

Crash-neutral Currency Carry Trades

Jakub W. Jurek*

Abstract

Currency carry trades implemented within G10 currencies have historically delivered significant excess returns with annualized Sharpe ratios in excess of one. This paper investigates whether these excess returns reflect compensation for exposure to crash risk by analyzing the time-series dynamics of the moments of the risk-neutral distribution extracted from currency options, and by examining returns to crash-neutral currency carry trades in which exposure to crashes has been hedged by combining positions in currencies and currency options. Risk-neutral and realized skewness are shown to move in opposite directions in response to realized currency returns such that insurance against currency crashes is cheapest precisely when it is needed most. Although excess returns to crash-neutral strategies decline relative to their unhedged counterparts, they remain positive and highly statistically significant. The results indicate that crash risk premia can explain 30-40% of the total excess return to currency carry trades. Rationalizing the entirety of the excess return via a crash risk premium would require implied volatilities of out-of-the-money currency options to be roughly four times greater than those observed in the data.

FIRST DRAFT: OCTOBER 2007

THIS DRAFT: AUGUST 2008

*Jurek: Princeton University, Bendheim Center for Finance, e-mail: jjurek@princeton.edu. I thank John Campbell, Steve Edelstein, Mark Mueller, Erik Stafford, and seminar participants at the Harvard Finance Lunch for providing valuable comments. I am especially grateful to J.P. Morgan for providing the FX option data.

Uncovered interest parity predicts that high interest rate currencies should depreciate relative to low interest rate currencies, such that investors would be indifferent between holding the two. Empirically, not only do high interest rate currencies not depreciate relative to their low interest rate counterparts, they tend to appreciate. This finding, often referred to as the *forward premium anomaly*, is one of the most prominent features of exchange rate data. Using data on foreign exchange options, I show that 30-40% of the excess returns to currency carry trades exploiting this anomaly are attributable to a crash risk premium attaching to currencies with high interest rates. However, unlike in equity option markets, the cost of protection against large adverse currency moves appears to be relatively cheap. In fact, a complete resolution of the forward premium anomaly via a crash risk premium argument requires implied volatilities of deep out-of-the-money currency options that are roughly four times greater than what is actually observed in the data. If true, this would imply a massive mispricing in the foreign exchange option market.

The currency carry trade exploits the forward premium anomaly by borrowing funds in currencies with low interest rates and lending them in currencies with high interest rates. This strategy captures the interest rate differential (carry) between the two currencies, but leaves the investor exposed to fluctuations in the exchange rate between the high interest rate currency and the low interest rate currency. Historically, since high interest rate currencies have tended to appreciate relative to their low interest rate counterparts, the currency exposure has actually turned out to be an additional source of return. As a result, currency carry trades have been characterized by an extremely attractive risk-return tradeoff. For example, a simple strategy which equal-weighted nine individual carry trades implemented in currency pairs involving the U.S. dollar and one of the remaining G10 currencies earned an annualized excess return of 4.78% with an annualized volatility of 5.07% (Sharpe ratio = 0.91) over the period from January 1990 to March 2007.¹ In the latter part of the sample (1999-2007), simple equal- and spread-weighted portfolios of carry trades delivered even higher Sharpe ratios of 1.26 and 1.48, respectively. To illustrate this, Figure 1 plots the total return indices for the carry strategy *vis a vis* the total return indices for the three Fama-French risk factors and the momentum portfolio. As can be readily seen from the plot, the returns to the carry strategy dominate the returns to each of the alternatives when the strategies are scaled to have the same ex post volatility.

The returns to currency carry trades are characterized by two main features. First, the volatility of the strategy is low and stands at roughly one half the volatility of any given X/USD exchange rate, suggesting that exchange rate innovations are largely uncorrelated in the cross-section. Second, although the low volatility allows the total return index to have an extremely smooth upwards progression, the strategy is punctuated with infrequent, but severe, losses. The skewness of the monthly returns to the equal-weighted carry strategy over the period from January 1990 to March 2007 is negative (-0.95) and its magnitude exceeds the skewness of excess returns on the U.S. equity market and the momentum portfolio. The very high realized Sharpe ratio of the carry trade and

¹The set of G10 currencies consists of: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD), Swedish krona (SEK), and the U.S. dollar (USD).

the prominence of negative skewness have prompted arguments that the excess returns earned by currency carry trade strategies represent compensation for exposure to rare, but severe, crashes in currencies with relatively higher interest rates.² The focus on rare events broadly parallels the approach first proposed by Rietz (1988), and recently revisited by Barro (2006), Martin (2008) and Weitzman (2007), in the context of research on the equity premium.³

This paper contributes to the literature by investigating whether crash risk premia can account for the empirically observed violations of uncovered interest parity (UIP). Using data on foreign exchange options, I extract a complete time-series of risk-neutral variance, skewness and kurtosis for each exchange rate, facilitating an explicit comparison of option-implied and realized moments. Since the distributions of currency returns embedded in option prices are forward-looking, they provide a convenient approach for assessing the market's *ex ante* perceptions regarding the likelihood of crashes, as well as, the cost of insuring against such events. Contrary to the predication of the crash risk hypothesis, I find that risk-neutral skewness does not forecast excess currency returns. Nonetheless, the dynamics of risk-neutral skewness exhibit very interesting behavior. In particular, I show that option-implied and realized skewness respond in opposite directions to lagged currency returns, such that the option-implied skewness forecasts future realized skewness negatively. These facts suggests that insurance against crashes is actually cheapest precisely when the risk of crashes is largest.

Finally, I investigate the returns to carry trades in which the risk of currency crashes has been completely eliminated through hedging in the option market. Consistent with the previous findings on the price of crash insurance, I find that crash-neutral currency carry trades continue to deliver positive and statistically significant excess returns, indicating that a crash risk premium alone cannot reconcile the violations of UIP. Simple equal- and spread-weighted portfolios of crash-neutral currency carry trades hedged using out-of-the-money (10δ) options deliver excess returns of 3.18% (t-stat: 3.13) and 5.31% (t-stat: 3.69), respectively. Although mean excess returns decline as the strategies are hedged with options that are progressively closer to at-the-money, they continue to remain statistically significant. When compared with the excess returns to their unhedged counterparts, the mean excess returns to crash-neutral carry trades are statistically significantly lower. The difference in the mean excess returns suggests that crash risk premia account for 30-40% of the excess returns to currency carry trades. In order to drive the excess returns to carry trades down to zero, the implied volatilities of out-of-the-money options hedging against currency crashes would have to have been nearly four times as large as actually observed in the data.

²Currency carry trade returns generally appear to be unrelated to risk factors proposed by traditional asset pricing models (Burnside, et. al (2006)), although Lustig and Verdelhan (2007a, 2007b) dispute these findings and argue in favor of a consumption risk factor. Brunnermeier, Gollier, and Parker (2007) argue that high returns of negatively skewed assets may be a general phenomenon.

³An extensive literature documents the presence of high volatility and crash risk premia in the equity option market; see for example, Coval and Shumway (2001), Pan (2002), Bakshi and Kapadia (2003), Driessen and Maenhout (2006) and references therein.

1 The Currency Carry Trade

The expectations hypothesis of exchange rates, also known as *uncovered interest parity (UIP)*, postulates that investors should be indifferent between holding riskless deposits denominated in various currencies. Equivalently, high interest rate currencies are expected to depreciate relative to low interest rate currencies, such that the currency return exactly offsets the interest rate differential. If this were the case, the forward exchange rate for a currency – the rate at which one can contract to buy/sell a foreign currency at a future date – would provide an unbiased estimate of the future exchange rate. This relationship is strongly violated in the data leading to a *forward premium anomaly*, whereby high interest rate currencies actually tend to appreciate, rather than depreciate, against their low interest rate counterparts.

The currency carry trade is designed to exploit this anomaly and involves borrowing funds in a currency with a low interest rate and lending them in a currency with a high interest rate. At some future date the proceeds from lending in the high-interest rate currency are converted back into the funding currency, and used to cover the low-interest rate loan. The balance of the proceeds constitutes the profit/loss from the carry trade, and can be thought of as a combination of the interest rate differential (carry) and the realized currency return. However, since the carry is known *ex ante*, and is riskless in the absence of counterparty risk, the sole source of risk in the carry trade stems from uncertainty regarding future exchange rates. Principally, the carry trade exposes the arbitrageur to rapid depreciations (crashes) of the currency which he is long *vis a vis* the funding currency. The next section documents the violations of UIP in the data and characterizes the historical returns to the currency carry trade implemented in the set of G10 currencies.

1.1 Uncovered Interest Parity

Suppose we denote the exchange rate – expressed as the price of one unit of foreign currency in terms of domestic currency – by S_t , and the foreign and domestic interest rates for τ -period loans as $r_{f,t}$ and $r_{d,t}$, respectively. Then, the forward rate, $F_{t,\tau}$, which is the time t price for one unit of foreign currency to be delivered τ periods later, is determined through no arbitrage and satisfies:

$$F_{t,\tau} = S_t \cdot \exp((r_{d,t} - r_{f,t}) \cdot \tau) \tag{1}$$

This relationship is known as *covered interest parity* and is essentially never violated in the data (Burnside, et al. (2006)). When the foreign interest rate, $r_{f,t}$, exceeds the domestic interest rate, $r_{d,t}$, S_t is above $F_{t,\tau}$ and the foreign currency is said to trade at a premium to its forward. Conversely, when $r_{d,t}$ exceeds $r_{f,t}$, the foreign currency trades at a discount ($F_{t,\tau} > S_t$). Under UIP, the forward rate provides an unbiased estimate of the future spot exchange rate, a condition which is stated either in levels or logs,

$$F_{t,\tau} = E_t[S_{t+\tau}] \quad \text{or} \quad f_{t,\tau} = E_t[s_{t+\tau}] \tag{2}$$

If true, investors should be indifferent between buying the foreign currency at time t in the forward market and converting their domestic currency to the foreign currency, investing in the foreign riskless bond and re-converting their investment proceeds back to the domestic currency at the future date. Empirical work, starting with Hansen and Hodrick (1980) and Fama (1984), tests UIP by regressing currency returns on the forward premium, defined as the difference between the prevailing forward and spot prices.⁴ When expressed in logs the regression test takes on the following form,

$$s_{t+1} - s_t = a_0 + a_1 \cdot (f_t - s_t) + \varepsilon_{t+1} \quad H_0 : a_0 = 0, a_1 = 1 \quad (3)$$

The null hypothesis under UIP is that the currency return is, in expectation, equal to the forward premium, which is given by the interest rate differential, $(r_{d,t} - r_{f,t}) \cdot \tau$. Although the empirical prediction of this theory only holds in the absence of currency risk premia (or, in the presence of risk-neutral investors), it constitutes a useful benchmark for examining the data.

Table I presents the results of UIP regressions for nine currency pairs, each containing one of the G10 currencies and the U.S. dollar, which is assumed to be the investor's domestic currency. Currency returns are computed for 21-day rolling windows and span the period from January 1990 to March 2007. The right panel presents the results for the January 1999 to March 2007 subperiod, which is the focus of the ensuing sections due to the simultaneous availability of foreign exchange option data. The (log) forward spread is measured using the differential between the U.S. eurocurrency rate and the foreign eurocurrency rate for one-month deposits. For the full subperiod, the null of UIP ($H_0 : a_0 = 0, a_1 = 1$) is rejected at the 5% significance level for four of the nine countries. Of the remaining five, another two have negative slope coefficients, \hat{a}_1 , indicating that high-interest rate currencies tend to *appreciate* relative to their low-interest rate counterparts. For three currencies – the British pound, the Norwegian krone and and Swedish krona – UIP cannot be rejected at conventional significance levels. Consequently, when UIP is evaluated in the context of a panel regression with fixed country effects, the null cannot be rejected, albeit the slope coefficient once again has the wrong sign. UIP fares considerably worse in the second part of the sample (1999-2007). The null hypothesis is rejected at the 5% (10%) significance level in six (all) of the G10 countries, with negative point estimates for the slope coefficients in all countries. Unsurprisingly, the panel regression, resoundingly rejects UIP in this subperiod.

Cross-sectional regressions also consistently indicate that the further a currency's lending rate is above (below) the U.S. lending rate, the greater the anticipated appreciation (depreciation). Unlike in the time series regressions, in which the adjusted R^2 of the forecasting regression is rarely above 1%, the cross-sectional R^2 is an order of magnitude higher. This suggests that although a trade aimed at exploiting violations of UIP in a single currency may experience quite variable performance over time, a portfolio trade exploiting the entire cross section of currency pairs is likely to be quite lucrative. Moreover, a portfolio of carry trades in which the weights of the individual currencies

⁴Froot and Thaler (1990), Lewis (1995) and Engel (1996) survey the vast empirical literature on tests of UIP.

are set in proportion to the interest rate differential is predicted to outperform an equal-weighted strategy.

Although the amount of predictability in the foreign exchange rate return at the one-month horizon is generally small, as evidenced by the low adjusted R^2 values, the regressions indicate that investors can earn excess returns by borrowing funds in the relatively low interest rate currency and lending them in the currency with the relatively high interest rate. This strategy, known as the *currency carry trade*, in reference to the interest rate differential (carry) it earns, has historically been one of the most popular foreign exchange strategies. The next sections describe the carry trade strategy in detail and examine its historical performance to assess the economic significance of deviations from UIP.

1.2 Exploiting deviations from UIP

In the standard carry trade an investor borrows money in a currency with a low interest rate and invests the proceeds in a currency with a high interest rate. We denote the price of the foreign currency in terms of the domestic currency – taken to be the U.S. dollar – by S_t , such that an increase (decrease) in S_t corresponds to an appreciation (depreciation) of the foreign currency relative to the dollar. To illustrate the payoff to this strategy consider an investor aiming to exploit the τ -month interest rate differential between the domestic currency, bearing a low interest rate, and a foreign currency, bearing a high interest rate.

The investor begins by borrowing S_t dollars at a rate of $r_{d,t}$ for a period of τ months in order to finance the purchase of one unit of the foreign currency. After converting the funds to the foreign currency he lends them out for a period of τ months at the then prevailing foreign interest rate $r_{f,t}$. When the domestic currency is the relatively high interest rate currency, the investor employs a symmetric strategy and shorts one unit of the foreign currency. At time $t + \tau$ the *payoff* to the strategy is given by:

$$\widetilde{CT}_{t+1} = \begin{cases} r_{f,t} > r_{d,t} : & \exp(r_{f,t} \cdot \tau) \cdot \tilde{S}_{t+\tau} - \exp(r_{d,t} \cdot \tau) \cdot S_t \\ r_{d,t} > r_{f,t} : & \exp(r_{d,t} \cdot \tau) \cdot S_t - \exp(r_{f,t} \cdot \tau) \cdot \tilde{S}_{t+\tau} \end{cases} \quad (4)$$

If uncovered interest parity (UIP) held, the expected payoff to the carry trade would be zero, since $S_{t+\tau} = F_{t,\tau}$, and the change in the exchange rate would exactly offset the interest differential (carry). However, in the presence of the empirically observed violations of UIP, the expected excess return to the currency carry trade is positive. The corresponding carry trade *return* can be obtained by standardizing the payoff by the funding capital – in this case given by the value of one unit of foreign currency at initiation, S_t – to obtain:

$$\tilde{R}_{t+1}^{CT} = \begin{cases} r_{f,t} > r_{d,t} : & \exp(r_{f,t} \cdot \tau) \cdot \left(\frac{\tilde{S}_{t+\tau}}{S_t}\right) - \exp(r_{d,t} \cdot \tau) \\ r_{d,t} > r_{f,t} : & \exp(r_{d,t} \cdot \tau) - \exp(r_{f,t} \cdot \tau) \cdot \left(\frac{\tilde{S}_{t+\tau}}{S_t}\right) \end{cases} \quad (5)$$

Whenever $r_{d,t} < r_{f,t}$, the carry trader is long exposure to the foreign currency, and loses money if the foreign currency *depreciates* relative to the funding currency by more than the interest rate differential. Symmetrically, when $r_{f,t} < r_{d,t}$, the carry trade investor borrows in the foreign currency and is short exposure, and stands to earn a negative excess return if the foreign currency *appreciates*. Put differently, if exchange rates were defined as the price of the high interest rate currency in terms of the low interest rate currency, carry traders stand to earn a negative return in the event of a rapid depreciation, or crash, of the high interest rate currency.

1.3 Historical performance

The historical performance of the currency carry trade is summarized in Tables II (1990-2007) and III (1999-2007). The strategy is implemented at a monthly frequency, with positions established at the end of month $t - 1$ and held until the end of month t . Within each pair, to determine which currency will be the long (short) leg of the trade I compare the one-month Eurocurrency rates prevailing at the end of month $t - 1$. Finally, since exchange rates are expressed in terms of units of domestic currency (U.S. dollar) per unit of foreign currency, the excess returns computed from (5) are interpretable as U.S. dollar returns.

Over the full sample period the carry trade delivers positive excess returns in all nine currency pairs. For five of the individual currency pairs the excess returns are statistically different from zero at the 5% significance level. The annualized Sharpe ratios vary from 0.16 (CHF/USD pair) to 0.79 (SEK/USD pair). Notably, however, the carry trade returns exhibit pronounced departures from normality, with significant negative skewness and excess kurtosis. The minimum and maximum monthly returns are large in magnitude and represent roughly 3.5 standard deviation moves, when compared with the the historically realized monthly volatility. These features are broadly representative of the the carry trade and parallel the results reported in the existing literature (e.g. Burnside, et al. (2006), Brunnermeier, Nagel and Pedersen (2008)).

Panel B of Tables II and III presents the analogous summary statistics for portfolios of the individual currency carry trades. Portfolios weighting the individual carry trades equally (*EQL*), or by the absolute interest rate spread at initiation (*SPR*), deliver positive and statistically significant returns with annualized Sharpe ratios close to (1990-2007) or above (1999-2007) one. Of the total return, between one-third and one-half comes from the carry – computed as the absolute value of the annualized interest rate differential – with the remainder attributable to the currency return. Interestingly, the annualized volatility of the portfolio strategy is less than the mean annualized volatility of the individual currency pairs suggesting that the currency returns from the individual pairs are generally uncorrelated. However, the lack of correlation does not facilitate diversification of skewness or kurtosis. In fact, the returns to portfolios of carry trades are actually even more skewed and heavy-tailed than the underlying pairs. To summarize the data, Figure 2 plots the returns to the portfolio strategies and the nine underlying currency pairs in mean-standard deviation space for the 1999-2007 period. The portfolios lie very close to the mean-standard deviation frontier, suggesting

that even a simple portfolio construction method delivers considerable benefits due to diversification. Moreover, the Sharpe ratios of the the equal- (1.26) and spread-weighted (1.48) strategies are very close to the Sharpe ratio of the the *ex post* mean-standard deviation frontier (1.54), indicating that these naive portfolio strategies are surprisingly efficient.⁵

Finally, to ensure that these characteristics of the currency carry trade are not related to the net dollar exposure of the portfolio strategies, the bottom panel also presents the results for an equal- and spread-weighted carry trade, in which the weights are constructed such that the exposure to the U.S. dollar nets to zero. These strategies are labeled *EQL-\$N* and *SPR-\$N*, respectively. The *EQL-\$N* strategy is constructed by assigning weights of $\frac{1}{m_s}$ to the m_s carry trades which are short the U.S. dollar and long the foreign currency ($r_{f,t} > r_{d,t}$), and weights of $\frac{1}{m_l}$ to the m_l carry trades which are long the U.S. dollar and short the foreign currency ($r_{d,t} > r_{f,t}$). The *SPR-\$N* strategy is constructed analogously, but with weights that are assigned in proportion to the absolute interest rate differential computed within the two sets. As can be readily seen, these strategies continue to deliver positive and statistically significant returns, as well as, retain the high negative skewness and kurtosis of the other non-dollar-neutral strategies. This suggests that the main characteristics of the carry trade strategy are not attributable to net dollar exposure.⁶ Notably, the return volatilities of the dollar-neutral strategies tend to be somewhat higher than those of the standard equal- and spread-weighted portfolio strategies. This feature may be due to excess co-movement of the relatively high- (low-) interest rate currencies, which effectively limits the amount of attainable diversification.

Aside from its high historical excess returns, the defining characteristic of the currency carry trade is the high negative skewness of the realized returns. Brunnermeier, Nagel and Pedersen (2008), for example, argue that realized skewness is related to rapid unwinds of carry trade positions, precipitated by shocks to funding liquidity. They show that funding liquidity measures predict exchange rate movements, and that controlling for the supply of arbitrage capital helps explain violations of UIP. In a related paper, Plantin and Shin (2008) provide a game-theoretic motivation of how strategic complementarities, which lead to crowding in carry trades, can endogenously generate currency crashes. In their model, the simultaneous entry of currency speculators causes currencies with relatively high (low) interest rates to appreciate (depreciate). Consequently, the greater the mass of speculators entering the carry trade, the more likely it is to deliver a positive excess return, and the greater the potential for a future unwind, or crash. This prediction of their model indeed finds empirical support, in that the change in monthly realized skewness (computed from daily carry trade returns) is negatively related to past carry trade returns.⁷ However, since the traders are

⁵DeMiguel, Garlappi and Uppal (2007) find that equal-weighted strategies involving portfolios of equities perform favorably relative to portfolios prescribed by more sophisticated portfolio construction models, due to errors in estimating the parameters required by those models (e.g. means and variances).

⁶Carry trades constructed for investors whose domestic currency is one of the other nine G10 currencies have features that are qualitatively comparable to those presented in Tables II and III. Results are available from the author by request.

⁷In unreported results, a panel regression of the monthly change in realized skewness on the lagged carry trade return produces a coefficient of -7.15 (t-stat: -7.07) and an adjusted R^2 of 4.89%.

risk-neutral in Plantin and Shin’s (2008) model, there are no excess returns (risk premia) to holding high interest rate currencies. Long periods of gradual appreciation of high interest rate currencies are followed by rapid crashes, such that traders – in equilibrium – are indifferent between currencies with high and low interest rates. Consequently, while their model helps explain the observed negative skewness, it is silent on the source of the observed violations of uncovered interest parity.

2 Assessing and Hedging Exposure to Currency Crashes

The persistent profitability of currency carry strategies has led to a search for underlying risk factors responsible for the high excess returns garnered by currencies with high interest rates (Farhi and Gabaix (2008), Lustig, Roussanov and Verdelhan (2008)). Importantly, risk factors postulated by traditional asset pricing models (CAPM, C-CAPM, Fama-French, etc.) appear to be entirely unrelated to the returns on carry trades (Burnside, et. al (2006)).⁸ Consequently, although the theoretical features of the currency risk premium necessary to match the empirical data on the violation of UIP have been known since Fama (1985), the asset pricing literature has struggled with providing a plausible rational model of this anomaly.⁹

Intuitively, the challenge to rational models arises from the extremely high historical Sharpe ratio of the carry trade, which suggests that the price of risk on the “hidden” underlying risk factor is nearly twice that of the equity market. This has led rational theories to consider the importance of rare, but extreme, crashes. For example, Farhi and Gabaix (2008) present a theoretical model in which the forward premium anomaly is generated by a time-varying exposure to disaster risk. This approach mirrors that taken in the equity premium literature (Rietz (1988), Barro (2006), Weitzman (2007), Martin (2008)) and the literature seeking to reconcile the prices of deep out-of-the-money puts with empirical return distributions (Pan (2002)). Indeed, Coval and Shumway (2001), Bakshi and Kapadia (2003), and Driessen and Maenhout (2006) report high Sharpe ratios for various delta-neutral option strategies, which can be interpreted as being consistent with large volatility and crash risk premia. Consequently, exposure to a highly priced crash risk factor attaching to currencies with relatively high interest rates provides a plausible mechanism for explaining the observed violations of uncovered interest parity.

In order to examine the crash risk hypothesis, I turn to data on foreign exchange (FX) options and begin by introducing methodologies used to assess and hedge exposure to currency crashes. The price of options protecting against the risk of rapid devaluations provides valuable information regarding the probability of currency crashes, as well as, the risk premia demanded by investors for being

⁸Lustig and Verdelhan (2007a, 2007b), dispute these claims and argue that returns to carry strategies represent compensation for exposure to real U.S. consumption risk.

⁹Backus, Telmer and Foresi (2001) show that in order to account for the anomaly in an affine model, one has to either allow for state variables to have asymmetric effects on state prices in different currencies or abandon the requirement that interest rates be strictly positive. Verdelhan (2007) argues that the forward premium anomaly is consistent with time-varying, countercyclical risk premia generated by country-specific habit processes. The model, however, requires high effective risk aversion levels, as well as, significant restrictions on risk sharing.

exposed to those risks. To examine the market's *ex ante* perceptions of crash risk, I first use the option data to extract the moments of the risk-neutral distribution. A comparison of the dynamics of the risk-neutral moments – in particular risk-neutral skewness – *vis a vis* realized skewness, provides surprising insights regarding the market's perceptions of tail risk. Finally, I introduce the construction of crash-neutral currency carry trades, in which exposure to rapid depreciations (appreciations) of the high (low) interest rate currencies has been hedged in the option market. To assess whether violations of UIP are attributable to crash risk premia, the empirical sections examine the excess returns to the crash-neutral strategies and compare them to the returns obtained from the unhedged carry strategy.

2.1 The market's perception of crash risk

Breeden and Litzenberger (1978) were the first to show that an asset's entire risk-neutral distribution (i.e. state price density) can be recovered from the prices of a complete set of options on that asset. Following the logic of state-contingent pricing (Arrow (1964), Debreu (1959)), the risk-neutral distribution, $q(S)$, enables one to value arbitrary state contingent payoffs, $H(S)$, via the following pricing equation:

$$p_t = \exp(-r_{d,t} \cdot \tau) \cdot \int_0^\infty H(S_{t+\tau}) \cdot q(S_{t+\tau}) dS_{t+\tau} \quad (6)$$

In particular, if we denote the continuously compounded return by, $R_{t,\tau} \equiv \ln S_{t+\tau} - \ln S_t$, the values of the (non-central) τ -period moments under the risk-neutral measure can be simply computed by setting $H(S_{t+\tau}) = (R_{t,\tau}(S_{t+\tau}))^n$ and removing the discounting. The corresponding discounted values can be interpreted as the prices of contracts paying the realized (non-central) moments of the distribution. After a few simple transformations these values can then be converted to the prices of contracts paying the realized central moments (variance, skewness, etc.). Consequently, this approach to deriving risk-neutral moments, developed in Bakshi, Kapadia, and Madan (2003), can be thought of as an extension of the early results of Britten-Jones and Neuberger (2000) and Carr and Madan (2001) on the pricing of variance swaps.

Rather than begin by extracting the entire risk-neutral distribution, $q(S)$, Bakshi and Madan (2000) show that any payoff function with bounded expectation can be spanned by a continuum of out-of-the-money call and put payoffs. This implies that the price, p_t , of an asset paying, $H(S_{t+\tau})$, can be conveniently obtained by valuing the relevant replicating portfolio of options. Specifically, if the payoff function is twice-differentiable, the asset's price can be obtained from:

$$p_t = \exp(-r_{d,t} \cdot \tau) \cdot (H(\bar{S}) - \bar{S}) + H_S(\bar{S}) \cdot S_t + \int_{\bar{S}}^\infty H_{SS}(K) \cdot C_t(K, \tau) dK + \int_0^{\bar{S}} H_{SS}(K) \cdot P_t(K, \tau) \quad (7)$$

where $H_S(\cdot)$ and $H_{SS}(\cdot)$, denote the first and second derivatives of the state-contingent payoff, and \bar{S} is some future value of the underlying, typically taken to be the forward price. Intuitively, this expression states that the payoff $H(S)$ can be synthesized by buying $(H(\bar{S}) - \bar{S})$ units of a riskless bond, $H_S(\bar{S})$ units of the underlying security and a linear combination of puts and calls with positions given by $H_{SS}(K)$.

Since the state-contingent payoffs, $H(S) = (R_{t,\tau}(S_{t+\tau}))^n$, satisfy the above technical conditions, we can obtain the prices of the *non-central* moment swaps, by substituting the relevant derivatives into (7). The expressions for the discounted values of the first three non-central moments, denoted by $V_t(\tau)$, $W_t(\tau)$ and $X_t(\tau)$, are provided in the Appendix. To fix intuition regarding the associated option portfolios, Figure 3 plots the option positions, $H_{SS}(\cdot)$, for each of the moment contracts as a function of the option strike value. For example, replication of the second moment requires establishing long positions in options at all moneyness levels, with sizes that scale inversely with the square of the option strike value. As a result, the price of the second moment contract is strictly positive, and is particularly high whenever the prices of deep out-of-the-money puts are high. The third non-central moment is replicated by a combination of negative positions in out-of-the-money puts and positive positions in out-of-the-money calls. Consequently, whenever the underlying distribution is negatively (positively) skewed and the prices of puts are greater than (less than) the prices of calls, the price of the replicating portfolio will be negative (positive). Finally, replication of the contract paying the realized fourth moment, once again entails strictly positive positions in options of all moneyness values, but now more heavily weighed in the tails. Once the prices of the three non-central moment contracts have been computed, the risk-neutral variance, skewness, and kurtosis can be obtained from:

$$\text{VAR}_t^Q(\tau) = \exp(r_{d,t} \cdot \tau) \cdot V_t(\tau) - \mu_t(\tau)^2 \quad (8)$$

$$\text{SKEW}_t^Q(\tau) = \frac{\exp(r_{d,t} \cdot \tau) \cdot (W_t(\tau) - 3 \cdot \mu_t(\tau) \cdot V_t(\tau)) + 2 \cdot \mu_t(\tau)^3}{(\exp(r_{d,t} \cdot \tau) \cdot V_t(\tau) - \mu_t(\tau)^2)^{\frac{3}{2}}} \quad (9)$$

$$\text{KURT}_t^Q(\tau) = \frac{\exp(r_{d,t} \cdot \tau) \cdot (X_t(\tau) - 4 \cdot \mu_t(\tau) \cdot W_t(\tau) + 6 \cdot \mu_t(\tau)^2 \cdot V_t(\tau)) - 3 \cdot \mu_t(\tau)^4}{(\exp(r_{d,t} \cdot \tau) \cdot V_t(\tau) - \mu_t(\tau)^2)^2} \quad (10)$$

where:

$$\mu_t(\tau) = -\exp(r_{d,t} \cdot \tau) \cdot \left(\frac{V_t(\tau)}{2} + \frac{W_t(\tau)}{6} + \frac{X_t(\tau)}{24} \right) \quad (11)$$

The Q superscripts are used to denote the fact that the moments are computed under the risk-neutral measure, in contrast to realized moments computed from historical data, which will be denoted with P superscripts.

When applied to data on foreign exchange options, these expressions allow me to dynamically extract the risk-neutral moments of the option-implied currency return distribution, resulting in a daily time series of option-implied variance, skewness and kurtosis observations. Importantly,

the time variation in option-implied skewness provides a direct way for assessing the market’s time-varying perceptions of crash risk, and the cost of insuring against extreme currency moves. In Section 4, I characterize the behavior of risk-neutral skewness in relation to the actual realized skewness, and relate both to measures of the attractiveness of currency carry trades (e.g. interest rate differentials), as well as, recent currency moves. I find that the realized and option-implied skewness measures exhibit dramatically different behavior. While future realized skewness is *negatively* related to past currency appreciations, suggesting that appreciated currencies are more likely to crash, the risk-neutral skewness is *positively* related to past currency appreciations, suggesting that the market perceives appreciated currencies as less likely candidates for a crash. As a result, insurance against crash risk becomes least expensive, precisely when it is needed most.

2.2 Crash-neutral currency carry trades

In order to provide a returns-based measure of the crash risk premium, I also construct carry trades in which the spot currency positions of the standard currency carry trade are combined with a position in a foreign exchange option. The option position is chosen such that the risk of extreme negative outcomes stemming from a depreciation (appreciation) of the high (low) interest rate currency is entirely eliminated. More precisely, whenever the foreign short-term instrument is the long leg of the trade, an investor seeking to limit exposure to sudden depreciations purchases a put option on the foreign currency. Conversely, if the carry trade is funded in the foreign currency, an investor seeking to limit downside exposure purchases a call option, limiting the risk from a sudden appreciation. I refer to these downside-protected trades as *crash-neutral carry trades*. They are constructed to have two features: (1) conditional on the option protection expiring in-the-money all currency risk exposure is eliminated; and, (2) the currency exposure of the crash-neutral portfolio matches that of the standard carry trade (i.e. the delta exposure of the option is hedged at initiation). Intuitively, the first condition ensures that exposure to crash risk is entirely eliminated, while the second, ensures that the returns from the crash-neutral carry strategies are directly comparable with those from the standard strategy presented in Section 1. To test whether the excess returns to the currency carry trade can be attributed to a crash risk premium, I compare the standard carry trade to three variants of the crash-neutral carry trade, differing in the amount of downside protection offered by the option overlay.¹⁰ The returns to portfolio strategies combining crash-neutral trades complement the analysis of risk-neutral moments, and indicate that 30-40% of the excess returns to standard carry trades may indeed be interpreted as compensation for exposure to currency crashes. Before turning to the data, however, I provide a detailed description of how the crash-neutral carry trades are constructed.

¹⁰Burnside, et al. (2008) construct similar currency strategies, however, their panel contains a smaller cross-section of countries and only examines carry trades hedged using *at-the-money* options. Since the focus is on eliminating exposure to extreme moves, deep out-of-the-money options are more relevant, as they provide the most direct measure of the cost of insuring against crashes. Moreover, since their strategies do not hedge the delta exposure of the option overlay the returns cannot be directly compared with returns from the unhedged carry trade.

2.2.1 Portfolio construction

First, consider the situation when the foreign interest rate, $r_{f,t}$, exceeds the domestic interest rate, $r_{d,t}$. In order to take advantage of the deviation from UIP, the trader would like to establish a long position in the foreign currency. In the standard carry trade, this long position exposes the carry trader to losses in the event of a sudden depreciation of the foreign currency. To protect against these losses the carry trader can purchase FX puts with a strike price of K_p at a cost of $P_t(K_p, \tau)$ dollars per put. If the carry trader purchases q_p puts, he must also purchase an additional $-q \cdot \delta_p$ units of the foreign currency, to hedge the negative delta of the put options. Consequently, if the trader started by buying one unit of the foreign currency – as in (5) – he must now buy an additional $q_p \cdot \delta_p$ units of the foreign currency. To fund this position he must borrow an additional $q_p \cdot \delta_p \cdot S_t$ in his domestic currency. Finally, we assume the purchase price of the puts is covered by borrowing additional funds in the domestic currency. At time $t + 1$ the payoff to this portfolio is given by:

$$\begin{aligned} \widetilde{CT}_{t+1}^{CN}(r_{f,t} > r_{d,t}) &= \exp(r_{f,t} \cdot \tau) \cdot (1 - q_p \cdot \delta_p) \cdot \tilde{S}_{t+1} + q_p \cdot \max(K_p - \tilde{S}_{t+1}, 0) - \\ &\quad - \exp(r_{d,t} \cdot \tau) \cdot ((1 - q_p \cdot \delta_p) \cdot S_t + q_p \cdot P_t(K_p, \tau)) \end{aligned} \quad (12)$$

In order to eliminate all currency exposure below the strike price of the option, K_p , the quantity of puts must satisfy,

$$q_p = \exp(r_{f,t} \cdot \tau) \cdot (1 - q_p \cdot \delta_p) \quad \rightarrow \quad q_p = \frac{\exp(r_{f,t} \cdot \tau)}{1 + \exp(r_{f,t} \cdot \tau) \cdot \delta_p} \quad (13)$$

With the above quantity restriction, the payoff equation can be re-expressed as:

$$\widetilde{CT}_{t+1}^{CN}(r_{f,t} > r_{d,t}) = q_p \cdot \max(K_p, \tilde{S}_{t+1}) - \exp(r_{d,t} \cdot \tau) \cdot ((1 - q_p \cdot \delta_p) \cdot S_t + q_p \cdot P_t(K_p, \tau)) \quad (14)$$

This expression makes transparent that the payoff to the strategy is bounded from below, and that for terminal realizations of the exchange rate that are above K_p , the strategy payoff response is steeper than in the standard carry trade. In the standard carry trade the sensitivity to changes in \tilde{S}_{t+1} is equal to $\exp(r_{f,t} \cdot \tau)$, whereas in the crash-neutral strategy it is given by q_p which is strictly greater than $\exp(r_{f,t} \cdot \tau)$ since $\delta_p < 0$. The payoff to the crash-neutral carry trade is illustrated *vis a vis* the payoff to the standard carry trade in the left panel of Figure 4. The crash-neutral carry trade is assumed to include a put option struck 2.5% out-of-the-money relative to the prevailing forward rate ($\frac{K_p}{F_t} = 0.9750$).

By simultaneously decreasing exposure to depreciations (crashes) of the high interest rate currency and increasing exposure to its appreciations, the crash-neutral strategy is able to maintain the same unconditional *ex ante* exposure as the standard carry trade, as desired. Moreover, as the put is struck progressively further out-of-the-money and offers less protection, the delta of the put converges to zero, causing the upside exposure of the crash-neutral trade, q_p , to converge to that of the standard carry trade. In this sense, the crash-neutral strategy nests the payoff to the standard carry strategy.

Now consider the situation when the domestic interest rate, $r_{d,t}$, exceeds the foreign short rate, $r_{f,t}$. In order to take advantage of the UIP violations, the investor borrows \$1 in foreign currency and invests the proceeds in a combination of the short-term domestic bonds and FX calls. Because the position involves borrowing at the foreign short rate it is effectively short the foreign currency, and the call options limit its downside exposure. Suppose the investor were to buy, q_c calls with a strike price of K_c , at a price of $C_t(K_c, \tau)$ dollars per call. In order to eliminate the additional exposure stemming from the long position in the calls ($\delta_c > 0$), the investors shorts $q \cdot \delta_c$ units of the foreign currency in addition to the baseline short position of one unit (as in the standard currency carry trade). Finally, the investor funds the purchase of the call options by borrowing the funds at the domestic interest rate, $r_{d,t}$.¹¹ At time $t + 1$ the payoff to the portfolio is:

$$\begin{aligned} \widetilde{CT}_{t+1}^{CN}(r_{d,t} > r_{f,t}) &= \exp(r_{d,t} \cdot \tau) \cdot ((1 + q_c \cdot \delta_c) \cdot S_t - q_c \cdot C_t(K_c, \tau)) + q_c \cdot \max(\tilde{S}_{t+1} - K_c, 0) - \\ &\quad - \exp(r_{f,t} \cdot \tau) \cdot (1 + q_c \cdot \delta_c) \cdot \tilde{S}_{t+1} \end{aligned} \quad (15)$$

Once again, the requirement that all currency exposure be eliminated above K_c implies that the call quantity satisfy:

$$q_c = \exp(r_{f,t} \cdot \tau) \cdot (1 + q_c \cdot \delta_c) \quad \rightarrow \quad q_c = \frac{\exp(r_{f,t} \cdot \tau)}{1 - \exp(r_{f,t} \cdot \tau) \cdot \delta_c} \quad (16)$$

which allows us to simplify the payoff to:

$$\widetilde{CT}_{t+1}^{CN}(r_{d,t} > r_{f,t}) = \exp(r_{d,t} \cdot \tau) \cdot ((1 + q_c \cdot \delta_c) \cdot S_t - q_c \cdot C_t(K_c, \tau)) - q_c \cdot \min(K_c, \tilde{S}_{t+1}) \quad (17)$$

The payoff makes clear that the portfolio is protected against appreciations of the low interest foreign currency beyond the strike price of the option K_c . By contrast, relative to the standard carry trade, the crash-neutral strategy increases exposure to depreciations of the funding currency. While the payoff of the standard carry trade responds by $\exp(r_{f,t} \cdot \tau)$ to moves in the foreign exchange rate, the crash-neutral strategy responds by q_c , whenever \tilde{S}_{t+1} remains below K_c . Since $\delta_c > 0$, one can see that $q_c > \exp(r_{f,t} \cdot \tau)$. The payoff to this crash-neutral carry trade is illustrated *vis a vis* the payoff to the standard carry trade in the right panel of Figure 4. The crash-neutral carry trade is assumed to include a call option struck 2.5% out-of-the-money relative to the prevailing forward rate ($\frac{K_c}{F_t} = 1.0250$). Once again, as the call option is struck at a progressively higher price, offering less protection, its delta converges to zero, such that the payoff to the crash-neutral strategy converges to the payoff of the standard carry trade.

Finally, to compute the *returns* to the carry trade we divide the payoffs by the dollar value of the capital necessary to establish the positions. In the case of the portfolio which includes puts and is long the foreign currency, the funding capital is $(1 - q_p \cdot \delta_p) \cdot S_t + q_p \cdot P_t(K_p, \tau)$. For the portfolio which

¹¹Formally, the investor could fund the purchase of the calls at the lower, foreign interest rate. The assumption of funding at the domestic rate is made to preserve the symmetry of the solution, which is lost when additional currency risk is borne by the investor when he funds the purchase of the calls at the foreign rate.

is long the domestic currency and includes calls, the funding capital is $(1 + q_c \cdot \delta_c) \cdot S_t - q_c \cdot C_t(K_c, \tau)$. When the quantity of options in the portfolio goes to zero, the funding capital goes to S_t in both cases, as in the case for the standard carry trade. Intuitively, the standard carry trade is simply the limiting case of the crash-neutral trade as we let the strike price of the put (call) diverge to zero (infinity). In these cases, the prices of the FX options go to zero, and the two strategies offer identical payoffs.

3 Data

The results in this paper are based on two datasets. The first dataset contains information on one-, three-, six-month and one-year Eurocurrency (LIBOR) rates and was obtained from Datastream. Eurocurrency rates are the interest rates at which banks are willing to borrow and lend foreign currency deposits, and constitute the effective interest rates at which a carry trade investor would be able to borrow and lend. The LIBOR rates published by Datastream are from daily fixings of the British Bankers Association (BBA). The data are daily and cover the period from January 1990 through March 2007. In the event that BBA interest rate data are unavailable for the entire time period, the BBA series is spliced with the corresponding interbank lending rate published by the country's central bank, provided through Datastream or Global Financial Data. Time series means of the one-month LIBOR rates are reported in Panel A of Table IV. Daily exchange rates for the nine G10 currencies versus the U.S. dollar are obtained from Reuters via Datastream.

The second dataset, the foreign exchange (FX) option dataset, is comprised of daily implied volatility quotes at five strikes and four maturities for options on the G10 currencies and fifteen emerging currency pairs. The exchange rate options are European and give their owners the right to buy or sell a foreign currency at a pre-specified exchange rate measured in U.S. dollars per unit of foreign currency. The data on exchange rate options was obtained from J.P. Morgan and covers the period from January 1999 to March 2007. I begin by introducing some formalisms specific to the foreign exchange option markets, characterizing the structure of implied volatilities, and describing the procedure for converting the implied volatility data into observations of risk-neutral moments.

3.1 Foreign exchange options

FX option prices are quoted in terms of their Garman-Kohlhagen (1983) implied volatilities, much like equity options are quoted in terms of their Black-Scholes (1973) implied volatilities. In fact, the Garman-Kohlhagen valuation formula is equivalent to the Black-Scholes formula adjusted for the fact that both currencies pay a continuous yield given by their respective interest rates. The price of a call and put option can be recovered from the following formulas:

$$C_t(K, \tau) = e^{-r_{d,t} \cdot \tau} \cdot \left[F_{t,\tau} \cdot N(d_1) - K \cdot N(d_2) \right] \quad (18a)$$

$$P_t(K, \tau) = e^{-r_{d,t} \cdot \tau} \cdot \left[K \cdot N(-d_2) - F_{t,\tau} \cdot N(-d_1) \right] \quad (18b)$$

where:

$$d_1 = \frac{\ln F_{t,\tau}/K}{\sigma_t(K,\tau) \cdot \sqrt{\tau}} + \frac{1}{2} \cdot \sigma_t(K,\tau) \cdot \sqrt{\tau} \quad d_2 = d_1 - \sigma_t(K,\tau) \cdot \sqrt{\tau} \quad (19)$$

and $F_{t,\tau}$ is the forward rate for currency to be delivered τ periods forward, and $r_{f,t}$ and $r_{d,t}$ are the foreign and domestic interest rates for τ -period loans, respectively. The forward rate is determined through the covered interest parity condition, a no-arbitrage relationship which must hold at time t , and is equal to $S_t \cdot \exp\{(r_{d,t} - r_{f,t}) \cdot \tau\}$. The implied volatilities necessary to match the price of the τ -period options will generally depend on the option's strike value, K , and are denoted by $\sigma_t(K,\tau)$.

Unlike equity options which have fixed calendar expiration dates and are quoted at fixed strike prices, foreign exchange options are generally quoted at constant maturities and fixed deltas. More precisely, market makers quote prices of 0.25 and 0.10 delta risk reversals and butterfly spreads, as well as, an at-the-money delta-neutral straddle. The strike price of the straddle, for any given maturity, is chosen such that the deltas of a put and call at that strike are equal, but of opposite sign. From these data, one can compute implied volatilities at five strike values. The time-series averages of the option-implied volatilities at the five quoted strikes are reported in Panel A of Table IV.

The most frequently traded options have maturities of 1M, 3M, 6M and 1Y, and include at-the-money options, as well as, calls and puts with deltas of 0.25 and 0.10 (in absolute value). The option deltas, obtained by differentiating the option value with respect to the spot exchange rate, S_t , are given by,

$$\delta_c(K) = e^{-r_{f,t} \cdot \tau} \cdot N(d_1) \quad (20a)$$

$$\delta_p(K) = -e^{-r_{f,t} \cdot \tau} \cdot N(-d_1) \quad (20b)$$

allowing for conversion between the strike price of an option and its corresponding delta. Specifically, the strike prices of puts and calls with delta values of δ_p and δ_c , respectively, are given by:

$$K_{\delta_c} = F_t \cdot \exp\left(\frac{1}{2}\sigma_t(\delta_c)^2 \cdot \tau - \sigma_t(\delta_c) \cdot \sqrt{\tau} \cdot N^{-1}[\exp(r_{f,t} \cdot \tau) \cdot \delta_c]\right) \quad (21a)$$

$$K_{\delta_p} = F_t \cdot \exp\left(\frac{1}{2}\sigma_t(\delta_p)^2 \cdot \tau + \sigma_t(\delta_p) \cdot \sqrt{\tau} \cdot N^{-1}[-\exp(r_{f,t} \cdot \tau) \cdot \delta_p]\right) \quad (21b)$$

The strike price of the delta-neutral straddle is obtained by setting $\delta_c(K) + \delta_p(K) = 0$ and solving for K . It is straightforward to see that the options in this portfolio must both have deltas of 0.50 (in absolute value), and the corresponding strike value is:

$$K_{\text{ATM}} = S_t \cdot \exp\left((r_{d,t} - r_{f,t}) \cdot \tau - \frac{1}{2}\sigma_t(\text{ATM})^2 \cdot \tau\right) = F_t \cdot \exp\left(\frac{1}{2}\sigma_t(\text{ATM})^2 \cdot \tau\right) \quad (22)$$

Consequently, although the straddle volatility is described as “at-the-money,” the corresponding option strike is neither equal to the spot price or the forward price. In the data, the one-month 0.25 delta options are roughly 1.5-2.5% out-of-the-money, and the one-month 0.10 delta options are roughly 3.0-4.5% out-of-the-money (Table IV, Panel B). When normalized by the at-the-money implied volatility (converted to monthly units), the strikes of the 0.25 delta options are 0.70 standard deviations away from the forward price, and the 0.10 delta options are about 1.40 standard deviations away from the forward price. Finally, in standard FX option nomenclature an option with a delta of δ is typically referred to as a $|100 \cdot \delta|$ option. For example, a put with $\delta = -0.10$, is referred to as a 10δ put. I use this convention from hereon.

3.2 Extracting the risk-neutral moments

The formulas for the risk-neutral moments derived in Section 2 assume the existence of a continuum of out-of-the-money puts and calls. In reality, of course, the data are available only at a discrete set of strikes spanning a bounded range of strike values, $[K_{min}, K_{max}]$, such that any implementation of the moment formulas provides only an approximation to the true risk-neutral moments. It is therefore important to ensure that the available data are adequate to obtain a credible estimate of the underlying moments.

Jiang and Tian (2005) investigate these types of approximation errors in the context of computing estimates of the risk-neutral variance from observations of equity index option prices. To address this issue the authors simulate data from various types of models for the underlying asset and then seek to reconstruct the risk-neutral variance from a discrete set of observed option prices. They examine the impact of having observations on a finite number of options with a bounded range of strikes, as well as, the impact of various interpolation and extrapolation procedures. They conclude that the discreteness of available strikes is not a major issue, and that estimation errors decline to 2.5% (0.5%) of the true volatility when the most deep out-of-the-money options are struck at 1 (1.5) standard deviations away from the forward price. With options struck at two standard deviations away from the forward price, approximation errors essentially disappear completely. Moreover, their results indicate that approximation errors are minimized by interpolating the option implied volatilities within the observed range of strikes, and extrapolating the option implied volatilities below K_{min} and above K_{max} by appending flat tails at the level of the last observed implied volatility. Consistent with intuition, they find that this form of extrapolation is preferred to simply truncating the range of strikes used in the computation. Carr and Wu (2008) follow a similar protocol in their study of variance risk premia in the equity market, and combine linear interpolation between observed implied volatilities with appending flat tails beyond the last observed strikes.

Guided by the sensitivity results in Jiang and Tian (2005), the available cross-section of foreign exchange options is deemed to be sufficiently broad to ensure that the error in extracting the risk-neutral moments is likely to be very small. The furthest out-of-the-money puts and calls are struck at roughly 1.4 times the at-the-money implied volatility away from the prevailing forward prices.

Before extracting the risk-neutral moments, I augment the data by interpolating the implied volatility functions between the observed data points, and append flat tails beyond the last observed strike. I interpolate implied volatilities using the *vanna-volga method* (Castagna and Mercurio (2007)), which is the standard approach used by participants in the FX option market.¹² The resulting risk-neutral moments turn out to be largely unaffected by the precise details of the interpolation scheme, and similar results obtain if a standard linear interpolation is used, e.g. as in Carr and Wu (2008).

The vanna-volga method is based on a static hedging argument, and essentially prices a non-traded option by constructing and pricing a replicating portfolio, which matches all partial derivatives up to second order. In a Black-Scholes world, only first derivatives are matched dynamically, so the replicating delta-neutral portfolio is comprised only of a riskless bond and the underlying. However, in the presence of time-varying volatility, it is necessary to also hedge the *vega* $\left(\frac{\partial C^{BS}}{\partial \sigma}\right)$, as well as, the *volga* $\left(\frac{\partial^2 C^{BS}}{\partial \sigma^2}\right)$ and *vanna* $\left(\frac{\partial^2 C^{BS}}{\partial \sigma \partial S_t}\right)$. In order to match these three additional moments, the replicating portfolio must now also include an additional three traded options. Consequently, to the extent that at least three FX options are available, the implied volatilities of the remaining options can be obtained by constructing the relevant replicating portfolio, and then inverting its price to obtain the corresponding implied volatility. Castagna and Mercurio (2007) show that the interpolated implied volatility for a τ -period option at strike K obtained from the vanna-volga method is approximately related to the implied volatilities of three other traded option with the same maturity and strikes $K_1 < K_2 < K_3$ through:

$$\tilde{\sigma}_t(K, \tau) \approx \frac{\ln \frac{K_2}{K} \cdot \ln \frac{K_3}{K}}{\ln \frac{K_2}{K_1} \cdot \ln \frac{K_3}{K_1}} \cdot \sigma_t(K_1, \tau) + \frac{\ln \frac{K}{K_1} \cdot \ln \frac{K_3}{K}}{\ln \frac{K_2}{K_1} \cdot \ln \frac{K_3}{K_2}} \cdot \sigma_t(K_2, \tau) + \frac{\ln \frac{K}{K_1} \cdot \ln \frac{K}{K_2}}{\ln \frac{K_3}{K_1} \cdot \ln \frac{K_3}{K_2}} \cdot \sigma_t(K_3, \tau) \quad (23)$$

This formula provides a convenient shortcut for carrying out the interpolation and is known to provide very accurate estimates of the implied volatilities whenever K is between K_1 and K_3 (Castagna and Mercurio (2007)). Extrapolations based on this formula, however, lead to spurious results. Since the above approximation is essentially quadratic in the log strike, it violates the technical conditions for the existence of moments under the risk-neutral measure when extrapolated to infinity (Lee (2004)). As mentioned earlier, to avoid these issues I append flat implied volatility tails beyond the last observed strikes.

4 Results

In order to investigate the hypothesis that the empirically observed violations of uncovered interest parity are attributable to the exposure of high interest rate currencies to rapid depreciations, or crashes, I turn to data on foreign exchange options. Options provide a valuable tool for assessing market participants' perceptions of the underlying currency return distributions, as well as, the risk

¹²Common approaches in the equity option literature either rely on non-parametric methods (Ait-Sahalia and Lo (1998)) or fit *ad hoc* functional specifications to the observed data (Shimko (1993), Coval, Jurek and Stafford (2008)).

premia that are demanded for insurance against rapid currency moves. If high interest rate currencies are indeed more prone to crashes, one would expect any of the following to hold in the option data: (a) option-implied skewness should forecast currency excess returns with a negative sign; (b) option-implied skewness should be negatively correlated with interest rate differentials in the time series, as well as, the cross section; and (c) strategies in which the exposure to currency crashes has been hedged in the option market should earn zero excess returns. I find that: (a) option-implied skewness does not forecast currency excess returns; (b) option-implied skewness is positively related to interest rate differentials in the panel in univariate specifications, and is insignificant in a multivariate specification; and (c) carry trades in which crash risk has been hedged continue to deliver positive, albeit significantly smaller, excess returns. Taken together, the results for crash-hedged strategies indicate that at most 30-40% of the excess returns to currency carry trades can be interpreted as compensation for exposure to currency crashes. Simply put, the price of crash insurance in the currency markets appears to be relatively low, especially when contrasted with equity markets. In fact, in order for the currency excess returns on crash-hedged strategies to be driven to zero, the implied volatilities of options hedging against crashes would have to have been roughly *four* times higher than what is actually observed in the data.

4.1 Implied volatility functions

Before turning to formal tests of the crash risk hypothesis, it is useful to begin with a brief summary of the stylized features of the foreign exchange option data. Much like equity options, the implied volatility functions of foreign exchange options – the plots of volatility as a function of strike price – exhibit a pronounced smile. Unlike in equities though, where the smile is essentially strictly downward sloping, the smile can take on a variety of shapes, suggesting that the risk-neutral distribution can be either positively or negatively skewed. To summarize these features Figure 5 plots the time-series means of the implied volatility functions for the nine G10 currencies. The red (blue) lines correspond to periods in which the foreign short-term interest rate was above (below) the US short-term interest rate. Before taking means the volatilities were re-scaled by the contemporaneous at-the-money values to ensure a scale free representation. As can be readily seen, the shape of the implied volatility function exhibits significant variation across countries and time. For example, the implied volatility functions for the Swiss franc (CHF) and the Japanese yen (JPY) exhibit a right-skewed smile indicating a positively skewed risk-neutral distribution, suggestive of the potential for rapid appreciations against the U.S. dollar. During the 1999-2007 period, both of these currencies were characterized by low interest rates (time series means: 1.40% (CHF), 0.14% (JPY)), and are anecdotally known to have been popular funding currencies for the carry trade. Conversely, the implied volatility functions for the Australian dollar (AUD) and New Zealand dollar (NZD), which had relatively high interest rates during the sample (time series means: 5.28% (AUD), 6.06% (NZD)) and were the target currencies for carry traders, exhibit left-skewed smiles, consistent with a negatively skewed risk-neutral distributions. In sum, the cross-sectional evidence points to the fact that the exchange rates of relatively high (low) interest rate currencies expose investors to the risk of large depreciations (appreciations), consistent with the data in Brunnermeier, Nagel and

Pedersen (2008). Unconditionally, the risk-neutral distributions of high (low) interest rate currencies are negatively (positively) skewed, consistent with the crash risk hypothesis.

However, theories of the forward premium puzzle invoking crash risk as the source of the risk premium attaching to currencies with relatively higher interest rates, also make a prediction about the time series behavior of the implied-volatility smile. In particular, the skewness of the risk-neutral distribution should change sign conditional on the sign of the interest rate differential. Whenever a currency features an interest rate that is above (below) the U.S. interest rate, the risk-neutral exchange rate distribution should be negatively (positively) skewed. If this were the case, the blue lines would exhibit a steeper slope for moneyness values below at-the-money (depreciation) than for moneyness values above at-the-money (appreciation). Symmetrically, the red lines would exhibit a steeper slopes for moneyness values above at-the-money, than for values below at-the-money. In general, the data do not seem to be supportive of the *conditional* prediction of the crash risk hypothesis. Of course, one can plausibly argue that a currency’s risk-neutral skewness should be a function of the country’s interest rate differential relative to *all* available currencies, rather than just the U.S. dollar. After all, carry traders are not constrained to trade only in foreign/USD currency pairs. Consequently, one cannot make much of this finding, beyond emphasizing that carry trades are likely to be a global phenomenon.

4.2 Risk-neutral moments

In order to facilitate regression-based tests of the crash risk hypothesis, I summarize the currency option data by extracting the risk-neutral moments of the option-implied distribution. To do this I first interpolate the implied volatility data on each day using the vanna-volga method, and then construct and price the option portfolios replicating variance, skewness and kurtosis swaps, using the interpolated implied volatility function. The undiscounted values of these swaps correspond to the first three moments of the risk-neutral distribution. Since I use currency options with a one month maturity in the procedure, the resulting risk-neutral moments describe the one month ahead distribution. This procedure yields a daily time series of forward looking moments, which are plotted in Figure 6 for a representative set of currencies (AUD, CHF, EUR, GBP, JPY, NZD). The table accompanying Figure 6 provides the time series means and standard errors of the risk-neutral moments, computed from non-overlapping, monthly observations. The mean skewness values for essentially all currencies are statistically distinguishable from zero at conventional significance levels, and range from -0.16 to 0.32. Overall these absolute magnitudes are not extreme, as one would expect in the presence of the risk of large, but rare, crashes. Finally, the mean kurtosis values are all greater than three, consistent with heavy-tailed risk-neutral distributions.

The top panel of Figure 6 graphs the risk-neutral volatility ($\sqrt{\text{Var}_t^Q}$), and shows a clear common trend in the riskiness of the individual currency pairs.¹³ More interesting, is the behavior of risk-

¹³A principal components decomposition of the risk-neutral variance time series, reveals that the first principal component – roughly an equal-weighted combination of the individual series excluding the Canadian dollar – explains

neutral skewness (Skew_t^Q) which exhibits pronounced time-series and cross-sectional variation. The skewness values for the two high interest rates (AUD, NZD) move closely in tandem and are generally negative, with a time series mean of -0.16, rejecting the null of no skewness with considerable significance. By contrast, the mean skewness values for the two low interest rate currencies (CHF, JPY) are 0.10 and 0.32, respectively, both of which are statistically distinguishable from zero. The skewness of the Swiss franc moves in tandem with the Euro, which is perhaps somewhat surprising given the currencies' dramatically different interest rate policies. The risk-neutral skewness of the Japanese Yen is particularly volatile during the sample period reaching values as high as 1, and as low as -0.5. Overall, a principal components analysis reveals that 45% of the total variation in risk-neutral skewness is attributable to a common trend (equal-weighted portfolio). The next principal component explains roughly 30% of the variation and is essentially uniquely represented by the skewness of the Japanese Yen, reflecting its special position as a global funding currency. Finally, the risk-neutral kurtosis of all currency pairs (Kurt_t^Q) is plotted in the bottom panel of Figure 6. The risk-neutral kurtosis values are consistently above three, and average to about 3.6 for all currencies with the exception of the Japanese yen, whose mean kurtosis stands at 4.1.

4.3 Forecasting crashes and currency returns

The stochastic nature of risk-neutral skewness is suggestive of the fact that the risk of crashes and/or the price of insuring against those events is highly time-varying. To disentangle these two effects I turn to the data on the realized skewness of currency excess returns. The currency excess return for time t is defined as:

$$xs_t = \ln \frac{S_t}{S_{t-1}} - (r_{d,t} - r_{f,t}) \cdot \tau \quad (24)$$

and is essentially the continuously compounded return to buying a currency forward. Under UIP, the excess return is equal to zero (in expectation). To compute the realized skewness (Skew_t^P) for month t , I first construct the daily excess returns during that month and then compute their skewness. Panel A of Table V reports the results of panel regression using monthly observations; all panel regression specifications include country fixed effects.¹⁴ Much like Brunnermeier, Nagel and Pedersen (2008), I find that future values of realized skewness are negatively related to past excess returns and the lagged values of the interest rate spread ($r_{f,t} - r_{d,t}$). This indicates that currencies that have high interest rates and have recently appreciated – i.e. have been targets of successful carry trades – are more likely to experience large negative moves, or crashes. By contrast, the first regression in Panel B, indicates that risk-neutral skewness is strongly positively related to the realized currency return. The existence of a positive relationship is consistent with a risk-based story in which skewness acts

65% of total variation. The second principal component – which is long high interest rate currencies (AUD and NZD) and short low interest rate currencies (CHF, JPY) – explains an additional 20% of the variation.

¹⁴Results based on panel regressions using rolling 21-day windows for all specifications in Table V are available from the author upon request. After adjusting standard errors for within time-period correlation and serial correlation, the results are qualitatively and quantitatively indistinguishable from those obtained in the non-overlapping, monthly panel.

as a proxy for a priced risk factor (e.g. as in Farhi and Gabaix (2008)). Namely, as currencies become more negatively (positively) skewed investors charge a greater (smaller) risk premium, and the currency experiences a contemporaneous depreciation (appreciation). However, the evidence on future realized skewness appears to suggest that appreciated currencies become riskier, not safer. As a result, I find that risk-neutral skewness predicts future realized skewness with a *negative* sign and is highly statistically significant in a univariate specification. Taken together, the evidence suggests that the cost of hedging crashes following large appreciations is low, precisely when the risk of a crash is high.

Another result emerging from the panel regressions reported in Panel B is that risk-neutral skewness is *positively* related to the interest rate spread in the univariate panel regression, despite a exhibiting a strong negative relationship in the cross section. This indicates that periods of abnormally high spreads are associated with more positive values of risk-neutral skewness, contrary to the predictions of the crash risk hypothesis. Consistent with Brunnermeier, Nagel and Pedersen (2008), however, I find that this relationship becomes negative, but statistically insignificant in a multivariate specification. Judging by the dramatic difference between the values of the adjusted R^2 gross and net of the fixed effects during this time period, risk-neutral skewness appears to be more a fixed feature of a country, rather than a feature of its time-varying interest rate environment.

Finally, if violations of uncovered interest parity (UIP) are attributable to crash risk premia, the magnitude of risk-neutral skewness – a proxy of the market’s crash expectation – would be expected to forecast currency excess returns with a negative sign. Panel C of Table V investigates the forecastability of one-month excess returns, using lagged excess returns, interest rate spreads, as well as, realized and option-implied measures of skewness. I find a minute amount of momentum in excess returns at the monthly horizon and positive predictability on the basis of the interest rate differential. Not only does the option-implied skewness measure *not* forecast currency excess returns, it appears in the forecasting regression with a positive, albeit insignificant, regression coefficient.

4.4 Crash-neutral carry trade strategies

In order to determine whether the excess returns to carry trades can be attributed to compensation for exposure to currency crashes, I turn to an analysis of crash-neutral carry trades. Since these trades eliminate exposure to currency crashes by establishing a protective position in foreign exchange options, their mean excess returns should be lower than those of the standard carry trade. In the event that crash risk accounts for the entirety of the observed violations of UIP, excess returns to crash-neutral carry trades would be statistically indistinguishable from zero. To construct a time series of returns for the crash-neutral carry trades – implemented in each of the nine individual currency pairs – I proceed by analogy to the approach used in the standard carry trade. At each month end, I compare the prevailing one-month interbank lending rates, and establish the relevant positions in the spot markets and one-month foreign exchange options prescribed by (14) and (17). These positions are then held until the end of the following month, when the option expires. I construct

three variants of the crash-neutral strategy, each offering a different amount of crash protection, as reflected by the strike price (delta) of the included FX option. The summary statistics for these strategies are presented in Table VI. Panel A presents the results for crash-neutral strategies using “deep” out-of-the-money options (10δ calls and puts), which only provide protection against moves that are greater than roughly 1.4 times the magnitude of the monthly standard deviation. Panel B presents results for crash-neutral strategies employing options with intermediate moneyness levels (25δ calls and puts), and finally, Panel C presents the results for strategies using at-the-money options. Crucially, note that since the crash-neutral strategies only employ options for which tradable price data are available in the J. P. Morgan dataset, the accrued strategy returns represent returns that were attainable in the marketplace (before transaction costs). The results in this section do not rely on the implied volatility interpolation procedure used in the previous sections whatsoever.

4.4.1 Individual currency pairs

It is immediate from Panel A of Table VI that the addition of crash protection using 10δ options has a limited effect on the realized risk-return profile of the currency carry trade. Consistent with intuition, the addition of the crash-protection via the FX option causes the realized returns to be more positively skewed, and the magnitude of the minimal monthly returns to decline relative to standard carry trades. Similarly, one also observes a minor decline in the annualized standard deviation of the returns, although the level of volatility is still high and close to 10%. Of the four individual currency pairs (AUD, EUR, NZD, SEK) that delivered positive and statistically significant excess returns in the standard carry trade strategy, three continue to do so in the crash-neutral variant, and the t-statistic on the fourth pair (NZD) is only slightly below the requisite threshold for significance at the 5% level. However, tests comparing returns of the 10δ crash-neutral carry trade with the standard carry trade, also indicate a statistically significant decline in the mean realized return. For the four currency pairs which individually delivered positive and statistically significant returns, the mean returns decline by -1.36% (AUD; t-stat: -4.94), -1.28% (EUR; t-stat: -4.44), -1.03% (NZD; t-stat: -1.92) and -1.21% (SEK; t-stat: -3.19) per year. Shifting to crash-neutral strategies with more protection (25δ ; Panel B) causes a further decline in the realized returns and a coincident increase in skewness. Now all nine individual currency pairs have positively skewed excess returns, with monthly skewness values ranging from 0.13 (NZD) to 0.73 (NOK). Eight of those continue delivering positive excess returns, three of which are statistically significant at the 5%-level. When compared with the standard carry trades, the mean returns decline by -2.25% (AUD; t-stat: -2.29), -3.06% (EUR; t-stat: -3.78), -1.97% (NZD; t-stat: -1.52) and -3.32% (SEK; t-stat: -3.78), respectively.

Finally, Panel C of Table VI presents the corresponding summary statistics for crash-neutral currency strategies constructed from positions in the spot currency markets and at-the-money options. With ATM options, the strike price of the option is slightly above the forward exchange rate, and the carry trader effectively pays up front for protection against all negative deviations from uncovered interest rate parity. In other words, the investor locks in the interest rate differential (carry) and retains the upside from potential currency moves, in exchange for the option premium. By construc-

tion, the returns to these strategies are extremely positively skewed, with monthly skewness values generally above one. However, there is a dramatic decline of excess returns in comparison to the standard carry trades, with excess returns declining anywhere from -1.08% (CHF) to -6.44% (SEK) per year. Nonetheless, even with ATM hedging, the excess returns to the crash-neutral carry trade implemented in the AUD and SEK remain statistically significant.

4.4.2 Portfolio strategies

To summarize the findings for crash-neutral currency carry trades, I once again construct equal- and spread-weighted portfolios (*EQL* and *SPR*), as well as, their dollar-neutral counterparts (*EQL-\$N* and *SPR-\$N*). The returns to the non-dollar-neutral strategies are reported in Panel A of Table VII, and their dollar-neutral counterparts are reported in Panel B. This time an interesting dichotomy emerges depending on the method of construction. While the excess returns to the non-dollar-neutral portfolio remain positive and highly statistically significant, even when hedged with ATM options, the returns to the dollar-neutral portfolio are statistically indistinguishable from zero, even when hedged with 10δ options. However, this divergence in significance appears primarily attributable to the greater volatility of the non-dollar-neutral portfolios, which makes accurate measurement of their mean excess returns more difficult. The proportionate declines in excess returns for dollar-neutral and non-dollar-neutral portfolio are in fact comparable.

The excess returns on the non-dollar-neutral portfolios remain positive and highly statistically significant independent of the moneyness of the options used to hedge exposure to currency crashes. For example, the equal-weighted portfolio of crash-neutral carry trades implemented in the nine individual currency pairs delivers an annualized excess return of 3.18% (10δ ; t-stat: 3.13). As expected, the mean return declines with the level of protection offered by the option to 2.55% (25δ ; t-stat: 2.77) and 1.34% (ATM; t-stat: 1.81). As with the standard carry trades, the returns on the spread-weighted strategies are somewhat higher and more volatile than the returns to the equal-weighted portfolios. The combination of portfolio diversification and the benefit of option protection drive down the annualized volatilities to around 3-4% per year, such that corresponding portfolio Sharpe ratios continue to be large and range from 0.63 to 1.28. The cumulative performance of the three equal-weighted crash-neutral currency carry strategies is illustrated in Figure 7. As is immediate from the plot, even though excess returns remain positive with the most aggressively crash-hedged strategy (ATM), nearly half of the cumulative profits are eliminated, even when the least aggressive crash-hedging (10δ) is employed. The declines in the mean realized returns for the non-dollar neutral strategies range from -2.42% (t-stat: -4.73) to -6.05% (t-stat: -5.03) per year. Consequently, while excess returns on portfolio of crash-neutral trades remain positive, providing evidence against the hypothesis that crash risk premia can account for violations of UIP, they are also considerably smaller. When measured relative to the excess returns earned by the corresponding non-dollar-neutral portfolios of standard carry trades, these declines represent 30-40% of the total return to the unhedged strategy. Therefore, while UIP violations cannot be attributed solely to exposure to currency crashes, the crash risk explanation goes a considerable distance towards rationalizing the

forward premium anomaly.

4.5 Robustness

Although hedging currency crash risk using 10δ options eliminates 30-40% of the excess returns to the standard carry trade, one may wish to know by how much this result would improve once option transaction costs were accounted for. In particular, the implied volatilities used to price the options used in the crash-neutral portfolio represented midquotes and did not account for the existence of a bid-ask spread. To address this question, I replicate the returns to the crash-neutral strategies with perturbed values of the implied volatilities used to price the crash protection. In particular, I apply a simple multiplicative transformation to the implied volatilities, thus increasing the price of all options being purchased.

To assess the impact of transaction costs I apply a volatility multiplier of 1.1, since bid-ask spreads in foreign exchange option markets are between $\pm 0.5 - 1\%$ volatility points, and the mean option implied volatilities in the sample are on the order of 10%. With this modification, the mean excess returns to the 10δ crash-neutral carry trades remain essentially unchanged. The equal-weighted portfolio strategy delivers an excess return of 2.89% (t-stat: 2.81), and the spread-weighted portfolio delivers an excess return of 5.03% (t-stat: 3.46), corresponding to a 30 basis point decline in realized excess returns to the crash-neutral strategies versus the baseline specification without transaction costs. Therefore, typical transaction costs are insufficient to eliminate the profitability of crash-neutral currency carry trades.

Relatedly, one would also like to know, what would it have taken – in terms of a shift in option prices (implied volatilities) – in order for the excess returns to the crash-neutral strategies to have been equal to zero in the historical sample? By varying the implied volatility multiplier, I find that in order to eliminate the excess returns to crash-neutral carry trades requires implied volatilities that are between twice and four times as large, as the values actually observed in the data. For example, with a multiplier of two, excess returns on the equal- and spread-weighted portfolios are 1.38% (t-stat: 1.29) and 3.45% (t-stat: 2.32), respectively, and with a multiplier of *four*, they are –1.12% (t-stat: -1.04) and 0.69 (t-stat: 0.46), respectively. Needless to say, these magnitudes for the multiplier imply enormous mispricings in the foreign exchange option market. Consequently, while hedging exposure to crash risk helps provide a partial resolution of the UIP puzzle, it is unlikely to resolve it entirely.

5 Conclusion

Using option data on foreign exchange rates between the U.S. dollar and the other G10 currencies, covering the period from January 1999 to March 2007, I show that carry trade portfolios in which the exposure to currency crashes has been completely hedged, continue to deliver positive and statistically significant excess returns. The decline in the excess returns delivered by the crash-neutral portfolios

strategies *vis a vis* their unhedged counterparts, suggests that exposure to currency crashes can account for between 30-40% of the observed excess currency returns stemming from violations of UIP.

When implemented in the nine currency pairs involving the U.S. dollar and one of the remaining nine G10 currencies, simple equal- and spread-weighted portfolios of crash-neutral carry trades deliver Sharpe ratios between 0.63 and 1.28, while simultaneously exhibiting positively skewed excess returns. The mean excess returns to these strategies are highly statistically significant, and the annualized volatilities range from 2 to 4%. These findings are robust to various portfolio construction methodologies (weightings, dollar neutrality), as well as, the choice of the option strike used to hedge exposure to crash risk. As expected, the mean excess returns decline as investors hedge exposure to currency crashes more aggressively ($10\delta > 25\delta > \text{ATM}$), although excess returns continue to remain statistically significant even when at-the-money options are used. Consequently, unlike in equities, protection against the crashes in high interest rate currencies appears to be relatively cheap.

An analysis of the behavior of the risk-neutral and realized skewness presents evidence in support of the hypothesis that crash protection may be abnormally cheap. Specifically, while realized skewness tends to be negatively related to lagged currency returns, the relationship for risk-neutral skewness is positive. Therefore, while an appreciation of a high-interest rate currency – delivering positive carry trade returns – heralds an increased risk of a crash, the risk-neutral skewness simultaneously becomes more positive, implying that the prices of puts (calls) decrease (increase). However, even after accounting for this unusual behavior of option prices, crash risk is unlikely to be able to provide a complete resolution of the forward premium anomaly. In order to drive excess returns on the crash-hedged currency carry trades to zero, implied volatilities of out-of-the-money puts (calls) on high (low) interest rate currencies would have to be two to four times their actual observed values, implying significant pricing errors in the foreign exchange option market.

A Non-central moments of the risk-neutral distribution

The expressions for the second, third and fourth non-central moments are:

$$\begin{aligned}
 V_t(\tau) &= \int_{\bar{S}}^{\infty} \frac{2 \cdot \left(1 - \ln \frac{K}{\bar{S}}\right)}{K^2} \cdot C_t(K, \tau) dK + \int_0^{\bar{S}} \frac{2 \cdot \left(1 + \ln \frac{\bar{S}}{K}\right)}{K^2} \cdot P_t(K, \tau) dK \\
 W_t(\tau) &= \int_{\bar{S}}^{\infty} \frac{6 \cdot \ln \frac{K}{\bar{S}} - 3 \cdot \left(\ln \frac{K}{\bar{S}}\right)^2}{K^2} \cdot C_t(K, \tau) dK - \\
 &\quad - \int_0^{\bar{S}} \frac{6 \cdot \ln \frac{\bar{S}}{K} + 3 \cdot \left(\ln \frac{\bar{S}}{K}\right)^2}{K^2} \cdot P_t(K, \tau) dK \\
 X_t(\tau) &= \int_{\bar{S}}^{\infty} \frac{12 \cdot \left(\ln \frac{K}{\bar{S}}\right)^2 - 4 \cdot \left(\ln \frac{K}{\bar{S}}\right)^3}{K^2} \cdot C_t(K, \tau) dK + \\
 &\quad + \int_0^{\bar{S}} \frac{12 \cdot \left(\ln \frac{\bar{S}}{K}\right)^2 + 4 \cdot \left(\ln \frac{\bar{S}}{K}\right)^3}{K^2} \cdot P_t(K, \tau) dK
 \end{aligned}$$

References

- [1] Ait-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance*, 53, 499-547.
- [2] Arrow, Kenneth J., 1964, The Role of Securities in the Optimal Allocation of Risk Bearing, *Review of Economic Studies* 31, 91-96.
- [3] Backus, David K., Silverio Foresi and Chris Telmer, 2001, Affine Term Structure Models and the Forward Premium Anomaly, 56(1), p. 279-304.
- [4] Bakshi, Gurdip and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies*, 16(2), p. 527-566.
- [5] Bakshi, Gurdip, Nikunj Kapadia and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws and the Differential Pricing of Individual Equity Options, *Review of Financial Studies*, 16(1), p. 101-143.
- [6] Barro, Robert, 2006, Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics*, 121, p. 823-866.
- [7] Black, Fischer and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy*, 81, p. 617-654.
- [8] Breeden, Douglas and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business*, p. 621-651.
- [9] Britten-Jones, M. and A. Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance*, 55, p. 839-866.
- [10] Brunnermeier, Markus, Christian Gollier, and Jonathan Parker, 2007, Optimal Beliefs, Asset Prices and the Preference for Skewed Returns, *American Economic Review (Paper and Proceedings)*, 97(2), p. 159-165.

- [11] Brunnermeier, Markus, Stefan Nagel and Lasse Pedersen, 2008, Currency Crashes and the Carry Trade, *NBER Macroeconomics Annual*, 23, forthcoming.
- [12] Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo, 2006, The Returns to Currency Speculation, NBER working paper.
- [13] Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo, 2008, Do Peso Problems Explain the Returns to the Carry Trade?, NBER working paper.
- [14] Campa, Jose M, Kevin Chang and Robert Reider, 1998, Implied Exchange Rate Distributions: Evidence from OTC Option Markets, *Journal of International Money and Finance*, 17(1), p. 117-160.
- [15] Carr, Peter and Dilip Madan, 2001, Optimal Positioning in Derivative Securities, *Quantitative Finance*, 1, p. 19-37.
- [16] Carr, Peter and Liuren Wu, 2007, Stochastic Skew in Currency Options, *Journal of Financial Economics*, 86(1), p.213-247
- [17] Carr, Peter and Liuren Wu, 2008, Variance Risk Premiums, *Review of Financial Studies*, forthcoming.
- [18] Coval, Joshua D. and Tyler Shumway, 2001, Expected Option Returns, *Journal of Finance*, 56(3), p. 983-1009.
- [19] Coval, Joshua D., Jakub W. Jurek and Erik Stafford, 2008, Economic Catastrophe Bonds, *American Economic Review*, forthcoming.
- [20] Castagna, Antonio and Fabio Brigo 2007, Consistent Pricing of FX Options, *Risk*
- [21] Debreu, Gerard, 1959, *Theory of Value*, New York: Wiley.
- [22] DeMiguel, V., L. Garlappi and R. Uppal, 2007, Optimal versus Naive Diversification: How Inefficient Is the $1/N$ Portfolio Strategy?, *Review of Financial Studies*, forthcoming.
- [23] Driessen, Joost and Pascal Maenhout, 2006, The World Price of Jump and Volatility Risk, working paper.
- [24] Engel, Charles, 1996, The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence, *Journal of Empirical Finance*, 3, 123-192.
- [25] Fama, Eugene F., Forward and Spot Exchange Rates, 1984, *Journal of Monetary Economics*, 14, p. 319-338.
- [26] Farhi, Emmanuel and Xavier Gabaix, 2008, Rare Disasters and Exchange Rates, working paper.
- [27] Froot, Kenneth.A, and Richard Thaler, 1990, Anomalies: Foreign Exchange, *Journal of Economic Perspectives*, 4, p. 179-192.
- [28] Garman, Mark B. and Steven W. Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance*, 2, p. 231-237.
- [29] Hansen, Lars P. and Robert J. Hodrick, 1980, Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *Journal of Political Economy*, 88, p. 829-53.

- [30] Jiang, George J. and Yisong S. Tian, 2005, The Model-Free Implied Volatility and Its Information Content, *Review of Financial Studies*, 18(4), p. 1305-1342.
- [31] Jorion, Philippe, 1995, Predicting Volatility in the Foreign Exchange Market, *Journal of Finance*, 50(2), p. 507-528.
- [32] Lee, Roger. W., 2004, The Moment Formula for Implied Volatility at Extreme Strikes, *Mathematical Finance*, 14(3), p. 469-480.
- [33] Lewis, Karen K., Puzzles in International Financial Markets, in Gene Grossman and Kenneth Rogoff, eds., *Handbook of International Economics*, Volume 3 (Amsterdam; New York and Oxford: Elsevier, North- Holland, 1995).
- [34] Lustig, Hanno, and Adrien Verdelhan, 2007a, The Cross-Section of Foreign Currency Risk Premia and US Consumption Growth Risk, *American Economics Review*, 97(1), p. 89-117.
- [35] Lustig, Hanno, and Adrien Verdelhan, 2007b, Note on the Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk, working paper.
- [36] Lustig, Hanno, Nick Roussanov and Adrien Verdelhan, 2008, Common Risk Factors in Currency Markets, working paper.
- [37] Martin, Ian, 2008, Consumption-Based Asset Pricing with Higher Cumulants, working paper.
- [38] Pan, Jun, 2002, The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics*, 63, p. 3-50.
- [39] Plantin, Guillaume, and Hyun S. Shin, 2008, Carry Trades and Speculative Dynamics, working paper.
- [40] Rietz, Thomas A., 1988, The Equity Risk Premium: A Solution, *Journal of Monetary Economics*, 22, p. 117-131.
- [41] Shimko, David, 1993, Bounds of probability, *Risk* 6, 33-37.
- [42] Thompson, Samuel, 2006, Simple Formulas for Standard Errors that Cluster by Both Firm and Time, working paper.
- [43] Verdelhan, Adrien, 2008, A Habit-Based Explanation of the Exchange Rate Risk Premium, working paper.
- [44] Weitzman, Martin, 2007, Subjective Expectations and Asset-Return Puzzles, *American Economic Review*, 97, p. 1102-30.

Table I
UIP Regressions.

This table reports coefficient estimates from the regression of the time $t+1$ log currency return on the time t log forward premium,

$$s_{i,t+1} - s_{i,t} = a_0 + a_1 \cdot (f_{i,t} - s_{i,t}) + \varepsilon_{i,t+1} \quad H_0 : a_0 = 0, a_1 = 1$$

Currency returns are computed using 21-day rolling windows and span the period from January 1990 to March 2007 (all exchange rates are expressed in terms of dollars per unit of foreign currency). The forward premia are measured using the spread between one-month eurocurrency (LIBOR) rates for loans denominated in U.S. dollars and loans denominated in the foreign currency. The table reports regression coefficients, standard errors (in parentheses), and the χ^2 test statistic for the null hypothesis of UIP (p -values in parentheses). Standard errors in individual regressions are adjusted for serial correlation using a Newey-West covariance matrix with 21 lags. The pooled (panel) regression is run with country-fixed effects; the reported standard errors are robust to within time-period correlation and are adjusted for serial correlation. The pooled regression χ^2 statistic is computed for the null that all country fixed effects are zero and the intercept is equal to one. R_{NFE}^2 reports the adjusted R^2 from the panel regression net of the fixed effects ($N = 40,500$). XS reports the time series means and standard errors of the regression coefficients from cross-sectional regressions performed for each t . For the cross-sectional regressions R^2 is the mean adjusted R^2 .

Currency	1990-2007				1999-2007			
	\hat{a}_0	\hat{a}_1	R_{NFE}^2	χ^2 test	\hat{a}_0	\hat{a}_1	R_{NFE}^2	χ^2 test
AUD	-0.0025 (0.0023)	-1.7483 (1.0522)	0.0105	8.87 (0.01)	-0.0028 (0.0036)	-3.9018 (1.9520)	0.0310	9.34 (0.01)
CAD	-0.0001 (0.0009)	-0.5077 (0.5104)	0.0019	9.13 (0.01)	0.0027 (0.0015)	-2.5012 (2.1091)	0.0115	5.06 (0.08)
CHF	0.0026 (0.0024)	-1.2815 (1.0008)	0.0069	5.60 (0.06)	0.0096 (0.0041)	-4.5238 (1.8485)	0.0350	9.03 (0.01)
EUR	0.0002 (0.0016)	-0.0320 (0.9072)	-0.0002	1.34 (0.51)	0.0036 (0.0024)	-4.4836 (1.6590)	0.0447	11.19 (0.00)
GBP	0.0021 (0.0019)	0.7061 (1.2755)	0.0020	3.65 (0.16)	0.0001 (0.0025)	-1.7371 (1.7738)	0.0061	4.55 (0.10)
JPY	0.0058 (0.0025)	-2.0823 (0.8787)	0.0165	12.33 (0.00)	0.0048 (0.0045)	-1.8183 (1.4003)	0.0099	6.20 (0.05)
NOK	0.0013 (0.0017)	0.6255 (0.6351)	0.0042	1.43 (0.49)	0.0007 (0.0026)	-1.4005 (1.2175)	0.0090	5.46 (0.07)
NZD	-0.0047 (0.0034)	-2.4128 (1.1975)	0.0147	15.46 (0.00)	-0.0067 (0.0045)	-4.7728 (1.6837)	0.0482	17.15 (0.00)
SEK	0.0004 (0.0017)	0.6081 (0.6046)	0.0046	0.51 (0.77)	0.0026 (0.0024)	-3.5247 (1.3764)	0.0405	11.09 (0.00)
Pooled	FE	-0.1795 (0.6589)	0.0002	2.59 (0.99)	FE	-3.0503 (1.1190)	0.0248	24.17 (0.01)
XS	0.0005 (0.0003)	-0.1883 (0.0836)	0.1070	-	0.0012 (0.0005)	-0.5994 (0.1087)	0.0966	-

Table IIa
Standard Currency Carry Trade (1990-2007).

Panel A of the table reports summary statistics for returns from implementing the standard carry trade in currency pairs involving the U.S. dollar and one of the remaining nine G10 currencies. Returns are in U.S. dollars and are computed monthly for the period from January 1990 to March 2007 ($N = 207$). Means, standard deviations and Sharpe ratios (SR) are annualized. *Carry* is the mean absolute interest differential between the one-month Eurocurrency rate for country X and the one-month U.S. dollar rate, $|r_{f,t} - r_{d,t}|$. *Min* and *Max* report the smallest and largest observed monthly return. Panel B presents analogous summary statistics for portfolios of individual carry trades. The *EQL* strategy equal-weights the nine underlying carry trades, whereas *SPR* weights them by the absolute interest rate spread at initiation of the trade. The *EQL-\$N* and *SPR-\$N* are constructed analogously, by with the additional requirement that the portfolio be neutral with respect to exposure to the U.S. dollar.

Panel A: Individual currency pairs

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	0.0690	0.0319	0.0168	0.0637	0.0248	0.0482	0.0519	0.0399	0.0839
t-stat	3.19	2.30	0.65	2.78	1.12	1.79	2.16	1.80	3.28
Std. dev.	0.0898	0.0577	0.1074	0.0951	0.0916	0.1120	0.1000	0.0921	0.1061
Skewness	-0.30	-0.28	-0.47	-0.37	-0.68	-0.97	-0.21	-0.44	-0.73
Kurtosis	2.81	3.33	3.21	3.82	5.56	6.84	3.46	3.55	5.35
Min	-0.0685	-0.0435	-0.1005	-0.1010	-0.1188	-0.1640	-0.0933	-0.0790	-0.1452
Max	0.0649	0.0464	0.0645	0.0755	0.0727	0.1054	0.0756	0.0738	0.0826
Carry	0.0206	0.0141	0.0248	0.0210	0.0219	0.0314	0.0263	0.0289	0.0304
SR	0.77	0.55	0.16	0.67	0.27	0.43	0.52	0.43	0.79

Panel B: Portfolio strategies

	EQL	SPR	EQL-\$N	SPR-\$N
Mean	0.0478	0.0655	0.0399	0.0709
t-stat	3.91	4.10	2.74	3.65
Std. dev.	0.0507	0.0663	0.0603	0.0807
Skewness	-0.95	-0.82	-0.58	-0.73
Kurtosis	5.50	4.47	5.76	5.27
Min	-0.0580	-0.0648	-0.0737	-0.0892
Max	0.0375	0.0468	0.0588	0.0651
Carry	0.0244	0.0355	0.0369	0.0514
SR	0.94	0.99	0.66	0.88

Table IIb
Standard Currency Carry Trade (1999-2007).

Panel A of the table reports summary statistics for returns from implementing the standard carry trade in currency pairs involving the U.S. dollar and one of the remaining nine G10 currencies. Returns are in U.S. dollars and are computed monthly for the period from January 1999 to March 2007 ($N = 99$). Means, standard deviations and Sharpe ratios (SR) are annualized. *Carry* is the mean absolute interest differential between the one-month Eurocurrency rate for country X and the one-month U.S. dollar rate, $|r_{f,t} - r_{d,t}|$. *Min* and *Max* report the smallest and largest observed monthly return. Panel B presents analogous summary statistics for portfolios of individual carry trades. The *EQL* strategy equal-weights the nine underlying carry trades, whereas *SPR* weights them by the absolute interest rate spread at initiation of the trade. The *EQL-\$N* and *SPR-\$N* are constructed analogously, by with the additional requirement that the portfolio be neutral with respect to exposure to the U.S. dollar.

Panel A: Individual currency pairs

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	0.0879	0.0340	-0.0019	0.0910	0.0165	0.0341	0.0496	0.0784	0.1142
t-stat	2.51	1.43	-0.06	2.94	0.62	1.07	1.43	2.01	3.38
Std. dev.	0.1005	0.0686	0.0955	0.0888	0.0766	0.0914	0.0994	0.1119	0.0970
Skewness	-0.41	-0.26	-0.53	-0.20	-0.01	-0.04	-0.02	-0.56	-0.09
Kurtosis	2.72	2.71	2.63	3.24	2.64	2.94	3.04	2.96	2.82
Min	-0.0685	-0.0435	-0.0723	-0.0774	-0.0533	-0.0662	-0.0700	-0.079	-0.0608
Max	0.0649	0.0464	0.0458	0.0617	0.0526	0.0681	0.0721	0.0738	0.0826
Carry	0.0186	0.0084	0.0221	0.0151	0.0140	0.0348	0.0213	0.0263	0.0190
SR	0.87	0.50	-0.02	1.02	0.22	0.37	0.50	0.70	1.18

Panel B: Portfolio strategies

	EQL	SPR	EQL-\$N	SPR-\$N
Mean	0.0560	0.0844	0.0434	0.0699
t-stat	3.63	4.26	2.30	2.84
Std. dev.	0.0443	0.0569	0.0543	0.0707
Skewness	-0.42	-0.20	-0.52	-0.23
Kurtosis	3.73	2.97	4.31	4.33
Min	-0.0368	-0.0335	-0.0569	-0.0661
Max	0.0375	0.0455	0.0369	0.0651
Carry	0.0200	0.0307	0.0345	0.0478
SR	1.26	1.48	0.80	0.99

Table III
Summary Statistics.

Panel A reports the time-series means of the one-month interbank (LIBOR) lending rates and of the implied volatilities of one-month options on the exchange rates of individual currencies versus the U.S. dollar. The foreign exchange option data is comprised of implied volatilities on five options with standardized Black-Scholes deltas: puts and calls with deltas of 0.10 and 0.25 in absolute value (denoted by $10\delta p$, $25\delta p$, $25\delta c$ and $10\delta c$, respectively), and a delta-neutral straddle (denoted by ATM). All values are annualized. Panel B presents the time series means of the corresponding option strikes in terms of option moneyness (defined as the ratio of the option strike to the prevailing forward rate) and a standardized moneyness (defined as the log moneyness scaled by the at-the-money implied volatility per month). The data are daily and span the period from January 1999 to March 2007 ($N = 2149$).

Panel A: LIBOR and Implied Volatilities

Currency	$r_{f,t}$	$10\delta p$	$25\delta p$	ATM	$25\delta c$	$10\delta c$
AUD	0.0528	0.1191	0.1116	0.1071	0.1075	0.1114
CAD	0.0367	0.0799	0.0758	0.0736	0.0751	0.0785
CHF	0.0140	0.1098	0.1058	0.1051	0.1089	0.1155
EUR	0.0305	0.1062	0.1020	0.1011	0.1046	0.1109
GBP	0.0480	0.0890	0.0843	0.0824	0.0843	0.0886
JPY	0.0014	0.1068	0.1024	0.1039	0.1117	0.1238
NOK	0.0484	0.1129	0.1087	0.1079	0.1114	0.1176
NZD	0.0606	0.1322	0.1243	0.1195	0.1200	0.1241
SEK	0.0317	0.1141	0.1098	0.1088	0.1122	0.1184
USD	0.0361	-	-	-	-	-

Panel B: FX Option Strike Values

Currency	Moneyness $\left(\frac{K}{F_t}\right)$					Standardized moneyness $\left(\frac{1}{\sigma_{ATM}} \cdot \ln \frac{K}{F_t}\right)$				
	$10\delta p$	$25\delta p$	ATM	$25\delta c$	$10\delta c$	$10\delta p$	$25\delta p$	ATM	$25\delta c$	$10\delta c$
AUD	0.9580	0.9793	1.0005	1.0214	1.0423	-1.4005	-0.6824	0.0153	0.6903	1.3508
CAD	0.9714	0.9857	1.0002	1.0149	1.0295	-1.3812	-0.6830	0.0106	0.6952	1.3698
CHK	0.9611	0.9803	1.0005	1.0217	1.0437	-1.3223	-0.6625	0.0150	0.7137	1.4222
EUR	0.9623	0.9810	1.0004	1.0208	1.0420	-1.3293	-0.6637	0.0145	0.7116	1.4207
GBP	0.9682	0.9841	1.0003	1.0167	1.0335	-1.3668	-0.6745	0.0119	0.6993	1.3918
JPY	0.9624	0.9811	1.0005	1.0220	1.0464	-1.3108	-0.6528	0.0147	0.7394	1.5371
NOK	0.9600	0.9798	1.0005	1.0222	1.0446	-1.3236	-0.6611	0.0154	0.7099	1.4135
NZD	0.9532	0.9769	1.0006	1.0241	1.0475	-1.3927	-0.6789	0.0172	0.6913	1.3496
SEK	0.9596	0.9796	1.0005	1.0224	1.0450	-1.3238	-0.6622	0.0156	0.7104	1.4114

Table V
Forecasting Skewness and Currency Excess Returns.

Panel A reports panel regressions forecasting one-month ahead realized skewness (Skew_t^P), computed from daily currency excess returns during month t . Panel B reports the results of the corresponding panel regressions, but with risk-neutral skewness (Skew_t^Q) as the dependent variable. Risk-neutral skewness is extracted from one-month foreign exchange options and is measured on the last day of month t . Panel C reports panel regressions forecasting the one-month ahead currency excess return, (xs_t). The differential between the annualized foreign and domestic LIBOR rates is denoted by $r_{f,t} - r_{d,t}$ and is measured at the close of business on the last day of month t . The tables report the magnitudes of the regression coefficients and the associated standard errors adjusted for within time-period correlation (in parentheses). R^2 is the adjusted R^2 from the panel regression; R_{NFE}^2 is the adjusted R^2 net of the fixed effects. The data form a non-overlapping series of monthly observations spanning the period from January 1999 to March 2007 ($N = 882$ currency/months).

Panel A: Realized skewness						
	xs_t	$r_{f,t} - r_{d,t}$	Skew_t^P	Skew_t^Q	R^2	R_{NFE}^2
Skew_{t+1}^P	-2.9845 (0.7471)				0.0522	0.0199
Skew_{t+1}^P		-5.4556 (1.8684)			0.0590	0.0269
Skew_{t+1}^P			-0.0185 (0.0483)		0.0382	-0.0008
Skew_{t+1}^P				-0.4746 0.1557	0.0542	0.0220
Skew_{t+1}^P	-1.2211 (0.9328)	-4.6725 (1.9204)	-0.0281 (0.0506)	-0.2742 (0.1597)	0.0732	0.0416

Panel B: Option-implied skewness						
	xs_t	$r_{f,t} - r_{d,t}$	Skew_t^P	Skew_t^Q	R^2	R_{NFE}^2
Skew_{t+1}^Q	3.7361 (0.4601)				0.5776	0.3198
Skew_{t+1}^Q		2.144 (1.0279)			0.4044	0.0410
Skew_{t+1}^Q			0.0212 (0.0171)		0.3810	0.0032
Skew_{t+1}^Q				0.5865 (0.0293)	0.5912	0.3418
Skew_{t+1}^Q	3.3864 (0.4922)	-0.1951 (0.5574)	0.0270 (0.0078)	0.5606 (0.0336)	0.7628	0.6180

Panel C: Currency excess returns						
	xs_t	$r_{f,t} - r_{d,t}$	Skew_t^P	Skew_t^Q	R^2	R_{NFE}^2
xs_{t+1}	0.0935 (0.0515)				0.0067	0.0076
xs_{t+1}		0.3411 (0.1023)			0.0446	0.0456
xs_{t+1}			-0.0005 (0.0026)		-0.0020	-0.0010
xs_{t+1}				0.0129 (0.0085)	0.0052	0.0062
xs_{t+1}	0.0295 (0.0540)	0.3268 (0.1027)	0.0009 (0.0024)	0.0038 (0.0059)	0.0441	0.0450

Table VI
Crash-neutral Currency Carry Trades (1999-2007): Individual currency pairs.

This table reports summary statistics for returns from implementing crash-neutral carry trades in currency pairs involving the U.S. dollar and one of the remaining nine G10 currencies. Returns are computed monthly for the period from January 1999 to March 2007 ($N = 99$). Means, standard deviations and Sharpe ratios (SR) are annualized. Panel A presents returns for carry trades implementing crash-protection using 10δ options; Panel B panel – using 25δ options, and Panel C – using ATM options. *Min* and *Max* report the smallest and largest observed monthly return. Mean (diff) reports the time series mean of the differences in returns between the crash-neutral carry trade and its standard counterpart, along with the corresponding t-statistic.

Panel A: Crash-neutral strategy using 10δ -options.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	0.0743	0.0275	-0.0071	0.0782	0.0099	0.0180	0.0422	0.0681	0.1021
t-stat	2.20	1.22	-0.23	2.62	0.39	0.58	1.31	1.86	3.15
Std. dev.	0.0971	0.0650	0.0881	0.0858	0.0723	0.0894	0.0929	0.1050	0.0933
Skewness	-0.28	-0.06	-0.28	0.00	0.24	0.06	0.32	-0.32	0.08
Kurtosis	2.47	2.41	2.17	2.73	2.36	2.79	2.68	2.51	2.62
Min	-0.0557	-0.0366	-0.0581	-0.0617	-0.0433	-0.0587	-0.0507	-0.0624	-0.0598
Max	0.0634	0.0453	0.0445	0.0601	0.0513	0.0666	0.0708	0.0722	0.0805
SR	0.76	0.42	-0.08	0.91	0.14	0.20	0.45	0.65	1.10
Mean (diff)	-0.0136	-0.0066	-0.0052	-0.0128	-0.0067	-0.0162	-0.0074	-0.0103	-0.0121
t-stat (diff)	-4.94	-2.02	-0.91	-4.44	-1.77	-7.86	-1.28	-1.92	-3.19

Panel B: Crash-neutral strategy using 25δ -options.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	0.0654	0.0230	-0.0021	0.0605	0.0046	0.0123	0.0284	0.0587	0.0812
t-stat	2.25	1.19	-0.08	2.23	0.21	0.46	0.98	1.88	2.73
Std. dev.	0.0835	0.0558	0.0710	0.0779	0.0633	0.0766	0.0834	0.0897	0.0854
Skewness	0.19	0.38	0.18	0.36	0.61	0.59	0.73	0.13	0.40
Kurtosis	2.01	2.33	2.05	2.43	2.49	2.85	2.79	2.25	2.40
Min	-0.0335	-0.0239	-0.0374	-0.0380	-0.0262	-0.0377	-0.0367	-0.0412	-0.0347
Max	0.0604	0.0431	0.0417	0.0569	0.0486	0.0636	0.0679	0.0692	0.0762
SR	0.78	0.41	-0.03	0.78	0.07	0.16	0.34	0.65	0.95
Mean (diff)	-0.0225	-0.0110	-0.0002	-0.0305	-0.0119	-0.0219	-0.0212	-0.0197	-0.0330
t-stat (diff)	-2.29	-1.49	-0.01	-3.78	-1.41	-2.38	-1.92	-1.52	-3.78

Panel C: Crash-neutral strategy using ATM-options.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	0.0448	0.0118	-0.0127	0.0326	0.0034	0.0013	0.0084	0.0340	0.0500
t-stat	2.08	0.84	-0.76	1.56	0.22	0.06	0.37	1.50	2.15
Std. dev.	0.0621	0.0403	0.0479	0.0602	0.0453	0.0570	0.0652	0.0650	0.0667
Skewness	0.77	1.11	0.96	1.04	1.36	1.35	1.38	0.93	0.98
Kurtosis	2.58	3.39	2.86	3.07	4.19	4.36	4.01	3.08	3.07
Min	-0.0189	-0.0114	-0.0180	-0.0196	-0.0127	-0.0211	-0.0171	-0.0213	-0.0177
Max	0.0534	0.0375	0.0351	0.0494	0.0425	0.0562	0.0607	0.0619	0.0660
SR	0.72	0.29	-0.27	0.54	0.08	0.02	0.13	0.52	0.75
Mean (diff)	-0.0430	-0.0223	-0.0108	-0.0583	-0.0131	-0.0329	-0.0412	-0.0445	-0.0642
t-stat (diff)	-2.31	-1.65	-0.50	-3.76	-0.88	-1.84	-2.18	-1.95	-3.96

Table VII
Crash-neutral Currency Carry Trades (1999-2007): Portfolio strategies.

This table reports summary statistics for returns on portfolios of crash-neutral carry trades implemented in each of the currency pairs involving the U.S. dollar and one of the remaining nine G10 currencies. The component crash neutral strategies are denoted by *CNCT (10 δ)*, *CNCT (25 δ)* and *CNCT (ATM)*, to reflect the level of protection demanded. The *EQL* portfolio strategy equal-weights the crash-neutral carry trades in the nine underlying currency pairs, whereas *SPR* weights them by the absolute interest rate spread at initiation of the trade. Panel A presents returns to portfolios that are not constrained to be neutral with respect to dollar exposure. Panel B presents the corresponding returns for dollar-neutral strategies, *EQL-\$N* and *SPR-\$N*. Returns are computed monthly for the period from January 1999 to March 2007 ($N = 99$). Means, standard deviations and Sharpe ratios (SR) are annualized. *Min* and *Max* report the smallest and largest observed monthly return. Mean (diff) reports the time series mean of the differences in returns between the crash-neutral carry trade and its standard counterpart, along with the corresponding t-statistic.

Panel A: Non-dollar-neutral portfolios

	CNCT(10 δ)		CNCT(25 δ)		CNCT(ATM)	
	EQL	SPR	EQL	SPR	EQL	SPR
Mean	0.0318	0.0532	0.0255	0.0430	0.0134	0.0239
t-stat	3.13	3.69	2.77	3.32	1.81	2.35
Std. dev.	0.0291	0.0414	0.0265	0.0372	0.0212	0.0292
Skewness	-0.23	0.01	0.22	0.38	0.89	0.97
Kurtosis	3.51	3.40	3.37	3.41	3.63	3.79
Min	-0.0216	-0.0251	-0.0201	-0.0248	-0.0106	-0.0142
Max	0.0248	0.0391	0.0226	0.0356	0.0200	0.0272
SR	1.09	1.28	0.97	1.16	0.63	0.82
Mean (diff)	-0.0242	-0.0313	-0.0305	-0.0414	-0.0426	-0.0605
t-stat (diff)	-4.37	-5.36	-4.08	-5.03	-4.10	-5.03

Panel B: Dollar-neutral portfolios

	CNCT(10 δ)		CNCT(25 δ)		CNCT(ATM)	
	EQL-\$N	SPR-\$N	EQL-\$N	SPR-\$N	EQL-\$N	SPR-\$N
Mean	0.0257	0.0490	0.0195	0.0386	0.0009	0.0106
t-stat	1.48	2.13	1.23	1.87	0.06	0.59
Std. dev.	0.0500	0.0660	0.0455	0.0593	0.0416	0.0517
Skewness	-0.18	0.12	-0.06	0.31	0.32	0.57
Kurtosis	2.91	3.44	2.52	3.33	2.85	3.16
Min	-0.0379	-0.0406	-0.0279	-0.0330	-0.0239	-0.0254
Max	0.0341	0.0624	0.0291	0.0571	0.0310	0.0469
SR	0.51	0.74	0.43	0.65	0.02	0.21
Mean (diff)	-0.0177	-0.0209	-0.0239	-0.0313	-0.0426	-0.0592
t-stat (diff)	-3.90	-4.48	-2.23	-2.81	-2.45	-3.14

Figure 1. Total Return Indices. This figure illustrates the total return indices for the Fama-French factors, momentum and an equally-weighted carry trade implemented in G10 currency pairs involving the U.S. dollar. The total return indices are computed by compounding the per period excess returns to each strategy. The returns to each strategy are scaled *ex post* to match the volatility of the returns to the currency carry trade. Return are monthly and cover the period from January 1990 to March 2007 ($N = 207$). The table below reports summary statistics for the strategy returns; means, standard deviations and Sharpe ratios (SR) are annualized.

	MKT	SMB	HML	UMD	FX Carry
Mean	0.0730	0.0227	0.0477	0.0985	0.0478
t-stat	2.13	0.75	1.72	2.51	3.91
St. dev.	0.1422	0.1261	0.1153	0.1630	0.0507
Skewness	-0.68	0.81	0.11	-0.66	-0.95
SR	0.51	0.18	0.41	0.60	0.94

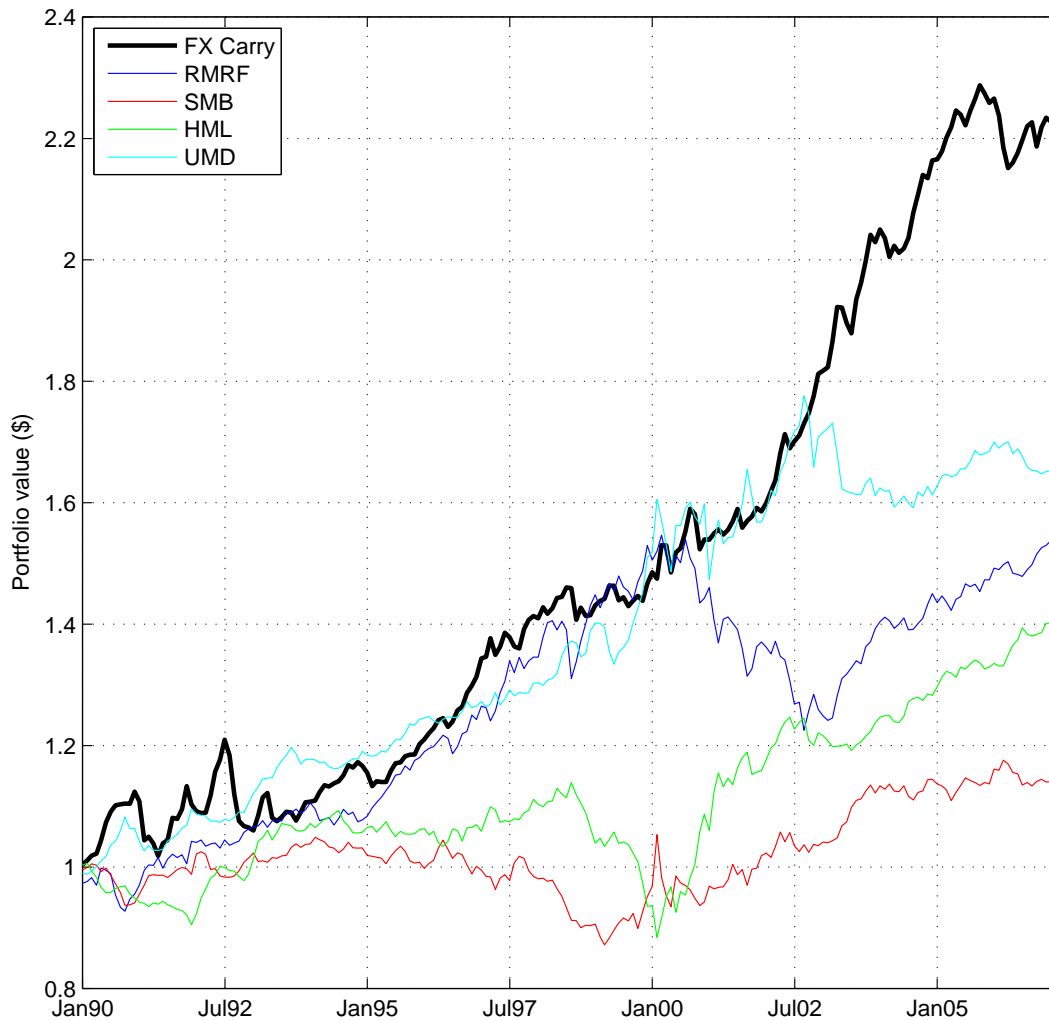


Figure 2. Returns to Standard Currency Carry Trades in Mean-Standard Deviation Space (1999-2007).

This figure depicts the returns to carry trades implemented in individual currency pairs involving the U.S. dollar and one of the nine remaining G10 currencies in mean/standard deviation space. *Equal Wgt* denotes a portfolio of standard carry trades which weights each of the underlying currency pairs equally. *Spread Wgt* denotes a portfolio of standard carry trades which weights the underlying currency pairs in proportion to their one-month interest rate differential. The dashed blue line represents the *ex post* mean variance frontier for zero-investment portfolios; the red-line represents the *ex post* mean variance frontier for portfolios combining the tangency portfolio with a riskless asset. The plot is constructed using monthly returns on the underlying pairs for the period from January 1999 to March 2007 ($N = 99$). Means and standard deviations are expressed in annualized units.

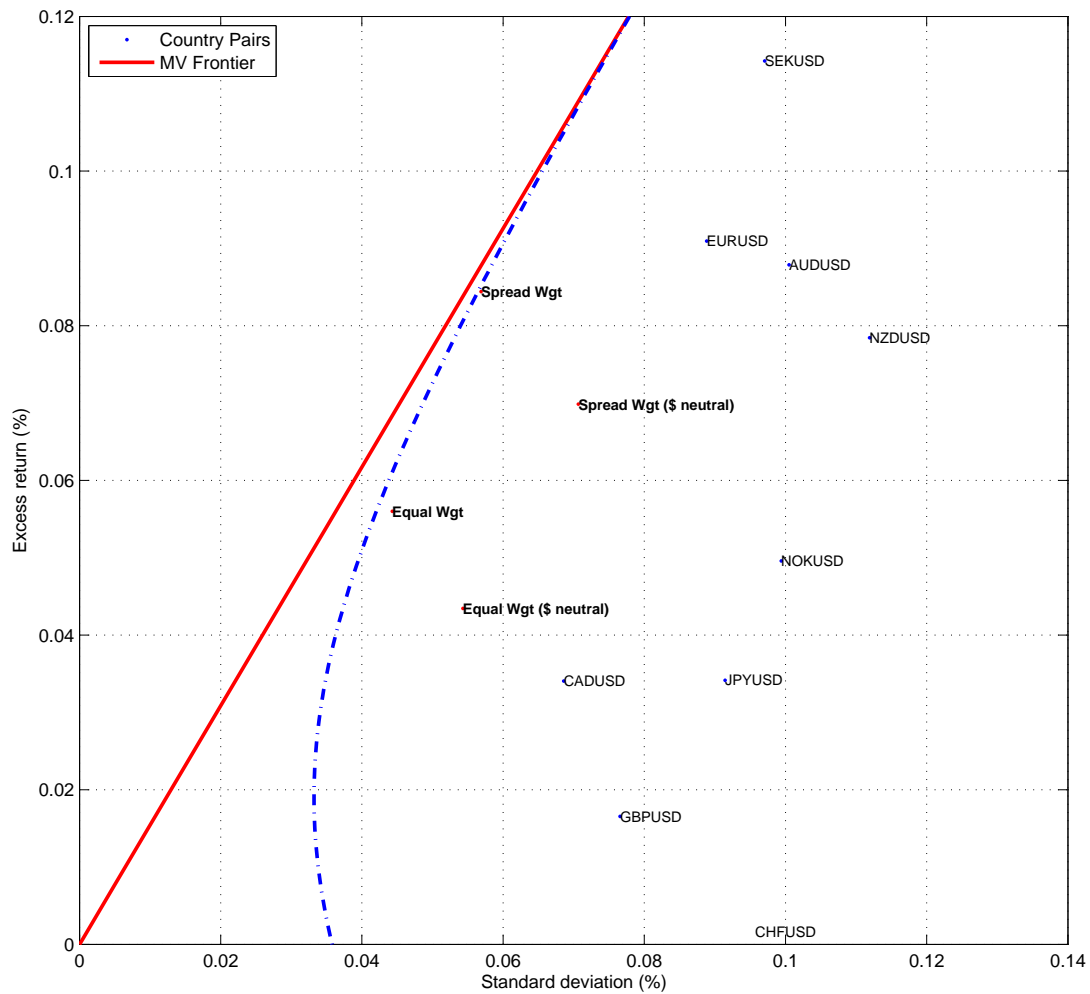


Figure 3. Option Portfolio Positions. This figure plots the positions of the replicating option portfolios used in computing the (non-central) moments of the risk-neutral distribution. Portfolio positions are plotted as a function of option moneyness ($\frac{K}{F_t}$). Positions for moneyness values less than (greater than) one apply to put (call) options, such that the replicating portfolio is constructed entirely of out-of-the-money options.

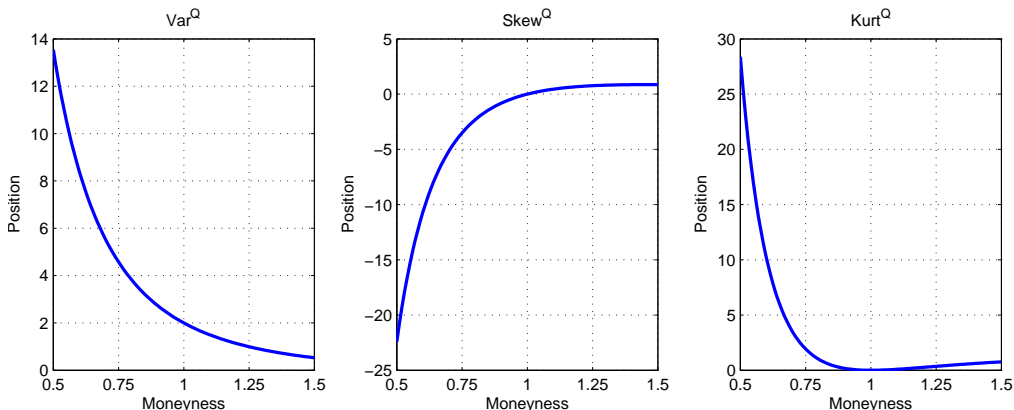


Figure 4. Carry Trade Payoff Diagram. This figure plots the payoff diagrams for the standard carry trade (dotted red) and the crash-neutral carry trade (solid blue). The left panel illustrates the strategy payoff when the foreign interest rate exceeds the domestic interest rate, and the investor is long the foreign currency. The right panel illustrates the symmetric case, when the domestic interest rate exceeds the foreign interest rate, and the investor is short the foreign currency. The crash-neutral carry trades use foreign exchange options struck 2.5% out-of-the-money relative to the prevailing forward rate.

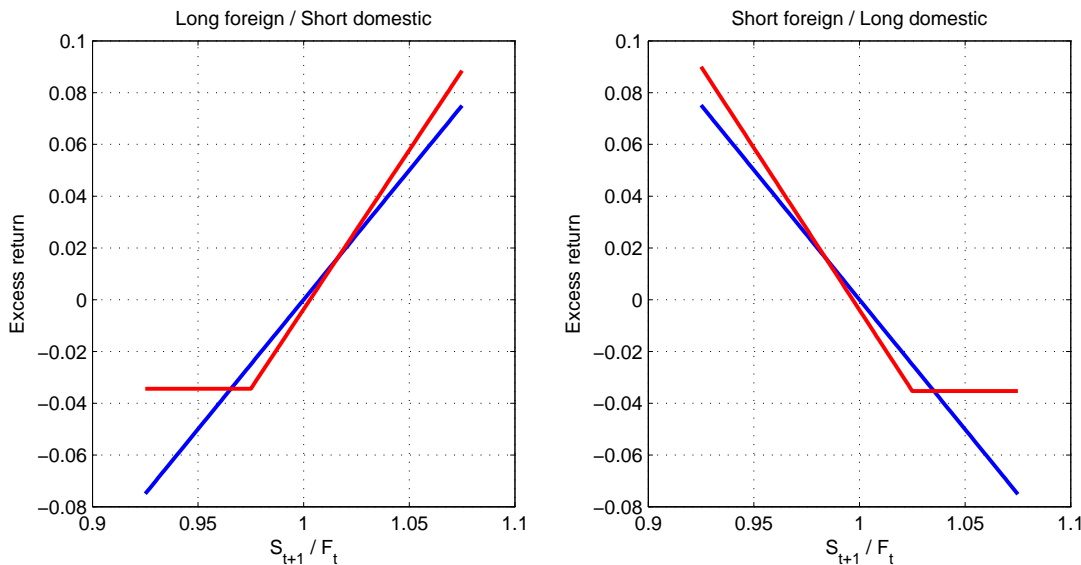


Figure 5. Implied Volatility Functions. The figure illustrates the average implied volatility functions for foreign exchange options on nine G10 currencies versus the U.S. dollar. The implied volatilities have been scaled by the contemporaneous at-the-money implied volatilities, and are presented for five standardized delta values (10δ put, 25δ put, ATM, 25δ call, 10δ call). The implied volatilities at the five standardized value are actual observed volatilities. All other volatilities were obtained by interpolating the data using the vanna-volga method. The red (blue) lines correspond to time series means for periods in which the foreign one-month Eurocurrency rate was above (below) the one-month U.S. interest rate. The underlying data are daily and cover the period from January 1999 to March 2007 ($N = 2149$).

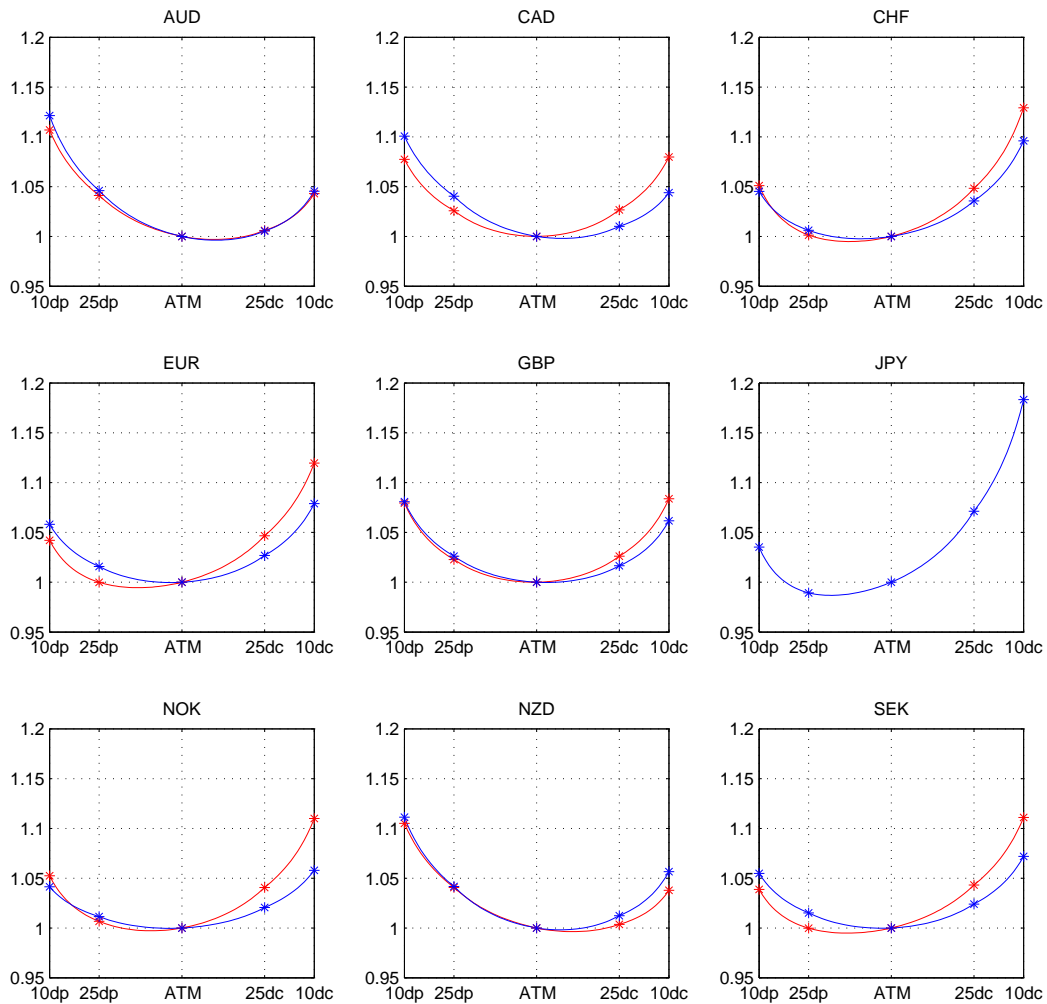


Figure 6. Risk-Neutral Moments of Currency Returns. The panels in this figure illustrate the time series of one-month option-implied values of risk-neutral volatility (square root of implied variance), skewness, and kurtosis. The currencies were selected to depict two currencies with relatively low interest rates versus the U.S. dollar (CHF, JPY) two with relatively high interest rates (AUD, NZD), and the two other major global currencies (EUR, GBP). The data in the figure are sampled weekly and cover the period from January 1999 to March 2007. The table reports the time series means of the risk-neutral moments and their standard errors (in parentheses) computed from non-overlapping, monthly observations.

Time series means and standard errors

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
$\sqrt{\text{Var}}^Q$	0.1102 (0.0023)	0.0761 (0.0013)	0.1082 (0.0018)	0.1043 (0.0022)	0.0855 (0.0013)	0.1079 (0.0026)	0.1113 (0.0017)	0.1237 (0.0023)	0.1121 (0.0017)
Skew ^Q	-0.1630 (0.0149)	-0.0701 (0.0200)	0.1015 (0.0156)	0.0807 (0.0160)	-0.0147 (0.0163)	0.3156 (0.0296)	0.0741 (0.0147)	-0.1625 (0.0139)	0.0594 (0.0143)
Kurt ^Q	3.6645 (0.0154)	3.6764 (0.0237)	3.6425 (0.0199)	3.6516 (0.0204)	3.6846 (0.0208)	4.1010 (0.0339)	3.6204 (0.0205)	3.6344 (0.0143)	3.6084 (0.0167)

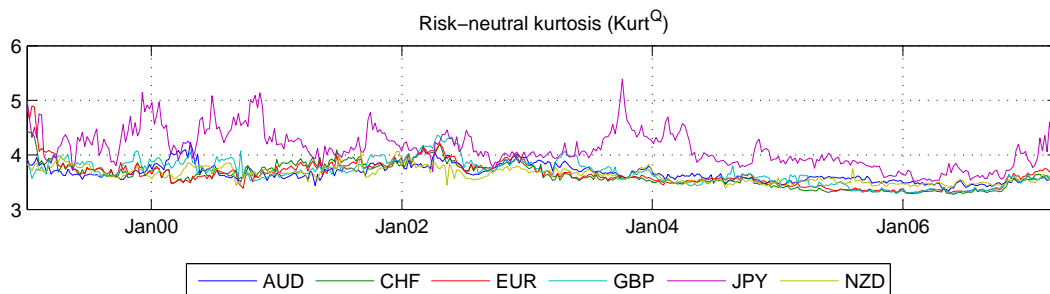
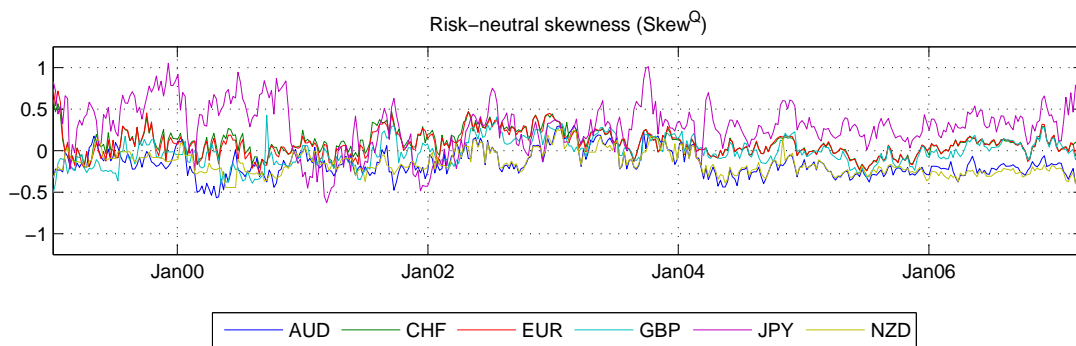
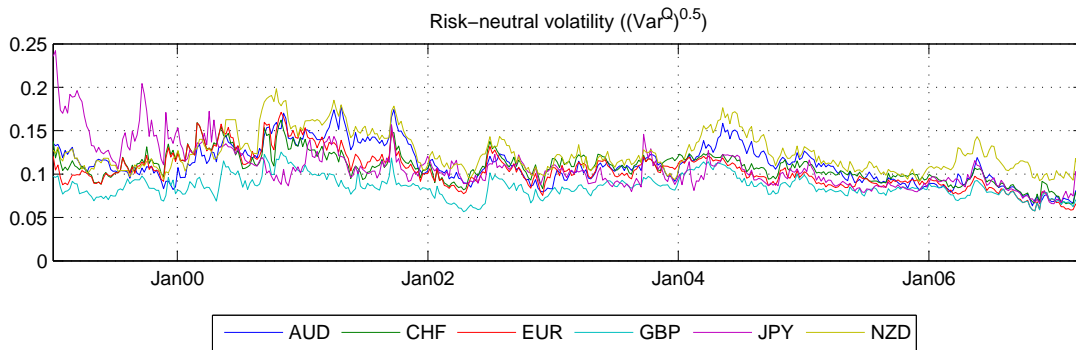


Figure 7. Crash-neutral Currency Carry Trade Strategy Returns. The figure depicts the total return indices for the standard carry trade and crash-neutral carry trades with various degrees of downside protection. The strategy returns are an equally-weighted combination of the carry trade returns for nine G10 currency pairs involving the U.S. dollar. Returns are computed monthly and cover the period from January 1999 to March 2007 ($N = 99$).

