Crash Risk in Currency Markets*

Emmanuel Farhi
Harvard University and NBER

Samuel P. Fraiberger
NYU

Xavier Gabaix
NYU Stern and NBER

Romain Ranciere
IMF and CEPR

Adrien Verdelhan
Boston University and NBER

June 1, 2009

Abstract

How much of carry trade excess returns can be explained by the presence of disaster risk? To answer this question, we propose a simple structural model that includes both Gaussian and disaster risk premia and can be estimated even in samples that do not contain disasters. The model points to a novel estimation procedure based on currency options with potentially different strikes. We implement this procedure on a large set of countries over the 1996–2008 period, forming portfolios of hedged and unhedged carry trade excess returns by sorting currencies based on their forward discounts. We find that disaster risk premia account for about 25% of carry trade excess returns in advanced countries.

*First version June 2008. Farhi: Department of Economics, Harvard University, and NBER, efarhi@fas.harvard.edu. Fraiberger: Department of Economics, New York University, spf248@nyu.edu. Gabaix: Stern School of Business, New York University, and NBER, xgabaix@stern.nyu.edu. Ranciere: IMF Research Department and CEPR, RRanciere@imf.org. Verdelhan: Department of Economics, Boston University, and NBER, av@bu.edu. Robert Tumarkin provided excellent research assistance. For helpful discussions and comments we thank Philippe Bacchetta, Eduardo Borenzstein, Robin Brooks, Markus Brunnermeier, Nicolas Coeurdacier, Chris Crowe, Francois Gourio, Bob King, Hanno Lustig, Borghan Narajabad, Jean-Charles Rochet, Hyun Shin, Kenneth Singleton and Stijn van Nieuwerburgh as well as seminar participants at Boston University, HEC Paris, HEC Lausanne, NY Fed, NYU, Paris School of Economics, Princeton, and the Texas Finance Festival. Farhi thanks the University of Chicago for its hospitality while this paper was being completed. Farhi and Gabaix gratefully acknowledge support from the NSF under grant 0820517. Ranciere gratefully acknowledges support from the IMF Research Grant Initiative.

1 Introduction

Currency carry trades offer large expected excess returns, challenging the benchmark models in international macroeconomics. In this paper, we explore whether a class of disaster-based models that postulate the existence of rare but large adverse aggregate shocks to stochastic discount factors can explain these excess returns. This class of models, pioneered by Rietz (1988) and Barro (2006), has received much attention recently in the macroeconomics and finance literature. However, this class of models is difficult to estimate because of the small number of disasters in sample. To address this difficulty, we provide a new method for estimating disaster risk premia even in samples that do not contain any disasters. We find that disaster risk premia are statistically significant and account for about one fourth of carry trade excess returns.

Currency carry trades are investment strategies where one borrows in low–interest rate currencies and invests in high–interest rate currencies. The value of the exchange rate at the end of the investment period is the sole source of risk. If investment currencies depreciate or funding currencies appreciate, then investors' returns decrease because they lose on their investment or must reimburse larger amounts. With risk-neutral and rational investors, high–interest rate currencies should depreciate on average against low–interest rate currencies and carry trade excess returns should be zero. Yet, in the data, these excess returns are large and positive on average. A natural explanation is that investors are risk averse and demand to be compensated for taking on such risk.

Carry trade investors, however, can use currency options to hedge this currency risk. For example, a domestic investor who is long in the foreign currency may buy a put contract that offers a large payoff in case the foreign currency depreciates. The investor thereby protects himself against adverse movements in the exchange rate. Likewise, a domestic investor who is short in the foreign currency may buy a call contract, protecting herself against an appreciation of the foreign currency. Using different currency option contracts, investors can tailor their exposure to exchange rate risk, buying protection against adverse exchange rate movements beyond any chosen cutoff. Intuitively, different hedged investment strategies should offer returns that are commensurate with their levels of risk. For example, the difference in returns between a strategy that is immune to large adverse changes in exchange rates and one that is not immune reflects the compensation for bearing the risk of a large currency depreciation. Yet a simple comparison across unhedged and hedged returns does not allow a precise estimation of disaster risk premia. The simple reason is that hedged strategies protect investors against large exchange rate changes of two types: those due to jumplike disasters and those that might occasionally happen without any jump in a world of Gaussian shocks.

In this paper, we propose a parsimonious exchange rate model to disentangle disaster from Gaussian risk premia. Following Backus, Foresi and Telmer (2001), we start off with the law of motion of the stochastic discount factor (SDF) in each country. These SDFs incorporate both a traditional log-normal component, as in Lustig, Roussanov and Verdelhan (2008), and a disaster component, as in Farhi and Gabaix (2008). We assume that financial markets are complete and thus define the change in exchange rate as the log difference between the domestic and foreign SDFs. In our model, expected currency excess returns are simply the sum of Gaussian and disaster risk premia. The former arise from random shocks observed every period, while the latter are due to rare disasters. We assume that these disasters do not occur in sample. As a consequence, changes in exchange rates follow a normal distribution in sample. Our model delivers closed-form solutions for short-dated put and call currency options, hedged currency excess returns, and risk reversals (traded option pairs that replicate a long out-of-the-money put position and a short out-of-the-money call position).[1] We use these expressions to establish a simple empirical procedure to measure the compensation for disaster risk. The decomposition of risk premia presented in this paper is a methodological contribution that could be useful in other asset markets.

We turn to currency data to implement our procedure and test the model's implications. To do so, we rely on currency spot, forward, and option contracts collected by JP Morgan for 32 countries. The data start in January 1996 and end in December 2008. Based on exchange rate normality tests, we restrict our sample in two dimensions: we focus on advanced countries and we exclude the fall of 2008. We take the view that the fall of 2008 corresponds to a unique disaster in our sample period, and we devote a final section to it. As a robustness check, we report in a separate appendix the results obtained with both advanced and emerging countries. Our data set comprises the prices of one-month options on bilateral exchange rates with different degrees of moneyness: far out-of-the-money puts (denoted 10 delta puts), out-of-the money puts (denoted 25 delta puts), at-the-money puts and calls, out-of-the-money calls (denoted 25 delta calls) and far out-of-the-money calls (denoted 10 delta calls).[2]

Following Lustig and Verdelhan (2007), we form portfolios of currency excess returns by sorting currencies based on their interest rates. We consider zero-investment strategies that go long in the highest–interest rate currencies and short in the lowest–interest rate currencies. We apply this methodology to both hedged and unhedged excess returns. Unhedged carry trades yield an average annual excess return of 6.5% in our sample. Carry trades hedged at 10 delta and 25 delta yield

[1]An option is said to be *at-the-money* if its strike price is equal to the forward exchange rate. A put (call) option is said to be *out-of-the-money* if its strike price is below (above) the forward rate—that is, if it takes a large depreciation (appreciation) to make the option worthwhile exercising. Figure 1 presents the payoffs of three option-based strategies considered throughout this paper: (i) being long an out-of-the-money put option, (ii) being long an out-of-the-money call option and (iii) being long a risk-reversal (i.e., being long an out-of-the-money put option and short an out-of-the-money call option with symmetric strikes.)

[2]The delta of an option represents its sensitivity to changes in the spot exchange rate. The delta of a put varies between 0 for extremely out-of-the-money options to -1 for extremely in-the-money-options. A 10 delta (25 delta) put is an option with a delta of 10% (25%). Figure 2 presents the deltas of put options as a function of their prices.

4.8% and 3.7% per annum, respectively, while carry trades hedged at the money yield 1.7% per annum. Hedged (except at-the-money) and unhedged returns and their differences are statistically all significant. Using Hansen's (1982) generalized method of moments (GMM) with at the money, 25-delta, and 10-delta options, we obtain a disaster risk premium of 1% per annum. This estimate is significantly different from zero, even after taking into account the small sample size. It represents approximately one fifth of unhedged carry excess returns. To maximize statistical precision, GMM puts relatively more weight on the deep out-of- the-money options. However, those out-of-the money options are likely to be the least liquid. A simple "equal weighted" estimator (with equal weights on the 10-delta, 25-delta, and at-the-money options) puts more weight than the GMM estimator on the more liquid options, which is preferable if the liquidity of out-of-the money options is a major concern.

We investigate the robustness of our results to the presence of transaction costs and counterparty risk. Bid–ask spreads are easily available on currency forward rates but not on options. We thus assume that bid–ask spreads are equal to 5% of implied volatilities for advanced countries and 10% for other countries.[3] As a result, our simulated bid–ask spreads increase in bad times. Their values are lower than the ones observed during the recent subprime mortgage crisis but correspond to market estimates. Taking into account bid–ask spreads and using GMM, we obtain a significant estimate of the disaster risk premium, which in this case is equal to 1.3% and represents one fourth of carry excess returns. This is our benchmark estimate. It is a lower bound because it does not take into account counterparty risk and because the GMM procedure puts relatively more weight on options that are deep out-of-the-money. We derive the sensitivity of this estimate to default probabilities on currency options markets.

The model also implies strong links between interest rates, contemporaneous and future changes in exchange rates, and the price of risk reversals – that is, the difference between the price of an out-of-the-money put option and the price of an out-of-the-money call option with symmetric strikes. Risk reversals capture the presence of asymmetric downside or upside risk. If the foreign currency is expected to depreciate, then out-of-the money puts should be more expensive than symmetric out-of-the-money calls. On the other hand, if exchange rates were normally distributed then symmetric puts and calls should have the same prices. The model predicts that: (i) risk reversals increase with interest rates; (ii) an increase in risk reversals is associated with a contemporaneous exchange rate depreciation reflecting the higher riskiness of the currency; and (iii) high values for risk reversals predict high average future currency returns because high exposure to disaster risk must be compensated by high returns. We check these predictions on individual countries, panel data, and

[3]The implied volatility is defined as the volatility necessary to match the observed option price using a standard Black–Scholes formula.

currency portfolios. Empirically, risk reversals increase with interest rates, as in the model. Protection against crash risk is more expensive for high–interest rate currencies than for low–interest rate ones. We find, as in the model, that increases in risk reversals and foreign currency depreciations tend to occur simultaneously. However, evidence is mixed as to whether risk reversals predict future exchange rates. Overall, risk reversals appear to contain useful information on potential disasters. Building portfolios on the basis of risk reversals delivers a monotonic cross-section of currency excess returns. The implied disaster risk premia is in line with our previous estimates.

We also examine the implications of our model for the implied volatility smile.[4] We present a simple calibration of the model that simultaneously matches our estimate of the disaster risk premium and provides a good fit for the smile observed in the data.

Overall, our model is not rejected by the data. We reach this conclusion by performing a J-test of the model's pricing errors. This validates our strategy of using a parsimonious and tractable model. In our view, resorting to a richer but more complex model would be justified only if we had access to a larger data set.

As a case study of a disaster episode, we use the fall of 2008. This period certainly represented bad times – corresponding to a high SDF – as evidenced by the deterioration in a large set of conventional risk measures. For example, during the fall 2008, the U.S. stock market index declined by 33% in terms of the MSCI index. Consistent with the disaster hypothesis, we document that the carry trade performed very poorly during that period: the cumulative loss amounted to 17.8% from September to December. This also represents an extreme drop from a statistical perspective, since the standard deviation of monthly carry trade returns over the whole sample is just 2%.

Our estimates of disaster risk premia and carry trade losses during fall 2008 are broadly consistent with the findings and calibration of Barro (2006) and Barro and Ursua (2008, 2009). In our model, the disaster risk premium depends on two main components: (i) the probability of disasters and the impact of disasters on SDFs, and (ii) the carry trade payoffs in times of disaster. We use the fall 2008 episode to calibrate the latter and the values in Barro and Ursua (2008) to characterize the former. These parameters imply a disaster risk premium of 2.8%, which is higher than but comparable to our estimate of 1.3%. This exercise should be viewed as a back-of-the-envelope calculation rather than a rigorous estimate, given that our inference relies on a single disaster.

Our paper is related to two different literatures: the forward premium puzzle and its potential explanations; and option pricing with jumps. Since the pioneering work of Hansen and Hodrick (1980) and Fama (1984), many papers have reported deviations from the uncovered interest rate parity (UIP) condition. These deviations are also known as the *forward premium puzzle*. In a

[4]The implied volatility of an option is a convenient normalization of the price of this option as a function of its strike. The smile refers to the relationship between the implied volatility and the strike. We provide formal definitions in Section 3 of the paper.

recent contribution, Lustig et al. (2008) build a cross-section of currency excess returns and show that it can be explained by covariances between returns and return-based risk factors. In order to replicate this result, stochastic discount factors must have not only a common component across countries but also heterogenous loadings on this common component. This paper builds on the disaster risk literature to satisfy this condition.[5] Our model derives from Farhi and Gabaix (2008), who augment the standard consumption-based model with disaster risk following Rietz (1988) and Barro (2006). World disaster risk is a common component, but countries differ in their exposures to world disasters. As a result, this paper contributes to the large literature on *peso problems* in international finance.[6]

Our paper also belongs to a recent literature using options to investigate the quantitative importance of disasters in currency markets. Bhansali (2007) was the first to document the empirical properties of hedged carry trade strategies. Brunnermeier, Nagel and Pedersen (2008) show that risk reversals increase with interest rates. In their view, the crash risk of the carry trade is due to a possible unwinding of hedge fund portfolios. This is consistent with one interpretation of disasters. Most closely related to this paper, Jurek (2008) provides a comprehensive empirical investigation of hedged carry trade strategies. He uses deep-out-of-the-money currency options to derive currency crash risk. Jurek's main result – that disaster risk explains 30% to 40% of carry trade returns – is consistent with the findings of this paper, but our approach differs in several dimensions. First, our model-based empirical strategy leads to a structural interpretation of the results. Second, the model allows us to use a variety of option strikes, including more-liquid at-the-money options, in order to disentangle Gaussian and disaster risk premia. Using at-the-money options, Burnside, Eichenbaum, Kleshchelski and Rebelo (2008) also find that disaster risk can account for the carry trade premium, where disaster risk comes in the form of a high value of the stochastic discount factor rather than large carry trade losses. In contrast to our approach, in their framework the only source of risk priced in carry trade returns is disaster risk.

---

[5] Other consumption-based models replicate the forward premium puzzle. Verdelhan (2009) uses habit preferences in the vein of Campbell and Cochrane (1999). Bansal and Shaliastovich (2008) build on the long-run risk model pioneered by Bansal and Yaron (2004). Guo (2007) presents a disaster-based model with monetary frictions. Ait-Sahalia, Wang and Yared (2001), Backus, Chernov and Martin (2009), Barro and Ursua (2009), Bates (2009), Bollerslev and Todorov (2009), Gabaix (2008), Gourio (2008), Julliard and Ghosh (2008), Liu, Pan and Wang (2005), Martin (2008), Pan (2002), Santa-Clara and Yan (2009) and Wachter (2008) all study disaster risk on equity and bond markets. Using swap rates, exchange rate returns, and prices of at-the-money currency options, Graveline (2006) estimates a two-country term structure model that replicates the forward premium anomaly. Barro (2009) studies the welfare costs of rare disasters.

[6] See Lewis (1995) for a survey. For example, Kaminsky (1993), extending the work of Engel and Hamilton (1990), considers the possibility of rare events explaining investors' expectations about exchange rates. Rare events in her model are infrequent switches from contractionary to expansionary monetary policy, and she provides evidence that investors' expectations are consistent with the model. However, she does not examine the forward premium puzzle and considers only one exchange rate (dollar–sterling) and a short time period.

A related literature studies high-frequency data and option pricing with jumps, following pioneering work by Bates (1996a, 1996b), who shows that exchange rate jumps are necessary to explain option smiles. More recently, Carr and Wu (2007) find great variations in the riskiness of two currencies (the yen and the British pound) against the U.S. dollar, and they relate it to stochastic risk premia. Campa, Chang and Reider (1998) document similar results for some European cross-rates. Bakshi, Carr and Wu (2008) find evidence that jump risk is priced in currency options. However, the jumps they consider are high-frequency jumps, whereas the disasters we have in mind are of very low frequency; in Barro (2006), disasters happen every 60 years. As a result, the economic analysis and our econometric technique are very different: we cannot directly measure disasters because they do not occur in our sample – unlike the small jumps that occur in studies such as Bakshi et al. (2008).

The paper is organized as follows. Section 2 presents our model and derives its main implications. Section 3 reports our empirical results and Section 4 concludes. A separate appendix reports proofs and empirical robustness checks.

2 Theory

We provide a simple model that serves as the basis for our empirical strategy. In the model, expected carry trade returns X^e correspond to the sum of two risk premia, a "normal times" or Gaussian risk premium π^G, and a disaster risk premium π^D:

$$X^e = \pi^D + \pi^G.$$

Here and in what follows, G refers to Gaussian and D refers to disaster.

Our main objective is to devise a simple structural estimation procedure to determine π^G, π^D and the fraction of carry trade returns due to disaster risk. To accomplish this, we use additional information from hedged carry trade returns. Hedged carry trades are zero-investment trades where the investor borrows in the funding currency and then uses the proceeds to invest in the investment currency and to purchase protection against a large depreciation of the investment currency through currency put options.[7] In the model, we derive closed-form solutions for the expected returns of hedged carry trades as a function of the option strikes. The expected return X^e_{hedged} of a hedged carry trade is

$$X^e_{\text{hedged}} = (1 + \Delta)\pi^G.$$

[7]In this simple overview, returns are computed in units of the funding currency. Later in the paper, we also treat the more general case where returns are computed in units of the investment currency.

In this formula, $\Delta \in (-1, 0)$ denotes the delta of the put option hedging the trade. The delta, which we define shortly, is increasing in the option strike. This is intuitive: the further away from the money, the more depreciation risk the investor bears and the higher the expected return of the hedged carry trade. We will make use of several strikes, with corresponding delta equal to -0.1 for deep-out-of-the-money options, -0.25 for out-of-the-money options, and -0.5 for at-the-money options. Hence the expected returns of a carry trade hedged deep out-of-the-money (10-delta), out-of-the-money (25-delta), and at-the-money (ATM) are respectively:

$$X^e_{\text{hedged, 10-delta}} = 0.9\pi^G, \quad X^e_{\text{hedged, 25-delta}} = 0.75\pi^G, \quad X^e_{\text{hedged, ATM}} = 0.5\pi^G.$$

To the best of our knowledge, this simple decomposition of hedged and unhedged returns is novel.

The rest of the section is devoted to setting up a model and deriving this result. Our modeling strategy follows Backus et al. (2001): we specify a stochastic discount factor for each country. These SDFs incorporate both a traditional log-normal component as in Lustig et al. (2008) and a disaster component as in Farhi and Gabaix (2008). This is enough to compute all relevant quantities, returns, and asset prices.

2.1 Model Setup

We focus on two countries, home and foreign, and develop a two-period model. In order to develop our empirical application, in Section 3 we explain how to incorporate this building block in a multi-country, multi-period extension. There, we introduce a state variable Ω_t that describes the state of the world. The parameters of our two-country, two-period model depend on Ω_t. All the results in this section should be understood as returns conditional on Ω_t, but for notational simplicity we do not make this dependence explicit. In particular, all the expectations in this section are conditional on Ω_t.

We assume that financial markets are complete but that some frictions prevent perfect risk sharing across countries.[8] Because we have data only for options on nominal exchange rates, we choose to consider only nominal returns. Therefore, our SDFs should be thought of as nominal SDFs (i.e., in units of local currency).[9]

[8] An example of such a friction often used in the literature is the assumption that some goods are not traded. The assumption of complete markets is not necessary. Technically, our theory requires only the absence of arbitrage and that risk-free bonds and options with enough strikes be traded. In other words, we rely on the existence of SDFs but do not need these SDFs to be unique.

[9] The link with real pricing kernels is well known. If $Q_{t,t+\tau}$ is the change in the quantity of real goods bought by one unit of the local currency and if $M^R_{t,t+\tau}$ is the real SDF, then the nominal SDF is $M_{t,t+\tau} = M^R_{t,t+\tau} Q_{t,t+\tau}$.

In the home country, the log SDF evolves as:

$$
\begin{aligned}
\log M_{t,t+\tau} &= -g\tau + \varepsilon\sqrt{\tau} - \frac{1}{2}\operatorname{var}(\varepsilon)\,\tau \\
&\quad + \left\{ \begin{array}{ll} 0 & \text{if there is no disaster at time } t+\tau \\ , \log(J) & \text{if there is a disaster at time } t+\tau \end{array} \right\}.
\end{aligned}
$$

We use a superscript star to denote foreign variables. The log of SDF in the foreign country evolves as:

$$
\begin{aligned}
\log M^{\star}_{t,t+\tau} &= -g^{\star}\tau + \varepsilon^{\star}\sqrt{\tau} - \frac{1}{2}\operatorname{var}(\varepsilon^{\star})\,\tau \\
&\quad + \left\{ \begin{array}{ll} 0 & \text{if there is no disaster at time } t+\tau \\ \log(J^{\star}) & \text{if there is a disaster at time } t+\tau \end{array} \right\}.
\end{aligned}
$$

Observe that the SDFs have two components. The first one, $-g\tau + \varepsilon\sqrt{\tau} - \frac{1}{2}\operatorname{var}(\varepsilon)\,\tau$, is a country-specific Gaussian risk with an arbitrary degree of correlation across countries. The second component, $\log(J)$, captures the impact of a disaster on the country's SDF.

The probability of a disaster between t and $t+\tau$ is given by $p\tau$. Note that disasters are perfectly correlated across the two countries: disasters are world disasters. Here, g and $g^{\star}$ are constants. The random variables $(\varepsilon, \varepsilon^{\star})$ are jointly normally distributed with mean 0 and may be correlated. However, $(\varepsilon, \varepsilon^{\star})$ are independent of the nonnegative random variables J and $J^{\star}$, which measure the magnitudes of the disaster event. All these variables are independent of the realization of the disaster event.

The "disaster" can have several interpretations. One, championed by Rietz (1988) and Barro (2006), is that of a macroeconomic drop in aggregate consumption, perhaps due to a war or a major economic crisis that affects many countries. Another interpretation is that of a financial stress or crisis affecting participants in world financial markets, perhaps via a drastic liquidity shortage and a violent drop in asset valuations. Both interpretations have merit, and we do not need to take a stand on the precise nature of a disaster.

This model is extremely tractable. Indeed, it yields closed-form solutions for a number of key moments of interest. However, this tractability does not come for free. It relies on a few important assumptions: that ϵ and ϵ^{*} are jointly normal and independent of the realization of the disaster. As we shall soon see, our model implies that, conditional on no disasters, the change in the exchange rate between home and foreign is an affine transformation of $\epsilon^{*} - \epsilon$. In Section 3 it is shown that, within our sample, we cannot reject the hypothesis that the distribution of monthly log exchange rate changes conditional on no disaster being lognormal.[10] This validates our assumption that

[10] At very high frequencies, exchange rates exhibit fat-tailed distributions. In line with the central limit theorem,

$\epsilon^* - \epsilon$ is normally distributed and independent of the realization of disasters. However, our model presumes not only that $\epsilon^* - \epsilon$ is normal but also that ϵ and ϵ^* are both normal.[11] This assumption on pricing kernels is harder to confront directly with the data. Section 3.2 provides an overall test of the fit of the model and fails to reject it. This result validates our overall strategy of building a simple and parsimonious model that is consistent with the data.

2.2 Interest Rates and Exchange Rates

In a complete markets economy such as ours, the change in the (nominal) exchange rate is given by the ratio of the SDFs (Backus et al., 2001):

$$\frac{S_{t+\tau}}{S_t} = \frac{M^{\star}_{t,t+\tau}}{M_{t,t+\tau}},$$

where S is measured in home currency per foreign currency. An increase in S represents an appreciation of the foreign currency. The exchange rate moves both in normal times and in disasters. In normal times, the exchange rate increases following a good realization of the home Gaussian risk ε or a bad realization of the foreign Gaussian risk $\varepsilon^{\star}$. In disasters, the exchange rate increases following a good realization of J or a bad realization of $J^{\star}$.

It is important to note that a low realization of $J^{\star}$ corresponds to a depreciation of the foreign currency. Hence, a country's exposure to disaster risk increases when the distribution of $J^{\star}$ decreases in the first-order stochastic dominance sense. Actually, we will see shortly that a summary statics for the foreign country's exposure to disaster risk is $-pE[J^{\star} - 1]$.

The home interest rate r is determined by the Euler equation $1 = E\left[M_{t,t+\tau}e^{r\tau}\right]$:

$$r = g - \log\left(1 + p\tau E\left[J - 1\right]\right)/\tau. \tag{1}$$

A similar expression determines the foreign interest rate. In the limit of small time intervals, this expression takes a very simple form.

Proposition 1. *In the limit of small time intervals $\tau \to 0$, the interest rate r in the home country is given by*

$$r = g - pE\left[J - 1\right].$$

A similar formula holds for the foreign interest rate. Ceteris paribus, if the foreign country has a higher average disaster risk or a lower $pE\left[J^{*} - 1\right]$, then it also has a higher interest rate. This

however, monthly changes in exchange rates very often appear to be Gaussian.

[11] In Section 3, we return to this issue and discuss how relaxing this hypothesis could potentially help us reduce the sensitivity of the estimated disaster risk premium on the strikes of the options used for its estimation.

higher interest can be understood as compensation for the risk of holding a currency that tends to depreciate in disasters when the SDF is high.

2.3 Options

To determine the payoffs of hedged carry trades, we need to specify some option-related notation. We denote by $P_{t,t+\tau}(K)$ and $C_{t,t+\tau}(K)$ the prices of one-period puts and calls on the home–foreign currency pair: $P_{t,t+\tau}(K)$ is the home currency price of a put yielding $(K - S_{t+\tau}/S_t)^+$ in the home currency, and $C_{t,t+\tau}(K)$ is the home currency price of a call yielding $(S_{t+\tau}/S_t - K)^+$ in the home currency.[12]

The Black–Scholes formula. Our closed-form solutions for hedged carry trade returns build on a version of the Black-Scholes formula. This formula, developed originally by Black and Scholes (1973) in the context of stocks, was adapted to a foreign exchange setting by Garman and Kohlhagen (1983). We denote by $V^P_{BS}(S, K, \sigma, r, r^\star, \tau)$ and $V^C_{BS}(S, K, \sigma, r, r^\star, \tau)$ the Black–Scholes price for a put and a call, respectively, when the spot is S, the strike is K, the volatility is σ, the time to maturity is τ, the home interest rate is r, and the foreign interest rate is $r^\star$. For example, the Black–Scholes price of a put is given by

$$V^P_{BS}(S, K, \sigma, r, r^\star, \tau) = Ke^{-r\tau}\mathbb{N}(-d_2) - Se^{-r^\star\tau}\mathbb{N}(-d_1),$$

where $\mathbb{N}$ is the cumulative distribution function of a Gaussian and where

$$d_1 = \frac{\log(S/K) + (r - r^\star + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \qquad d_2 = d_1 - \sigma\sqrt{\tau}.$$

The Black–Scholes formula has a simple scaling property with respect to the time to maturity τ and the interest rates r and $r^\star$:

$$V^P_{BS}(S, K, \sigma, r, r^\star, \tau) = V^P_{BS}(Se^{-r^\star\tau}, Ke^{-r\tau}, \sigma\sqrt{\tau}, 0, 0, 1).$$

This scaling property means that we can use the formula whenever the time to maturity is equal to 1 and both interest rates are 0. For notational convenience, we will omit the arguments 0 and 1 and simply write

$$V^P_{BS}(S, K, \sigma) \equiv V^P_{BS}(S, K, \sigma, 0, 0, 1).$$

[12] We use the notation: $y^+ \equiv \max(0, y)$.

The delta of options. The delta of an option is the sensitivity (or the partial derivative) of the option price to a change in the underlying exchange rate. The delta of a put is negative because the value of a put increases when the underlying currency depreciates. The delta decreases with the strike of a put: a deep-out-of-the-money put has a delta close to 0, while a deep-in-the-money has a delta close to $-e^{-r^{\star}\tau}$. For example, in the Black–Scholes model, the delta of a put is given by

$$\partial V_{BS}^{P}(S, K, \sigma, r, r^{\star}, \tau)/\partial S = -e^{-r^{\star}\tau}\mathbb{N}(-d_1).$$

We will often consider the limit of short time to maturity. The delta of the option then has a simple interpretation: it is the probability that the put will be exercised. More formally, the delta of a put option with time to maturity τ and strike $Se^{\kappa\sqrt{\tau}}$ has the following limit:[13]

$$\Delta_{BS}^{P}(\kappa) \equiv \lim_{\tau\to 0} \partial V_{BS}^{P}(S, Se^{\kappa\sqrt{\tau}}, \sigma, r, r^{\star}, \tau)/\partial S = -\mathbb{N}(\kappa/\sigma) \in (-1, 0),$$

where the partial derivative is taken with respect to the first argument.

For example, $\kappa = 0$ for at-the-money options and so the delta of an ATM put is $-1/2$.

2.4 Hedged and Unhedged Carry Trade Returns

We compute returns in units of the home currency. However, we want to allow for the possibility that home might be both the funding currency (if $r < r^{\star}$) and the investment currency (if $r > r^{\star}$). Hence we define two carry trade payoffs X and Y that correspond to these two cases:

$$X_{t,t+\tau} = e^{r^{\star}\tau}\frac{S_{t+\tau}}{S_t} - e^{r\tau} \ ;$$
$$Y_{t,t+\tau} = -X_{t,t+\tau}.$$

The payoff $X_{t,t+\tau}$ corresponds to the following trade: at date t, borrow one unit of the home currency at rate r and invest the proceeds in the foreign currency at rate $r^{\star}$. At the end of the trade, at date $t+\tau$, convert the proceeds back into the home currency. The payoff $Y_{t,t+\tau} = -X_{t,t+\tau}$ corresponds to the opposite trade.

In the main text, we treat the case where the home currency is the funding currency ($r < r^{*}$). The corresponding derivations can be found in Appendix A. In Appendix B, we derive the corresponding results for the case where home is the investment currency.

We now construct the hedged carry trade returns, $X_{t,t+\tau}(K)$. The return $X_{t,t+\tau}(K)$ is the payoff of the following zero-investment trade: borrow one unit of the home currency at interest

[13] In this equation, κ is a normalized measure of the moneyness of the option.

rate r ; use the proceeds to buy $\lambda^P_{t,t+\tau}(K)$ puts with strike K, protecting against a depreciation in the foreign currency; and invest the remainder $\left(1-\lambda^P_{t,t+\tau}(K)P_{t,t+\tau}(K)\right)$ in the foreign currency at interest rate $r^\star$. Here $P_{t,t+\tau}(K)$ is the home currency price of a put yielding $(K - S_{t+\tau}/S_t)^+$ in the home currency,

$$X_{t,t+\tau}(K) = \left(1-\lambda^P_{t,t+\tau}(K)\,P_{t,t+\tau}(K)\right)e^{r^\star\tau}\frac{S_{t+\tau}}{S_t} + \lambda^P_{t,t+\tau}(K)\left(K-\frac{S_{t+\tau}}{S_t}\right)^+ - e^{r\tau},$$

where we choose the hedge ratio $\lambda^P_{t,t+\tau}(K)$ to eliminate disaster risk:

$$\lambda^P_{t,t+\tau}(K) = e^{r^\star\tau}/\left(1+P(K)\,e^{r^\star\tau}\right).$$

Of foremost interest to us is the annualized expected returns, conditional on no disasters, of two strategies: the unhedged carry trade, X^e, and the hedged carry trades, $X^e(\kappa)$, at strike $e^{\kappa\sqrt{\tau}}$ over short horizons τ. These returns correspond to the following limiting cases:

$$X^e = \lim_{\tau\to 0} E^{ND}\left[X_{t,t+\tau}\right]/\tau,$$
$$X^e(\kappa) = \lim_{\tau\to 0} E^{ND}\left[X_{t,t+\tau}\left(e^{\kappa\sqrt{\tau}}\right)\right]/\tau.$$

To summarize our notation: $X_{t,t+\tau}$ denotes the carry trade return and X^e is its expected value; $X_{t,t+\tau}(e^{\kappa\sqrt{\tau}})$ denotes the hedged carry trade return with strike $K = e^{\kappa\sqrt{\tau}}$ and $X^e(\kappa)$ is the expected value of that hedged carry trade return. E^{ND} denotes expectations under the assumption of no disaster.

The following proposition offers a decomposition of these returns in terms of disaster and Gaussian risk premia.

Proposition 2. *In the limit of small time intervals ($\tau \to 0$), carry trade expected returns (conditional on no disasters) are given by*

$$X^e = pE\left[J - J^\star\right] + \operatorname{cov}\left(\varepsilon, \varepsilon - \varepsilon^\star\right). \tag{2}$$

In the same limit, hedged carry trade expected returns (conditional on no disasters) are given by

$$X^e(\kappa) = -pE\left[(J^\star - J)^+\right] + \operatorname{cov}\left(\varepsilon, \varepsilon-\varepsilon^\star\right)\left(1+\Delta^P_{BS}(\kappa)\right). \tag{3}$$

The first term in equation (2) is the risk premium associated with disaster risk:

$$\pi^D \equiv pE\left[J - J^\star\right].$$

If the foreign country is riskier, then $E[J - J^{\star}] > 0$ and the expected return due to disaster risk is positive. The second term in (2) is the risk premium associated with "Gaussian risk" à la Backus et al. (2001):[14]

$$\pi^{G} \equiv \operatorname{cov}(\varepsilon, \varepsilon - \varepsilon^{\star});$$

this is the covariance between the home SDF and the bilateral exchange rate $S_{t+\tau}/S_t$. In our model, the expected return of the carry trade compensates for the exposure to these two sources of risk.

The purchase of protection against extreme depreciation affects the loading of the carry trade payoff on the two sources of risk in the model. This is reflected in the expression for the expected value of the hedged carry trade return in equation (3). The disaster risk premium π^D is reduced to $pE[(J^{\star} - J)^{+}]$, which equals zero if $J > J^{\star}$ almost surely. The Gaussian risk premium π^G is reduced to $\operatorname{cov}(\varepsilon, \varepsilon - \varepsilon^{\star})\left(1 + \Delta^{P}_{BS}(\kappa)\right)$. This can be understood as follows: because the put option has a sensitivity to currency changes that is equal to the option delta $\Delta^{P}_{BS}(\kappa)$, hedging reduces the risk premium corresponding to Gaussian risk by $\operatorname{cov}(\varepsilon, \varepsilon - \varepsilon^{\star})\,|\Delta^{P}_{BS}(\kappa)|$. We will expand on the intuition for this term in Section 2.5.

Implied volatilities. To put Proposition 2 to work, we use implied volatilities. The implied volatility $\hat{\sigma}_{t,t+\tau}(K)$ of a put with strike K is defined implicitly as the volatility that would make the Black–Scholes price match the observed price of the option:

$$P_{t,t+\tau}(K) = e^{-r^{*}\tau} V^{P}_{BS}\left(1, K e^{(r^{\star}-r)\tau}, \hat{\sigma}_{t,t+\tau}(K)\sqrt{\tau}\right).$$

A similar definition holds for call options. By the put–call parity formula, the implied volatility of a put and a call having the same strike and maturity are equal. We now state a lemma that will simplify the empirical analysis.

Lemma 1. *In the limit of small time intervals ($\tau \to 0$), the Black–Scholes implied volatility $\hat{\sigma}_{t,t+\tau}\left(e^{\kappa\sqrt{\tau}}\right)$ of a put or a call with strike $e^{\kappa\sqrt{\tau}}$ is given by* $\operatorname{var}(\varepsilon^{\star} - \varepsilon)^{1/2}$.

Lemma 1 states that, in the limit of small time intervals, the implied volatility is equal to the

[14] Backus et al. (2001) show that, if markets are complete and SDFs are log normal, then expected log currency excess returns are equal to $E(\log R^e) = 1/2Var(\log M) - 1/2Var(\log M^{\star})$. We focus here instead on the log of expected currency excess returns, but the two expressions are naturally consistent. Starting from Backus et al. (2001), we obtain:

$$\begin{aligned} \log E(R^e) &= E(\log R^e) + \frac{1}{2}var(R^e) = \frac{1}{2}var(\varepsilon) - \frac{1}{2}var(\varepsilon^{\star}) + \frac{1}{2}var(\varepsilon - \varepsilon^{\star}) \\ &= var(\varepsilon) - cov(\varepsilon, \varepsilon^{\star}); \end{aligned}$$

physical Gaussian volatility of the bilateral exchange rate, $\operatorname{var}(\varepsilon^\star - \varepsilon)^{1/2}$. This is true even though our model contains both normal-times risk and disaster risk. The intuition is as follows. For options close to the money, the value of the option due to disasters is proportional to $p\tau$, the probability that the disaster will occur during the lifetime τ of the option. This is very small compared to the value of the option due to normal-times volatility, which is proportional to $\sqrt{\tau}$. Hence, for small maturities and strikes close to the money, most of the value of the option comes from Gaussian risk rather than disaster risk. Correspondingly, the implied volatility of the option is well approximated by the physical volatility of the exchange rate.

In the case of short-dated options with near-the-money strikes, Lemma 1 implies that we can use the Black–Scholes implied volatilities $\hat{\sigma}_{t,t+\tau}\left(e^{\kappa\sqrt{\tau}}\right)$ instead of the physical Gaussian volatility $\operatorname{var}(\varepsilon^\star - \varepsilon)^{1/2}$ when computing $\Delta^P_{BS}(\kappa)$ in equation (3). This is true even though – owing to the presence of disasters – the assumptions of the Black–Scholes model do not hold.

As a result, we need not to forecast future volatility country by country (which would be difficult given that market participants have more information than we do). We can instead rely on option-implied volatilities. The quality of this approximation deteriorates for out-of-the-money options, in which cases the implied volatility is larger than the physical volatility. Our procedure will then bias our estimates of option deltas away from 0, leading to an overestimation of Gaussian risk premia and an underestimation of disaster risk premia.

Rather than using the underlying options strike, traders in practice routinely use its Black–Scholes delta, which is a conventional quantity computed as

$$-e^{-r^*\tau}\mathbb{N}\left(\frac{\kappa\sqrt{\tau} + (r - r^* - \hat{\sigma}^2/2)\,\tau}{\hat{\sigma}\sqrt{\tau}}\right).$$

Note that this quantity might differ from the true sensitivity of the option with respect to the fundamental. However, it converges to $\Delta^P_{BS}(\kappa) = -e^{-r^\star\tau}\mathbb{N}(-d_1)$ in the limit of small time intervals.

Using Lemma 1 therefore provides us with a useful simplification: in the limit of small time intervals, the conventional deltas that traders use to quote currency options coincide not only with the true deltas of the options but also with the quantity $\Delta^P_{BS}(\kappa)$ featured in our model.

In practice, this approximation is valid when the disaster risk premium $p(J^* - J)\tau$ is small in absolute value compared to the option price, which is of order $\xi\sigma\sqrt{\tau}$ (where $\xi > 0$ depends on κ). Therefore, our approximation will be valid only if $\tau \ll \left(\xi\sigma/\left(p\left|J - J^\star\right|\right)\right)^2$. Numerically, with yearly units volatility is about 10% so $\sigma \simeq 0.1$. The disaster part of the carry trade risk premium is, in order of magnitude, 1.5%, so $p\left|J^* - J\right| \simeq 0.015$.[15] Thus we need $\tau \ll 44\xi^2$. For at-the-money options, $\xi = 1/\sqrt{2\pi}$ and the condition is $\tau \ll 44\xi^2 = 6.9$ years. Because we use one-month options

[15] For this analysis we need not to decompose the relative contributions of p and $J^* - J$, as Farhi and Gabaix (2008) do. Only the value of the disaster risk premium, $p(J^* - J)\tau$, matters.

($\tau = 1/12$), our approximation is expected to be valid in practice. Furthermore, in practice the ratio of the implied volatility of 10-delta and 25-delta options to the implied volatility of ATM options typically lies between 1 and 1.2. Hence, using the volatility ATM rather than the implied volatility at 10-delta would change the factor $1 + \Delta$ of 10-delta options from 0.9 to 0.94; for the 25-delta options, the $1 + \Delta$ factor would be equal to 0.79 instead of 0.75.[16] These corrections would imply only trivial modifications to our empirical estimates, much below their reported standard errors.

2.5 Estimating the Contribution of Disasters

The expected return of the unhedged carry trade in equation (2) can be re-expressed as

$$X^e = \pi^D + \pi^G. \tag{4}$$

Assume that $J^\star < J$ almost surely: this means that the exchange rate of the foreign country will depreciate with respect to the home country in case of a disaster. A put option protects the investor against a large depreciation associated with disasters and also against more a modest depreciation resulting from Gaussian risk. As a consequence, the hedged carry trade is less risky and commands a lower risk premium. The further out of the money the put option is, the more risk the investor bears, and so the higher the hedged carry trade return. Indeed, we can re-express (3) as

$$X^e(\kappa) = \pi^G \left(1 + \Delta^P_{BS}(\kappa)\right).$$

For instance, take the carry trade hedged with at-the-money options ($\kappa = 0$). In this case, $\Delta^P_{BS}(\kappa) = -1/2$ and $X^e(\kappa) = 0.5\pi^G$. The expected return of the carry trade hedged at the money is equal to half of the no-disaster risk premium π^G.[17]

The intuition here is that the hedge eliminates all the disaster risk and half the Gaussian risk. That exactly half of the Gaussian risk is eliminated might seem surprising, given that the SDF puts more weight on depreciation of the foreign currency than on its appreciation. The intuition is as follows. In the limit of small time horizons $\tau \to 0$, the shape of the distribution is a Gaussian with standard deviation $\sigma\sqrt{\tau}$, whereas the adjustments for risk that govern the difference between the physical and risk-adjusted probability are much smaller – of the order of magnitudes of τ. Together with the fact that the Gaussian distribution is symmetric around 0, this implies $X^e(0) = 0.5\pi^G$.

[16] With an upper bound of 1.1, the numbers are 0.92 and 0.77; with an upper bound of 1.3, they are 0.95 and 0.81.

[17] An informal explanation runs as follows. The carry trade has a "disaster beta" of 1, and a "Gaussian" beta of 1. Hence, its risk premium is $\pi^D + \pi^G$. On the other hand, the carry trade hedged at the money has a disaster beta of 0 and a Gaussian risk beta of 1/2 (as we saw earlier, it eliminates half the Gaussian risk). Hence, its risk premium is $0.5\pi^G$. Likewise, the carry trade hedged at 10 delta has a disaster beta of 0 and a Gaussian risk beta of 0.9 (because it eliminates 10% of the Gaussian risk), so its risk premium is $0.9\pi^G$.

Next, take the carry trade hedged with a put option at 25 delta. In the language of currency traders, this means that the strike is such that the delta of the put is -0.25; thus $X^e(\kappa) = 0.75\pi^G$. Likewise, for the carry trade hedged at 10 delta, we get $X^e(\kappa) = 0.9\pi^G$. Again, the intuition is that, since that the hedge uses a relatively deep-out-of-the-money put, investors bear much of the Gaussian risk but not all of it: they bear 90% of the risk, so that the expected return of the carry trade at 10 delta is 0.9 times the Gaussian risk premium.

The method behind our estimation procedure is to use expected returns of different investment strategies with different loadings on disaster and Gaussian risks to derive π^G and π^D. Alternatively, option prices can also be used directly to make some inference about those risk premia. We turn to this issue in the next section.

2.6 Risk Reversals

Roughly speaking, if the foreign currency is riskier than the home currency, then out-of-the-money put prices on the currency pair (home, foreign) should be higher than out-of-the-money call prices, since the price of protection against a devaluation of the foreign currency should be high. In this section we construct a simple metric – risk reversals – to measure the gap between the out-of-the-money puts and out-of-the money calls.

One tradition is to construct risk reversals as the implied volatility of an out-of-the-money put, minus the implied volatility of a symmetric out-of-the-money call. A more theoretically appealing definition for our purposes involves looking at the difference between the prices of put and calls rather than at the difference between their implied volatilities. More precisely, we call $\mathcal{F} = e^{(r-r^*)\tau}$ the *forward rate* of the bilateral exchange rate $S_{t+\tau}/S_t$. We use k, which in practice is close to 1, in order to indicate the moneyness of the options. For instance, for puts and calls corresponding to movements of 10% from the forward rate, $k = 1.1$. We define the risk reversal to be

$$RR(\mathcal{F}k) = P\left(\mathcal{F}k^{-1}\right) - k^{-1}C\left(\mathcal{F}k\right). \tag{5}$$

Risk reversals are the price of one put with strike $\mathcal{F}k^{-1}$ minus k^{-1} calls with strike $\mathcal{F}k$, which is symmetric with respect to the money forward rate $\mathcal{F}$. For instance, in the previous case where $k = 1.1$, the risk reversal is the price of a put protecting against a 10% depreciation of the foreign currency *minus* 0.9 units of a call paying off symmetrically (i.e., if the foreign currency appreciates by 10%).

The next lemma gives the reason for the definition in equation (5): if there is only Gaussian risk, then the risk reversal is exactly 0.

Lemma 2. *If there is no disaster risk, then the risk reversal is exactly zero, for all strikes: $RR(\mathcal{F}k) = 0$ for all $k > 0$.*

On the other hand, if there is disaster risk then the risk reversal is basically the price of an out-of-the-money put (in the previous example, protecting against a 10% depreciation of the foreign currency) minus the price of a symmetric call (e.g., protecting against a 10% appreciation of the foreign currency). Hence, if the foreign country has more crash risk than the home country, its risk reversal is positive.

In the next proposition, we characterize the limit price of risk reversals for strikes in the parametric class $e^{\kappa\sqrt{\tau}}$.

Proposition 3. *In the limit of small time intervals, the price of risk reversals is given by*

$$\lim_{\tau\to 0} RR(\mathcal{F}e^{\kappa\sqrt{\tau}})/\tau = pE\left[(J - J^\star)^+ - (J^\star - J)^+\right] + 2(1 + \Delta^P_{BS}(\kappa))pE\left[(J^\star - J)\right]. \qquad (6)$$

Consider a risk reversal at-the-money forward ($\kappa = 0$) in the case where $J > J^\star$ almost surely. Then, $\Delta^P_{BS}(0) = -1/2$ and $\lim_{\tau\to 0} RR(\mathcal{F}e^{\kappa\sqrt{\tau}})/\tau = 0$. In other words, disaster risk generates non trivial risk reversals only for strikes away from the money.

Risk reversals on the currency pair (home, foreign) essentially capture the relative loadings on disaster risk of the home currency and the foreign currency in the following sense. If the distribution of $J^\star$ decreases in a first-order stochastic dominance sense (i.e., if the foreign currency bears more crash risk), then the value of the risk reversal is weakly higher ($\lim_{\tau\to 0} RR(\mathcal{F}e^{\kappa\sqrt{\tau}})/\tau$ is weakly higher).

We can also consider strikes that do not scale as $\kappa\sqrt{\tau}$ in the limit of short time horizons. If instead the strike is constant at $K > 0$, then the delta of the corresponding put option is equal to -1. In this case, the price of deep-out-of-the money risk reversals is

$$\lim_{\tau\to 0} RR(K)/\tau = pE\left[\left(K^{-1}J - J^\star\right)^+ - \left(K^{-1}J^\star - J\right)^+\right]. \qquad (7)$$

We conclude this section with a proposition linking risk reversals to interest rates.

Proposition 4. *In the domain where the foreign country has more disaster risk than the home country ($J > J^*$), ceteris paribus, the more the foreign country is exposed to disaster risk (the lower is $J^\star$ in the sense of first-order stochastic dominance), the higher are the interest rate differential $r^* - r$ and the short-maturity risk reversal.*

Proposition 4 is natural. Riskier countries should have higher interest rates as we have already seen, and they should have higher prices of put premia because they bear important crash risk: their risk reversals are higher. An analogous proposition naturally holds if the foreign country has less disaster risk than the home country.

3 Estimation

The theoretical results presented in the previous section guide our empirical work on carry trade returns. From a methodological perspective, the model has two main implications: currency excess returns increase with interest rates, and currency options allow the estimation of disaster risk premia. We follow these two insights. Because the forward premium puzzle implies that risk premia are time-varying, we build *portfolios* of currency excess returns by sorting countries based on their interest rates. By doing so, we obtain currency excess returns that are significantly different from zero and capture *expected* excess returns from currency markets. We apply this methodology to unhedged and hedged currency excess returns. As a result, we obtain the empirical counterparts to the expected excess returns described in the previous section. Using the closed-form expressions derived there, we estimate the market compensation for crash risk.

3.1 Data

We first describe our data set and how we build currency portfolios, and then turn to our results on disaster risk premia. We start off with spot, forward, and option contracts on currency markets.

Spot, forward, and currency options. All exchange rates in our sample are in U.S. dollar per foreign currency. As a result, an increase in the exchange rate corresponds to an appreciation of the foreign currency and a decline of the U.S. dollar. For each currency, our sample presents spot and forward exchange rates at the end of the month and implied volatilities from currency options for the same dates. We consider one-month forward rates and options with one-month maturity. Longer-term contracts are available but much less traded. We construct foreign interest rates using forward currency rates and the U.S. LIBOR, assuming that the covered interest rate parity condition holds.[18]

Options are quoted using their Black and Scholes implied volatilities for five different deltas.[19] Our sample comprises deep-out-of-the-money puts (denoted 10 delta puts), out-of-the-money puts

[18] In normal conditions, forward rates satisfy the covered interest rate parity condition (CIP): forward discounts (e.g., the log differences between forward and spot rates) equal the interest rate differentials between two countries. Akram, Rime and Sarno (2008) study high-frequency deviations from CIP and conclude that CIP holds at daily and lower frequencies.

[19] Jorion (1995), Carr and Wu (2007) and Corte, Sarno and Tsiakas (2009) study the features of these currency options.

(25 delta puts), at-the-money puts and calls, out-of-the money calls (25 delta calls) and deep-out-of-the money calls (10 delta calls) for the 1996–2008 period.[20] Figure 3 presents, as an example, the implied volatilities of the currency options in our sample at the end of August 2008. If the underlying risk-neutral distributions of exchange rates were purely log-normal, then these lines would be flat: implied volatilities would not differ across strike prices. This is clearly not the case here. Note for instance that the implied volatility curve is decreasing for Australia or New Zealand (two high–interest rate countries at that time) and increasing for Japan or Switzerland (two low–interest rate countries). These curves signal departures from the normality assumption. Let us take a simple example. A high implied volatility for an out-of-the-money call option implies that the probability of a foreign currency appreciation is higher than in a normal distribution. At the end of August 2008, option prices reflect large probabilities of appreciation for the Japanese yen and Swiss franc as well as large probabilities of depreciation for the Australian and New Zealand dollars. These expected changes actually occurred in the next months.

Using these spot, forward, and option contracts, we now build unhedged and hedged currency excess returns following the definitions presented in Section 2.4.

Portfolios of unhedged and hedged currency excess returns. For each individual currency, we construct the corresponding excess return from the perspective of a U.S. investor. We consider two cases: the investor goes either long or short in the foreign currency. In each case, we build the hedged excess return obtained by buying protection on the option market against an unfavorable change in the foreign currency. When the U.S. investor is long in the foreign currency he buys a put contract, thereby protecting himself against a depreciation of the foreign currency. When he is short, he buys a call contract. Again, the strike price of these options contracts is either far out of the money (at 10 delta), out of the money (at 25 delta), or at the money.

We sort currencies on their forward discounts and allocate them into three portfolios, rebalancing every month. The first portfolio contains the lowest–interest rate currencies while the last portfolio contains the highest–interest rate currencies. By sorting currencies on their risk characteristics, we focus on sources of aggregate risk and average out idiosyncratic variations. When computing portfolio averages, we use equal weights for all currencies. We obtain average currency excess returns, average implied volatilities, and average risk reversals for each portfolio.[21]

[20] By using data from the Chicago Mercantile Exchange, we could have extended the sample to 1986 for three currencies (Canadian dollar, Swiss franc, and yen) and to 1994 for two others (Australian dollar and British pound). Unfortunately CME data do not provide at each date a constant variety of option strikes, which is crucial for our estimation procedure.

[21] Note that the hedge strategy requires buying one option for every currency in the portfolio. In essence, this amounts to buying protection against adverse movements of every currency in the portfolio against the U.S. dollar. Another potentially interesting strategy consists of buying a single option to protect against an adverse movement of

The connection with the theory developed in Section 2 is as follows. The different countries are indexed by $i \in I$. A state variable Ω_t describes the state of the world at date t. This state variable follows an arbitrary stationary stochastic process. All the parameters of the model are arbitrary functions of Ω_t: p, g_i, J_i and $\text{cov}(\varepsilon_i, \varepsilon_j)$. Correspondingly all the computed variables r_i, X_i^e, $X^e(\kappa)_i$, π_i^D and π_i^G depend on Ω_t. Underlying our three portfolios are the three state-dependent sets $I_1(\Omega_t)$, $I_2(\Omega_t)$, and $I_3(\Omega_t)$.

High interest rates r_i can be due to high values of g^i or to low values of $pE[J_i - 1]$. If disaster risk is an important determinant of cross-country variations in interest rates, then a portfolio formed by selecting countries with high–interest rates will, on average, select countries that feature high disaster risk, $-E[J_i]$. The empirical analysis that follows indeed confirms that.

Sample. Our data set comes from JP Morgan. It contains 32 currencies: Argentina, Australia, Brazil, Canada, Chile, China, Columbia, Czech Republic, Denmark, Euro area, Hong Kong, India, Indonesia, Israel, Japan, Malaysia, Mexico, New Zealand, Norway, Peru, Philippines, Poland, Singapore, South Africa, South Korea, Sweden, Switzerland, Taiwan, Thailand, Turkey, United Kingdom, and Venezuela. Following the World Economic Outlook (IMF, 2008) classification, we split the sample between advanced countries and emerging countries.[22]

There are two main reasons to focus on advanced countries: the higher liquidity of their option markets and the normality of their returns. We focus here on normality tests and investigate later the impact of transaction costs.

Our model implies that, as long as a currency crash does not occur in sample, changes in exchange rate are conditionally normally distributed. We check this implication in our data, limiting first our attention to the 1/1996 – 8/2008 period. We exclude the last four months of our sample because, during the fall of 2008, high–interest rate currencies depreciated and low–interest rate currencies appreciated sharply. Carry trades thus paid very badly in the fall of 2008, when stock markets tumbled worldwide and liquidity dried up. We take the view that this period represents an example of disasters in our sample and will pay special attention to this particular period in the next section. For now, we exclude it from our sample.

Table 9 in Appendix C reports higher moments of changes in exchange rates along with the standard Jarque and Bera (1980) and Lilliefors (1967) normality tests for each currency available over this period. The left panel focuses on advanced countries. Bootstrapping the skewness and kurtosis statistics, we find that the sample values are not significantly different from the Gaussian

the basket of currencies in this portfolio. However, we do not have data on basket options and so we do not pursue that route.

[22] The Word Economic Outlook classification combines three criteria: (i) per capita GDP, (ii) export diversification, and (iii) integration into the global financial system.

ones for all countries, except for South Korea and Singapore. The Lilliefors test leads to the same conclusion. The Jarque–Bera test rejects normality more often (adding the United Kingdom and Japan to the list), but the test is known to over-reject in short samples. The comparison with the right panel, which focuses on emerging countries, is striking. There, most exchange rate distributions differ from normality. Most rejections come from high kurtosis.[23] If we include fall 2008 in our sample, the recent large changes in exchange rates lead to rejection of the normal distribution even for many advanced countries.

Our model implies that *conditional* changes in exchange rates are normal. Yet the normality tests reported so far are unconditional, and exchange rates tend to exhibit time-varying volatility. To take into account such heteroscedasticity, we estimate a GARCH (1, 1) model for each currency. We then run normality tests on exchange rate changes normalized by their volatility. To save space, we report results in Table 10 in Appendix C. After the GARCH (1, 1) correction, all advanced countries, except South Korea, exhibit conditionally Gaussian exchange rates in our sample. Most emerging countries, however, still fail normality tests.

As a result, we focus here on our sample of advanced countries (excluding South Korea) over the 1/1996–8/2008 period.[24] We turn now to our main empirical results. Note that results obtained with the whole sample of advanced and emerging countries are reported in Appendix C as robustness checks. In that appendix we also consider a smaller sample of the nine most advanced countries as in Jurek (2008).

3.2 Results

We first present the key characteristics of our currency portfolios and then focus on measures of disaster risk premia.

Portfolio characteristics. Forming portfolios is a way to compute moments conditional on the three sets I_1, I_2, and I_3. Of particular interest to us will be three of these moments: carry trade returns and the corresponding disaster and Gaussian risk premia. For instance, the expected return on portfolio k is simply the average return over the countries in the portfolio:

$$\overline{X}_k^e = E\left[\frac{\sum_{i \in I_k(\Omega_t)} X_i^e(\Omega_t)}{\# I_k(\Omega_t)}\right],$$

[23] We also report, in Appendix C, higher moments and normality tests for our portfolios of currency excess returns. In our benchmark sample of advanced countries, the Lilliefors test cannot reject the normality assumption for any of our portfolios. In our large sample of advanced and emerging countries, however, the high–interest rate portfolios exhibit fat tails and thus clearly depart from normality.

[24] Our sample consists of Canada, Czech Republic, Denmark, Euro area, Israel, Japan, New Zealand, Norway, Poland, Singapore, Sweden, Switzerland, Thailand, and United Kingdom

where I_k denotes the set of currencies in portfolio k. Similarly, the expected hedged return on portfolio k is:

$$\overline{X}^e_k(\kappa) = E\left[\frac{\sum_{i\in I_k(\Omega_t)} X^e_i(\Omega_t)(\kappa)}{\# I_k(\Omega_t)}\right].$$

Table 1 reports average currency excess returns that are either unhedged, hedged at 10 delta, hedged at 25 delta, or hedged at the money. Average currency excess returns increase monotonically from the first to the last portfolio. This is not a surprise: we know from the empirical literature on the uncovered interest rate parity that high–interest rate currencies tend to appreciate on average. As a result, investors in these currencies gain both the interest rate differential and the foreign exchange rate appreciation. Hedging downside risks decreases average returns. A hedge at 10 delta protects the investor against large drops in foreign currencies, whereas a hedge at the money protects the investor against any depreciation of the foreign currency: the latter insurance is obviously more expensive because it covers more states of nature and thus leads to lower excess returns.

For each portfolio, we also report in Table 2 the average implied volatility at different strikes. One result stands out: the average implied volatility of high–interest rate currencies (e.g., portfolio 3) is much higher for out-of-the-money put options than for other strikes and other portfolios. Option markets price a large depreciation risk for high–interest rate currencies. The same insight is apparent in risk reversals.

The last panel of Table 2 presents average risk reversals at delta 10 and 25 delta:

$$\overline{RR}_k = E\left[\frac{\sum_{i\in I_k(\Omega_t)} RR_i(\Omega_t)}{\# I_k(\Omega_t)}\right].$$

Recall that risk reversals correspond to positions that are long put and short call options. As a result, higher levels of risk reversals indicate higher probabilities of depreciation for the foreign currency. We report risk reversals quoted in terms of implied volatilities. As in the model, risk reversals increase monotonically with interest rates. Higher–interest rate currencies have higher probabilities of depreciation. This result is in line with the premises of our model, which introduces the risk of large-scale depreciation in currency markets.

The strong link between interest rates and risk reversals suggests a comparable sorting that uses risk reversals instead of interest rates. Underlying this construction are three different portfolio sets with their corresponding conditional moments. Here again we obtain a monotonically increasing cross-section of excess returns. Table 3 reports hedged and unhedged average excess returns. Countries with higher levels of risk reversals tend to offer higher currency returns on average. The difference in unhedged returns between the last and first portfolio is lower than in our previous portfolios, but it is still significant.

We now turn to the direct estimation of the market's compensation for bearing disaster risk.

Disaster risk premia. In order to estimate disaster risk premia, we focus on a zero-investment strategy that goes long on high–interest rate currencies and short on low–interest rate currencies. This strategy corresponds to usual currency carry trades.

The expected return of the carry trade is $\overline{X}^e = \overline{X}_3^e - \overline{X}_1^e$. It can be decomposed as the sum of a disaster risk premium $\overline{\pi}^D$ and a Gaussian risk premium $\overline{\pi}^G$. The disaster risk premium is the difference between the average disaster risk premium in portfolio 3 and the average disaster risk premium in portfolio 1:

$$\overline{\pi}^D = E\left[\frac{\sum_{i \in I_3(\Omega_t)} \pi_i^D(\Omega_t)}{\# I_3(\Omega_t)}\right] - E\left[\frac{\sum_{i \in I_1(\Omega_t)} \pi_i^D(\Omega_t)}{\# I_1(\Omega_t)}\right].$$

Similarly, the Gaussian risk premium is the difference between the average Gaussian risk premium in portfolio 3 and the average Gaussian risk premium in portfolio 1:

$$\overline{\pi}^G = E\left[\frac{\sum_{i \in I_3(\Omega_t)} \pi_i^G(\Omega_t)}{\# I_3(\Omega_t)}\right] - E\left[\frac{\sum_{i \in I_1(\Omega_t)} \pi_i^G(\Omega_t)}{\# I_1(\Omega_t)}\right].$$

The average unhedged return of this strategy is equal to 6.5% per year in our sample. It corresponds to the sum of the average return on the third portfolio in the left panel of Table 1 (when the investor is long on the foreign currency) and the first portfolio in the right panel (when the investor is short on the foreign currency). We also report hedged carry trades at 10 delta, 25 delta, and at-the-money (ATM). They correspond to $\overline{X}^e(\kappa) = \overline{X}_3^e(\kappa) - \overline{X}_1^e(\kappa)$. The first panel of Table 4 presents these average carry excess returns and their standard errors. The latter are obtained by bootstrapping the monthly excess returns under the assumption that they are independent and identically distributed (i.i.d.). As a result, these standard errors take into account the short sample size. Carry excess returns that are either unhedged or hedged at 10 delta and 25 delta are statistically different from zero. Carry returns hedged at the money are positive but not significant. The differences between unhedged and hedged returns are all positive and significant.

The second panel of Table 4 reports structural estimates of the disaster risk component ($\overline{\pi}^D$) and the Gaussian risk component ($\overline{\pi}^G$). We start with simple estimates that only require computing averages, and then we turn to GMM estimates.

As derived in the previous section, unhedged excess returns correspond to the sum of $\overline{\pi}^D$ and $\overline{\pi}^G$. Hedged excess returns are approximately equal to $\overline{\pi}^G$ multiplied by a correction factor related to the delta of the option. To estimate $\overline{\pi}^D$ and $\overline{\pi}^G$, we first correct each average hedged return for its delta component:

$$\widetilde{X^e(\kappa)} = \overline{X}^e(\kappa)/(1 + \Delta_\kappa),$$

where $\overline{X}^e(\kappa)$ corresponds to the average carry return hedged at delta κ ($\kappa = 10$, 25, or at-the-money) and Δ_κ denotes the option delta (respectively equal to -0.1, -0.25, and -0.5). Section 2.5 shows that the expected value of each $\widetilde{X^e(\kappa)}$ is simply $\overline{\pi}^G$. So, we form our estimate of the Gaussian risk premium as a simple weighted average of the delta-corrected hedged carry trade returns:[25]

$$\widehat{\pi}^G = \frac{\sum_{\kappa \in I} \widetilde{X^e(\kappa)}}{N}, \tag{8}$$

where N is the number of hedged excess returns considered. For instance, $N = 1$ when we use ATM options only and $N = 3$ when we use 10 delta, 25 delta and ATM options.

As warranted by the analysis in Section 2.5, our estimate of the disaster risk premium is the average unhedged carry trade return, $\overline{X}^e$, minus the estimate of the Gaussian premium:

$$\widehat{\pi}^D = \overline{X}^e - \widehat{\pi}^G. \tag{9}$$

We report four sets of estimates obtained using the methodology just described and four different sets of hedged returns: 10 delta (first column), 25 delta (second column), at-the-money (third column) hedged returns along with the previous three hedged returns combined together (fourth column). Note that we estimate two risk premia, $\overline{\pi}^D$ and $\overline{\pi}^G$, using either two (first, second, and third columns) or four moments (fourth column). Again, standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Depending on the specification, Gaussian risk premia range from 3.4% to 5.3%; disaster risk premia amount to 1.2% to 3.1% annually. The latter account for approximately 20% to 50% of the average carry trade returns in our sample. The lower estimate is obtained when using only deep-out-of-the-money options. Disaster risk premia are significantly different from zero in all cases, except when using solely at-the-money options.

Our previous estimates of disaster risk premia, obtained with simple averages, correspond to the minimization of the sum of squared differences between empirical and theoretical excess returns. We now turn to Hansen's (1982) GMM estimates of disaster risk premia. We use all the available unhedged and hedged excess returns and thus have four moments to estimate two parameters. The other cases reported previously are just-identified with two moments to determine two parameters.[26]

[25] This estimate corresponds to the minimization of:

$$(\overline{X}^e - \overline{\pi}^D - \overline{\pi}^G)^2 + \sum_{\kappa \in I} (\widetilde{X^e(\kappa)} - \overline{\pi}^G)^2.$$

[26] This estimate corresponds to the minimization of $g_T' W^{-1} g_T$, where W is the variance–covariance matrix of all hedged and unhedged returns and where g_T describes all moment conditions: $g_T = [(\overline{X}^e - \overline{\pi}^D - \overline{\pi}^G), (\widetilde{X^e(\kappa_1)} - \overline{\pi}^G), ..., (\widetilde{X^e(\kappa_3)} - \overline{\pi}^G)]$. If $W^{-1} = A'A$ then the estimate minimizes $g_T' A' A g_T$; this corresponds to the "square" of

In order to weight the different moments, we use the covariance matrix of all hedged and unhedged returns. We do not use a spectral density matrix because of the short length of our sample. We obtain a disaster risk premium of 1% (with a standard error of 0.4) and a Gaussian risk premium of 4.8% (with a standard error of 1.9). The disaster risk premium obtained with all hedged returns is close to the one obtained with 10-delta returns. This happens because the standard deviation of delta-corrected ATM hedged returns is much higher than the other ones. As a result, the GMM estimation underweights this moment, which previously delivered the higher estimate of disaster risk premia. This procedure thus gives a lower bound on disaster risk premia. Note also that the GMM estimation does not impose the condition that unhedged excess returns are the sum of disaster and Gaussian risk premia.

We check our results on different portfolios that feature either different sorts or different countries. We obtain similar results on portfolios of currency excess returns sorted on risk reversals. Recall that these portfolios deliver a monotonic cross-section of returns and offer a carry excess return of 3.2% annually. Table 5 reports estimates of the corresponding Gaussian and disaster risk premia. The former varies from 1.3% to 1.7%, and the latter ranges from to 1.4% to 1.9%. Again, all estimates except the one using solely at-the-money options, are statistically significant. Disaster risk premia account for approximately 40%–60% percent of the long–short returns on these risk reversal–based portfolios.

As robustness checks, we consider two additional samples: either all the developed and emerging countries in our data set or a subset of nine developed countries (Australia, Canada, Euro area, Japan, New Zealand, Norway, Sweden, Switzerland, and United Kingdom). To save space, we report all tables in Appendix C.[27] We obtain very similar estimates on the small sample of nine developed countries as before on our larger sample of advanced countries. Using GMM, we obtain a disaster risk premium of 1.1%, which accounts for 25% of the carry trade returns. We obtain somehow lower disaster risk premia on our large sample of advanced and emerging countries. Emerging markets, however, present lower liquidity and higher bid–ask spreads as we have seen; moreover most fail normality tests. Taking transaction costs into account helps reconcile the results obtained on both samples.

We view these estimates of disaster risk premia as the main empirical contribution of this paper because they are derived within a theoretical framework that allows us to incorporate a variety of options. We draw two clear conclusions from this experiment. First, disaster risk is priced on

linear combinations of our original moments. As a result, the minimization does not imply that $\overline{X}^e = \overline{\pi}^D + \overline{\pi}^G$. The J-statistic is equal to $g_T var(g_T)^{-1} g_T \sim \chi^2(\#moments - \#parameters)$; cf Cochrane (2005).

[27] Table 13 presents disaster risk premia for the nine developed countries. Table 14 reports average currency excess returns across portfolios when we sort developed and emerging countries based on interest rates. Table 15 presents implied volatilities and risk reversals for developed and emerging countries. Table 16 reports estimates of disaster risk premia in the same sample.

currency markets. Second, there are significant differences in the amounts of disaster risk across countries. If all countries bore the same amount of disaster risk, then it would cancel out in our long–short excess returns.

The estimate of disaster risk premia $\overline{\pi}^D$ is higher when using at-the-money options than out-of-the-money options. In light of the model, out-of-the-money options seem "too cheap" compared to at-the-money options. Note, however, that differences in disaster risk premia across these options are not statistically significant. Take for example the GMM estimate as a benchmark. Then the other estimates, obtained using simple averages, differ by 0.15, 0.62, 2.09, and 0.95 percentage points (see Table 4). But the corresponding standard errors on these differences are 0.57, 0.96, 1.72, and 1.01 percentage points. Therefore, the estimates of disaster premia are not statistically different across strikes. With this caveat in mind, we turn to potential explanations for these different point estimates. We see three possible explanations: illiquidity, counterparty risk, and model misspecification.

The illiquidity explanation runs as follows. The JP Morgan market maker simply gives indicative prices by using the Black–Scholes formula (which generates a low option price), but there is little trading of out-of-the-money options. If someone wanted to aggressively buy these options, then she would end up moving prices against herself and paying higher prices. So the potential trading prices are higher than the indicative prices we have in our data.

In the counterparty risk explanation, the seller of a put might actually default during a disaster. Put premia take that risk into account and are lower than in the model. This issue, of course, affects not only currency options but also stock options, credit default swaps, and the like. We expand on this issue in Section 3.4.

Finally, the model may simply be misspecified. The model might generate too small a risk-neutral probability for small depreciations. One way to incorporate this possibility in our model would be to allow for two kinds of disasters: large disasters and small disasters. In such a specification, out-of-the-money options offer no protection against small disasters and would therefore be cheaper than at-the-money options.

We do not attempt to enrich the model to capture liquidity and counterparty risks or small disasters, leaving this for future research. In this paper, we focus on the simplest model that is not rejected by the data. We can formally test if the model is rejected with our GMM estimation. Following Hansen (1982), we compute the J-test of the model's pricing errors. This statistic is distributed as a chi-square with two degrees of freedom. The J-statistic is 2.51, leading to a p-value of 0.28. The model is thus not rejected in our sample.

3.3 Transaction Costs

So far, our estimates of disaster risk premia do not take into account bid–ask spreads on currency markets. Transaction costs on forward and spot contracts would reduce unhedged excess returns. Transaction costs on currency options would increase insurance costs against disasters. As a result, these costs would increase the share of disaster risk premia. In this respect, the numbers previously reported in this paper constitute a lower bound.

Bid and ask spreads are not available in the JP Morgan dataset. For the spot and forward markets, we rely on Reuters daily quotes available on Datastream. Measured in our sample, these quotes imply average spreads (divided by the mid rate) of 9 basis points for forwards and 8 basis points for spot rates. When implementing carry trades through forward markets, investors who go long on high–interest rate currencies buy forward contracts at the ask price. When they receive the corresponding foreign currencies at the end of the contract, they convert their proceeds back into U.S. dollars at the bid price. As a result, they incur half the bid–ask spread on both the forward and spot contracts. Assuming a spread of 8 basis points and 12 trades per year, the annual cost is equal to about 100 basis points or 1%. Gilmore and Hayashi (2008) argue that such spreads overstate transaction costs on currency markets because investors might roll over their positions each month instead of closing them to re-open them the next day. With an example based on the South African rand, they show that forward markets imply an annual carry cost of 192 basis points whereas rolling over positions would cost only 13 basis points i.e., 15 times less; cf. Appendix 2 of their paper. This estimate, however, assumes that a given currency remains in the carry portfolio for five years, and thus it underestimates the costs due to portfolio rebalancing. As a result, we assume that the average actual transaction costs on our unhedged carry portfolio are in between these two estimates. We take an annual value of 0.25% for advanced countries and 2% for emerging countries.

We should like to assess transaction costs on currency option markets but unfortunately we do not have access to time–series of bid–ask spreads on these markets. To obtain an order of magnitude, we collected bid–ask spreads on November 10, 2008 and January 20, 2009 for different currency pairs.[28] Table 12 in Appendix C presents these bid–ask spreads on currency options quoted in terms of implied volatilities. Because of the global financial crisis, implied volatilities are much higher than in the rest of our sample. For most currency pairs, implied volatilities in November 2008 are more than twice their sample means. According to market participants, bid–ask spreads in November 2008 were also much higher than in our sample. These spreads reached 30% of the underlying midpoint (mean of bid and ask) values for out-of-the-money options on emerging market currencies. Bid–ask spreads are much tighter for the currencies of the most advanced countries.

[28]We thank the Bank of France for sharing these data with us.

In January 2009, most implied volatilities were lower but spreads remained around 10%. According to market participants, these spreads are abnormally large. To estimate the impact of transaction costs on our results, we assume bid–ask spreads of 5% for advanced countries and 10% for the others. As a result, spreads widen when implied volatilities increase, but not fully to the levels observed during fall 2008. We convert these implied volatilities spreads into bid–ask prices and then re-estimate hedged excess returns.

We test the robustness of our results to the inclusion of these transaction costs. As expected, transaction costs increase the share of disaster risk; the results are reported in Table 6. Using simple averages, Gaussian risk premia now range from 1.6% to 4.7%. Disaster risk premia also range from 1.6% to 4.7% annually, accounting for approximately 25%–70% of the average carry trade in our sample. Disaster risk premia are significantly different from zero. Using GMM, we obtain a disaster risk premium of 1.3% It is three standard errors away from zero and represents one fourth of the carry trade excess returns. We consider this value as our best estimate of the compensation for disaster risk considering the data available. It is, however, a lower bound because it does not take into account default probabilities on option markets.

3.4 Counterparty Risk

So far we have assumed that there is no counterparty risk for options. However, it is reasonable to think that the seller of a put might default with some probability ϕ if a disaster occurs. In that case, an agent engaging in hedged carry trade still bears some disaster risk. Indeed, the expected excess return of the hedged carry trade is then:

$$X^e_{\text{hedged}} = (1+\Delta)\pi^G + \phi\pi^D.$$

Since with probability ϕ the agent is exposed to disasters, the compensation for the disaster risk is then $\phi\pi^D$. Our estimation procedure to uncover disaster risk premia must now be amended as follows:

$$\pi^D = \frac{\overline{X}^e - \overline{X}^e(\kappa)/(1+\Delta_\kappa)}{1-\phi/(1+\Delta_\kappa)}. \tag{10}$$

For instance, take the case of deep-out-of-the-money options ($\Delta = -0.1$). Equation (10) shows that the estimate of π^D that does not take into account counterparty risk must now be multiplied by approximately $1/(1-1.1\phi)$. When $\phi = 0.1$, π^D is multiplied by 1.12; when $\phi = 0.25$, it is multiplied by 1.38. For ATM options ($\Delta = -0.5$) the adjustment is even larger: when $\phi = 0.1$, π^D is multiplied by 1.25; when $\phi = 0.25$, it is multiplied by 2.

This section demonstrates that counterparty risk can substantially increase our estimate of disaster risk premia. However, we lack data to pin down default probabilities on option markets. As a result, our estimate of disaster risk premia should be considered as a lower bound. One approach to estimate default probabilities could be to use information from the credit default swap or corporate bond markets, but this is beyond the scope of this paper and we leave it for further research. Instead, we now compare our estimate of disaster risk premia to the macroeconomic literature on disasters, starting with a case study of fall 2008.

3.5 Fall 2008 and Comparison with Barro and Ursua (2008)

We view this recent period as the unique example of disaster in our data. As noted earlier, its inclusion in our sample is enough to reject the normality assumption for many countries. In this section, we provide a brief description of what happened in currency markets. Both spot and option markets support the characterization of this period as a financial disaster.

Fall 2008. In our sample, fall 2008 stands out as the worst time for carry traders. This is obvious for specific currencies, but it also holds for currency portfolio returns. We start with a simple example using two bilateral exchange rates; in the recent period, the New Zealand dollar has been a high–interest rate currency while the Japanese yen has been a low–interest rate one. Figure 4 plots monthly changes in these exchange rates against the U.S. dollar. We start our graph at the beginning of the subprime crisis on financial markets; the sample period is thus 7/2007 – 12/2008. Clearly, the Japanese yen appreciated and the New Zealand dollar depreciated during that period, with both movements hurting carry traders. The same figure also reports the return index on a carry trade strategy that borrows in yen to invest in the New Zealand dollar. The index starts at 100 in July 2007. At the end of December 2008, the index is slightly above 60, and most of the losses have occurred in the last four months of the sample. These losses are not specific to the New Zealand dollar–Japanese yen pair; we obtain similar results with our baskets of currencies. The average return of our carry trade strategy was *minus* 4.5% in the fall 2008, for a cumulative decline from September to December that amounts to 17.8%. This is a large drop, as the standard deviation of monthly returns over the whole sample is just 2%. Almost all of the 17.8% decline is due to losses on high–interest rate currencies, which depreciated sharply.

Similar conclusions arise in the case of currency options. Large changes in exchange rates triggered exercise of currency options embedded in our portfolios. Figure 5 plots the frequency of call and put options exercised on currencies allocated in the first and last portfolios, respectively. At each moment in time, the frequency is obtained as the number of options exercised divided by the number of currencies in the portfolio at that time. Recall that the first portfolio contains low–interest rate currencies and thus funding currencies. Investors want to buy call options to insure themselves against large appreciation of such currencies. The last portfolio contains high–interest rate currencies. There, investors consider put options. The figure shows clearly that the frequency of 10-delta put options exercised reaches an all-time high in the fall of 2008. The proportion of call options triggered was also high, but not at its maximum value in the sample.

These very low returns on currency markets occurred in bad times for U.S. investors. During fall 2008, the U.S. stock market declined by 33% in terms of the MSCI index.[29] Figure 6 compares equity and currency excess returns over our sample. The correlation between these excess returns is particularly high, reaching 0.7 since the start of the subprime mortgage crisis in July 2007.

Standard risk measures beyond those from equity markets point in the same direction in our sample: the equity option–implied volatility index VIX, its bond equivalent MOVE, and credit spreads were at their all-time high in the fall of 2008. Figure 7 presents all these variables in a standardized way: currency returns and risk measures are all de-meaned and then divided by their standard deviations. The events of fall 2008 represent up to five standard deviations in these series. Very low currency excess returns (four standard deviations below their means) happened exactly when volatilities in equity and bond markets and credit spreads were high (four standard deviations above their means) — that is, in bad times. Our sample in this paper is short, but our findings are in line with the literature. As Lustig et al. (2008) show, carry trades tend to pay poorly during times of crises, exactly when stock markets tank. This high correlation between stock and currency markets also occurred during the 1987 stock market crash and during the Mexican, Asian, and Russian crises. These market-based indices offer real-time measures of risk that complement the approach based on marginal utilities and real consumption growth rates. Figure 8 focuses on consumption growth, and the same conclusion emerges here. Preliminary estimates of U.S. national account statistics point toward an annualized decrease of 4.3% in real personal consumption expenditures in the fourth quarter of 2008, following an annualized decrease of 3.8% in the third quarter. These shocks represent declines of more than three standard deviations in the mean consumption growth rate. As reported in Lustig and Verdelhan (2007) on an earlier sample, low carry trade excess returns tend to occur in times of low consumption growth.

Finally, note that the link between risk reversals and subsequent currency appreciations differ during crisis and normal times. In normal times, according to the model, high levels of risk reversals should predict foreign currency appreciations. Using actual data, however, we did not find much significant predictability though. During times of crisis, high risk reversals should predict foreign currency depreciation. This is what happened during the fall of 2008: foreign currency depreciation seemed to follow high risk reversals. This behavior is line with the model if we interpret the fall of

[29] The closest event to this very strong decline in equity and currency returns is the 1987 stock market crash. From September to November 1987, the U.S. stock market lost 32.6%. This period is not in our sample because we do not have currency option data before January 1996.

2008 as a disaster. The evidence is, of course, very limited because we have only one disaster in our sample. As a consequence, we do not attempt to quantify this point and instead simply present, in Figure 9, exchange rate appreciations and risk reversals for each month and each currency in the fall of 2008.

According to many markets and risk factors, the fall of 2008 constitutes a disaster. We use this example to connect our findings with the previous macroeconomic literature on disasters.

Comparison with Barro and Ursua (2008). In a disaster, the stochastic discount factor is multiplied by an amount J. To relate this J to more primitive economic quantities, we use the model of Farhi and Gabaix (2008). In that model, $J = B^{-\gamma}F$ where $B^{-\gamma}$ is the growth of real marginal utility during a disaster and F is the growth of the value of one unit of the local currency in terms of international goods during the same disaster. Hence $\overline{\pi}^D = \overline{pE[J]}_1 - \overline{pE[J]}_3 = \overline{pE[B^{-\gamma}(F)]}_1 - \overline{pE[B^{-\gamma}(F)]}_3$. Therefore, the disaster risk premium depends on the probability of disasters p, the relative value of the SDF $B^{-\gamma}$, and the payoff of the carry trade in disasters through the sufficient statistic $\overline{pE[B^{-\gamma}(F)]}_1 - \overline{pE[B^{-\gamma}(F)]}_3$. Using the episode of fall 2008 to calibrate the value of $\overline{F}_1 - \overline{F}_3$ and assuming away a potential correlation between $B^{-\gamma}$ and $\overline{F}_1 - \overline{F}_3$, we can shed some light on the typical value of $pB^{-\gamma}$. This exercise should be viewed as a back-of-the-envelope calculation rather than a rigorous estimate, since our inference of $\overline{F}_1 - \overline{F}_3$ relies on a single disaster that is still unfolding as this paper is written. Thus we cannot observe the full path to recovery and, as Gourio (2008) shows, we might overestimate the impact of disasters. With this caveat in mind, if we retain a value for $\overline{F}_1 - \overline{F}_3$ of 20% then a value for $\overline{pE[B^{-\gamma}]}$ of 6.5% is necessary to generate a disaster risk premium $\overline{\pi}^D$ of the order of magnitude that we estimate in the data.

We compare this value to Barro and Ursua's (2008b) estimates. These authors use long samples of consumption series for a large set of countries.[30] Their findings are broadly consistent with the estimates from Barro (2006), which are based on GDP disasters. Barro and Ursua (2008b) estimate a probability of disasters p equal to 3.63%. A coefficient of relative risk aversion $\gamma = 3.5$ then implies that $\overline{E[B^{-\gamma}]} = 3.88$, leading to a value of $\overline{pE[B^{-\gamma}]}$ equal to 14%. The authors show that these values can rationalize the equity premium.

Using a value of 14% for $\overline{pE[B^{-\gamma}]}$ and a value of 20% for $\overline{F}_1 - \overline{F}_3$ leads to a disaster risk premium of $0.14 \times 0.2 = 2.8\%$, which is higher than but still comparable to our point estimate. Therefore, we view our estimates as being broadly consistent with Barro and Ursua (2008b)'s findings. We end this paper with a review of the link between volatility smiles, risk reversals, and exchange rates.

[30] Note, however, that interpreting our pricing kernel strictly as a simple function of consumption growth would open a large debate that is beyond the scope of this paper. Constant relative risk aversion and complete markets imply, for example, a very high correlation between consumption growth and exchange rates, a high correlation that is not evident in the data (Backus and Smith, 1993).

3.6 Volatility Smiles, Risk Reversals, and Exchange Rates

We first provide a simple calibration of the model in order to check that it simultaneously accounts for the volatility smile observed in the data and the disaster risk premium that we have estimated. We then test the contemporaneous relationship between risk reversals and exchange rates, and the predictive content of risk reversals for currencies.

Accounting for the smile. In this section we examine the implications of our model for the volatility smile – that is, the relationship between the implied volatility and the strike of currency options. The exact value of a put with strike K is given by

$$\begin{aligned} P_{t,t+\tau}(K) &= (1-p\tau)\, e^{-g^*\tau} V_{BS}^{P}\left(1, K e^{-(g-g^*)\tau}, \sigma\sqrt{\tau}\right) \\ &+ p\tau e^{-g^*\tau} E\left[J^* V_{BS}^{P}\left(1, K e^{-(g-g^*)\tau} J/J^*, \sigma_{t,t+\tau}\sqrt{\tau}\right)\right], \end{aligned}$$

where $\sigma_{t,t+\tau} = \sqrt{var(\varepsilon - \varepsilon^*)}$ and the expectation operator E is over the joint distribution of J and J^*.

The implied volatility $\widehat{\sigma}_{t,t+\tau}$ is computed by solving the following implicit equation:

$$P_{t,t+\tau}(K) = e^{-r^*\tau} V_{BS}^{P}\left(1, K e^{-(r-r^*)\tau}, \widehat{\sigma}_{t,t+\tau}\sqrt{\tau}\right),$$

where $r = g - \log(1 + p\tau E[J-1])/\tau$ and $r^* = g^* - \log(1 + p\tau E[J^*-1])/\tau$. Recall that when quoting options, traders routinely use the delta of the underlying option rather than its strike, which is a conventional quantity computed as

$$-e^{-r^*\tau}\mathbb{N}\left(\frac{\log(K) - \left(r - r^* + \widehat{\sigma}_{t,t+\tau}^2/2\right)\tau}{\widehat{\sigma}_{t,t+\tau}\sqrt{\tau}}\right).$$

Note that this quantity might differ from the true sensitivity of the option with respect to the fundamental.

All our currency options are options on exchange rates against the U.S dollar. It is therefore most natural to attempt to calibrate our model to fit the average volatility smile of a given portfolio. We choose to focus on portfolio 3 which represents a carry trade where the funding currency is the U.S. dollar. To calibrate the model, we choose the parameters as follows. We take J and J^* to be deterministic. We assume that the values of p and J for the United States are consistent with the estimation of Barro and Ursua: $J = B^{-\gamma} = 3.88$ and $p = 3.63\%$. We choose J^* to match a value of $\pi^D = 1.6\%$, a number that is roughly in the middle of our range of estimates. We shall investigate the sensitivity of the calibration to the exact value retained for π^D, which implies that

$J^* = J\left(1 - \pi^D / \left(pB^{-\gamma}\right)\right) = 3.44$. We choose the physical volatility of the exchange rate to match an implied volatility at the money in portfolio 3 of 10%. This leads us to pick $\sigma_{t,t+\tau} = 9.6\%$. We pick $g = 13.4$ and $g^* = 14.6\%$ in order to match the average U.S. interest rate ($r = 3\%$) and the average interest rate in portfolio 3 ($r^* = 5.8\%$) over the sample.

The resulting implied volatilities as a function of the delta of the option in this calibration are as follows. For a 10-delta put, the implied volatility is 11.4%. For a 25-delta put, the implied volatility is 10.4%. At the money, the implied volatility is 10.0%. For a 25-delta call, the implied volatility is 9.9%. Finally, for a 10-delta call, the implied volatility is 9.8%.[31]

These values should be compared with the implied volatilities for portfolio 3 in the data. For a 10-delta put, the implied volatility is 11.5%. For a 25-delta put, the implied volatility is 10.6%. At the money, the implied volatility is 10.0%. For a 25-delta call, the implied volatility is 10.0%. Finally, for a 10-delta call, the implied volatility is 10.4%. The overall fit of our model is quite good. It is better for out-of-the-money puts than for out-of-the-money calls. Note, however, that we obtain these values by assuming constant J and $J^\star$. The fit could be further improved by choosing an appropriate probability distribution for J and J^*.

Risk reversals and exchange rates. The model implies that (i) increases in risk reversals are associated with *contemporaneous* exchange rate depreciations, and (ii) high levels of risk reversals *predict* future currency returns. We test these predictions both on panel data and on portfolio series.

In order to test for the first prediction, we first regress monthly changes in nominal exchange rates on monthly changes in risk reversals. We use risk reversals measured in prices at 10 and 25 deltas. Because these deltas imply different deviations from forward rates across countries, we also check our findings on risk reversals that are normalized: these risk reversals correspond to strikes that are 5% or 10% away from forward rates. We de-mean both the regressor and the dependant variable so as to remove the central role played by the U.S. dollar. The results on portfolios are reported in Table 7. Tables 18 and 19 in Appendix C report panel results for advanced economies and the whole sample, respectively. All panel specifications include currency fixed effects, and standard errors are obtained by bootstrap. We find a highly robust negative correlation between changes in risk reversals and changes in exchange rates. This negative relationship is robust to alternative

[31] Following the same calibration procedure but using a value of 2% for π^D leads to the following implied volatilities. For a 10-delta put, the implied volatility is 12.1%. For a 25-delta put, the implied volatility is 10.6%. At the money, the implied volatility is 10.0%. For a 25-delta call, the implied volatility is 9.9%. Finally, for a 10-delta call, the implied volatility is 9.8%. We also report the implied volatilities when the retained value for π^D is 1%. For a 10-delta put, the implied volatility is 10.5%. For a 25-delta put, the implied volatility is 10.2%. At the money, the implied volatility is 10.0%. For a 25-delta call, the implied volatility is 10.0%. Finally, for a 10-delta call, the implied volatility is 9.9%.

risk-reversal measures and to controlling for the effect of the dollar.[32] With portfolios and risk reversals at either 10 or 25 deltas, R^2 values range from 25% to 40%. In our panel estimates using demeaned country-level exchange rates, R^2 values are close to 5%. In both cases, risk reversals are statistically significant. Their effect is also economically significant: a one-standard-deviation change in risk reversals is associated with a 1% to 2.3% variation in exchange rates, which is slightly below the monthly standard deviation of nominal exchange rate changes (2.8%).

In order to test for the second prediction, we augment standard UIP regressions with risk reversals. Equivalent regressions are run against excess returns instead of changes in exchange rates. The null hypothesis of UIP is a coefficient of 1 for the interest differential (defined as the difference between domestic and foreign interest rate in the specification with exchange rate change) and a coefficient of 0 in the specification with excess returns. We recover the usual negative coefficient on the interest rate differential. Adding risk reversals to the usual UIP regressions does not improve one-month-ahead exchange rate forecasts, and no risk reversal significantly predicts currency excess returns or changes in nominal exchange rates in panel data, as shown in Table 8. To save space, we report equivalent panel results in Tables 20 and 21 in Appendix C. Currency portfolios suggest a clear positive relationship between average currency excess returns and average risk reversals across portfolios. As previously noted, the last panel of Table 2 reports an increase in average risk reversals from the first to the last portfolio. Equivalent results are obtained for other measures of risk reversals and for the whole sample of advanced and emerging countries. However, within portfolios, there is no one-month-ahead predictability of risk reversals on currency excess returns; this is shown in Table 8.

Overall we find strong evidence in favor of a contemporaneous link between exchange rates and risk reversals, but we find more limited evidence of exchange rate predictability.

4 Conclusion

The objective of this paper is to provide a simple model-based estimation of the share of carry trade returns that can be attributed to disaster risk. Our main empirical result shows that disaster premia explain one fourth of carry trade returns. This result suggests that the introduction of a time-varying disaster risk in exchange rate models, as in Farhi and Gabaix (2008), is empirically relevant.

Although we find that disaster risk plays a significant role in explaining currency returns, we fall short of fully solving the carry trade puzzle though disasters. In fact, our findings suggest that

[32] Carr and Wu (2007) also report high contemporaneous correlation between currency excess returns and risk reversals for the yen and the British pound against the U.S. dollar.

a typical investor can still obtain significant carry trade returns while being hedged against large currency crashes. Several interpretations of these hedged excess returns are possible. First, the investor naturally expects to be compensated for the remaining Gaussian, non–disaster risk. In bad times high–interest rate currencies tend to depreciate and low–interest rate currencies tend to appreciate. Second, out-of-the-money options might be relatively cheap in our sample. These options are not default-free, and counterparty risk might push their prices downward.

References

Ait-Sahalia, Yacine, Yubo Wang, and Francis Yared, "Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?," *Journal of Econometrics*, 2001, *102*, 67–110.

Akram, Q. Farooq, Dagfinn Rime, and Lucio Sarno, "Arbitrage in the Foreign Exchange Market: Turning on the Microscope," *Journal of International Economics*, 2008, *forthcoming.*

Backus, David and Gregor Smith, "Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods," *Journal of International Economics*, 1993, *35*, 297–316.

___ **, Mikhail Chernov, and Ian Martin**, "Disasters Implied by Equity Index Options," 2009. Working Paper.

___ **, Silverio Foresi, and Chris Telmer**, "Affine Models of Currency Pricing: Accounting for the Forward Premium Anomaly," *Journal of Finance*, 2001, *56*, 279–304.

Bakshi, Gurdip, Peter Carr, and Liuren Wu, "Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies," *Journal of Financial Economics*, 2008, *87*, 132156.

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 2004, *59* (4), 1481 – 1509.

___ **and Ivan Shaliastovich**, "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," June 2008. Working Paper Duke University.

Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 2006, *121*, 823–866.

___ , "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 2009, *99* (1), 243–264.

___ **and José F. Ursua**, "Consumption Disasters in the Twentieth Century," *American Economic Review*, 2008, *98* (2), 58–63.

___ **and** ___ , "Macroeconomic Crises since 1870," *Brookings Papers on Economic Activity*, 2008, *Spring*, 255–335.

___ **and** ___ , "Stock Market Crashes and Depressions," 2009. Working Paper NBER 14760.

Bates, David, "Dollar jump fears, 1984-1992: distributional abnormalities implicit in currency futures options," *Journal of International Money and Finance*, 1996, *15* (1), 65–93.

____, "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options," *Journal of Financial Economics*, 1996, *9*, 69 – 107.

____, "U.S. Stock Market Crash Risk, 1926 - 2006," 2009. Working Paper.

Bhansali, Vineer, "Volatility and the Carry Trade," *The Journal of Fixed Income*, Winter 2007, pp. 72–84.

Black, Fischer and Myron Scholes, "The Pricing of Options and Corporate Liabilities," *The Journal of Political Economy*, 1973, *81* (3), 637–654.

Bollerslev, Tim and Viktor Todorov, "Tails, Fears and Risk Premia," 2009. Working Paper.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, *NBER Macroeconomics Annual*, University of Chicago Press, 2008.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, "Can Peso Problems Explain the Returns to the Carry Trade?," May 2008. Working Paper.

Campa, Jose M., P.H. Kevin Chang, and Robert L. Reider, "Implied Exchange Rate Distributions: Evidence from OTC Option Markets," *Journal of International Money and Finance*, 1998, *17*, 117–160.

Campbell, John and John Cochrane, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.," *Journal of Political Economy*, 1999, *107* (2), 205–251.

Carr, Peter and Liuren Wu, "Stochastic Skew in Currency Options," *Journal of Financial Economics*, 2007, *86*, 213–247.

Cochrane, John H., *Asset Pricing*, Princeton, N.J.: Princeton University Press, 2005.

Corte, Pasquale Della, Lucio Sarno, and Ilias Tsiakas, "The Forward Volatility Bias in Foreign Exchange," 2009. Working Paper.

Engel, Charles and James D. Hamilton, "Long Swings in the Dollar: Are They in the Data and Do Markets Know It?," *American Economic Review*, 1990, *80* (4), 689–713.

Fama, Eugene, "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 1984, *14*, 319–338.

Farhi, Emmanuel and Xavier Gabaix, "Rare Disasters and Exchange Rates," 2008. NBER Working Paper 13805.

Gabaix, Xavier, "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," 2008. NBER Working Paper 13724.

Garman, Mark B. and Steven W. Kohlhagen, "Foreign Currency Option Values," *Journal of International Money and Finance*, December 1983, *2* (3), 231–237.

Gilmore, Stephen and Fumio Hayashi, "Emerging Market Currency Excess Returns," 2008. Working Paper.

Gourio, Francois, "Disasters and Recoveries," *American Economic Review, Papers and Proceedings*, 2008, *98* (2), 68–73.

Graveline, Jeremy J., "Exchange Rate Volatility and the Forward Premium Anomaly," 2006. Working Paper.

Guo, Kai, "Exchange Rates and Asset Prices in An Open Economy with Rare Disasters," 2007. Working Paper, Harvard University.

Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 1982, *50*, 1029–54.

___ and Robert J. Hodrick, "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, October 1980, *88* (5), 829–853.

IMF, *World economic outlook : a survey by the staff of the International Monetary Fund*, IMF, 2008.

Jarque, Carlos M. and Anil Bera, "Efficient tests for normality, homoscedasticity and serial independence of regression residuals," *Economics Letters*, 1980, *6* (3), 255–259.

Jorion, Philippe, "Predicting Volatility in the Foreign Exchange Market," *Journal of Finance*, 1995, *50*, 507–528.

Julliard, Christian and Anisha Ghosh, "Can Rare Events Explain the Equity Premium Puzzle?," 2008. Working Paper.

Jurek, Jakub W., "Crash-neutral Currency Carry Trades," 2008. Working Paper.

Kaminsky, Graciela, "Is There a Peso Problem? Evidence from the Dollar/Pound Exchange Rate, 1976-1987," *American Economic Review*, 1993, *83*, 552–76.

Lewis, Karen K., *Handbook of International Economics*, Vol. 3, Amsterdam; New York and Oxford: Elsevier, North-Holland, 1995.

Lilliefors, H., "On the Kolmogorov-Smirnov test for normality with mean and variance unknown," *Journal of the American Statistical Association*, 1967, *62*, 399–402.

Liu, Jun, Jun Pan, and Tan Wang, "An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks," *Review of Financial Studies*, 2005, *18*, 131–164.

Lustig, Hanno and Adrien Verdelhan, "The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk," *American Economic Review*, March 2007, *97* (1), 89–117.

___, Nick Roussanov, and Adrien Verdelhan, "Common Risk Factor in Currency Returns," 2008. Working Paper.

Martin, Ian, "Consumption-Based Asset Pricing with Higher Cumulants," 2008. Working Paper.

Pan, Jun, "The jump-risk premia implicit in options: evidence from an integrated time-series study," *Journal of Financial Economics*, 2002, *63* (1), 3–50.

Rietz, Thomas A., "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 1988, *22*, 117–131.

Santa-Clara, Pedro and Shu Yan, "Crashes, Volatility, and the Equity Premium: Lessons from S&P500 Options," *Review of Economic Studies, forthcoming*, 2009.

Verdelhan, Adrien, "A Habit-Based Explanation of the Exchange Rate Risk Premium," *Journal of Finance, forthcoming*, 2009.

Wachter, Jessica A., "Can time-varying risk of rare disasters explain aggregate stock market volatility?," 2008. Working Paper.

Table 1: Excess Returns: Advanced Countries Sorted on Interest Rates

Portfolios	1	2	3	1	2	3
	Going Long			Going Short		
			Panel I: Unhedged			
Mean	−1.37	1.45	5.13	1.37	−1.45	−5.13
	[2.08]	[2.25]	[2.08]	[2.02]	[2.14]	[1.99]
Sharpe Ratio	−0.19	0.19	0.71	0.19	−0.19	−0.71
			Panel II: Hedged at 10 delta			
Mean	−2.30	0.65	4.06	0.74	−1.58	−5.33
	[1.93]	[1.99]	[1.90]	[1.86]	[1.94]	[1.87]
Sharpe Ratio	−0.33	0.09	0.60	0.11	−0.23	−0.81
			Panel III: Hedged at 25 delta			
Mean	−2.14	0.59	3.03	0.62	−1.21	−4.68
	[1.72]	[1.82]	[1.66]	[1.48]	[1.59]	[1.53]
Sharpe Ratio	−0.36	0.09	0.51	0.12	−0.21	−0.86
			Panel IV: Hedged ATM			
Mean	−1.33	0.61	1.68	0.02	−0.86	−3.47
	[1.27]	[1.40]	[1.26]	[1.07]	[1.13]	[1.10]
Sharpe Ratio	−0.31	0.13	0.39	0.00	−0.21	−0.91

Notes: This table reports average currency excess returns that are unhedged or hedged at 10 delta, at 25 delta, and at-the-money for our three portfolios. In the left section, we assume that the U.S. investor goes long in the foreign currency; in the right section, we assume that the U.S. investor goes short in the foreign currency. In each case, we report the mean excess return, its standard error, and the corresponding Sharpe ratio. The mean and standard deviations are annualized (multiplied respectively by 12 and $\sqrt{12}$). The Sharpe ratio corresponds to the ratio of the annualized mean to the annualized standard deviation. Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Portfolio 1 contains currencies with the lowest interest rates; portfolio 3 contains currencies with the highest interest rates. The horizon of the excess returns and the option maturity is one month for each. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008.

Table 2: Implied Volatilities and Risk Reversals: Advanced Countries Sorted on Interest Rates

Portfolios	1	2	3
		Panel I: Implied Volatilities	
10δ Put	9.78	10.09	11.50
	[0.14]	[0.17]	[0.20]
25δ Put	9.38	9.56	10.60
	[0.15]	[0.16]	[0.17]
ATM	9.33	9.31	10.02
	[0.14]	[0.16]	[0.17]
25δ Call	9.78	9.55	10.02
	[0.15]	[0.16]	[0.15]
10δ Call	10.51	10.05	10.39
	[0.16]	[0.17]	[0.16]
		Panel II: Risk Reversals (Implied Volatilities)	
Mean RR10	−0.73	0.05	1.12
	[0.06]	[0.05]	[0.06]
Mean RR25	−0.40	0.01	0.58
	[0.03]	[0.03]	[0.03]

Notes: This table reports average implied volatilities and risk reversals by portfolios. The first panel reports average implied volatilities on put and call contracts for strike prices at 10 delta, at 25 delta, and at-the-money. The second panel reports risk reversals at 10 delta and 25 delta measured in terms of implied volatilities. The figures are quoted in annual percentages. Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Portfolio 1 contains currencies with the lowest interest rates; portfolio 3 contains currencies with the highest interest rates. The horizon of the excess returns and the option maturity is one month for each. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008.

Table 3: Excess Returns: Advanced Countries Sorted on Risk Reversals

Portfolios	1	2	3	1	2	3
		Going Long			Going Short	
			Panel I: Unhedged			
Mean	0.48	1.22	3.70	−0.48	−1.22	−3.70
	[2.10]	[2.11]	[1.95]	[2.06]	[2.05]	[1.87]
Sharpe Ratio	0.06	0.16	0.54	−0.06	−0.16	−0.54
			Panel II: Hedged at 10 delta			
Mean	−0.38	0.47	2.57	−1.00	−1.39	−3.96
	[2.02]	[2.05]	[1.83]	[1.98]	[1.90]	[1.76]
Sharpe Ratio	−0.05	0.07	0.39	−0.14	−0.20	−0.62
			Panel III: Hedged at 25 delta			
Mean	−0.21	0.05	1.83	−0.68	−1.29	−3.45
	[1.68]	[1.70]	[1.51]	[1.66]	[1.61]	[1.45]
Sharpe Ratio	−0.03	0.01	0.33	−0.12	−0.23	−0.65
			Panel IV: Hedged ATM			
Mean	−0.03	−0.09	1.17	−0.53	−1.33	−2.55
	[1.28]	[1.31]	[1.10]	[1.12]	[1.16]	[1.06]
Sharpe Ratio	−0.01	−0.02	0.29	−0.13	−0.32	−0.69

Notes: This table reports average currency excess returns that are unhedged or hedged at 10 delta, at 25 delta, and at the money for our three portfolios. In the left section, we assume that the U.S. investor goes long in the foreign currency; in the right section, we assume that the U.S. investor goes short in the foreign currency. In each case, we report the mean excess return, its standard error, and the corresponding Sharpe ratio. The mean and standard deviations are annualized (multiplied respectively by 12 and $\sqrt{12}$). The Sharpe ratio corresponds to the ratio of the annualized mean to the annualized standard deviation. Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Portfolio 1 contains currencies with the lowest risk reversals at 10 delta; portfolio 3 contains currencies with the highest risk reversals at 10 delta. The horizon of the excess returns and the option maturity is one month for each. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008.

Table 4: Disaster Risk Premia - Advanced Countries Sorted on Interest Rates

Panel I: Carry Excess Returns					
	Unhedged Carry	Hedged at 10δ	Hedged at 25δ	Hedged ATM	
Mean	6.50	4.80	3.65	1.70	
	[1.88]	[1.59]	[1.41]	[1.12]	
Mean Spread		1.70	2.85	4.80	
		[0.41]	[0.85]	[1.32]	
Panel II: Estimations					
	10δ	25δ	ATM	10δ, 25δ, and ATM	GMM 2^{nd} Stage
$\overline{\pi}^D$	1.16	1.63	3.10	1.96	1.01
	[0.41]	[0.87]	[1.68]	[0.93]	[0.36]
$\overline{\pi}^G$	5.33	4.87	3.40	4.53	4.77
	[1.79]	[1.87]	[2.21]	[1.87]	[1.92]
$\overline{\pi}^D - \overline{\pi}^G$	-4.17	-3.23	-0.30	-2.57	-3.76
	[1.90]	[2.31]	[3.51]	[2.35]	[2.02]

Notes: This first panel of this table reports average returns on hedged and unhedged currency carry trades and their standard errors. We use the currency portfolios presented in Table 1. Carry trades correspond to returns on the last minus returns on the first portfolio. We consider different hedges: 10-delta, 25-delta and at-the-money. We also report the average difference between unhedged and hedged carry trades. The second panel reports structural estimates. Here $\overline{\pi}^D$ denotes the part of the carry excess return linked to disaster risk and $\overline{\pi}^G$ corresponds to the Gaussian, non-disaster part of the same excess return. These estimates are obtained using hedged returns at 10 delta (first column), 25 delta (second column), and ATM (third column) or at 10 delta, at 25 delta, and ATM combined (fourth and fifth columns). Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008.

Table 5: Disaster Risk Premia : Advanced Countries Sorted on Risk Reversals

Panel I: Carry Excess Returns					
	Unhedged Carry	Hedged at 10δ	Hedged at 25δ	Hedged ATM	
Mean	3.22	1.57	1.15	0.64	
	[1.66]	[1.53]	[1.29]	[1.14]	
Mean Spread		1.65	2.07	2.58	
		[0.36]	[0.80]	[1.32]	
Panel II: Estimations					
	10δ	25δ	ATM	10δ, 25δ, and ATM	GMM 2^{nd} Stage
$\overline{\pi}^D$	1.48	1.68	1.94	1.70	1.41
	[0.36]	[0.87]	[1.72]	[0.94]	[0.32]
$\overline{\pi}^G$	1.74	1.54	1.28	1.52	1.67
	[1.67]	[1.74]	[2.11]	[1.74]	[1.78]
$\overline{\pi}^D - \overline{\pi}^G$	-0.26	0.14	0.66	0.18	-0.27
	[1.79]	[2.22]	[3.49]	[2.28]	[1.90]

Notes: This first panel of this table reports average returns on hedged and unhedged currency carry trades and their standard errors. We use the currency portfolios presented in Table 3. Carry trades correspond to returns on the last minus returns on the first portfolio. We consider different hedges: 10-delta, 25-delta and at-the-money. We also report the average difference between unhedged and hedged carry trades. The second panel reports structural estimates. Here $\overline{\pi}^D$ denotes the part of the carry excess return linked to disaster risk and $\overline{\pi}^G$ corresponds to the Gaussian, non-disaster part of the same excess return. These estimates are obtained using hedged returns at 10 delta (first column), 25 delta (second column), and ATM(third column) or at 10 delta, at 25 delta and ATM combined (fourth and fifth columns). Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008.

Table 6: Disaster Risk Premia - Advanced Countries Sorted on Interest Rates with Transaction Costs

	Panel I: Carry Excess Returns				
	Unhedged Carry	Hedged at 10δ	Hedged at 25δ	Hedged ATM	
Mean	6.25	4.21	2.83	0.78	
	[1.83]	[1.67]	[1.44]	[1.14]	
Mean Spread		2.04	3.42	5.47	
		[0.43]	[0.85]	[1.34]	
	Panel II: Estimations				
	10δ	25δ	ATM	10δ, 25δ, and ATM	GMM 2^{nd} Stage
$\overline{\pi}^D$	1.57	2.47	4.69	2.91	1.28
	[0.41]	[0.87]	[1.68]	[0.93]	[0.37]
$\overline{\pi}^G$	4.67	3.78	1.56	3.34	4.02
	[1.81]	[1.91]	[2.29]	[1.91]	[1.96]
$\overline{\pi}^D - \overline{\pi}^G$	-3.10	-1.31	3.14	-0.42	-2.74
	[1.91]	[2.35]	[3.60]	[2.41]	[2.04]

Notes: This first panel of this table reports average returns on hedged and unhedged currency carry trades and their standard errors. We use the currency portfolios presented in Table 1. Carry trades correspond to returns on the last minus returns on the first portfolio. We consider different hedges: 10-delta, 25-delta and at-the-money. We also report the average difference between unhedged and hedged carry trades. The second panel reports structural estimates. Here $\overline{\pi}^D$ denotes the part of the carry excess return linked to disaster risk and $\overline{\pi}^G$ corresponds to the Gaussian, non-disaster part of the same excess return. These estimates are obtained using hedged returns at 10 delta (first column), 25 delta (second column), and ATM (third column) or at 10 delta, at 25 delta, and ATM combined (fourth and fifth columns). Standard errors are obtained by bootstrapping the monthly excess returns under the assumptions that they are i.i.d. Data are monthly, from JP Morgan. The sample period is 1/1996 – 8/2008. We assume annual transaction costs of 0.25% on unhedged returns and bid–ask spreads of 5% on implied volatilities.

Table 7: Changes in Risk Reversals and Exchange Rates: Contemporaneous Specifications within Portfolios

Dependant Variable:	Exchange Rates					
	Panel I: Raw Variables			Panel II: De-meaned Variables		
Portfolios	P1	P2	P3	P1	P2	P3
Risk Reversals	-126.63	-131.82	-105.18	-119.95	-132.09	-145.43
Strike: Delta 10	[12.93]***	[24.22]***	[28.46]***	[27.30]***	[18.09]***	[17.87]***
Observations	155	155	155	155	155	155
R^2	0.4	0.28	0.41	0.37	0.42	0.35
Risk Reversals	-77.56	-62.66	-49.29	-54.95	-62	-74.57
Strike: Delta 25	[8.46]***	[18.28]***	[16.76]***	[19.08]***	[17.25]***	[14.26]***
Observations	155	155	155	155	155	155
R^2	0.38	0.25	0.36	0.32	0.36	0.31
Risk Reversals	-61.64	-39.38	-30.31	-96.83	-45.76	-69.08
Strike: Forward +/- 10%	[14.66]***	[36.52]	[13.61]**	[60.45]	[12.88]***	[30.00]**
Observations	96	125	133	96	125	133
R^2	0.22	0.14	0.28	0.05	0.25	0.16
Risk Reversals	-40.08	-48.97	-46.8	-50.99	-52.8	-47.9
Strike: Forward +/- 5%	[4.69]***	[6.05]***	[7.66]***	[7.51]***	[5.08]***	[6.80]***
Observations	147	155	144	147	155	144
R^2	0.39	0.3	0.46	0.42	0.44	0.32

Notes: This table documents contemporaneous relationships between changes in nominal exchange rates and changes in risk reversals. Constant terms are included but not reported. Panel I presents results based on raw variables; panel II uses cross-sectionally de-meaned variables to control for the specific role of the U.S. dollar. Changes in exchange rates correspond to monthly log changes; changes in risk reversals correspond to first differences. Each horizontal panel presents the results of regressions including a different risk-reversal measure. Standard errors obtained from bootstrap procedures using 1000 replications are presented below the point estimates. The symbols ***, **, and * indicate statistical significance at 1%, 5%, and 10% confidence levels respectively. The sample comprises currencies from advanced countries (excluding observations with non floating exchange rate according to the IMF de facto classification). Data are monthly, from JP Morgan. The sample period is 02/1996 – 08/2008.

Table 8: Risk Reversals, Exchange Rate Changes and Currency Excess Returns: Predictive Specifications within Portfolios

Dependant Variable:	Panel I: Exchange Rates			Panel II: Currency Excess Returns		
Portfolios	P1	P2	P3	P1	P2	P3
Interest Rate Differentials	-1.27	-4.16	-0.97	-2.27	-5.17	-1.97
	[1.52]	[1.77]**	[1.08]	[1.49]	[1.74]***	[1.06]*
Risk Reversals: (+/- 10%)	13.1	-1.12	-3.7	13.11	-1.14	-3.72
	[13.36]	[37.33]	[19.30]	[14.94]	[40.95]	[19.38]
Observations	109	129	138	109	129	138
R^2	0.02	0.04	0.01	0.04	0.06	0.03
Interest Rate Differentials	-2.78	-3.49	-0.96	-3.78	-4.5	-1.97
	[1.28]**	[1.72]**	[1.15]	[1.27]***	[1.79]**	[1.16]*
Risk Reversals: (+/- 5%)	0.81	-2.37	-3.44	0.81	-2.39	-3.47
	[5.52]	[9.54]	[7.53]	[5.55]	[9.69]	[7.26]
Observations	109	129	138	109	129	138
R^2	0.03	0.04	0.01	0.05	0.06	0.02
Interest Rate Differentials	-2.5	- 3.48	-0.7	-3.5	-4.49	-1.71
	[1.21]**	[1.71]**	[1.02]	[1.22]***	[1.65]***	[1.06]
Risk Reversals: Delta 10	4.18	-8.18	-7.39	4.17	-8.23	-7.44
	[16.66]	[25.22]	[18.81]	[17.10]	[26.06]	[18.55]
Observations	155	155	155	155	155	155
R^2	0.02	0.04	0.01	0.05	0.06	0.02
Interest Rate Differentials	-2.51	-3.49	-0.76	-3.52	-4.5	-1.76
	[1.26]**	[1.69]**	[1.07]	[1.23]***	[1.68]***	[1.12]
Risk Reversals: Delta 25	0.39	-5.32	-5.06	0.38	-5.35	-5.09
	[9.31]	[13.27]	[10.02]	[9.41]	[14.19]	[10.90]
Observations	155	155	155	155	155	155
R^2	0.02	0.04	0.01	0.05	0.06	0.02

Notes: This table presents results of predictability tests. We regress monthly changes in nominal exchange rates (panel I) or monthly currency excess returns (panel II) on risk reversals and interest differentials. The interest differential is defined as the difference between the domestic and the foreign interest rate. The null hypothesis of UIP not being rejected is a coefficient of 1 for the interest rate differential in panel I and a coefficient of 0 in panel II. Constant terms are included but not reported. Standard errors obtained from a bootstrap procedure using 1000 replications are presented below their respective point estimates. The symbols ***, **, and * indicate statistical significance at 1%, 5%, and 10% confidence levels respectively. The sample comprises currencies from advanced countries (excluding observations with non floating exchange rate according to the IMF de facto classification.) Data are monthly, from JP Morgan. The sample period is 02/1996 – 08/2008.

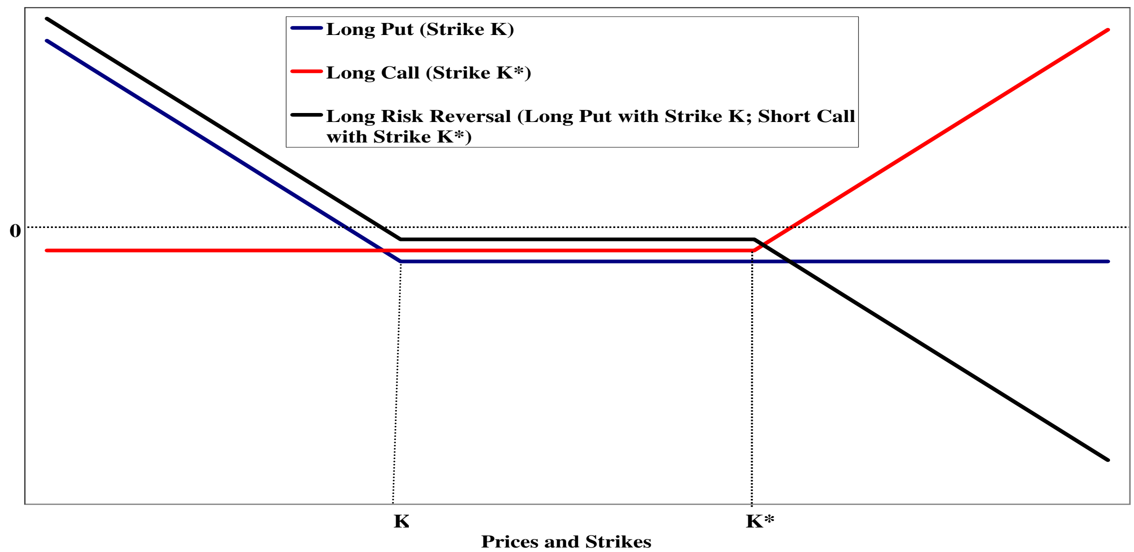

Figure 1: Option Payoffs

This figure presents the payoffs of different option investments as a function of the underlying asset prices and strikes. We consider the payoff of buying a call (with strike $K^{\star}$) or buying a put option (with strike K). Finally, we consider a risk reversal that corresponds to selling a call (with strike $K^{\star}$) and simultaneously buying a put (with strike K).

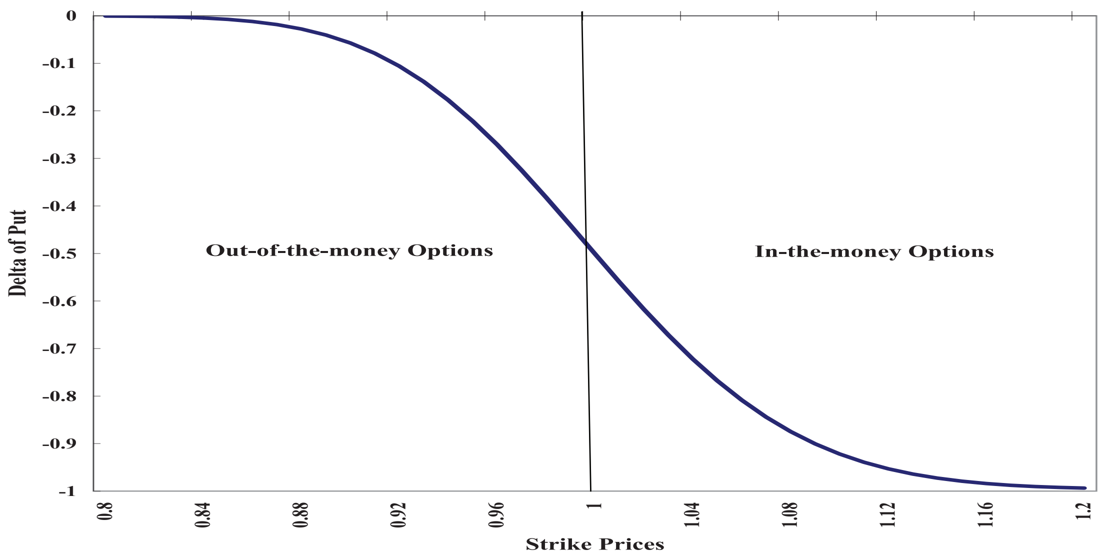

Figure 2: Deltas

This figure presents the deltas of put options as a function of their prices. The delta of an option is defined as the rate of change of the option price with respect to the price of the underlying asset. The delta of a put varies between 0 for the most deep-out-of-the-money options and -1 for the most deep-in-the-money options. The figure is computed for a currency put option with a one-month maturity, an annualized implied volatility of 10%, and foreign and domestic interest rates both set equal to 4% per annum.

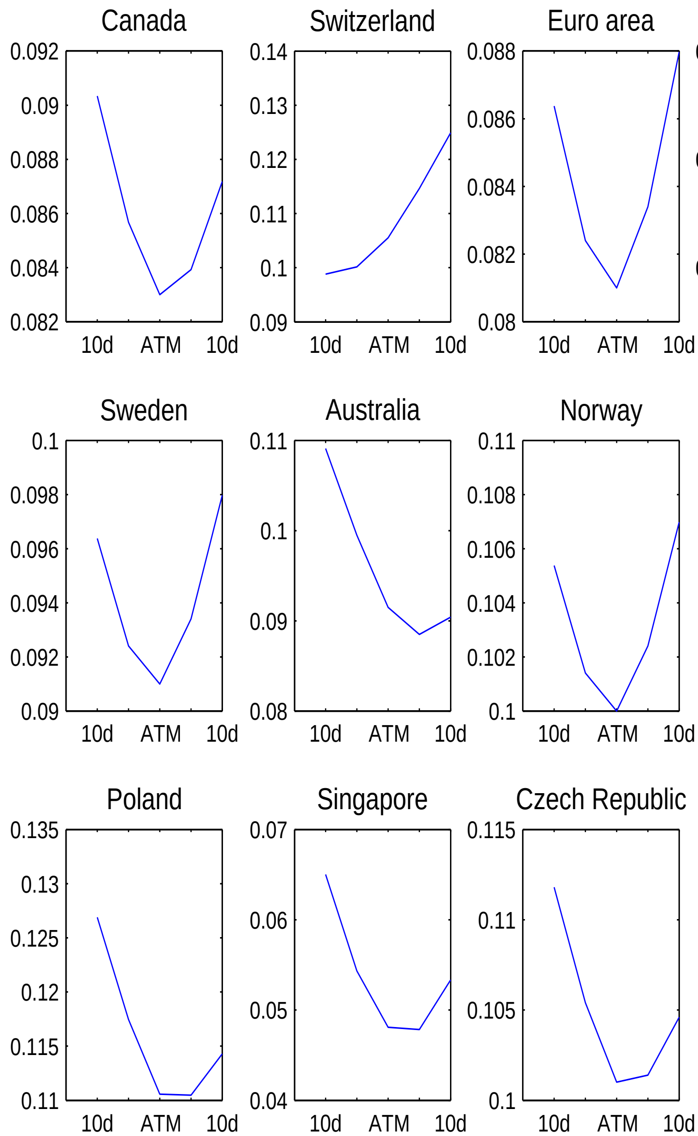

Figure 3: One-Month Option-Implied Volatility Smiles, August 2008

This figure plots, for each currency in our sample, implied volatilities for different strike prices. Implied volatilities are in percentages; strike prices are scaled by spot rates.

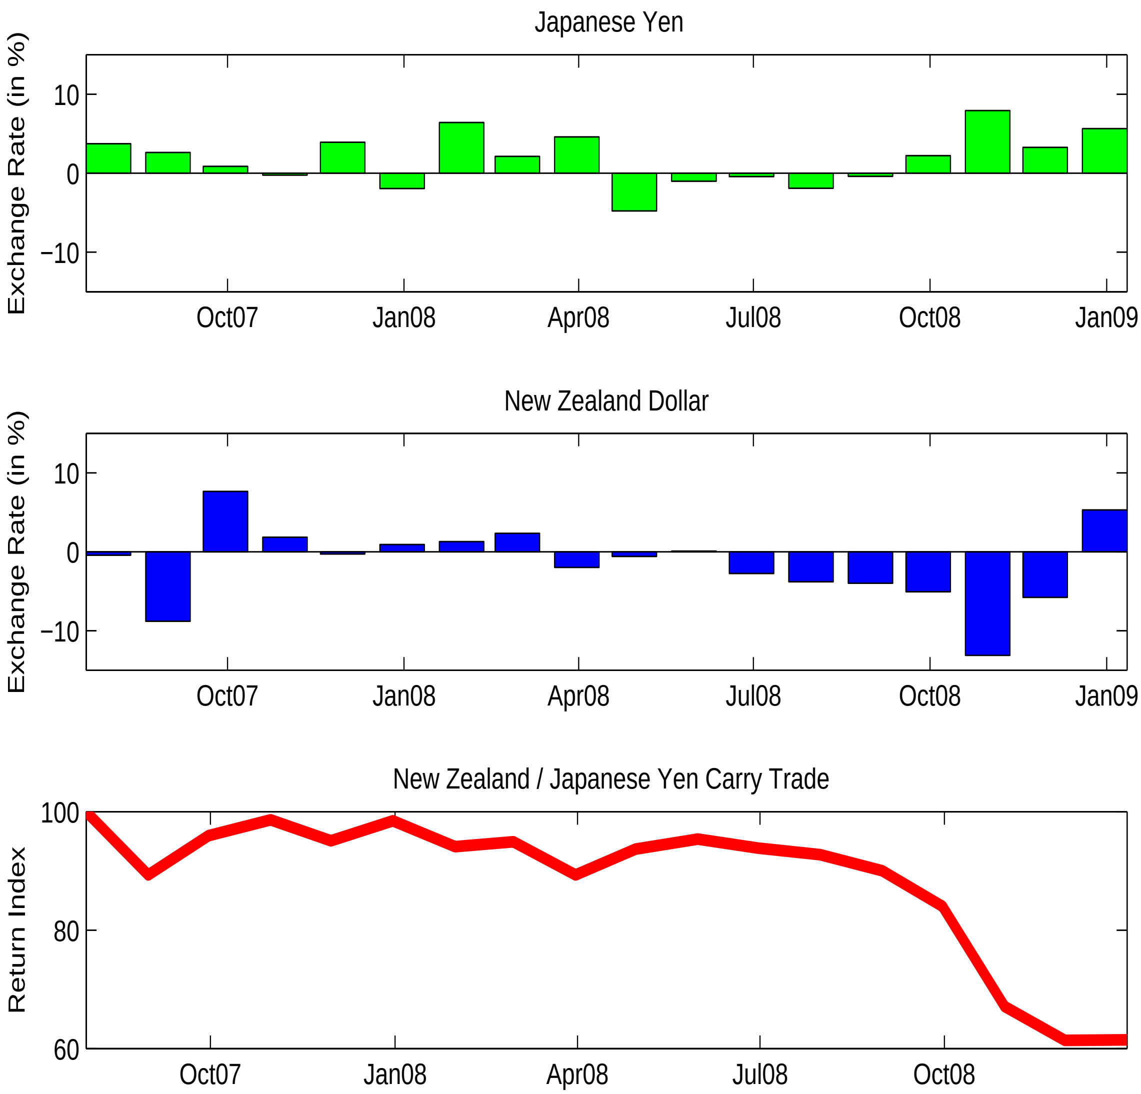

Figure 4: New Zealand Dollar and Japanese Yen

This figure plots monthly changes in exchange rates for the New Zealand dollar and the Japanese yen as well the return index on a carry trade strategy that borrows in yen to invest in New Zealand dollars. The sample period is 7/2007–12/2008.

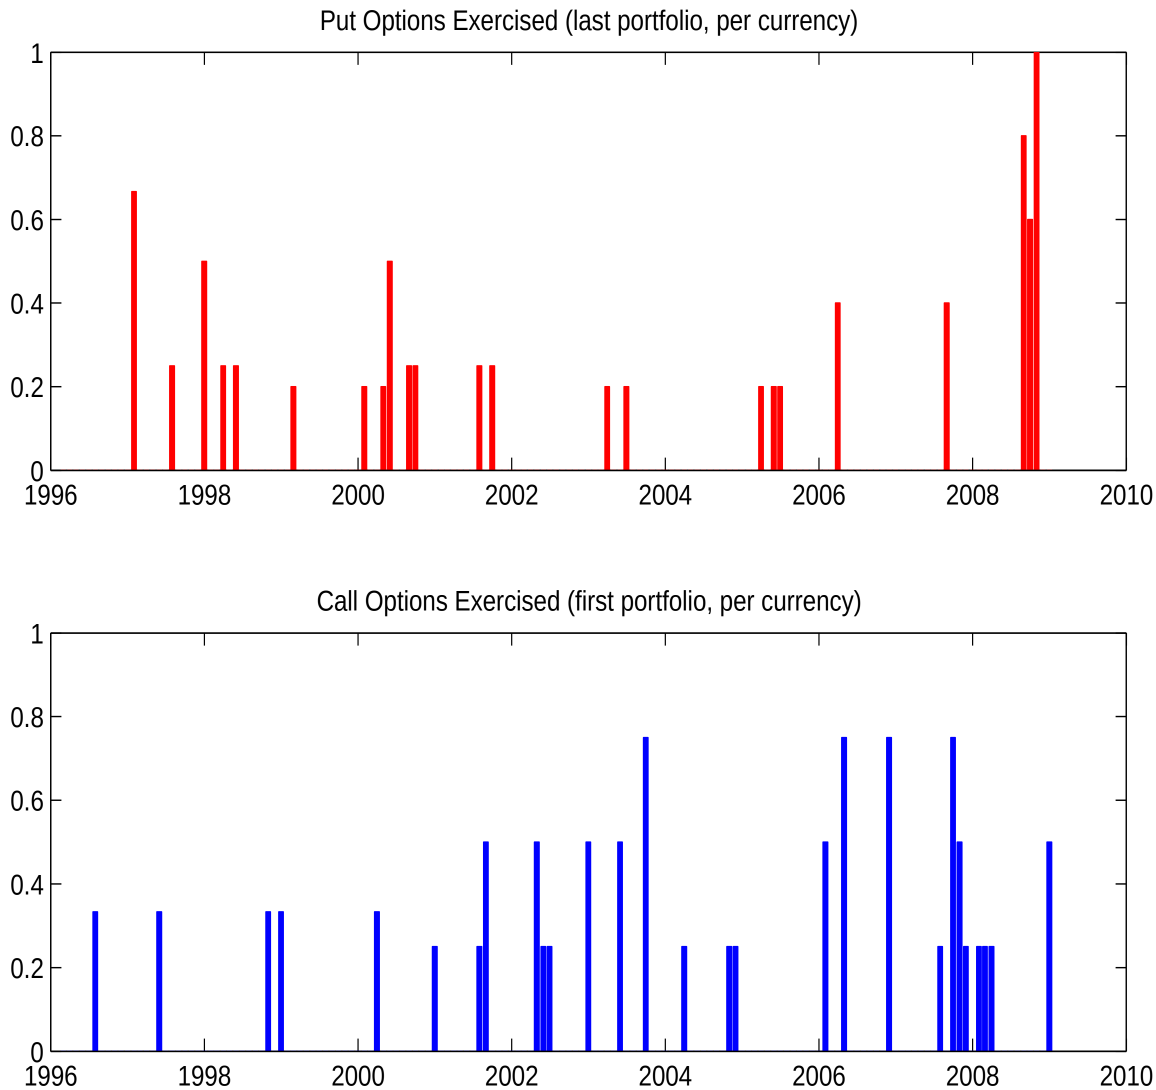

Figure 5: Options Exercised

This figure plots the frequency of call and put options exercised (respectively) in the first and last portfolios. At each moment in time, the frequency is obtained as the number of options exercised divided by the number of currencies in the portfolio at that time. We consider only options at 10 delta. The sample period is 2/1996 – 12/2008.

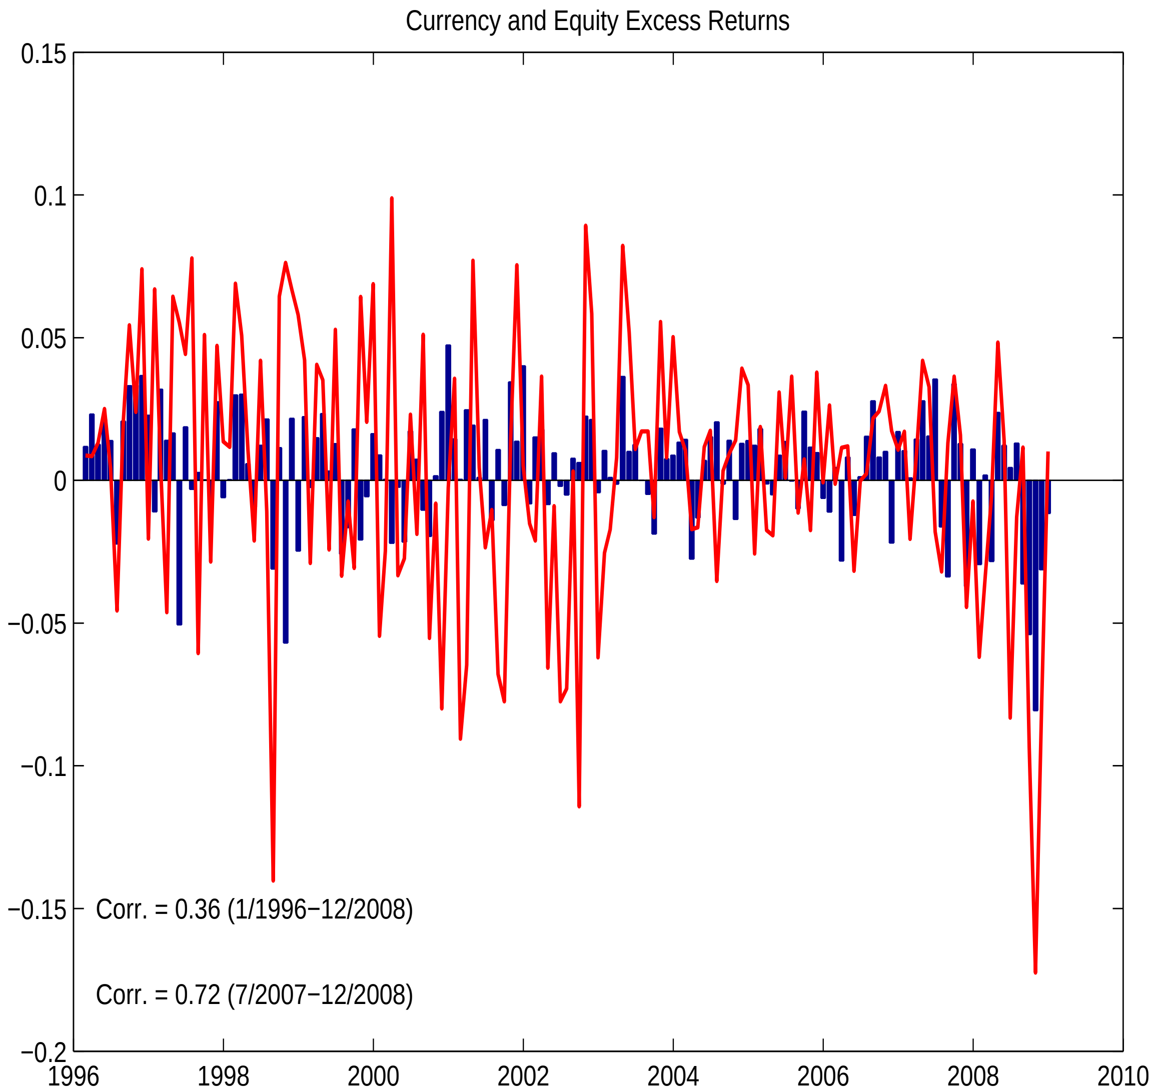

Figure 6: Currency Carry Trades and Equity Returns

This figure plots monthly currency carry trades and U.S. equity returns. Carry excess returns (blue bars) correspond to our sample of advanced countries. Data are monthly, from JP Morgan (IMF). Equity returns (red line) correspond to the U.S. MSCI index. The sample period is 2/1996 – 12/2008.

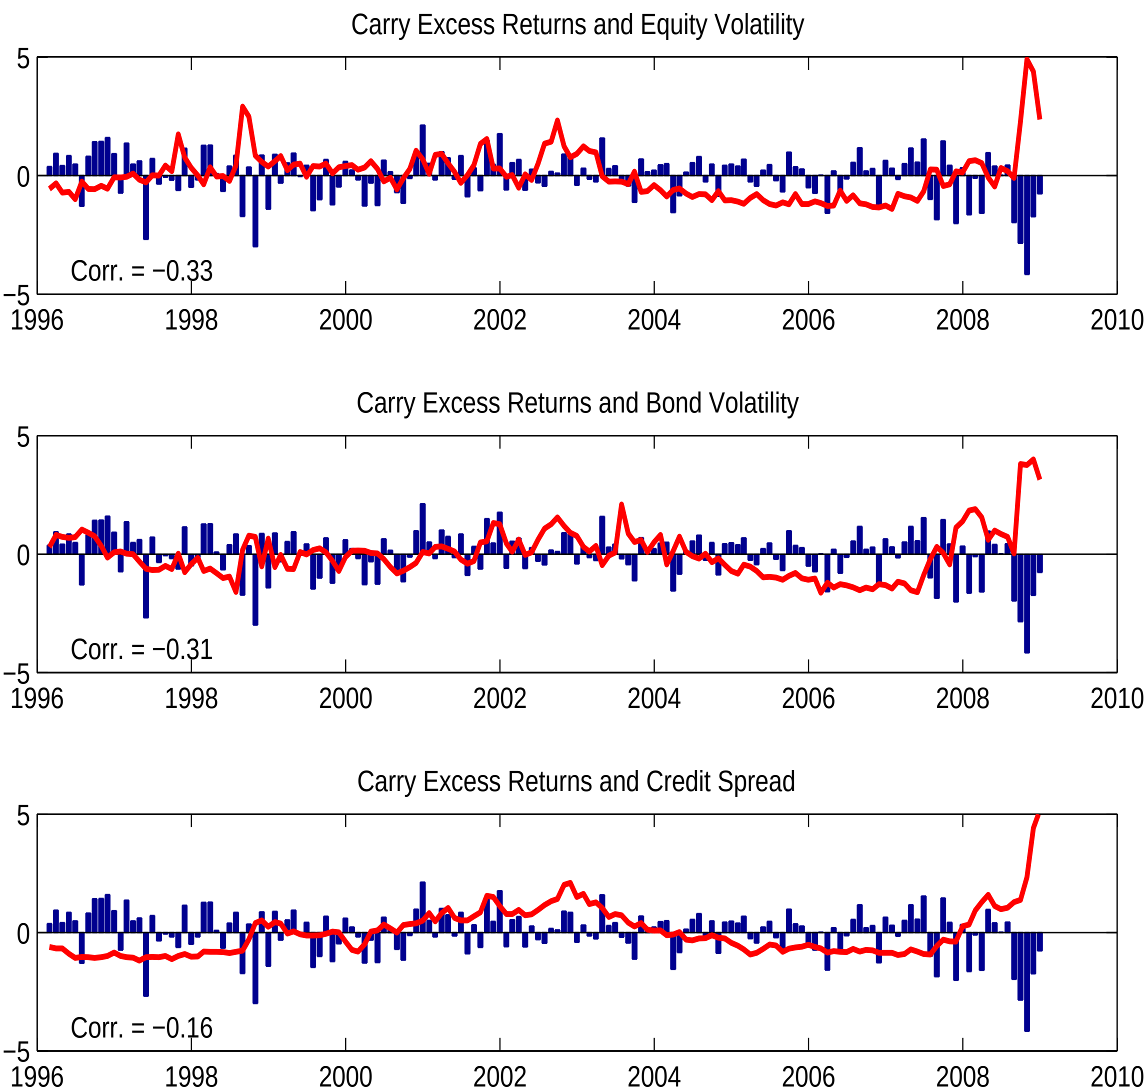

Figure 7: Carry Returns and Risk Measures

This figure plots carry excess returns and different risk measures. The upper panel uses the equity option–implied volatility index VIX; below are the bond option–implied volatility MOVE index and the credit spread (measured as the yield spreads between BAA bonds and 10-year U.S. Treasury bonds). Currency returns (blue bars) and risk measures (red lines) are all de-meaned and then divided by their standard deviations. The sample period is 2/1996–12/2008.

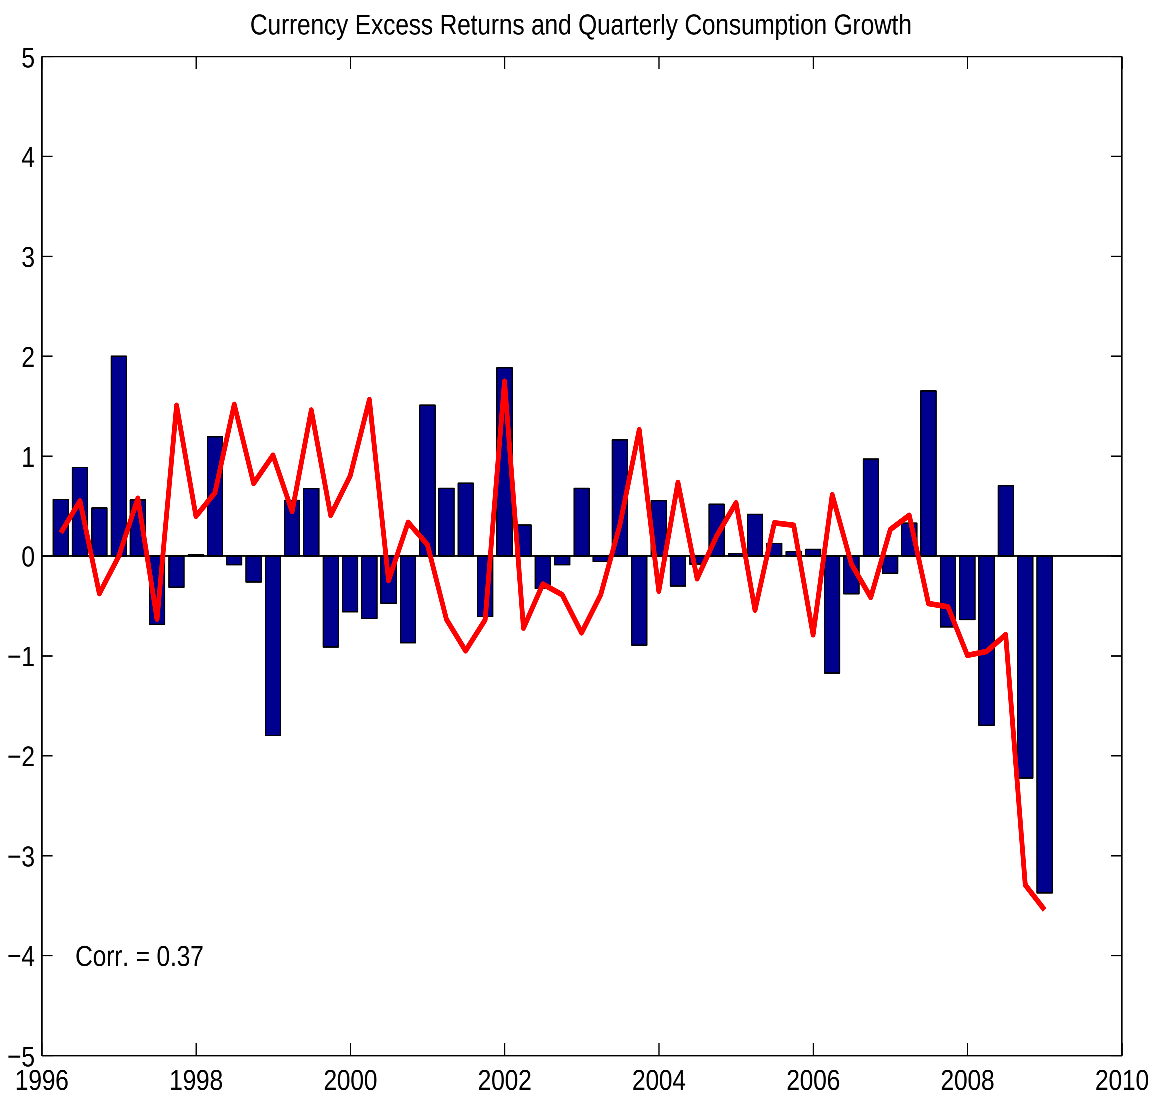

Figure 8: Carry Returns and Consumption Growth

This figure presents quarterly carry excess returns and real consumption growth per capita. Currency returns (blue bars) and consumption growth (red line) are all de-meaned and then divided by their standard deviations. The sample period is 2/1996–12/2008.

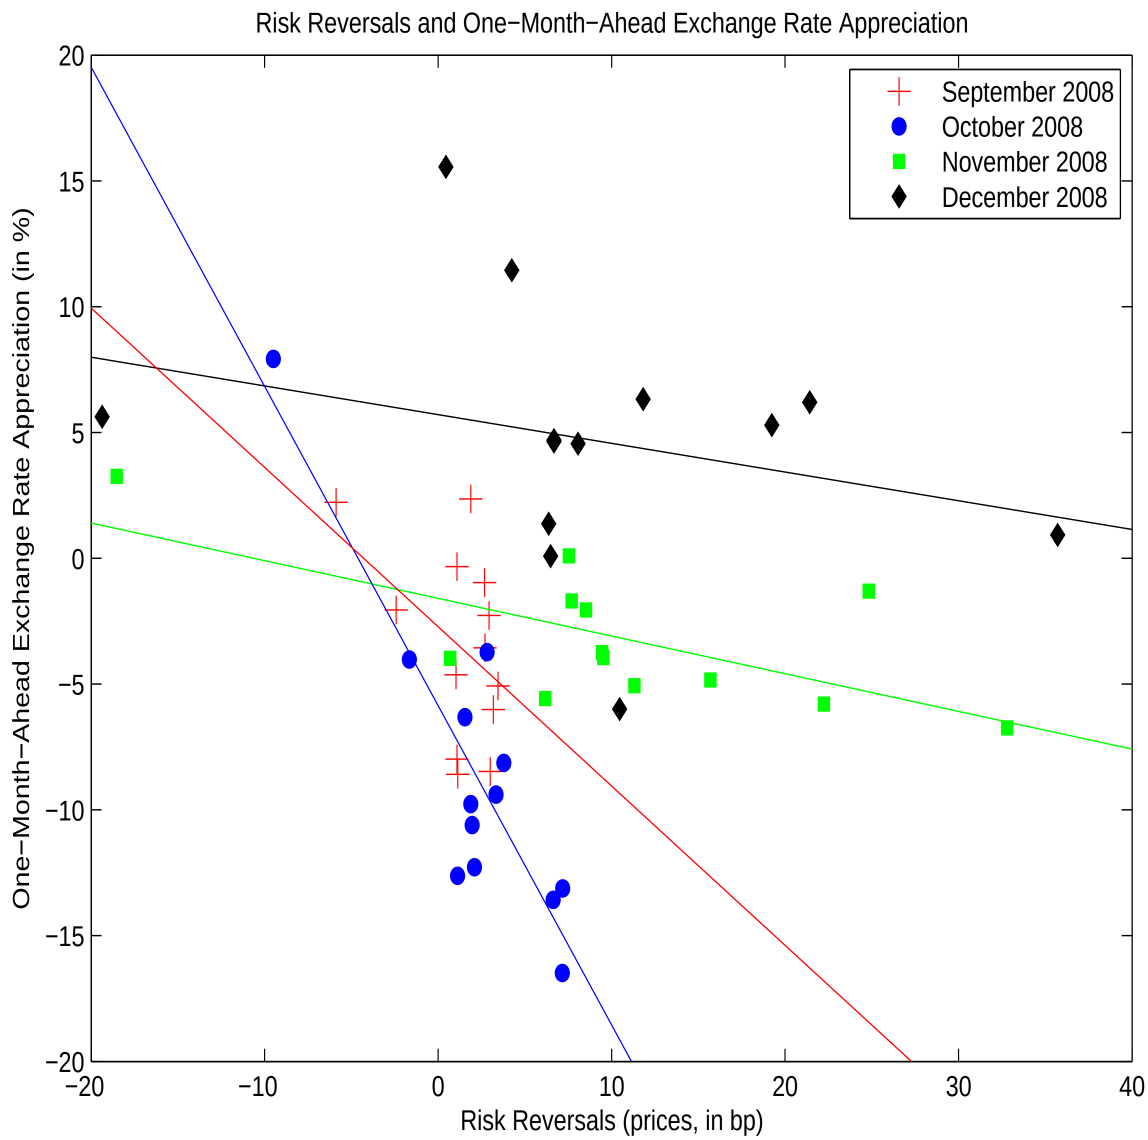

Figure 9: Risk Reversals and Changes in Exchange Rates, Fall 2008

This figure plots risk reversals at 10 delta and subsequent one-month changes in exchange rates for each month of fall 2008. Risk reversal prices are in basis points; changes in exchange rates are in percentages. Increases in exchange rates correspond to depreciation of the U.S. dollar. Exchange rate changes between date t and $t+1$ are dated $t+1$. The sample focuses on advanced countries and covers the period from 9/2008 to 12/2008.