



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



International Correlation Risk

Philippe Mueller
Andreas Stathopoulos
Andrea Vedolin

SRC Discussion Paper No 26

December 2014



Systemic Risk Centre

Discussion Paper Series

Abstract

We document that cross-sectional FX correlation disparity is countercyclical, as exchange rate pairs with high average correlation become more correlated in bad times whereas pairs with low average correlation become less correlated. We show that currencies that perform badly (well) during periods of high cross-sectional disparity in conditional FX correlation yield high (low) average excess returns, suggesting that correlation risk is priced in currency markets. Furthermore, we find a negative cross-sectional relationship between average FX correlations and average FX correlation risk premia. Finally, we propose a no-arbitrage model that can match salient properties of FX correlations and correlation risk premia.

Keywords: Correlation risk, carry trade, international finance, exchange rates.

This paper is published as part of the Systemic Risk Centre's Discussion Paper Series. The support of the Economic and Social Research Council (ESRC) in funding the SRC is gratefully acknowledged [grant number ES/K002309/1].

Philippe Mueller, Assistant Professor of Finance, London School of Economics and Political Science, Research Associate of the Systemic Risk Centre.

Andreas Stathopoulos, Assistant Professor of Finance and Business Economics, USC Marshall School of Business.

Andrea Vedolin, Assistant Professor of Finance, London School of Economics and Political Science, Research Associate of the Systemic Risk Centre.

Published by
Systemic Risk Centre
The London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© Philippe Mueller, Andreas Stathopoulos, Andrea Vedolin submitted 2014

International Correlation Risk*

Philippe Mueller

London School of Economics[†]

Andreas Stathopoulos

USC Marshall School of Business[‡]

Andrea Vedolin

London School of Economics[§]

Abstract

We document that cross-sectional FX correlation disparity is countercyclical, as exchange rate pairs with high average correlation become more correlated in bad times whereas pairs with low average correlation become less correlated. We show that currencies that perform badly (well) during periods of high cross-sectional disparity in conditional FX correlation yield high (low) average excess returns, suggesting that correlation risk is priced in currency markets. Furthermore, we find a negative cross-sectional relationship between average FX correlations and average FX correlation risk premia. Finally, we propose a no-arbitrage model that can match salient properties of FX correlations and correlation risk premia.

Keywords: Correlation risk, carry trade, international finance, exchange rates

First Version: December 2011

This Version: November 2014

*Philippe Mueller and Andrea Vedolin acknowledge financial support from STICERD at LSE. We would like to thank Dante Amengual, Andrew Ang, Svetlana Bryzgalova, Joe Chen, Mike Chernov, Ram Chivukula, Robert Dittmar, Dobrislav Dobrev, Anh Le, Angelo Ranaldo, Paul Schneider, Ivan Shaliastovich, Adrien Verdelhan, Hao Zhou, and seminar and conference participants of LUISS Guido Carli, Rome, University of Lund, University of Piraeus, University of Bern, Stockholm School of Economics, Federal Reserve Board, University of Essex, London School of Economics, Ohio State University, Southern Methodist University, University of Pennsylvania (Wharton), Bank of England, Chicago Booth Finance Symposium 2011, 6th End of Year Meeting of the Swiss Economists Abroad, Duke/UNC Asset Pricing Conference, UCLA-USC Finance Day 2012, the Bank of Canada–Banco de España Workshop on “International Financial Markets”, Financial Econometrics Conference at the University of Toulouse, SED 2012, the 8th Asset Pricing Retreat at Cass Business School, the CEPR Summer Meeting 2012, and the EFA 2012 for thoughtful comments.

[†]Department of Finance, Houghton Street, WC2A 2AE London, UK, Email: p.mueller@lse.ac.uk

[‡]Marshall School of Business, Los Angeles, CA 90089, USA, Email: astathop@marshall.usc.edu

[§]Department of Finance, Houghton Street, WC2A 2AE London, UK, Email: a.vedolin@lse.ac.uk

The importance of correlation in financial markets has spawned a broad literature documenting that asset return correlations are stochastic and countercyclical. In particular, there is evidence from the equity market that correlation carries a significant risk premium, arguably due to the deterioration of investors' investment opportunities that results from a reduction in diversification benefits when asset return correlations increase. Yet, the existing literature has largely ignored the foreign exchange (FX) market. In this paper, we provide novel empirical evidence that correlation risk is priced in FX markets and we propose a reduced-form, no-arbitrage model that is consistent with our empirical findings.

Empirically, we start by documenting large cross-sectional differences in both average FX correlations and average FX correlation risk premia, defined as the difference between FX correlations under the risk-neutral and objective measures. Moreover, we show that there is a negative cross-sectional relationship between average FX correlations and average correlation risk premia: on average, FX pairs characterized by low average correlation exhibit high correlation risk premia whereas FX pairs that are highly correlated on average have low correlation risk premia.

We then explore the time series properties of FX correlations and FX correlation risk premia. First, we find a negative relationship between the average level of FX correlations and cyclicality. In particular, using several business cycle proxies, we show that currencies with high average correlations become more correlated in adverse economic times, whereas FX pairs with low average correlations become even less correlated. This can be best illustrated by focusing on the USD exchange rate of the Japanese Yen (JPY), a typical low interest rate currency, and the USD exchange rates of the Australian and the New Zealand Dollar (AUD and NZD), two high interest rate currencies. The average correlation between the JPY exchange rate and either the AUD or the NZD exchange rate is countercyclical and fairly low at 16% and 15%, respectively. On the other hand, the average correlation of the USD exchange rates of the two high interest rate currencies is 76% and procyclical. Thus, the cross-sectional dispersion of FX correlations widens in bad states of the world and tightens in good states of the world. Second, we show that there is a very strong negative time series relationship between FX correlations and FX correlation risk premia, both in levels and in changes, for virtually all FX

pairs. Furthermore, we find that FX pairs with high average correlation risk premia have countercyclical correlation risk premia, whereas pairs with low correlation risk premia have procyclical correlation risk premia. Thus, bad states amplify the magnitude of FX correlation risk premia, increasing their cross-sectional dispersion.

We exploit the countercyclical behavior of cross-sectional FX correlation dispersion by defining a novel FX correlation factor, FXC , which measures cross-sectional differences in the conditional FX correlations of the G10 currencies. To construct our factor, we sort FX pairs into deciles based on their conditional FX correlation and subtract the average conditional FX correlation of the bottom decile from the average conditional FX correlation of the top decile. As suggested by the simple three currency example, we find that the resulting factor is strongly countercyclical.

We use our FX correlation factor in order to quantify the compensation for exposure to FX correlation risk using a portfolio sorting approach, in line with the recent international finance literature (see, e.g., Lustig and Verdelhan (2007), Lustig, Roussanov, and Verdelhan (2011), Burnside (2011), and Menkhoff, Sarno, Schmeling, and Schrimpf (2012)). We find that correlation risk is priced in currency markets, yielding a negative price of risk: currencies with low FX correlation betas have higher average excess returns, whereas currencies that appreciate strongly when FX correlation differentials increase yield lower excess returns. Shorting the high correlation beta currencies and investing in the low correlation beta currencies generates an average annual excess return of between 4% (all countries) and 5.6% (developed countries) and Sharpe ratios of 0.46 and 0.59, respectively. Using various test assets, we estimate the price of FX correlation risk to be about minus 48 basis points (bps) per month. This means that shorting an asset with a beta of one would result in an annual excess return of 5.7%.

We rationalize our empirical findings with a no-arbitrage model of exchange rates. The main tension we address is between the dynamics of FX correlation in the physical and the risk-neutral measure. Under the physical measure, the FX correlation differential between high and low average correlation FX pairs is countercyclical. At the same time, high correlation exchange rate pairs on average have low or negative correlation risk premia whereas low average correlation pairs exhibit rather high correlation risk premia,

suggesting that U.S. investors attach significant prices to states in which the cross section of FX correlations tightens.

Our model builds on the work of Lustig, Roussanov, and Verdelhan (2011), Lustig, Roussanov, and Verdelhan (2014) and Verdelhan (2013). Their models feature both common (global) and country-specific (local) pricing factors. The former price purely global SDF shocks, common to all countries, whereas the latter price purely local SDF shocks, independent across countries. Importantly, they assume that innovations in the local pricing factors are uncorrelated across countries, as they are perfectly negatively correlated with the corresponding local shocks. They also assume that innovations in the global pricing factors are countercyclical, perfectly negatively correlated with the corresponding global shocks. In order to address the tension mentioned above we deviate in one key aspect from this setup: we allow the local pricing factors to also be exposed to global shocks, an assumption which generates comovement in local pricing factors across countries. Indeed, in our benchmark model we consider the polar case that all local pricing factors are solely exposed to global shocks and, thus, are identical across countries. In that case, there are only two pricing factors, common to all countries: a global one and a local one.

In our benchmark model, exchange rates are exposed to two kinds of shocks: (i) local, country-specific shocks, each of which is priced by the common local pricing factor, and (ii) a single global shock that is priced by the global pricing factor. Countries are assumed to have heterogeneous exposure to the global shock and identical exposure to their corresponding local shock, so the only source of heterogeneity across countries is the heterogeneous loading on the global innovation. The cross section of average FX correlations is determined by the cross section of exposures to the global shock: FX pairs that correspond to foreign countries with similar exposure to global risk (called similar FX pairs) are more correlated on average than FX pairs of countries with dissimilar global risk exposure (called dissimilar FX pairs).

The dynamics of conditional FX correlations in the physical measure depend on the relative importance of changes in the local and the global pricing factor. In particular, an increase of the local risk factor increases the variance of all exchange rates, as country-specific shocks are not offset across countries. As a result, when local shocks become

more highly priced, the cross section of conditional FX correlations tightens, with high correlation FX pairs becoming less correlated and low correlation FX pairs becoming more correlated. On the other hand, an increase in the price of the first global shock amplifies the significance of the exchange rate fluctuations that are driven by global risk exposure, increasing the correlation of similar FX pairs and decreasing the correlation of dissimilar FX pairs. Thus, when the price of global risk increases, the cross section of conditional FX correlations widens. As a result, for the dynamics of FX correlation in the model to replicate the empirical FX dynamics, with dispersion in conditional FX correlation being countercyclical, it must be the case that the dominant fluctuations in the physical measure are those of the global pricing factor, with changes in the local pricing factor having a second-order effect on conditional FX correlation. This provides the rationale for the construction of our empirical FX correlation factor FXC .

The key to resolving the tension between the objective and the risk-neutral measure lies in the different pricing of innovations in the two pricing factors by the domestic (United States) investor. In particular, the U.S. investor prices fluctuations of the local pricing factor more than those of the global pricing factor, thus attaching a higher price to states characterized by high value of the local pricing factor and low values of the global pricing factor. As we have seen, those are exactly the states characterized by low cross-sectional dispersion in FX correlations. Calibrating our model, we find that it generates realized FX correlations, implied FX correlations and FX correlation risk premia that match the joint time series and cross sectional properties of their empirical counterparts, all the while matching the standard moments of exchange rates, interest rates and inflation.

Related literature: This paper builds on the literature addressing the risk–return relationship in FX markets. Lustig, Roussanov, and Verdelhan (2011) identify two risk factors: the average forward discount of the U.S. dollar against foreign currencies (*DOL* factor) and the return to the carry trade portfolio itself (*HML* factor). They show that the cross section of interest rate sorted currency portfolio returns can be mapped to differential exposure to the *HML* factor. They interpret the *HML* factor to be a global risk factor, so carry trade excess returns compensate U.S. investors for exposure to global risk. Lustig, Roussanov, and Verdelhan (2014) study a different investment strategy

that exploits the time series variation in the average U.S. interest rate differential vis-à-vis the rest of the world. This dollar carry trade is shown to be predictable by the average forward discount. In a similar vein, Verdelhan (2013) studies beta-sorted dollar portfolios and finds that similar to the carry factor, the dollar factor is a priced risk factor in the cross section of currencies. We contribute to this literature by presenting a new priced risk factor and by showing that distinguishing the relative importance of country-specific variation from global variation is key for understanding time-varying correlations in currency markets.

In other recent work on the carry trade, Cenedese, Sarno, and Tsiakas (2014) find that higher (lower) average currency excess return variance (correlation) leads to larger carry trade losses (gains). Menkhoff, Sarno, Schmeling, and Schrimpf (2012) show that the carry trade can be explained as compensation for global FX volatility risk. Mancini, Ranaldo, and Wrampelmeyer (2013) study the impact of FX liquidity on carry returns, and Lettau, Maggiori, and Weber (2014) and Dobrynskaya (2014) argue that the high returns are due to high conditional exposure to the equity market return in bad periods. Jurek (2014) shows that returns to selling put options which are exposed to downside risk can explain carry returns.

The rest of the paper is organized as follows. Section 1 describes the data and details the construction of the implied variance and correlation measures. Section 2 contains our empirical findings regarding the cross section and time series properties of FX correlations and FX correlation risk premia, as well as the pricing of correlation risk in currency markets. Section 3 presents our no-arbitrage model, and Section 4 concludes. We provide additional material on our model in the Appendix. Finally, additional results and robustness checks are presented in an Online Appendix.

1 Data and measures of FX correlation

We start by describing the data and the procedure to estimate realized and implied correlations. To construct measures of implied FX variances and correlations we use daily FX options data for the G10 currencies. Realized variances and correlations are

calculated using daily spot exchange rates. Our benchmark sample period starts in 1996 and is dictated by the availability of the options data.

1.1 Data description

Currency options: We use daily over-the-counter (OTC) currency options data from J.P.Morgan for the G10 currencies (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD, SEK, USD). In addition to the nine FX pairs versus the U.S. dollar we also have options data for the 36 cross rates. Using OTC options data has several advantages over exchange-traded options 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, and this leads to more competitive quotes in the OTC market. Second, the conventions for writing and quoting options in the OTC markets exhibit several features that are appealing when performing empirical studies. In particular, new option series with fixed times to maturity and fixed strike prices, defined by sticky deltas, are issued daily; in comparison, the time to maturity of an exchange-traded option series gradually declines with the approaching expiration date and so the moneyness continually changes as the underlying exchange rate moves. As a result, OTC options data allows for better comparability over time because the series' main characteristics do not change from day to day. The options used in this study are plain-vanilla European calls and puts with five option series per currency pair. Specifically, we consider a one-month maturity and a total of five different strikes: at-the-money (ATM), 10-delta and 25-delta calls, as well as 10-delta and 25-delta puts. The options data is available starting in 1996.

Spot and forward rates: To calculate realized variances and correlations we use daily spot exchange rates from WM/Reuters obtained through Datastream. In order to form currency portfolios we also collect one-month forward rates from WM/Reuters.¹ Following the previous literature (see, e.g., Fama, 1984), we work with the log spot and

¹The spot and forward rates are fixed at 4 p.m. UK time, which is standard in the FX market. For our benchmark sample period, we can use spot rates from a single source. This is important as we use daily exchange rates to calculate realized correlations and we therefore require that rates are measured in a consistent way. In order to conduct robustness checks we extend our sample to January 1984 by combining various data sources. We verify that for the overlapping period the spot rates from the additional sources are virtually identical.

one-month forward exchange rates, denoted $s_t^i = \ln(S_t^i)$ and $f_t^i = \ln(F_t^i)$, respectively.² We use the U.S. dollar as the base currency, so superscript i always denotes the foreign currency. In addition to the G10 currencies for which we have options data, our sample includes the following countries: Austria, Belgium, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Kuwait, Malaysia, Mexico, Netherlands, Philippines, Poland, Portugal, Singapore, South Africa, South Korea, Spain, Taiwan and Thailand.³ We run a separate analysis using only the sample of developed countries that includes Australia, Austria, Belgium, Canada, Denmark, France, Finland, Germany, Greece, Italy, Ireland, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom, as well as the Euro.

To summarize, we conduct our analysis for three samples: The full sample consists of a maximum of 34 currencies before the introduction of the Euro and of 24 currencies thereafter. The sample of developed currencies consists of 20 (10 after the introduction of the Euro) currencies and the sample of G10 currencies consists of a total of nine currencies (throughout the paper we use the German mark before the introduction of the Euro).

To illustrate the cross-sectional properties of our nine G10 currencies, we rank them according to their average nominal interest rate differential (forward discount) against the USD and report the average excess returns over the options data sample period starting in 1996. As can be seen from Panel A in Table 1, the NZD, AUD and NOK are characterized by high nominal interest rates, as well as high average currency excess returns. The reverse is true for the JPY, CHF and EUR/DEM. In line with the extant literature on the FX carry trade, currencies with high nominal interest rates achieve higher average dollar excess returns.

²WM/Reuters forward rates are available since 1997. For 1996 (and for the extended sample used for robustness checks we either use forward rates from alternative sources or we construct ‘implied’ forward rates using the interest rate differential between the U.S. and the foreign country (again, we collect interest rate data from Datastream). We thus exploit the fact that during normal conditions covered interest rate parity holds and, hence, $f_t - s_t \approx r_t^{i,\$} - r_t^{0,\$}$, where $r_t^{i,\$}$ and $r_t^{0,\$}$ denote the foreign and domestic nominal risk-free rates over the maturity of the contract, respectively. We verify that all results are robust to using the WM/Reuters data only.

³We start with the same set of currencies used in Lustig, Roussanov, and Verdelhan (2011). However, we exclude some currencies such as the Hong Kong dollar as they are pegged to the U.S. dollar. We also exclude the Danish krone after the introduction of the Euro.

[Insert Table 1 here.]

Carry portfolios: At the end of each month t , we allocate currencies into four portfolios (three portfolios for the G10 sample) based on their end-of-the-month forward discount and re-balance every month. Since covered interest rate parity holds in the data at the monthly frequency, sorting on forward discounts is equivalent to sorting on interest rate differentials. Portfolio 1 contains currencies with the lowest interest rates (smallest forward discounts) while Portfolio 4 contains currencies with the highest interest rates (largest forward discounts). Monthly excess returns from holding the foreign currency i are computed as $rx_{t+1}^i \approx f_t^i - s_{t+1}^i$.

We follow Lustig, Roussanov, and Verdelhan (2011) and build a long–short carry factor (*HML*) by investing in Portfolio 4 and shorting Portfolio 1. We also build a zero-cost dollar portfolio (*DOL*), which is an equally weighted average of the four currency portfolios and, thus, consists of borrowing U.S. dollars and investing in global money markets outside the United States in equal weights.

Summary statistics on the interest rate sorted currency portfolios, the *HML* factor, and the *DOL* factor are presented in Table 1. Panels B and C report the summary statistics for portfolios over the full sample period based on all countries and developed countries only, respectively. In line with previous findings, there is a monotonic increase in the average excess return from the lowest to the highest forward discount-sorted portfolio. The unconditional average excess return from holding an equally weighted average carry portfolio (*DOL*) is about 1.5% per annum for the full set of countries and 1.1% for the sample of developed countries. The *HML* portfolio is highly profitable; it has an average annual return of 6.1% (5.4%) for the set of all (developed) countries with an associated annualized Sharpe ratio of 0.79 (0.55). For the sample of the nine G10 currencies sorting on the average forward discount and sorting on average excess returns is nearly the same as can be seen from Panel A in Table 1. Sorting the G10 currencies into three different bins, we find that the average annual return to the G10 *HML* portfolio is 5.2% with a Sharpe ratio equal to 0.61.⁴

⁴The summary statistics for the extended sample starting in 1984 are very similar. The Sharpe ratios for the *HML* portfolios are 0.83 (all countries) and 0.54 (developed and G10 countries), respectively.

1.2 Construction of realized and implied volatility and correlation measures

In this section, we construct empirical measures of realized and implied exchange rate correlation for the G10 currencies sample. For the former we use daily exchange rate data, and for the latter we rely on a cross section of options written on those exchange rates.

1.2.1 Realized variance and correlation

We use daily spot exchange rates to calculate measures of realized variances and correlations. Denote $\Delta s_t^i = \ln(S_t^i) - \ln(S_{t-1}^i)$ the daily log change for currency i . The annualized realized variance observed at t is then calculated as follows:

$$\text{RV}_t = \frac{252}{K} \sum_{k=0}^{K-1} \Delta s_{t-k}^2,$$

where K refers to a three month window to estimate the rolling realized variances. Following Bollerslev, Tauchen, and Zhou (2009) we use this rolling estimate to proxy for the expected variance over the next month.⁵

In a similar spirit, we derive the annualized realized covariance between exchange rates s^i and s^j :

$$\text{RCov}_t^{i,j} = \frac{252}{K} \sum_{k=0}^{K-1} \Delta s_{t-k}^i \Delta s_{t-k}^j.$$

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

$$\text{RC}_t^{i,j} = \text{RCov}_t^{i,j} / \sqrt{\text{RV}_t^i} \sqrt{\text{RV}_t^j}.$$

1.2.2 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. They show that if

⁵Della Corte, Ramadorai, and Sarno (2014) use a similar approach in their analysis of the relationship between FX volatility risk premia and exchange rate predictability.

the underlying asset price is continuous, then the risk-neutral expectation over a horizon $T - t$ of total return variance is defined as an integral of option prices over an infinite range of strike prices:

$$\mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T (\sigma_u^i)^2 du \right) = 2e^{r(T-t)} \left(\int_0^{S_t^i} \frac{1}{K^2} P(K, T) dK + \int_{S_t^i}^{\infty} \frac{1}{K^2} C(K, T) dK \right), \quad (1)$$

where S_t is the underlying spot exchange rate and $P(K, T)$ and $C(K, T)$ are the respective put and call prices with maturity date T and strike K . In practice, the number of traded options for any underlying asset is finite; hence the available strike price series is a finite sequence. Calculating the model-free implied variance involves the entire cross section of option prices: for each maturity T , all five strikes are taken into account. These are quoted in terms of the option delta. In addition, we use daily spot rates and one-month Eurocurrency (LIBOR) rates from Datastream. Following the conventions in the FX market we use the Garman and Kohlhagen (1983) valuation formula to extract the relevant strike prices and to calculate the corresponding option prices.⁶

To approximate the integral in equation (1), we adopt a trapezoidal integration scheme over the range of strike prices covered by our dataset. 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 that arise due to the unavailability of a continuum of available options. We find that both errors are extremely small when currency options are used. For example, the size of the errors totals only half a percentage point in terms of volatility.

Model-free implied correlations are constructed from the available model-free implied volatilities.⁷ For the construction we require all cross rates for three currencies, S_t^i , S_t^j ,

⁶See, e.g., Wystup (2006) or, more recently, Jurek (2014) for the specifics of FX options conventions.

⁷Brandt and Diebold (2006) use the same approach to construct realized covariances of exchange rates from range-based volatility estimators.

and S_t^{ij} . The absence of triangular arbitrage then implies that:⁸ $S_t^{ij} = S_t^i/S_t^j$. Taking logs, we derive the following equation:

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

Finally, taking variances yields:

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

where $\gamma_t^{i,j}$ denotes the covariance of returns between exchange rate pairs s_t^i and s_t^j . Solving for the covariance term, we obtain:

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

Using the standard replication arguments, we find that:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \gamma_u^{i,j} du \right) &= e^{r(T-t)} \left(\int_t^{S_t^i} \frac{1}{K^2} P^i(K, T) dK + \int_{S_t^i}^{\infty} \frac{1}{K^2} C^i(K, T) dK \right. \\ &\quad + \int_t^{S_t^j} \frac{1}{K^2} P^j(K, T) dK + \int_{S_t^j}^{\infty} \frac{1}{K^2} C^j(K, T) dK \\ &\quad \left. - \int_t^{S_t^{ij}} \frac{1}{K^2} P^{ij}(K, T) dK - \int_{S_t^{ij}}^{\infty} \frac{1}{K^2} C^{ij}(K, T) dK \right). \end{aligned} \quad (2)$$

⁸Recent studies report that the average violation of triangular arbitrage is about 1.5 basis points with an average duration of 1.5 seconds (Kozhan and Tham, 2012). However, we observe that most papers examining violations of triangular arbitrage use indicative quotes, which give only an approximate price at which a trade can be executed. Executable prices can differ from indicative prices by several basis points. Using executable FX quotes, Fenn, Howison, McDonald, Williams, and Johnson (2009) report that triangular arbitrage is less than 1 basis point and the duration less than 1 second. Our data also indicate that triangular arbitrage is less than 1 basis point. We therefore conclude that these violations have no effect on calculated quantities.

The model-free implied correlation can then be calculated using expression (2) and the model-free implied variance expression (1):⁹

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

The summary statistics for the implied and realized FX variances for the G10 currencies are provided in Panels A and B in Table 2 (expressed as monthly numbers in squared percent). Given the nine available currencies, we have a total of 36 correlations. The first four columns in Table 3 provide the mean and the standard deviations for the pairwise implied and realized correlations where the FX pairs are ordered alphabetically.

[Insert Tables 2 and 3 here.]

2 Empirical analysis

In this section, we first calculate the correlation risk premia as the difference between the risk-neutral (implied) and objective (realized) correlation measures and we document their properties in the cross section and the time series. We then proceed by documenting

⁹Our expression for the risk-neutral correlation does not imply that the correlations need to be bounded between minus and plus one. One way to ensure that the absolute implied correlations never reach unity would be to impose a normalization in the spirit of the Dynamic Conditional Correlation model of Engle (2002). Because we find no implied correlations as high as one, we do not apply this normalization.

the link between correlation risk premia and average correlations, and we study the market price of correlation risk in FX markets.

2.1 Correlation and variance risk premia in the cross section

Consistent with the literature on variance and correlation risk premia in other asset markets, we define exchange rate correlation risk premia as the difference between the risk-neutral and objective measures of the FX correlation:¹⁰

$$\text{CRP}_{t,T}^{i,j} \equiv \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \rho_u^{i,j} du \right) - \mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \rho_u^{i,j} du \right).$$

We only consider one-month premia, i.e., $T = t+1$ for a monthly frequency. Variance risk premia are defined analogously as the difference between the risk-neutral and objective measures of FX variance.

Given the availability of FX options, we calculate correlation and variance risk premia for the nine G10 currencies during the sample period from 1996 to 2013 for a total of 216 monthly observations.¹¹ Panel C in Table 2 provides the summary statistics for the variance risk premia (expressed as monthly numbers in squared percent), whereas the last three columns in Table 3 provide the mean and standard deviations as well as the corresponding t-statistics for the pairwise correlation risk premia.

Despite the evidence for significant variance risk premia in equity markets (see, e.g., Bollerslev, Tauchen, and Zhou (2009)), the results in currency markets are mixed. On the one hand, the variance risk premia for a number of exchange rates are either not statistically significant (AUD, CHF, NZD) or only marginally significant (NOK), while, on the other hand, variance risk premia for the CAD and the SEK are significant at the 5% level and the variance risk premia for the EUR, GBP and JPY are significant at the 1% level. For all exchange rates for which the average variance risk premium is significant, the mean is positive. The average variance risk premium across all exchange rates is 0.54, which is smaller by a factor of more than ten compared to the equity

¹⁰We follow for example Bollerslev, Tauchen, and Zhou (2009) who define the variance risk premium in the same way. Alternatively, the variance risk premium is often defined as the difference between the physical and the risk-neutral measure (i.e. by reversing the sign), thus being consistent with the payoff to a long position in a variance swap.

¹¹For the Euro, the options data starts in 1999 for a total of 181 observations.

variance risk premium and by a factor of 4.5 compared to the Treasury variance risk premium.¹² Furthermore, there is an almost even split between exchange rates that have left- and right-skewed distributions for the variance risk premia. The two most negatively skewed USD exchange rates are the AUD and the NZD, which reflects the implicit crash risk in these currencies as they are often used as investment currencies in the carry trade (see Brunnermeier, Nagel, and Pedersen (2009)), whereas the EUR (starting in 1999), the JPY and the GBP are the most positively skewed USD exchange rates.

In contrast, we find that correlation risk premia can be substantial in the foreign exchange market. Moreover, the size and the sign of the correlation risk premia can vary greatly in the cross section of FX pairs. Roughly two thirds of the pairwise correlation risk premia are positive and one third are negative; overall, three quarters of all pairwise average correlation risk premia are significant at the 5% level of significance. The average of the bottom quartile of correlation risk premia is -3.7% , whereas the top quartile average is 7.2% .

A comparison of the average realized correlation (RC) and the average correlation risk premia (CRP) in Table 3 furthermore suggests a link between correlation risk premia and levels of correlation. While FX pairs that are less correlated exhibit positive correlation risk premia, the opposite holds for pairs characterized by high average correlations. For example, the JPY/AUD pair has a very low average realized correlation (15.5%) but a positive and highly significant correlation risk premium of 8.3% (t-statistic of 7.58). On the other hand, the AUD/NZD pair has a very high realized correlation (75.5%) but a negative and significant correlation risk premium of -1.6% (t-statistic of -2.97).

[Insert Figure 1 here.]

Figure 1 plots the average correlation risk premia of the 36 G10 exchange rate pairs against their average realized correlations, illustrating the cross-sectional relationship between correlations and correlation risk premia. Across all 36 FX pairs, the correlation

¹²The corresponding numbers are calculated using the methodology reported in Mueller, Vedolin, and Choi (2014) for a slightly longer sample. The corresponding values are 6 and 2.4 for the equity and Treasury variance risk premia, respectively.

between average FX correlation risk premia and average FX realized correlations is -55% . In sum, we conclude that unconditional correlation risk premia are economically large and positive for FX pairs with low average realized correlations and significantly negative for pairs with high average realized correlations.

2.2 The time series of FX correlations and correlation risk premia

We now turn to the time series properties of FX correlations and FX correlation risk premia. The first two columns of Table 4 present the time series correlations between levels and changes in conditional realized and implied correlations for every FX pair in our sample; unsurprisingly, the correlation between realized and implied correlations in levels is very high. Columns three and four show the time series correlations between conditional realized correlations and correlation risk premia: they are negative both for levels and changes for almost all FX pairs. For changes, the correlations range between -35% and -79% with an average of -62% . For levels, the correlations are slightly smaller but still average -52% .

[Insert Table 4 here.]

Next, we turn to the cyclical properties of FX correlations in the cross section of FX pairs. In particular, for each exchange rate pair we calculate the unconditional correlation between the time series of its conditional correlations and the time series of a number of macroeconomic and market variables that are well-known to exhibit countercyclical behavior; we take this unconditional correlation to be our FX correlation cyclicity measure. The macroeconomic and market variables we consider are a global equity volatility measure ($GVol$), a global funding illiquidity measure (GFI), the TED spread (TED), and the CBOE VIX (VIX).¹³ Then, we explore the cross-sectional properties of our FX correlation cyclicity measure by plotting the cyclicity measure

¹³ $GVol$ is constructed as in Lustig, Roussanov, and Verdelhan (2011). GFI is constructed based on the method proposed by Hu, Pan, and Wang (2013) but calculated using an international sample of Treasury securities as in Malkhozov, Mueller, Vedolin, and Venter (2014). TED is from FRED and is the spread between the three month USD LIBOR and the three month Treasury Bill rate. VIX is backed out from options on the S&P 500 stock index. TED and VIX are U.S. specific measures but are often used as global indicators. $GVol$ and GFI are calculated using international data in local currencies.

of all 36 FX pairs against either their average conditional correlation coefficient or their average correlation risk premium.

Our findings are presented in Figure 2. For all four cyclicity proxies, there is a significant positive relationship in the cross section of FX pairs between our FX correlation cyclicity measure and average FX correlations (left hand side panels), which implies a negative relationship between FX correlations and FX correlation cyclicity. Given the negative relationship between average FX correlations and average FX correlation risk premia, it is not surprising that we also find a significant negative relationship between the cyclicity measures and average correlation risk premia (right hand side panels), which implies a positive relationship between average correlation risk premia and FX correlation cyclicity. All univariate slope coefficients are highly statistically significant. The cross-sectional R^2 from regressing the average FX correlations on our FX correlation cyclicity measure range between 14% (*VIX*) and 50% (*TED*), whereas the R^2 for the same regressions using average correlation risk premia as the regressand range between 50% (*GFI*) and 61% (*VIX*). For example, the three FX pairs with the highest average realized correlations (CHF/EUR, EUR/NOK, and EUR/SEK) and those with the lowest average correlation risk premia (AUD/CAD, CAD/NZD and CAD/SEK) exhibit either acyclical correlations (i.e., the unconditional correlation between the realized correlations and the cyclical variables is not significant) or slightly countercyclical correlations. In the other extreme, the three FX pairs with the highest correlation risk premia (AUD/JPY, JPY/NOK, and JPY/NZD) and the lowest realized correlations (AUD/JPY, CAD/JPY and JPY/NOK) are characterized by strongly procyclical FX correlations.

[Insert Figures 2 and 3 here.]

Figure 3 plots the relationship between measures of cyclicity of correlation risk premia (instead of FX correlations) and average correlation risk premia. We find a positive cross-sectional association between the measures of cyclicity of correlation risk premia and unconditional correlation risk premia: FX pairs with high average correlation risk premia have countercyclical correlation risk premia, whereas pairs with low correlation risk premia have procyclical correlation risk premia.

In sum, FX pairs with high average correlations or low average correlation risk premia exhibit countercyclical correlations and procyclical correlation risk premia, whereas FX pairs with low average correlations or high average correlation risk premia tend to have procyclical correlations and countercyclical correlation risk premia. Our findings imply that in periods characterized by adverse economic conditions or market stress, the cross section of conditional FX correlations widens, as high correlation FX pairs become more correlated and low correlation FX pairs become less correlated. Thus, the difference in conditional correlations between high correlation FX pairs and low correlation FX pairs is also countercyclical, increasing during crises and declining during booms. Similarly, bad times also amplify the magnitude of correlation risk premia, increasing their cross-sectional dispersion, as the risk premia of FX pairs with high average risk premia increase and the risk premia of FX pairs with low average risk premia decline.

Our findings have important implications for models of exchange rate determination. Under the assumption that investors are risk-averse, we may expect that FX pairs with countercyclical (procyclical) correlations are characterized by positive (negative) correlation risk premia. Since the former are the high average correlation pairs and the latter are the low average correlation pairs, we may thus anticipate a positive cross-sectional relationship between average FX correlation coefficients and average FX correlation risk premia, which would imply that investors place a high price on states characterized by a tightening of cross-sectional differences in conditional FX correlation. However, exactly the reverse is true in the data. In Section 3 we present a no-arbitrage model of exchange rates that addresses this apparent inconsistency.

2.3 The FX correlation risk factor and the cross section of currency returns

In this Section we describe the construction of our correlation factor FXC and we document that our factor is priced in the cross section of currency returns. In particular, we find that our correlation factor exhibits a negative price, so there is a negative relationship between FXC betas and currency returns.

2.3.1 Construction of the FX correlation risk factor

As noted in Section 2.2, the difference in conditional correlation between high correlation FX pairs and low correlation FX pairs increases in bad times. We use this finding to construct our FX correlation factor as follows: every period t , we sort all FX pairs according to their conditional correlation, defined as the realized correlation over the past three months. We then calculate the average conditional correlation for the top and bottom decile (which consists of four pairs each) and take the difference of the two values as our FX correlation factor at time t , FXC_t .¹⁴ Due to the time variation in conditional FX correlations, there is turnover in both the top and bottom deciles; in order to abstract from composition effects, we also compute an alternative correlation factor (FXC^{UNC}) by using the top and bottom deciles of FX pairs based on the unconditional realized correlations.

We plot the time series of the level of the two FX correlation factors in Panel A of Figure 4. The correlation between the two series is 87%, indicating that the two factors are very similar.¹⁵ In Panel B, we plot the (standardized) macroeconomic and market variables used to measure the cyclicity of correlations in the previous Section. Table 5 reports the unconditional correlations between our two correlation factors and the cyclicity proxies. All correlations are significantly positive, confirming the countercyclicality of our correlation factors.

[Insert Figure 4 and Table 5 here.]

2.3.2 Portfolios sorted on exposure to correlation risk

Next, we sort currencies into portfolios based on their exposure to correlation risk. In particular, we measure correlation risk exposure by the currency return beta with

¹⁴Given the strong negative correlation between correlation risk premia and realized correlation, we could also sort based on correlation risk premia and then construct the correlation factor as the difference in correlations between the decile of low CRP currencies and the decile of high CRP currencies. However, using realized correlations directly allows us to expand the sample for robustness checks.

¹⁵The Online Appendix presents additional results for alternative construction methods. Overall, we find that results using the alternative factors remain qualitatively the same.

respect to innovations in the FX correlation risk factor FXC .¹⁶ Our currency portfolios are rebalanced monthly; this means each month t we calculate rolling betas using 36 monthly observations and, hence, portfolios are formed using only information available at time t . Since we only have nine G10 currencies, we sort them into three portfolios, whereas for the full set of countries and the set of developed countries we form four portfolios. The first portfolio (Pf1) contains the currencies with the low correlation risk factor betas while the last portfolio (Pf3 or Pf4) contains the high beta currencies. We also construct an HML^C portfolio that is long in the high correlation beta currencies (Pf3 or Pf4) and short in the low correlation beta currencies (Pf1). Table 6 reports the summary statistics for the various portfolios and the three sets of currencies. There is an inverse relationship between exposure to the FX correlation factor and average portfolio returns, implying that currencies that depreciate when cross-sectional disparity in conditional FX correlation increases (those with low or negative FXC betas) are risky, whereas currencies that appreciate when our FX correlation factor increases (those with high or positive betas) provide hedging benefits. The average return to the HML^C is negative and highly statistically significant for all three currency samples. For the sample of G10 currencies, shorting the HML^C portfolio yields an annualized average excess return of 6.4% with an associated Sharpe ratio of 0.82. For the sample of all countries and developed countries, the average annualized returns are 4% (HML^{C-ALL}) and 5.5% (HML^{C-DEV}), respectively, and the Sharpe ratios are 0.46 and 0.59. In terms of magnitudes, the results are comparable to those presented in Table 1 for the carry portfolios.¹⁷

[Insert Table 6 and Figure 5 here.]

¹⁶In order to calculate innovations in FXC we cannot simply take first differences in the level of the factor, as the composition of the deciles changes over time. Thus, innovations of period $t + 1$ are calculated using the average of first differences in conditional FX correlation for the FX pairs that belong to the top and bottom decile in period t . On the other hand, since the FX pairs used to calculate FXC^{UNC} are fixed, innovations in FXC^{UNC} can be simply defined as first differences in the level of the factor.

¹⁷Extending the sample back to 1984, the results are weaker but qualitatively similar. Shorting the HML^C portfolio yields average annual excess returns of 3.7%, 2.2% and 3.1% for the sample of G10, all and developed countries, respectively (with the associated Sharpe ratios being 0.44, 0.24 and 0.36, respectively). On the other hand, the results are stronger if the credit crisis is excluded and the sample ends in July 2007. The average excess returns to shorting the HML^C portfolio are then 7.4%, 5.5% and 7.2% for the sample of G10, all and developed countries, respectively (with associated Sharpe ratios of 1.1, 0.7 and 1.0, respectively). Further details on the subsamples are provided in the Online Appendix.

Figure 5 summarizes the average excess portfolio returns for the three sets of currencies (G10, all countries and developed countries) and four subsamples (starting in either January 1996 or January 1984 and ending in December 2013 or just before the credit crisis in July 2007). While the average return to the HML^C portfolio varies depending on the sample period, there is a consistently negative relationship between average portfolio excess returns and exposure to correlation risk: for all samples and sample periods we find that portfolios of currencies with low FXC betas have higher excess returns than portfolios of currencies with high FXC betas.

2.4 Cross-sectional pricing and the price of correlation risk

Finally, we turn to estimating the price of correlation risk. Similarly to Lustig, Roussanov, and Verdelhan (2011), we consider a linear pricing model with two factors. The first factor is the dollar factor DOL , defined as the simple average of all available currency returns. Lustig, Roussanov, and Verdelhan (2011) have shown that DOL acts as level factor for currency returns. As a second factor, we use HML^C , the return difference between the high and the low correlation beta portfolio for the sample of G10 currencies and, thus, a traded factor that captures correlation risk.¹⁸ Hence, our model is:

$$E[rx^i] = \beta_i^{DOL} \lambda^{DOL} + \beta_i^{HML^C} \lambda^{HML^C},$$

where rx^i denotes the excess return in levels (i.e., corrected for the Jensen term) instead of logs as in Lustig, Roussanov, and Verdelhan (2011). To estimate the factor prices λ we follow the traditional two-stage procedure of Fama and MacBeth (1973): in the first stage, we run a time series regression of returns on the factors and then we run a cross-sectional regression of average portfolio returns on the betas. We do not include a constant in the cross-sectional regression of the second stage.¹⁹

Table 7 reports the asset pricing results on the currency portfolios sorted on exposure to the correlation risk factor FXC . The left hand side of the table presents the results for the correlation portfolios using all currencies whereas the right hand side of the table

¹⁸The results remain qualitatively the same if we use HML^{C-ALL} or HML^{C-DEV} instead.

¹⁹The dollar factor DOL essentially performs the function of a constant to allow for average mispricing (see Lustig, Roussanov, and Verdelhan (2011)).

presents the results for the developed currencies only. The first stage regression results are presented in Panel A. The HML^C betas are monotonically increasing for both sets of currencies and with one exception they are highly statistically significant. In line with the countercyclical nature of the correlation risk factor we find a (significantly) negative price of correlation risk of 48bps (all countries) and 49bps (developed countries) per month.²⁰ This corresponds to almost 90% of the average HML^C excess return of minus 54bps per month and is, hence, not significantly different. The second stage R^2 are very high for both sets of currencies (95% and 96%). HML^C betas for the correlation portfolios are also monotonically increasing when we consider the alternative sample periods 1984–2013 and 1996–2007 and the estimates for the market price of HML^C remain significantly negative with slightly different magnitudes in line with different average excess returns of HML^C during those alternative sample periods.²¹

[Insert Figure 6 and Table 7 here.]

For robustness, we repeat the analysis above using alternative sets of test assets. Figure 6 illustrates the performance of our linear two-factor model by plotting the predicted excess returns for various test assets against the actual annualized mean excess returns: in Panel A, the test assets are the individual G10 currencies, whereas in Panels B and C we consider the four interest-rate-sorted (i.e., carry) and FXC -beta-sorted portfolios for all countries and developed countries, respectively. The second stage R^2 for the three sets of test assets are 92%, 81% and 90%, respectively. We consistently estimate a significantly negative market price of FX correlation risk that is never significantly different from the mean return of the HML^C portfolio. For the set of individual G10 currencies for example, the estimated market price of risk is minus 48bps per month (−5.7% annualized).

Overall, we find that correlation risk is priced in the cross section of currency returns with a negative price, in line with the intuition that investors want to be compensated for investing in currencies that perform badly during periods of increased cross-sectional disparity in conditional FX correlations.

²⁰Using the non-traded correlation risk factor FXC we also estimate a significantly negative price of correlation risk.

²¹Further details are provided in the Online Appendix.

3 A no-arbitrage model of exchange rates

In this section, we introduce a reduced-form, no-arbitrage model of exchange rates that is consistent with our empirical findings. Our model builds on the reduced-form models in Lustig, Roussanov, and Verdelhan (2011, 2014) and Verdelhan (2013). In contrast to these models, which assume that innovations in the local pricing factors are purely driven by local shocks and, thus, are uncorrelated across countries, we assume that there is significant cross-country comovement in the price of local risk. This key assumption allows our model to match the joint empirical properties of FX correlations and FX correlation risk premia.

3.1 Model setup

The global economy comprises $I+1$ countries ($i = 0, 1, \dots, I$), each with a corresponding currency. Without loss of generality, we will call country $i = 0$ the domestic country and countries $i = 1, \dots, I$ the foreign countries. We assume that financial markets are frictionless and complete, so that there is a unique stochastic discount factor (SDF) for each country, but that frictions in the international market for goods induce non-identical stochastic discount factors across countries. In particular, the SDF of country i , denoted by m^i , is exposed to two global shocks, u^w and u^g , and a country-specific shock u^i , and satisfies

$$-m_{t+1}^i = \alpha^i + \chi^i z_t^i + \varphi^i z_t^w + \sqrt{\kappa^i z_t^i} u_{t+1}^i + \sqrt{\gamma^i z_t^w} u_{t+1}^w + \sqrt{\delta^i z_t^i} u_{t+1}^g,$$

where z^i and z^w are the local pricing factor of country i and the global pricing factor, respectively. The price of the local shock in country i is $\sqrt{\kappa^i z_t^i}$, so the relevant pricing factor is the local one. The first global shock u^w is priced with the global pricing factor, with its price in country i being $\sqrt{\gamma^i z_t^w}$, while the second global shock u^g is locally priced and its price in country i is $\sqrt{\delta^i z_t^i}$. Therefore, the relative price of risk across countries for the local shock and the second global shock u^g exhibits time-variation, whereas the relative price of risk for the first global shock u^w is constant.

The pricing factors are time-varying, stationary processes driven by the aforementioned shocks. In particular, the country-specific pricing factor z^i is exposed to both the local shock u^i and the second global shock u^g , with law of motion

$$\Delta z_{t+1}^i = \lambda^i(\bar{z}^i - z_t^i) - \xi^i \sqrt{z_t^i} \left(\sqrt{\rho^i} u_{t+1}^i + \sqrt{1 - \rho^i} u_{t+1}^g \right). \quad (3)$$

Thus, all local pricing factors are stationary processes, reverting to their unconditional mean \bar{z}^i at speed λ^i . The exposure of the local pricing factors to the global shock u^g introduces cross-country comovement in local pricing factors. Importantly, all local pricing factors are countercyclical, as adverse shocks increase their value.

The global pricing factor z^w is only exposed to the global shock u^w ; it is also a stationary process, with law of motion

$$\Delta z_{t+1}^w = \lambda^w(\bar{z}^w - z_t^w) - \xi^w \sqrt{z_t^w} u_{t+1}^w$$

and also features countercyclical pricing of risk. To ensure that all pricing factors are strictly positive, we further assume that the Feller conditions $2\lambda^i \bar{z}^i > (\xi^i)^2$ for all i and $2\lambda^w \bar{z}^w > (\xi^w)^2$ are satisfied.

Finally, we specify an exogenous inflation process given by

$$\pi_{t+1}^i = \bar{\pi}^i + \psi^i z_t^i + \zeta^i z_t^w + \sqrt{\sigma^i} \eta_{t+1}^i.$$

Inflation has a local and a global component and is exposed to unpriced inflation innovations η^i . Given that inflation shocks are i.i.d., there is no inflation risk premium in our model (see, e.g., Lustig, Roussanov, and Verdelhan (2011)). The nominal stochastic discount factor of country i , $m^{i,\$}$, satisfies

$$m_{t+1}^{i,\$} = m_{t+1}^i - \pi_{t+1}^i.$$

All shocks are i.i.d. standard normal.

In order to illustrate the basic economic mechanism of our model with convenient simplicity, it will be advantageous to focus on a simpler version of our full model, which

will be referred to as the benchmark model. In particular, in the benchmark model the only source of heterogeneity across countries is their differential exposure γ to the first global shock u^w . Apart from γ , all other parameters are assumed to be identical across countries, so we can drop the i superscript for notational convenience. We also assume that $\rho = 0$, which means that all local pricing factors are driven solely by the second global shock u^g . Further, we assume that all local pricing factors have the same initial value. As a result, all local pricing factors are identical and we can consider a single local pricing factor denoted by z . In the larger part of the remainder of this section, we focus on the benchmark model. We consider the full model in the last part of our analysis.

3.2 *The properties of conditional FX moments*

We denote the real log exchange rate between foreign country i and the domestic country by q^i (units of foreign currency per units of domestic currency, in real terms). As a result of financial market completeness, real exchange rate changes equal the SDF differential between the two countries,

$$\Delta q_{t+1}^i = m_{t+1}^0 - m_{t+1}^i,$$

which implies that real exchange rate changes can be decomposed into a part driven by local shocks and a part that reflects exposure to global risk:

$$\Delta q_{t+1}^i = E_t(\Delta q_{t+1}^i) + \sqrt{\kappa z_t} u_{t+1}^i - \sqrt{\kappa z_t} u_{t+1}^0 + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right) \sqrt{z_t^w} u_{t+1}^w.$$

If the foreign country has a higher (lower) exposure to the first global shock u^w than the domestic country, its currency depreciates (appreciates) when a positive u^w realization occurs. On the other hand, exposure to the second global shock u^g drops out of exchange rate changes since all countries have the same loading on u^g , and, thus, the only global shock that is priced in foreign exchange markets is u^w . Therefore, in the remainder of this section, global FX risk always refers to the first global shock u^w .

We now turn to conditional FX moments. The conditional variance of changes in the log real exchange rate i is increasing in both the local pricing factor z and the global pricing factor z^w :

$$\text{var}_t(\Delta q_{t+1}^i) = 2\kappa z_t + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2 z_t^w. \quad (4)$$

The first effect arises from the country-specific component of stochastic discount factors: given the independence of local shocks across countries, the larger those shocks are, the more the two SDFs diverge and, hence, the more volatile the exchange rate is. The second effect arises from the global component of SDFs: the higher the heterogeneity between the global risk exposure of country i and the domestic country, and the more severely those differences are priced, the higher the real exchange rate volatility is.

The conditional covariance of changes in log real exchange rates i and j is

$$\text{cov}_t(\Delta q_{t+1}^i, \Delta q_{t+1}^j) = \kappa z_t + D^{i,j} z_t^w, \quad (5)$$

where we define the constant $D^{i,j}$ as follows:

$$D^{i,j} \equiv \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right) \left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right).$$

We call exchange rate pairs (i, j) that satisfy $D^{i,j} > 0$ “similar” and exchange rate pairs that satisfy $D^{i,j} < 0$ “dissimilar”. Thus, similar exchange rates correspond to foreign countries which both have either more or less exposure to global risk than the domestic country, whereas dissimilar exchange rates correspond to pairs of foreign countries in which one country has a higher and the other country a lower exposure to global risk compared to the domestic country.

The first component of conditional FX covariance reflects the price of the domestic local shock, as the two exchange rates are mechanically correlated through their common exposure to the domestic SDF. When z increases, this “domestic currency effect” becomes more prevalent, increasing the covariance between the two exchange rates, as both foreign currencies appreciate or depreciate together against the domestic currency.

The second component captures the comovement of exchange rate changes that arises from exposure to global FX risk. On average, foreign countries with similar exposure

to the global shock u^w (i.e., that satisfy $D^{i,j} > 0$) have exchange rates that covary more than the exchange rates of countries that have dissimilar exposure to global FX risk. Furthermore, the effect of fluctuations in z^w on the conditional covariance of log exchange rate changes depends on whether those fluctuations increase or decrease the similarity of the pricing of global risks vis-à-vis the domestic country. In particular, for similar exchange rates an increase in the global pricing factor increases their conditional covariance. On the other hand, dissimilar exchange rates comove less when the global pricing factor increases. Therefore, although an increase in z^w increases the conditional volatility of all exchange rates, it can either increase or decrease the conditional covariance of an FX pair depending on the global risk loadings of the countries involved.

Combining equations (4) and (5), we get the conditional FX correlation between the log exchange rate changes of countries i and j , given by:

$$\text{corr}_t(\Delta q_{t+1}^i, \Delta q_{t+1}^j) = \frac{\kappa z_t + D^{i,j} z_t^w}{\sqrt{2\kappa z_t + (\sqrt{\gamma^i} - \sqrt{\gamma^0})^2 z_t^w} \sqrt{2\kappa z_t + (\sqrt{\gamma^j} - \sqrt{\gamma^0})^2 z_t^w}}.$$

As happens for FX covariances, country heterogeneity in exposure to the global shock u^w generates cross-sectional heterogeneity in average conditional FX correlations: similar FX pairs have higher unconditional correlations than dissimilar ones. In fact, the higher the similarity or dissimilarity in exposures (measured by the magnitude of the absolute value of $D^{i,j}$), the higher the cross-sectional dispersion of average FX correlations. In the time series, an increase in the global pricing factor z^w increases the dispersion in the cross section of conditional FX correlations, as it raises the correlation of exchange rates with high average correlation (those of similar FX pairs) and decreases the correlation of countries with low average correlation (those of dissimilar FX pairs). In the limit, as $z^w \rightarrow \infty$, similar exchange rates become perfectly positively correlated and dissimilar exchange rates become perfectly negatively correlated.

On the other hand, an increase of the local pricing factor z increases both the FX variance of all exchange rates and the FX covariance of all exchange rate pairs, the latter due to the domestic currency effect. As $z \rightarrow \infty$, the correlation of all FX pairs converges to $\frac{1}{2}$. This happens because all cross-sectional differences in global risk exposure become

second-order and what ultimately drives comovement across countries is the domestic currency effect. Intuitively, the behavior of real exchange rate changes is described by

$$\Delta q_{t+1}^i \rightarrow E_t(\Delta q_{t+1}^i) + \sqrt{\kappa z_t} u_{t+1}^i - \sqrt{\kappa z_t} u_{t+1}^0.$$

All country-specific shocks have equal volatility, so the domestic shock, which accounts for half of the conditional FX variance and generates all the FX comovement, pushes all FX correlations towards $\frac{1}{2}$. Thus, as the local pricing factor increases, the conditional correlation of similar exchange rates (which have high unconditional correlations) declines, whereas the conditional correlation of dissimilar exchange rates (with low unconditional correlations) increases, leading to a tightening of the cross section of conditional FX correlations in bad times.

To quantify the effects of the pricing factors on the conditional FX correlations, we consider a world of $I = 3$ foreign countries. Countries 1 and 2 are less exposed to global FX risk than the domestic country, while country 3 is more exposed than the domestic country. This implies that the FX pair (1,2) is similar whereas FX pair (1,3) is dissimilar. To ensure symmetry, we set the values of the country exposures to global risk such that the condition $D^{1,2} = -D^{1,3} > 0$ is satisfied.

[Insert Figure 7 here.]

We first consider the impact of the global pricing factor z^w ; the left panels of Figure 7 present the results. In particular, Panels A, C and E plot the conditional FX correlation as a function of z^w for different values of the local pricing factor (in particular, $z = 0.2\bar{z}$, \bar{z} and $5\bar{z}$, respectively). Panel A refers to the similar exchange rate pair (1,2), Panel C refers the dissimilar exchange rate pair (1,3) and Panel E plots the difference in the conditional FX correlations between the two FX pairs. We see that regardless of the value of the local pricing factor, the conditional correlation of the similar FX pair is increasing in z^w as the similarity of the two exchange rates to global risk exposure increases their comovement when global fluctuations become more highly priced. For large values of the global pricing factor, the similar exposure of the two exchange rates to global risk leads to a perfect positive correlation between the two exchange rates. Exactly the

opposite occurs for the dissimilar exchange rate pair: an increase in z^w always reduces their conditional correlation and for large values of the global pricing factor they become perfectly negatively correlated. Taken together, these results imply that the correlation of similar and dissimilar FX pairs diverges as z^w increases. Thus, similar FX pairs become more correlated, whereas dissimilar FX pairs become less correlated, and, hence, the disparity in conditional FX correlations is increasing in z^w , as illustrated in Panel E.

This pattern is reversed once we consider the effects of the local pricing factor z . The results are presented in the right panels of Figure 7, where Panels B, D and F plot the sensitivity of the conditional FX correlation to the value of the local pricing factor z for different values of the global pricing factor (i.e., $z^w = 0.2\bar{z}$, \bar{z} and $5\bar{z}$, respectively), with Panel B referring to the similar FX pair, Panel D to the dissimilar FX pair and Panel F to the difference in their conditional FX correlations. As seen in Panel B, the conditional correlation between the two exchange rates is decreasing in z regardless of the value of z^w . This is due to the domestic currency effect, which for large values of z pushes the correlation of the two exchange rates towards $\frac{1}{2}$. On the other hand, the domestic currency effect induces a negative relationship between the value of the local pricing factor and the conditional FX correlation of the dissimilar pair. As a result, the difference in the two FX correlations is decreasing in z , regardless of the value of z^w : as z increases, what matters for exchange rates are domestic currency fluctuations, not differences across individual currencies. In those states, all currencies appreciate and depreciate together against the domestic currency.

In sum, what drives the behavior of conditional FX correlations and their dispersion is the relative importance of the global against the local pricing factor: fluctuations in the global pricing factor z^w generate countercyclical FX correlation differentials, whereas changes in the local pricing factor z lead to procyclical FX correlation differences. Empirically, we find that conditional FX correlations become more cross-sectionally dispersed in bad economic times (see Figure 4).²² Therefore, an important corollary of our model is that the *countercyclical* widening of FX correlation differentials observed empirically

²²As outlined in Section 2.3, we construct our empirical correlation factor FXC on the basis of this finding.

suggests that the time series of conditional FX correlation is primarily driven by fluctuations in the *global* pricing factor.

Nominal exchange rate changes satisfy

$$\Delta s_{t+1}^i = m_{t+1}^{0,\$} - m_{t+1}^{i,\$} = \Delta q_{t+1}^i + \pi_{t+1}^i - \pi_{t+1}^0.$$

Given the symmetry assumption for all inflation parameters, expected inflation is identical across countries, so nominal expected FX changes are identical to real expected FX changes. However, inflation differentials add unpriced volatility due to the idiosyncratic nature of the inflation shocks η . In particular, the conditional variance of nominal log exchange rate changes is given by

$$\text{var}_t(\Delta s_{t+1}^i) = \text{var}_t(\Delta q_{t+1}^i) + \text{var}_t(\pi_{t+1}^i - \pi_{t+1}^0) = \text{var}_t(\Delta q_{t+1}^i) + 2\sigma.$$

Furthermore, domestic inflation shocks enhance the domestic currency effect in FX co-movement:

$$\text{cov}_t(\Delta s_{t+1}^i, \Delta s_{t+1}^j) = \text{cov}_t(\Delta q_{t+1}^i, \Delta q_{t+1}^j) + \sigma.$$

Given the homoskedasticity of inflation innovations, conditional nominal FX moments equal their real FX counterparts adjusted by constants, so their behavior has the same properties as that of real exchange rate moments.

3.3 Interest rates and currency returns

The real interest rate of country i is determined by the local pricing factor z and the global pricing factor z^w , each generating both a smoothing and a precautionary savings motive:

$$r_t^i = \alpha + \left(\chi - \frac{1}{2}\kappa - \frac{1}{2}\delta \right) z_t + \left(\varphi - \frac{1}{2}\gamma^i \right) z_t^w.$$

All cross-sectional heterogeneity in real interest rates is due to cross-sectional differences in global risk exposure γ : each period, countries with high (low) exposure to global risk have a relatively low (high) average interest rate, due to a higher (lower) precautionary savings motive. As a result, interest rate differentials against the domestic country

solely reflect the difference in global risk exposure between each foreign country and the domestic country.

The nominal risk-free rate of country i is

$$r_t^{i,\$} = r_t^i + \bar{\pi} + \psi z_t + \zeta z_t^w - \frac{1}{2}\sigma.$$

Hence, due to symmetry, the inflation-induced component is identical across currencies and, thus, nominal interest rate differentials are identical to their real counterparts.

The log excess return that the domestic investor gets by investing in the foreign currency i is given by

$$rx_{t+1}^i = r_t^i - r_t^0 - \Delta q_{t+1}^i$$

and the associated conditional risk premium (including the Jensen term) is

$$\begin{aligned} rp_t^i &\equiv E_t(rx_{t+1}^i) + \frac{1}{2}var_t(rx_{t+1}^i) = -cov_t(m_{t+1}^0, -\Delta q_{t+1}^i), \\ &= cov_t(\sqrt{\kappa z_t} u_{t+1}^0, \sqrt{\kappa z_t} u_{t+1}^0) + cov_t\left(\sqrt{\gamma^0 z_t^w} u_{t+1}^w, \left(\sqrt{\gamma^0} - \sqrt{\gamma^i}\right) \sqrt{z_t^w} u_{t+1}^w\right), \\ &= \kappa z_t + \left(\sqrt{\gamma^0} - \sqrt{\gamma^i}\right) \sqrt{\gamma^0} z_t^w. \end{aligned}$$

FX risk premia have two components: a part that compensates domestic investors for the fact that investing in a foreign currency essentially entails shorting the domestic SDF, and a part that reflects compensation for exposure to global risk. The first component is identical across currencies, so any cross-sectional variation in FX risk premia is solely due to heterogeneity in global risk exposure γ . For example, if the foreign country is less exposed to global risk than the domestic country, then its currency depreciates against the domestic currency when a bad realization of the global shock u^w occurs. Assuming that the domestic country has a positive loading γ^0 to the global shock u^w , domestic investors require a positive risk premium in order to hold that currency. Conversely, currencies of countries with high exposure to u^w have a negative compensation for global FX risk, as they appreciate in bad global times, providing a hedge to domestic investors.

It follows that any variable that is exposed to u^w shocks should be priced in the cross section of currency returns. In our model, the cross-sectional disparity in conditional FX correlations is such a variable: a negative u^w shock increases the global pricing factor z^w

but has no effect on the local pricing factor z , so it leads to a widening of FX correlation differentials. As a result, changes in our empirical correlation factor FXC should price the cross section of FX returns with a negative price of risk, as they are negatively associated with global shocks u^w .

Notably, interest rate differentials against the domestic currency and currency risk premia can be used to infer the sign and magnitude of the difference in global risk exposures between each country and the domestic country: there is generally a positive relationship between both the product of interest rate differentials $(r_t^i - r_t^0)(r_t^j - r_t^0)$ and the product of conditional currency risk premia $rp_t^i rp_t^j$ with the conditional FX correlations of exchange rates i and j . This relationship also holds on average: our model suggests that average conditional FX correlations are generally positively associated with the product of average interest rate differentials $E(r_t^i - r_t^0)E(r_t^j - r_t^0)$ and the product of average currency risk premia $E(rp_t^i)E(rp_t^j)$ in the cross section. This positive relationship is present in the data: in our sample of G10 exchange rates, the cross-sectional correlation of average FX correlations with the product of corresponding interest rate differentials is 35% and with the product of average currency risk premia is 42%.

3.4 Pricing the cross section

We can contrast our FX correlation factor FXC , which captures the cross-sectional disparity in conditional FX correlation, with factors that have been shown in earlier research to be priced in the cross section of currency returns.

We begin with the carry trade factor HML . The excess return to the carry trade HML portfolio is defined as:

$$rx_{t+1}^{HML} = \frac{1}{N} \sum_{i \in H} rx_{t+1}^i - \frac{1}{N} \sum_{i \in L} rx_{t+1}^i,$$

with high interest currencies in set H and low interest rate currencies in set L . Provided that the long and the short end of the portfolio contain enough currencies so that the local shocks cancel out, i.e.

$$\frac{1}{N} \sum_{i \in H} u_{t+1}^i \rightarrow 0 \quad \text{and} \quad \frac{1}{N} \sum_{i \in L} u_{t+1}^i \rightarrow 0,$$

the return innovations of the HML portfolio are solely due to the FX global shock u^w :

$$rx_{t+1}^{HML} - E_t(rx_{t+1}^{HML}) = \frac{1}{N} \left(\sum_{i \in L} \sqrt{\gamma^i} - \sum_{i \in H} \sqrt{\gamma^i} \right) \sqrt{z_t^w} u_{t+1}^w.$$

In our model, the global pricing factor z^w is countercyclical and driven by u^w , so a positive innovation in u^w is also translated into a negative innovation in z^w . As a result, there is a negative relationship between HML returns and changes in our FX correlation factor FXC . However, our FX correlation factor is not fully subsumed by HML returns, as it also reflects fluctuations in the local pricing factor z . In fact, the correlation between changes in the FXC correlation factor and HML returns is -18% for our sample period.

Another factor shown to be priced is the FX variance factor (see Menkhoff, Sarno, Schmeling, and Schrimpf (2012)), defined as the cross-sectional average of all FX conditional variances:

$$FXV = \frac{1}{N} \sum_{i \in H} var_t(\Delta q_{t+1}^i) = 2\kappa z_t + \frac{1}{N} \sum_{i \in H} \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right)^2 z_t^w.$$

The FX variance factor is increasing in both the global and the local pricing factors. Since increases in the global pricing factor z^w increase both the FX variance factor and our FX correlation factor, the dependence on the global pricing factor generates a positive association between the two factors in both levels and changes. However, changes in the local pricing factor z engender a negative relationship between the two FX factors, as they increase the conditional variance of all FX pairs due to country-specific risk and therefore raise FXV , but they reduce the importance of heterogeneity in global risk exposure, tightening the cross section of FX correlation and, thus, decreasing our FX

correlation factor. In the data, the correlation between FXC and FCV is roughly 30% both in levels and changes.

Lastly, if we consider the domestic country to be the United States, we can characterize the U.S. dollar factor DOL , the returns to which are defined as:

$$rx_{t+1}^{DOL} = \frac{1}{N} \sum_i rx_{t+1}^i.$$

The innovations are given by

$$rx_{t+1}^{DOL} - E_t(rx_{t+1}^{DOL}) = -\frac{1}{N} \sum_i \left(\sqrt{\kappa z_t} u_{t+1}^i - \sqrt{\kappa z_t} u_{t+1}^0 + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right) \sqrt{z_t^w} u_{t+1}^w \right).$$

Invoking the law of large numbers $\left(\frac{1}{N} \sum_i u_{t+1}^i \rightarrow 0 \right)$, we can show that DOL return innovations reflect both U.S. local innovations and global innovations, the latter provided that for the U.S. global risk exposure γ^0 it holds that $\sqrt{\gamma^0} \neq \frac{1}{N} \sum_i \sqrt{\gamma^i}$:

$$x_{t+1}^{DOL} - E_t(x_{t+1}^{DOL}) = \sqrt{\kappa z_t} u_{t+1}^0 + \left(\sqrt{\gamma^0} - \frac{1}{N} \sum_i \sqrt{\gamma^i} \right) \sqrt{z_t^w} u_{t+1}^w.$$

The first component of DOL return innovations is purely U.S.-specific and, thus, uncorrelated with any of the innovations in HML returns, the FX variance factor or the FX correlation factor. This is because HML , FXV and FXC only reflect innovations in the two pricing factors, both of which are solely driven by global shocks. Thus, if $\sqrt{\gamma^0} = \frac{1}{N} \sum_i \sqrt{\gamma^i}$, DOL return innovations and innovations in the other three factors should be uncorrelated. In that case, and given that all currency returns have the same loading on U.S.-specific innovations, DOL returns constitute a level factor for currency returns. However, if $\sqrt{\gamma^0} \neq \frac{1}{N} \sum_i \sqrt{\gamma^i}$, the global pricing factor induces a positive association between DOL and the three aforementioned factors.

To summarize, DOL returns mostly act as a level factor in the cross section of currency returns, in congruence with earlier research. Exposure to global risk, which determines the cross section of currency returns, is best captured by exposure to HML returns, as this is the only factor completely uncontaminated by fluctuations in the local pricing factor z , the innovations to which are unpriced in currency markets. On the

other hand, innovations in both our FX correlation factor and the FX variance factor partly capture innovations in the global pricing factor z^w , which are priced in currency markets, and innovations in the local pricing factor z , which are unpriced in currency markets, and, thus, add noise for pricing purposes. Therefore, our model suggests that the cross section of currency returns line up with *HML* return betas better than with the betas of any of the two FX moment factors, and that both FX moment factors have a negative price of risk in line with our empirical findings.

3.5 The risk premia of FX moments

The negative cross-sectional relationship between average FX correlations and average FX correlation risk premia implies that domestic investors are attaching a high price to states of the world in which the cross section of FX correlations tightens; as discussed above, those are states characterized by high values of the local pricing factor. We now show that if the domestic agent prices fluctuations in the local pricing factor sufficiently more than fluctuations in the global pricing factor, then correlation risk premia identify bad states of the world as states in which adverse local pricing factor realizations occur. Intuitively, domestic agents risk-adjust by overweighing states in which the local pricing factor z is high and underweighting states in which the local pricing factor z is low. As we have seen, those are exactly the states in which similar FX pairs are less correlated than average and dissimilar FX pairs are more correlated than average.

In order to explore the properties of the FX correlation risk premia, we first need to characterize the law of motion of the pricing factors under the risk-neutral measure. From the perspective of the domestic investor, the law of motion for the global pricing factor z^w is

$$\Delta z_{t+1}^w = \lambda^w (\bar{z}^w - z_t^w) + \xi^w \sqrt{\gamma^0} z_t^w - z_t^w \xi^w \sqrt{z_t^w} u_{t+1}^{w,\mathbb{Q}}, \quad (6)$$

so the drift adjustment is positive and equal to $\xi^w \sqrt{\gamma^0} z_t^w$. We can rewrite equation (6) as a standard square root process,

$$\Delta z_{t+1}^w = \lambda^{w,\mathbb{Q}} (\bar{z}^{w,\mathbb{Q}} - z_t^w) - \xi^w \sqrt{z_t^w} u_{t+1}^{w,\mathbb{Q}},$$

where $\lambda^{w,\mathbb{Q}} \equiv \lambda^w - \xi^w \sqrt{\gamma^0}$ and $\bar{z}^{w,\mathbb{Q}} \equiv \frac{\lambda^w}{\lambda^{w,\mathbb{Q}}} \bar{z}^w$. Thus, under the risk-neutral measure the global pricing factor z^w has a higher unconditional mean ($\bar{z}^{0,\mathbb{Q}} \geq \bar{z}^0$) and is more persistent ($\lambda^{0,\mathbb{Q}} \leq \lambda^0$) than under the physical measure. Similarly, the risk-neutral measure law of motion for the local pricing factor z is given by

$$\Delta z_{t+1} = \lambda^{\mathbb{Q}}(\bar{z}^{\mathbb{Q}} - z_t) - \xi \sqrt{z_t} u_{t+1}^{g,\mathbb{Q}},$$

where $\lambda^{\mathbb{Q}} \equiv \lambda - \xi \sqrt{\delta}$ and $\bar{z}^{w,\mathbb{Q}} \equiv \frac{\lambda}{\lambda^{\mathbb{Q}}} \bar{z}$, so the local pricing factor also has a higher unconditional mean and persistence under the risk-neutral measure than under the physical measure. Notably, the relative risk adjustment of the two factors depends crucially on the volatility parameters ξ and ξ^w , as well as the exposure parameters δ and γ^0 . The higher ξ is compared to ξ^w , and the higher δ is relative to γ^0 , the higher the relative drift adjustment of the local pricing factor over the global pricing factor, as the shocks of the former become more highly priced compared to the shocks of the latter.

We can now turn to the risk premia of the conditional FX moments. First, we need to determine the risk-neutral FX moments. For the period $[t, T]$, the expected variance of the changes in the log exchange rate change for currency i is given by

$$E_t^{\mathbb{Q}} \left(\sum_{s=0}^{T-t-1} \text{var}_{t+s} (\Delta q_{t+s+1}^i) \right) = \sum_{s=0}^{T-t-1} E_t^{\mathbb{Q}} \left[2\kappa z_{t+s} + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right)^2 z_{t+s}^w \right],$$

and the expected covariance of the changes in log exchange rates i and j is

$$E_t^{\mathbb{Q}} \left(\sum_{s=0}^{T-t-1} \text{cov}_t (\Delta q_{t+1}^i, \Delta q_{t+1}^j) \right) = \sum_{s=0}^{T-t-1} E_t^{\mathbb{Q}} \left[\kappa z_{t+s} + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right) \left(\sqrt{\gamma^j} - \sqrt{\gamma^0} \right) z_{t+s}^w \right].$$

Finally, the expected FX correlation is defined as the corresponding expected FX covariance, adjusted by the product of the squared root of the two FX variances, as in the empirical section of our paper.

Note that for the local pricing factor we have

$$E_t^{\mathbb{Q}}(z_{t+s}) = \left(1 - (1 - \lambda^{\mathbb{Q}})^s \right) \bar{z}^{\mathbb{Q}} + (1 - \lambda^{\mathbb{Q}})^s z_t$$

under the risk-neutral measure, compared to

$$E_t(z_{t+s}) = (1 - (1 - \lambda)^s) \bar{z} + (1 - \lambda)^s z_t$$

under the physical measure. Given the higher steady-state and higher persistence of the local pricing factor under the risk-neutral measure, the wedge $E_t^{\mathbb{Q}}(z_{t+s}) - E_t(z_{t+s})$ is an affine function of z_t , with both the constant and the slope coefficient being positive: the wedge between the future value of the local pricing factor under the risk-neutral and the physical measure is always positive and increasing in its current value z_t .²³ Exactly the same is true for the global pricing factor z^w . We can write the FX correlation risk premium as

$$CRP_t^{i,j} = \frac{\kappa(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t) + (\sqrt{\gamma^i} - \sqrt{\gamma^0})(\sqrt{\gamma^j} - \sqrt{\gamma^0})(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w)}{\sqrt{2\kappa(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t) + (\sqrt{\gamma^i} - \sqrt{\gamma^0})^2(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w)}\sqrt{2\kappa(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t) + (\sqrt{\gamma^j} - \sqrt{\gamma^0})^2(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w)}} \cdot \frac{\kappa(A + Bz_t) + (\sqrt{\gamma^i} - \sqrt{\gamma^0})(\sqrt{\gamma^j} - \sqrt{\gamma^0})(A^w + B^wz_t^w)}{\sqrt{2\kappa(A + Bz_t) + (\sqrt{\gamma^i} - \sqrt{\gamma^0})^2(A^w + B^wz_t^w)}\sqrt{2\kappa(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t) + (\sqrt{\gamma^j} - \sqrt{\gamma^0})^2(A^w + B^wz_t^w)}}.$$

Thus, the magnitude of the correlation risk premium is determined by the disparity between the positive risk-neutral measure parameters $A^{\mathbb{Q}}$, $B^{\mathbb{Q}}$, $A^{w,\mathbb{Q}}$ and $B^{w,\mathbb{Q}}$ and the positive physical measure parameters A , B , A^w and B^w , with all the risk-neutral measure parameters being higher than their physical measure counterparts.

Of particular relevance is the case in which the domestic agent prices fluctuations in the local pricing factor more heavily than fluctuations in the global pricing factor, i.e., where $\xi\sqrt{\delta} \gg \xi^w\sqrt{\gamma^0}$. Then, the local pricing factor parameters will be considerably higher under the risk-neutral measure than under the physical measure, whereas the difference across measures will be small for the global pricing factor parameters. This has implications for both the cross section and the time series of FX correlation risk premia.

²³The constant in the wedge is positive because the constant is higher in the risk-neutral measure than in the physical measure due to the fact that the function $f(x) = \frac{1-(1-x)^s}{x}$ for $s > 1$ is decreasing for $x \in (0, 1)$.

First, the relative importance of the global pricing factor will be smaller under the risk-neutral measure than under the physical measure. Mathematically, if

$$(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t) - (A + Bz_t) \gg (A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w) - (A^w + B^wz_t^w),$$

then the FX correlation under the risk-neutral measure is essentially the physical measure FX correlation with a large upwards adjustment for the local pricing factor and a much smaller upwards adjustment for the global pricing factor. As Figure 7 suggests, this implies that the risk-neutral FX correlation is lower than the physical correlation for similar exchange rates, and higher than the physical correlation for dissimilar exchange rates. Thus, the cross section of average risk-neutral FX correlations is tighter than the cross section of average physical FX correlations, and high average correlation exchange rates tend to have a negative average CRP and low average correlation exchange rates tend to have positive average CRP. In sum, average FX correlations are negatively associated with average CRP in line with our empirical findings presented in Figure 1.

Second, although fluctuations in both the global pricing factor and the local pricing factor are amplified under the risk-neutral measure, the amplification is very small for changes in the global pricing factor z^w ; changes in z^w in the physical measure translate to roughly equal changes in z^w in the risk-neutral measure. Given the higher relative importance of the local pricing factor under the risk-neutral measure, same-size fluctuations in z^w typically have a smaller effect on risk-neutral measure conditional FX correlations than on physical measure FX correlations. This is because, as seen in the left panels of Figure 7, conditional FX correlation is a less steep function of z^w for higher values of z . Thus, changes in z^w , which mainly drive conditional FX correlations, move risk-neutral conditional FX correlations less than physical FX correlations. As a result, changes in physical FX correlations are associated with opposite sign changes in correlation risk premia for all FX pairs.

Conversely, we can use the same reasoning to determine that if the domestic agent attaches a higher relative price to z^w fluctuations than z^0 fluctuations, there will be a counter-factually positive cross-sectional relationship between average FX correlations and average FX correlation risk premia.

3.6 Calibration

We now turn to the calibration of our no-arbitrage model and show that it can match key moments of currency and correlation risk premia, as well as the standard interest rate and exchange rate moments.

The benchmark version of our model has $15 + (I + 1)$ parameters: five common SDF parameters (α , χ , ϕ , κ and δ), $I + 1$ heterogeneous parameters (the loading γ for each country), six pricing factor parameters—three for the local pricing factor (λ , \bar{z} and ξ) and three for the global pricing factor (λ^w , \bar{z}^w and ξ^w)—and, finally, four common inflation parameters ($\bar{\pi}$, ψ , ζ and σ). In our calibration, we follow Lustig, Roussanov, and Verdelhan (2011) and reduce the set of parameters by imposing the constraint that the loadings γ^i are equally spaced across all foreign countries. In particular, we assume that foreign country $i = 1$ has loading γ^{min} , foreign country $i = I$ has loading γ^{max} and each intermediate foreign country i , for $i = 2, \dots, I - 1$, has loading $\gamma^i = \frac{\gamma^{max} - \gamma^{min}}{I - 1}$. Therefore, we calibrate 18 parameters in total, the 15 common SDF, inflation and pricing factor parameters, as well as γ^{min} , γ^{max} and γ^0 . Given that we have options data (and therefore FX correlation risk premia time series) for only nine exchange rates, we limit our calibration to ten countries ($I=9$ foreign countries and the U.S.), consistent with our empirical analysis. Furthermore, we simulate the model at the monthly frequency, with conditional FX moments (realized and implied) calculated using conditional expectations over a period of 21 days (i.e., one month) into the future, with the model parameters appropriately adjusted to the daily frequency for the calculations.

In our calibration, we set the values of all SDF and inflation parameters, as well as the unconditional means of the local and the global pricing factor (\bar{z} and \bar{z}^w , respectively) equal to the corresponding values in Lustig, Roussanov, and Verdelhan (2011), which are set to match specific interest rate and exchange rate moments. Notably, the calibration in Lustig, Roussanov, and Verdelhan (2011) does not involve any moments related to FX correlations or FX correlation risk premia. We depart from that calibration as regards the values of λ , ξ , λ^w and ξ^w . The reason is that, as discussed above, for our model to generate correlation risk premia that exhibit behavior consistent with our empirical findings, it is necessary that shocks in the local pricing factor are priced much

more severely than shocks in the global pricing factor. In terms of our calibration, we need that ξ is significantly higher than ξ^w . To that end, we calibrate the four aforementioned pricing factor parameters using a grid search in a parameter space that has two properties: i) it contains relatively high values of ξ and relatively low values of ξ^w , and ii) it adjusts λ so that the unconditional volatility of the global pricing factor is identical to that in the Lustig, Roussanov, and Verdelhan (2011) calibration. Our calibrated parameters are those in the grid space that minimize the sum of squared percentage deviations of three moments: U.S. real exchange rate volatility, the cross-sectional average of foreign real interest rate volatility, and the cross-sectional average of real exchange rate volatility. In the final step, we adjust the constants α and $\bar{\pi}$ in order to exactly match the average U.S. real interest rate and the average U.S. inflation rate, respectively. Importantly, we do not use any of the moments related to either FX correlation or FX correlation risk premia in our calibration. We provide further details about the calibration in Appendix A.²⁴

The calibrated annualized values of our parameters are reported in Table 8. Note that the calibrated value of the local pricing factor volatility parameter ξ is 6.8 times its global pricing factor counterpart ξ^w . We use the calibrated parameter values to simulate our model for 50,000 monthly periods. We initialize the economy at the pricing factor steady-state values \bar{z} and \bar{z}^w and discard the first 5,000 observations in order to reduce the effect of initial conditions. Tables 9 and 10 present the simulation results.

[Insert Tables 8, 9 and 10 here.]

Table 9 reports moments for inflation, real and nominal interest rates and real and nominal exchange rates for the U.S. and the foreign countries. Given that the calibration of Lustig, Roussanov, and Verdelhan (2011), on which our calibration is largely based, targets inflation, interest rate and exchange rate moments and since our calibration adjustments, detailed above, also consider interest rate and exchange rate moments, it is not surprising that our model is doing very well on those dimensions. The unconditional mean and volatility of the U.S. and foreign real interest rates are matched almost

²⁴Interest rate differentials against the USD are proxied by the corresponding forward discounts. The nominal USD interest rate is proxied by the Fama-French 1-month Treasury Bill rate. Inflation in each country is constructed using the corresponding CPI.

perfectly, although the model undershoots the cross-sectional average of foreign mean real interest rates (0.29% versus 1.15%). In addition, the model produces an annualized real exchange rate volatility of 12.09%, in line with its empirical counterpart (10.82%). Moreover, simulated exchange rate changes exhibit no autocorrelation, as is true empirically. Inflation moments are also well-matched. Since the model matches both real and inflation moments, it follows that it is successful in quantitatively addressing nominal interest rate and exchange rate moments: indeed, all nominal moments are well-matched, although the model slightly overshoots both U.S. and foreign average nominal interest rate volatility (1.28% for the U.S. and 1.29% for the foreign countries, against empirical volatilities of 0.62% and 0.44%, respectively). Further, the model generates a carry trade effect: the return on the FX carry portfolio has an average excess return of 3.33%, compared to 5.41% in the data. Finally, our calibration generates realistic SDF properties: the cross-sectional average of the standard deviation of the log SDF (both real and nominal) is 0.54.

In Table 10, we turn to FX correlation moments, none of which have been used to calibrate the parameters. The first panel reports the cross-sectional minimum, mean, and maximum of average realized exchange rate correlations. Our model generates a reasonable spread of average FX correlation coefficients, albeit somewhat tighter than the empirical spread: the model values range from 11% to 61%, against a minimum of 5% and a maximum of 89% for the empirical distribution. Importantly, the cross-sectional average of mean FX correlations (42%) is very close to its empirical counterpart (45%). Not surprisingly, the same is true for implied FX correlations: the cross-sectional average of mean implied FX correlations is 43% in the model and 48% in the data. However, the model is able to capture the empirical fact that the cross-sectional distribution of average implied FX correlations is tighter than the cross-sectional distribution of realized FX correlations, a key feature for correlation risk premia. Furthermore, the model is able to replicate the empirically observed positive relationship between realized and implied correlations both in the cross section of time series averages and in the time series of individual realized and implied correlations, both in levels and changes. In particular, the cross-sectional correlation of average realized and implied correlations is almost perfect in the model, as it is in the data. Additionally, the model generates realized and implied

correlations that move together almost perfectly, both in levels and in changes—this captures both the qualitative and the quantitative aspect of their empirical comovement in levels (79% in the data), but overstates it in changes (28% in the data).

The third panel of Table 10 reports moments for FX correlation risk premia. The model is able to successfully replicate the negative association observed in the data between realized correlations and correlation risk premia in two dimensions: the negative cross-sectional correlation between average realized correlations and average correlation risk premia, and the negative time series correlation between realized correlations and correlation risk premia both in levels and in changes for virtually all FX pairs. The only weakness of the calibrated model regards the magnitude of correlation risk premia: the model-implied correlation risk premium are lower (in absolute terms) than their empirical counterparts. However, our model is able to successfully generate both positive and negative correlation risk premia. In particular, high average correlation FX pairs have negative average correlation risk premia and low average correlation FX pairs have positive correlation risk premia in the model as in the data. Finally, the model generates a negative price of risk for exposure to the correlation risk factor FXC : the annualized average excess return for the currency portfolio that is long currencies with high exposure to the correlation risk factor and short currencies with a low exposure is -1.08% .

3.7 The importance of local pricing factor comovement

The key innovation in our reduced-form model compared to earlier research is that we allow for the local pricing factors to comove across countries. Indeed, in our benchmark model we shut down any cross-sectional heterogeneity in the pricing of local risk. To illustrate the importance of comovement in local pricing factors, we relax the assumption that all local pricing factors are identical, achieved by imposing that $\rho = 0$, and allow heterogeneity in local pricing factors across countries.

Recall the dynamics of the local pricing factor given in equation (3). If $\rho > 0$, local pricing factors, although ex ante identical, have different realizations in different countries due to the independence of the local shocks. As a result, countries have

different conditional SDF loadings in the global innovation u^g and the exposure to u^g now enters the expression for real exchange rate changes:

$$\Delta q_{t+1}^i = E_t(\Delta s_{t+1}^i) + \sqrt{\kappa z_t^i} u_{t+1}^i - \sqrt{\kappa z_t^0} u_{t+1}^0 + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right) \sqrt{z_t^w} u_{t+1}^w + \sqrt{\delta} \left(\sqrt{z_t^i} - \sqrt{z_t^0} \right) u_{t+1}^g.$$

Under the risk-neutral measure, the law of motion for the global pricing factor z^w is given by equation (6), as in the benchmark model, whereas the local pricing factors z^i , for $i = 0, 1, \dots, I$ satisfy

$$\Delta z_{t+1}^i = \lambda^{i,\mathbb{Q}} (\bar{z}^{i,\mathbb{Q}} - z_t^i) - \xi \sqrt{z_t^i} \left(\sqrt{\rho} u_{t+1}^{i,\mathbb{Q}} + \sqrt{1-\rho} u_{t+1}^{g,\mathbb{Q}} \right),$$

where $\bar{z}^{i,\mathbb{Q}} \equiv \frac{\lambda^{i,\mathbb{Q}}}{\lambda^{0,\mathbb{Q}}} z^i$. Note that $\lambda^{0,\mathbb{Q}} \equiv \lambda - \xi \left(\sqrt{\kappa} \sqrt{\rho} + \sqrt{1-\rho} \sqrt{\delta} \right)$, as both components of the innovations in the domestic pricing factor z^0 are priced by the domestic investor, whereas for $i = 1, \dots, I$ we have $\lambda^{i,\mathbb{Q}} \equiv \lambda - \xi \sqrt{1-\rho} \sqrt{\delta}$ as only the global component of the foreign pricing factor innovations is priced by the domestic investor.

For convenience, consider the polar case of $\rho = 1$, in which case local pricing factors are i.i.d. random variables; this is the assumption in Lustig, Roussanov, and Verdelhan (2011, 2014) and Verdelhan (2013). In that case, the domestic investor prices the shocks to z^w and z^0 , but not the innovations to the foreign local pricing factors. Thus, to understand implied FX correlations and FX correlation risk premia, we only need to consider the dependence of conditional FX correlations on z^w and z^0 . Appendix B provides a detailed discussion of that dependence.

In order to explore the quantitative aspects of the full model with independent local pricing factors, we run a simulation with 55,000 monthly periods and eliminate the first 5,000, and set all parameter values equal to the calibrated values in Table 8, with the only exception being that we set $\rho = 1$ instead of $\rho = 0$. FX correlation moments of interest are reported in Table 11. As expected, all correlation risk premia are positive (and small) and there is an almost perfect positive association between average FX correlation and average FX correlation risk premia. In general, as the degree of local pricing factor comovement declines, so do FX correlation risk premia. For example, setting $\rho = 0.5$ while keeping the rest of the parameters equal to their calibrated values generates FX

correlation risk premia ranging from -0.12% to 0.38% , although the relative pricing of the local factor is still strong enough to generate a negative cross-sectional relationship between average FX correlation and average FX correlation risk premia.

In sum, introducing cross-country comovement in local pricing factors allows for explaining the joint time series and cross-sectional behavior of FX correlations and FX correlation risk premia, while preserving all the previously studied properties of earlier models, as those do not crucially depend on the relative pricing of local pricing factor innovations over global pricing factor innovations.

[Insert Table 11 here.]

4 Conclusion

We document large cross-sectional differences in average FX correlations and average FX correlation risk premia, and show that there is a negative association between the two in the cross section of FX pairs. Furthermore, we show there is a negative relationship between FX correlations and FX correlation risk premia, both in levels and in changes, in the time series of virtually all FX pairs, suggesting that implied FX correlation is less responsive to shocks than physical FX correlation. Finally, we find a negative cross-sectional relationship between average FX correlations and correlation cyclicity, implying that FX pairs that are highly correlated on average become even more correlated in bad times while pairs characterized by low average correlation become even less correlated. Thus, FX correlations become more dispersed in adverse economic states.

We capture the countercyclicality of cross-sectional dispersion in conditional FX correlation by constructing the FX correlation factor FXC , defined as the difference between the average conditional correlation of the most and least conditionally correlated FX pairs. We then sort currencies into portfolios based on their exposure to our correlation factor and show that the spread between high and low FXC beta currency portfolios is economically and statistically large, ranging between 4% and 6.4% , depending on the set of currencies, for our benchmark sample period of January 1996 to December 2013. For the same time period, we estimate the price of FX correlation risk

to be almost -6% per year. In short, we find that investors want to be compensated for investing in currencies that perform badly during periods of increased cross-sectional disparity in conditional FX correlations.

We rationalize our findings with a no-arbitrage model of exchange rates that is able to replicate the salient empirical time series and cross sectional properties of FX correlations and FX correlation risk premia, and show the importance of cross-country comovement in the price of local risk. Thus, richer models that feature endogenously determined stochastic discount factors and aim to explain the joint dynamics of FX correlations under the physical and under the risk-neutral measure need to feature comovement in the pricing of not just common, but also country-specific shocks.

References

- BOLLERSLEV, T. G., G. TAUCHEN, AND H. ZHOU (2009): “Expected Stock Returns and Variance Risk Premia,” *Review of Financial Studies*, 22, 4463–4492.
- BRANDT, M. W., AND F. X. DIEBOLD (2006): “A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations,” *Journal of Business*, 79, 61–73.
- 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, 313–347.
- BURNSIDE, C. (2011): “Carry Trades and Risk,” NBER Working Paper 17278.
- CENEDESE, G., L. SARNO, AND I. TSIAKAS (2014): “Foreign Exchange Risk and the Predictability of Carry Trade Returns,” *Journal of Banking and Finance*, 42, 302–313.
- DELLA CORTE, P., T. RAMADORAI, AND L. SARNO (2014): “Volatility Risk Premia and Exchange Rate Predictability,” Working Paper, Imperial Business School.
- DEMETERFI, K., E. DERMAN, M. KAMAL, AND J. ZHOU (1999): “A Guide to Volatility and Variance Swaps,” *Journal of Derivatives*, 6, p. 9 – 32.
- DOBRYNSKAYA, V. (2014): “Downside Market Risk of Carry Trades,” *Review of Finance*, 24, 1–29.
- ENGLE, R. F. (2002): “Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models,” *Journal of Business and Economic Statistics*, 20, 339–350.
- FAMA, E., AND J. MACBETH (1973): “Risk, Return, and Equilibrium: Empirical Tests,” *Journal of Political Economy*, 81, 607–636.
- FAMA, E. F. (1984): “Forward and Spot Exchange Rates,” *Journal of Monetary Economics*, 14, 319–338.
- FENN, D. J., S. D. HOWISON, M. McDONALD, S. WILLIAMS, AND N. F. JOHNSON (2009): “The Mirage Of Triangular Arbitrage In The Spot Foreign Exchange Market,” *International Journal of Theoretical and Applied Finance*, 12, 1105–1123.
- GARMAN, M. B., AND S. W. KOHLHAGEN (1983): “Foreign Currency Option Values,” *Journal of International Money and Finance*, 2, 231–237.
- HU, G. X., J. PAN, AND J. WANG (2013): “Noise as Information for Illiquidity,” *Journal of Finance*, 68, 2341–2382.
- JIANG, G., AND Y. TIAN (2005): “Model-Free Implied Volatility and Its Information Content,” *Review of Financial Studies*, 18, 1305–1342.
- JUREK, J. W. (2014): “Crash-Neutral Currency Carry Trades,” *Journal of Financial Economics*, 113, 325–347.
- KOZHAN, R., AND W. W. THAM (2012): “Execution Risk in High-Frequency Arbitrage,” *Management Science*, 58, 2131–2149.

- LETTAU, M., M. MAGGIORI, AND M. WEBER (2014): “Conditional risk premia in currency markets and other asset classes,” *Journal of Financial Economics*, 114, 197–225.
- LUSTIG, H., N. ROUSSANOV, AND A. VERDELHAN (2011): “Common Risk Factors in Currency Markets,” *Review of Financial Studies*, 24, 3731–3777.
- (2014): “Countercyclical Currency Risk Premia,” *Journal of Financial Economics*, 111, 527–553.
- LUSTIG, H., AND A. VERDELHAN (2007): “The Cross-Section of Foreign Currency Risk Premia and Consumption Growth Risk,” *American Economic Review*, 97, 89–117.
- MALKHOZOV, A., P. MUELLER, A. VEDOLIN, AND G. VENTER (2014): “International Funding Illiquidity,” Working Paper, London School of Economics.
- MANCINI, L., A. RANALDO, AND J. WRAMPMEYER (2013): “Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums,” *Journal of Finance*, 68, 1805–1841.
- MENKHOFF, L., L. SARNO, M. SCHMELING, AND A. SCHRIMPF (2012): “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance*, 67, 681–718.
- MUELLER, P., A. VEDOLIN, AND H. CHOI (2014): “Bond Variance Risk Premiums,” Working Paper, London School of Economics.
- NEWKEY, W. K., AND K. D. WEST (1987): “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55(3), 703–708.
- SHANKEN, J. (1992): “On the Estimation of Beta Pricing Models,” *Review of Financial Studies*, 5, 1–34.
- VERDELHAN, A. (2013): “The Share of Systematic Risk in Bilateral Exchange Rates,” Working Paper, MIT.
- WYSTUP, U. (2006): *FX Options and Structured Products*. Chichester: John Wiley and Sons.

Appendix A

In the first step of our calibration, we perform a grid search that considers values of the local pricing factor volatility parameter ξ ranging from 0.005561 to 0.031893 and values of the local pricing factor volatility parameter ξ^w ranging from 0.000900 to 0.005561. Thus, we restrict our search in the subset of the parameter space that ensures that ξ is at least as high as ξ^w , following our model intuition. For comparison, Lustig, Roussanov, and Verdelhan (2011) use a common value of 0.002500 for both ξ and ξ^w . We also vary the value of the mean-reversion parameters λ and λ^w in order to ensure that: i) the unconditional volatility of the local pricing factor z ranges between 1 and 5 times its unconditional volatility in the Lustig, Roussanov and Verdelhan (2011) calibration, and that ii) the unconditional volatility of the global pricing factor z^w is exactly equal to its unconditional volatility in the Lustig, Roussanov and Verdelhan (2011) calibration. Therefore, although we vary λ , we severely constrain the range of values it can achieve, given ξ . Furthermore, the value of λ^w is exactly determined by the value of ξ^w though the expression for the unconditional volatility of the global pricing factor, so variation in λ^w provides no extra degree of freedom in our calibration.

In the second step of our calibration, we select the set of parameter values that minimizes the expression

$$J = \left(\frac{\hat{\beta}_1 - \beta_1}{\beta_1} \right)^2 + \left(\frac{\hat{\beta}_2 - \beta_2}{\beta_2} \right)^2 + \left(\frac{\hat{\beta}_3 - \beta_3}{\beta_3} \right)^2 \quad (\text{A-1})$$

where β_1 , β_2 and β_3 refer to the empirical values of U.S. real exchange rate volatility, the cross-sectional average of foreign real interest rate volatility, and the cross-sectional average of real exchange rate volatility, respectively, and $\hat{\beta}_1$, $\hat{\beta}_2$ and $\hat{\beta}_3$ refer to the simulated values of the aforementioned moments.

Appendix B

We illustrate the impact of changes in z^w and z^0 on the conditional FX correlations in the full model in Figure 8. As in Figure 7, we study a world of three foreign countries. Countries 1 and 2 are less exposed to the first global shock u^w than the domestic country, while country 3 is more exposed. The left panels of Figure 8 depict the conditional FX correlations as a function of the global pricing factor z^w for different values of the domestic local pricing factor z^0 , holding all the foreign local pricing factors equal to their common steady-state value \bar{z} . Not surprisingly, in the full model the impact of changes in the global pricing factor z^w is the same as in the benchmark model that features identical local pricing factors: as z^w increases, similarities and dissimilarities in exposure to global risk get amplified.

The right panels of Figure 8 present the conditional FX correlations as a function of z^0 for different values of z^w , assuming that all other pricing factors are equal to their steady-state values. Notably, the relationship between z^0 and the conditional FX correlation is not monotonic. For small values of z^0 , the conditional FX correlation is higher than its steady-state value for all FX pairs. This is because all FX pairs are similar regarding their exposure to the second global shock u^g : the loading of all foreign countries is lower than the domestic loading. In that range, as the value of z^0 increases, the value of the FX correlation decreases, since higher values of z^0 reduce the FX correlation arising from exposure to u^g . When $z^0 = \bar{z}$, all local factors have identical values, so exposure to u^g does not affect FX moments. Finally, for large values of z^0 , increases in z^0 increase the similarity of all FX pairs regarding the exposure to u^g , as the domestic loading becomes much higher than all foreign loadings.

[Insert Figure 8 here.]

Crucially, changes in z^0 push the conditional correlation of all FX pairs, similar and dissimilar, in the same direction. This implies that, in the absence of comovement in local pricing factors across countries, the model's ability to match the negative cross-sectional relationship between average FX correlations and average FX correlation risk premia is severely hindered. Recall that in the benchmark model, the desired cross-sectional pattern was achieved by overweighing states of the world characterized by high values of the common local pricing factor. However, Figure 8 shows that overweighing states in which the domestic pricing factor z^0 has a high value does not generate the desired cross-sectional pattern when only z^0 fluctuations are priced. Indeed, in the absence of global pricing effects, such a risk adjustment leads to positive correlation risk premia for all FX pairs. As a result, local pricing effects are not likely to be able to overturn the positive cross-sectional relationship between FX correlations and FX correlation risk premia generated by global pricing effects.²⁵

²⁵The Online Appendix considers a global economy in which the loading of all countries to the second global shock is equal to zero, i.e., $\delta = 0$, and shows that local pricing effects do not have the desired properties in that case either.

Appendix C Tables

Table 1
Summary statistics for G10 currencies and carry trade portfolios

This table reports summary statistics for the G10 currencies (Panel A) and currency portfolios sorted on forward discounts observed at the end of period t (Panels B and C). In Panel A currencies are sorted on the average forward discount (first row, annualized expressed in percent). Portfolio 1 (Pf1) contains 25% of all the currencies with the lowest forward discounts whereas Portfolio 4 (Pf4) contains currencies with the highest forward discounts. Panel B reports the summary statistics for the full set of currencies, whereas in Panel C only developed countries are considered. All returns are excess returns in USD, annualized and expressed in percent. *DOL* denotes the average return of the four currency portfolios, *HML* denotes a long-short portfolio that is short in Pf1 and invests in Pf4. Data is monthly and runs from January 1996 through December 2013. Before 1999 we use the DEM instead of the EUR.

PANEL A: G10 COUNTRIES									
	JPY	CHF	EUR	SEK	CAD	GBP	NOK	AUD	NZD
$f_t - s_t$	-3.01	-2.00	-0.60	-0.10	-0.04	0.91	0.98	2.12	2.70
Mean	-2.74	-0.39	-0.46	0.22	1.12	1.37	1.17	3.01	3.73
StDev	10.78	10.91	10.25	11.22	8.50	8.50	11.15	12.78	13.09
Skewness	0.48	0.13	-0.15	-0.08	-0.60	-0.50	-0.36	-0.60	-0.37
Kurtosis	5.22	4.40	3.80	3.61	7.26	4.73	4.10	5.29	4.85

PANEL B: CARRY PORTFOLIOS ALL COUNTRIES						
	Pf1	Pf2	Pf3	Pf4	DOL	HML
Mean	-1.76	0.83	2.51	4.32	1.47	6.08
StDev	7.00	7.24	8.17	9.81	7.33	7.66
Skewness	-0.07	-0.46	-0.80	-0.92	-0.64	-0.54
Kurtosis	3.01	4.66	5.66	6.14	4.89	4.34
Sharpe Ratio	-0.25	0.11	0.31	0.44	0.20	0.79

PANEL C: CARRY PORTFOLIOS DEVELOPED COUNTRIES						
	Pf1	Pf2	Pf3	Pf4	DOL	HML
Mean	-1.30	0.29	1.39	4.07	1.11	5.37
StDev	8.58	8.40	9.86	11.56	8.48	9.71
Skewness	0.24	-0.11	-0.58	-0.39	-0.26	-0.53
Kurtosis	3.19	3.95	6.19	5.15	4.15	4.97
Sharpe Ratio	-0.15	0.03	0.14	0.35	0.13	0.55

Table 2
Summary Statistics for variances and variance risk premia

This table reports summary statistics for implied and realized variances (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 using past daily log exchange rate changes over a three month window. Variances and variance risk premia are monthly and expressed in squared percent. Data is monthly and runs from January 1996 to December 2013 (options data for EUR starts in January 1999).

PANEL A: IMPLIED VARIANCE									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	13.32	6.97	10.49	10.54	7.71	11.94	12.73	14.84	12.97
StDev	14.63	7.65	5.62	7.87	7.33	8.59	10.05	11.97	10.52
Skewness	6.74	4.26	3.25	3.91	5.22	3.95	4.17	4.15	3.93
Kurtosis	65.86	26.34	19.89	24.87	37.06	28.87	26.29	29.47	22.95
AC(1)	0.69	0.86	0.73	0.77	0.80	0.68	0.80	0.74	0.82

PANEL B: REALIZED VARIANCE									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	14.52	6.55	10.07	8.34	7.05	10.44	11.97	15.20	12.02
StDev	24.57	8.57	7.28	5.70	7.67	7.86	11.53	16.60	12.01
Skewness	6.51	4.48	3.88	3.49	4.69	3.05	3.78	4.48	3.97
Kurtosis	49.68	28.15	20.78	19.48	26.91	15.17	19.44	27.16	21.11
AC(1)	0.84	0.91	0.84	0.92	0.93	0.86	0.93	0.89	0.94

PANEL C: VARIANCE RISK PREMIUM									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	-1.19	0.42	0.43	1.72	0.66	1.50	0.77	-0.36	0.95
StDev	14.74	3.16	5.35	4.58	3.77	5.81	5.79	9.09	6.47
Skewness	-8.68	-1.32	-1.96	5.47	3.05	1.90	-0.19	-5.66	0.18
Kurtosis	84.82	34.11	20.07	54.18	45.93	32.81	19.86	47.94	30.79
AC(1)	0.60	0.45	0.58	0.45	0.46	0.33	0.57	0.52	0.54

Table 3
Summary statistics for correlation risk premia

This table reports means and standard deviations for implied and realized correlations, as well as correlation risk premia for all FX pairs. Correlation risk premia (CRP) are defined as the difference between the implied and realized correlations. Implied correlations (IC) are calculated from daily option prices on the underlying exchange rates. Realized correlations (RC) are calculated using past daily log exchange rate changes over a three month window. Correlations and correlation risk premia are expressed in decimals. Data is monthly and runs from January 1996 to December 2013 (options data for EUR starts in January 1999).

Currency		IC		RC		CRP		t-stat
		Mean	StDev	Mean	StDev	Mean	StDev	
AUD	CAD	0.430	0.274	0.471	0.246	-0.041	0.148	-4.07
AUD	CHF	0.405	0.200	0.357	0.273	0.048	0.149	4.73
AUD	EUR	0.544	0.162	0.450	0.276	0.019	0.089	2.81
AUD	GBP	0.453	0.186	0.422	0.241	0.031	0.119	3.86
AUD	JPY	0.238	0.257	0.155	0.341	0.083	0.161	7.58
AUD	NOK	0.431	0.294	0.467	0.258	-0.036	0.201	-2.64
AUD	NZD	0.739	0.146	0.755	0.164	-0.016	0.081	-2.97
AUD	SEK	0.480	0.202	0.474	0.255	0.005	0.127	0.61
CAD	CHF	0.283	0.211	0.233	0.283	0.050	0.149	4.94
CAD	EUR	0.405	0.187	0.307	0.298	0.024	0.133	2.45
CAD	GBP	0.307	0.231	0.281	0.269	0.025	0.153	2.34
CAD	JPY	0.136	0.190	0.054	0.259	0.082	0.164	7.33
CAD	NOK	0.341	0.279	0.340	0.276	-0.002	0.184	-0.17
CAD	NZD	0.352	0.342	0.413	0.234	-0.061	0.215	-4.19
CAD	SEK	0.287	0.290	0.352	0.261	-0.069	0.170	-5.96
CHF	EUR	0.875	0.117	0.888	0.135	-0.010	0.078	-1.69
CHF	GBP	0.605	0.146	0.580	0.190	0.025	0.110	3.32
CHF	JPY	0.456	0.179	0.405	0.255	0.051	0.145	5.15
CHF	NOK	0.731	0.119	0.726	0.160	0.006	0.110	0.73
CHF	NZD	0.370	0.199	0.358	0.230	0.012	0.162	1.06
CHF	SEK	0.712	0.133	0.707	0.159	0.004	0.105	0.58
EUR	GBP	0.683	0.101	0.644	0.153	0.003	0.079	0.54
EUR	JPY	0.364	0.195	0.324	0.271	0.067	0.154	5.84
EUR	NOK	0.798	0.071	0.825	0.095	-0.025	0.065	-5.20
EUR	NZD	0.501	0.173	0.440	0.230	0.005	0.125	0.55
EUR	SEK	0.817	0.077	0.816	0.113	-0.022	0.064	-4.64
GBP	JPY	0.293	0.188	0.217	0.264	0.076	0.152	7.29
GBP	NOK	0.638	0.120	0.577	0.155	0.059	0.161	5.39
GBP	NZD	0.404	0.219	0.415	0.225	-0.011	0.136	-1.15
GBP	SEK	0.598	0.133	0.560	0.165	0.037	0.129	4.26
JPY	NOK	0.347	0.213	0.248	0.260	0.099	0.158	9.22
JPY	NZD	0.233	0.237	0.146	0.320	0.087	0.180	7.09
JPY	SEK	0.294	0.197	0.241	0.271	0.052	0.155	4.95
NOK	NZD	0.413	0.274	0.449	0.222	-0.036	0.198	-2.65
NOK	SEK	0.780	0.106	0.796	0.102	-0.016	0.081	-2.93
NZD	SEK	0.403	0.270	0.439	0.231	-0.036	0.181	-2.89

Table 4
Time series correlations for RC, IC, and CRP

This table reports the time series correlations between realized correlations (RC) and implied correlations (IC) on one hand (columns one and two), and between realized correlations and correlation risk premia (CRP) on the other hand (columns three and four) for all FX pairs. Correlation risk premia (CRP) are defined as the difference between the implied and realized correlations. Implied correlations (IC) are calculated from daily option prices on the underlying exchange rates. Realized correlations (RC) are calculated using past daily log exchange rate changes over a three month window. Time series correlations are expressed in decimals and they are measured using monthly data from January 1996 to December 2013 (options data for EUR starts in January 1999).

Currency		Correlation RC/IC		Correlation RC/CRP	
		Level	Changes	Level	Changes
AUD	CAD	0.843	0.248	-0.102	-0.676
AUD	CHF	0.844	0.348	-0.695	-0.773
AUD	EUR	0.923	0.424	-0.714	-0.698
AUD	GBP	0.876	0.211	-0.656	-0.710
AUD	JPY	0.892	0.503	-0.695	-0.641
AUD	NOK	0.744	0.101	-0.213	-0.721
AUD	NZD	0.872	0.319	-0.457	-0.686
AUD	SEK	0.870	0.191	-0.618	-0.618
CAD	CHF	0.856	0.180	-0.684	-0.769
CAD	EUR	0.864	0.247	-0.702	-0.791
CAD	GBP	0.825	0.266	-0.518	-0.719
CAD	JPY	0.777	0.361	-0.680	-0.686
CAD	NOK	0.780	0.142	-0.316	-0.633
CAD	NZD	0.784	0.066	0.161	-0.574
CAD	SEK	0.813	0.029	-0.137	-0.598
CHF	EUR	0.846	0.495	-0.603	-0.354
CHF	GBP	0.816	0.315	-0.640	-0.634
CHF	JPY	0.835	0.323	-0.733	-0.520
CHF	NOK	0.725	0.232	-0.671	-0.600
CHF	NZD	0.724	0.270	-0.532	-0.733
CHF	SEK	0.757	0.284	-0.560	-0.425
EUR	GBP	0.774	0.381	-0.592	-0.621
EUR	JPY	0.858	0.423	-0.760	-0.567
EUR	NOK	0.704	0.350	-0.632	-0.426
EUR	NZD	0.770	0.224	-0.467	-0.692
EUR	SEK	0.721	0.368	-0.549	-0.537
GBP	JPY	0.824	0.405	-0.713	-0.594
GBP	NOK	0.332	0.293	-0.710	-0.595
GBP	NZD	0.812	0.234	-0.350	-0.597
GBP	SEK	0.644	0.273	-0.615	-0.551
JPY	NOK	0.795	0.394	-0.572	-0.620
JPY	NZD	0.831	0.396	-0.680	-0.664
JPY	SEK	0.825	0.379	-0.699	-0.582
NOK	NZD	0.699	0.051	-0.157	-0.659
NOK	SEK	0.701	0.219	-0.347	-0.403
NZD	SEK	0.750	0.088	-0.158	-0.699

Table 5
Correlations for risk factors

This table reports the correlation coefficients between the correlation factors FXC and FXC^{UNC} and the global equity volatility measure used in Lustig, Roussanov, and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin, and Venter (2014) (GFI), the CBOE VIX (VIX), and the TED spread (TED). Data is monthly and runs from January 1996 through December 2013.

	FXC	FXC^{UNC}	$GVol$	GFI	TED	VIX
FXC	1.00	0.86	0.35	0.48	0.42	0.45
FXC^{UNC}	0.86	1.00	0.26	0.44	0.41	0.39
$GVol$	0.35	0.26	1.00	0.53	0.59	0.81
GFI	0.48	0.44	0.53	1.00	0.57	0.61
TED	0.42	0.41	0.59	0.57	1.00	0.43
VIX	0.45	0.39	0.81	0.61	0.43	1.00

Table 6
Summary statistics for correlation portfolios

Panel A reports summary statistics for G10 currency portfolios sorted on exposure to the correlation risk factor FXC observed in period t . The exposure is measured by regressing currency excess returns on innovations in the correlation risk factor over the preceding 36 months. Portfolio 1 (Pf1) contains the three currencies with the lowest pre-sort correlation betas whereas Portfolio 3 (Pf3) contains the three currencies with the highest pre-sort correlation betas. Panels B and C report the summary statistics for the four correlation sorted portfolios using full set of currencies and the set developed currencies, respectively. All returns are excess returns in USD, annualized and expressed in percent. HML^C , HML^{C-ALL} and HML^{C-DEV} denote the long-short portfolios that invest in the high correlation beta currencies (Pf3 or Pf4, respectively) and short the low correlation beta currencies (Pf1) for the different sets of currencies. Data is monthly and runs from January 1996 through December 2013.

PANEL A: G10 COUNTRIES					
	Pf1	Pf2	Pf3	HML^C	
Mean	4.04	0.99	-2.38	-6.42	
StDev	10.26	9.11	7.86	7.83	
Skewness	-0.66	0.06	0.01	0.44	
Kurtosis	6.57	3.53	3.09	4.75	
Sharpe Ratio	0.39	0.11	-0.30	-0.82	

PANEL B: ALL COUNTRIES					
	Pf1	Pf2	Pf3	Pf4	HML^{C-ALL}
Mean	3.92	1.63	0.28	-0.12	-4.04
StDev	10.52	8.60	6.87	6.13	8.72
Skewness	-0.95	-0.58	-0.30	-0.43	0.70
Kurtosis	6.23	5.69	3.77	4.08	5.09
Sharpe Ratio	0.37	0.19	0.04	-0.02	-0.46

PANEL C: DEVELOPED COUNTRIES					
	Pf1	Pf2	Pf3	Pf4	HML^{C-DEV}
Mean	4.29	2.31	-1.88	-1.36	-5.64
StDev	11.02	9.96	8.93	8.03	9.64
Skewness	-0.58	-0.20	-0.14	0.19	0.80
Kurtosis	7.18	4.02	3.73	3.61	6.80
Sharpe Ratio	0.39	0.23	-0.21	-0.17	-0.59

Table 7
Estimating the price of correlation risk: HML^C

Test assets are the four portfolios sorted based on exposure to the correlation risk factor FXC from Table 6 based on either all countries or developed countries only. HML^C is the high minus low portfolio constructed using the G10 currencies. Panel A reports factor betas and Newey and West (1987) standard errors (in parentheses) while Panel B reports the Fama and MacBeth (1973) factor prices and standard errors (in parentheses). Shanken (1992)-corrected standard errors are reported in brackets. Data is monthly and runs from January 1996 through December 2013.

PANEL A: FACTOR BETAS								
	ALL COUNTRIES				DEVELOPED COUNTRIES			
	α	DOL	HML^C	R^2	α	DOL	HML^C	R^2
Pf1	0.07	1.17	-0.34	0.89	0.00	1.11	-0.52	0.88
	(0.06)	(0.03)	(0.03)		(0.08)	(0.04)	(0.04)	
Pf2	-0.02	1.05	-0.11	0.88	0.04	1.15	-0.06	0.74
	(0.06)	(0.04)	(0.03)		(0.10)	(0.06)	(0.05)	
Pf3	0.02	0.87	0.14	0.79	-0.06	1.15	0.17	0.80
	(0.06)	(0.04)	(0.03)		(0.09)	(0.06)	(0.04)	
Pf4	0.09	0.76	0.31	0.75	0.11	0.96	0.53	0.75
	(0.06)	(0.04)	(0.03)		(0.09)	(0.05)	(0.04)	

PANEL B: FACTOR PRICES					
ALL COUNTRIES			DEVELOPED COUNTRIES		
DOL	HML^C	R^2	DOL	HML^C	R^2
0.14	-0.49	0.96	0.12	-0.48	0.95
(0.15)	(0.22)		(0.15)	(0.16)	
[0.15]	[0.23]		[0.15]	[0.17]	

Table 8
Parameter values

This table reports the annualized parameter values for the calibrated version of the model. All countries share the same parameter values except for γ : γ^0 is the parameter for the home country, whereas the values for γ^i are linearly spaced on the interval $[\gamma^{min}, \gamma^{max}]$.

SDF PARAMETERS							
α	χ	ϕ	κ	δ	γ^0	γ^{min}	γ^{max}
0.0077	2.78	2.78	0.65	16.04	12.84	8.35	17.34
FACTOR DYNAMICS							
λ	\bar{z}	ξ	λ^w	\bar{z}^w	ξ^w		
0.93	0.000781	0.025453	0.19	0.000781	0.003741		
INFLATION DYNAMICS							
$\bar{\pi}$	ψ	ζ	$\sqrt{\sigma}$				
-0.0054	0	9.41	0.0027				

Table 9
Simulated moments: Interest rates, inflation, and exchange rates

In the first panel, we report annualized means and standard deviations of real US interest rates and cross-sectional averages of mean and standard deviation of foreign (FGN) interest rates in the data and in the model. The second panel reports cross-sectional averages of exchange rate volatility and autocorrelation. The third panel reports average and standard deviation of US inflation and cross-sectional averages of mean and standard deviation of foreign inflation. The fourth panel reports annualized means and standard deviations of nominal US interest rates and cross-sectional averages of mean and standard deviation of foreign (FGN) interest rates. The fifth panel reports cross-sectional averages of the volatility and autocorrelation of nominal exchange rates. The sixth panel reports the average excess return on the carry factor and the last panel reports the cross-sectional average of the standard deviation of the real and nominal log SDF.

MOMENT	DATA	MODEL
$E(r^{U.S.})$	0.28%	0.24%
$Std(r^{U.S.})$	1.35%	1.25%
$E_{\text{cross}}(E(r^{FGN}))$	1.15%	0.24%
$E_{\text{cross}}(Std(r^{FGN}))$	1.19%	1.26%
$E_{\text{cross}}(Std(\Delta q_{t+1}))$	10.82%	12.13%
$E_{\text{cross}}(AC(\Delta q_{t+1}))$	-0.01	0.00
$E(\pi^{U.S.})$	2.32%	2.32%
$Std(\pi^{U.S.})$	1.27%	1.10%
$E_{\text{cross}}(E(\pi^{FGN}))$	1.56%	2.33%
$E_{\text{cross}}(Std(\pi^{FGN}))$	1.12%	1.10%
$E(r^{\$,U.S.})$	2.60%	2.56%
$Std(r^{\$,U.S.})$	0.62%	1.28%
$E_{\text{cross}}(E(r^{\$,FGN}))$	2.70%	2.55%
$E_{\text{cross}}(Std(r^{\$,FGN}))$	0.44%	1.29%
$E_{\text{cross}}(Std(\Delta s_{t+1}))$	10.76%	12.20%
$E_{\text{cross}}(AC(\Delta s_{t+1}))$	0.01	0.00
$E(rx_{t+1}^{HML})$	5.41%	3.33%
$E_{\text{cross}}(Std(m_{t+1}))$	-	0.54
$E_{\text{cross}}(Std(m_{t+1}^{\$}))$	-	0.54

Table 10
Simulated moments: FX correlations and risk premia

For all moments, the left column presents the empirical moments, whereas the right column presents the simulated moments. All moments refer to nominal exchange rates. The first panel reports the cross-sectional minimum, mean and maximum of average realized FX correlation. The second panel reports the cross-sectional minimum, mean and maximum of average implied FX correlation, the cross-sectional correlation between average realized and average implied FX correlation, the cross-sectional average of the correlation between realized and implied FX correlation and the cross-sectional average of the correlation between changes in realized and implied FX correlation. The third panel reports the cross-sectional minimum, mean and maximum of average FX correlation risk premia, the cross-sectional correlation between average realized correlation and average correlation risk premia, the cross-sectional minimum, mean and maximum of the correlation between realized correlation and correlation risk premia, and the cross-sectional minimum, mean and maximum of the correlation between changes in realized correlation and changes in correlation risk premia. The last panel reports the average excess return of a monthly rebalanced portfolio that has long position on the 3 foreign currencies with the highest conditional loading on the correlation risk factor FXC and a short position on the 3 foreign currencies with the lowest conditional loading on FXC .

MOMENT	DATA	MODEL
$\text{Min}_{\text{cross}} (E(RC^{i,j}))$	0.05	0.11
$E_{\text{cross}} (E(RC^{i,j}))$	0.45	0.42
$\text{Max}_{\text{cross}} (E(RC^{i,j}))$	0.89	0.61
$\text{Min}_{\text{cross}} (E(IC^{i,j}))$	0.14	0.12
$E_{\text{cross}} (E(IC^{i,j}))$	0.48	0.43
$\text{Max}_{\text{cross}} (E(IC^{i,j}))$	0.88	0.61
$\text{corr}_{\text{cross}} (E(RC^{i,j}), E(IC^{i,j}))$	0.98	1.00
$E_{\text{cross}} (\text{corr}(RC^{i,j}, IC^{i,j}))$	0.79	1.00
$E_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(IC^{i,j})))$	0.28	1.00
$\text{Min}_{\text{cross}} (CRP^{i,j})$	-6.94%	-0.23%
$E_{\text{cross}} (CRP^{i,j})$	1.58%	0.17%
$\text{Max}_{\text{cross}} (CRP^{i,j})$	9.91%	0.77%
$\text{corr}_{\text{cross}} (E(RC^{i,j}), E(CRP^{i,j}))$	-0.55	-1.00
$\text{Min}_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	-0.76	-0.98
$E_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	-0.52	-0.87
$\text{Max}_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	0.16	-0.54
$\text{Min}_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.79	-0.98
$E_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.62	-0.85
$\text{Max}_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.35	-0.33
$E (rx_{t+1}^{HML^C})$	-6.42%	-1.08%

Table 11
Simulated moments: FX correlations and risk premia no comovement in local pricing factors

The first panel reports the minimum, average and maximum cross-sectional average of realized nominal exchange rate correlation in the model. The second panel reports cross-sectional average of mean implied correlation, the unconditional cross-sectional correlation between average realized and implied exchange rate correlation, the cross-sectional average of the correlation between realized and implied exchange rate correlation and the cross-sectional average of the correlation between changes in realized and implied exchange rate correlation. The third panel reports the cross-sectional average of the correlation risk premia, the cross-sectional correlation of average realized correlation and correlation risk premia, the cross-sectional minimum of the correlation between realized correlation and correlation risk premia and changes thereof, the cross-sectional average of the correlation of realized correlation and the correlation risk premia and the changes thereof and the cross-sectional maximum of the correlation between realized correlation and correlation risk premia and the changes thereof. Lastly, we report the average excess return on the correlation portfolio.

MOMENT	MODEL
$\text{Min}_{\text{cross}} (E(RC^{i,j}))$	0.27
$E_{\text{cross}} (E(RC^{i,j}))$	0.43
$\text{Max}_{\text{cross}} (E(RC^{i,j}))$	0.52
$E_{\text{cross}} (E(IC^{i,j}))$	0.43
$\text{corr}_{\text{cross}} (E(RC^{i,j}), E(IC^{i,j}))$	1.00
$E_{\text{cross}} (\text{corr}(RC^{i,j}, IC^{i,j}))$	1.00
$E_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(IC^{i,j})))$	1.00
$\text{Min}_{\text{cross}} (CRP^{i,j})$	0.08%
$E_{\text{cross}} (CRP^{i,j})$	0.12%
$\text{Max}_{\text{cross}} (CRP^{i,j})$	0.14%
$\text{corr}_{\text{cross}} (E(RC^{i,j}), E(CRP^{i,j}))$	0.98
$\text{Min}_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	-0.11
$E_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	-0.04
$\text{Max}_{\text{cross}} (\text{corr}(RC^{i,j}, CRP^{i,j}))$	0.02
$\text{Min}_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.19
$E_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.12
$\text{Max}_{\text{cross}} (\text{corr}(\Delta(RC^{i,j}), \Delta(CRP^{i,j})))$	-0.08
$E (rx_{t+1}^{\text{HML}^C})$	-0.27%

Appendix D Figures

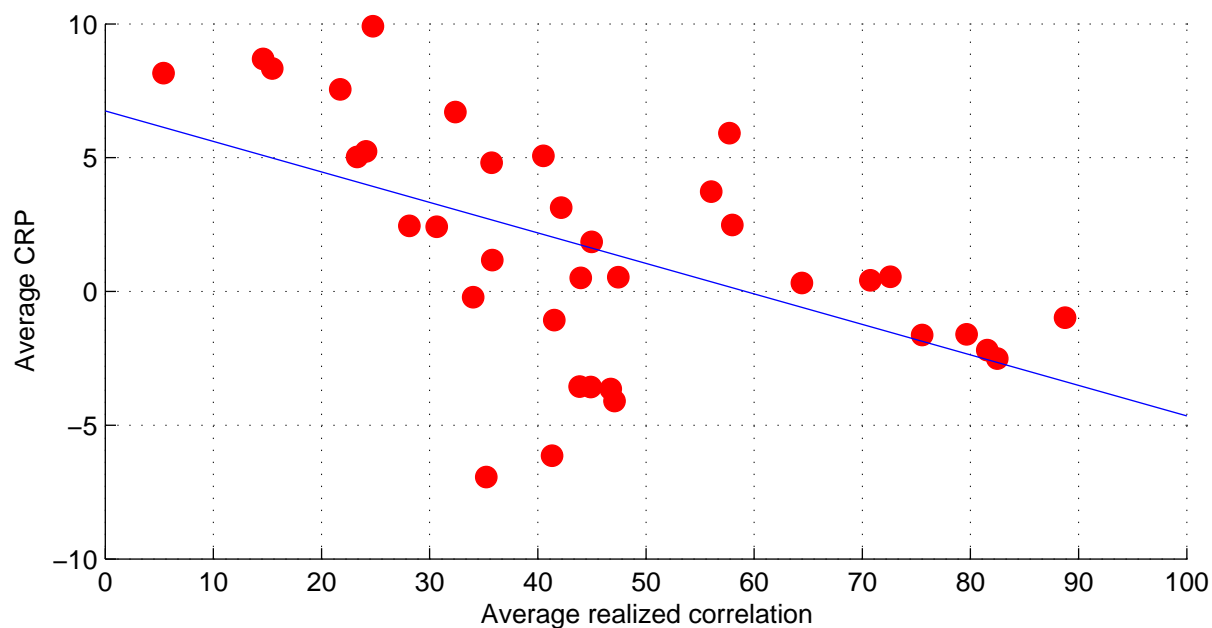


Figure 1. G10 realized correlation and correlation risk premia

This figure plots the average correlation risk premia for all 36 G10 exchange rate pairs against the average realized correlations. The unconditional correlation between the average correlation risk premia and average correlation is -55% . Correlation risk premia and correlations are expressed in percent. Data is monthly and runs from January 1996 (EUR since January 1999) to December 2013.

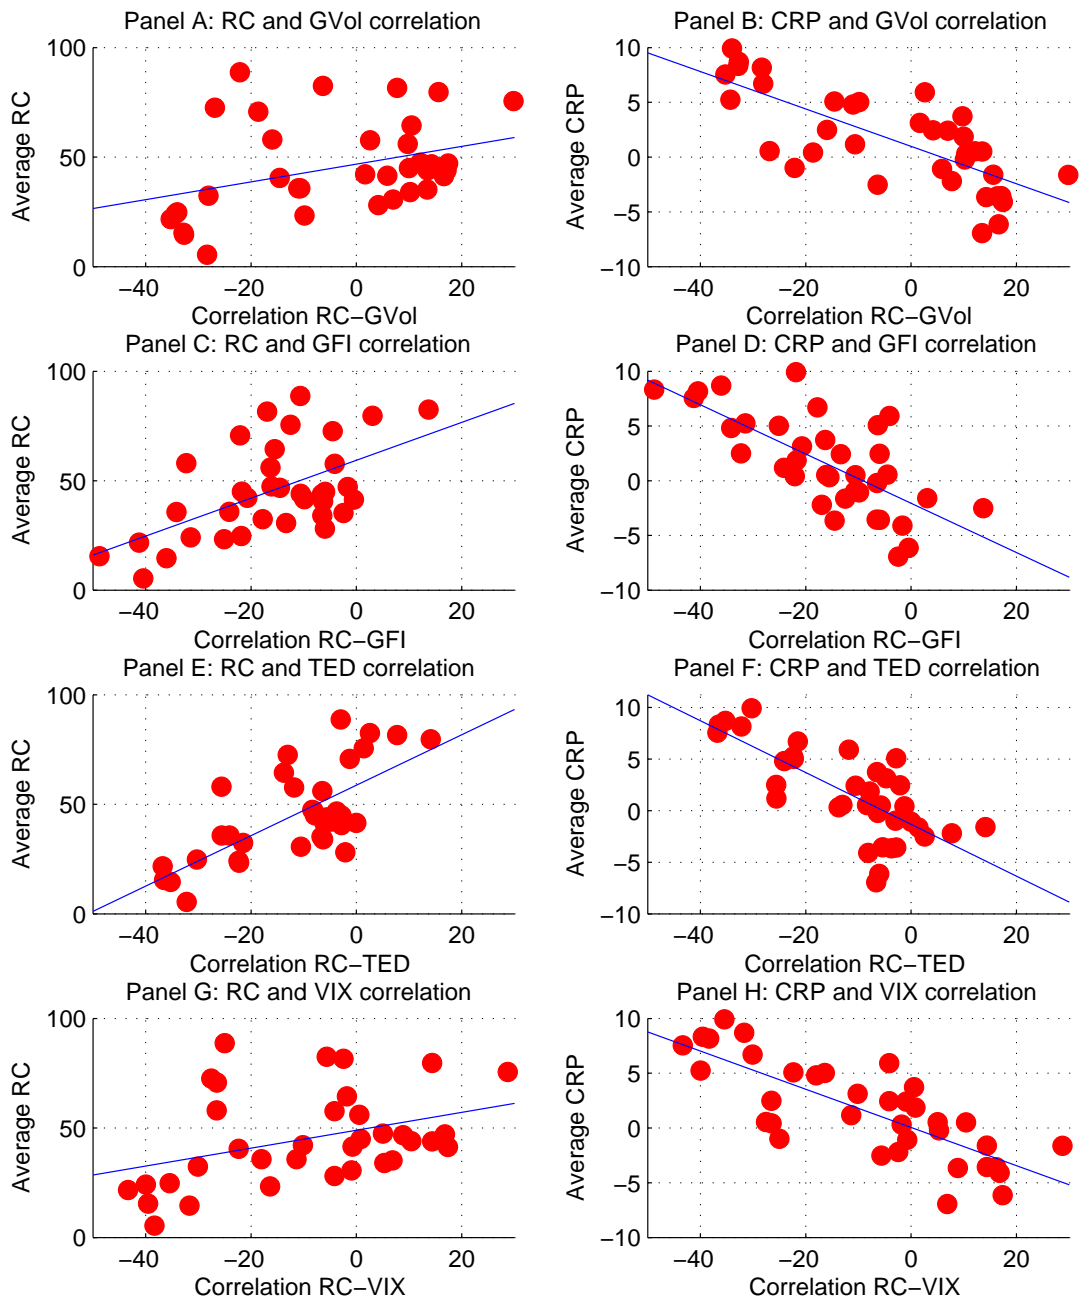


Figure 2. Realized correlations, correlation risk premia, and cyclicity

This figure illustrates the relationship between measures of cyclicity of FX correlations and average realized correlations and average correlation risk premia, respectively. Cyclicity is measured by the correlation between the realized correlation time series for a FX pair and a proxy for a global risk. The proxies considered are the global equity volatility measure from Lustig, Roussanov, and Verdelhan (2011) (*GVol*, Panels A and B), the global funding illiquidity measure (*GFI*, Panels C and D) from Malkhozov, Mueller, Vedolin, and Venter (2014), the TED spread (*TED*, Panels E and F), and the CBOE VIX (*VIX*, Panels G and H). Data is monthly and runs from January 1996 to December 2013.

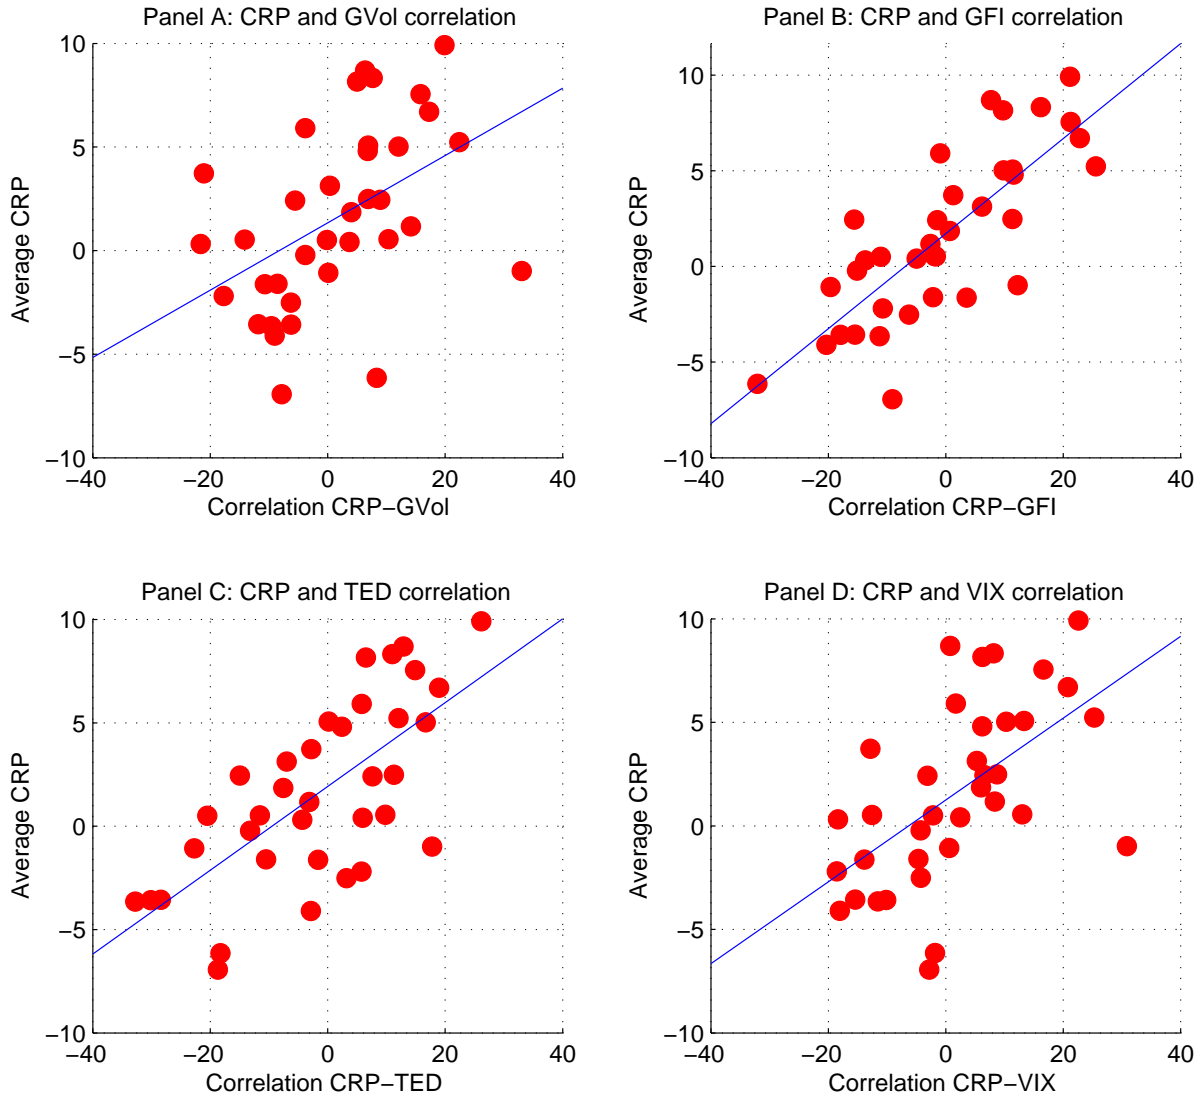


Figure 3. Correlation risk premia and cyclicity

This figure illustrates the relationship between measures of cyclicity of correlation risk premia and average correlation risk premia. Cyclicity is measured by the correlation between the correlation risk premia for a FX pair and a proxy for a global risk. The proxies considered are the global equity volatility measure from Lustig, Roussanov, and Verdelhan (2011) (*GVol*, Panel A), the global funding illiquidity measure (*GFI*, Panels B) from Malkhozov, Mueller, Vedolin, and Venter (2014), the TED spread (*TED*, Panels C), and the CBOE VIX (*VIX*, Panels D). The cross-sectional correlations between the average correlation and the proxies are 47%, 79%, 70% and 58%, respectively. Data is monthly and runs from January 1996 to December 2013.

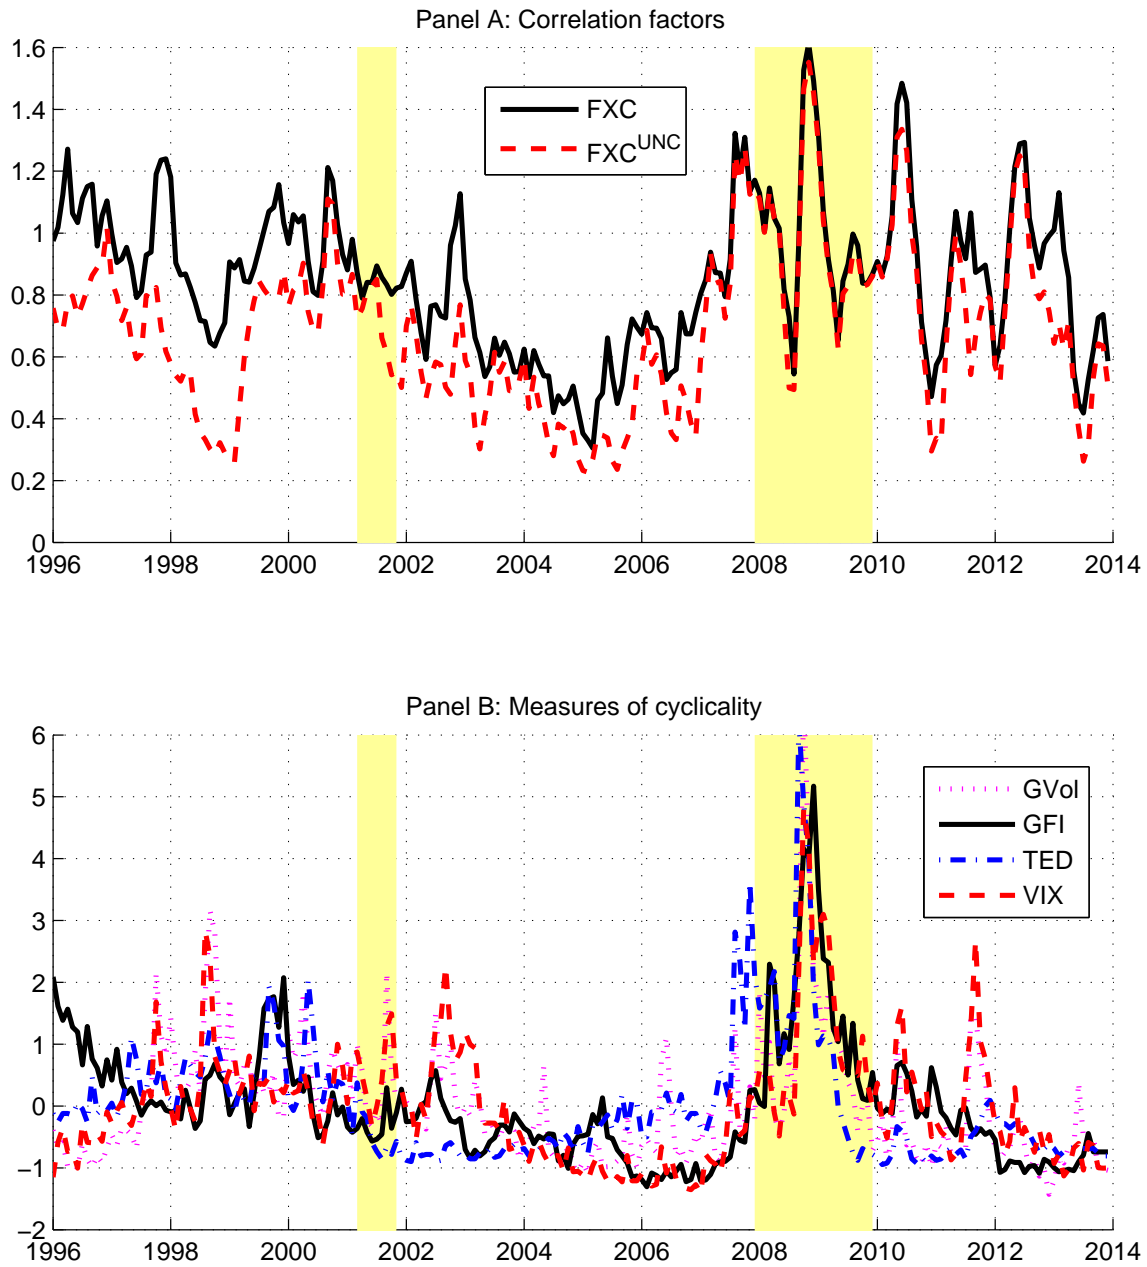


Figure 4. FX correlation risk factor and measures of cyclicity

Panel A plots the correlation risk factor FXC calculated as the difference between the average high and low correlation FX pairs (solid line). The two groups consist of the highest and lowest decile of realized correlations across all 36 G10 FX pairs. The deciles are rebalanced every month. The correlation risk factor is calculated for the period from January 1996 to December 2013. The alternative correlation risk factor FXC^{UNC} is calculated as the difference of correlations between the decile of high correlation pairs and the decile of low correlation pairs measured over the whole sample period. Panel B plots the global equity volatility measure used in Lustig, Roussanov, and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin, and Venter (2014) (GFI), the CBOE VIX (VIX), and the TED spread (TED). All series are standardized to have zero mean and a standard deviation of one. The shaded areas depict NBER recessions.

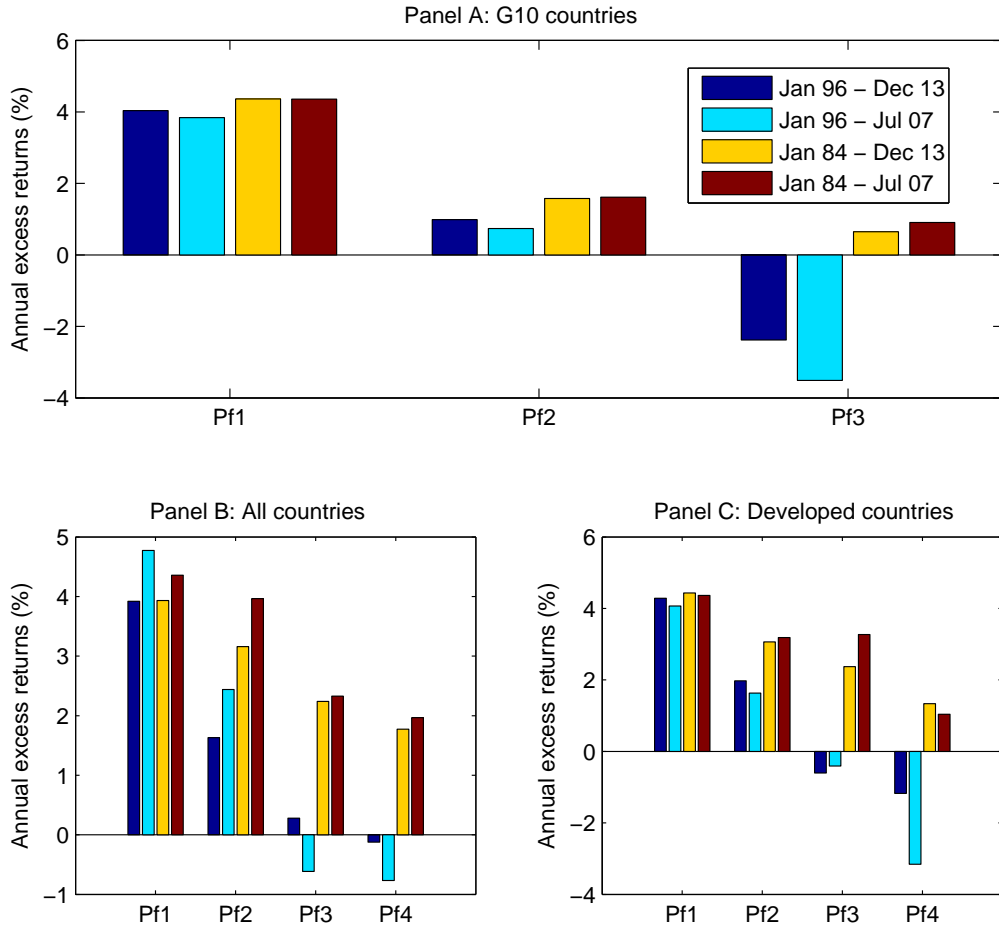


Figure 5. Portfolios sorted on exposure to correlation risk

The figure displays average portfolio excess returns for different subsamples. Currencies are sorted at time t into portfolios based on exposure to correlation risk at the end of period $t - 1$. The exposure is measured by regressing currency excess returns on innovations in the correlation risk factor over the preceding 36 months. Portfolio 1 (Pf1) contains the currencies with the lowest pre-sort correlation beta whereas Portfolio 3 or 4 (Pf3 or Pf4) contains the currencies with the highest pre-sort correlation beta. The average portfolio excess returns are calculated for the various sample periods starting either in January 1984 or January 1996 and ending in December 2013 or July 2007 (i.e., excluding the financial crisis). Panel A presents the results for the three G10 currency portfolios sorted based on the exposure to innovations in the correlation risk factor FXC . Panels B and C present the portfolio excess returns for portfolio sorts using an extended set of currencies (either developed currencies only or the full set as described in Section 1.1).

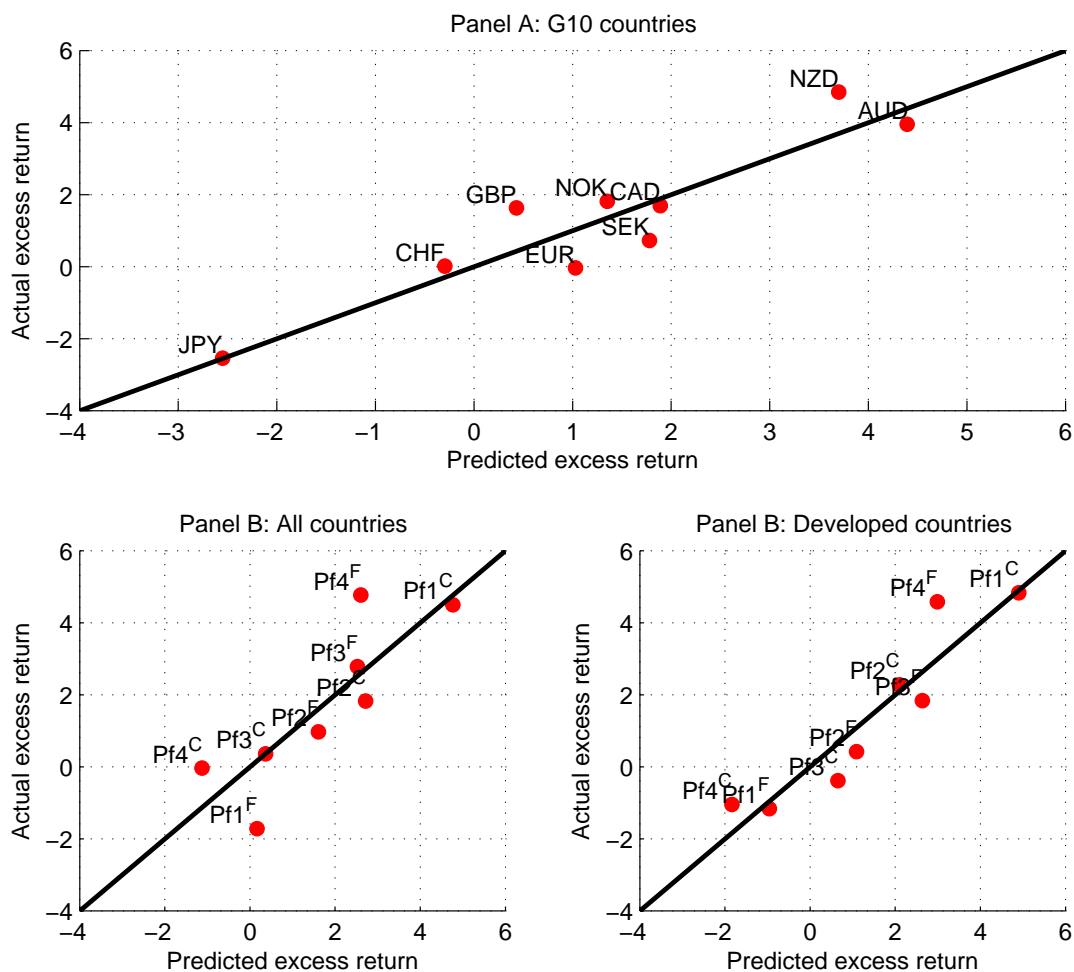


Figure 6. Model performance with various test assets

The figure plots the actual annualized mean excess returns in percent versus the predicted excess returns for various test assets using a linear pricing model that includes the dollar factor DOL and the HML^C correlation factor. Panel A displays the results for the nine G10 currencies and Panels B and C display the results for the four carry (Pf1^F to Pf4^F) and correlation (Pf1^C to Pf4^C) portfolios constructed using all or only developed currencies, respectively. Data is monthly and runs from January 1996 to December 2013.

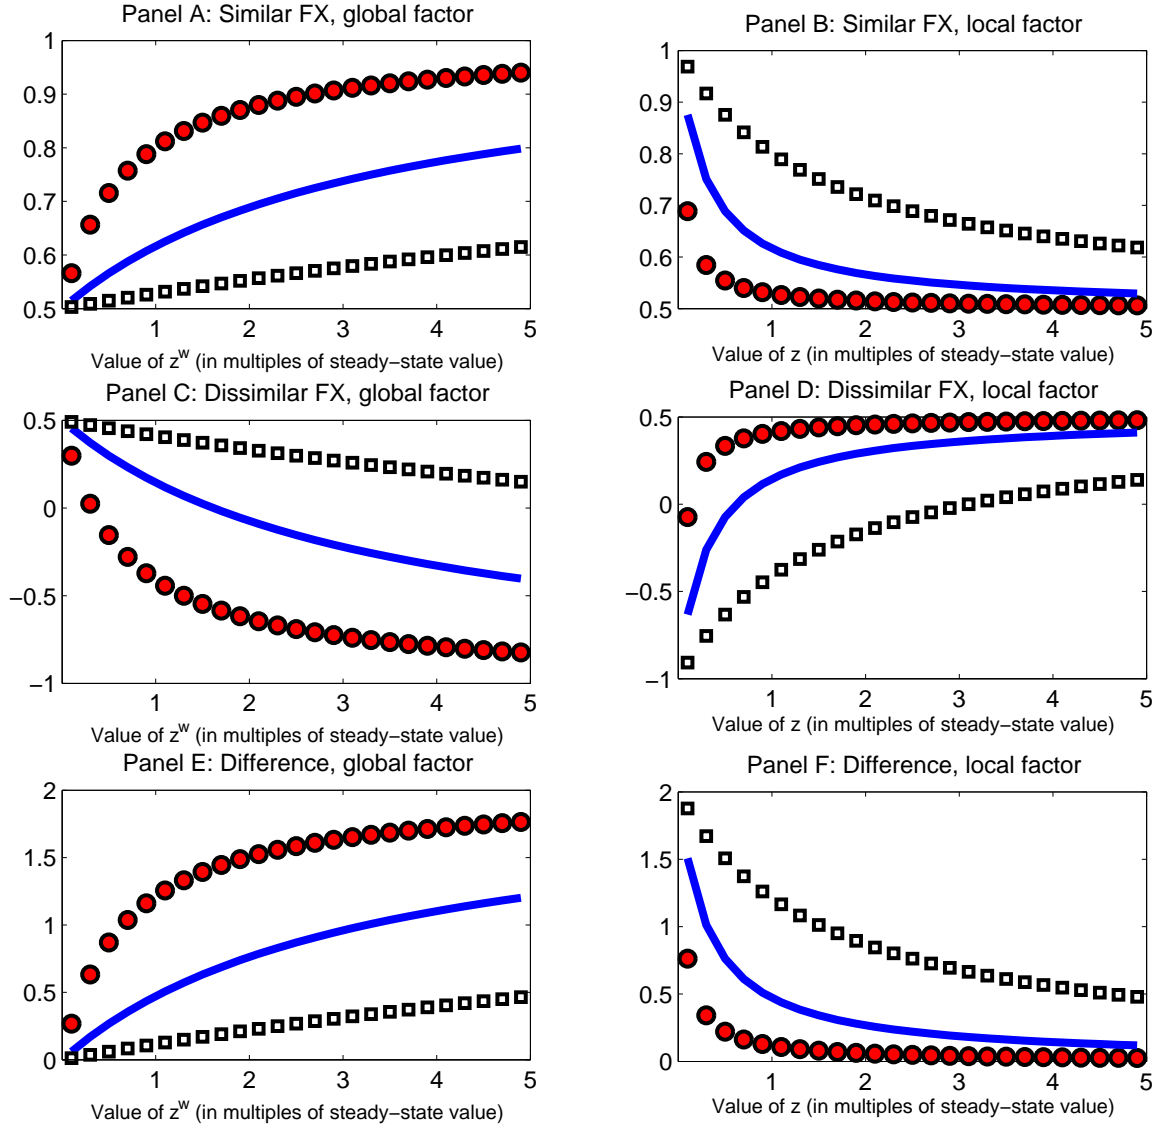


Figure 7. Model-implied FX correlations

The figure displays the properties of conditional real FX correlation in the benchmark model, which features identical local pricing factors. Panels A, C and E plot conditional FX correlation as a function of the global pricing factor z^w , holding the local pricing factor z constant: Panel A refers to the conditional FX correlation of the similar FX pair (1,2), Panel C refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel E refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the local pricing factor z being equal to 0.2, 1, and 5 times its steady-state value \bar{z} , respectively. Panels B, D and F plot conditional FX correlation as a function of the local pricing factor z , holding the global pricing factor z^w constant: Panel B refers to the conditional FX correlation of the similar FX pair (1,2), Panel D refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel F refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the global pricing factor z^w being equal to 0.2, 1, and 5 times its steady-state value \bar{z}^w , respectively. To plot the figures, we set the relevant model parameters equal to their calibrated values in Table 8.




Figure 8. Model-implied FX correlations: no comovement in local pricing factors

The figure displays the properties of conditional real FX correlation in the full model. Panels A, C and E plot conditional FX correlation as a function of the global pricing factor z^w , holding all the local pricing factors constant: Panel A refers to the conditional FX correlation of the similar FX pair (1,2), Panel C refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel E refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the domestic local pricing factor z^0 being equal to 0.2, 1, and 5 times its steady-state value \bar{z} , respectively, and all the foreign local pricing factors being equal to their common steady-state value \bar{z} . Panels B, D and F plot conditional FX correlation as a function of the local pricing factor z^w , holding the global pricing factor z^w constant: Panel B refers to the conditional FX correlation of the similar FX pair (1,2), Panel D refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel F refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the global pricing factor z^w being equal to 0.2, 1, and 5 times its steady-state value \bar{z}^w , respectively, and all the foreign local pricing factors being equal to their steady-state value \bar{z} . To plot the figures, we set the relevant model parameters equal to their calibrated values in Table 8.



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



The London School of Economics
and Political Science
Houghton Street
London WC2A 2AE
United Kingdom

Tel: +44 (0)20 7405 7686
systemicrisk.ac.uk
src@lse.ac.uk