International Correlation Risk*

Philippe Mueller
London School of Economics†

Andreas Stathopoulos
USC Marshall School of Business‡

Andrea Vedolin
London School of Economics§

Abstract

Foreign exchange correlation is a key driver of risk premia in the cross-section of carry trade returns. First, we show that the correlation risk premium, defined as the difference between the risk-neutral and objective measure correlation is large (15% per year) and highly time-varying. Second, sorting currencies according to their exposure with correlation innovations yields portfolios with attractive risk and return characteristics. We also find that high (low) interest rate currencies have negative (positive) loadings on the correlation risk factor. To address our empirical findings, we consider a multi-country general equilibrium model with time-varying risk aversion generated by external habit preferences. In the model, currency risk premia mostly compensate for exposure to global risk aversion, defined as a weighted average of country risk aversions. Given countercyclical real interest rates, the model can also address the forward premium puzzle, as high interest rate currencies are exposed to (while low interest rate currencies provide a hedge to) global risk aversion risk. We also show that high global risk aversion is associated with high conditional exchange rate variance and covariance, providing theoretical justification for sorting currencies on their exposure to fluctuations of exchange rate conditional second moments.

JEL Classification Codes: F31, G12, G13.

Keywords: Carry Trade, Correlation Risk, Habit Formation, International Finance, Exchange Rates.

First Version: December 2011.

This Version: February 2012.

---

*Preliminary, please do not distribute. Comments are very welcome. Philippe Mueller and Andrea Vedolin acknowledge financial support of STICERD at LSE. We would like to thank Mike Chernov and seminar participants of London School of Economics, LUISS Guido Carli, Rome, University of Lund, University of Piraeus, Stockholm School of Economics, Chicago Booth Finance Symposium, and the 6th End of Year Meeting of the Swiss Economists Abroad for thoughtful comments.

†Department of Finance, Houghton Street, London WC2A 2AE, UK, Phone: +44 20 7955 7012, Email: p.mueller@lse.ac.uk.

‡Bridge Hall 308, 3670 Trousdale Parkway, Los Angeles, CA 90089, USA, Phone: +1 213 740 7628, Email: astathop@marshall.usc.edu.

§Department of Finance, Houghton Street, London WC2A 2AE, UK, Phone: +44 20 7107 5017, Email: a.vedolin@lse.ac.uk.

It is only very recently that currencies have hit the front page of the financial press due to so called currency wars. However, the FX market has been in the thick of the action since the onset of the financial crisis in 2007. After a prolonged phase of carry trade out-performance in a low and decreasing volatility environment, the emergence of financial stress culminating with the Lehman Brothers bankruptcy led to a spike in FX market volatility. UBS, the world's second largest currency dealer, coined the notion of a "decade of super volatility" amongst global currencies.[1] Carry trades unwind during periods of high levels of volatility, which go in lock-step with high levels of correlation. Risk reversals provide a good market indicator of perceived risks in carry trades, i.e., the probability of large swings in exchange rates.

The behavior of the Swiss franc illustrates these risks. The Swiss franc (together with the Japanese Yen) constitutes a common funding currency due to its low interest rates. Moreover, it is often referred to as a safe haven. The Euro, on the other hand is perceived as riskier, especially on the eve of the largest financial crisis since more than half a century. Therefore, we would expect that a large devaluation of the Euro vis-à-vis the Swiss franc is a more likely event than vice versa. In Figure 1 we plot the EURCHF risk reversal, which is the difference between the 1 month implied volatility of an out-of-the-money call and put option. In addition, we depict the implied covariance between the EURCHF calculated using 1 month at-the-money implied volatilities. Note that the EURCHF risk reversal turns positive the first time on the date that the Swiss National Bank (SNB) intervenes in the currency markets to stop the appreciation of the Swiss franc (the first intervention since 1995). The second time the risk reversal turns positive is right before the SNB decided to abandon the exchange rate target of 1.50. Since then, the risk reversal has remained negative. It is interesting to note that the implied covariance between the two currencies exhibits pronounced spikes at these exact dates and these are also the points in time, when the carry trade or risk reversal do not work anymore.

[Insert Figure 1 approximately here.]

[1] See "FX Mega-Trends 2010–2020," UBS Investment Research, November 2010.

Motivated by this example, in this paper we study the link between correlation risk and carry returns both empirically and theoretically. We denote a safe haven currency as a rainy day asset, meaning, safe haven currencies provide a hedge in times of high correlation by yielding positive returns whereas high interest rate currencies are negatively related to global correlation risk and thus deliver low returns during periods of distress. We address two main questions. First, we ask whether correlation risk is priced in the cross-section of foreign exchange risk premia. Second, we ask what the underlying economic underpinnings are that drive this risk premium.

We tackle the first question in two different ways. First, we start by constructing the currency correlation risk premium from a cross-section of currency option prices and high frequency data on the underlying spot exchange rates. We find that the size of the estimated correlation risk premium is economically large and comparable to the one documented in the equity options literature. Driessen, Maenhout, and Vilkov (2009) estimate an equity correlation risk premium of approximately 18%, with an average realized correlation of 29% and an average implied correlation of 47%. In currency markets, the numbers are similar: The average implied correlation amounts to 55% and the average realized correlation to 48%. However, while in the equity market the correlation risk premium is predominantly positive except for the most recent crisis period (see Buraschi, Kosowski, and Trojani, 2011), the correlation risk premium in currency markets changes its sign quite often.

Second, we then proceed by quantifying the price of correlation risk in the cross-section of currency excess returns using a portfolio sorting approach. Intuitively, if correlation risk is priced in currency markets, then currencies with low correlation betas (i.e. that co-move weakly with correlation) should carry higher returns, whereas low correlation risk currencies (i.e. currencies that co-move strongly with correlation and thus hedge against correlation risk) should yield lower returns. To this end, we broadly follow the most recent literature (Lustig and Verdelhan, 2007, Lustig, Roussanov, and Verdelhan, 2010; Menkhoff, Sarno, Schmeling, and Schrimpf, 2011; and Burnside, 2011) and sort currencies into portfolios according to their correlation risk exposure at the end of each month. We then form four such portfolios and invest in the portfolio with the

lowest relative correlation risk exposure (i.e. a portfolio consisting of currencies that provide a good hedge against correlation) while shorting the portfolio with the highest correlation risk exposure. We find that high interest rate currencies are negatively related to correlation risk and deliver lower returns in bad states of the world, i.e. when correlation is high, whereas low interest rate currencies (safe havens) provide a good hedge by yielding positive returns. The spread portfolio which is long the low correlation risk (high correlation beta) currencies and short the high correlation risk currencies leads to significant average excess returns of 4% and to attractive Sharpe ratios over the period from 1999 to 2011. In a next step, we estimate the cross-sectional price of correlation risk. Following the two step procedure by Fama and MacBeth (1973), we find a negative price of correlation risk of almost 1% per annum. These results are consistent with the notion that the cross-section of currency returns reflect exposure to correlation risk and that the price of global correlation risk is significantly negative. When we add other risk factors, the results remain essentially unchanged.

To address some of our empirical findings, we explore the implications of time variation in conditional risk aversion for currency risk premia. For that purpose, we consider a multi-country, multi-good general equilibrium model in which preferences are characterized by external habit formation (of the Menzly, Santos and Veronesi, 2004 type) and home bias, as in Stathopoulos (2011a, 2011b). In the model, currency risk premia compensate investors mainly for exposure to global risk aversion fluctuations: the price of the global risk aversion risk factor is negative, so agents are willing to accept lower returns for assets that have negative global risk aversion betas and, thus, provide a hedge against increases in global risk aversion. We also show that global risk aversion is positively associated with conditional exchange rate variances and covariances, so currencies that hedge against adverse global risk aversion fluctuations can be empirically identified as currencies that appreciate when conditional exchange rate second moments are high. Finally, we show that our model is able to link currency risk premia to real interest rate differentials and, thus, address the forward premium puzzle. If real interest rates are pro-cyclical, which is true if the precautionary savings motive is sufficiently strong, being long a high interest rate currency and short a low interest rate currency (i.e. engaging

in the carry trade) entails holding a position with a negative global risk aversion beta and, thus, very high exposure to the global risk aversion factor. As a result, investors require a very high compensation in terms of expected return in order to engage in the carry trade.

Related Literature:

This paper builds on the extant literature on the risk-return relationship of excess returns in currency markets. Lustig, Roussanov, and Verdelhan (2011) identify two new risk factors: The average forward discount of the US dollar against developed market currencies and the return to the carry trade portfolio itself. They then study the predictive content of these two factors and find that the average forward discount is the best predictor of average currency excess returns even when controlling for the forward discount. The paper closest to ours is by Menkhoff, Sarno, Schmeling, and Schrimpf (2011) who study whether currency excess returns can be explained by a compensation for global currency volatility risk. They find that high interest rate currencies are negatively related to innovations in global FX volatility and thus deliver low returns in times of unexpectedly high volatility, when low interest rate currencies provide a hedge by yielding positive returns. Della Corte, Sarno, and Tsiakas (2011) study the predictive power of forward implied volatility for spot volatility in foreign exchange markets. They find that the forward implied volatility is a biased predictor that overestimates movements in future spot implied volatility.

In contrast to the aforementioned papers, which focus solely on volatility risk, we expand our focus to correlation risk. Our intuition is as follows: If investments in currencies with high interest rates deliver low returns during "bad times" (i.e. when correlation is high), then carry trade profits are merely a compensation for higher risk-exposure by investors. Safe haven currencies (good hedges in bad states) yield lower returns than other currencies on average. This is in line with our finding that high interest-rate currencies are negatively related to innovations in correlation and thus deliver low returns in times of unexpectedly high correlation, when low interest rate currencies provide a hedge by yielding positive returns.

Our paper is part of the recent literature that addresses the failure of the expectations hypothesis for exchange rates. Brunnermeier, Nagel, and Pedersen (2009), Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2009), Jurek (2009), Burnside, Eichenbaum, Kleshchelski and Rebelo (2011), and Farhi and Gabaix (2011) emphasize the importance of disaster risk for currency risk premia, Yu (2011) studies the effect of investor sentiment, while Colacito and Croce (2009, 2010) and Bansal and Shaliastovich (2011) focus on long-run risks. The paper closest to ours is Verdelhan (2010): he proposes a two-country, single-good model with trade frictions and time-varying risk aversion generated by external habit formation and illustrates the importance of procyclical real interest rates for addressing the forward premium puzzle. In our multi-country model, we endogenize consumption and focus on the relationship between global risk aversion and conditional exchange rate variance and covariance.

Finally, different versions of our theoretical setup have been used to address the Brandt, Cochrane and Santa-Clara (2006) international risk sharing puzzle (Stathopoulos, 2011a) and the portfolio home bias puzzle (Stathopoulos, 2011b); our paper extends the insights of that work by considering the effects of time variation in risk aversion for currency returns.

The remainder of the paper is organized as follows. Section I. describes the data and details how we construct the correlation and variance risk premia. Section II. describes how we build the currency portfolios contains the empirical results with regards to the priced correlation risk in currency markets. Section III. sets up a multi country model with external habit and Section IV. concludes. All tables and figures are in the Appendix. A separate Online Appendix reports additional results.

I. Data and Risk Premium Construction

We start by describing the data and how to construct the correlation and variance risk premia. We use daily option prices on the most heavily traded currency pairs and high frequency data on the underlying spot exchange rates to construct the two risk

premia. Risk-neutral variances and correlations are calculated using the option prices while we use the underlying spot rates to construct the realized counterparts.

A. Data Description

High Frequency Currency Data:

The high frequency spot exchange rates EURUSD, JPYUSD, GBPUSD, and CHFUSD are from Olsen & Associates and cover the period from January 1990 to February 2011. Given that the high liquidity in foreign exchange markets prevents triangular arbitrage opportunities in the five most heavily traded currencies, calculating the remaining cross rates using the four exchange rates is common practice. The raw data contains all interbank bid and ask indicative quotes for the exchange rates for the nearest even second. After filtering the data for outliers, the log price at each 5 minute tick is obtained by linearly interpolating from the average of the log bid and log ask quotes for the two closest ticks. As options are traded continuously throughout the day, this results in a total of 288 observations over a 24 hour period.[2]

Currency Option Data:

We use over-the-counter (OTC) currency options data from JP Morgan for the four currency pairs EURUSD, JPYUSD, GBPUSD, and CHFUSD plus the six cross rates (i.e. we have have options data on a total of ten exchange rates). The frequency is daily and the data runs from January 1998 to February 2011. The use of OTC option data has several advantages over exchange traded option data. First, the trading volume in the OTC FX options market is several times larger than the corresponding volume on exchanges such as the Chicago Mercantile Exchange. As a direct consequence, this leads to more competitive quotes in the OTC market. Second, the conventions for writing and quoting options in the OTC markets have several features that are appealing when performing empirical studies: Every day, new option series with fixed times to maturity and fixed strike prices, defined by sticky deltas, are issued. In comparison, the time to maturity of exchange-traded option series gradually declines with the approaching

[2]We follow the empirical literature and take five minute intervals opposed to higher frequencies to mitigate the effect of spurious serial correlation due to microstructure noise (see Andersen and Bollerslev, 1998).

expiration date, and the moneyness continually changes as the underlying exchange rate moves. Therefore, the OTC option data allows for better comparability over time, as the series' main characteristics do not change from day to day. The options used in this study are plain-vanilla European calls and puts and encompass 35 option series per exchange rate. We consider a total of seven maturities: one week, one, two, three, six, and nine months, and one year. For each of the maturities, there are five different strikes available: (forward) at-the-money (ATM), 10-delta call and 25-delta call, 10-delta-put and 25-delta put.

Spot and Forward Rates:

To form our portfolios, we use daily data for spot exchange rates and one month forward rates versus the US dollar obtained from Datastream. We start from daily data in order to construct the correlation risk exposure. In line with the previous literature (see Fama, 1984), we work with the log spot and forward exchange rates, denoted as $s_t^i = \ln(S_t^i)$ and $f_t^i = \ln(F_t^i)$, respectively. We use the US dollar as the home currency and thus the superscript i always denotes the foreign currency. Our total sample consists of the same 35 countries as in Lustig, Roussanov, and Verdelhan (2011).

Portfolio Construction, HML_{FX} and DOL Factors:

At the end of each period t, we allocate currencies into four portfolios based on their forward discounts at the end of period t. Sorting on forward discounts is the same as to sorting on interest rate differentials since covered interest parity holds closely in the data at the frequency analyzed in this paper. We re-balance portfolios at the end of each month. This is repeated month by month. Currencies are ranked from low to high interest rate differentials. Portfolio 1 contains currencies with the lowest interest rate (or smallest forward discounts) and portfolio 4 contains currencies with the highest interest rates (or largest forward discounts). Monthly excess returns for holding foreign currency k, say, are computed as:

$$rx_{t+1}^k \approx f_t^k - s_{t+1}^k.$$

We follow Lustig, Roussanov, and Verdelhan (2011) and build a long-short factor based on carry trade portfolios (HML_{FX}). We also build a zero-cost dollar portfolio (DOL),

which is an equally weighted average of the different currency portfolios, i.e. the average return of a strategy that consists of borrowing money in the U.S. and investing in the global money markets outside the U.S.

[Insert Table 1 approximately here.]

Summary statistics of the carry trade, HML_{FX}, and DOL factor are presented in Table 1. Inline with previous findings, there is a monotonic increase from the lowest to the highest forward discount sorted portfolio. The unconditional average excess return from holding an equally weighted average carry portfolio is 1% per annum similar to the 2% reported in Menkoff, Sarno, Schmeling, Schrimpf (2011). The HML_{FX} portfolio is highly profitable with an average return of 22% and a Sharpe ratio of 1.7.

B. Construction of Variance and Correlation Risk Premia

The variance and correlation risk premium are defined as the difference between the risk-neutral and physical expectations of the variance and correlation, respectively. Thus, $\text{VRP}^i_{t,T}$, the $(T-t)$-period variance risk premium for the log exchange rate s^i at time t is defined as:

$$\text{VRP}^i_{t,T} \equiv E^{\mathbb{Q}}_t\left(\int_t^T \left(\sigma^i_u\right)^2 du\right) - E^{\mathbb{P}}_t\left(\int_t^T \left(\sigma^i_u\right)^2 du\right) \tag{1}$$

In a similar vein, the expression for the correlation risk premium between exchange rates S^i and S^j, $\text{CRP}^{i,j}_{t,T}$, is defined as:

$$\begin{aligned}\text{CRP}^{i,j}_{t,T} &\equiv E^{\mathbb{Q}}_t\left(\int_t^T \rho^{i,j}_u du\right) - E^{\mathbb{P}}_t\left(\int_t^T \rho^{i,j}_u du\right) \\ &\equiv \frac{E^{\mathbb{Q}}_t\left(\int_t^T \gamma^{i,j}_u du\right)}{\sqrt{E^{\mathbb{Q}}_t\left(\int_t^T \left(\sigma^i_u\right)^2 du\right)}\sqrt{E^{\mathbb{Q}}_t\left(\int_t^T \left(\sigma^j_u\right)^2 du\right)}} \\ &\quad - \frac{E^{\mathbb{P}}_t\left(\int_t^T \gamma^{i,j}_u du\right)}{\sqrt{E^{\mathbb{P}}_t\left(\int_t^T \left(\sigma^i_u\right)^2 du\right)}\sqrt{E^{\mathbb{P}}_t\left(\int_t^T \left(\sigma^j_u\right)^2 du\right)}},\end{aligned} \tag{2}$$

where $\rho_t^{i,j}$ and $\gamma_t^{i,j}$ are the conditional correlation and covariance between the two exchange rates, respectively.

Realized Variance and Correlation:

Currencies are traded continuously throughout the day all over the world. To match the time when we measure the daily option prices, we record the spot exchange rate at 4pm GMT. We thus have 288 intra-day currency returns over five minute intervals:

$$r_{k,5\text{min}} = \ln(S_k) - \ln(S_{k-5\text{min}}).$$

We follow Andersen, Bollerslev, Diebold, and Labys (2000) and compute the realized variance by summing the squared 5-minute frequency returns over the day:[3]

$$RV_t = \sum_{k=1}^{K} r_{k,5\text{min}}^2.$$

In a similar spirit, we derive the realized covariance between exchange rates S_i and S_j, respectively:

$$RCov_t^{i,j} = \sum_{k=1}^{K} r_{k,5\text{min}}^i r_{k,5\text{min}}^j.$$

The realized correlation is then simply the ratio between the realized covariance and the product of the respective standard deviations:

$$RCorr_t^{i,j} = RCov_t^{i,j} / \sqrt{RV_t^i} \sqrt{RV_t^j}.$$

We use realized measures observable at time t to proxy for the expectation under the physical measure for the period $T - t$.

Implied Variance and Correlation:

We follow Demeterfi, Derman, Kamal, and Zhou (1999) and Britten-Jones and Neuberger (2000) to obtain a model-free measure of implied volatility. The authors show that if the underlying asset price is continuous, the risk-neutral expectation of total

[3]We also use more refined measures of realized variance for robustness checks. The results, however, are not sensitive to how we measure the expected variance under the physical probability.

return variance is defined as an integral of option prices over an infinite range of strike prices:

$$E_t^{\mathbb{Q}}\left(\int_t^T \sigma_u^2 du\right) = 2e^{rT}\left(\int_0^{S_t} \frac{1}{K^2}\mathrm{P}(K,T)dK + \int_{S_t}^{\infty} \frac{1}{K^2}\mathrm{C}(K,T)dK\right), \tag{3}$$

where S_t is the underlying spot exchange rate and $\mathrm{P}(K,T)$ and $\mathrm{C}(K,T)$ are the put and call prices with maturity date T and strike K, respectively. Since, in practice, the number of traded options for any underlying asset is finite, the available strike price series is a finite sequence. Suppose the available strike prices of the call options belong to $\left[\underline{K}^{\mathrm{C}}, \overline{K}^{\mathrm{C}}\right]$, where $\overline{K}^{\mathrm{C}} \geq \underline{K}^{\mathrm{C}} \geq 0$. We use the trapezoidal rule to numerically calculate the integral:

$$2\int_{\underline{K}^{\mathrm{C}}}^{\overline{K}^{\mathrm{C}}} \frac{\mathrm{C}\left(K,T\right) - \max\left(0, S_t - K\right)}{K^2}dK \approx \frac{\overline{K}^{\mathrm{C}} - \underline{K}^{\mathrm{C}}}{m}\sum_{k=1}^{m}\left[g_{t,T}\left(K_k^{\mathrm{C}}\right) + g_{t,T}\left(K_{k-1}^{\mathrm{C}}\right)\right], \tag{4}$$

where

$$g_{t,T}\left(K_k^{\mathrm{C}}\right) = \frac{\mathrm{C}\left(K_k^{\mathrm{C}}, T\right) - \max\left(0, S_t - K_k^{\mathrm{C}}\right)}{\left(K_k^{\mathrm{C}}\right)^2}, \tag{5}$$

and K_k^{C} is the k^{th} largest strike price for the call option. The corresponding spot rates are extracted from the high-frequency dataset in accordance with the exact time of the daily options quote, i.e. at 4pm GMT. The risk-free interest rates for USD, EUR, JPY, GBP, and CHF are represented by the London Interbank Offered Rates (LIBOR). Jiang and Tian (2005) report two types of implementation errors: (i) Truncation errors due to the non availability of an infinite range of strike prices and (ii) discretization errors due to the fact that there is no continuum of options available. We find that both errors are extremely small using currency options. For example, the size of the errors totals only half a percentage point in volatility.

Model-free implied correlations are constructed in the spirit of Carr and Madan (1999). To this end, we assume that there exist three possible pairings of three currencies, S_t^i, S_t^j, and S_t^{ij}. The absence of triangular arbitrage then yields:

$$S_t^{ij} = S_t^i / S_t^j.$$

Taking logs, we derive the following relationship:

$$\ln\left(\frac{S_T^{ij}}{S_t^{ij}}\right) = \ln\left(\frac{S_T^{i}}{S_t^{i}}\right) - \ln\left(\frac{S_T^{j}}{S_t^{j}}\right).$$

Finally, taking variances yields:

$$\int_t^T \left(\sigma_u^{ij}\right)^2 du = \int_t^T \left(\sigma_u^{i}\right)^2 du + \int_t^T \left(\sigma_u^{j}\right)^2 du - 2\int_t^T \gamma_u^{i,j} du,$$

where $\gamma_t^{i,j}$ denotes the realized covariance of returns between currency pairs S_t^i and S_t^j. Solving for the covariance term, we get:

$$\int_t^T \gamma_u^{i,j} du = \frac{1}{2}\int_t^T \left(\sigma_u^{i}\right)^2 du + \frac{1}{2}\int_t^T \left(\sigma_u^{j}\right)^2 ds - \frac{1}{2}\int_t^T \left(\sigma_u^{ij}\right)^2 du.$$

Using the standard replication arguments, we find that:

$$\begin{aligned} E_t^{\mathbb{Q}}\left(\int_t^T \gamma_u^{i,j} du\right) \;=\; & e^{rT}\Bigg(\int_t^{S_t^i} \frac{1}{K^2}\mathrm{P}^i(K,T)dK + \int_{S_t^i}^{\infty} \frac{1}{K^2}\mathrm{C}^i(K,T)dK \\ & + \int_t^{S_t^j} \frac{1}{K^2}\mathrm{P}^j(K,T)dK + \int_{S_t^j}^{\infty} \frac{1}{K^2}\mathrm{C}^j(K,T)dK \\ & - \int_t^{S_t^{ij}} \frac{1}{K^2}\mathrm{P}^{ij}(K,T)dK - \int_{S_t^{ij}}^{\infty} \frac{1}{K^2}\mathrm{C}^{ij}(K,T)dK\Bigg). \end{aligned} \tag{6}$$

The model-free implied correlation can then be calculated using expression (6) and the model-free implied variance expression (3):

$$E_t^{\mathbb{Q}}\left(\int_t^T \rho_u^{i,j} du\right) \equiv \frac{E_t^{\mathbb{Q}}\left(\int_t^T \gamma_u^{i,j} ds\right)}{\sqrt{E_t^{\mathbb{Q}}\left(\int_t^T \left(\sigma_u^{i}\right)^2 du\right)}\sqrt{E_t^{\mathbb{Q}}\left(\int_t^T \left(\sigma_u^{j}\right)^2 du\right)}}. \tag{7}$$

Tables 2 and 3 provide summary statistics of the realized and implied volatility and correlation measures together with their corresponding risk premia while Figure 3 plots the realized and implied correlation measures. On average, implied volatility exceeds realized volatility for the Euro and the Japanese Yen but realized volatility is larger on average than the implied counterpart for the British Pound and the Swiss Franc.

Variance risk premia are small on average and statistically not different from zero.[4] We also note that the variance risk premia are extremely left skewed, which echoes the findings in Brunnermeier, Nagel, and Pedersen (2009) that carry trades are subject to crash risk.

[Insert Figure 3 and Tables 2 and 3 approximately here.]

Implied correlations exceed realized correlation on average for all currencies. The correlation risk premium is economically significant. The average correlation risk premium is 15%. The correlation risk premium is positively skewed on average but much more persistent than the variance risk premium with a first order autocorrelation coefficient of 0.96.

Figure 3 reveals that both realized and implied correlation remain quite stable until 2006 but then suddenly drop for most currency pairs. It is also interesting to note that the conditional correlations are mostly positive except for the safe haven currencies. One feature that is common to all currency pairs is that correlations exhibit very high volatility during the recent financial crisis.

The summary statistics of the risk premia are reported in Panel C of Tables 2 and 3. The variance risk premia are not statistically different from zero due to the high volatility of the series themselves. Variance risk premia in currency markets are switching sign quite often and on average the variance risk premia are positive only 60% of the time. In contrast, correlation risk premia are mostly positive and economically large: The average correlation risk premium is 15%, which is comparable to what is observed in the equity market.[5]

Global Variance and Correlation Risk:

To construct our global volatility and correlation risk factors, we average implied volatility and correlation over all different currency pairs at any given day. As explained above, we have options on EUR, GBP, JPY and CHF versus the USD and currency

[4] See also recent evidence in Chernov, Graveline, and Zviadadze (2011).

[5] Driessen, Maenhout, and Vilkov (2009) report an equity correlation risk premium of around 18%.

options for the ten cross pairs of the five currencies. This allows us to calculate four implied volatility and correlation measures. Hence,

$$\mathrm{IC}_{t,T}^{G} = 1/6 \times \sum_{k=1}^{6} \mathrm{IC}_{t,T}^{k}, \tag{8}$$

and

$$\mathrm{IV}_{t,T}^{G} = 1/4 \times \sum_{i=1}^{4} \mathrm{IV}_{t,T}^{i}, \tag{9}$$

where IC_t^k is the implied correlation for currency pair $k = i, j$ and IV_t^i is the implied volatility for exchange rate i. The two Hodrick and Prescott filtered series are plotted in Figure 4.

[Insert Figure 4 approximately here.]

The global volatility factor shows a distinct spike in the most recent financial crisis where volatility increased from almost 0 to 0.15. The global correlation factor is more volatile with distinct spikes also at the recent financial crisis but also the crisis in 2001. Overall, the volatility factor shows little movement except for the most recent years whereas the correlation risk factor changes its sign quite often.

For our empirical analysis, we use innovations in the correlation and variance risk premia, denoted ΔCorr and ΔVol, respectively. We construct the innovations by estimating an AR(1) process for the correlation and variance risk premia and taking the residuals from this regression. As an alternative way to construct the non-traded correlation and variance risk proxy we could take first differences from the corresponding level time series. Due to the high autocorrelation in the two series however, we prefer the first method.[6]

[6]We run portfolio sorts with both first differences and the AR(1) innovations and find that the results remain robust to the chosen method.

II. Empirical Analysis

In this section, we study the empirical relation between the global correlation risk proxy and the risk-return profile of currency portfolios. The previous section has demonstrated the large correlation risk premium in currency markets. If correlation risk is indeed priced in currency markets, then sorting currencies according to their exposure to correlation risk should yield a significant spread in average returns. We start by sorting a cross-section of different currencies according to their exposure to correlation risk and variance risk. We then ask whether our proposed risk factor can significantly explain carry trade portfolios. To this end, we run time series regressions of each portfolio's excess return on a set of potential risk factors:

$$z_{it} = \alpha_i + f_t' \beta_i + \epsilon_{it},$$

where f_t' constitute the matrix of risk factors. Using these factor betas, we assess the price of correlation risk using the two stage methodology from Fama and MacBeth (1973).

A. Correlation Risk Sorted Portfolios

We first construct monthly portfolios sorted according to the correlation and variance risk exposure. Intuitively, we expect those currencies to yield lower returns that hedge well against correlation risk, whereas we expect currencies that have only weak comovement with the correlation risk premium to yield high returns.

At the end of each period t, we build four currency portfolios based on the correlation (variance) risk exposure of the respective currencies. We estimate pre-ranking betas from rolling regressions of currency excess returns on the global correlation (variance) risk using 36 month windows (as in Lustig, Roussanov, and Verdelhan, 2010 and Menkhoff, Sarno, Schmeling, and Schrimpf, 2011):

$$rx_{t+1}^i = \alpha + \beta_{\mathrm{IC}^G}^i \mathrm{IC}_t^G + \epsilon_t^i,$$

where rx^i_{t+1} is the one month excess return of currency i, defined as $rx^i_{t+1} \equiv f^i_t - s^i_{t+1}$ and IC^G denotes innovations in the correlation risk factor. We repeat the same regression using the variance risk premium factor. Then, following Fama and French (1992), we estimate the betas as the sum of the slopes, i.e. $\beta_{\text{IC}^G} = \sum_{k=1}^{K} \hat{\beta}^k_{\text{IC}^G}$ and $\beta_{\text{IV}^G} = \sum_{k=1}^{K} \hat{\beta}^k_{\text{IV}^G}$. Descriptive portfolio statistics are reported in Table 4.

[Insert Table 4 approximately here.]

In Panel A of Table 4 we report summary statistics for the correlation risk sorted currency portfolios and in Panel B we report the statistics for the variance portfolios. The numbers speak for themselves: Correlation sorted portfolios yield attractive Sharpe ratios of 0.66 on average. We find that investing in currencies with high correlation betas leads to significantly lower returns compared to investing in low correlation beta currencies. Longing low correlation beta currencies and shorting high correlation beta currencies yields an average return of nearly 4% and an annualized Sharpe ratio of 0.4. When we move to Panel B, we find that compensation for volatility risk in terms of return risk trade off is very similar to the one of correlation risk. Low and high volatility beta currencies both yield average returns of between 2% and 10%. There is, however, no monotonic behavior between the low and high volatility exposure currencies. The long/short portfolio yields an average excess return of 3.6% with a Sharpe ratio of 0.4.

B. Factor Mimicking Portfolios

The portfolio sorting exercise has provided some evidence that global correlation risk is priced in the cross-section of currency returns, in a next step, we assess the cross-sectional price of correlation risk. To this end, we estimate a factor premium, λ_{FIC} (λ_{FIV}) on the mimicking correlation (volatility) factor, denoted by FIC (FIV). Following Ang, Hodrick, Xing, and Zhang (2006), we construct a factor-mimicking portfolio of correlation and volatility risk innovations. This allows us to naturally asses the factor prices of correlation and volatility risk vis-à-vis other factors.

To this end, we regress innovations in the global correlation and volatility risk proxies on the four excess carry return portfolios:

$$\mathrm{IC}_t^G = c + b'\mathbf{rx}_t + u_t,$$

where $\mathbf{rx}_t$ is the vector of excess returns. The factor mimicking portfolio excess return is then the product of the estimated slope coefficients and the excess returns, i.e. $FIC \equiv \hat{b}'\mathbf{rx}_t$. Again, we repeat the analysis using the global variance risk factor, FIV.

In the first step of the Fama and MacBeth (1973) regressions, we estimate betas using the full sample, in the second stage, we use the cross-sectional regressions to estimate the factor premia. Panel B of Table 5 shows the premia. The price of the carry trade risk factor is positive, inline with previous findings. In contrast, the price of correlation risk is -0.63% and statistically significant. Using the global volatility risk factor in the second regression, we estimate a cross-sectional price of volatility of 0.62% while the price of correlation risk is -0.48%. The negative factor prices are in line with our previous findings that portfolios, which co-move positively with correlation innovations require lower risk premia. The question then is which portfolios provide a good hedge against correlation and variance risk? To this end, we estimate the factor betas for the different currency portfolios. The results are reported in Table 5, Panel A.

[Insert Table 5 approximately here.]

Low interest rate currencies have a high correlation beta and thus provide a good hedge against correlation risk as they have a large negative exposure to correlation risk. Currencies with low correlation betas have a have a large positive exposure to correlation risk and thus command a correlation risk premium, which is reflected in higher interest rate. Furthermore, the estimated coefficients are highly significant for the dollar factor (DOL), which is not surprising given previous results in Lustig, Roussanov, and Verdelhan (2011).

III. Model

A. *Setup of the Model*

The world economy comprises $n+1$ countries, indexed by i: the domestic country ($i = 0$) and n foreign countries ($i = 1, ..., n$), each of which is populated by a single representative agent. There are $n+1$ distinct perishable goods in the world economy, indexed by j: the domestic good ($j = 0$) and n foreign goods ($j = 1, ..., n$). Uncertainty in the economy is represented by a filtered probability space $(\Omega, \mathcal{F}, \mathbf{F}, P)$, where $\mathbf{F}=\{\mathcal{F}_t\}$ is the filtration generated by the standard m-dimensional Brownian motion $\mathbf{B}_t$, $t \in [0, \infty)$, augmented by the null sets.

Each agent i is initially endowed with a claim on the entirety of the world endowment of the corresponding good ($j = i$). The endowment stream of good j is denoted by $\{\widetilde{X}_t^j\}$; all endowment processes are Itô processes satisfying:

$$d \log \widetilde{X}_t^j = \mu_t^{j,X} dt + \sigma_t^{j,X\prime} \mathbf{dB}_t,\ j = 0, 1, ..., n$$

with $\sigma_t^{j,X} \neq \mathbf{0}$ for all j. We can fix a global numéraire and denote the numéraire price of each good by Q_t^j. Since all goods are frictionlessly traded internationally, the price of each good, in units of the global numéraire, is the same in all countries.

Preferences:

The representative agent i has expected discounted utility:

$$E_0 \left[\int_0^\infty e^{-\rho t} \log(C_t^i - H_t^i) dt \right],$$

where $\rho > 0$ is her subjective discount rate, $X^{i,j}$ is the quantity of good j that agent i consumes,

$$C^i \equiv \left(\prod_{j=0}^{n} \left(X^{i,j}\right)^{a^{i,j}} \right)$$

is the domestic consumption basket and H^i is the habit level associated with consumption basket C^i. Notably, preferences are not symmetric regarding goods; the preferences

of agent i with respect to the $n+1$ goods are described by the vector of preference parameters $\mathbf{a}^i = [a^{i,0}, a^{i,1}, ..., a^{1,n}]$, such that $\sum_{j=0}^{n} a^{i,j} = 1$ and $a^{i,j} > 0$ for all i and j.

The external habit is of the Menzly, Santos and Veronesi (2004) form. Specifically, the inverse surplus consumption ratio $G^i = \left(\frac{C^i - H^i}{C^i}\right)^{-1}$ solves the stochastic differential equation:

$$dG_t^i = \varphi\left(\bar{G} - G_t^i\right) dt - \delta\left(G_t^i - l\right)\left(\frac{dC_t^i}{C_t^i} - E_t\left(\frac{dC_t^i}{C_t^i}\right)\right)$$

The inverse surplus consumption ratio is a stationary process, reverting to its long-run mean of $\bar{G}$ at speed φ and is driven by consumption growth shocks. The parameter $\delta > 0$ scales the impact of a consumption growth shock and the parameter $l \geq 1$ is the lower bound of the inverse surplus ratio G_t. The local curvature of the utility function is given by:

$$-\frac{u_{CC}(C_t^i, H_t^i)}{u_C(C_t^i, H_t^i)} C_t^i = -\frac{-\frac{1}{\left(C_t^i - H_t^i\right)^2}}{\frac{1}{C_t^i - H_t^i}} C_t^i = \frac{1}{C_t^i - H_t^i} C_t^i = G_t^i.$$

For that reason, in a slight abuse of terminology, in the remainder of this paper G^i will be referred to as the conditional risk aversion of country i.

B. Prices and Real Exchange Rates

The price of the country i consumption basket C^i (in units of the global numéraire) is:

$$P_t^i = \prod_{j=0}^{n} \left(\frac{Q_t^j}{a^{i,j}}\right)^{a^{i,j}}$$

and is defined as the minimum expenditure required to buy a unit of the consumption basket C^i.

The time t real exchange rate S^i (for all foreign countries $i = 1, ..., n$) is the price of the domestic consumption basket expressed in units of the consumption basket of country i:

$$S_t^i = \frac{P_t^0}{P_t^i} = \prod_{j=0}^{n} \left(\frac{\left(a^{i,j}\right)^{i,j}}{\left(a^{0,j}\right)^{0,j}}\right) \prod_{j=0}^{n} \left(Q_t^j\right)^{a^{0,j} - a^{i,j}}$$

so an increase of S^i denotes real appreciation of the domestic currency. The real exchange rate of country i is constant and equal to 1 (purchasing power parity) only if the two countries' preferences are identical ($a^{i,j} = a^{0,j}$ for all goods j), so that the two consumption baskets have the same composition. In the case of differences in preferences, purchasing power parity is violated.

Agents can frictionlessly trade $m+1$ non-redundant securities; of those, m are risky and 1 is locally riskless in terms of the global numéraire. The risky asset returns, in units of the global numéraire good, are given by the m-dimensional process

$$\mathbf{dR}_t = \mu_t dt + \sigma_t \mathbf{dB}_t$$

where μ_t is the $m \times 1$ vector of expected returns and σ_t is a non-singular $m \times m$ asset return diffusion matrix. Securities markets are dynamically complete. The world economy is a securities market economy, so the aggregate dividend of the risky assets equals the world endowment for each good and each state.

The price of the locally riskless asset, in units of the global numéraire, is denoted by D_t and satisfies

$$dD_t = r_t^f D_t dt$$

where r_t^f is the continuously compounded numéraire riskless rate. Then, we can define excess returns as

$$\mathbf{dR}_t^e = \mu_t^e dt + \sigma_t \mathbf{dB}_t$$

where $\mu_t^e = \mu_t - r_t^f \mathbf{1}$, with $\mathbf{1}$ representing a $m \times 1$ vector of ones.

The Agent's Problem:

Let the $m \times 1$ vector π_t^i describe the investment decision of agent i, with each element of the vector describing the amount, in units of the global numéraire, that agent i invests in each risky asset in period t. Thus, agent i chooses consumption shares $X_t^{i,j}$ and amounts π_t^i so as to maximize her expected discounted utility

$$\max_{\left\{X_t^{i,j}, \pi_t^i\right\}} E_0 \left[\int_0^\infty e^{-\rho t} \log\left(C_t^i - H_t^i\right) dt \right]$$

subject to the intertemporal flow budget constraint:

$$dW_t^i = \pi_t^{i\prime} \left(\mu_t^e dt + \sigma_t \mathbf{dB}_t \right) + W_t^i r_t^f dt - C_t^i P_t^i dt$$

where W_t^i is the period t wealth of agent i in units of the global numéraire; consequently, the investment of agent i in the riskless asset is $W_t^i - \pi_t^{i\prime}\mathbf{1}$.

C. Equilibrium

We can now solve for the competitive equilibrium. The solution details can be found in Stathopoulos (2011b). The competitive equilibrium is equivalent to the solution of the planner's problem:

$$\max_{\left\{X_t^{i,j}\right\}} E_0 \left[\int_0^\infty e^{-\rho t} \left(\sum_{i=0}^n \mu^i \log \left(C_t^i - H_t^i \right) \right) dt \right],$$

subject to the family of resource constraints:

$$\widetilde{X}_t^j = \sum_{i=0}^n X_t^{i,j} \text{ , for } j = 0, ..., n$$

for all t, where μ^i, $i = 0, ..., n$ is the welfare weight of country i. We normalize the welfare weights to sum to one: $\sum_{i=0}^n \mu^i = 1$.

The state-price density of the global numéraire satisfies:

$$\frac{\Lambda_t}{\Lambda_0} = e^{-\rho t} \mu^i \frac{G_t^i}{C_t^i} \frac{1}{P_t^i}.$$

Rearranging and summing over countries, we get:

$$\frac{\Lambda_t}{\Lambda_0} \left(\sum_{i=0}^n C_t^i P_t^i \right) = e^{-\rho t} \left(\sum_{i=0}^n \mu^i G_t^i \right).$$

Simplifying notation and rearranging, we can show that the global numéraire state-price density is increasing in global risk aversion and decreasing in global consumption expenditure:

$$\frac{\Lambda_t}{\Lambda_0} = e^{-\rho t}\frac{G_t^W}{\mathcal{C}_t^W},$$

where we define global conditional risk aversion G^W as the welfare-weighted average of all countries' conditional risk aversions,

$$G_t^W \equiv \sum_{i=0}^{n} \mu^i G_t^i$$

and the global consumption expenditure is:

$$\mathcal{C}_t^W \equiv \sum_{i=0}^{n} \left(C_t^i P_t^i\right).$$

D. Exchange Rates and Currency Risk Premia

We assume that the *local* numéraire in each country i is the local consumption basket C^i. We also assume that the *global* numéraire is the local numéraire of the domestic country, taken to be the United States. Thus, the global numéraire is the US consumption basket C^0, so $P_t^0 = 1$ for all t. The state-price density for the local numéraire of country i, Λ^i, is:

$$\Lambda_t^i = e^{-\rho t}\mu^i \frac{G_t^i}{C_t^i},$$

with law of motion

$$\frac{d\Lambda_t^i}{\Lambda_t^i} = -r_t^i dt - \eta_t^{i\prime}\mathbf{dB}_t,$$

where r_t^i is the real risk-free rate in country i and η^i is the market price of risk, both in units of the local numéraire of country i. Note that, since the global numéraire is the US local numéraire, it holds that $\Lambda \equiv \Lambda^0$. As a result, the global numéraire real risk-free rate is the US real risk-free rate ($r = r^0$) and the global numéraire market price of risk is the US market price of risk ($\eta = \eta^0$).

The real money market account in country i has price process D^i, expressed in units of the local numéraire, with law of motion:

$$dD_t^i = r_t^i D_t^i dt$$

while the country i market price of risk, which equals the conditional volatility of the stochastic discount factor (SDF) of country i, is:

$$\eta_t^i = \left(1 + \delta\left(\frac{G_t^i - l}{G_t^i}\right)\right)\sigma_t^{i,C}.$$

The conditional market price of risk has two components: the first component is increasing in conditional risk aversion, while the second component equals the conditional consumption growth volatility of country i. The first, risk aversion-related component, is increasing in conditional risk aversion G^i; in the absence of external habit formation, it would be constant and equal to one, the relative risk aversion implied by log utility. The second component, depends not only on the conditional risk aversion of country i, G^i, but also on the risk aversion of all other countries: as detailed in Stathopoulos (2011a), the conditionally less risk averse countries insure the conditionally more risk averse countries.

E. Currency Risk Premia

Since the domestic country is the United States, S^i denotes the real exchange rate of country i $(i = 1, ..., n)$ against the USD. It is easy to show that S^i equals the scaled ratio of two local numéraire state-price densities:

$$S_t^i = \frac{P_t^0}{P_t^i} = \frac{\mu^0}{\mu^i}\frac{\Lambda_t}{\Lambda_t^i}.$$

Under the objective probability measure ($\mathbb{P}-$measure), the real exchange rate satisfies the law of motion:

$$\frac{dS_t^i}{S_t^i} = \left[\left(r_t^i - r_t\right) + \eta_t^{i\prime}\left(\eta_t^i - \eta_t\right)\right]dt + \left(\eta_t^i - \eta_t\right)' \mathbf{dB}_t$$

Consider the portfolio that comprises a long position in the foreign money market account and a short position in the domestic money market account. The position of the domestic agent has a real USD value of $D_t^i/S_t^i - D_t^0$. The real USD return of this portfolio is

$$dR_t^{i,FX} = \frac{d(D_t^i/S_t^i)}{D_t^i/S_t^i} - \frac{dD_t^0}{D_t^0} = \eta_t' \left(\eta_t - \eta_t^i\right) dt + \left(\eta_t - \eta_t^i\right)' \mathbf{dB}_t$$

The currency risk premium is determined by the covariance of this return with the USD state-price density:

$$E_t\left(dR_t^{i,FX}\right) = -E_t\left(dR_t^{i,FX}\frac{d\Lambda_t}{\Lambda_t}\right) = \lambda_t^C \beta_t^{i,C} + \lambda_t^G \beta_t^{i,G}$$

where λ^C is the price of the exposure to global consumption expenditure shocks:

$$\lambda_t^C \equiv var_t\left(\frac{d\mathcal{C}_t^W}{\mathcal{C}_t^W}\right)$$

and λ^G is the price of the exposure to global risk aversion shocks:

$$\lambda_t^G \equiv -var_t\left(\frac{dG_t^W}{G_t^W}\right)$$

Importantly, the conditional betas of currency return i, $\beta^{i,C}$ and $\beta^{i,G}$ are time-varying.

F. Simulation

To build intuition for our results, we consider a global economy of three countries, the US and two foreign countries. The annualized parameters used for the simulation are presented in Table 6.

For all countries, the expected endowment growth rate is 1.5% per year, the standard deviation of the endowment growth rate is 8.12% per year and the correlation between any pair of endowment growth rates is 0.4. Since the preference parameter a is set equal to 0.004, 86% of each country's consumption expenditure corresponds to the purchase

of the locally endowed good and the balance is spent on foreign goods (0.4% on each of the foreign goods).

[Insert Table 6 approximately here.]

F.1. Exchange Rate Volatility and Correlation

We first consider the effects of an increase in all countries' conditional risk aversion. The results are presented in Figure 5. The horizontal axis represents the value of G_t^i $(i = 0, 1, 2)$ ranging from 20 to 150, imposing the restriction that all three countries have equal conditional risk aversion. We consider a symmetric parameter calibration, so all the moments of interest are identical for all three countries.

[Insert Figure 5 approximately here.]

Panels A, B, C and D of Figure 5 regard the variance and correlation of consumption growth rates and SDFs across countries. As global risk aversion increases, international risk sharing increases, sharply initially, more slowly afterwards: cross-country consumption growth and SDF correlations increase (Panels B and D, respectively). Furthermore, increased international risk sharing decreases the conditional variance of consumption growth rates (Panel A), but the reduction in consumption risk is not enough to balance the increase in the risk aversion-related component of the SDF, so conditional SDF variances increase (Panel C).

Panels E, F, G and H focus on the conditional exchange rate variances and covariances. We consider not only the two exchange rates with regard to the US dollar, S^1 and S^2, but also the exchange rate between the two foreign countries' currencies, denoted by $S^{1,2}$, where $S^{i,j}$ is defined as the price of the consumption basket of country j in units of the consumption basket of country i:

$$S_t^{i,j} \equiv \frac{P_t^j}{P_t^i} = \frac{S_t^i}{S_t^j}$$

Trivially, $S^{i,0} = S^i$ for all $i = 1, ..., n$. We first consider the conditional variance of real exchange rate changes. The conditional variance of changes in exchange rate $S^{i,j}$ is:

$$\left(\sigma_t^{i,j}\right)^2 \equiv \frac{1}{dt} var_t \left(\frac{dS_t^{i,j}}{S_t^{i,j}} \right) = \left(\eta_t^i - \eta_t^j\right)' \left(\eta_t^i - \eta_t^j\right)$$

We also adopt the notation $\left(\sigma_t^i\right)^2 \equiv \left(\sigma_t^{i,0}\right)^2$ for the conditional variance of changes in S^i. The conditional exchange rate variance equals the variance of the difference of the two countries' SDF; in that sense, the conditional variance of the real exchange rate measures the amount of risk not shared between the two countries. The amount of unshared risk can be decomposed into two parts: the amount of aggregate risk and the proportion of the aggregate risk that is not internationally shared. The amount of aggregate risk of countries i and j is defined as:

$$RP_t^{i,j} \equiv var_t \left(\frac{d\Lambda_t^i}{\Lambda_t^i} \right) + var_t \left(\frac{d\Lambda_t^j}{\Lambda_t^j} \right) = \eta_t^{i\prime} \eta_t^i + \eta_t^{j\prime} \eta_t^j$$

To measure the proportion of aggregate risk that is internationally shared, we use the Brandt, Cochrane and Santa-Clara (2006) measure:

$$RS_t^{i,j} = 1 - \frac{var_t \left(\frac{dS_t^{i,j}}{S_t^{i,j}} \right)}{var_t \left(\frac{d\Lambda_t^i}{\Lambda_t^i} \right) + var_t \left(\frac{d\Lambda_t^j}{\Lambda_t^j} \right)}$$

The international risk sharing index equals 1 if international risk sharing is perfect between countries i and j (and, thus, their bilateral exchange rate is constant), and 0 if there is no risk sharing between the two countries. Thus, conditional exchange rate volatility - which, as mentioned above, is the amount of unshared global risk - is the product of the proportion of aggregate risk not shared times the amount of aggregate risk:

$$\left(\sigma_t^{i,j}\right)^2 = (1 - RS_t^{i,j}) RP_t^{i,j}$$

Exchange rate covariance comprises three components:

$$\begin{aligned}\gamma_t^{i,j} &= \frac{1}{2}\left(\sigma_t^i\right)^2 + \frac{1}{2}\left(\sigma_t^j\right)^2 - \frac{1}{2}\left(\sigma_t^{i,j}\right)^2, \\ &= \frac{1}{2}(1 - RS_t^i)RP_t^i + \frac{1}{2}(1 - RS_t^j)RP_t^j - \frac{1}{2}(1 - RS_t^{i,j})RP_t^{i,j},\end{aligned}$$

so it depends on all bilateral risk pricing (RP^i, RP^j, $RP^{i,j}$) and risk sharing (RS^i, RS^j, $RS^{i,j}$) terms.

Since, as we see in Panel C, conditional SDF variance is increasing in global risk aversion, the risk pricing component of all exchange rates also increases when countries become more risk averse (Panel G). On the other hand, as seen in Panel D, more of the global risk is shared, leading to a reduction of unshared international risk (Panel H). The risk pricing component is the dominant one, overwhelming the effect of the declining risk sharing component, leading to an increase of the conditional variance of exchange rates (Panel E). Since all three exchange rate conditional variances are identical, the conditional exchange rate covariance $\gamma^{1,2}$ is also increasing in global risk aversion (Panel F).

While Figure 5 focused on the effects of an equal increase in all countries' conditional risk aversion, Figure 6 explores the effects of conditional risk aversion heterogeneity by considering the sensitivity of the variables of interest to changes in US conditional risk aversion G^0. Specifically, the horizontal axis of all Figure 2 panels represents the value of G_t^0, ranging from 20 (the lowest possible value) to 150. We fix the conditional risk aversion of the first foreign country, G_t^1, to 29 and the conditional risk aversion of the second foreign country, G_t^2, to 39, which constitutes 5 units below and above, respectively, their steady-state value of 34.

[Insert Figure 6 approximately here.]

An increase in domestic risk aversion is associated with more international risk sharing: consumption growth rates become more highly correlated internationally (Panel B), generating increasing SDF correlation as well (Panel D). However, the cross-sectional dispersion in countries' conditional risk aversion also introduces a cross-country insurance

effect. Panel A shows that, as domestic risk aversion increases, the global consumption risk is reallocated, with more risk assumed by the two foreign countries, and less risk assumed by the domestic country. Note that when domestic risk aversion is below foreign risk aversion, the foreign countries insure the domestic country, with the situation reversing when the domestic country is more risk averse than its trading partners. Furthermore, the first foreign country, which is the less risk averse of the two, is always characterized by more consumption growth volatility than the second, more risk averse, foreign country.

Panel C shows what happens to the conditional variance of the countries' SDF. As domestic risk aversion increases, there are two opposing effects on the domestic SDF variance: an increase in domestic risk aversion increases the price that the domestic country assigns to consumption risk (given by the first term of the SDF), but, as a result of international risk sharing, the domestic country assumes less consumption risk. The first effect always dominates, so domestic SDF variance is increasing in domestic risk aversion. Regarding the two foreign countries, the first term of the SDF is constant (since both countries' risk aversions are constant), but the insurance effect increases the amount of consumption risk they are undertaking as domestic risk aversion is raised, thus increasing the conditional variance of their SDF.

Panel E presents the conditional variance of real exchange rates S^1, S^2 and $S^{1,2}$. Although the conditional variance of the cross exchange rate, $\sigma_t^{1,2}$, monotonically falls as domestic risk aversion increases, this is not true for the conditional variance of S^1 and S^2: for very low and very high values of domestic risk aversion, the two conditional variances are decreasing in domestic risk aversion, but for intermediate values of G_t^0 they are increasing in G_t^0. The conditional exchange rate covariance has the same behavior as exchange rate variance, decreasing in domestic risk aversion for small values of G_t^0 and increasing in G_t^0 for larger values of domestic risk aversion (Panel F). To understand those results, we can consider the two components of exchange rate variance, the risk pricing component and the risk sharing component, presented in Panels G and H, respectively.

Panel G presents the behavior of the risk pricing component of all three exchange rates. As a result of the findings of Panel C, the risk pricing component of all three

exchange rates is increasing in domestic risk aversion. On the other hand, as implied by Panels B and D, the risk sharing component of all three exchange rates is declining in domestic risk aversion. Given that the risk pricing and the risk sharing effects are in conflict, the magnitude of the conditional variance of each exchange rate depends on the relative importance of the two opposing forces. Regarding S^1 and S^2, small increases in domestic risk aversion initially engender a very rapid increase in international risk sharing, so the decrease in the risk sharing term overwhelms the slower increase in the risk pricing term. At higher values of domestic risk aversion, however, international risk sharing becomes almost perfect, so the feasible gains from further increase of the domestic risk aversion are small and, thus, the risk pricing effect dominates, increasing exchange rate variance. For very high values of G_t^0, both effects are small, but the risk sharing effect dominates. Regarding the cross rate $S^{1,2}$, the risk sharing effect dominates the risk pricing effect, so exchange rate volatility declines in G_t^0.

F.2. Simulated Currency Portfolios

While the previous section has studied the effect of both global risk aversion and dispersion in risk aversion on the currency risk premia, here, we assess whether our model is able to replicate the features in the data and in particular of Table 4. We simulate a global economy of 36 countries, the domestic one and 35 foreign ones. The calibration is symmetric. Specifically, the endowment processes are specified as follows:

$$d \log \widetilde{X}_t^j = \mu dt + \sigma^j \mathbf{dB}_t, \; j = 0, 1, ..., 36$$

where σ^j is a 36×1 vector such that $\sigma^j[j, 1] = \sigma$ and $\sigma^j[j', 1] = \sigma\rho$ for all $j' \neq j$. Thus, all log endowment growth rates have constant and equal conditional first and second moments: their conditional mean equals μ and their conditional variance equals $\sigma(1 + 35\rho^2)$. Furthermore, the conditional correlation between any two log endowment growth rates is also constant and equal to $\frac{2\rho+34\rho^2}{1+35\rho^2}$. Regarding preferences, we set $a^{i,j} = a$ for $i \neq j$ and $a^{i,j} = 1 - 35a$ for $i = j$.

As outlined above, risk premia compensate investors for exposure to the global consumption expenditure risk factor and the global risk aversion risk factor. However, the price of the latter risk factor is an order of magnitude higher than the price of the former one, so the cross-section of currency returns largely mirrors the cross-section of global risk aversion betas $\beta^{i,G}$. Table 7 illustrates that point: We do a monthly sort of the 35 foreign currencies into 4 portfolios according to their conditional global risk aversion beta, with Portfolio 1 comprising long positions on currencies in the lowest $\beta^{i,G}$ quartile and Portfolio 4 containing the currencies in the highest$\beta^{i,G}$ quartile. As expected, there is a monotonic negative relationship between global risk aversion betas and average currency portfolio returns: the riskiest portfolio, Portfolio 1, which has the lowest risk aversion beta and, thus, the highest adverse exposure to the global risk aversion factor outperforms Portfolio 4, which provides the best hedge against global risk aversion, by about 37% in annual terms.

[Insert Table 7 approximately here.]

In the empirical part of our paper, we sorted currencies according to their exposure to the variance (IV^G) and correlation (IC^G) risk factors, which correspond to equally weighted indexes of conditional implied exchange rate variance and correlation, respectively. As shown in Figure 5, higher global risk aversion is associated with higher conditional exchange rate variance and covariance. Although this connection is not monotonic in the presence of cross-country heterogeneity in conditional risk aversion due to insurance effects, as illustrated in Figure 6, our simulation results show that global risk aversion is positively (and significantly) correlated with the average conditional exchange rate variance and covariance, which provides a justification for constructing portfolios sorted by their exposure to IV^G and IC^G, as we do in Table 4.

[Insert Table 8 approximately here.]

Finally, we explore the ability of our model to address the forward premium puzzle, as illustrated in Table 1. Table 8 reports the summary statistics on portfolios sorted on

interest rate differentials (forward discounts): Portfolio 1 comprises a long position on the currencies ranked in the bottom forward discount quartile (low interest rate currencies), while Portfolio 4 contains the high interest rate currencies. Notably, our model is able to match both the first and second moments of the carry trade portfolios. Given the results in Table 7, Table 8 implies that high (low) interest rate currencies have low (high) risk aversion betas, i.e. they depreciate (appreciate) in bad times, when global risk aversion is high. The negative relationship between the global risk aversion beta and the interest rate of a given currency is due to the fact that, given our calibration, real interest rates are pro-cyclical, due a very strong precautionary savings motive. As a result, a positive interest rate differential in favor of the foreign currency implies low foreign risk aversion and high domestic risk aversion, in relative terms. Given that the conditional SDF volatility of each country, thus, the global risk aversion beta of its currency, is positively associated with risk aversion, and that a long position in a foreign currency entails a long position in the foreign SDF and a short position in the domestic SDF, a negative relationship between global risk aversion betas and interest rates arises. Specifically, the high interest rate currency portfolio (Portfolio 4) has high average returns because it has a low (negative, to be precise) global risk aversion beta and, thus, is very exposed to global risk aversion risk, while the low interest rate portfolio (Portfolio 1) provides a good hedge against increases in global risk aversion and, thus, has a negative average return.

IV. Conclusion

We uncover new evidence of priced correlation risk in the cross section of foreign exchange risk premia. The correlation risk premium is economically large and highly time-varying. In contrast, volatility risk premia in currency markets are not statistically significant from zero. When we sort a broad cross-section of currency excess returns according to their exposure to correlation risk, we find that high interest rate currencies (investment currencies) co-move negatively with correlation risk whereas low interest rate currencies (funding currencies) provide a good hedge against correlation changes. A portfolio which is long low correlation exposure currencies and short high correlation

exposure currencies yields attractive returns and Sharpe ratios. We also show that the price of correlation risk is negative and cannot be subsumed by a carry trade factor.

We then proceed by setting up a multi country economy with imperfect risk sharing. We then identify a global risk aversion parameter, which is defined as the welfare weighted average of all countries' conditional risk aversions. We study two different cases. First, if there is no cross-sectional dispersion in conditional risk aversion, when risk aversion goes up homogeneously for all countries, unshared risk increases. The increase is non linear as countries' stochastic discount factors also become more correlated. In this case, we can show that exchange rate volatility and covariance increases as global risk aversion goes up. Second, when countries' risk aversions differ, there is an insurance effect at play. Less risk averse countries insure more risk averse countries. Given a country's risk aversion, the more insurance it provides, the higher its stochastic discount factor volatility will be. The opposite happens to the country that receives insurance: High cross-sectional dispersion means that the country's stochastic discount factor volatility falls, ceteris paribus. So, in the case where countries' risk aversion differ, exchange rate moments depend on not just the global risk aversion itself, but about the cross-sectional distribution of conditional risk aversions. However, in equilibrium, countries' risk aversions tend to not differ very much: as a result, the effect coming from the global risk aversion dominates.

References

Andersen, T., and T. Bollerslev (1998): "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, p. 885 – 905.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2000): "Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian," *Multinational Finance Journal*, 4, p. 159 – 179.

Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006): "The Cross-Section of Volatility and Expected Returns," *Journal of Finance*, 61, p. 259 – 299.

Bansal, R., and I. Shaliastovich (2011): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," Working Paper, Duke University.

Brandt, M., J. Cochrane, and P. Santa-Clara (2006): "International Risk Sharing Is Better Than You Think, or Exchange Rates are too Smooth," *Journal of Monetary Economics*, 53, p. 671 – 698.

Britten-Jones, M., and A. Neuberger (2000): "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839–866.

Brunnermeier, M., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," *NBER Macroeconomics Annual 2008*, 23, p. 313 – 347.

Buraschi, A., R. Kosowski, and F. Trojani (2011): "When There is No Place to Hide: Correlation Risk and the Cross-Section of Hedge Fund Returns," Working Paper, Imperial College London.

Burnside, C. (2011): "Do Peso Problems Explain the Returns to the Carry Trade?," *forthcoming, Review of Financial Studies*.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011): "Do Peso Problems Explain the Returns to the Carry Trade?," *Review of Financial Studies*, 24, p. 853 – 891.

Carr, P., and D. Madan (1999): "Currency Covariance Contracting," *Risk*, February, p. 47 – 51.

Chernov, M., J. Graveline, and I. Zviadadze (2011): "Sources of Risk in Currency Returns," Working Paper, London School of Economics.

Colacito, R., and M. M. Croce (2009): "Six Anomalies looking for a model: A consumption based explanation of International Finance Puzzles," Working Paper, UNC.

——— (2010): "Risk Sharing for the Long-Run: A General Equilibrium Approach to International Finance with Recursive Preferences and Long-Run Risks," Working Paper, UNC.

Corte, P. D., L. Sarno, and I. Tsiakas (2011): "Spot and Forward Volatility in Foreign Exchnage," *Journal of Financial Economics*, 100, p. 496 – 513.

Demeterfi, K., E. Derman, M. Kamal, and J. Zhou (1999): "A Guide to Volatility and Variance Swaps," *Journal of Derivatives*, 6, p. 9 – 32.

Driessen, J., P. Maenhout, and G. Vilkov (2009): "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance*, 64, p. 1377 – 1406.

Fama, E., and J. MacBeth (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, p. 607 – 636.

Fama, E. F. (1984): "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14, p. 319 – 338.

Fama, E. F., and K. R. French (1992): "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, p. 427 – 465.

Farhi, E., S. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009): "Crash Risk in Currency Markets," Working Paper, Harvard University.

Farhi, E., and X. Gabaix (2011): "Rare Disasters and Exchange Rates," Working Paper, Harvard University.

Jiang, G., and Y. Tian (2005): "Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies*, 18, p. 1305 – 1342.

Jurek, J. (2009): "Crash-Neutral Currency Carry Trades," Working Paper, Princeton University.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Countercylical Currency Risk Premia," Working Paper, UCLA.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2011): "Carry Trades and Global Foreign Exchnage Volatility," *forthcoming, Journal of Finance*.

Menzly, L., T. Santos, and P. Veronesi (2004): "Understanding Predictability," *Journal of Political Economy*, 112, p. 1 – 47.

Stathopoulos, A. (2011a): "Asset Prices and Risk Sharing in Open Economies," Working Paper, University of Southern California.

——— (2011b): "Portfolio Home Bias and External Habit Formation," Working Paper, University of Southern California.

Verdelhan, A. (2010): "A Habit-Based Explanation of the Exchange Rate Risk Premium," *Journal of Finance*, 65, p. 123 – 146.

Yu, J. (2011): "A Sentiment-Based Explanation of the Forward Premium Puzzle," Working Paper, University of Minnesota.

Appendix

Appendix A. Countries

We use the same 35 currencies as in Lustig, Roussanov, and Verdelhan (2011): Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Euro area, Finland, France, Germany, Greece, Hong Kong, Hungary, India, Indonesia, Ireland, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, United Kingdom.

Appendix B. Data

The exchange rate series consist of the last available mid-quote in each five-minute interval, resulting in 288 observations a day, as trading on the FX markets takes places 24h a day due to the fact that market opening hours rotate around the globe, from Asia to Europe to America to Asia and so on. My choice of these specific series has the following reasons:

1. Equally spaced intervals of five minutes: constructing an artificial time series of equally spaced observations, called "sampling in calender time", is the only time scale lending itself to multivariate applications. Choosing five-minute intervals is standard in empirical studies (see Andersen and Bollerslev, 1998). Although the theory of quadratic variation asks for the highest possible sampling frequency to yield the most accurate volatility measurement, returns measured at intervals shorter than five minutes heavily suffer from spurious serial correlation due to market-microstructure effects. These effects, as well as the optimal sampling frequency will be addressed below.

2. Last quote in every interval: in constructing the equally spaced time series, this so-called "previous-tick interpolation" only uses data available at this very point in time and hence respects causality. The increase in the number of quotes collected in the O&A database diminishes the average time interval between two consecutive quotes from 10-20 seconds in 1997 to less than a second in 2010. This gives a hint at how liquid the FX market is and how close to the five-minute time-stamps the last quotes in every interval are. Hence, problems arising from non-synchronous quotations or stale quotes are insignificant.

3. Mid-quotes (instead of transaction prices): as currency trading is mainly OTC, it is hard to get records of true transaction prices. However, Goodhart, Itô, and Payne (1995) and Danielsson and Payne (2001) verify by means of a short data sample from the Reuters 2000-2 electronic FX trading platform that the characteristics of the merely indicative quotes of the O&A data closely match the ones of the true transaction prices. In addition, using mid-quotes avoids contamination of the data series with market-microstructure noise arising from transactions being executed at either the quoted bid or ask price (bid-ask bounces), see e.g. Roll (1984).

To determine periods of inactive trading and hence possible problems with missing data, we find that a quote was recorded for each of the $7 \times 24 \times 12 = 2016$ five-minute intervals of a week. In the calculation of volatilities and correlations, the days are set to end at 16:00 GMT to match the time of the daily option quotes (all times in this study refer to Greenwich Mean

Time). Monday to Friday is treated as business days, while the reduced market activity on Saturdays and Sundays calls for a special treatment of the weekend period, which is defined as usual to last from Friday 21:00 to Sunday 21:00, see Andersen and Bollerslev (1998). Hence, "Mondays" consist of the remaining five hours on Friday from 16:00 to 21:00 and the 19 hours from Sunday 21:00 to Monday 16:00. In the 3% of cases that a quote is not available on a business day, the missing value is filled with the previous quote, which induces a slight bias towards zero in the realized (co-)volatility measures.

Appendix C. Realized Correlation and Volatility Sorting

In this subsection, we check whether our sorting procedure is also robust to realized measures of correlation and volatility. We apply the same procedure as in Section *A.*. The results are reported in Table 9. We note that the results remain virtually unchanged. Low exposure currencies yield higher returns than high exposure currencies.

Insert Table 9 approximately here.

Appendix D. Real Carry Returns

We proxy the real forward discount as the nominal forward discount minus the inflation differential. We retrieve monthly inflation rates from the International Labour Office (see http://laborsta.ilo.org/). We define the Euro inflation as the geometric average between Austria, Belgium, Netherlands, Finland, France, Germany, Ireland, Italy, Portugal and Spain. We use monthly data from January 1999 to June 2007 as after this date, many series have missing values. We then adjust the nominal forward discount by subtracting the inflation differential. Inflation differentials on average are very small and mostly not significantly different from zero: The average inflation differential is 0.1% with a standard deviation of 0.6%. The largest inflation differential has Indonesia with an average differential of 0.5%.

Insert Table 10 approximately here.

Appendix E. Proofs

Equilibrium Prices and Quantities:

Under the assumption of market completeness, there is a unique global numéraire state-price density, Λ, which satisfies the SDE

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t dt - \eta_t' \mathbf{dB}_t$$

where r is the global numéraire risk-free rate and η is the market price of risk process.

Using Λ, the intertemporal budget constraint of agent i can be written in static form as follows:

$$E_0\left[\int_0^\infty \frac{\Lambda_t}{\Lambda_0} C_t^i P_t^i dt\right] \leq E_0\left[\int_0^\infty \frac{\Lambda_t}{\Lambda_0} \widetilde{X}_t^i Q_t^i dt\right]$$

or

$$E_0\left[\int_0^\infty \frac{\Lambda_t}{\Lambda_0}\left(\sum_{j=0}^{n+1} X_t^{i,j} Q_t^j\right) dt\right] \leq E_0\left[\int_0^\infty \frac{\Lambda_t}{\Lambda_0}\widetilde{X}_t^i Q_t^i dt\right]$$

After replacing each agent's intertemporal dynamic budget constraint with her static budget constraint, we can solve for the competitive equilibrium. The first order conditions (FOCs) of agent i are:

$$e^{-\rho t} a^i \frac{G_t^i}{X_t^{i,j}} = \frac{1}{\mu^i}\frac{\Lambda_t}{\Lambda_0} Q_t^j, \text{ for all } j$$

where $\frac{1}{\mu^i}$ is the Lagrange multiplier associated with the budget constraint of agent i holding with equality. Combining the FOCs with the market clearing conditions:

$$\sum_{k=0}^{n} X_t^{i,j} = \tilde{X}_t^j, \text{ for all } j$$

we get the equilibrium consumption allocation:

$$X_t^{i,j} = \frac{a^{i,j}\mu^i G_t^i}{\sum_{k=0}^n a^{k,j}\mu^k G_t^k}\widetilde{X}_t^j$$

To calculate the Lagrange multipliers $\frac{1}{\mu^i}$, we substitute equilibrium quantities and prices in the static budget constraint of agent i (holding with equality). After some algebra, we get:

$$\mu^i\left(\varphi\bar{G} + \rho G_0^i\right) = \sum_{k=0}^{n} a^{k,i}\mu^k\left(\varphi\bar{G} + \rho G_0^k\right)$$

This system of equations has solutions of the form

$$\frac{\mu^i}{\mu^0} = b^i \frac{\varphi\bar{G} + \rho G_0^0}{\varphi\bar{G} + \rho G_0^i}$$

where the vector $\mathbf{b} = [b^1, b^2, ..., b^n]'$ is the unique solution of

$$\mathbf{b} = \begin{bmatrix} a^{0,1} & a^{1,1} & ... & a^{n,1} \\ a^{0,2} & a^{1,2} & ... & a^{n,2} \\ ... & ... & ... & ... \\ a^{0,n} & a^{1,n} & ... & a^{n,n} \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{b} \end{bmatrix}$$

The budget constraint determines only the ratios $\frac{\mu^i}{\mu^0}$. To pin down the values for the Lagrange multipliers, we impose the normalization $\sum_{i=0}^n \mu^i = 1$.

Equilibrium consumption processes:

Since equilibrium consumption $C = [C^0, ..., C^n]'$ is a function of the vector of conditional risk aversion $G = [G^0, ..., G^n]'$, we need to solve for the fixed point that satisfies both the

equilibrium consumption allocations and the law of motion for G. By the definition of the consumption baskets, we have:

$$C^i \equiv \left(\prod_{j=0}^{n} \left(X^{i,j} \right)^{a^{i,j}} \right), \text{ for all } i$$

so, applying Itô's lemma and equating the diffusion terms, we get, after some algebra:

$$\sigma_t^C = \left(\Psi_t^{-1} \mathbf{A} \right) \sigma_t^X$$

where σ_t^C is the $(n+1) \times m$ consumption volatility matrix

$$\sigma_t^C = \left[\begin{array}{c} \sigma_t^{0,C\prime} \\ ... \\ \sigma_t^{n,C\prime} \end{array} \right]$$

σ_t^X is the $(n+1) \times m$ endowment volatility matrix

$$\sigma_t^X = \left[\begin{array}{c} \sigma_t^{0,X\prime} \\ ... \\ \sigma_t^{n,X\prime} \end{array} \right]$$

and, finally, $\mathbf{\Psi}$ is the $(n+1) \times (n+1)$ matrix defined as

$$\Psi_t = [\psi_{i,j}] = \psi_t^{i-1,j-1}$$

where

$$\psi_t^{i,i} \equiv 1 + \left(1 - \sum_{j=0}^{n} \frac{a^{i,j} a^{i,j} \mu^i G_t^i}{\sum_{k=0}^{n} a^{k,j} \mu^k G_t^k} \right) \delta \left(\frac{G_t^i - l}{G_t^i} \right)$$

and

$$\psi_t^{i,i'} \equiv - \left(\sum_{j=0}^{n} \frac{a^{i,j} a^{i',j} \mu^{i'} G_t^{i'}}{\sum_{k=0}^{n} a^{k,j} \mu^k G_t^k} \right) \delta \left(\frac{G_t^{i'} - l}{G_t^{i'}} \right), \; i \neq i'$$

Table 1
Summary Statistics Carry Trade Portfolios

This table reports summary statistics for portfolios sorted on time $t-1$ forward discounts. We also report annualized Sharpe Ratios (SR) and the first order autocorrelation coefficient (AC(1)). Portfolio 1 contains 25% of all the currencies with the lowest forward discounts whereas Portfolio 4 contains currencies with the highest forward discounts. All returns are excess returns in USD. DOL denotes the average return of the four currency portfolios, HmL denotes a long-short portfolio that is long in Pf1 and short in Pf4. Data is sampled monthly and runs from January 1999 to February 2011.

	Pf1	Pf2	Pf3	Pf4	DOL	HmL
Mean	-0.1175	0.0202	0.0483	0.1043	0.0138	0.2218
StDev	0.1267	0.0649	0.0763	0.0988	0.0736	0.1300
Skew	-2.2605	-0.2741	-0.3402	-1.4752	-0.5363	1.3767
Kurtosis	8.6690	4.2208	5.8966	8.9847	3.9290	6.7569
SR	-0.9274	0.3112	0.6331	1.0551	0.1876	1.7058
AC(1)	0.7180	0.0625	0.1573	0.2619	0.3116	0.6625

Table 2
Summary Statistics Variance

This table reports summary statistics for implied and realized volatilities (i.e. the square root of variance, Panels A and B) and the variance risk premium, which is defined as the difference between the implied and realized variance (Panel C). Implied variances are calculated from daily option prices on the underlying exchange rates. Realized variances are calculated from five minute tick data on the underlying spot exchange rates. All numbers are annualized. Data is daily and runs from February 1998 to February 2011.

Panel A: Implied Volatility

	EUR	JPY	GBP	CHF
	EUR	JPY	GBP	CHF
Mean	0.1089	0.1199	0.0962	0.1111
StDev	0.0319	0.0379	0.0339	0.0250
Kurtosis	9.0464	8.5052	13.5283	8.5036
Skewness	1.8192	1.7933	2.8592	1.4667
Min	0.0485	0.0612	0.0482	0.0531
Max	0.3018	0.4039	0.3208	0.2637
AC(1)	0.9900	0.9800	0.9923	0.9826

Panel B: Realized Volatility

	EUR	JPY	GBP	CHF
Mean	0.1061	0.1148	0.0980	0.1145
StDev	0.0370	0.0502	0.0406	0.0356
Kurtosis	13.7172	33.3336	15.3902	13.5564
Skewness	2.1475	3.6054	2.7993	2.0435
Min	0.0344	0.0267	0.0254	0.0374
Max	0.5264	0.9162	0.4937	0.5038
AC(1)	0.7050	0.7377	0.8074	0.6074

Panel C: Variance Risk Premium

	EUR	JPY	GBP	CHF
Mean	0.0001	-0.0001	-0.0009	-0.0015
StDev	0.0084	0.0173	0.0082	0.0080
Kurtosis	267.2600	743.4876	168.3246	49.9990
Skewness	-10.5019	-21.3103	-9.1316	-4.8326
Min	-0.2484	-0.6764	-0.1975	-0.1188
Max	0.0416	0.0366	0.0412	0.0295
AC(1)	0.2713	0.3868	0.4237	0.3145

Table 3
Summary Statistics Correlation

This table reports summary statistics for implied and realized correlations (Panels A and B) and the correlation risk premium, which is the difference between the implied and realized correlation (Panel C). Implied correlations are calculated from daily option prices on the underlying exchange rates. Realized correlations are calculated from five minute tick data on the underlying spot exchange rates. Data is daily and runs from February 1998 to February 2011.

Panel A: Implied Correlation

	EURJPY	EURGBP	EURCHF	JPYGBP	JPYCHF	GBPCHF
Mean	0.3957	0.6742	0.8906	0.2966	0.4667	0.6252
StDev	0.1825	0.1072	0.0835	0.2045	0.1646	0.1466
Kurtosis	3.6678	2.6582	10.4716	3.3690	3.2698	2.3661
Skewness	-0.7809	-0.5134	-2.5060	-0.4875	-0.5750	-0.4432
Min	-0.4392	0.2559	0.4462	-0.5075	-0.1582	0.1952
Max	0.7622	0.8746	0.9887	0.7387	0.8257	0.8872
AC(1)	0.9832	0.9797	0.9909	0.9851	0.9813	0.9842

Panel B: Realized Correlation

	EURJPY	EURGBP	EURCHF	JPYGBP	JPYCHF	GBPCHF
Mean	0.2930	0.5238	0.7025	0.1858	0.2989	0.4463
StDev	0.2190	0.1564	0.1468	0.1944	0.1905	0.1574
Kurtosis	3.0352	2.8465	2.7149	2.7828	2.4153	2.4422
Skewness	-0.5820	-0.8063	-0.6756	-0.2130	-0.1400	-0.3117
Min	-0.4211	0.0970	0.2556	-0.3806	-0.2661	0.0487
Max	0.6613	0.7920	0.9401	0.6032	0.7176	0.7677
AC(1)	0.9980	0.9985	0.9968	0.9982	0.9981	0.9979

Panel C: Correlation Risk Premium

	EURJPY	EURGBP	EURCHF	JPYGBP	JPYCHF	GBPCHF
Mean	0.1027	0.1506	0.1883	0.1091	0.1678	0.1779
StDev	0.1349	0.1220	0.1434	0.1360	0.1189	0.1404
Kurtosis	4.6117	2.6784	2.7289	4.1417	3.2539	2.7480
Skewness	0.5151	0.2028	0.7345	-0.4918	0.4372	-0.3144
Min	-0.3767	-0.1853	-0.1133	-0.3971	-0.1579	-0.3251
Max	0.6850	0.4748	0.6656	0.5827	0.6010	0.5264
AC(1)	0.9618	0.9808	0.9929	0.9614	0.9569	0.9787

Table 4
Portfolios Sorted on Betas with Correlation and Variance Risk

This table reports summary statistics for portfolios sorted on correlation (variance) risk betas, i.e. currencies are sorted according to their betas in a rolling time-series regression of individual currencies' daily excess returns on daily innovations in the correlation (variance) risk. Correlation (variance) risk is defined as the residual from an AR(1) process of implied correlation (variance). Portfolio 1 contains currencies with the lowest betas whereas portfolio 4 contains currencies with the highest betas. LmH is long Portfolio 1 and short Portfolio 4. The mean, standard deviation, and Sharpe Ratios are annualized, the rest is per month. We also report pre-formation betas, Pre β. Data is monthly and runs from January 1999 to February 2011.

Panel A: Correlation Risk

	Pf1	Pf2	Pf3	Pf4	LmH
Mean	0.0890	0.0681	0.0534	0.0497	0.0393
StDev	0.1259	0.0978	0.0862	0.0769	0.0993
Skew	-1.1805	-0.5635	-0.0415	-0.1305	-0.3926
Kurtosis	6.3188	6.2227	3.0149	3.2757	3.4086
SR	0.7071	0.6968	0.6192	0.6456	0.3962
AC(1)	0.2209	0.1026	0.0573	0.1262	0.1477
Pre β	-26.814	-1.205	14.499	24.983	

Panel B: Variance Risk

	Pf1	Pf2	Pf3	Pf4	LmH
Mean	0.1082	0.0296	0.0543	0.0721	0.0361
StDev	0.1206	0.0936	0.0869	0.0842	0.0922
Skew	-1.0073	-0.7308	-0.8184	-0.2168	-1.2587
Kurtosis	6.1168	7.2339	4.3523	3.2091	10.3332
SR	0.8972	0.3167	0.6240	0.8562	0.3919
AC(1)	0.2355	0.0693	0.1742	0.0767	0.2071
Pre β	-29.779	-6.365	3.805	20.862	

Table 5
Estimating the Price of Correlation Risk

FIC and FIV are the mimicking factors for global correlation and volatility innovations, DOL the average carry trade portfolio as in Lustig, Roussanov, and Verdelhan (2011). In Panel A, we report factor betas. Panel B reports the Fama and MacBeth (1973) factor prices on the carry return portfolios. Newey-West standard errors are reported in parentheses. Data is monthly and runs from January 1999 to February 2011.

Panel A: Factor Betas

Pf	α	DOL	FIC	FIV	R^2
1	-0.01	0.96	4.71		0.92
	(0.01)	(0.59)	(0.05)		
	-0.01	1.04	4.99	8.49	0.94
	(0.01)	(0.04)	(0.40)	(1.50)	
2	-0.01	0.65	2.04		0.92
	(0.01)	(0.09)	(0.95)		
	-0.01	0.38	1.07	-29.33	0.99
	(0.01)	(0.01)	(0.12)	(0.45)	
3	-0.01	1.21	-1.98		0.90
	(0.01)	(0.06)	(0.73)		
	-0.01	1.32	-1.60	11.60	0.97
	(0.01)	(0.05)	(0.45)	(1.87)	
4	-0.01	0.83	-3.94		0.90
	(0.01)	(0.03)	(0.42)		
	-0.01	0.75	-4.23	-8.59	0.92
	(0.01)	(0.03)	(0.37)	(1.67)	

Panel B: Factor Prices

DOL	FIC	FIV	χ^2
0.163	-0.629		4.25
(0.02)	(0.00)		
0.172	-0.483	0.623	0.39
(0.02)	(0.00)	(0.00)	

Table 6
Choice of Parameter Values and Benchmark Values of State Variables

This table lists the parameter values used for all figures in the paper. All parameters are annualized. If not mentioned otherwise, we study a symmetric economy, where parameters for foreign countries are assumed to be the same.

PARAMETERS FOR ENDOWMENT		
Endowment expected growth rate	μ	0.015
Endowment volatility parameter	σ	0.07
Endowment correlation parameter	ρ	0.10
PREFERENCE PARAMETERS		
Preference parameter	α	0.004
Subjective rate of time preference	ρ	0.04
Speed of G mean reversion	k	0.05
G sensitivity to consumption growth shocks	δ	120
Lower bound of G	l	20
Steady-state value of G	$\bar{G}$	34

Table 7
Summary Statistics Simulated Conditional RA Sorted

This table reports summary statistics for portfolios sorted on global risk aversion betas using 500 simulations of 240 months for 35 currencies. We also report annualized Sharpe Ratios (SR) and the first order autocorrelation coefficient (AC(1)). Portfolio 1 contains 25% of all the currencies with the lowest global risk aversion beta whereas Portfolio 4 contains currencies with the highest global risk aversion beta. All returns are excess returns in USD. HmL denotes a long-short that is long in Pf1 and short in Pf4.

	Pf1	Pf2	Pf3	Pf4	HmL
Mean	0.1895	0.0906	-0.0189	-0.1778	0.3674
StDev	0.1347	0.1304	0.1355	0.1642	0.1476
Skew	0.1951	0.1700	0.1490	0.4251	-0.7685
Kurtosis	6.4318	6.9504	6.4408	5.6454	6.7157
SR	1.4340	0.7161	-0.1358	-1.1006	2.5620
AC(1)	0.0040	-0.0329	0.0019	0.0590	0.2318

Table 8
Summary Statistics Simulated Carry Trade Portfolios

This table reports summary statistics for simulated portfolios sorted on time $t-1$ forward discounts. We also report annualized Sharpe Ratios (SR) and the first order autocorrelation coefficient (AC(1)). Portfolio 1 contains 25% of all the currencies with the lowest forward discounts whereas Portfolio 4 contains currencies with the highest forward discounts. All returns are excess returns in USD. DOL denotes the average return of the four currency portfolios, HmL denotes a long-short that is long in Pf1 and short in Pf4.

	Pf1	Pf2	Pf3	Pf4	DOL	HmL
Mean	-0.1554	-0.0223	0.0752	0.1370	0.0086	0.2925
StDev	0.1429	0.1264	0.1270	0.1407	0.1212	0.1186
Skew	0.4643	0.3234	0.0158	-0.5081	0.0405	-1.2438
Kurtosis	6.7152	6.5610	6.3384	6.1254	6.0466	8.6005
SR	-1.1142	-0.1783	0.6126	1.0070	0.0812	2.5176
AC(1)	0.0440	0.0256	-0.0026	0.0294	-0.0118	0.2955

Table 9
Portfolios Sorted on Betas with Realized Correlation and Volatility Risk

This table reports summary statistics for portfolios sorted on correlation (variance) risk betas, i.e. currencies are sorted according to their betas in a rolling time-series regression of individual currencies' daily excess returns on daily innovations in the correlation (variance) risk. Correlation (variance) risk is defined as the residual from an AR(1) process of realized correlation (variance). Portfolio 1 contains currencies with the lowest betas whereas portfolio 4 contains currencies with the highest betas. LmH is long Portfolio 1 and short Portfolio 4. The mean, standard deviation, and Sharpe Ratios are annualized, the rest is per month. We also report pre-formation betas, Pre β. Data is monthly and runs from January 1999 to February 2011.

Panel A: Correlation Risk

	Pf1	Pf2	Pf3	Pf4	LmH
Mean	0.0815	0.0573	0.0228	0.0346	0.0469
StDev	0.0787	0.0903	0.1007	0.0938	0.0721
Skew	0.1255	-0.4209	-0.6100	-0.4753	0.5221
Kurtosis	3.4037	7.1746	5.4467	3.6043	3.6421
SR	1.0360	0.6351	0.2264	0.3690	0.6507
AC(1)	0.1629	0.1439	0.1967	0.0879	0.0791
Pre β	-104.529	-8.323	47.771	131.896	

Panel B: Variance Risk

	Pf1	Pf2	Pf3	Pf4	LmH
Mean	0.0920	0.0616	0.0595	0.0490	0.0431
StDev	0.1099	0.1078	0.0947	0.0754	0.0842
Skew	-0.2985	-1.5170	-0.2099	0.1411	-0.1717
Kurtosis	3.8256	9.0365	5.8127	3.4792	4.0753
SR	0.8375	0.5711	0.6278	0.6496	0.5114
AC(1)	0.1281	0.2813	0.0579	0.0127	0.1859
Pre β	-48.699	-31.298	-19.164	7.997	

Table 10
Summary Statistics Real Carry Trade Portfolios

This table reports summary statistics for portfolios sorted on time $t-1$ real forward discounts. We also report annualized Sharpe Ratios (SR) and the first order autocorrelation coefficient (AC(1)). Portfolio 1 contains 25% of all the currencies with the lowest forward discounts whereas Portfolio 4 contains currencies with the highest forward discounts. All returns are excess returns in USD. DOL denotes the average return of the four currency portfolios, HmL denotes a long-short that is long in Pf1 and short in Pf4. Data runs from January 1999 to June 2007.

	Pf1	Pf2	Pf3	Pf4	DOL	HmL
Mean	-0.1778	0.0074	0.0250	0.1640	0.0046	0.3419
StDev	0.1678	0.0743	0.0640	0.0803	0.0744	0.1624
Skew	-2.0982	-0.0791	-0.0894	0.3695	-0.4326	1.8471
Kurtosis	6.7164	3.3703	2.4553	3.1241	2.8587	5.9423
SR	-1.0601	0.0996	0.3898	2.0418	0.0622	2.1050
AC(1)	0.8035	0.1614	0.1253	0.2857	0.4538	0.7677

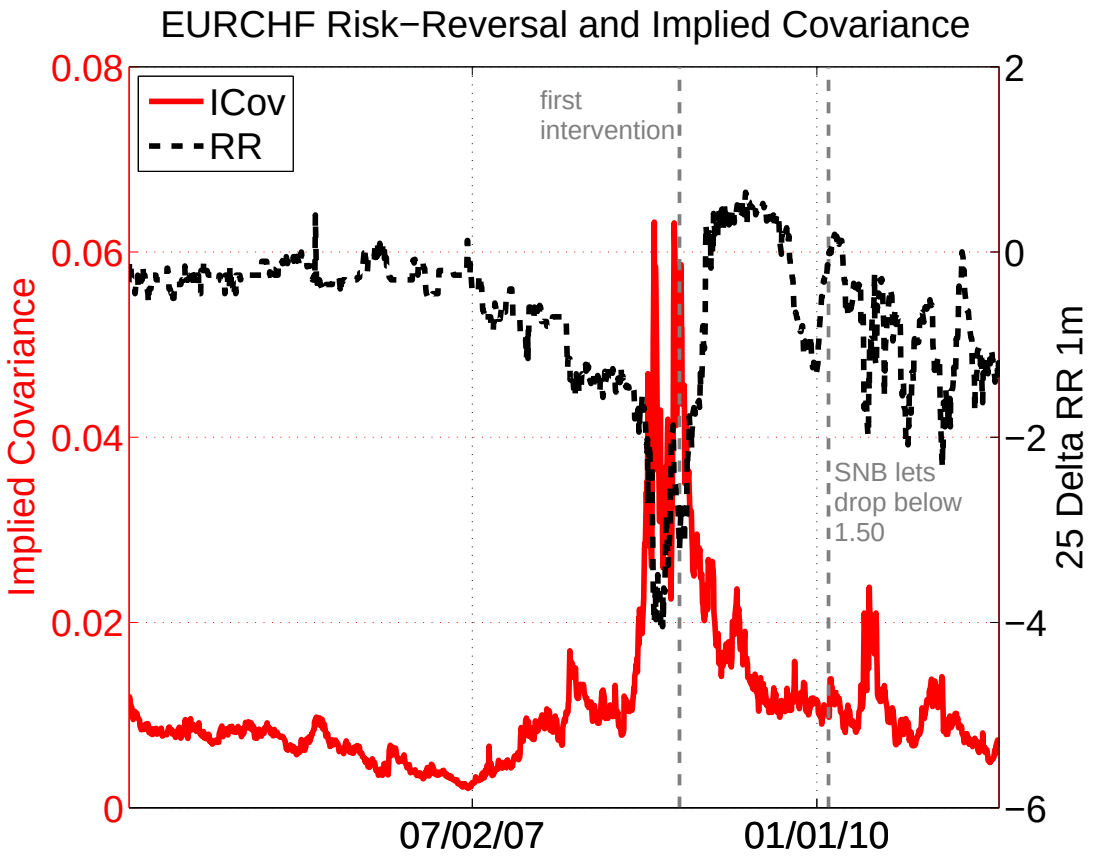


Figure 1. EURCHF Risk Reversal and Implied Covariance

This figure plots the 1 month risk reversal and implied covariance from daily options data on the EURCHF spot exchange rate. Risk reversal is defined as the difference between a 25% Delta call and a 25% Delta put option. The implied covariance is calculated using daily implied variance measures from options on the EURCHF, CHFUSD, and EURUSD spot exchange rates. Data is daily and runs from February 2007 to end of April 2011.

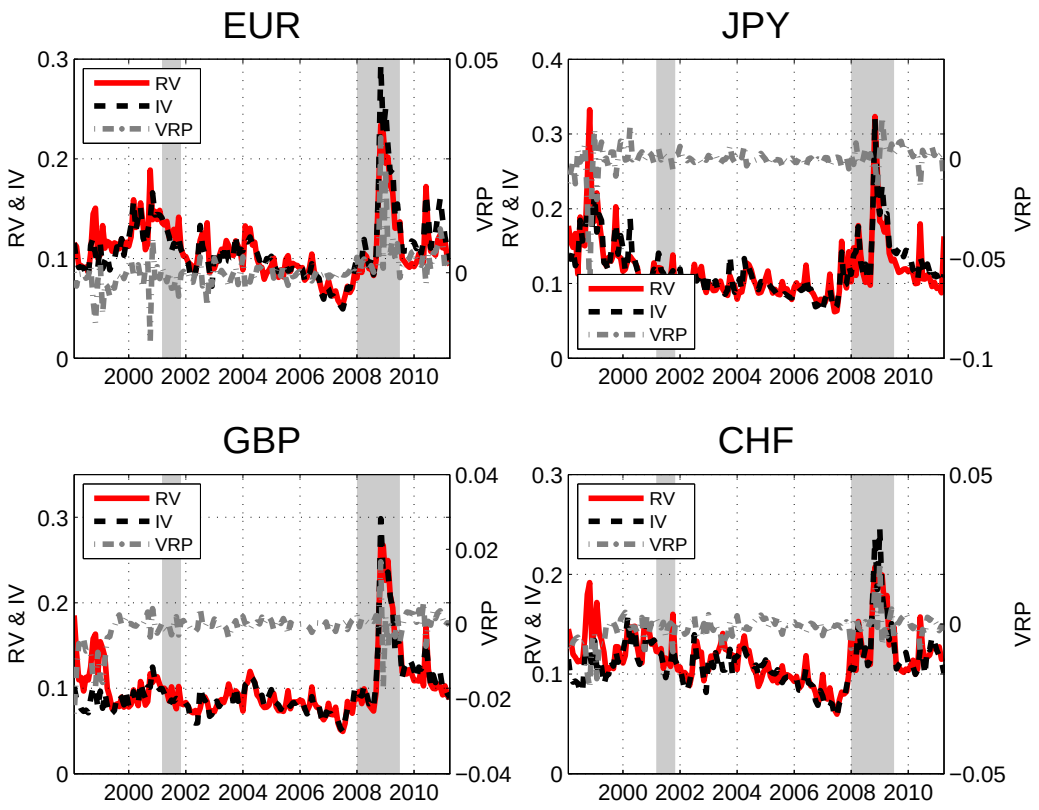


Figure 2. Implied and Realized Variance and Variance Risk Premium

This figure plots the daily time series of realized (left axis) and implied correlation (right axis) for six currency pairs. Realized and implied correlation are calculated from five minute tick data and daily option prices. All currencies are with respected to the USD. Gray shaded areas are recessions as defined by the NBER. Data runs from January 1999 to February 2011.

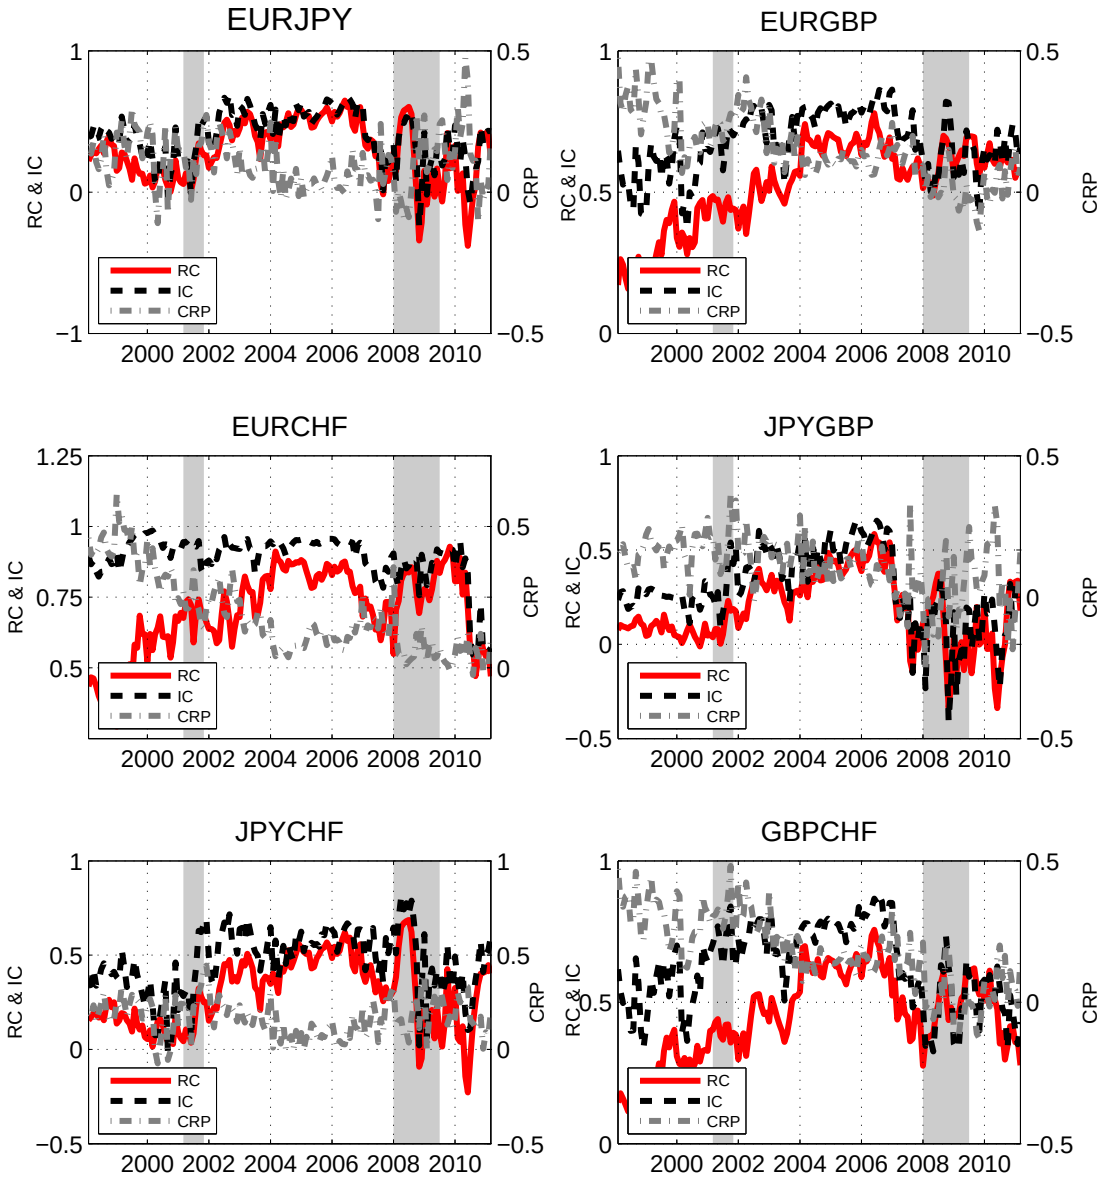


Figure 3. Implied and Realized Correlation

This figure plots the daily time series of realized (left axis) and implied correlation (right axis) for six currency pairs. Realized and implied correlation are calculated from five minute tick data and daily option prices. All currencies are with respected to the USD. Gray shaded areas are recessions as defined by the NBER. Data runs from January 1999 to February 2011.

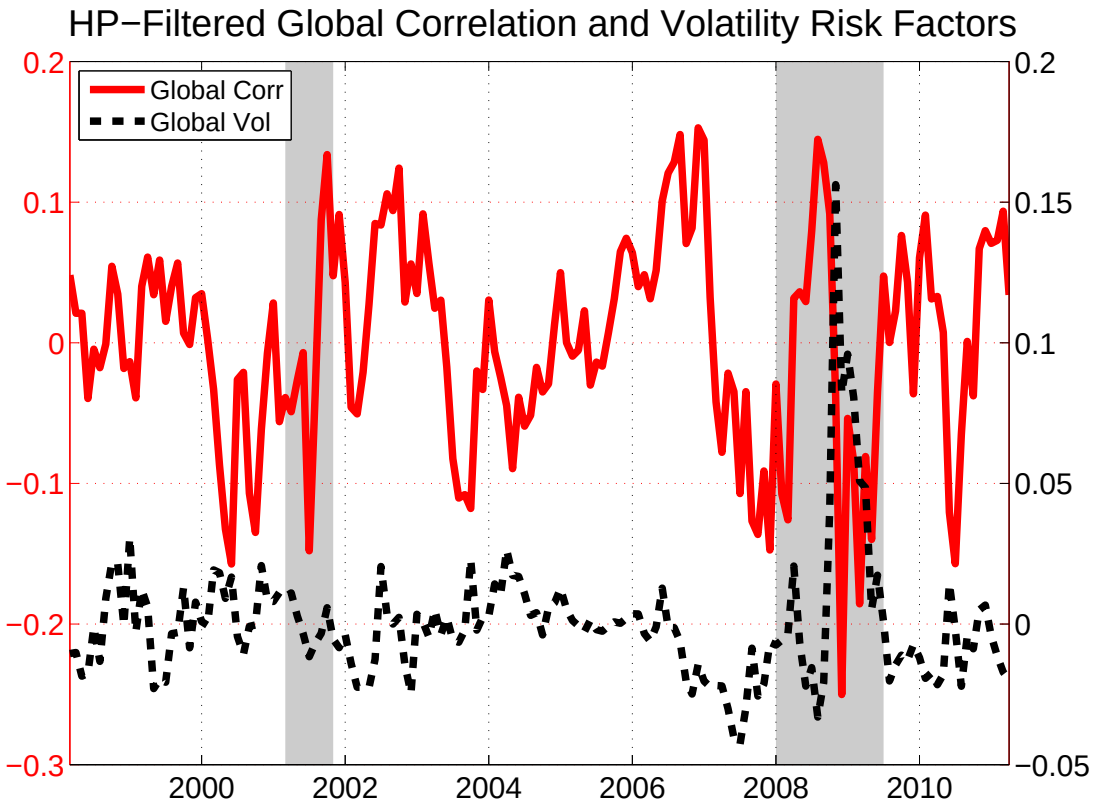


Figure 4. HP-Filtered Global Volatility and Correlation Risk

This figure plots the Hodrick and Prescott filtered volatility and correlation risk factor defined as an equal weighted average of individual implied volatility and correlation measures. Implied volatilities and correlations are calculated using daily options on different currency pairs. Gray shaded areas are recessions as defined by the NBER. Data runs from February 1998 to March 2011.

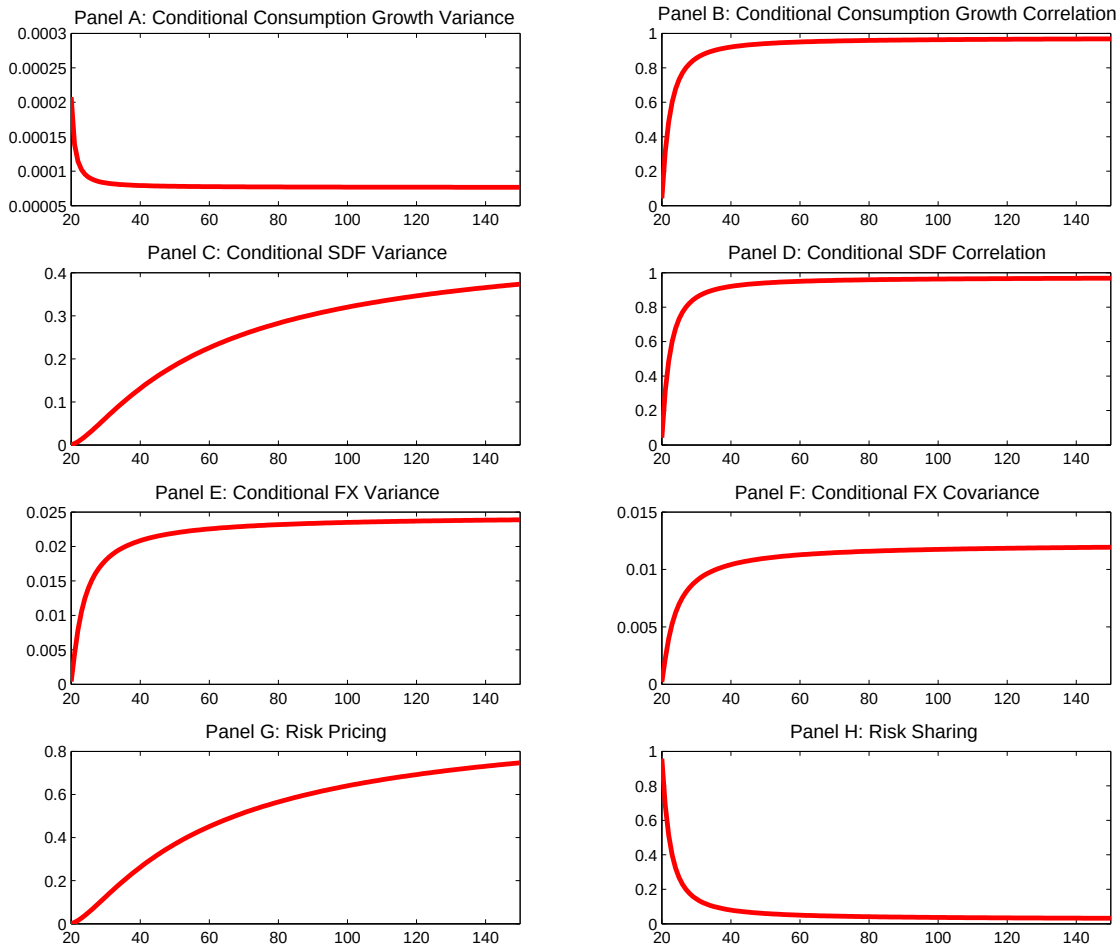


Figure 5. Conditional Moments and Countries' Conditional Risk Aversion

These figures plot the conditional consumption growth, SDF and exchange rate variance (Panels A to C) as a function of countries' conditional risk aversion. Panels B to F plot the conditional covariances and Panel G represents the risk pricing and Panel H the risk sharing which are defined in the main text.

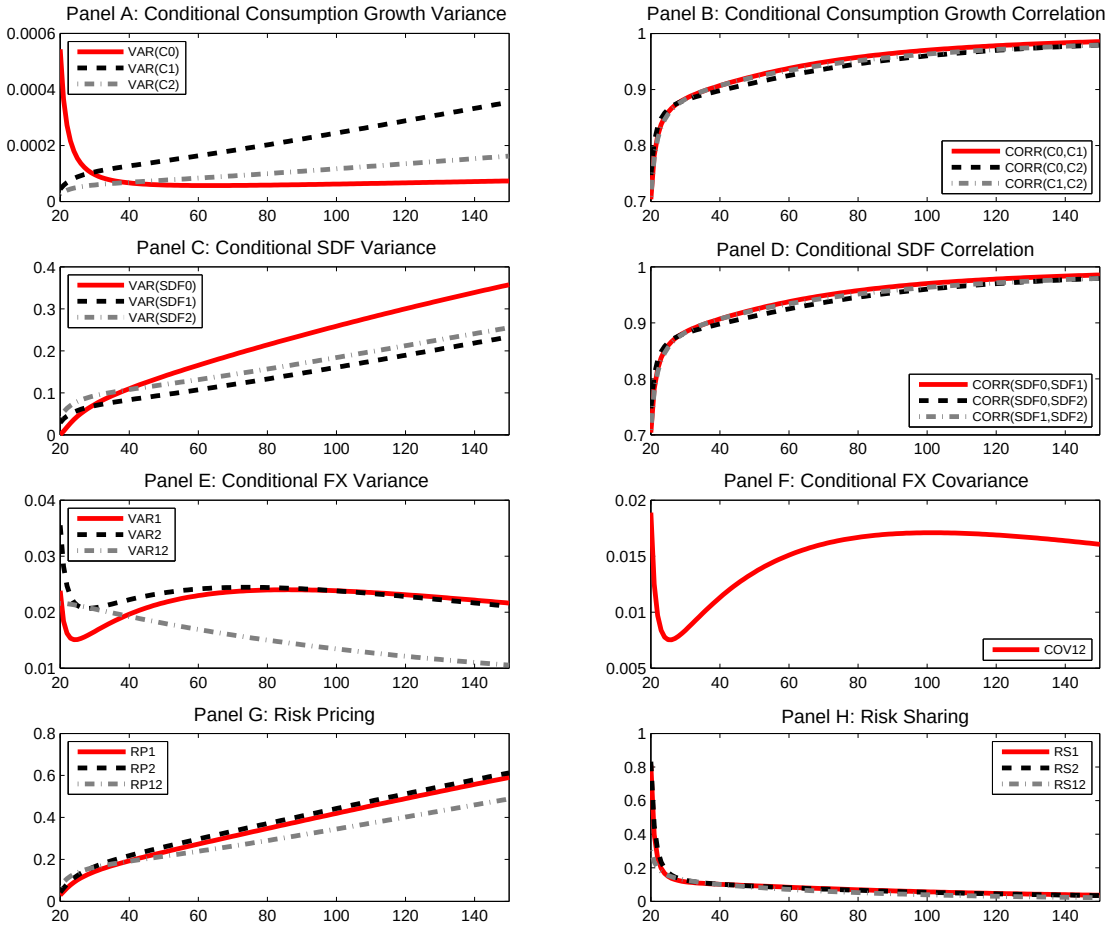


Figure 6. Conditional Moments and Countries' Dispersion in Conditional Risk Aversion

These figures plot the conditional consumption growth, SDF and exchange rate variance (Panels A to C) as a function of countries' dispersion in conditional risk aversion. Panels B to F plot the conditional covariances and Panel G represents the risk pricing and Panel H the risk sharing which are defined in the main text.