

NBER WORKING PAPER SERIES

CARRY TRADE AND SYSTEMIC RISK:
WHY ARE FX OPTIONS SO CHEAP?

Ricardo J. Caballero
Joseph B. Doyle

Working Paper 18644
<http://www.nber.org/papers/w18644>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2012

We thank the National Science Foundation for financial support. The views expressed herein are those of the authors 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Ricardo J. Caballero and Joseph B. Doyle. 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.

Carry Trade and Systemic Risk: Why are FX Options so Cheap?
Ricardo J. Caballero and Joseph B. Doyle
NBER Working Paper No. 18644
December 2012
JEL No. F31,G01,G15

ABSTRACT

In this paper we document first that, in contrast with their widely perceived excess returns, popular carry trade strategies yield low systemic-risk-adjusted returns. In particular, we show that carry trade returns are highly correlated with the return of a VIX rolldown strategy —i.e., the strategy of shorting VIX futures and rolling down its term structure— and that the latter strategy performs at least as well as beta-adjusted carry trades, for individual currencies and diversified portfolios. In contrast, hedging the carry with exchange rate options produces large returns that are not a compensation for systemic risk. We show that this result stems from the fact that the corresponding portfolio of exchange rate options provides a cheap form of systemic insurance.

Ricardo J. Caballero
MIT
Department of Economics
Room E52-373a
Cambridge, MA 02142-1347
and NBER
caball@mit.edu

Joseph B. Doyle
Cornerstone Research
699 Boylston Street, 5th floor
Boston MA 02116
jdoyle@cornerstone.com

1 Introduction

The high returns of the forex carry trade —i.e., investing in high interest rate currencies and funding it with low interest currencies— has led to an extensive literature documenting the “puzzle” and its robustness to a wide variety of controls.¹ This carry trade premium is not explained by traditional risk factors, such as those suggested by Fama and French (1993).² Moreover, several recent works have examined whether carry returns can be explained by crash risk and concluded that it cannot. Most prominently, Burnside, Eichenbaum, Kleshchelski et al. (2011) find that hedging the carry with ATM FX options leaves its returns unchanged, and therefore conclude that the crash risk exposure is not the source of the premium.³

The main contribution of this paper is to turn the puzzle on its head. We reconcile the past findings by showing that while the standard carry trade *is essentially compensation for systemic risk*, the corresponding bundle of *crash protection FX options are puzzlingly cheap*. In particular, we show that carry trade returns are highly correlated with the returns of a VIX rolldown strategy —i.e., the strategy of shorting VIX futures and rolling down its term structure— for individual currencies as well as for diversified portfolios.⁴ We find that while typical carry trade strategies produce large returns, this is explained by its comovement with VIX rolldowns. On the other hand, portfolios of exchange rate options designed to hedge the carry provide a cheap form of systemic risk insurance. As a result, when the carry trade is hedged with exchange rate options, its average return remains strongly significant even after controlling for its exposure to VIX rolldowns.

¹The profitability of the carry trade strategy stems from the fact that high interest rate currencies tend to appreciate rather than depreciate, in contrast with the most basic implication of the uncovered interest parity condition.

²See, e.g., Table 4 from Burnside, Eichenbaum, Kleshchelski et al. (2011).

³Farhi et al. (2009) uses currency options data to estimate that crash risk may account for roughly 25% of carry returns in developed countries, leaving plenty of the carry return unexplained.

⁴The VIX is an S&P500 implied volatility index, which is often described as the “global (financial) fear” indicator.

As a preview of our results, we run regressions of the form

$$\bar{z}_t = \alpha + x_t\beta + e_t$$

where \bar{z}_t are the excess returns to a carry trade portfolio and x_t are the excess returns to VIX rolldowns.⁵ Figure 1 plots the cumulative returns to carry strategies that place equal weights on each of the 25 countries in our sample against the required returns based on their exposure to VIX rolldowns.⁶ Panel A shows the results for the standard carry, and Panel B is for the carry when hedged with at-the-money exchange rate options. Each panel includes two series: $\prod_{j=1}^t (1 + \bar{z}_j) - 1$ (“Realized”) and $\prod_{j=1}^t (1 + x_j \hat{\beta}_{j-1}) - 1$ (“Required”). While the unhedged carry does only marginally better than its required returns, the hedged carry beats its systemic counterpart by more than a factor of 10. For the unhedged carry, the annualized Sharpe ratio is 0.47, below the 0.57 of its required returns. But for hedged carry, the realized Sharpe ratio of 0.79 is nearly 40% higher than the 0.57 earned by its required returns. This suggests that exchange rate options earn a significant premium above that which would be required based on their systemic exposure. Indeed, we find that after controlling for its systemic exposure, the portfolio of carry protection options alone earns a Sharpe ratio of 1.19, and we thus conclude that such a portfolio of exchange rate options is a cheap form of systemic insurance.

⁵Our data are monthly from March 2004 (the starting date for VIX futures) to August 2012.

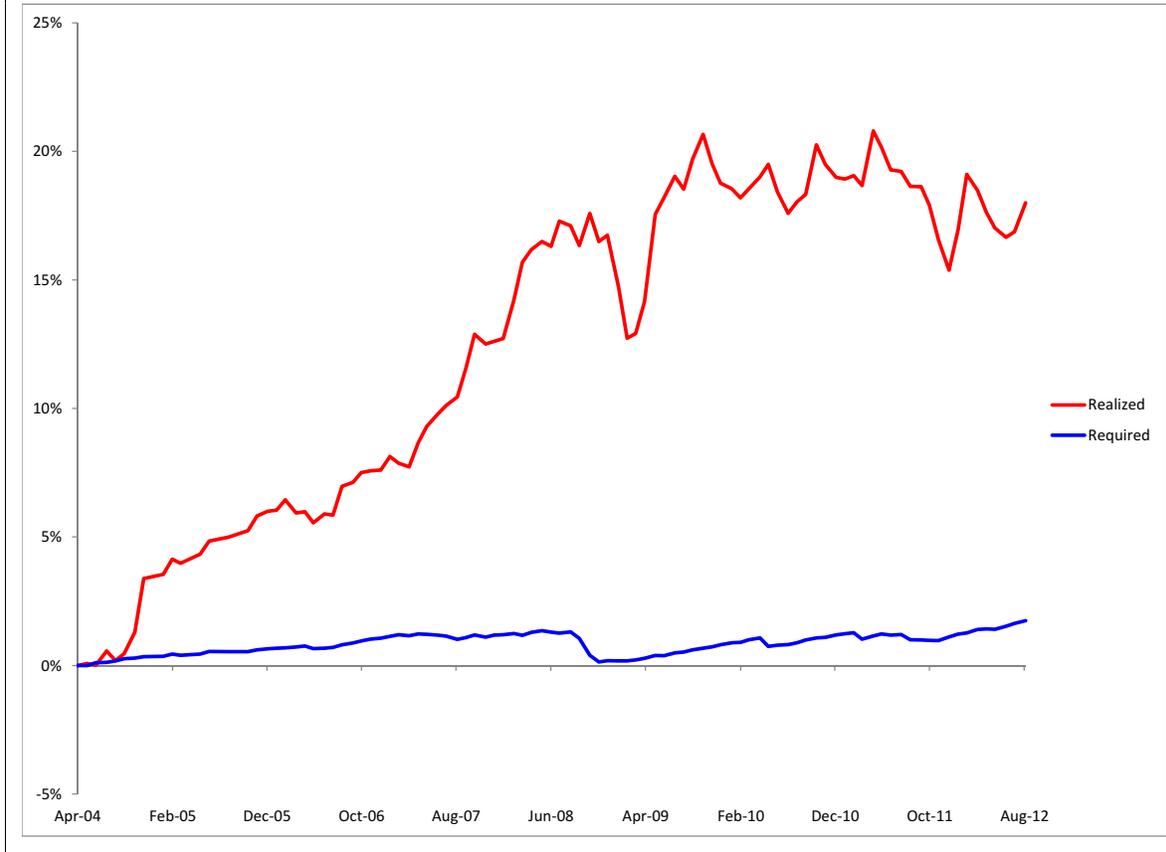
⁶ $\bar{z}_t = \frac{1}{N} \sum_{i=1}^N \text{sign}(r_{i,t} - r_{US,t}) z_{i,t}$ where $z_{i,t}$ are the (either hedged or unhedged) excess returns to the carry for currency i using USD as the base currency.

Figure 1: Cumulative Returns

Panel A: Unhedged



Panel B: Hedged



Our work is directly related to a long literature that documents the puzzlingly high returns to the carry trade. This literature includes dozens, if not hundreds, of papers. Perhaps the most influential paper is Fama (1984) while Engel (1996) provides an excellent survey. Burnside, Eichenbaum, Kleshchelski et al. (2011) confirms that these findings hold in more recent data and with updated models and risk-factors. Focusing on the cross-section, Lustig and Verdelhan (2007) show that low interest rate currencies provide a hedge for domestic consumption growth risk, and this can explain why these currencies do not appreciate as much as the interest rate differential. However, Burnside (2007) shows that such a model does not have significant β 's and leaves a large unexplained intercept term,

both of which are primary objects of our analysis.

Our findings also relate to a more recent literature that points to crash risk as a potential explanation for average carry returns. Jurek (2009) and Burnside, Eichenbaum, Kleshchelski et al. (2011) find that carry returns remain high even if one purchases crash protection via put options. Brunnermeier et al. (2009) shows that there is a strong relationship between carry return skewness and interest rate differentials, providing evidence that sudden exchange rate moves may be due to the reduction in carry positions as traders near constraints. Further, they find that increases in VIX coincide with reductions in carry positions, and a higher level of VIX predicts higher future carry returns. Focusing on the cross-section, Menkhoff et al. (2011) shows that there is a strong relationship between interest rates and global volatility measures. Finally, Lustig et al. (2010) shows that US business cycle indicators help predict foreign exchange returns.

The remainder of the paper is organized as follows. Section 2 introduces notation and discusses the basics of the carry trade. Section 3 describes the details of the theory that underlies our empirical approach. Section 4 provides a description of our data. Section 5 describes our results for the period leading up to the 2008 crisis and section 6 presents the results for the entire sample.

2 Carry Trade Basics

In this section, we describe standard carry trade strategies and introduce notation. For now, it is convenient to imagine that there are only two countries: the US (domestic) and foreign. Both countries issue sovereign risk-free bonds and have independent currencies. In the US, bonds pay an interest rate of r_t and abroad they pay r_t^* . Imagine also that there are only two periods, today and tomorrow. Let S_t be the exchange rate today and F_t be the forward exchange rate at which one can agree to make a currency exchange tomorrow. Both S_t and F_t are in terms of US dollars (USD) per foreign currency units (FCU). Without loss

of generality assume that $r_t < r_t^*$. An obvious trading strategy for a US investor today is to borrow \$1 at the low domestic interest rates, convert the borrowed USD to FCU, and lend those FCU to a foreigner at the higher interest rate. Since the investor borrows domestically, USD is referred to as the funding currency while FCU is the investment currency. Today the investor will have no net cash flows since he immediately lends all of his newly borrowed funds. Tomorrow the foreigner will pay him $\frac{(1+r_t^*)}{S_t}$ in FCU, and he will pay back $(1+r_t)$ in USD. Since both interest rates are risk-free, the investor only faces risk from changes in the exchange rate. In order to perfectly hedge against FX risk, he could simply purchase a forward contract. Doing that, tomorrow he would receive the risk-free payment of $F_t \frac{(1+r_t^*)}{S_t}$ in USD. Since this strategy is both costless and riskless, any profits would be pure arbitrage. Therefore the following arbitrage condition is expected to be satisfied most of the time:⁷

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

This is known as the covered interest parity condition (CIP). In particular, it states that

$$f_t - s_t = \tilde{r}_t - \tilde{r}_t^*$$

where $f_t = \log(F_t)$, $s_t = \log(S_t)$, $\tilde{r}_t = \log(1+r_t)$, and $\tilde{r}_t^* = \log(1+r_t^*)$. The quantity $f_t - s_t$ is commonly referred to as the average forward discount, and we write $AFD_t = f_t - s_t$.

The carry trade is a simple risky variant on this strategy where the investor elects not to purchase the forward FX contract in the hope that the USD will not appreciate by enough to eliminate his entire profits. Denote the net payoff from this strategy by

$$\Pi_{t+1} = S_{t+1} \frac{(1+r_t^*)}{S_t} - (1+r_t)$$

⁷In practice, this condition can break down during episodes of severe financial distress.

Using CIP, this can be rewritten as

$$\Pi_{t+1} = \frac{(1 + r_t^*)}{S_t}(S_{t+1} - F_t)$$

That is, the carry trade strategy is equivalent to buying $\frac{(1+r_t^*)}{S_t}$ FCU forward. The investor will make profits on average as long as $E_t(S_{t+1}) > F_t$, and will only lose money if the USD appreciates to the point that $S_{t+1} < F_t$. In what follows, we will focus on the simpler and more commonly employed version of the carry trade where the investor buys 1 FCU forward. We define the excess return to this strategy as

$$z_{t+1} = \frac{S_{t+1}}{F_t} - 1.$$

3 Risk-based Explanations

As discussed in the introduction, empirical research has found the carry trade to produce puzzlingly high returns to investors. In principle, these returns could simply represent compensation for some form of risk, but it has proven difficult to identify which risks are relevant. In this section we describe a simple and standard framework for assessing risk-based explanations, which we will rely upon for our analysis.

Standard models of asset prices can be reduced to the specification of an asset pricing kernel, M_{t+1} , such that

$$E_t[M_{t+1}R_{t+1}] = 0$$

holds for any excess return R_{t+1} denominated in USD. Plugging in $z_{t+1} = R_{t+1}$, we see that

$$\frac{E_t[M_{t+1}S_{t+1}]}{E_t[M_{t+1}]} = F_t$$

which implies

$$E_t[S_{t+1}] = F_t - \frac{Cov_t[M_{t+1}, S_{t+1}]}{E_t[M_{t+1}]}.$$

The historical literature, surveyed by Engle (1996), focuses on the special case where $Cov_t[M_{t+1}, S_{t+1}] = 0$. In that case, the forward rate is an unbiased predictor of the future spot rate, and the average excess carry return is zero. Both of these hypotheses have been consistently rejected empirically where researchers have estimated the regression

$$E_t\left[\frac{S_{t+1} - S_t}{S_t}\right] = \alpha + \beta \frac{F_t - S_t}{S_t}$$

and rejected the hypothesis that $\alpha = 0$ and $\beta = 1$. This is commonly referred to as the forward premium puzzle. In fact, estimates of β are often negative, implying that high interest rate currencies actually tend to appreciate, exactly the opposite from what would be expected.

Allowing for non-zero correlation between exchange rates and the pricing kernel opens the door for risk-based explanations. Rewriting things slightly, one can see that

$$E[Z_{t+1}] = \left(\frac{Cov[M_{t+1}, z_{t+1}]}{Var[M_{t+1}]} \right) \left(-\frac{Var[M_{t+1}]}{E[M_{t+1}]} \right)$$

or

$$E[z_{t+1}] = \tilde{\beta} \tilde{\lambda}$$

where $\tilde{\beta}$ is the slope coefficient from the regression of z_t on M_t and $\tilde{\lambda} = -\frac{Var[M_{t+1}]}{E[M_{t+1}]}$. Next if $M_{t+1} = a + bx_{t+1}$ where x_{t+1} is an excess return, then

$$\tilde{\lambda} = bE[x_{t+1}].$$

Therefore,

$$E[z_{t+1}] = \tilde{\beta}\tilde{\lambda} = \beta E[x_{t+1}]$$

where $\beta = b\tilde{\beta}$. In other words, the model predicts that if one runs the regression of $z_t = \alpha + \beta x_t + \epsilon_t$, one should find that $\alpha = 0$.⁸ Here α is the portion of the excess return z_t that remains after controlling for its exposure to the factors x_t . The quantities $z_t - \hat{\beta}x_t$ are commonly referred to as the pricing errors from the model where $\hat{\beta}$ is the coefficient estimated by OLS. The goal then becomes identifying a set of excess returns x_t that have this property. The classical example of this approach is Fama and French (1993) which shows that a set of three factors that do a very good job of pricing US equities. In this project, we consider using excess returns on VIX futures to price currency forwards.

4 Data

We obtained daily closing spot and forward exchange rates from Datastream for 67 countries in terms of USD/FCU. Data on VIX futures prices and the underlying index were also taken from Datastream. Our sample covers the period from March 26, 2004 when VIX futures started trading, to August 24, 2012. While FX forward rates trade for each horizon on every day, VIX futures only trade for fixed maturity dates. For each VIX futures contract, we find the trading day such that the maturity date of the contract is the same as the maturity date for the FX forward rates.⁹ Since there is never more than one VIX futures expiration date in a single month, this procedure creates approximately non-overlapping holding periods at the one month horizon. In total, we have 96 observations for each country. We also employ daily data on the Fama-French three factors for the same period, which were obtained from Kenneth French's website. Daily returns are compounded to match the holding periods

⁸Tests of this form were first proposed by Black, Jensen, and Scholes (1972).

⁹In cases when there is no exact match, we use the trading day such that the maturities are as close together as possible. They never differ by more than one day.

under consideration. Our primary analysis focuses on a set of relatively developed countries, the Expanded Majors as defined by Bloomberg, and we exclude the countries from this set which do not have floating exchange rates.¹⁰ This leaves us with a total of 25 countries in our primary sample.¹¹

We also collected daily data on the implied volatilities of at-the-money (spot) exchange rate options with a one month horizon from Datastream for 22 out of the 25 countries in our primary sample.¹² We convert implied volatilities into prices using the Black-Scholes model. Our options data covers the same period as our other data, and thus contains 96 monthly observations.¹³

We consider two different periods of time for our analysis. For comparability with the literature, our primary period of interest includes only the dates in our sample prior to the financial crisis of 2008 (pre-crisis).¹⁴ We also focus on this period because in all likelihood the full sample dramatically overstates the probability of a catastrophic downturn in financial markets. However, we will show that the results from the pre-crisis period carry over qualitatively into the full sample period.

5 Carry Trade and Systemic Risk: Pre-Crisis

5.1 Traditional Carry Trade Portfolios

In this section, we begin by presenting results for the typical diversified carry trade strategies that have been the focus of much of the recent literature. These are portfolios of

¹⁰We follow the IMF’s classification of currencies, which is available at its website. We include only currencies which are classified by the IMF as either “Independently floating” or “Managed Floating with no pre-determined path for the exchange rate.”

¹¹Our primary sample of currencies includes AUD, BRL, CAD, CHF, CLP, COP, CZK, EUR, GBP, HUF, IDR, ILS, INR, KRW, MXN, NOK, NZD, PEN, PLN, SEK, SGD, TRY, TWD, USD, ZAR.

¹²The three excluded currencies are CZK, HUF, and PEN. Options pricing data for NZD were not available after February 2012.

¹³Implied volatility data for TRY is missing for the first two observations in the sample.

¹⁴Specifically, we include monthly holding periods ending on May 19, 2004 through August 20, 2008 and only for those months when VIX futures contracts expired. This leaves a total of 48 observations.

dynamically optimized carry strategies that are long a currency’s forward whenever $r_t^* > r_t$ and short otherwise.¹⁵ We follow Burnside, Eichenbaum, Kleshchelski et al. (2011) and others in using equally weighted baskets (EQL) with USD as the base currency. However, as pointed out by Jurek (2009), using interest rate spreads as portfolio weights (SPD) tends to produce higher returns, and the results may vary substantially for different base currencies. Therefore, we also consider equal- and spread-weighted portfolios of JPY-based carry and currency neutral carry.¹⁶ Our primary goal is to examine to what extent these returns can be explained by the excess returns to VIX futures. To that end, we run regressions of the form

$$\bar{z}_t = \alpha + x_t\beta + e_t$$

where $\bar{z}_t = \sum_{i=1}^N \omega_i \text{sign}(r_{i,t}^* - r_t) z_{i,t}$ is the excess returns to the carry trade portfolio, i indexes currencies, ω denotes portfolio weights, and x_t is the excess returns to a long position in VIX futures.

Table 1 summarizes the results from these regressions for the pre-crisis period. In every case, the excess carry trade returns were highly significant, consistent with the findings of Burnside, Eichenbaum, Kleshchelski et al. (2011). However, the exposure of these strategies to systemic risk is evident from the also highly significant values of β in all regressions. Importantly, after correcting for this risk exposure, we find little excess returns, as the estimated α 's are statistically indistinguishable from 0 in all but one case.¹⁷

¹⁵We assess whether $r_t^* > r_t$ by looking at a currency’s average forward discount.

¹⁶Following Jurek (2009), currency neutral carry portfolios are computed in the following way. First, two sub-portfolios are formed one containing only those currencies with corresponding interest rates higher than that of the US and the other with only lower interest rate currencies. The final portfolio is equally weighted in these two sub-portfolios. The portfolio weighting scheme (EQL or SPD) refers to the weights used to construct the sub-portfolios.

¹⁷In the Appendix, we show that the results are robust to controlling for the Fama-French factors, and those traditional risk factors alone cannot explain the carry, in line with the conclusions of Burnside, Eichenbaum, Kleshchelski et al. (2011).

Table 1: Carry Trade Returns and Systemic Risk (Pre-Crisis)

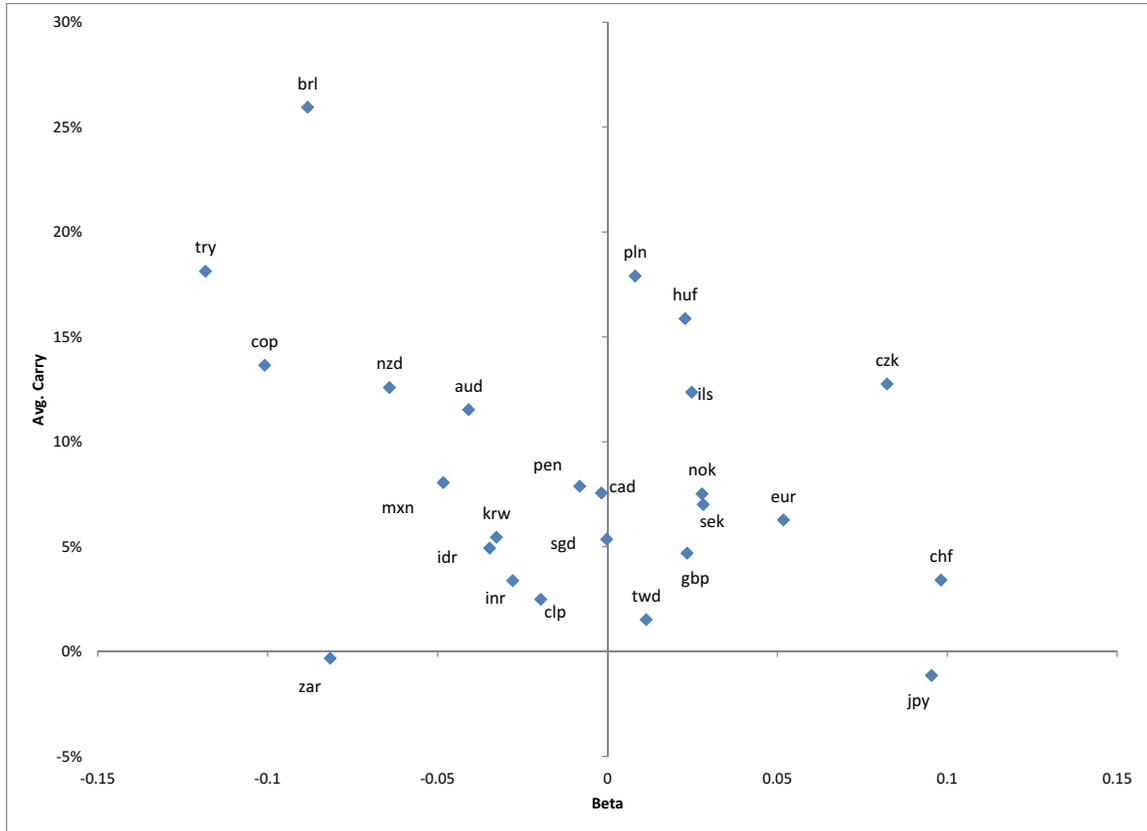
Currency	USD		JPY		Neutral	
Weights	EQL	SPD	EQL	SPD	EQL	SPD
Avg. Carry	4.86%***	12.07%***	10.63%**	15.42%***	3.79%***	7.92%***
	(2.85)	(4.51)	(2.31)	(2.82)	(2.76)	(3.62)
α	2.56%	7.86%**	2.70%	5.46%	0.94%	4.03%
	(1.28)	(2.62)	(0.56)	(0.95)	(0.61)	(1.58)
β	-0.028***	-0.052***	-0.097***	-0.122***	-0.035***	-0.048***
	(-2.76)	(-3.14)	(-3.37)	(-3.77)	(-4.09)	(-3.33)

Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

These findings also carryover to individual currency pairs.

Figure 2: Avg. Carry vs. Beta (Pre-Crisis)

Panel A: USD



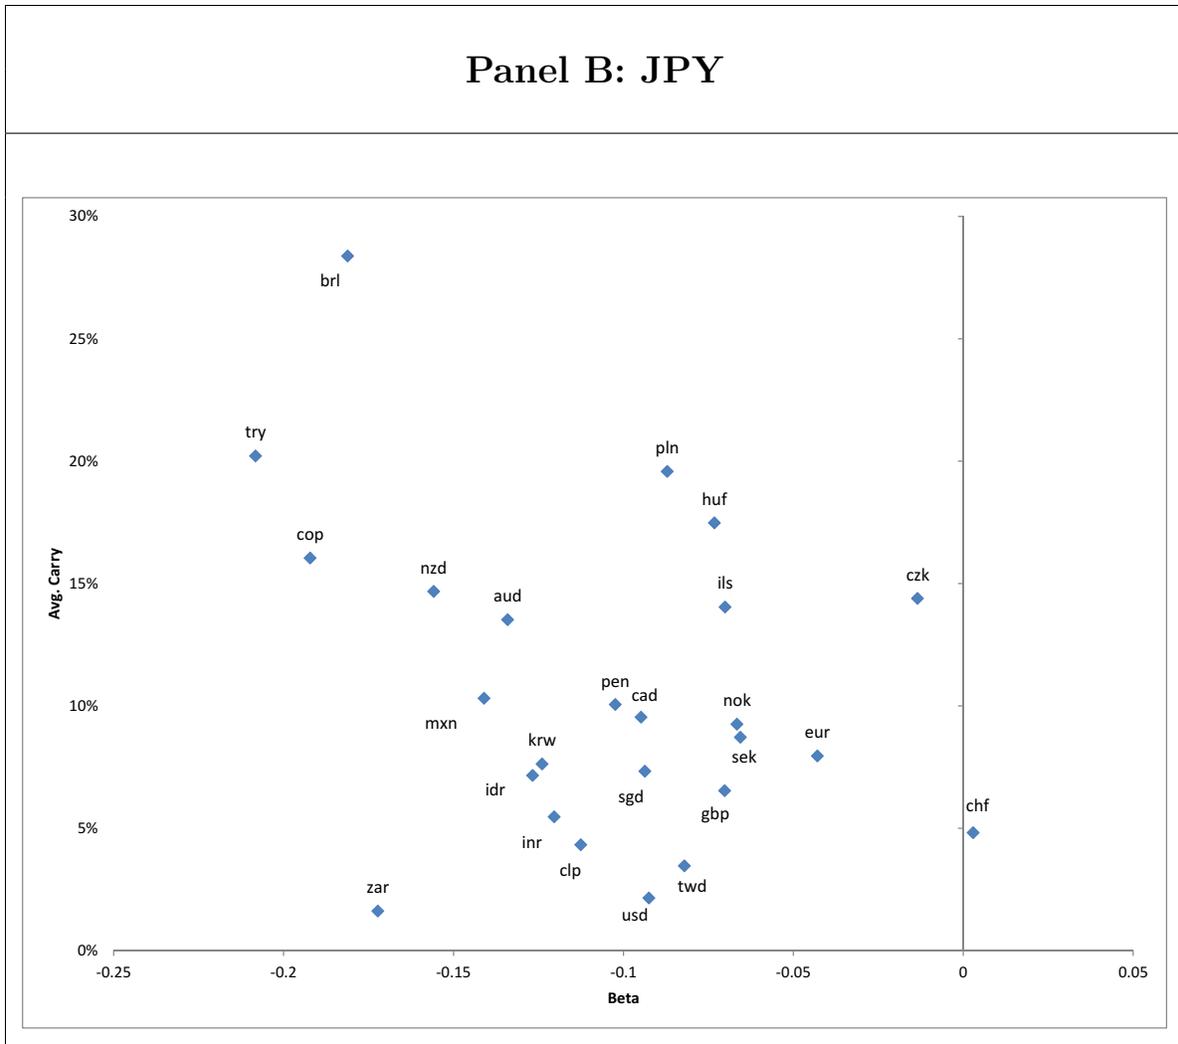
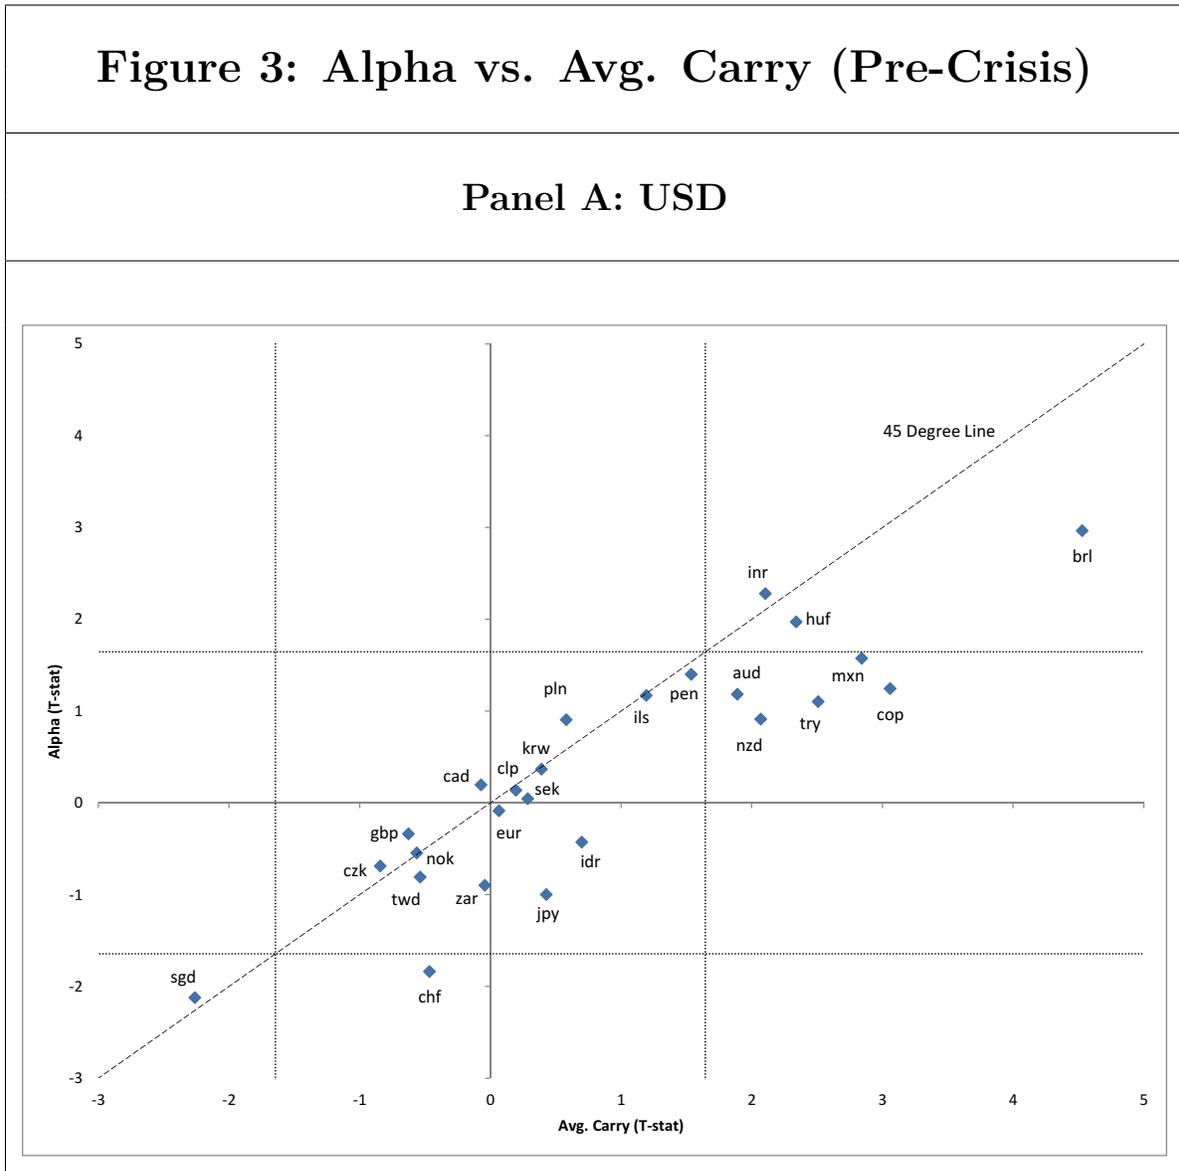


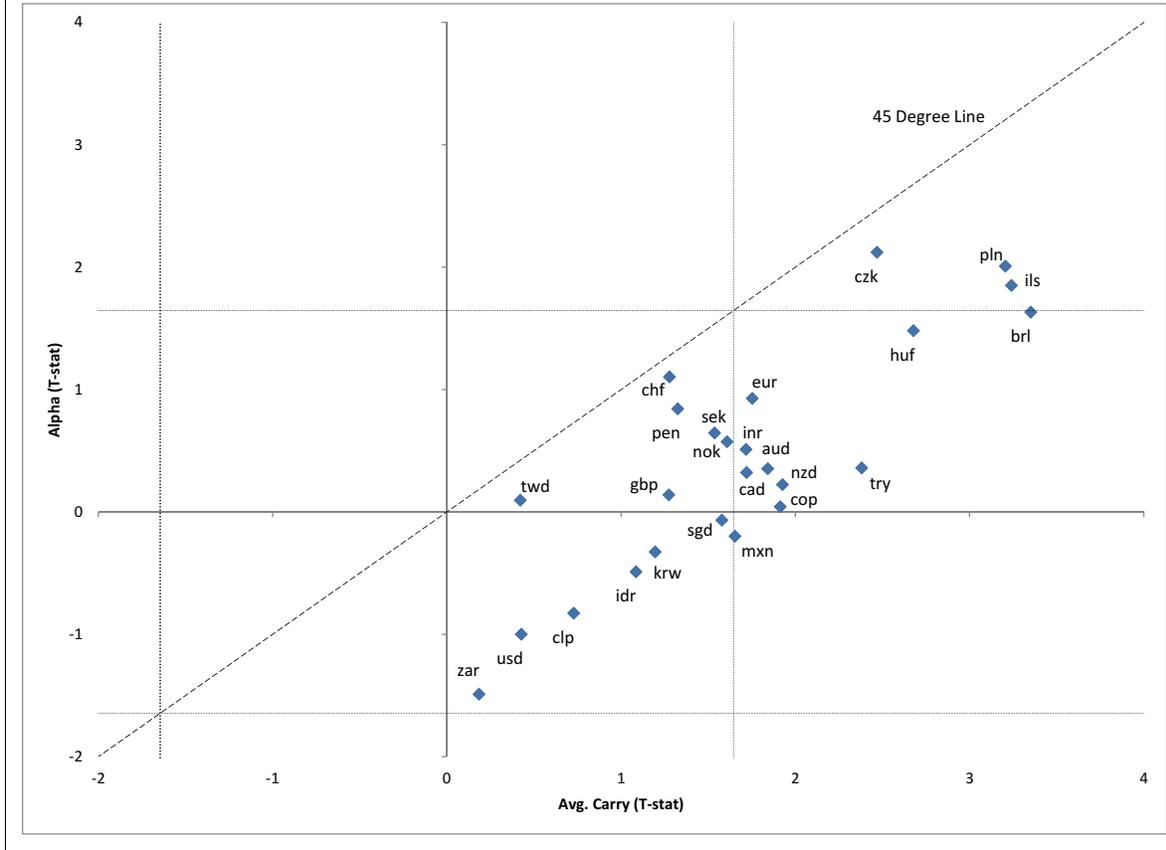
Figure 2 plots the average carry returns against the β_i 's for positions that are all short the base currency (ie, the base currency is used as the funding currency). Note that this differs from the main strategy that we consider which determines the investment currency dynamically depending on the sign of the interest rate differential. Currencies which should have been funding currencies against the base show up here with positive β_i 's. The evidence that carry is closely related to systemic risk exposure is visible.

Figure 3 plots the T-statistics for average carry returns against the corresponding α 's for each country in our sample. Here, funding currencies are determined dynamically by the sign of the interest rate differential. Results are shown separately when USD and JPY are

used as the base currency. In line with our portfolio regressions, the α 's are almost uniformly below the average carry returns across currencies. The pattern is most easily seen when JPY is used as the base where all currencies have positive carry returns with 13 significant at the 10% level, but only 3 α 's are significantly above 0 while 7 are actually negative.



Panel B: JPY



5.2 Options Hedged Carry Trade

Up to now, we have shown that hedging the carry trade with a long position in VIX rolldowns leaves investors with no significant excess returns. Next, we consider the natural alternative of hedging with exchange rate options. The strategy we examine involves purchasing an ATM call option on any exchange rate where the carry trader holds a short forward position and an ATM put on exchange rates where the carry trader is long. For 4 currencies in our sample (AUD, EUR, GBP, NZD), our options data is provided for the

USD/FCU rate. In those cases, we compute the options hedged carry trade returns as¹⁸

$$z_t^h = \begin{cases} \frac{S_{t+1} - F_t + \max\{0, S_t - S_{t+1}\} - (1+r)P_t}{F_t} & \text{if } S_t > F_t \\ \frac{F_t - S_{t+1} + \max\{0, S_{t+1} - S_t\} - (1+r)C_t}{F_t} & \text{if } S_t < F_t \end{cases}.$$

In all other cases, the options data is for the FCU/USD rate, so we compute returns as¹⁹

$$z_t^h = \begin{cases} \frac{S_{t+1}^{-1} - F_t^{-1} + \max\{0, S_t^{-1} - S_{t+1}^{-1}\} - (1+r^*)P_t}{S_{t+1}^{-1}} & \text{if } S_t < F_t \\ \frac{F_t^{-1} - S_{t+1}^{-1} + \max\{0, S_{t+1}^{-1} - S_t^{-1}\} - (1+r^*)C_t}{S_{t+1}^{-1}} & \text{if } S_t > F_t \end{cases}.$$

Our primary analysis involves the same regressions as before, but replacing z_t with z_t^h as the dependent variable. Table 2 displays the results. Consistent with the findings of Burnside, Eichenbaum, Kleshchelski et al (2011), hedging with FX options leaves a significant carry return. In fact, the T-statistics increase by over 70% on average. Furthermore, such hedging removes much of the trades' systemic exposure, reducing β by over 60% on average and leaving significant α 's.

¹⁸This is the same formula used to compute FX options hedged carry by Burnside, Eichenbaum, Kleshchelski et al (2011).

¹⁹We compute this expression with S_{t+1}^{-1} in the denominator so that removing the options components leaves the same unhedged returns from before.

Table 2: Options Hedged Carry Trade Returns (Pre-Crisis)

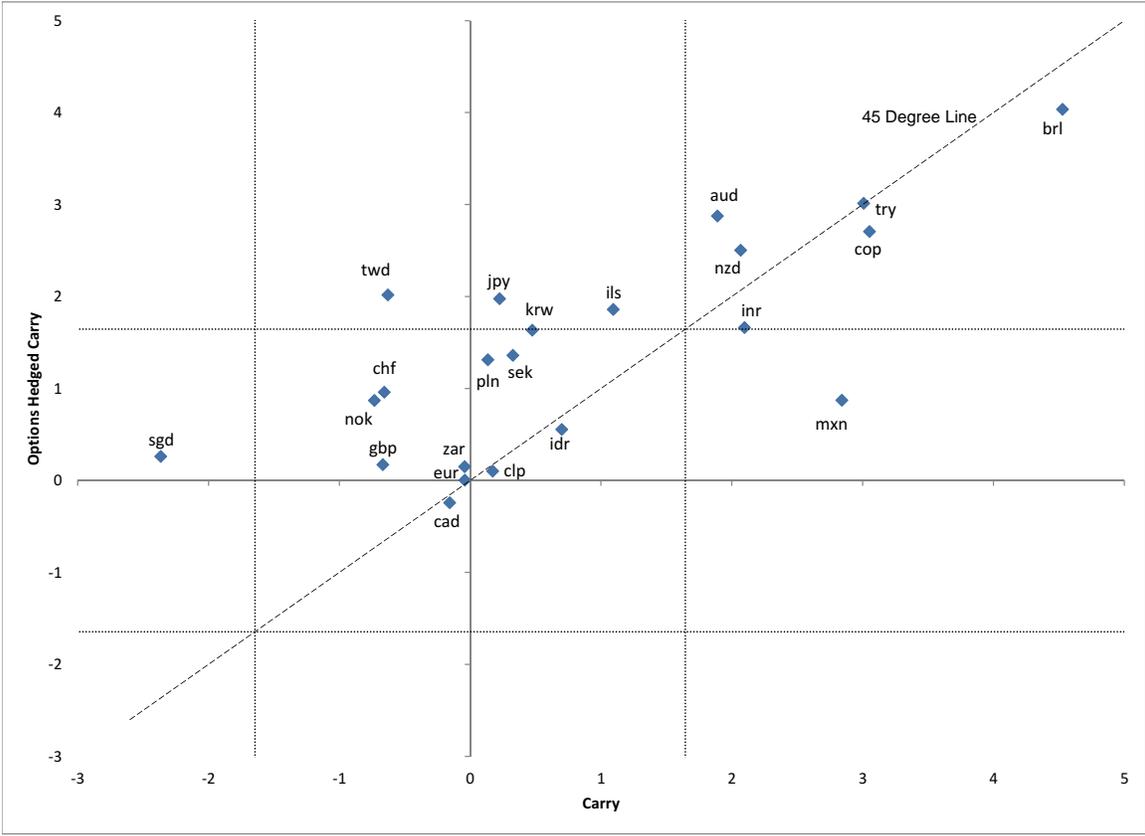
Currency	USD		Neutral	
Weights	EQL	SPD	EQL	SPD
Avg. Carry	4.47%** (2.61)	10.69%*** (3.81)	3.38%** (2.30)	6.77%*** (3.04)
Hedged Carry	4.22%*** (4.57)	7.24%*** (4.96)	3.85%*** (4.81)	5.70%*** (5.17)
α	3.53%*** (3.69)	5.43%*** (4.54)	2.89%*** (4.05)	4.11%*** (4.73)
β	-0.008 (-1.64)	-0.022*** (-3.08)	-0.012*** (-2.57)	-0.019*** (-3.24)

Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

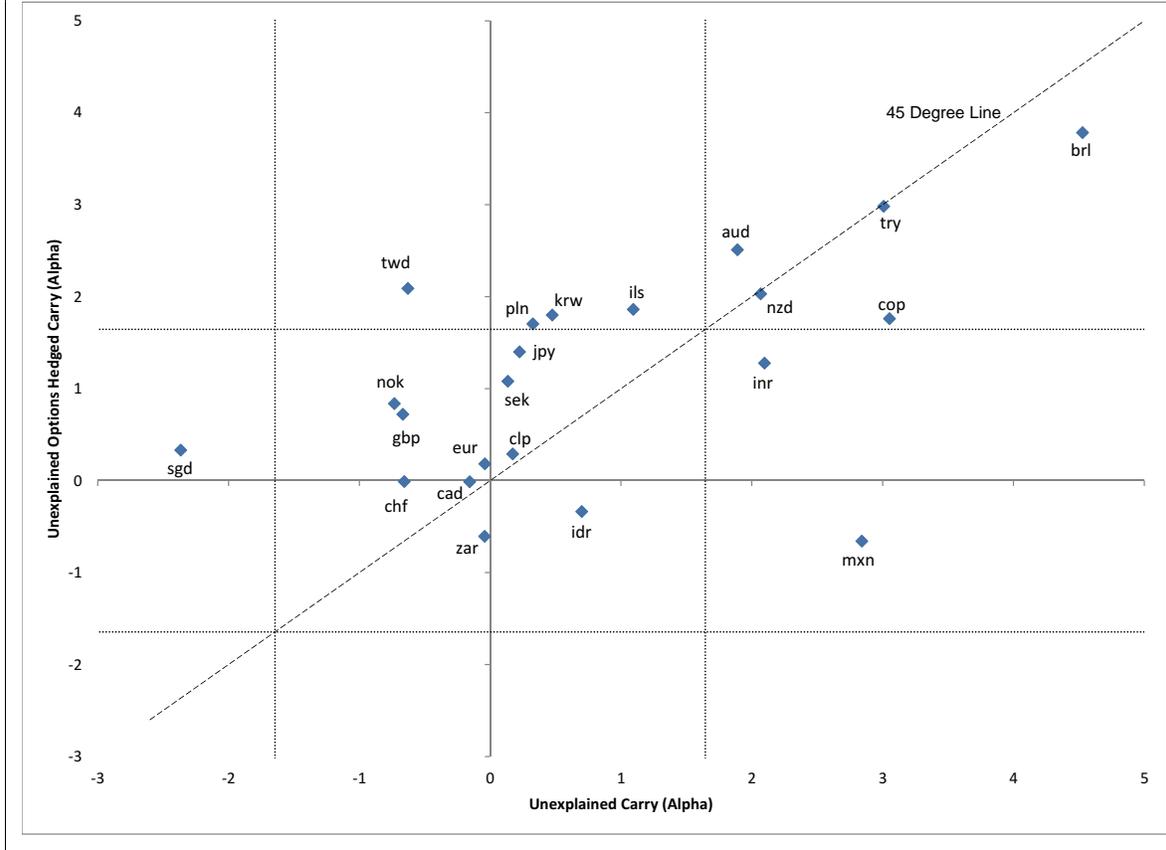
Figure 4 shows that the results hold for individual currencies as well. The T-statistics for all options hedged returns are either close to or above those of their unhedged counterparts with the sole exception of MXN. The same is true in $\alpha - space$.

Figure 4: Hedged vs. Unhedged Carry (Pre-Crisis)

Panel A: Returns (T-stats)



Panel B: Alphas (T-stats)

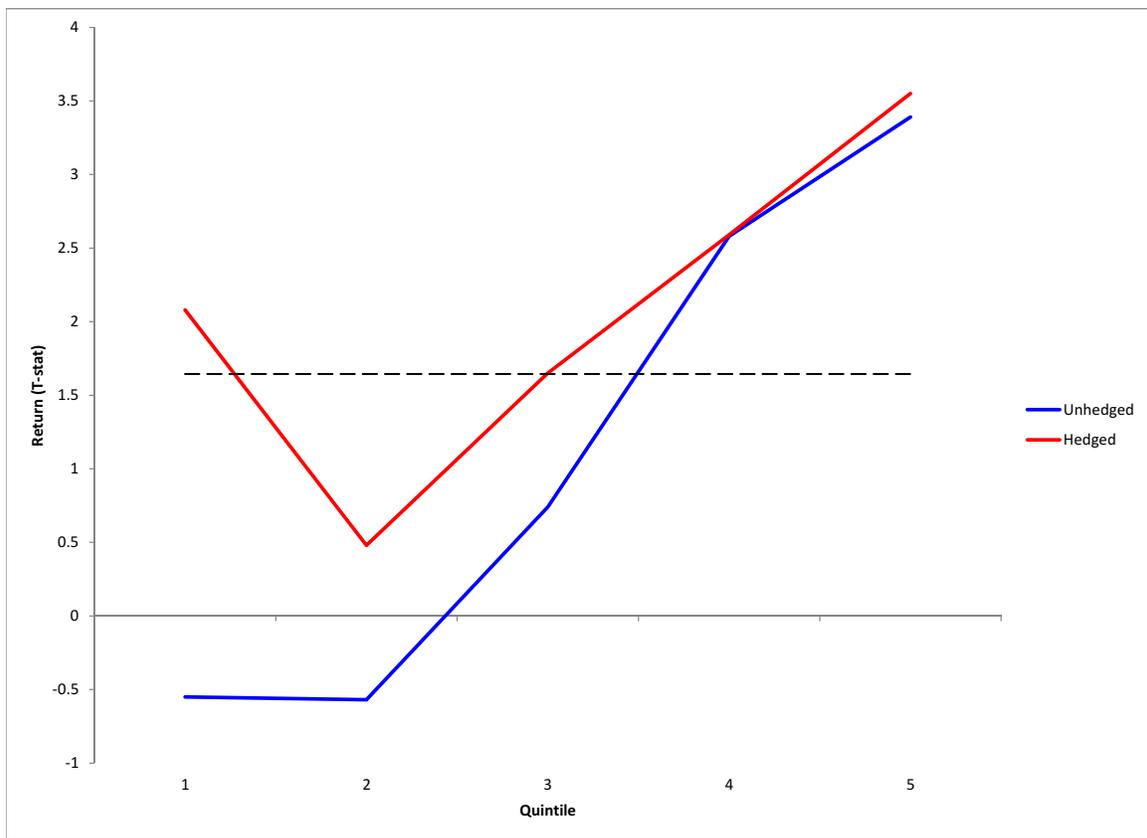


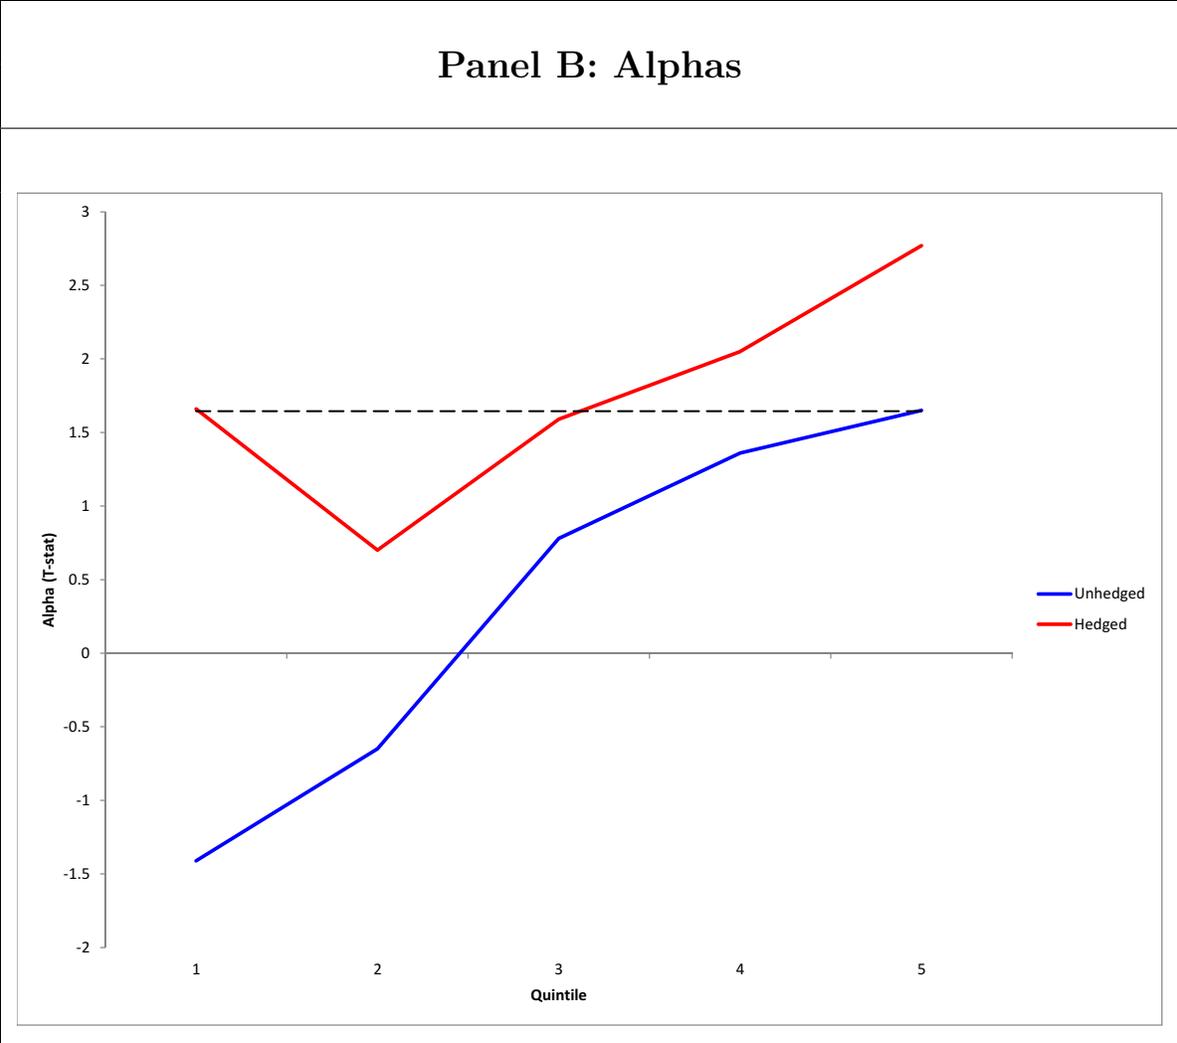
We dissect these results further by separating currencies into interest rate quintiles that are rebalanced every period. Figure 5 illustrates that the gains from hedging with options comes primarily from the low interest rate quintiles. For the two lowest quintiles, the carry trade returns were actually negative on average, but positive payoffs from options were enough to outweigh those losses. Meanwhile, at the two highest quintiles, Sharpe ratios were unchanged by hedging with options. In the α – space, the gains still come primarily at the lowest quintiles where α 's go from negative to positive. However, the highest quintiles also experienced some gains via a reduction in systemic risk exposure. Since the lowest quintile is comprised entirely of short positions in foreign currencies, we find that call options on

funding currencies are especially cheap as both overall (Panel A) and systemic (Panel B) insurance.

Figure 5: Carry Returns by Quintile (Pre-Crisis)

Panel A: Returns





5.3 A New Puzzle

Given our earlier results on the systemic exposure of unhedged carry returns, these findings on the options-hedged carry suggest that a strategy involving only foreign exchange options designed to hedge carry returns may provide a cheap form of systemic insurance. We test this claim in the context of our model by regressing the excess returns to FX options

on the excess returns to VIX rolldowns where the dependent variable is computed as

$$z_t^{opt} = \begin{cases} \frac{\max\{0, S_t - S_{t+1}\} - P_t}{P_t} - r & \text{if } S_t > F_t \\ \frac{\max\{0, S_{t+1} - S_t\} - C_t}{C_t} - r & \text{if } S_t < F_t \end{cases}$$

for those options quoted in terms of USD/FCU and

$$z_t^{opt} = \begin{cases} \left(\frac{S_{t+1}}{S_t}\right) \left(\frac{\max\{0, S_t^{-1} - S_{t+1}^{-1}\} - P_t}{P_t}\right) - r & \text{if } S_t < F_t \\ \left(\frac{S_{t+1}}{S_t}\right) \left(\frac{\max\{0, S_{t+1}^{-1} - S_t^{-1}\} - C_t}{C_t}\right) - r & \text{if } S_t > F_t \end{cases}$$

for those options quoted in FCU/USD. Estimates are contained in Table 3 for various portfolios of these options. In every case, the portfolio of options is significantly *positively* correlated with a long position in VIX rolldowns. This is not surprising, as the options strategies are designed to hedge the carry. However, the important result is that these portfolios have very large α 's that are strongly significant in all cases, indicating that they do indeed provide (excessively) cheap systemic insurance.

Table 3: FX Options Returns and Systemic Risk (Pre-Crisis)

Currency	USD		Neutral	
Weights	EQL	SPD	EQL	SPD
α	452.40%***	263.50%	584.27%***	536.35%***
	(2.78)	(1.37)	(3.54)	(2.75)
β	2.37**	3.32***	2.79***	3.38***
	(2.55)	(2.98)	(3.12)	(3.10)

Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

That is, the new puzzle is that there is premium to selling systemic risk insurance, hedged with currency options. The conventional carry trade is a form of selling systemic insurance,

which when hedged with FX-options generates an excess return. But the source of the excess return is in the low cost of the hedge, not the high return of the carry itself.

6 Carry Trade and Systemic Risk: Full Sample

Having examined the properties of typical carry trade strategies when hedged with either VIX futures or FX options during (relatively) normal times, it remains to be seen how these hedges perform in a sample that contains a major financial turndown. In this section, we explore this issue by using our full sample period that extends to August 2012 and includes the major asset market collapse surrounding the bankruptcy of Lehman Brothers in 2008. Although there are some nuances, the core message of the previous section remains unchanged.

Table 4 presents our factor regressions for the unhedged carry over this period. As noted earlier, our full sample overstates the empirical frequency of crises, and the result is that we do not see the strikingly positive carry trade returns as in the pre-crisis period or as has been documented for longer periods of time in previous work. On the other hand, as in the pre-crisis period, the exposure of the carry to systemic risk remains strongly significant in all cases, and taking this exposure into account leaves the carry with broadly reduced returns and T-statistics. The JPY based carry even underperforms relative to its systemic exposure. Figure 6 shows that these results are true at the currency level as well where 21 out of 25 JPY based and 11 out of 25 USD based currency carry trades underperform.

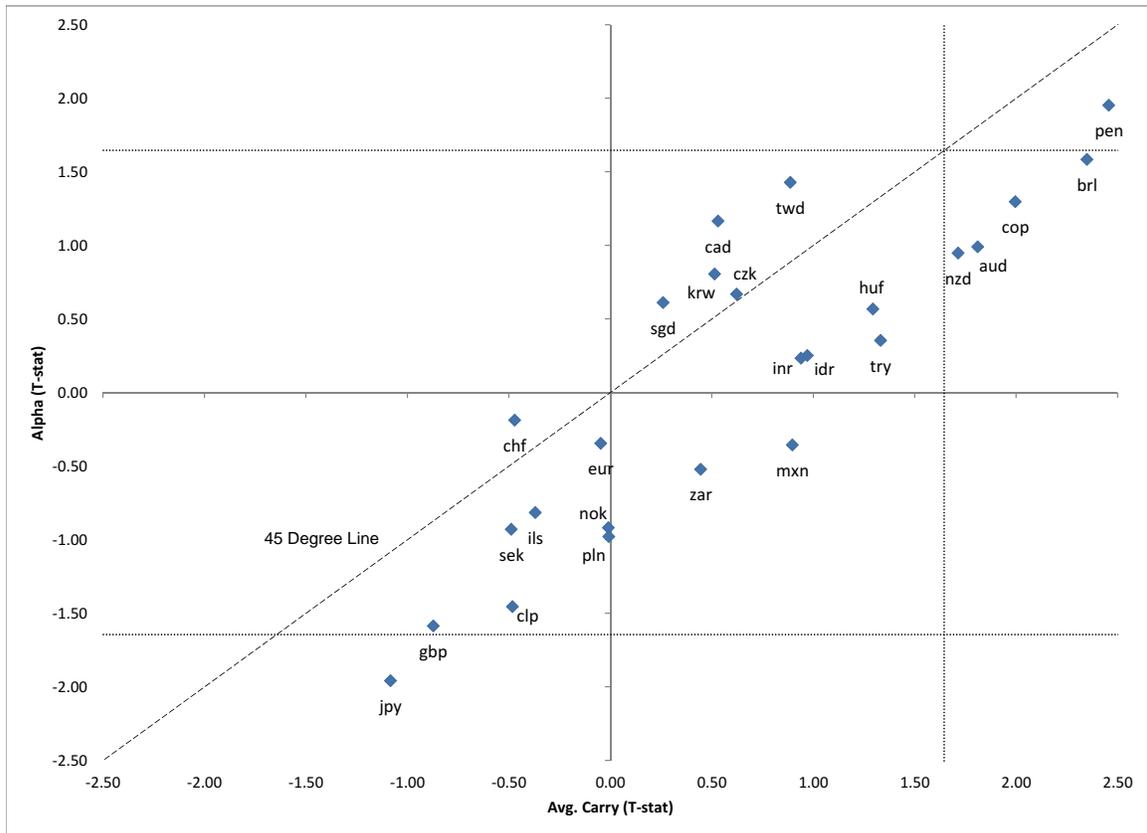
Table 4: Carry Trade Returns and Systemic Risk (Full Sample)

Currency	USD		JPY		Neutral	
	EQL	SPD	EQL	SPD	EQL	SPD
Avg. Carry	3.08%	8.08%**	2.13%	5.07%	2.16%**	5.48%***
	(1.33)	(2.31)	(0.48)	(0.95)	(1.98)	(3.30)
α	0.46%	4.14%	-3.13%	-1.40%	0.99%	3.70%***
	(0.27)	(1.61)	(-1.00)	(-0.37)	(1.12)	(2.79)
β	-0.048***	-0.073***	-0.097***	-0.119***	-0.022***	-0.033***
	(-9.06)	(-9.97)	(-10.73)	(-10.82)	(-7.43)	(-8.48)

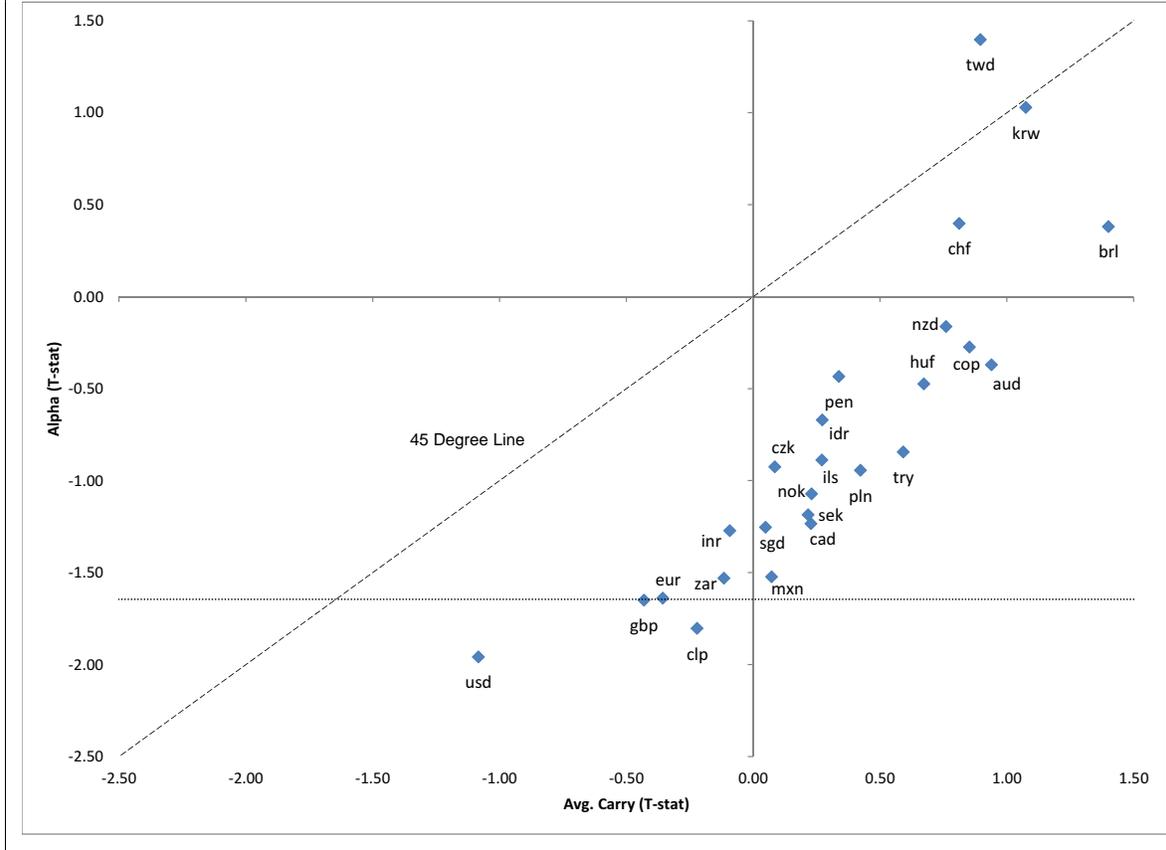
Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

Figure 6: Alpha vs. Avg. Carry (Full Sample)

Panel A: USD



Panel B: JPY



Turning next to the options hedged carry, regression results are shown in Table 5.²⁰ As before, the hedged carry is significant and remains so after correcting for its systemic exposure. This happens even though the unhedged carry is not broadly significant. On average, the T-statistics for the hedged carry are nearly double those for the unhedged carry. Currency level regressions confirm these results where all but 6 currencies' T-statistics were higher for their hedged strategy, as illustrated in Figure 7.

²⁰We drop NZD from the hedged carry sample, as its data series was not available after February 2012. Including NZD and data only through February 2012 has little impact on our results.

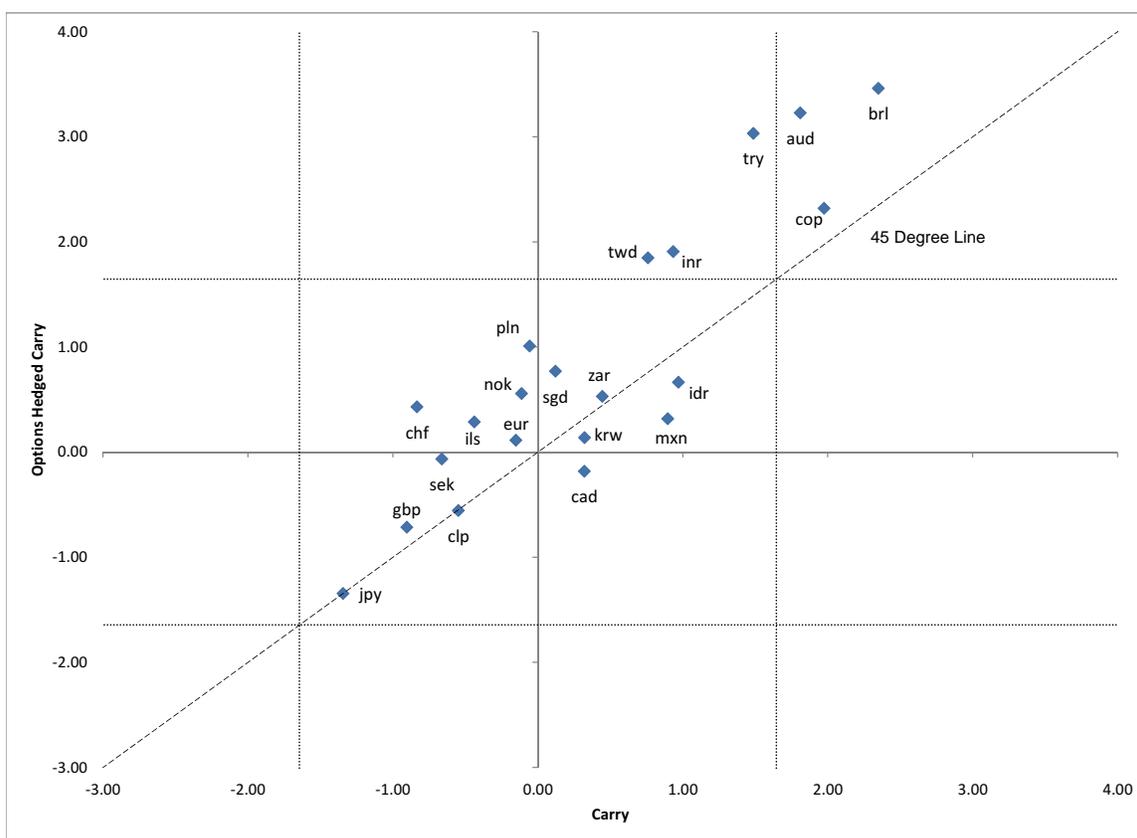
Table 5: Options Hedged Carry Trade Returns (Full Sample)

Currency	USD		Neutral	
Weights	EQL	SPD	EQL	SPD
Avg. Carry	1.94%	7.14%**	1.19%	4.32%**
	(0.87)	(2.07)	(0.96)	(2.43)
Hedged Carry	2.10%**	5.14%***	1.68%**	3.46%***
	(2.23)	(3.38)	(2.08)	(3.40)
α	1.89%*	4.43%***	1.95%**	3.52%***
	(1.94)	(2.93)	(2.16)	(3.13)
β	-0.004	-0.013*	0.005	0.001
	(-0.94)	(-1.88)	(-1.03)	(-0.18)

Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

Figure 7: Hedged vs. Unhedged Carry (Full Sample)

Panel A: Returns (T-stats)



Panel B: Alphas (T-stats)

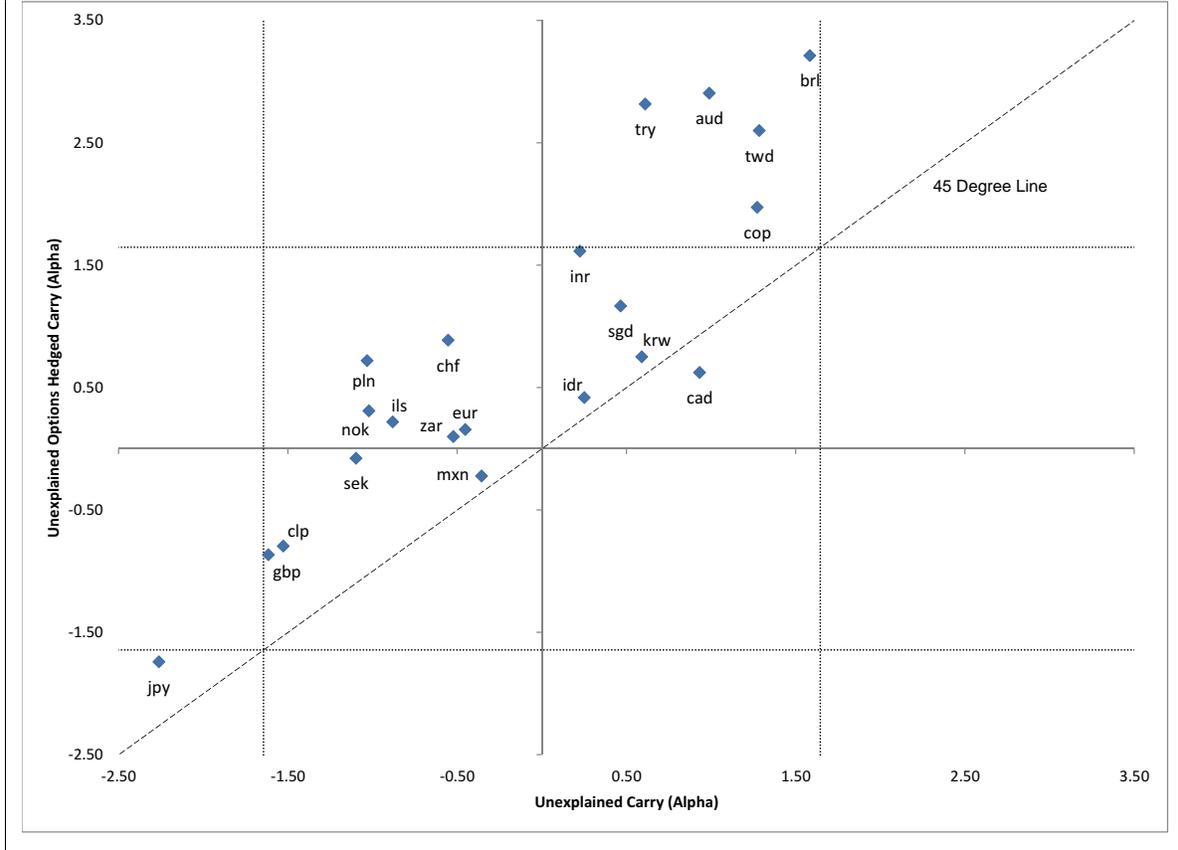
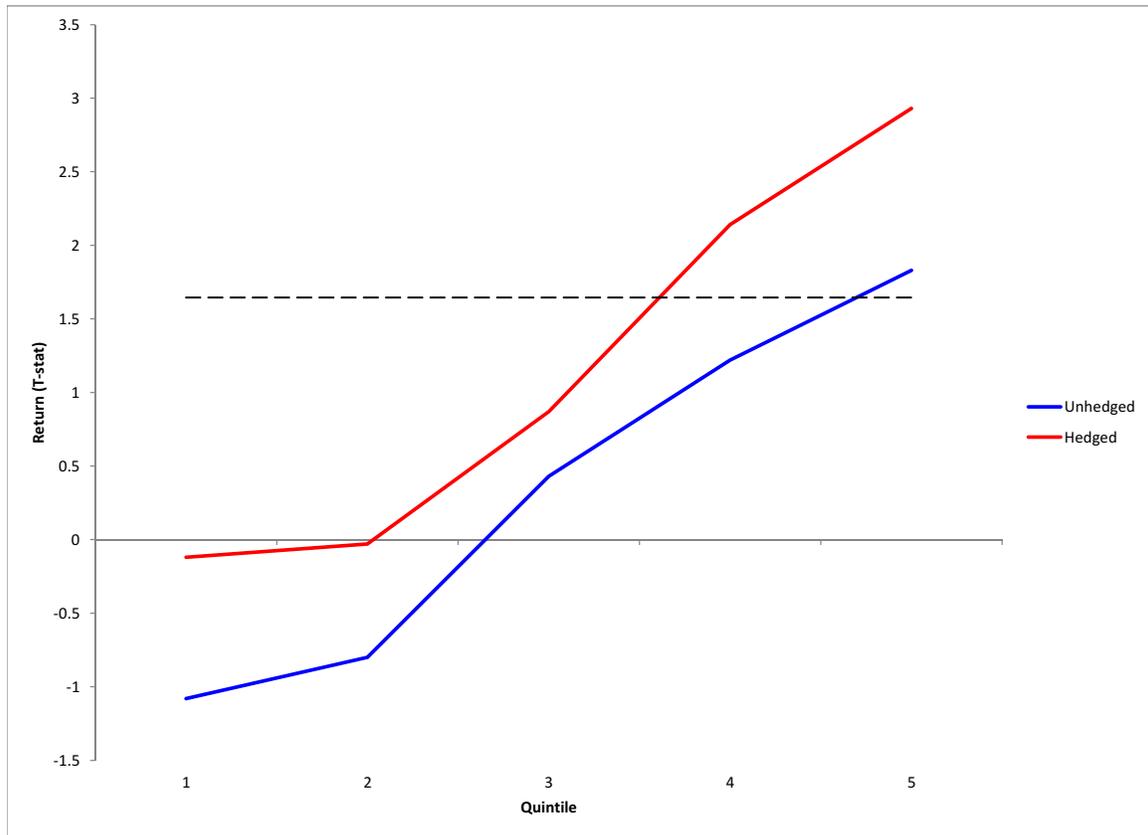


Figure 8 plots the returns and α 's separately for each interest rate quintile, illuminating a somewhat different pattern from pre-crisis. Whereas in pre-crisis the gains from hedging were only present for lower interest rate quintiles, in the full sample they exist for all interest rate quintiles. Panel B reveals a similar pattern in the α -space. While three quintiles of the unhedged carry actually underperformed their exposure, all quintiles of the hedged carry outperformed. We also plot the cumulative returns to the two strategies in Figure 9. Interestingly, while the gains from hedging came largely from increased expected returns for lower quintiles, the gains for high quintiles were entirely from reduced volatility.

Figure 8: Carry Returns by Quintile (Full Sample)

Panel A: Returns



Panel B: Alphas

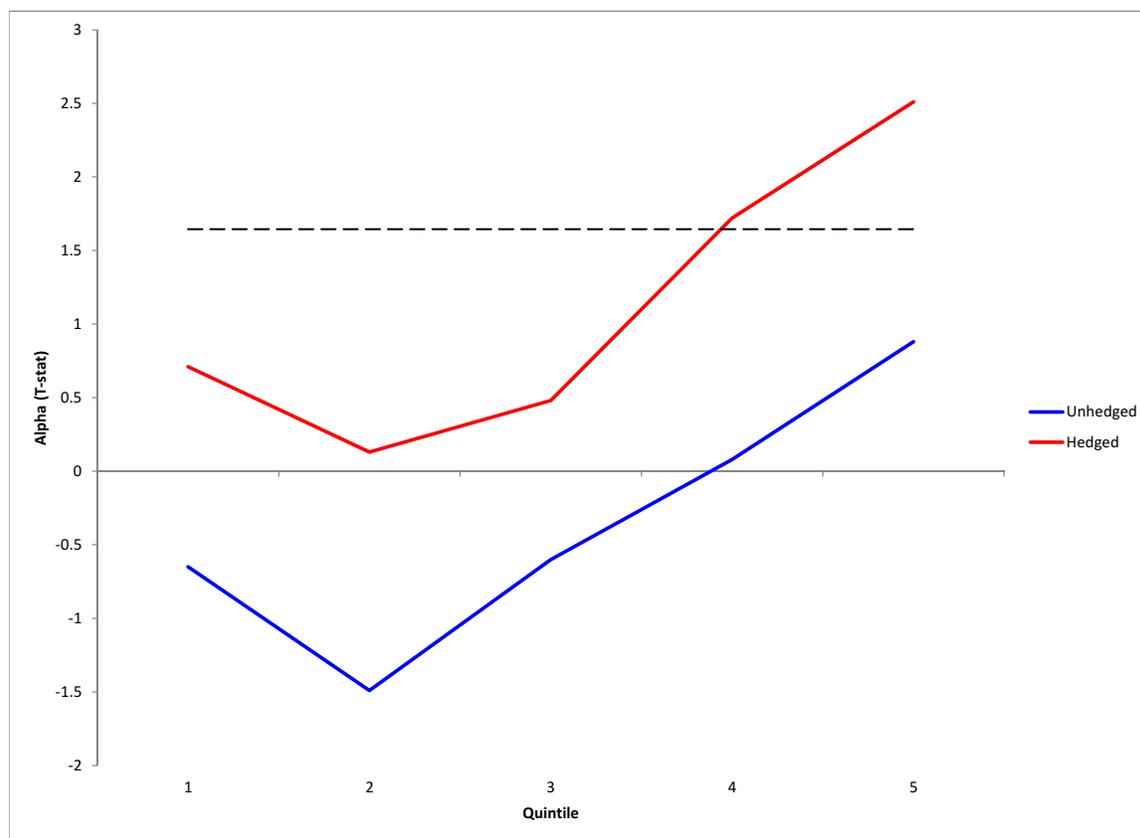
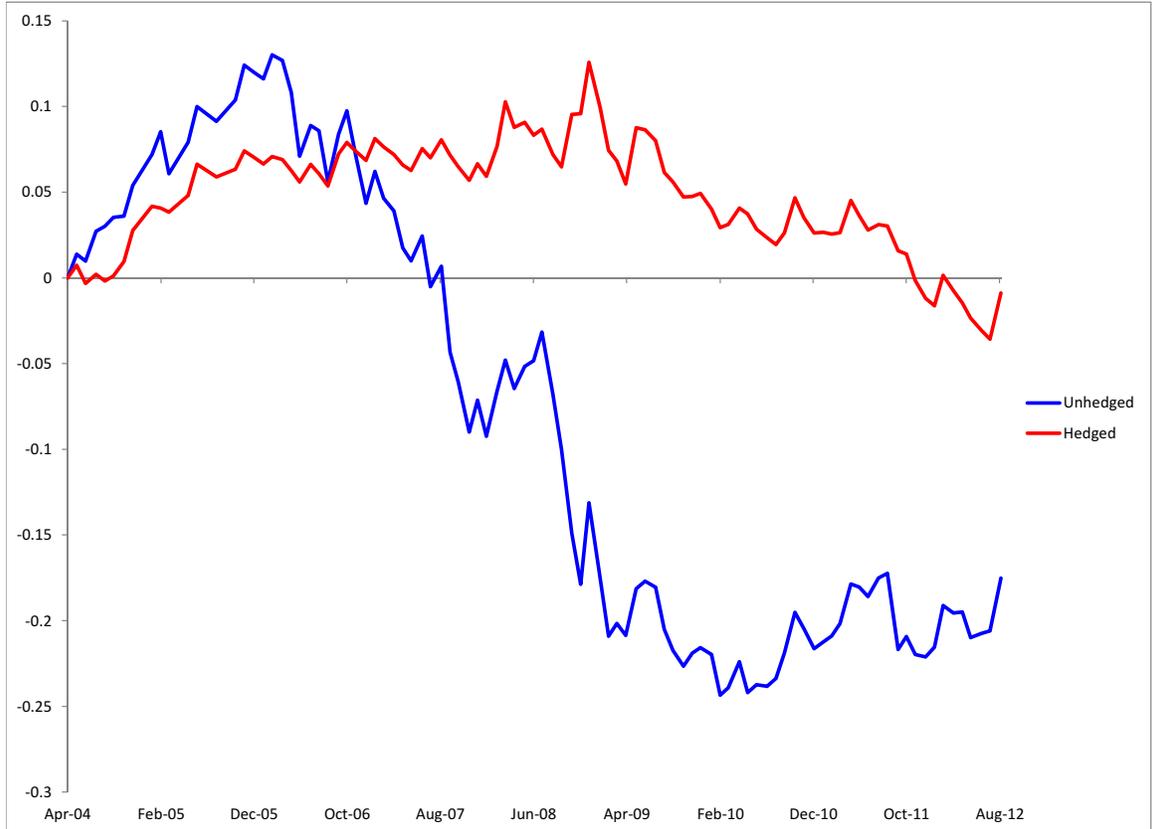


Figure 9: Cumulative Returns

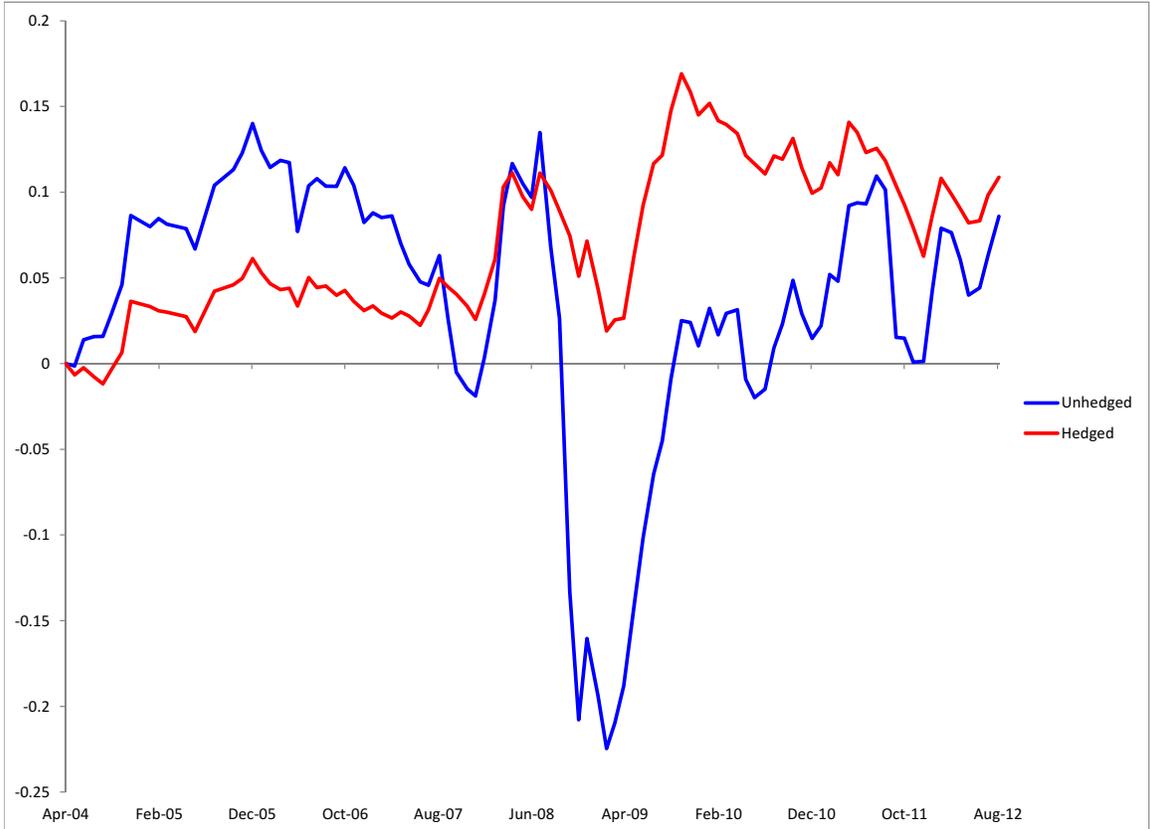
Quintile 1



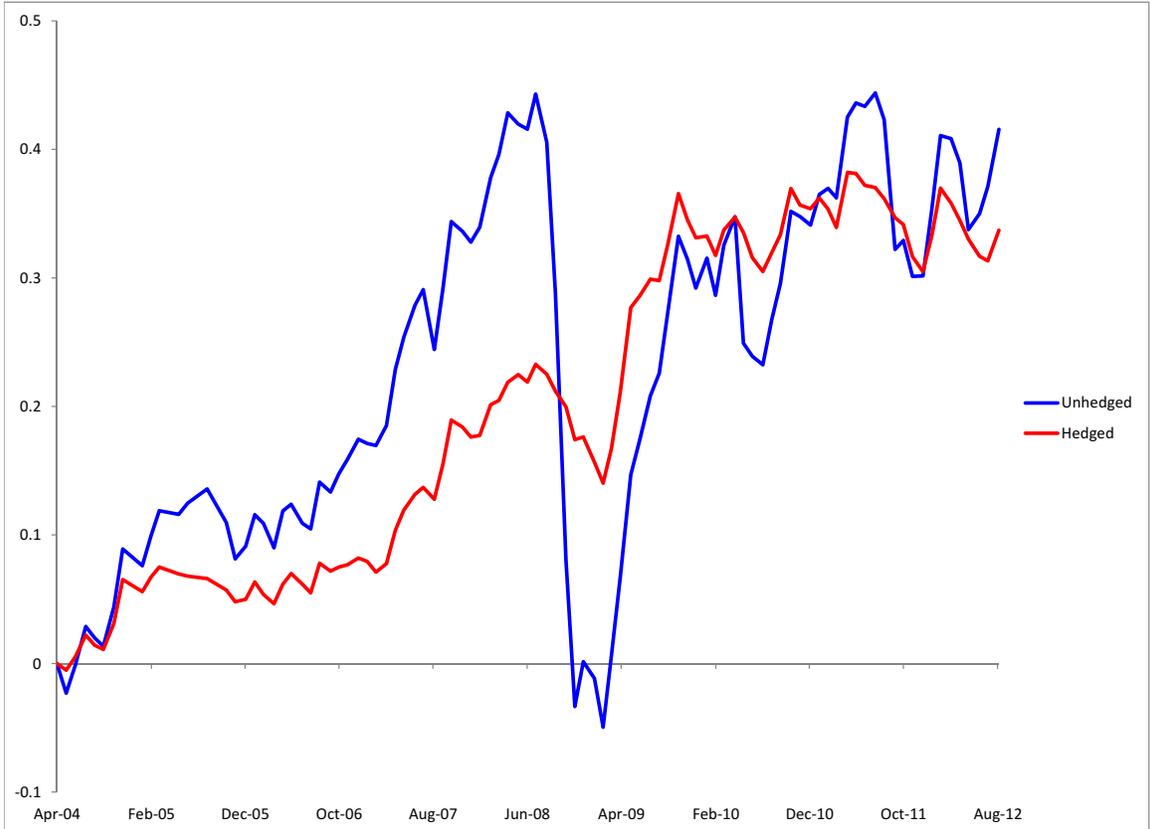
Quintile 2



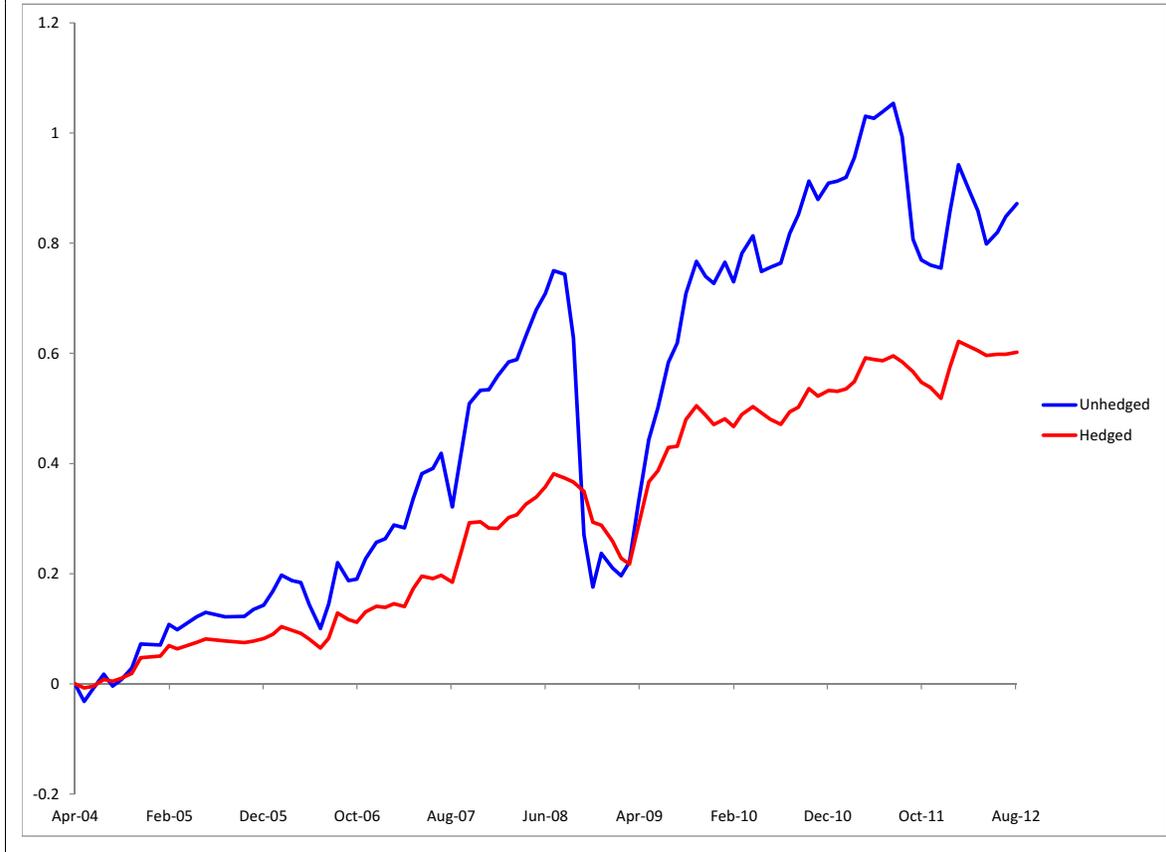
Quintile 3



Quintile 4



Quintile 5



Finally, we implement our direct test of whether the appropriately constructed bundle of FX options provide a cheap form of systemic insurance. The results presented in Table 6 confirm our findings from the pre-crisis.

Table 6: FX Options Returns and Systemic Risk (Full Sample)

Currency	USD		Neutral	
Weights	EQL	SPD	EQL	SPD
α	307.19%***	214.72%*	299.19%***	258.34%**
	(2.96)	(1.79)	(3.29)	(2.36)
β	2.71***	3.49***	1.91***	2.20***
	(5.66)	(6.53)	(5.85)	(5.06)

Notes: Standard errors are robust to heteroskedasticity. Avg. carry and α are annualized. T-statistics are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

7 Final Remarks

To summarize, we find that after appropriately hedging the carry trade with VIX roll-downs in both samples that exclude and include crises, there is no evidence that its return is particularly large, and in the latter, there is evidence that it may actually be too small. This contrasts with Burnside, Eichenbaum, Kleshchelski et al (2011) who find that the returns to the carry trade are not a compensation for risk. On the other hand, like those previous authors, we find that when the carry is hedged with FX options, it does indeed produce significantly positive returns for which we have no risk-based explanation. Taken jointly, our two sets of results suggest that portfolios of FX options designed to hedge the carry trade provide a relatively cheap means of hedging systemic risk, and our tests confirm this hypothesis.

Put differently, the new puzzle is that there is a premium to selling systemic risk insurance, hedged with currency options. The conventional carry trade is a form of selling systemic insurance, which when hedged with FX options generates an excess return. But the source of the excess return is in the low cost of the hedge, not the high return of the

carry itself.

References

- [1] BACKUS, D., FORESI, S., AND TELMER, C. Affine term structure models and the forward premium anomaly. *The Journal of Finance* 56, 1 (2001), 279–304.
- [2] BHANSALI, V. Volatility and the carry trade. *The Journal of Fixed Income* 17, 3 (2007), 72–84.
- [3] BRUNNERMEIER, M., NAGEL, S., AND PEDERSEN, L. Carry trades and currency crashes. In *NBER Macroeconomics Annual 2008*, vol. 23. University of Chicago Press, Chicago, IL, 2009, pp. 313–347.
- [4] BURNSIDE, C. The forward premium is still a puzzle. Working Paper 13129, National Bureau of Economic Research, 2007.
- [5] BURNSIDE, C. Carry trades and risk. Working Paper 17278, National Bureau of Economic Research, 2011.
- [6] BURNSIDE, C., EICHENBAUM, M., KLESHCHELSKI, I., AND REBELO, S. Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24, 3 (2011), 853–891.
- [7] BURNSIDE, C., EICHENBAUM, M., AND REBELO, S. Carry trade and momentum in currency markets. *Annual Review of Financial Economics* 3, 1 (2011), 511–535.
- [8] CLARIDA, R., DAVIS, J., AND PEDERSEN, N. Currency carry trade regimes: Beyond the fama regression. *Journal of International Money and Finance* 28, 8 (2009), 1375–1389.
- [9] COCHRANE, J. *Asset Pricing*. Princeton University Press, Princeton, NJ, 2001.
- [10] ENGEL, C. The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of Empirical Finance* 3, 2 (1996), 123–192.
- [11] FAMA, E. Forward and spot exchange rates. *Journal of Monetary Economics* 14, 3 (1984), 319–338.
- [12] FAMA, E., AND FRENCH, K. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 1 (1993), 3–56.
- [13] FARHI, E., FRAIBERGER, S., GABAIX, X., RANCIERE, R., AND VERDELHAN, A. Crash risk in currency markets. Working Paper 15062, National Bureau of Economic Research, 2009.

- [14] GOURIO, F., SIEMER, M., AND VERDELHAN, A. International risk cycles. Working Paper 17277, National Bureau of Economic Research, 2011.
- [15] JENSEN, M., SCHOLES, M., AND BLACK, F. The capital asset pricing model: Some empirical tests. In *Studies in the Theory of Capital Markets*. Praeger Publishers Inc., New York, NY, 1972, p. 79–121.
- [16] JUREK, J. Crash-neutral currency carry trades. Working paper, Princeton University, 2009.
- [17] LUSTIG, H., ROUSSANOV, N., AND VERDELHAN, A. Countercyclical currency risk premia. Working Paper 16427, National Bureau of Economic Research, 2010.
- [18] LUSTIG, H., ROUSSANOV, N., AND VERDELHAN, A. Common risk factors in currency markets. *Review of Financial Studies* 24, 11 (2011), 3731–3777.
- [19] LUSTIG, H., AND VERDELHAN, A. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97, 1 (2007), 89–117.
- [20] LUSTIG, H., AND VERDELHAN, A. The cross-section of foreign currency risk premia and consumption growth risk: A reply. Working Paper 13812, National Bureau of Economic Research, 2008.
- [21] MENKHOFF, L., SARNO, L., SCHMELING, M., AND SCHRIMPF, A. Carry trades and global foreign exchange volatility. CEPR Discussion Paper 8291, Center for Economic and Policy Research, 2011.
- [22] SHANKEN, J. On the estimation of beta-pricing models. *Review of Financial Studies* 5, 1 (1992), 1–55.
- [23] WHITE, H. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48, 4 (1980), 817.

A Appendix

A.1 Carry Trade and Fama-French Risk Factors: Pre-Crisis

In this section, we analyze the exposure of the equally weighted carry to the traditional Fama-French risk factors (MKT, SMB, HML), as studied in Burnside, Eichenbaum, Kleshchelski et al. (2011). In the pre-crisis period, the estimates confirm that the carry trade is not significantly exposed to these traditional factors, and controlling for them leaves a significant excess return. Furthermore, we find that even after controlling for the Fama-French factors, the carry remains significantly exposed to VIX rolldowns, and the resulting α is still insignificant.

Table A1: Fama-French Risk Factors and Carry (USD)

	(1)	(2)	(3)	(4)
MKT	0.075	0.070	-0.013	-0.013
SMB		0.051		0.031
HML		0.000		-0.009
β			-0.031*	-0.030*
α	4.44%**	4.39%**	2.43%	2.48%
	(2.49)	(2.58)	(1.20)	(1.22)
Adj.R2	0.061	0.030	0.130	0.094
N	48	48	48	48

Notes: Standard errors are robust to heteroskedasticity, and α 's are annualized. T-statistics for α 's are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

Table A2: Fama-French Risk Factors and Carry (JPY)

	(1)	(2)	(3)	(4)
MKT	0.324***	0.331***	0.103	0.106
SMB		-0.076		-0.129
HML		0.020		-0.005
β			-0.077*	-0.081*
α	8.81%**	8.81%*	3.77%	3.67%
	(2.06)	(1.99)	(0.73)	(0.67)
Adj.R2	0.190	0.157	0.250	0.224
N	48	48	48	48

Notes: Standard errors are robust to heteroskedasticity, and α 's are annualized. T-statistics for α 's are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.

Table A3: Fama-French Risk Factors and Carry (Neutral)

	(1)	(2)	(3)	(4)
MKT	0.083**	0.082**	-0.039	-0.039
SMB		0.025		-0.004
HML		-0.016		-0.029
β			-0.043***	-0.043***
α	3.32%**	3.36%**	0.53%	0.61%
	(2.40)	(2.43)	(0.32)	(0.37)
Adj.R2	0.135	0.101	0.380	0.356
N	48	48	48	48

Notes: Standard errors are robust to heteroskedasticity, and α 's are annualized. T-statistics for α 's are reported in parentheses. * - $p < 0.10$, ** - $p < 0.05$, *** - $p < 0.01$.