analogue of equation (19). Furthermore, the two equations are equivalent if equation (35) holds with equality.

In sum, in the presence of multiple peso states our empirical results can be reinterpreted as pertaining to the average value of payoffs and the SDF across peso states.

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 of 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 average payoff to the carry trade.

We argue that the positive average payoff to the unhedged carry trade reflects peso event risk. A peso event consists of a negative payoff to the carry trade and an associated value of the SDF. We develop and implement a strategy to characterize the peso event. This strategy uses the payoffs to a version of the carry trade that employs currency options to protect an investor from the downside risk associated with large, adverse movements in exchange rates. By construction, this hedged carry trade strategy eliminates the large negative payoffs associated with peso events. Our key finding is that a peso event is characterized by modest negative payoffs to the unhedged carry trade along with a large value of the SDF.

It is important to emphasize that we base our results on an asset pricing framework that is linear except for the peso state. So, we do not rule out the possibility that the payoffs to the carry trade can be explained using a non-linear SDF model.

Finally, we note that our analysis is based on unconditional covariances between risk factors and carry trade payoffs. It might be difficult to detect non-zero covariances in samples of our size when the conditional covariance is time varying. Suppose, for example, that the covariance is zero in most states of nature but it is strongly negative in states that occur with low probability. In this setting the sample size necessary to find a statistically significant covariance might be much greater than the size of our sample.

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.
- Bates, David (1996a) “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options,” *Review of Financial Studies* 9:1, 69–107.
- Bates, David (1996b) “Dollar Jump Fears, 1984–1992: Distribution Abnormalities Implicit in Currency Futures Options,” *Journal of International Money and Finance*, 15: 1, 65–93.
- 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.
- Bhansali, Vineer (2007) “Volatility and the Carry Trade,” *The Journal of Fixed Income*, 72–84, Winter.
- Broadie, Mark , Michael S. Johannes and Mikhail Chernov (2009) “Understanding Index Option Returns,” *Review of Financial Studies* 22 (4, May): 4493–4529.
- Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen (2008) “Carry Trades and Currency Crashes,” *NBER Macroeconomics Annual 2008*, 23: 313–347, 2009.
- 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 (1999) “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.
- 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 (1973) “Risk Return and Equilibrium: Empirical Tests,” *Journal of Financial Political Economy*, 71: 607-636.
- Farhi, Emmanuel and Xavier Gabaix (2008) “Rare Disasters and Exchange Rates,” National Bureau of Economic Research Working Paper No. 13805.
- Galati, Gabriele and Michael Melvin (2004) “Why has FX Trading Surged? Explaining the 2004 Triennial Survey,” *BIS Quarterly Review* (Dec.), 67–74.
- Garman, Mark B. and Steven W. Kohlhagen (1983) “Foreign currency option values,” *Journal of International Money and Finance*, 2, 231–7.
- Glosten, Lawrence and Ravi Jagannathan (1994) “A Contingent Claims Framework for Analyzing the Performance of Portfolio Managers,” *Journal of Empirical Finance*, 1, 133-160.
- Hansen, Lars Peter (1982) “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica* 50 (4, July), 1029–54.
- Ibbotson and Associates (2006) *Stocks, Bonds, Bills and Inflation Yearbook*. Chicago, Ill.: Ibbotson Associates.
- Jurek, Jakub W. (2008) “Crash-neutral Currency Carry Trades,” mimeo, Princeton University.
- Li, Haitao and Yewu Xu Xiaoyan Zhang (2010) “Evaluating Asset Pricing Models Using the Second Hansen-Jagannathan Distance,” forthcoming in *Journal of Financial Economics*.
- Lustig, Hanno and Adrien Verdelhan (2007) “The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk,” *American Economic Review*, 97, 89–117.
- Nakamura, Emi, Jon Steinsson, Robert J. Barro, and José F. Ursúa (2010) “Crises and Recoveries in an Empirical Model of Consumption Disasters,” NBER Working Paper 15920.
- Pastor, Lubos and Robert F. Stambaugh (2003) “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy* 111, 642–85.
- Verdelhan, Adrien (2010) “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance*, 65: 1, 123-145..
- 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: ANNUALIZED PAYOFFS OF THE EQUALLY-WEIGHTED CARRY-TRADE STRATEGY

	No Transactions Costs			With Transactions Costs		
	Mean	Standard Deviation	Sharpe Ratio	Mean	Standard Deviation	Sharpe Ratio
British Pound is the Base Currency February 1976 to July 2009	0.0319 (0.0080)	0.043 (0.003)	0.748 (0.194)	0.0288 (0.0111)	0.057 (0.004)	0.507 (0.203)
US Dollar is the Base Currency January 1997 to July 2009	0.0440 (0.0171)	0.051 (0.005)	0.865 (0.358)	0.0431 (0.0213)	0.062 (0.007)	0.694 (0.356)

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 the Appendix. Heteroskedasticity consistent GMM standard errors are in parentheses.

TABLE 2: ANNUALIZED PAYOFFS OF INVESTMENT STRATEGIES
 US Dollar is the Base Currency (February-1976 to July-2009)

	Mean	Standard Deviation	Sharpe Ratio	Skewness	Excess Kurtosis	Jarque-Bera Statistic
U.S. stock market	0.0582 (0.0281)	0.156 (0.010)	0.373 (0.192)	-0.808 (0.288)	2.53 (1.17)	150.9 (0.000)
Equally-weighted carry trade	0.0482 (0.0101)	0.053 (0.005)	0.911 (0.222)	-0.648 (0.520)	5.81 (2.02)	592.6 (0.000)
Average of individual-currency carry trades	0.0492	0.114	0.442	-0.229	1.57	67.4

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 the Appendix. Heteroskedasticity consistent GMM standard errors are in parentheses, except for the Jarque-Bera statistic for which the p-value is reported in parentheses. The mean payoff of the equally-weighted carry trade is not equal to the average mean payoff of the individual-currency carry trades because the sample periods for which the currencies are available varies (see Appendix A).

TABLE 3: FACTOR BETAS OF THE EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO

Factors	Intercept	Beta(s)			R^2
(A) Monthly Payoffs and Risk Factors (February 1976 to July 2009)					
CAPM	0.004 (0.001)	0.018 (0.018)			0.003
Fama-French factors	0.004 (0.001)	0.033 (0.019)	-0.045 (0.031)	0.028 (0.033)	0.016
(B) Quarterly Real Excess Returns and Risk Factors (1976Q2 to 2009Q2)					
C-CAPM	0.012 (0.005)	0.058 (0.840)			0.000
Extended C-CAPM	0.007 (0.008)	-0.210 (0.921)	0.607 (0.663)	0.011 (0.034)	0.009

Notes: Part (A) of the table reports estimates of the equation $z_t = a + f_t' \beta + \epsilon_{t+1}$, where z_t is the monthly nominal payoff 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 ($Mkt - Rf$), and the Fama-French factors are the $Mkt - Rf$, SMB and HML factors (available from Kenneth French's database). Heteroskedasticity-robust standard errors are in parentheses. Part (B) of 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 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. Details of the risk factors are provided in the Appendix. Heteroskedasticity-robust standard errors are in parentheses.

TABLE 4: GMM ESTIMATES OF LINEAR FACTOR MODELS

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

	μ	b	λ	R^2	J	α
(A) Monthly Data (February 1976 to July 2009)						
CAPM	0.0049 (0.0023)	2.01 (1.25)	0.41 (0.24)	-1.97	102 (0.00)	0.047 (0.010)
Fama-French Factors						
<i>Mkt-Rf</i>	0.0049 (0.0023)	3.45 (1.50)	0.44 (0.24)	0.40	89.4 (0.00)	0.046 (0.010)
<i>SMB</i>	0.0027 (0.0016)	3.21 (1.80)	0.22 (0.15)			
<i>HML</i>	0.0036 (0.0018)	7.25 (2.02)	0.40 (0.17)			
(B) Quarterly Data (1976Q2 to 2009Q2)						
C-CAPM	0.0044 (0.0006)	196 (83.7)	0.34 (0.13)	-2.33	35.8 (0.07)	0.049 (0.026)
Extended C-CAPM						
<i>Consumption</i>	0.0044 (0.0006)	11.2 (114)	0.02 (0.20)	-7.77	2.83 (1.00)	0.050 (0.031)
<i>Durables</i>	0.0100 (0.0028)	0.96 (61.5)	0.01 (0.17)			
<i>Market return</i>	0.0184 (0.0073)	0.23 (2.14)	0.22 (1.22)			

Notes: Part (A) of the table reports second stage GMM estimates of the SDF $m_t = 1 - (f_t - \mu)'b$ using the moment conditions $E(R_t m_t) = 0$ and $E(f_t - \mu) = 0$, where R_t is a 26×1 vector containing the nominal excess returns of the Fama-French 25 value-weighted portfolios of US stocks sorted on size and the book-to-market value ratio as well as the monthly nominal payoff of the equally-weighted carry-trade portfolio, and f_t is a scalar or vector of risk factors. Part (B) of the table uses the same moment conditions, but R_t is a 26×1 vector containing the real quarterly excess returns of the Fama-French 25 portfolios and the equally-weighted carry-trade portfolio. The risk factors are described in more detail in the footnote to Table 3. The GMM procedure is described in more detail in the Appendix. Estimates of the factor risk premia $\hat{\lambda} = \hat{V}_f \hat{b}$ are 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 R^2 is a measure of fit between the sample mean of R_t and the predicted mean returns, given by $T^{-1} \sum_{t=1}^T R_t (f_t' - \hat{\mu})' \hat{b}$. Tests of the overidentifying restrictions are also reported. The test statistic, J , is asymptotically distributed as a χ^2_{26-k} , where k is the number of risk factors. The p-value is in parentheses. The pricing error of the equally-weighted carry-trade portfolio (α) is reported annualized. Its standard error is in parentheses.

TABLE 5: ANNUALIZED PAYOFFS OF INVESTMENT STRATEGIES
 US Dollar is the Base Currency (February 1987 to April 2009)

	Mean	Standard Deviation	Sharpe Ratio	Skewness	Excess Kurtosis	Jarque-Bera Statistic
U.S. stock market	0.0452 (0.0349)	0.158 (0.014)	0.286 (0.239)	-1.141 (0.333)	3.31 (1.53)	180 (0.000)
S&P 100 stock index						
Unhedged	0.0687 (0.0347)	0.163 (0.013)	0.422 (0.233)	-0.593 (0.246)	2.00 (0.63)	60.0 (0.000)
Hedged	-0.0479 (0.0214)	0.098 (0.007)	-0.491 (0.225)	0.955 (0.339)	2.24 (1.44)	96.3 (0.000)
Equally-weighted carry trade						
Unhedged	0.0296 (0.0136)	0.062 (0.005)	0.476 (0.234)	-0.708 (0.154)	1.47 (0.44)	46.3 (0.000)
Hedged	0.0158 (0.0078)	0.035 (0.002)	0.449 (0.212)	0.722 (0.248)	1.14 (0.63)	37.6 (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 6: FACTOR BETAS OF THE HEDGED EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIO

Factors	Intercept	Beta(s)			R^2
(A) Monthly Payoffs and Risk Factors (February 1987 to April 2009)					
CAPM	0.001 (0.001)	0.021 (0.015)			0.009
Fama-French factors	0.001 (0.001)	0.037 (0.018)	-0.003 (0.020)	0.060 (0.025)	0.042
(B) Quarterly Real Excess Returns and Risk Factors (1987Q1 to 2009Q1)					
C-CAPM	0.004 (0.002)	-0.041 (0.059)			0.005
Extended C-CAPM	0.004 (0.002)	-0.038 (0.067)	-0.008 (0.078)	0.007 (0.095)	0.005

Notes: Part (A) of the table reports estimates of the equation $z_t^H = a + f_t' \beta + \epsilon_{t+1}$, where z_t^H is the monthly nominal payoff of the hedged equally-weighted carry-trade portfolio and f_t is a scalar or vector of risk factors. Part (B) of 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. The risk factors are described in the footnote to Table 3. Heteroskedasticity-robust standard errors are in parentheses.

TABLE 7: GMM ESTIMATES OF LINEAR FACTOR MODELS

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

	μ	b	λ	R^2	J	α
(A) Monthly Data (February 1987 to April 2009)						
CAPM	0.0038 (0.0028)	2.00 (1.49)	0.42 (0.28)	-0.49	83.4 (0.00)	0.015 (0.008)
Fama-French Factors						
<i>Mkt-Rf</i>	0.0038 (0.0028)	3.18 (1.73)	0.41 (0.28)	0.29	78.8 (0.00)	0.011 (0.008)
<i>SMB</i>	0.0009 (0.0021)	1.32 (2.02)	0.03 (0.20)			
<i>HML</i>	0.0025 (0.0020)	5.71 (2.31)	0.36 (0.22)			
(B) Quarterly Data (1987Q1 to 2009Q1)						
C-CAPM	0.0041 (0.0006)	153 (57.8)	0.21 (0.09)	-0.83	54.6 (0.00)	0.013 (0.011)
Extended C-CAPM						
<i>Consumption</i>	0.0041 (0.0008)	-8.16 (76.0)	-0.01 (0.09)	-6.11	2.67 (1.00)	0.016 (0.011)
<i>Durables</i>	0.0104 (0.0048)	-0.15 (42.1)	0.00 (0.07)			
<i>Market return</i>	0.0152 (0.0090)	-0.02 (1.07)	-0.07 (0.77)			

Notes: Part (A) of the table reports second stage GMM estimates of the SDF $m_t = 1 - (f_t - \mu)'b$ using the moment conditions $E(R_t m_t) = 0$ and $E(f_t - \mu) = 0$, where R_t is a 26×1 vector containing the nominal excess returns of the Fama-French 25 value-weighted portfolios of US stocks sorted on size and the book-to-market value ratio as well as the monthly nominal payoff of the hedged equally-weighted carry-trade portfolio, and f_t is a scalar or vector of risk factors. Part (B) of the table uses the same moment conditions, but R_t is a 26×1 vector containing the real quarterly excess returns of the Fama-French 25 portfolios and the hedged equally-weighted carry-trade portfolio. The risk factors are described in more detail in the footnote to Table 3. The GMM procedure is described in more detail in the Appendix. Estimates of the factor risk premia $\hat{\lambda} = \hat{V}_f \hat{b}$ are 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 R^2 is a measure of fit between the sample mean of R_t and the predicted mean returns, given by $T^{-1} \sum_{t=1}^T R_t (f_t' - \hat{\mu})' \hat{b}$. Tests of the overidentifying restrictions are also reported. The test statistic, J , is asymptotically distributed as a χ^2_{26-k} , where k is the number of risk factors. The p-value is in parentheses. The pricing error of the hedged equally-weighted carry-trade portfolio (α) is reported annualized. Its standard error is in parentheses.

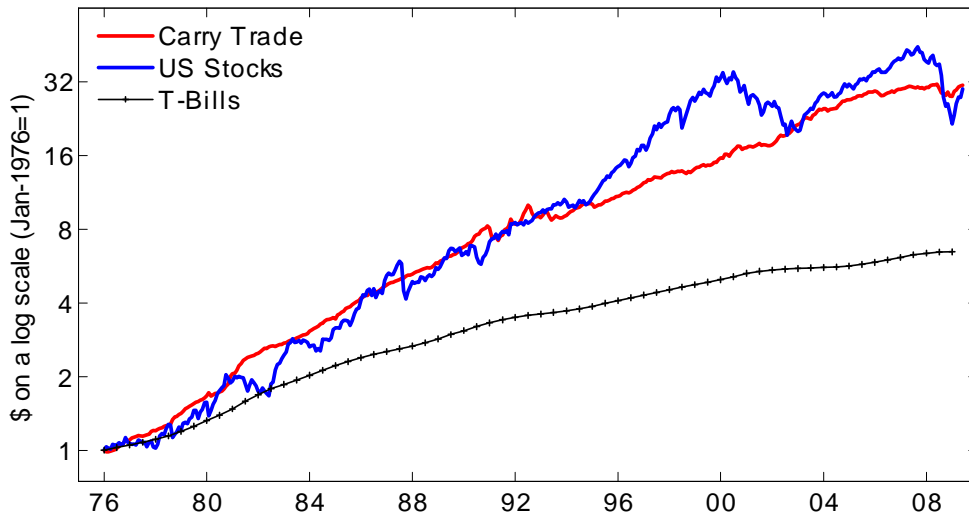
TABLE 8: ESTIMATES OF PAYOFFS AND THE SDF IN THE PESO EVENT

Equally-Weighted Carry Trade Portfolios				Stock Portfolio	
Options data:	CME	J.P. Morgan		S&P 100 / VXO	
Sample:	1987M2-2009M4	1995M2-2009M7		1986M2-2009M7	
Currencies:	6	6	10		
(A) Non-risk-corrected calculations					
$E^N(z)$	0.0025 (0.0011)	0.0031 (0.0010)	0.0035 (0.0012)	$E^N(x)$	0.0068 (0.0029)
$E^N(z^H)$	0.0013 (0.0007)	0.0011 (0.0006)	0.0011 (0.0007)	$E^N(x^H)$	-0.0032 (0.0018)
$E^N(h)$	-0.0105 (0.0004)	-0.0114 (0.0006)	-0.0120 (0.0006)	$E^N(h_x)$	-0.0257 (0.0015)
z'	-0.0198 (0.0055)	-0.0328 (0.0113)	-0.0375 (0.0146)		
$M'/E^N(M)$	87.1 (44.3)	66.0 (36.2)	64.9 (41.1)		
(B) Risk-corrected calculations					
$E^N(Mz)$	0.0029 (0.0011)	0.0033 (0.0010)	0.0040 (0.0011)	$E^N(Mx)$	0.0182 (0.0026)
$E^N(Mz^H)$	0.0014 (0.0007)	0.0010 (0.0006)	0.0011 (0.0007)	$E^N(Mx^H)$	0.0025 (0.0020)
z'	-0.0216 (0.0059)	-0.0386 (0.0173)	-0.0438 (0.0207)	x'	-0.188 (0.130)
$M'/E^N(M)$	92.9 (47.5)	60.0 (39.1)	63.3 (42.8)	$M'/E^N(M)$	67.6 (54.8)

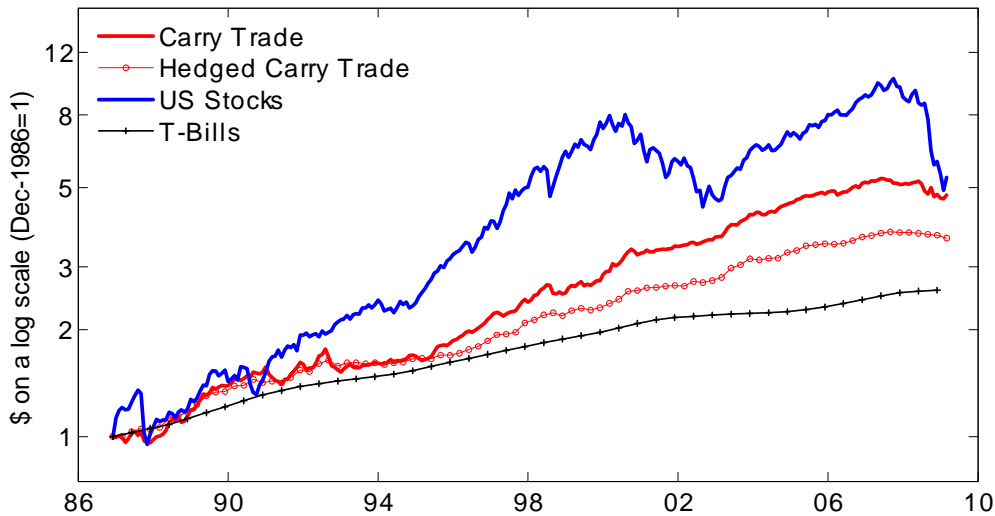
Notes: As in the text z and x denote, respectively, the unhedged payoffs of the equally-weighted carry trade and S&P 100 index portfolios. The hedged payoffs of these portfolios are denoted, respectively, z^H and x^H . The variables h and h_x denote, respectively, denote the minimum payoffs of the two hedged portfolios. The payoffs of the unhedged portfolios in the peso state are denoted, respectively, z' and x' . The variable M represents the stochastic discount factor, with M' denoting its payoff in the peso state. The operator E^N is the unconditional expectations operator that applies to non-peso states of the world. To compute z^H we use data from the Chicago Mercantile Exchange (CME) and J.P. Morgan. To compute x^H we use implied volatility data (VXO) corresponding to options on the level of the S&P 100 index. Our procedure for constructing M and our data are described in the text and the the Appendix. Heteroskedasticity consistent GMM standard errors are in parentheses.

FIGURE 1: CUMULATIVE RETURNS OF INVESTMENT STRATEGIES

(A) 20 Currency Carry Trade and US Stocks, Jan 1976–July 2009

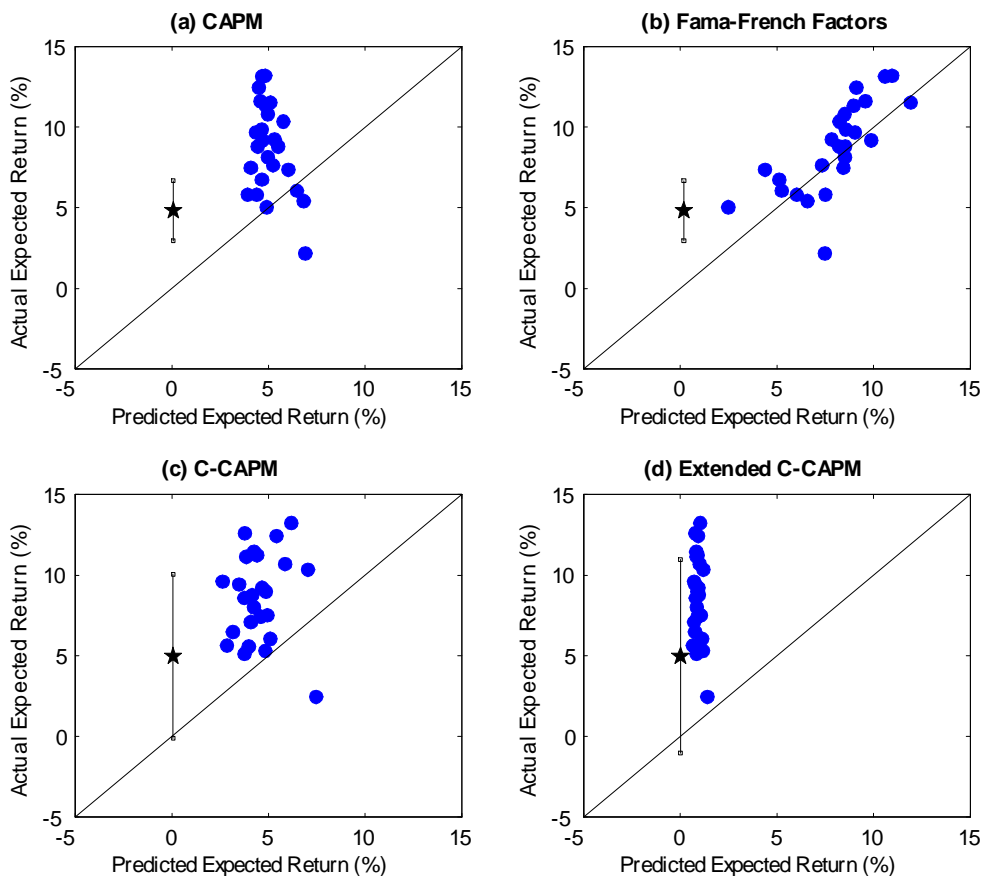


(B) 6 Currency Carry Trade (Hedged and Unhedged) and US Stocks, Jan 1987–April 2009



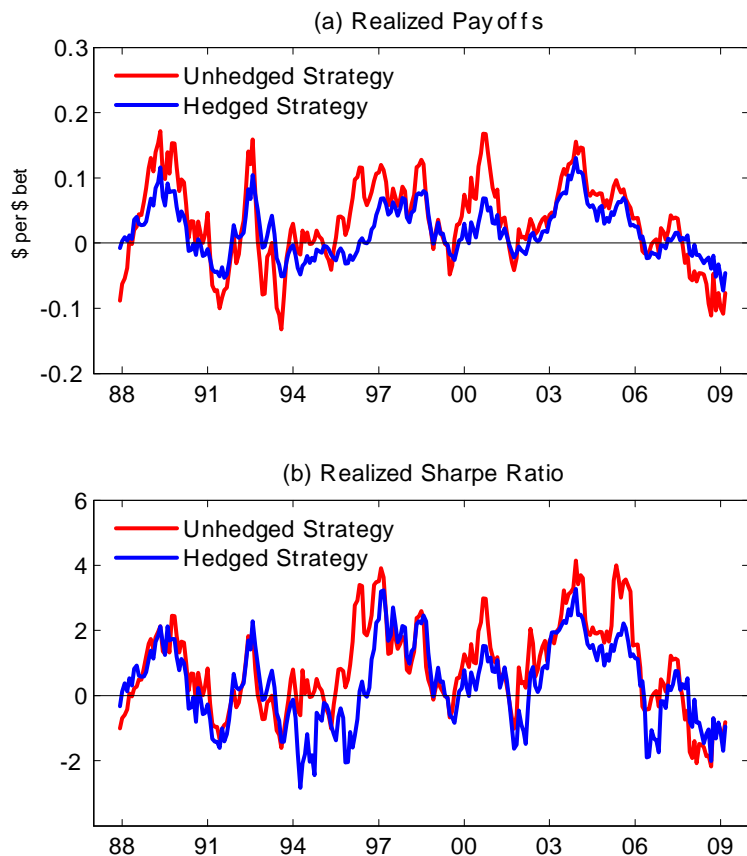
Note: The figure plots the cumulative returns of a trader who begins with 1 dollar in January 1976 (panel A) or January 1987 (panel B) and invests his accumulated earnings exclusively in one of four strategies. For T-bills and US stocks we use the risk free rate and value-weighted market return reported in Kenneth French’s database. For the carry trade strategies we assume that the trader invests the initial dollar in T-bills and bets the future nominal value of those T-bills in the carry trade. In each period all proceeds are deposited in the T-bill account, and the future value of the T-bill account is bet on the carry trade. Details of the strategies are provided in the text.

FIGURE 2: CROSS-SECTIONAL FIT OF ESTIMATED FACTOR MODELS
 Test Assets: Fama-French 25 Portfolios & the Equally-Weighted Carry-Trade Portfolio



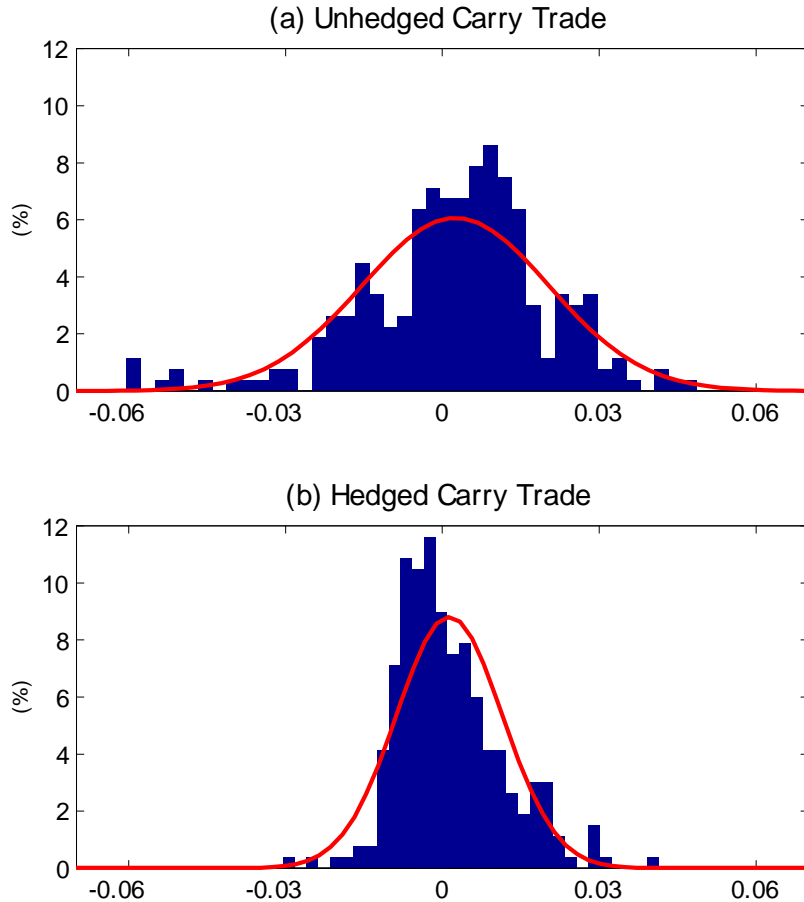
Note: In each case the parameters μ and b in the SDF $m_t = 1 - (f_t - \mu)'b$ are estimated by GMM using the method described in the text. The risk factors, f_t , are indicated by the title of each plot with details provided in the main text. The predicted expected return is $(1/T) \sum_{t=1}^T R_{it}(f_t - \hat{\mu})'\hat{b}$ for each portfolio's excess return, R_{it} . The actual expected return is $\bar{R}_i = (1/T) \sum_{t=1}^T R_{it}$. 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. For models (a) and (b) the sample period is 1976M2–2009M7. For models (c) and (d) the sample period is 1976Q2–2009Q2. Expected returns are annualized.

FIGURE 3: ANNUALIZED REALIZED AVERAGE PAYOFFS AND SHARPE RATIOS OF THE EQUALLY-WEIGHTED HEDGED AND UNHEDGED CARRY-TRADE PORTFOLIOS 12-Month Rolling Window, February 1987–April 2009



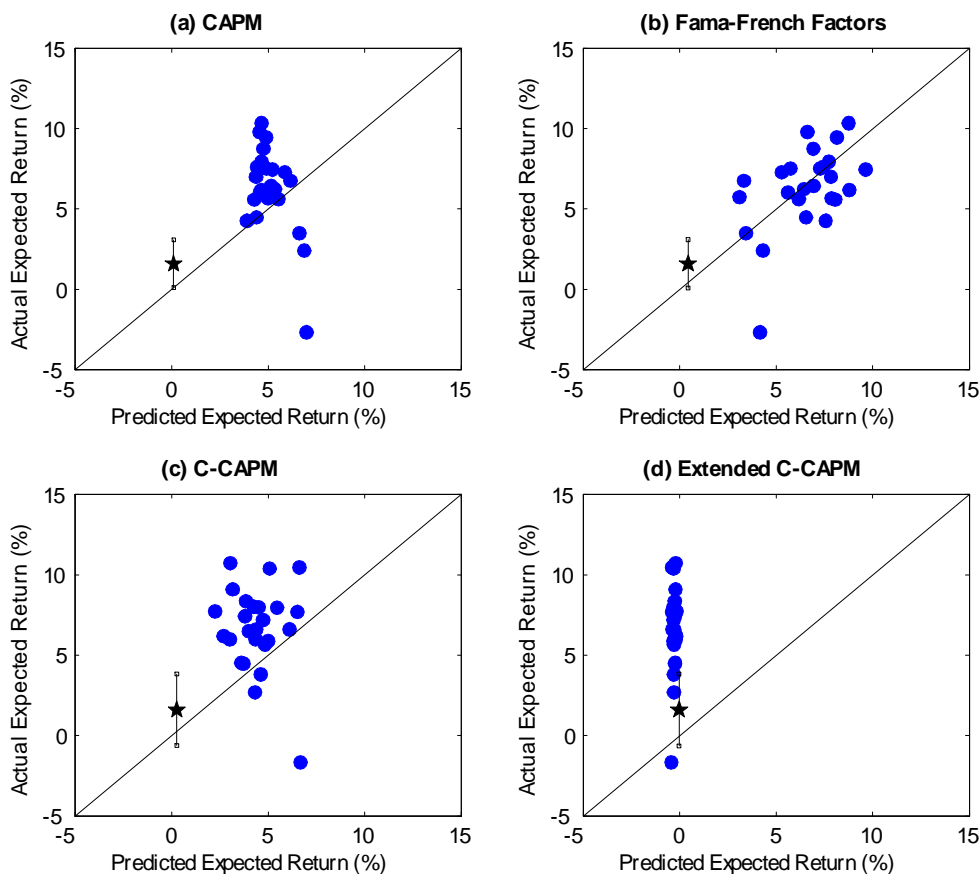
Note: Plot (a) shows the annualized average payoff from month $t - 11$ to month t , in US dollars, per dollar bet in the carry trade. Plot (b) shows the ratio of the annualized average payoff, to the annualized standard deviation of the payoff, both being measured from month $t - 11$ to month t . The 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 4: SAMPLING DISTRIBUTIONS OF THE PAYOFFS OF THE EQUALLY-WEIGHTED CARRY-TRADE PORTFOLIOS
February 1987–April 2009



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 5: CROSS-SECTIONAL FIT OF ESTIMATED FACTOR MODELS
 Test Assets: Fama-French 25 Portfolios & the Hedged Equally-Weighted Carry-Trade



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 $(1/T) \sum_{t=1}^T R_{it} (f_t - \hat{\mu})' \hat{b}$ for each portfolio's excess return, R_{it} . The actual expected return is $\bar{R}_i = (1/T) \sum_{t=1}^T R_{it}$. 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. For models (a) and (b) the sample period is 1987M2–2009M4. For models (c) and (d) the sample period is 1987Q1–2009Q1. Expected returns are annualized.

A: Spot and Forward Exchange Rate Data

We obtain our foreign exchange rate data 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 (Australian dollar, Austrian schilling, Belgian franc, Canadian dollar, Danish krone, euro, French franc, German mark, Irish punt, Italian lira, Japanese yen, Netherlands guilder, New Zealand dollar, Norwegian krone, Portuguese escudo, South African rand, Spanish peseta, Swedish krona, Swiss franc, U.S. dollar) quoted against the British pound. The data series begin in January 1976, with some exceptions (Ireland, April 1979; Japan, June 1978; euro, December 1998; Australia, NZ and South Africa, December 1996) and end in July 2009, with the exception of the euro legacy currencies (December 1998). The mnemonics for each currency are indicated in the online appendix. 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 second data set consists of spot exchange rates and one month forward exchange rates for twenty currencies quoted against the U.S. dollar. The currencies are the same as above, with the British pound replacing the U.S. dollar. The data series begin in December 1996, with the exception of the euro (December 1998), and end in July 2009, with the exception of the euro legacy currencies (December 1998). The mnemonics for each currency are indicated in the online appendix. 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.

When we ignore bid-ask spreads we obtain a data set running from January 1976 to July 2009 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. The original data set includes observations on all weekdays. In our analysis of the unhedged carry trade (Tables 1–4) we measure payoffs using last business day of the month observations.

B: Options Data and Options-Based Strategies

CME Options 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 April 2009. Due to sparse coverage in the early part of the sample we begin our analysis no earlier than January 1987.

CME options are options against currency futures. The options themselves expire early in each month (two Fridays prior to the third Wednesday in the month). The futures against which the options are written expire on the Monday prior to the third Wednesday (except for the Canadian dollar, for which expiry takes place on the Tuesday prior to the third Wednesday) of March, June, September and December. When we construct hedged positions using options, we use options written against the futures contract with the nearest expiry date that is at least one month ahead. For example, if, in December, we take a bet with a one month horizon the options we use are options on the March futures contract.

We use the following notation for variables measured at time t : the spot exchange rate (S_t), the one month forward exchange rate (F_t), the price of the futures contract with the nearest expiry date that is at least one month ahead (ϕ_t), the strike price on the options contract, K_t , the settlement price of the call option (C_t), and the settlement price of the put option (P_t). In the description that follows, the variables S_t , F_t , ϕ_t , and K_t are measured

in USD/FCU, while the variables C_t and P_t are measured in USD per foreign currency unit transacted. CME options contracts are quoted in the same units, and settlement prices on the options are provided directly in the data set and do not have to be obtained by converting implied volatilities.

To be concrete about how we construct hedged and unhedged positions using the CME data, consider the following example, where a trader takes a position in January 2006 that expires in February 2006. In February 2006 the third Wednesday was February 15th. Two Fridays prior to the third Wednesday was February 3rd. We therefore look for transactions that were initiated on January 4th 2006 with expiry 30 days later on February 3rd 2006.¹⁵ Suppose we consider a currency for which $F_t > S_t$. In this case, a trader executing the unhedged carry trade sells $1/F_t$ units of the foreign currency forward and obtains the payoff $(F_t - S_{t+1})/F_t$. In our example we measure S_t and F_t on January 4th and S_{t+1} on February 3rd. We take the values of these variables from the Datastream data set described in Appendix A.¹⁶ A trader executing the hedged carry trade takes the same position in the forward market as the unhedged trader and in addition purchases X_t/F_t call options on the foreign currency at strike price K_t . The hedged carry trade payoff gross of the cost of the option is, therefore,

$$\frac{F_t - S_{t+1}}{F_t} + \frac{X_t}{F_t} \max\{\phi_{t+1} - K_t, 0\}.$$

To complete our description of the hedged carry trade we next specify the values of X_t and K_t . We set $X_t = (S_t/F_t)^\delta$ where δ is the number of months (unrounded) between the expiry date of the option (February 3rd 2006) and the expiry date of the underlying future (March 13th 2006). We choose the call option with strike price, K_t , closest to $F_t^\delta S_t^{1-\delta}$. Our choices of X_t and K_t are motivated by two considerations. Since the underlying asset is a futures contract,

¹⁵The only exceptions to this rule for choosing dates is if date t or date $t + 1$ is a holiday with no data available. In this case we shift both dates back one day at a time until the data are no longer missing.

¹⁶Because we have CME options on the Australian dollar dating from 1994, and currency quotes on the Australian dollar sources from WMR are not available on Datastream prior to the end of 1996, we augment our forward and spot exchange rate data for Australia for the period 1994-1996 with data sourced by Datastream from Barclay's (BBAUDSP and BBAUD1F are the mnemonics for the spot rate and one month forward rate).

not the currency spot rate, perfect hedging is not possible unless interest rates are constant between date t and the expiry of the futures contract. However, if interest rates remained constant over this period, covered interest rate parity would imply that $(F_t/S_t)^\delta S_{t+1} = \phi_{t+1}$ and the hedged carry trade payoff gross of the cost of the option would be

$$\frac{F_t - S_{t+1}}{F_t} + \frac{1}{F_t} \max\{S_{t+1} - X_t K_t, 0\}.$$

Thus, when the approximation of constant interest rates holds the hedge is perfect in the sense that if the option is in the money the payoff does not depend on the realization of S_{t+1} . Second, if the strike price is exactly $K_t = F_t^\delta S_t^{1-\delta}$ the payoff to the hedged carry trade gross of the cost of the option is

$$\frac{F_t - S_{t+1}}{F_t} + \frac{1}{F_t} \max\{S_{t+1} - S_t, 0\} = \frac{1}{F_t} \max\{F_t - S_{t+1}, F_t - S_t\},$$

implying that the position in the option on the futures contract is equivalent to an option on a spot contract whose strike price is S_t .

We use one-month eurodollar deposit rates from the Federal Reserve Board interest rate database (H.15) to compute the ex-post prices of the options.

J.P. Morgan Options Data Our second source of options data is J.P. Morgan. We obtained daily one-month at-the-money implied volatility quotes, forward points, and spot exchange rates, for ten currencies against the U.S. dollar. These data are available from January 1995 until July 2009 for the following currencies: Australian dollar, Canadian dollar, Danish krone, euro (beginning January 1999), Japanese yen, Swiss franc, British pound, New Zealand dollar, Norwegian krone, and Swedish krone. In the J.P. Morgan data, “at-the-money” one-month options are at the money forward.

We convert the implied volatility quotes to option prices using the Garman and Kohlhagen (1983) formula in combination with the forward points and spot exchange rate data contained in the same data set. We use the last business day of the expiry month as $t + 1$, and 30 days prior as date t to compute payoffs. If this choice implies that date t is a Saturday, Sunday,

or otherwise missing observation in the data set, we shift both dates back one day at a time until we have a valid pair of business day observations.

VXO Options Data We also use data on options on the S&P 100 index, referred to as the VXO index. These data are available daily from Datastream (mnemonic CBOEVXO) as implied volatilities. We use VXO data rather than VIX data because they are available over a longer sample period (January 1986–July 2009), but the two series behave similarly over the common sample. To generate our monthly data we look for a trade initiation date within each month that is 30 days prior to the third Friday of the following month. We translate implied volatilities to option prices using the Black-Scholes formula. The price of a put option on the S&P 100 index is given by

$$P_t^x = V_t [\Phi(-D_{2t})/(1 + r_t) - \Phi(-D_{1t})/(1 + \delta_t)]$$

where $D_{1t} = (r_t - \delta_t + \frac{1}{24}\sigma_t^2) / (\sigma_t/\sqrt{12})$, $D_{2t} = D_{1t} - \sigma_t/\sqrt{12}$, V_t is the level of the S&P 100 index, the strike price of the option is V_t , δ_t is the dividend yield of the index (on a monthly basis), r_t is the one-month eurodollar rate (on a monthly basis), described above, and σ_t is the implied volatility quote (on an annual basis). We source daily S&P 100 index and dividend yield data from the Global Financial Database (mnemonics OEX, SPY100W).

We measure the unhedged excess return of the S&P 100 index using the total return on the S&P 100 index minus the one-month eurodollar rate (on a monthly basis). We source the total return from the Global Financial Database (mnemonic TRGSPOD). The hedged excess return of the S&P 100 index is

$$\max \left\{ \frac{V_{t+1} - V_t}{V_t}, 0 \right\} + \frac{D_t}{V_t} - r_t - (1 + r_t) \frac{P_t^x}{V_t}.$$

Here D_t/V_t is the total return to the S&P 100 index minus the rate of change of the index. That is, it is the component of the return due to the dividend.

C: Details of the Risk-Factor Analysis

Monthly Risk Factors When working with monthly data, we use nominal payoffs to strategies. 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). Results for additional monthly risk factors, and data sources, are described in the online appendix.

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 (6)) 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:

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

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 $R_s^e = R_s^{e,q}/(1 + \pi_s)$, 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 and 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 and described below in more detail.

When we work with currency options data, monthly payoffs are realized early in each month (the median day of the month is the 6th, with no payoffs occurring before the 2nd of the month, and no payoffs occurring after the 9th of the month). Therefore, when defining the returns for the first quarter we accumulate the monthly payoffs (as described above) that were realized early February, early March and early April. For the second, third and fourth quarters returns are defined analogously.

Quarterly Risk Factors Real per-capita consumption growth (used for the C-CAPM model) 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 extended C-CAPM model adds two factors to the C-CAPM model: the real growth rate of the per-capita service flow from the stock of consumer durables, and the market return (Mkt-Rf plus the risk free rate, Rf, in real terms). To estimate the former 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. We obtain Mkt and Rf from Kenneth French's data library. To obtain the quarterly real market return, we proceed as described above for our currency strategies.

Results for additional quarterly risk factors, and relevant data sources, are described in the online appendix.