Philippe Mueller, Andreas Stathopoulos, Andrea Vedolin

International correlation risk

Article (Accepted version) (Refereed)

Original citation:
Mueller, Philippe and Stathopoulos, Andreas and Vedolin, Andrea (2017) *International correlation risk.* Journal of Financial Economics. ISSN 0304-405X

DOI: 10.1016/j.jfineco.2016.09.012

International correlation risk

Philippe Mueller, Andreas Stathopoulos, Andrea Vedolin

PII: S0304-405X(17)30215-5
DOI: 10.1016/j.jfineco.2016.09.012
Reference: FINEC 2807

To appear in: *Journal of Financial Economics*

Received date: 26 January 2016
Revised date: 23 August 2016
Accepted date: 21 September 2016

Please cite this article as: Philippe Mueller, Andreas Stathopoulos, Andrea Vedolin, International correlation risk, *Journal of Financial Economics* (2017), doi: 10.1016/j.jfineco.2016.09.012

International correlation risk☆

Philippe Mueller[a], Andreas Stathopoulos[b,*], Andrea Vedolin[a]

[a]London School of Economics, Department of Finance, Houghton Street, WC2A 2AE London, UK
[b]University of Washington, Foster School of Business, 4277 E Stevens Way NE, Seattle, WA 98195, USA

Abstract

We show that the cross-sectional dispersion of conditional foreign exchange (FX) correlation is countercyclical and that currencies that perform badly (well) during periods of high dispersion yield high (low) average excess returns. We also find a negative cross-sectional association between average FX correlations and average option-implied FX correlation risk premiums. Our findings show that while investors in spot currency markets require a positive risk premium for exposure to high dispersion states, FX option prices are consistent with investors being compensated for the risk of low dispersion states. To address our empirical findings, we propose a no-arbitrage model that features unspanned FX correlation risk.

JEL classification: F31, G15

Keywords: Correlation risk, Exchange rates, International finance

1. Introduction

Existing literature has shown that stock return correlations are counter cyclical and correlation risk is priced, arguably due to the reduction of diversification benefits that occurs when stock return correlations increase. However, the literature has largely ignored the foreign exchange (FX) market. In this paper, we explore the properties of FX correlations using both spot and options market data and we propose a reduced-form no-arbitrage model that is consistent with our empirical findings.

We begin by exploring the empirical properties of conditional FX correlations. We consider exchange rates against the US dollar (USD) and find substantial cross-sectional heterogeneity in the average conditional correlation of FX pairs. Furthermore, using several business cycle proxies, we find that the cross-sectional dispersion of FX correlations is counter cyclical, as FX pairs with high (low) average correlation become more (less) correlated in adverse economic times. We exploit the cyclical properties of conditional FX correlation by defining an FX correlation dispersion measure, FXC, and sort currencies into portfolios based on the beta of their returns with respect to innovations in FXC, denoted by ΔFXC. We find that currencies with low ΔFXC betas have high average excess returns and that currencies with high ΔFXC betas yield low excess returns, suggesting that FX correlation risk has a negative price in

☆We would like to thank Dante Amengual, Andrew Ang, Svetlana Bryzgalova, Joe Chen, Mike Chernov, Ram Chivukula, Max Croce, Robert Dittmar, Dobrislav Dobrev, Anh Le, Angelo Ranaldo, Paul Schneider, Ivan Shaliastovich, Adrien Verdelhan, Hao Zhou, and seminar and conference participants in the Bank of England, Federal Reserve Board, London School of Economics, LUISS Guido Carli, Ohio State University, University of Bern, University of Essex, University of Lund, University of Pennsylvania, University of Piraeus, Southern Methodist University, Stockholm School of Economics, the 2011 Chicago Booth Finance Symposium, the Sixth End-of-Year Meeting of the Swiss Economists Abroad, the 2012 Duke University and University of North Carolina Asset Pricing Conference, the 2012 University of California at Los Angeles–University of Southern California Finance Day, the Bank of Canada–Banco de España Workshop on International Financial Markets, the Financial Econometrics Conference at the University of Toulouse, the 2012 Society for Economic Dynamics Meeting, the Eighth Asset Pricing Retreat at Cass Business School, the 2012 Centre for Economic Policy Research Summer Meeting, the 2012 European Finance Association Meeting, and the 2016 American Economic Association Meeting for thoughtful comments. Philippe Mueller and Andrea Vedolin acknowledge financial support from STICERD, the Systemic Risk Centre at London School of Economics and the Economic and Social Research Council (Grant ES/K002309/1).

*Corresponding author.

Email addresses: p.mueller@lse.ac.uk (Philippe Mueller), astath@uw.edu (Andreas Stathopoulos), a.vedolin@lse.ac.uk (Andrea Vedolin)

spot FX markets. In our benchmark sample of currencies, HML^C, a currency portfolio with a short position in the high ΔFXC beta currencies and a long position in the low ΔFXC beta currencies, generates a highly significant average annual excess return of 6.42% with a Sharpe ratio of 0.82.

We continue our empirical investigation by using currency option prices to extract conditional FX correlation dynamics under the risk-neutral measure. We calculate FX correlation risk premiums (CRPs), defined as the difference between conditional FX correlations under the risk-neutral measure and the physical measure, and we find a strongly negative cross-sectional association between average FX correlations and average FX correlation risk premiums, with FX pairs characterized by low (high) average correlations tending to exhibit positive (negative) correlation risk premiums. Thus, the cross-sectional dispersion of FX correlations is on average lower under the risk-neutral measure than under the physical measure. We also find a very strong negative time series association between FX correlations and FX correlation risk premiums for almost all FX pairs. As regards cyclicality, FX pairs with high average correlation risk premiums have countercyclical correlation risk premiums and pairs with low correlation risk premiums have procyclical premiums. Thus, bad states amplify the magnitude of FX correlation risk premiums, increasing their cross-sectional dispersion.

We rationalize our empirical findings with a no-arbitrage model of exchange rates. The main tension we address is between the physical and the risk-neutral measure FX correlation dynamics. Under the physical measure, the negative association between ΔFXC betas and currency returns suggests that US investors require a positive risk premium for being exposed to states in which the cross section of FX correlations widens. However, FX options are priced in a way that suggests that US investors worry about states in which the cross section of FX correlations tightens, as the risk-neutral measure FX correlation dispersion is on average lower than its physical measure counterpart. To address this apparent contradiction, we propose a model in which FX correlation risk is not spanned by exchange rates, as some shocks that affect the pricing kernel of US investors also affect conditional FX correlations, but do not impact exchange rate levels.

In the model, each country's stochastic discount factor (SDF) is exposed to two global shocks, as well as a single country-specific shock. A key assumption is that countries have heterogeneous loadings on the first global shock and identical loadings on the second global shock. As a result, the absence of arbitrage in international financial markets suggests that exchange rates are exposed only to the first global shock, whereas the second global shock cancels out and does not affect exchange rates at all. The steady state cross-sectional distribution of conditional FX correlations is determined by the cross section of exposures to the first global shock, so the USD exchange rates of foreign countries with similar exposure to the first global shock (called similar FX pairs) are more correlated on average than FX pairs of countries with dissimilar global risk exposure (called dissimilar FX pairs). Crucially, the cross section of conditional FX correlations exhibits time variation due to the fact that conditional FX correlations are determined by the relative importance of country-specific risk and global risk, which varies over time. When the relative magnitude of country-specific SDF shocks increases, the countries' heterogeneous exposure to the first global shock becomes less important quantitatively, and the cross section of conditional FX correlations tightens, with high correlation FX pairs becoming less correlated and low correlation FX pairs more correlated. Conversely, a relative increase in the magnitude of global risk increases the correlation of similar FX pairs and decreases the correlation of dissimilar FX pairs, widening the cross section of conditional FX correlations.

In turn, the relative magnitude of country-specific and global risk is determined by the relative magnitude of the local pricing factor, which prices country-specific risk and is exposed to the second global shock, and the global pricing factor, which prices global risk and is exposed to the first global shock. When the second global shock has an adverse realization, the local pricing factor increases, tightening the cross section of conditional FX correlations. Conversely, when the second global shock has a positive realization, the cross section of conditional FX correlation becomes more dispersed. The reverse occurs for realizations of the first global shock, as its adverse (positive) realizations increase (decrease) the global pricing factor, widening (tightening) the cross section of FX correlations. Thus, the cross section of conditional FX correlations is driven by both global shocks. In the model, both shocks are priced, but not symmetrically, as US investors price the second shock more severely than the first, so they attach a high price to states characterized by large relative values of the local pricing factor. Given that those are the states in which the cross-sectional dispersion of FX correlation is tight, our model is able to match the cross sectional properties of average correlation risk premiums implied by FX option prices.

As regards spot FX markets, exchange rate levels are unaffected by the second global shock, so exchange rate risk does not span FX correlation risk. This lack of spanning allows our model to generate a negative relation between ΔFXC betas and currency returns. This is because investing in exchange rates draws compensation solely for exposure to the first global shock and, due to the fact that negative realizations of that shock lead to a widening of the cross section of FX correlations, investors require high returns for holding negative ΔFXC beta currencies, which depreciate when the cross section of conditional FX correlations becomes more dispersed.

In sum, conditional FX correlation, which can be indirectly traded using currency options, is exposed to two global shocks. US investors price the second global shock more severely than the first one, so FX correlation risk premiums reflect the desire of currency option holders to primarily avoid states with negative realizations of the second shock. Those states are characterized by a tightening of the cross-sectional dispersion of FX correlation, and currency option prices reveal that feature. Investing in foreign currency is different, as it exposes investors only to the first global shock. Therefore, currency risk premiums reflect solely FX investors' desire to avoid the corresponding bad states, characterized by a widening of the cross-sectional dispersion of FX correlation, and compensate investors for exposure to those states. Thus, the lack of spanning of FX correlation risk by exchange rates and currency returns and, in particular, the lack of exposure of exchange rates to the second global shock allows our model to jointly address the empirical properties of FX correlations, currency risk premiums, and FX correlation risk premiums.

A simulated version of our model generates realized FX correlations, implied FX correlations, and FX correlation risk premiums that match the cross-sectional and time series properties of their empirical counterparts, all the while fitting the standard exchange rate, interest rate and inflation moments.

This paper is part of the literature addressing the salient empirical properties of FX markets. Our model builds on the work of Lustig, Roussanov and Verdelhan (2011, 2014) and Verdelhan (2015). Their models feature global SDF shocks, common across countries, and local SDF shocks, independent across countries, and assume that the price of country-specific shocks is uncorrelated across countries, as local pricing factors are perfectly negatively correlated with the corresponding country-specific shocks. We show that allowing for cross-country comovement of the local pricing factors is crucial for explaining the joint behavior of FX correlations under the physical and the risk-neutral measure.

Our model assumes ex ante heterogeneity across countries regarding their exposure to global shocks. Recent international finance models that address the cross section of currency risk premiums by assuming ex ante heterogeneity across countries include those of Hassan (2013), Tran (2013), Backus, Gavazzoni, Telmer and Zin (2013), Colacito and Croce (2013), Colacito, Croce, Gavazzoni and Ready (2015), and Ready, Roussanov and Ward (2016). In all those models, high interest rate currencies are risky because they depreciate in bad global states, as high interest rate countries are those with low exposure to global risk: small countries, countries with smooth non-traded output, countries with very procyclical monetary policy, commodity producers, or countries with low exposure to global long-run endowment shocks, depending on the model.

Finally, our paper is related to the literature on currency options. Most of that literature focuses on crash risk, especially in the context of the FX carry trade. Examples include Jurek (2014), Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2015), and Chernov, Graveline and Zviadadze (2016). Our aim is different, as we use option prices to study the properties of FX correlation risk premiums.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 reports our empirical findings regarding the cross-sectional and time series properties of FX correlations, as well as the pricing of correlation risk in currency markets. Our empirical findings concerning FX correlation risk premiums are presented in Section 4. Section 5 introduces our no-arbitrage model, and Section 6 concludes. The Appendix contains details on the construction of the realized and implied FX correlation measures, results on the price of FX correlation risk, and model details, including those on model calibration and simulation. Additional results and robustness checks are in an Online Appendix.

2. Data

Our benchmark sample period is dictated by the availability of the currency options data, so it starts in January 1996 and ends in December 2013.

To calculate physical measure FX moments, we use daily spot exchange rates from WM/Reuters obtained through Datastream. From the same source, we also collect one-month forward rates to calculate forward discounts.

Following the extant literature (see, e.g., Fama, 1984), we work with log spot and log one-month forward exchange rates, denoted $s_t^i = \ln(S_t^i)$ and $f_t^i = \ln(F_t^i)$, respectively. We use the US dollar as the base currency, and both spot and forward rates are expressed in units of foreign currency i per USD. WM/Reuters forward rates are available from 1997 onward. For 1996, either we use forward rates from alternative sources or we construct implied forward rates using the interest rate differential between the US and the foreign country using interest rate data from Datastream, exploiting the fact that covered interest rate parity holds during normal conditions. We verify that our results are robust to using only the WM/Reuters data. Monthly log excess returns from holding the foreign currency i are computed as $rx_{t+1}^i = f_t^i - s_{t+1}^i$. Our benchmark sample contains the nine G10 foreign currencies [Australian dollar (AUD), Canadian dollar (CAD), Swiss Franc (CHF), Euro (EUR), Pound Sterling (GBP), Japanese yen (JPY), Norwegian krona (NOK), New Zealand dollar (NZD), and Swedish krona (SEK)] from January 1996 to December 2013. For robustness checks, we consider the longer January 1984 to December 2013 sample period. Before the introduction of the EUR in January 1999, we use the German mark (DEM) in its place.

Table 1 presents the properties of the G10 currency excess returns. In line with the literature on the FX carry trade, currencies with high nominal interest rates, such as the NZD and the AUD, yield high average dollar excess returns, while low interest rate currencies, such as the JPY and the CHF, have low dollar excess returns on average.

[Insert Table 1 near here.]

For robustness, we extend the cross section of currencies and consider two additional currency sets: developed and emerging market currencies. The developed country sample, apart from the G10 currencies, contains the currencies of Austria, Belgium, Denmark, Finland, France, Greece, Italy, Ireland, Netherlands, Portugal, and Spain. To form our full sample, we start with the same set of currencies used in Lustig, Roussanov and Verdelhan (2011), but exclude some currencies, such as the Hong Kong dollar, as they are pegged to the USD. We also exclude the Danish krone after the introduction of the EUR. Therefore, the full sample includes all the developed country currencies, along with the currencies of the Czech Republic, Hungary, India, Indonesia, Kuwait, Malaysia, Mexico, Philippines, Poland, Singapore, South Africa, South Korea, Taiwan, and Thailand.

We use daily over-the-counter G10 currency options data from J. P. Morgan. In addition to the nine currency pairs versus the US dollar, we have options data for all 36 cross rates. The options used in this study are plain-vanilla European calls and puts, with five option series per currency pair. We focus on the one-month maturity and a total of five different strikes: at-the-money (ATM), 10-delta and 25-delta calls, and 10-delta and 25-delta puts.

3. Exchange rate correlations

In this section, we first show that the cross-sectional dispersion of conditional FX correlation is countercyclical. We then construct an FX correlation dispersion measure, FXC, and sort currencies into portfolios based on their return exposure to FXC innovations, denoted by ΔFXC. We find a negative association between ΔFXC betas and currency excess returns, suggesting that currency exposure to FX correlation risk is compensated with a positive risk premium.

3.1. Properties of exchange rate correlations

We proxy the conditional one-month correlation of each FX pair under the physical measure at time t with its realized correlation over a rolling three-month window of past daily observations. Appendix A provides the details. In the remainder of the paper, we often refer to physical measure conditional FX correlation as realized FX correlation, to distinguish it from the option-implied risk-neutral measure FX correlation (implied FX correlation). For robustness, we also proxy the conditional one-month correlation of each FX pair at time t with its realized correlation over a rolling one-month window of past daily observations, as well as with its realized correlation during the one-month-ahead period, i.e., from t to $t + 1$. Our empirical results are robust to those alternative specifications. Our findings for correlation risk premiums using the alternative realized correlation proxies are reported in the Online Appendix.

The first two columns of Table 2 report the time series mean and standard deviation of the conditional FX correlation of each of the 36 G10 FX pairs. The mean conditional correlation is positive for all 36 FX pairs, indicating that all pairs of USD exchange rates exhibit positive comovement on average. The cross-sectional average of the conditional correlation means is 0.45, but there is substantial cross-sectional heterogeneity, as the means range from almost zero (CAD/JPY with 0.05, indicating that fluctuations in the relative price of the CAD and the JPY against the USD are almost disconnected) to almost one (CHF/EUR with 0.89).[1] Furthermore, conditional FX correlations exhibit considerable variability across time. The cross-sectional average of the standard deviation of conditional FX correlations is 0.23, ranging from 0.09 (EUR/NOK pair) to 0.34 (AUD/JPY pair), suggesting nontrivial swings in the degree of exchange rate comovement across time for all FX pairs.

[Insert Table 2 near here.]

Given the time variation in conditional FX correlations, it is worth exploring whether that time variation is cyclical and, if so, whether there is any cross-sectional heterogeneity in its properties. To that end, we consider the comovement of conditional FX correlations with market variables that are known to exhibit countercyclical behavior. The market variables we consider are a global equity volatility measure ($GVol$), a global funding illiquidity measure (GFI), the TED spread (TED), and the Chicago Board Options Exchange (CBOE) Volatility Index (VIX). $GVol$ is constructed as in Lustig, Roussanov and Verdelhan (2011). GFI is constructed following the methodology of Hu, Pan and Wang (2013) but calculated using an international sample of government bond securities as in Malkhozov, Mueller, Vedolin and Venter (2016). TED is the spread between the three-month USD London Interbank Offered Rate (LIBOR) and the three-month Treasury bill rate and is available in the Federal Reserve Economic Data (FRED) database. VIX is calculated from the prices of options on the Standard & Poor's 500 stock index and is available from the CBOE. TED and VIX are US-specific measures, but they are often used as global market indicators. $GVol$ and GFI are calculated using international data in local currencies. For each FX pair and each market measure, we define the cyclicality measure to be the unconditional correlation of the market variable with the conditional correlation of the FX pair. Thus, we calculate four FX correlation cyclicality measures for each exchange rate pair, each corresponding to a market variable. We present the cyclicality measures for the 36 G10 FX pairs in the first four columns of Table 3.

[Insert Table 3 near here.]

Our results indicate substantial cross-sectional heterogeneity regarding the cyclicality properties of conditional FX correlations. To further explore the properties of that cross-sectional heterogeneity, we plot each cyclicality measure of the 36 FX pairs against their average conditional correlation. Panels A to D in Fig. 1 present the plots for the four cyclicality measures. We find that that FX pairs with high average correlation tend to exhibit countercyclical correlations and that FX pairs with low average correlation are characterized by procyclical FX correlations. Each panel also presents the line of best fit from the corresponding cross-sectional regression. We show the details of the four cross-sectional regressions in Panel A of Table 4. For each regression, the table reports the point estimate of the slope coefficient, its asymptotic t-statistic, the 95% bootstrapped confidence interval (2.5 and 97.5 bootstrap percentiles),

[1] Beginning in September 2011, the Swiss National Bank imposed a cap on the relative value of the CHF by establishing a floor of 1.2 CHF per EUR. The average correlation between the CHF/USD exchange rate and the EUR/USD exchange rate in the period before the cap (0.887) is almost identical to their average correlation during the cap period (0.895). Given that the cap does not seem to have changed the behavior of the CHF, we choose to retain the CHF in our sample after September 2011. We have verified that removing the CHF during the cap period does not materially affect our results.

and the regression R^2. The asymptotic t-statistic is calculated using White (1980) standard errors that adjust for cross-sectional heteroskedasticity, and the bootstrapped confidence interval accounts for potential small sample effects. All four slope coefficients are positive and statistically significant at the 5% level using either the asymptotic or the bootstrapped distribution, suggesting a positive cross-sectional association between average conditional FX correlation and FX correlation cyclicality.

[Insert Figure 1 and Table 4 near here.]

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. To further explore the time series properties of the cross-sectional dispersion in conditional FX correlation, we construct a conditional FX correlation dispersion measure, called FXC. To do so, at each period t we sort all FX pairs in deciles on their conditional correlation, calculate the average conditional correlation for the top and bottom deciles (which consist of four FX pairs each), and take the difference between the top and the bottom decile averages to be our dispersion measure at t, FXC_t. Due to the time variation in conditional FX correlations, there is turnover in both the top and bottom deciles, so the composition of the portfolios changes over time. To eliminate composition effects, we compute an alternative dispersion measure (FXC^{UNC}) by considering top and bottom deciles of FX pairs formed using average conditional correlations, and use this alternative measure for robustness exercises.

We plot the time series of the level of the two FX correlation dispersion measures in Panel A of Fig. 2. The correlation between FXC and FXC^{UNC} is 0.86, indicating that the two measures are very similar. During the recent financial crisis, the two measures are almost perfectly correlated, as there is little turnover in the extreme deciles of FX conditional correlation. To evaluate the cyclicality properties of the FX correlation dispersion measures, we explore their association with the market variables we use to measure the cyclicality of FX correlations. For reference, in Panel B of Fig. 2, we plot the (standardized) market variables. Panel A of Table 5 reports the unconditional correlations between our two FX correlation dispersion measures and the market variables, in the January 1996 to December 2013 sample period, along with their bootstrap standard errors. Both dispersion measures have a positive correlation with all four market variables, with bootstrap confidence intervals (which account for non-normality in small samples and are not reported in Table 5) indicating that the correlation is statistically significant at the 1% level in all eight cases. Panel B repeats the same exercise for the longer January 1984 to December 2013 period. Again, all eight correlations of interest are positive and significant at the 1% level.

[Insert Figure 2 and Table 5 near here.]

In the Online Appendix we consider alternative construction methods for FXC and we show that our portfolio results are robust to those alternative specifications.

3.2. Correlation risk and the cross section of currency returns

We can now explore how exposure to FX correlation risk relates to currency returns. To do so, we sort currencies into portfolios based on the exposure (beta) of currency excess returns to innovations in our dispersion measure FXC. Innovations between t and $t+1$ are denoted by ΔFXC_{t+1} and are defined as the average of changes (first differences) in conditional FX correlation for the FX pairs that belong to the top decile in period t minus the corresponding average for the bottom decile. Our currency portfolios are rebalanced monthly. To rebalance, we calculate, for each month t, rolling ΔFXC return betas using the last 36 monthly observations. Hence, currency portfolios are formed using information available at the time of rebalancing.

We sort the nine G10 currencies into three portfolios. The first portfolio (Pf1^{C}) contains the currencies with the lowest ΔFXC betas, and the last portfolio (Pf3^{C}) contains the highest ΔFXC beta currencies. Of particular interest is the HML^C portfolio, which takes a long position in Pf3^{C} and a short position in Pf1^{C}. Panel A of Table 6 reports the summary statistics for the three ΔFXC beta-sorted currency portfolios and the HML^C portfolio. Average portfolio returns are monotonically decreasing in the ΔFXC beta, suggesting that ΔFXC is a priced currency risk factor. As a result, the average return to HML^C is negative and highly statistically significant. Shorting the HML^C portfolio yields an annualized average excess return of 6.42%, with a t-statistic of 3.47 and an associated Sharpe ratio of 0.82.

[Insert Table 6 near here.]

Our finding of a strongly negative return for HML^C is robust to different sample periods. We consider the following periods, two of which do not include the recent financial crisis: January 1996 to July 2007, January 1984 to December 2013, and January 1984 to July 2007. Our findings are reported in Panels B to D of Table 6. Consistent with our results for the benchmark period, we find an inverse relation between exposure to the FX correlation factor ΔFXC and average currency portfolio excess returns in each of the three periods. Excluding the financial crisis increases the average excess return of shorting the HML^C portfolio to 7.35%, with an associated Sharpe ratio of 1.10 (Panel B). Return differences across portfolios somewhat attenuate when the sample period is extended back to January 1984 (Panels C and D), but shorting the HML^C portfolio still yields highly significant annualized average excess returns (3.72% and 3.45%, respectively). Overall, our results are very robust to different sample periods and do not appear to be driven by the recent financial crisis.

For further robustness, we explore extended cross sections of currencies. We consider a sample that includes other developed country currencies and a sample that includes the entirety of the developed sample and also some emerging currencies. The composition of those extended samples is discussed in detail in Section 2. For each of the two extended samples, we construct four ΔFXC beta-sorted portfolios. Fig. 3 presents the average excess returns of ΔFXC beta-sorted currency portfolios for each of three sets of currencies (G10, all countries, and developed countries) and each of the four periods. We find a consistently negative association between average portfolio excess returns and exposure to correlation risk, with negative average HML^C returns across the board. Furthermore, average HML^C returns are significant at the 5% level for all currency and period samples, with the sole exception of the samples starting in 1984 for the full set of currencies. For the benchmark period from January 1996 to December 2013, the average annualized return of shorting HML^C in the developed country sample is 5.46% (with a t-statistic of 2.42) and the associated Sharpe ratio is 0.57. For the full cross section of currencies, shorting HML^C yields 4.04% on average (with a t-statistic of 1.97) and a Sharpe ratio of 0.46.

[Insert Figure 3 near here.]

Finally, given the significant excess returns to the HML^C portfolio, we attempt to determine the market price of FX correlation risk. We follow the extant literature and consider a linear pricing model with two traded factors. The first factor is the dollar factor DOL, defined as the simple average of all available FX excess returns and shown by Lustig, Roussanov and Verdelhan (2011) to act as a level factor for currency returns, and the second factor is HML^C, the return difference between the high and low ΔFXC beta portfolios for the sample of G10 currencies. Our estimates for the market price of HML^C range from −51 to −67 basis points per month, depending on the set of test assets, so HML^C acts as a slope factor for pricing currency risk. The results are presented in detail in Appendix B.

4. Exchange rate correlation risk premiums

In this section, we present the cross-sectional and time series properties of FX correlation risk premiums and explore the relation between FX correlation risk premiums and FX correlations.

4.1. The cross-sectional properties of correlation risk premiums

Consistent with the literature on variance and correlation risk premiums in other asset classes, we define FX correlation risk premiums as the difference between expected conditional FX correlations under the risk-neutral ($\mathbb{Q}$) and the physical ($\mathbb{P}$) measure:

$$\mathrm{CRP}_{t,T}^{i,j} \equiv \mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \rho_u^{i,j} du\right) - \mathrm{E}_t^{\mathbb{P}}\left(\int_t^T \rho_u^{i,j} du\right). \tag{1}$$

We consider only one-month premiums, i.e., $T = t + 1$, as the maturity of the FX options we use to derive risk-neutral measure moments is one month. Variance risk premiums are defined analogously as the difference in expected conditional FX variance between the risk-neutral and the physical measure. The Online Appendix contains a discussion of the properties of variance risk premiums, as well as of physical measure and risk-neutral measure FX variance.

To calculate the risk-neutral (implied) conditional FX correlation, we follow the literature on model-free measures of implied volatility and covariance using daily FX option prices. The details of the calculations are presented in Appendix C. Given the availability of FX options, we calculate correlation risk premiums for each of the 36 FX pairs formed using the nine G10 exchange rates against the USD. For each FX pair not involving the EUR, our sample period starts in January 1996 and ends in December 2013, for a total of 216 monthly observations. For the EUR, the options data start in January 1999.

The time series mean and standard deviation of the implied FX correlations of each of the 36 G10 FX pairs are reported in Table 2. The cross-sectional average of implied FX correlation means is 0.48, slightly higher than its physical measure counterpart (0.45). Furthermore, implied FX correlation means exhibit less cross-sectional heterogeneity than physical measure ones, with the lowest implied FX correlation mean being 0.14 (CAD/JPY pair) and the highest being 0.88 (CHF/EUR pair). In contrast, realized correlation means range from 0.05 to 0.89. The volatility of implied FX correlations is of the same order of magnitude as the volatility of realized FX correlations, with standard deviations ranging from 0.07 to 0.34 and their cross-sectional average being 0.19.

Finally, the last five columns of Table 2 present the descriptive statistics for FX correlation risk premiums. From left to right, we report the time series mean and standard deviation of the correlation risk premium of each FX pair, followed by the asymptotic t-statistic and the bootstrapped 95% confidence interval of the CRP mean. We find that CRP means exhibit considerable cross-sectional heterogeneity, with their size and sign varying greatly across FX pairs, ranging from −0.069 (CAD/SEK) to 0.099 (JPY/NOK). Roughly two-thirds of CRP means are positive and one third are negative, with their cross-sectional average being 0.016. Furthermore, three-quarters of the means are statistically significant at the 5% level according to either the asymptotic or the bootstrapped distribution. Compared with equity correlation risk premia, FX correlation risk premia are small, with the maximum FX correlation risk premium we find being about half of the equity correlation risk premium reported by Driessen, Maenhout and Vilkov (2009). Furthermore, correlation risk premiums are very volatile. Despite the fact that premiums are much smaller than either realized or implied FX correlations, CRP standard deviations are of the same order of magnitude as those of realized or implied correlations (ranging from 0.06 to 0.22, with a cross-sectional average of 0.14), suggesting substantial time variation in the disparity between physical measure and risk-neutral measure FX correlations.

To explore whether average FX correlation risk premiums exhibit a cross-sectional pattern, we plot the average CRP of all G10 exchange rate pairs against their average realized correlations. Fig. 4 presents the scatterplot, along with the line of best fit. The cross-sectional correlation between average FX correlation risk premiums and average FX realized correlations is −0.55. For example, the AUD/JPY pair, characterized by a very low average realized FX correlation (0.16), has a positive and highly significant average CRP of 0.083. In contrast, the AUD/NZD pair has a very high average realized correlation (0.76) and a negative and significant average premium (−0.016). A cross-sectional regression of average correlation risk premiums on average realized correlations yields a statistically significant slope coefficient of −0.144. Its asymptotic t-statistic, calculated using White (1980) standard errors, is −5.80 and the bootstrapped 95% confidence interval is [−0.154, −0.076]. The strongly negative cross-sectional association between average realized FX correlations and average FX correlation risk premiums is what generates the tighter cross-sectional distribution of average implied FX correlations versus that of realized FX correlations.

[Insert Figure 4 near here.]

The relative tightness of the cross-sectional distribution of conditional FX correlation under the risk-neutral measure implies a potential tension regarding the pricing of FX correlation risk. This is because the negative association between ΔFXC betas and currency excess returns suggests that US investors require a risk premium for being exposed to states in which FXC increases, i.e., in which the cross section of FX correlations widens. However, FX options are priced in a way that indicates that US investors price states in which the cross section of FX correlations tightens.

4.2. The time series properties of correlation risk premiums

We now turn to the time series properties of implied FX correlations and FX correlation risk premiums. The first four columns of Table 7 provide summary statistics on the time series association between realized and implied FX correlations. For each FX pair, we report the unconditional correlation coefficient between the two time series and its asymptotic t-statistic and 95% bootstrapped confidence interval. Realized and implied correlations exhibit substantial comovement across time for all FX pairs, with the unconditional correlations between the two ranging from 0.28 to 0.92, all being statistically significant, and the cross-sectional mean being 0.79.

[Insert Table 7 near here.]

The last four columns of Table 7 report descriptive statistics on the unconditional correlation between realized FX correlations and FX correlation risk premiums. The cross-sectional average of those unconditional correlation coefficients is −0.52 across the 36 G10 FX pairs, suggesting that elevated FX correlation is typically associated with lower than usual CRP, i.e., with a lower than usual disparity between the physical measure and the risk-neutral measure FX correlation. This association is pervasive and robust. Thirty-five of the 36 unconditional correlation coefficients are negative, with all but one of them being statistically significant.

Finally, to assess the cyclicality of correlation risk premiums, we construct CRP cyclicality measures. As we did for FX correlations, we define our CRP cyclicality measures to be the unconditional correlations between FX correlation risk premiums and four market variables ($GVol$, GFI, TED, and VIX). The last four columns of Table 3 report the four CRP cyclicality measures for each of the 36 G10 FX pairs, and Panels A to D of Fig. 5 plot those cyclicality measures against average FX correlation risk premiums. We find a positive cross-sectional association, with FX pairs with high average CRP having countercyclical correlation risk premiums, and pairs with low average CRP having procyclical premiums. The regression results in Panel B of Table 4 suggest that this positive cross-sectional association is statistically significant for all four cyclicality measures. Thus, the cross-sectional dispersion in FX correlation risk premiums is countercyclical. In bad times, the premiums of FX pairs with high average CRP increase and the premiums of FX pairs with low average CRP decline, widening the cross-sectional distribution of FX correlation risk premiums.

[Insert Figure 5 near here.]

5. 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 (2015). In contrast to those models, which assume that innovations in the price of country-specific shocks are uncorrelated across countries, we assume that local risk is priced identically across countries. This assumption implies a lack of spanning of FX correlation risk by exchange rates, a feature that is crucial in jointly explaining the behavior of FX correlations and FX correlation risk premiums.

5.1. Model setup

There are $I+1$ countries ($i = 0, 1, \ldots, I$) in the global economy, with each country having its own currency. Without loss of generality, we call country $i = 0$ the domestic country and countries $i = 1, ..., I$ the foreign countries. We assume that financial markets are frictionless and complete, but there are frictions in the international market for goods. As a result, each country has a unique stochastic discount factor (SDF), but SDFs are not identical across countries. The log SDF of country i, denoted by m^i, is exposed to two global shocks, u^w and u^g, and a country-specific (local) shock u^i, and satisfies

$$-m^i_{t+1} = \alpha + \chi z_t + \varphi z^w_t + \sqrt{\kappa z_t} u^i_{t+1} + \sqrt{\gamma^i z^w_t} u^w_{t+1} + \sqrt{\delta z_t} u^g_{t+1}, \quad (2)$$

where z and z^w are the local and the global pricing factor, respectively. Both pricing factors are common to all countries. Countries are ex ante heterogeneous only with regard to their exposure γ to the first global shock u^w, with all other SDF parameters being identical across countries. In our model, global risk exposure γ is exogenous. Richer models that endogenize unconditional cross-sectional differences in global risk exposure include Hassan (2013), Tran (2013), Backus, Gavazzoni, Telmer and Zin (2013), Colacito, Croce, Gavazzoni and Ready (2015), and Ready, Roussanov and Ward (2016).

The local pricing factor z prices both the local shock u^i and the second global shock u^g. In all countries, the price of the local shock is $\sqrt{\kappa z_t}$ and the price of the second global shock is $\sqrt{\delta z_t}$. Due to differences in γ, the first global shock u^w is differentially priced across countries, with its price in country i being $\sqrt{\gamma^i z_t^w}$.

The two pricing factors are stationary processes. The local pricing factor z is driven by the second global shock u^g, and has law of motion

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

Thus, the local pricing factor is a square root process, reverting to its unconditional mean of $\bar{z}$ at speed λ. Importantly, the local pricing factor is countercyclical, as adverse u^g shocks increase its value.

The global pricing factor z^w is driven by the global shock u^w. It is also a square root process, with law of motion

$$\Delta z_{t+1}^w = \lambda^w(\bar{z}^w - z_t^w) - \xi^w \sqrt{z_t^w} u_{t+1}^w, \tag{4}$$

which also implies countercyclical pricing of risk. To ensure that both pricing factors are strictly positive, we impose the Feller conditions $2\lambda\bar{z} > \xi^2$ and $2\lambda^w\bar{z}^w > (\xi^w)^2$. All parameters except α, χ, and φ are strictly positive and all shocks are independently and identically distributed (i.i.d.) standard normal.

Finally, the inflation process for country i is given by

$$\pi_{t+1}^i = \bar{\pi} + \zeta z_t^w + \sqrt{\sigma}\eta_{t+1}^i. \tag{5}$$

Expected inflation rates are time varying and identical across countries. However, realized inflation rates differ across countries, as inflation shocks η^i are i.i.d. standard normal. Conditional inflation variance is constant and equal to σ and inflation shocks are unpriced, so the model does not feature any inflation risk premiums. As a result, all the salient economic mechanisms in the model arise from real variables, as nominal variables inherit all the conditional properties of their nominal counterparts. For that reason, we discuss the model intuition using real variables and consider nominal variables only in the simulation part (Subsection 5.5).

5.2. The properties of conditional FX moments

We denote the real log exchange rate between foreign currency i and the domestic currency by q^i (units of foreign currency per unit 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,

$$\Delta q_{t+1}^i = m_{t+1}^0 - m_{t+1}^i, \tag{6}$$

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

$$\Delta q_{t+1}^i = \sqrt{\kappa z_t} u_{t+1}^i - \sqrt{\kappa z_t} u_{t+1}^0 + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right) \sqrt{z_t^w} u_{t+1}^w. \tag{7}$$

Differences in γ capture an exchange rate fixed effect. If the foreign country has a higher (lower) exposure γ to global shock u^w than the domestic country, its currency appreciates (depreciates) against the domestic currency when a negative u^w realization occurs. Exposure to the second global shock u^g drops out of exchange rate changes because all countries have the same loading on u^g, so the only global shock that affects exchange rate changes directly is u^w. Thus, in the remainder of the paper, 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:

$$var_t\left(\Delta q^i_{t+1}\right) = 2\kappa z_t + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2 z^w_t. \quad (8)$$

The first effect arises from the country-specific component of stochastic discount factors. Given the independence of local shocks across countries, the higher the impact of local shocks on the SDF, 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 difference in global risk exposure between country i and the domestic country, and the more severely global risk exposure is priced, the more volatile the real exchange rate is.

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

$$cov_t\left(\Delta q^i_{t+1}, \Delta q^j_{t+1}\right) = \kappa z_t + D^{i,j} z^w_t, \quad (9)$$

where the constant $D^{i,j}$ is defined as

$$D^{i,j} \equiv \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)\left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right). \quad (10)$$

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 both of which have either more or less exposure to global risk than the domestic country. Dissimilar exchange rates correspond to pairs of foreign countries in which one country has higher, and the other country lower, exposure to global risk compared with the domestic country.

The first component of conditional FX covariance is due to the common exposure of the two exchange rates to the domestic local shock, as the two exchange rates are mechanically positively correlated through their relation 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 FX comovement that arises from exposure to global FX risk. Foreign countries with similar exposure to the global shock u^w (i.e., countries that satisfy $D^{i,j} > 0$) have exchange rates that comove more than the exchange rates of countries that have dissimilar exposure to global FX risk. Furthermore, fluctuations in z^w have different effects on conditional FX covariance, depending on the type of the FX pair. An increase in the global pricing factor amplifies the importance of exposure to global risk and, thus, increases the conditional covariance of similar exchange rates and reduces the covariance of dissimilar exchange rates.

We can now turn to conditional FX correlations. As happens for FX covariances, country heterogeneity in exposure to the global shock u^w generates cross-sectional heterogeneity in average conditional FX correlations, with similar FX pairs having higher correlations on average than dissimilar ones. Furthermore, the time variation in the pricing factors z^w and z introduces time variation in the conditional correlation of both similar and dissimilar FX pairs and, thus, in the cross-sectional distribution of conditional FX correlation.

To illustrate the effects of the two pricing factors on conditional FX correlations, we consider a world of $I = 3$ foreign countries. Country 1 and Country 2 are less exposed to global FX risk than the domestic country, and Country 3 is more exposed than the domestic country. This implies that the FX pair (1,2) is similar and the 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.

We first consider the impact of the global pricing factor z^w. Panels A, C, and E of Fig. 6 plot conditional FX correlations as a function of z^w for different values of the local pricing factor ($z = 0.2\bar{z}$, $\bar{z}$, and $5\bar{z}$, depicted with circles, solid lines, and squares, respectively). Panel A refers to the similar exchange rate pair (1,2), Panel C considers the dissimilar exchange rate pair (1,3), and Panel E plots the difference in the conditional FX correlations of the two FX pairs. An increase in the global pricing factor z^w raises the relative importance of exposure to the global shock u^w, amplifying similarities and dissimilarities, with similar FX pairs (Panel A) becoming more correlated, and dissimilar

FX pairs (Panel C) becoming less correlated. When $z^w \to \infty$, similar exchange rates become perfectly positively correlated and dissimilar exchange rates become perfectly negatively correlated. Taken together, these results imply that the disparity in conditional FX correlation across exchange rate pairs is increasing in z^w (Panel E).

[Insert Figure 6 near here.]

We now turn to the effects of the local pricing factor z. Panels B, D, and F of Fig. 6 plot the sensitivity of conditional FX correlations to the value of the local pricing factor z for different values of the global pricing factor ($z^w = 0.2\bar{z}$, $\bar{z}$, and $5\bar{z}$), with Panel B referring to the similar FX pair, Panel D to the dissimilar FX pair, and Panel F to the difference in the two pairs' conditional FX correlations. An increase in the local pricing factor z increases both the variance of all exchange rates and the covariance of all exchange rate pairs, due to the domestic currency effect. However, the impact of that effect on FX correlation depends on the type of FX pair. When $z \to \infty$, the correlation of all FX pairs converges to 0.5. This happens because all cross-sectional differences in global risk exposure become second-order and what ultimately drives FX comovement is the domestic currency effect. Consider the limit behavior of log exchange rate changes, described by

$$\Delta q^i_{t+1} \to \sqrt{\kappa z_t} u^i_{t+1} - \sqrt{\kappa z_t} u^0_{t+1}. \tag{11}$$

Exposure to the domestic local shock, which accounts for half of the conditional FX variance and generates all the FX comovement, pushes all FX correlations toward 0.5. Due to the domestic currency effect, when the local pricing factor increases, the importance of similar or dissimilar exposure to global risk is attenuated. As a result, the conditional correlation of similar exchange rates declines (Panel B), and the conditional correlation of dissimilar exchange rates increases (Panel D), leading to a tightening of the cross section of conditional FX correlations (Panel F).

In sum, the cross-sectional dispersion of conditional FX correlations is increasing in the global pricing factor z^w and decreasing in the local pricing factor z. Given that z^w increases after negative u^w shocks and z increases after negative u^g shocks, changes in FXC reflect both u^w shocks (with a positive sign) and u^g shocks (with a negative sign). Empirically, FXC is strongly positively correlated with four market variables that reflect credit risk, illiquidity, and stock market volatility, suggesting that those variables identify exposure to the first global shock u^w, and not to the second global shock u^g. Therefore, those business cycle variables can be proxied in our model by z^w.

5.3. Correlation risk and the cross section of FX returns

The USD excess return for investing in the currency of country i satisfies

$$rx^i_{t+1} - E_t(rx^i_{t+1}) = -\Delta q^i_{t+1} + E_t(\Delta q^i_{t+1}) = -\sqrt{\kappa z_t} u^i_{t+1} + \sqrt{\kappa z_t} u^0_{t+1} - \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right) \sqrt{z^w_t} u^w_{t+1}, \tag{12}$$

so FX excess returns are not exposed to u^g risk. As a result, the conditional risk premium that the domestic investor receives for investing in foreign currency i (including the Jensen term) is

$$rp^i_t \equiv E_t\left(rx^i_{t+1}\right) + \frac{1}{2} var_t(rx^i_{t+1}) = -cov_t(m^0_{t+1}, -\Delta q^i_{t+1}) = \kappa z_t + \left(\sqrt{\gamma^0} - \sqrt{\gamma^i}\right) \sqrt{\gamma^0} z^w_t. \tag{13}$$

FX risk premiums have two components: a part that compensates domestic investors for the fact that investing in a foreign currency essentially entails shorting the country-specific component of the domestic SDF, and a part that reflects compensation for exposure to the global shock u^w. The first component is identical across currencies, so all cross-sectional variation in FX risk premiums is solely due to heterogeneity in exposure to u^w, i.e., heterogeneity in γ. The compensation provided by currency i for exposure to u^w shocks is decreasing in the country loading γ^i. For example, if $\gamma^i < \gamma^0$, then currency i depreciates against the domestic currency when a bad realization of the global shock u^w occurs. Given that $\gamma^0 > 0$, i.e., that a bad realization of u^w increases domestic marginal utility, domestic investors require a positive risk premium to hold currency i. Conversely, currencies of countries with high exposure to u^w ($\gamma^i > \gamma^0$) have a negative premium for global FX risk, as they provide a hedge to domestic investors.

We can now turn to the determinants of the ΔFXC loadings of FX returns. Fluctuations in FXC, the cross-sectional dispersion in conditional FX correlation, reflect innovations in both the global pricing factor z^w (which are scaled multiples of the global shock u^w) and the local pricing factor z^w (scaled multiples of the global shock u^g). Both kinds of innovations are priced and have opposite effects on ΔFXC, so it is not trivial to establish whether a positive loading of an asset return on ΔFXC should be associated with a positive or a negative risk premium. Assets should earn a negative premium for a positive loading on ΔFXC that arises from exposure to u^w, and a positive premium for a positive loading that arises from exposure to u^g. However, no ambiguity exists in the case of FX returns, as the only global innovations to which they are exposed are u^w shocks. As a result, the conditional loading of FX returns on ΔFXC has the same sign as their conditional loading on Δz^w, so in the interests of tractability we can consider the latter. We have

$$\frac{cov_t(rx^i_{t+1}, \Delta z^w_{t+1})}{var_t(\Delta z^w_{t+1})} = \frac{cov_t\left(\left(\sqrt{\gamma^0} - \sqrt{\gamma^i}\right)\sqrt{z^w_t}u^w_{t+1}, -\xi^w\sqrt{z^w_t}u^w_{t+1}\right)}{var_t(-\xi^w\sqrt{z^w_t}u^w_{t+1})} = \frac{\sqrt{\gamma^i} - \sqrt{\gamma^0}}{\xi^w}. \tag{14}$$

Thus, countries i with a higher SDF exposure γ^i to global risk u^w than the domestic country have FX excess returns with a positive conditional loading on ΔFXC. Conversely, the FX returns of countries with $\gamma^i < \gamma^0$ have a negative loading on ΔFXC. Given the negative cross-sectional association between γ and currency risk premiums, those loadings imply a negative risk premium for high ΔFXC beta exchange rates and a positive premium for low ΔFXC beta exchange rates, in line with our empirical findings.

We finish with a note on the cross-sectional relation between interest rates and currency risk premiums. In the model, the real interest rate of country i is given by

$$r^i_t = \alpha + \left(\chi - \frac{1}{2}\kappa - \frac{1}{2}\delta\right)z_t + \left(\varphi - \frac{1}{2}\gamma^i\right)z^w_t, \tag{15}$$

so all cross-sectional heterogeneity in interest rates is due to cross-sectional differences in global risk exposure γ. In all periods, countries with high (low) exposure to global FX risk have a relatively low (high) interest rate, due to a stronger (weaker) precautionary savings motive. As a result, high interest rate currencies are associated with low γs and, thus, high risk premiums.

5.4. The properties of correlation risk premiums

To explore the properties of FX correlation risk premiums, 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 risk-neutral measure law of motion for the global pricing factor z^w is

$$\Delta z^w_{t+1} = \lambda^w(\bar{z}^w - z^w_t) + \xi^w\sqrt{\gamma^0}z^w_t - \xi^w\sqrt{z^w_t}u^{w,\mathbb{Q}}_{t+1}, \tag{16}$$

so the drift adjustment is positive and equal to $\xi^w\sqrt{\gamma^0}z^w_t$. We can rewrite Eq. (16) as a square root process,

$$\Delta z^w_{t+1} = \lambda^{w,\mathbb{Q}}(\bar{z}^{w,\mathbb{Q}} - z^w_t) - \xi^w\sqrt{z^w_t}u^{w,\mathbb{Q}}_{t+1}, \tag{17}$$

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 ($\bar{z}^{w,\mathbb{Q}} > \bar{z}^w$) and is more persistent ($\lambda^{w,\mathbb{Q}} < \lambda^w$) than under the physical measure. Similarly, the risk-neutral measure law of motion for the local pricing factor z is given by

$$\Delta z_{t+1} = \lambda^{\mathbb{Q}}(\bar{z}^{\mathbb{Q}} - z_t) - \xi\sqrt{z_t}u^{g,\mathbb{Q}}_{t+1}, \tag{18}$$

where $\lambda^{\mathbb{Q}} \equiv \lambda - \xi\sqrt{\delta}$ and $\bar{z}^{\mathbb{Q}} \equiv \frac{\lambda}{\lambda^{\mathbb{Q}}}\bar{z}$, so the local pricing factor also has a higher unconditional mean and is more persistent under the risk-neutral measure than under the physical measure. The drift adjustment of the two factors depends crucially on the volatility parameters ξ^w and ξ, which determine the sensitivity of the pricing factors to shocks

u^w and u^g, respectively, and on the exposure parameters γ^0 and δ, which regulate the pricing of shocks u^w and u^g, respectively, for the domestic agent. The higher ξ is relative to ξ^w, and the higher δ is relative to γ^0, the higher the drift adjustment of the local pricing factor is relative to the adjustment of the global pricing factor, as the shocks to the former are more highly priced compared with the shocks to the latter.

The conditional expectation of the global pricing factor is

$$E_t^{\mathbb{Q}}(z_{t+s}^w) = \left(1 - (1 - \lambda^{w,\mathbb{Q}})^s\right) \bar{z}^{w,\mathbb{Q}} + (1 - \lambda^{w,\mathbb{Q}})^s z_t^w \tag{19}$$

under the risk-neutral measure, compared with

$$E_t^{\mathbb{P}}(z_{t+s}^w) = (1 - (1 - \lambda^w)^s)\, \bar{z}^w + (1 - \lambda^w)^s z_t^w \tag{20}$$

under the physical measure, for $s > 0$. Given the higher steady state value and higher persistence of the global pricing factor under the risk-neutral measure, the wedge $E_t^{\mathbb{Q}}(z_{t+s}^w) - E_t^{\mathbb{P}}(z_{t+s}^w)$ is always positive and increasing in z_t^w.[2] Exactly the same is true for the local pricing factor z. Thus, the implied conditional FX correlations are calculated using higher expected values for both z and z^w than their physical counterparts. This is because states characterized by high values of z and z^w are bad states and, thus, receive an elevated probability weight under the risk-neutral measure. The expression for FX correlation risk premiums is derived in Appendix D. Intuitively, the wedge between implied and physical FX correlations is determined by the wedge in the expected values of z and z^w between the two measures, i.e., by the wedge between the risk-neutral and physical measure conditional distributions of z and z^w.

Of particular relevance is the case in which the domestic agent prices fluctuations in the local pricing factor z more heavily than fluctuations in the global pricing factor z^w, i.e., when $\xi \sqrt{\delta} >> \xi^w \sqrt{\gamma^0}$. In that case, the domestic investor risk-adjusts by assigning higher probabilities to states in which z has elevated values. States in which z^w is high also receive elevated importance under the risk-neutral measure, but risk adjustment mainly involves paying attention to high z states. This risk adjustment has implications both for the cross section and the time series of FX correlation risk premiums.

We start with the cross-sectional implications. When investors price z shocks more heavily than z^w shocks, risk adjustment involves paying more attention to states in which the cross-sectional dispersion of FX correlation tightens. High z states are associated with lower than usual FX correlations for similar FX pairs and higher than usual FX correlations for dissimilar pairs (consider Fig. 6), so focusing attention on high z states generates implied FX correlations that are on average lower than physical FX correlations for similar FX pairs. As a result, similar FX pairs (which have high average FX correlations) have negative average FX correlation risk premiums. Conversely, dissimilar FX pairs (which have low average FX correlations) have higher implied FX correlations than physical FX correlations on average and, thus, positive average FX correlation risk premiums. Our model therefore generates a negative cross-sectional association between average FX correlations and average FX correlation risk premiums, in line with the empirical findings presented in Fig. 4.

We now turn to the time series properties of FX correlation risk premiums. We begin by considering similar FX pairs. As discussed in Subsection 5.2, the correlation of similar FX pairs is increasing in the global pricing factor z^w. Although this is true for both implied and physical FX correlations, implied FX correlations are less sensitive to z^w than their physical counterparts. Panel A of Fig. 6 provides a useful visualization. Circles plot FX correlation as a function of z^w conditional on a low z value ($z = 0.2\bar{z}$), and squares plot FX correlation as a function of z^w conditional on a high z value ($z = 5\bar{z}$), with the high z curve (squares) being much flatter than the low z one (circles) in the region of the state space in which the economy spends most of the time (values of z^w between 0 and $2\bar{z}^w$). Because risk adjustment puts more weight on high z states, implied FX correlations are less sensitive to z^w than physical correlations for similar FX pairs. This sensitivity differential means that implied FX correlations increase less than physical correlations in

[2] The wedge is an affine function of z_t^w, with both the constant and the slope coefficient being positive. The constant is positive due to the fact that the function $f(x) = \frac{1-(1-x)^s}{x}$ for $s > 1$ is decreasing in x for $x \in (0, 1)$.

high z^w states (empirically mapped to recessions), reducing the correlation risk premiums of similar FX pairs in those states. Conversely, implied FX correlations drop less than physical FX correlations in low z^w states (booms), increasing the correlation risk premiums of similar FX pairs. In short, the model implies that similar FX pairs have procyclical FX correlation risk premiums and, because they also have countercyclical conditional correlations, the time series correlation between FX correlations and FX correlation risk premiums is negative for similar FX pairs. Similarly, we can use Panel C of Fig. 6 to show that dissimilar FX pairs have countercyclical FX correlation risk premiums, which also implies a negative time series correlation between FX correlations and FX correlation risk premiums for those FX pairs. In short, our model is able to address the key empirical time series properties of FX correlation risk premiums presented in Table 7 and Fig. 5.

In short, conditional FX correlation, which can be indirectly traded using currency options, is exposed to both u^w and u^g innovations. If the domestic agent is pricing z shocks (i.e., u^g innovations) more severely than z^w shocks (u^w innovations), then FX correlation risk premiums largely reflect the desire of currency option holders to avoid high z states, which feature a tightening of the cross-sectional dispersion of FX correlation. In contrast, investing in foreign currency exposes investors only to u^w innovations, so currency risk premiums reflect solely the desire to avoid high z^w states, which are characterized by a widening of the cross-sectional dispersion of FX correlation. Thus, the lack of spanning of FX correlation risk by currency returns, and in particular the lack of exposure of exchange rates to u^g innovations, allows the model to jointly address the empirical properties of FX correlations, FX correlation risk premiums, and currency risk premiums.

5.5. Model simulation

Finally, we assess the quantitative performance of our model and show that it can match key FX correlation moments, as well as the standard interest rate and exchange rate moments.

To illustrate the importance of unspanned FX correlation risk, we consider a nesting model that includes our model and the Lustig, Roussanov and Verdelhan (2014) model as special cases. The law of motion of the local pricing factor of country i, z^i, in the nesting model is

$$\Delta z^i_{t+1} = \lambda(\bar{z} - z^i_t) - \xi \sqrt{z^i_t} \left(\sqrt{\rho} u^g_{t+1} + \sqrt{1-\rho} u^i_{t+1} \right), \tag{21}$$

where $0 \leq \rho \leq 1$, so z^i is driven by both the global shock u^g and the local shock u^i. The nesting model allows for imperfect comovement of (and, thus, for heterogeneity in) local pricing factors across countries. As a result, countries can have different conditional loadings on the global innovation u^g and the exposure to u^g now enters the expression for real exchange rate changes:

$$\Delta q^i_{t+1} = E_t(\Delta q^i_{t+1}) + \sqrt{\kappa z^i_t} u^i_{t+1} - \sqrt{\kappa z^0_t} u^0_{t+1} + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0} \right) \sqrt{z^w_t} u^w_{t+1} + \sqrt{\delta} \left(\sqrt{z^i_t} - \sqrt{z^0_t} \right) u^g_{t+1}. \tag{22}$$

If $\rho = 1$ and all local pricing factors have the same initial value, then all local pricing factors are identical and we retrieve our model, which features unspanned risk. If $\rho = 0$, we retrieve the model in Lustig, Roussanov and Verdelhan (2014), which features independent local pricing factors and in which FX correlation is fully spanned by exchange rates.[3]

Given that our empirical results focus on G10 exchange rates, we simulate our model assuming a global economy with ten countries, the United States and $I = 9$ foreign countries. We simulate the model for different values of ρ, and we run two types of simulations: small sample and large sample. For a given value of ρ, a small sample simulation consists of one thousand simulation paths of 216 monthly observations each, matching the size of our empirical sample. For each simulated moment, the point estimate and the standard error of the moment is, respectively, the moment average across the one thousand simulations and the moment standard deviation across those simulations. We also calculate

[3]The empirical spanning properties of FX correlation are explored in the Online Appendix.

the 95% confidence interval for each moment using the 2.5 and 97.5 percentiles of the moment in the cross section of the one thousand simulation paths. For a given value of ρ, a large sample simulation consists of a single path of 50 thousand monthly observations. The calibration and simulation details are discussed in Appendix E, and the values of our model parameters can be found in Table 8. The output of our small sample simulations is reported in Tables 9 and 10 and Fig. 11, with all other simulation results referring to large sample simulations.

[Insert Table 8 near here.]

Our quantitative analysis starts with the benchmark model, which features perfectly correlated local pricing factors ($\rho = 1$). Table 9 reports empirical and simulated moments for inflation rates, interest rates, and exchange rates. For each empirical moment, we report the value of the moment in our sample, and its bootstrap standard error. The latter equals the standard deviation of the moment across one thousand block bootstrap samples of 216 monthly observations each, with a block length of three monthly observations. We find that our model matches all moments reasonably well.

[Insert Table 9 near here.]

We can now consider FX correlation moments. The first two columns of Table 10 contrast the empirical moments with the benchmark model moments. Our model generates a non trivial cross-sectional spread in average physical and implied FX correlations, in line with the empirical evidence, and is able to closely match their cross-sectional mean. One weakness of the model is that the model-implied premiums are lower (in absolute terms) than their empirical counterparts. As a result, the cross-sectional mean of average premiums in the model, while positive, is lower than the empirical mean (0.71% in the model, compared with 1.58% in the data) and the model is unable to match the wide cross-sectional dispersion in average correlation risk premiums that is observed empirically. Notably, though, the model is able to successfully generate both positive and negative FX correlation risk premiums, as in the data. The model is also able to match the almost perfect positive cross-sectional association between average realized and average implied FX correlations (0.98 in the data, 1.00 in the model) and, crucially, the strongly negative cross-sectional association between average realized correlations and average CRP. In the simulated data, FX pairs with high average FX correlation have negative average CRP and FX pairs with low average FX correlation have positive average CRP, which is consistent with the empirical evidence. Fig. 7 provides an illustration of that feature by plotting the average model-implied CRP against the average model-implied FX correlation for all 36 FX pairs. As regards time series properties, the model generates a perfect time series correlation between realized and implied correlation for all FX pairs, replicating the very high average correlation (0.79) observed in the data, and a negative time series correlation between realized correlation and CRP (−0.77), also in line with the empirical evidence (−0.52).

[Insert Table 10 and Fig. 7 near here.]

In our model, exchange rates are exposed only to the first global shock u^w, so bad states for investors in foreign currencies are those characterized by high values of the global pricing factor z^w. Thus, we explore the cyclicality of FX correlations and FX correlation risk premiums in the model by mapping the countercyclical market variables we use in the empirical part of our paper to z^w. Our aim is to match the empirical cyclicality findings in Figs. 1 and 5. To do so, we follow the same two-step approach we use for our empirical data. First, we calculate the correlation cyclicality measure of each exchange rate pair, equal to the time series correlation of its conditional FX correlation with z^w. Second, we calculate the cross-sectional correlation of the FX correlation cyclicality measures with average FX correlations (36 observations, one for each FX pair). We find that the FX correlation cyclicality measures range from −0.73 to 0.73 across FX pairs and that their cross-sectional correlation with average FX correlations is strongly positive (0.75), suggesting that high correlation FX pairs have countercyclical correlations and low correlation pairs have procyclical correlations, in line with empirical evidence. We repeat the same exercise for correlation risk premiums, and find that the FX CRP cyclicality measures range from −0.78 to 0.79 and that their cross-sectional correlation with average CRP is positive (0.81), again in line with the data.

Our model assumes only one dimension of ex ante heterogeneity across countries, their exposure γ to the global shock u^w. That heterogeneity generates cross-sectional differences in average FX correlations, average interest rates, and average currency excess returns and, thus, engenders cross-sectional linkages among those three measures. Those linkages imply that average correlations across FX pairs are positively associated with both the product of the corresponding foreign currencies' average interest rate differentials $E(r^i - r^0)E\left(r^j - r^0\right)$ and the product of their average currency excess returns $E(rx^i)E(rx^j)$. Those cross-sectional associations in simulated data are presented in Fig. 8. Panel A illustrates the relation between average FX correlations and the product of average nominal interest rate differentials, and Panel B shows the relation between average FX correlations and the product of average currency excess returns. In support of our model, both those model-implied positive cross-sectional associations are present in the data. In the sample of G10 exchange rates, the cross-sectional correlation of average nominal FX correlations with the product of corresponding nominal interest rate differentials is 0.35 and the correlation with the product of average currency excess returns is 0.42.

[Insert Fig. 8 near here.]

Finally, we consider the asset pricing implications of the model. First, we sort the nine currencies into three portfolios according to their nominal interest rate, and report the annualized average excess return of each portfolio in Panel A of Fig. 9. The model generates a strong carry trade effect, with the return on the FX carry portfolio having an annualized average excess return of 2.79%. In congruence with the extant literature, the Lustig, Roussanov and Verdelhan (2011) HML^{FX} factor is priced in the cross section of simulated interest rate-sorted portfolios. Our low, medium, and high interest rate currency portfolios have HML^{FX} betas of −0.41, 0.06, and 0.59, respectively.

[Insert Fig. 9 here.]

Next, we consider currency portfolios sorted on their ΔFXC beta. Their annualized average excess returns are presented in Panel B of Fig. 9. We find that the annualized average excess return for the currency portfolio that is long currencies with low ΔFXC beta and short currencies with a high ΔFXC beta is 1.27%, suggesting a negative price for exposure to FX correlation risk, consistent with our empirical findings. The Lustig, Roussanov and Verdelhan (2011) HML^{FX} factor is priced in the cross section of ΔFXC beta-sorted currency portfolio returns, with the low, medium, and high ΔFXC beta portfolios having an HML^{FX} beta of 0.32, 0.06, and −0.15, respectively. Furthermore, we find a negative cross-sectional association between nominal interest rates and ΔFXC betas. The low, medium, and high ΔFXC beta portfolios have an average interest rate differential (against the domestic country) of 0.81%, 0.16%, and −0.43%, respectively.

For comparison, we turn to the case of non identical local pricing factors across countries ($0 \leq \rho < 1$). Conditional FX variance is given by

$$var_t\left(\Delta q^i_{t+1}\right) = \kappa z^i_t + \kappa z^0_t + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2 z^w_t + \delta\left(\sqrt{z^i_t} - \sqrt{z^0_t}\right)^2, \tag{23}$$

and conditional FX covariance is

$$cov_t\left(\Delta q^i_{t+1}, \Delta q^j_{t+1}\right) = \kappa z^0_t + D^{i,j} z^w_t + \delta\left(\sqrt{z^i_t} - \sqrt{z^0_t}\right)\left(\sqrt{z^j_t} - \sqrt{z^0_t}\right). \tag{24}$$

When the local pricing factors differ across countries, exchange rates are more volatile than in the benchmark model, as differential exposure to u^g increases SDF disparity. As regards FX covariance, exposure to u^g risk has one key difference compared with exposure to u^w. Country exposure to u^w is regulated by the fixed parameter γ and thus is constant over time, so FX pairs are either always similar or always dissimilar regarding their u^w exposure. In contrast, the exposure of each country i to u^g is determined by $\sqrt{\delta z^i}$, so it is unconditionally equal across countries, but time-varying, implying that each FX pair can switch between being similar and being dissimilar with respect to u^g exposure.

To understand the behavior of the cross section of conditional FX correlations, we study the properties of the conditional correlations of similar and dissimilar FX pairs in the special case of independent local pricing factors ($\rho = 0$). The intuition is similar for other values of ρ less than one. Similar to Fig. 6, Fig. 10 illustrates the effect of z^w and z^0 on conditional FX correlations in a world of three foreign countries: Country 1 and Country 2 are less exposed to the first global shock u^w than the domestic country, and Country 3 is more exposed.

[Insert Figure 10 near here.]

Panels A, C, and E of Fig. 10 depict conditional FX correlations as a function of the global pricing factor z^w holding all local pricing factors (domestic and foreign) constant at their common steady state value $\bar{z}$. Not surprisingly, the impact of changes in the global pricing factor z^w is the same as in the model with identical local pricing factors. As z^w increases, similarities and dissimilarities in exposure to global risk get amplified, so the cross-sectional dispersion in FX correlation is increasing in z^w (Panel E).

Panels B, D, and F of Fig. 10 present conditional FX correlations as a function of the domestic local pricing factor z^0, assuming that the global pricing factor z^w and all foreign local pricing factors are equal to their steady state values. We find that the relation between z^0 and conditional FX correlation is not monotonic. For small values of z^0, conditional FX correlation is high for both similar and dissimilar FX pairs (Panel B and Panel D, respectively). In those states, all FX pairs are similar regarding their exposure to u^g, as the loading of all foreign countries is higher than the domestic loading. As the value of z^0 increases, conditional FX correlation decreases, because the component of FX correlation arising from exposure to u^g is attenuated. When z^0 reaches $\bar{z}$, all local factors have identical values, so exposure to u^g does not affect FX moments, as it drops out of exchange rates. Finally, for large values of z^0, all FX pairs are again similar regarding their exposure to u^g, this time because the domestic loading is higher than all foreign loadings, so all FX pairs are highly correlated. As $z^0 \rightarrow \infty$, all FX pairs become conditionally perfectly correlated. Thus, the cross-sectional dispersion of FX correlation is not monotonic in z^0 (Panel F).

The business cycle behavior of FXC, the cross-sectional dispersion of conditional FX correlation, depends on the relative importance of z^w and z^0 for FX correlation determination. The higher the correlation among the local pricing factors, the lower the importance of u^g exposure (and thus z^0) for conditional FX correlation, so high (low) values of ρ are associated with high (low) comovement between FXC and z^w. Panel A of Fig. 11 presents the correlation of FXC with z^w against different values of ρ, showing both the point estimate (solid line) and the 95% confidence interval (shaded area). We find that the correlation between FXC and z^w hovers around zero for almost the entirety of the ρ state space. Even for $\rho = 0.95$, the correlation between the two measures is only 0.02. That correlation jumps to 0.60 for $\rho = 1$, with an associated 95% confidence interval of [0.27,0.83], underscoring the importance of extremely high local pricing factor comovement. In sum, only very high values of ρ lead to empirically plausible and statistically significant correlation between FXC and z^w.

[Insert Fig. 11 near here.]

We now turn to correlation risk premiums, the details for which are discussed in Appendix D. In the special case of independent local pricing factors ($\rho = 0$), the domestic investor only prices z^0 and z^w shocks, and innovations in the foreign local pricing factors are foreign-specific shocks that do not enter the domestic investor's SDF and, thus, are unpriced. In that case, the risk-neutral measure overweighs states in which z^w and z^0 have elevated values. Assuming, as we did for our benchmark model, that the domestic agent prices local shocks more harshly than global shocks, risk adjustment mainly entails paying attention to high z^0 states. As seen in Panels B and D of Fig. 10, those states are characterized by high conditional FX correlations for both similar and dissimilar FX pairs. Thus, pricing states in which the domestic pricing factor z^0 has a high value tends to generate higher implied than physical FX correlations, and thus positive correlation risk premiums, for all FX pairs.

The simulated FX moments of the model with independent local pricing factors ($\rho = 0$) are reported in the third column of Table 10. The cross section of average physical FX correlations is much tighter now than in the benchmark model, as exchange rate exposure to u^g ameliorates the importance of differences in u^w exposure across countries.

The same is true for implied FX correlations. Average FX correlation risk premiums are small for all FX pairs, and, consistent with the discussion above, are positive. The left tail (i.e., the 2.5 percentile) of average CRP is 0.00%, and the right tail (i.e., the 97.5 percentile) is 0.08% and statistically significant. Furthermore, the model generates no cross-sectional association between average FX correlations and average FX correlation risk premiums, at odds with the empirical evidence. This is because the exposure to u^g (which tends to increase the correlation of all FX pairs as z^0 increases and thus generates positive CRP for all FX pairs) offsets the effects of the exposure to u^w, which tends to decrease the correlation of similar FX pairs and increase the correlation of dissimilar FX pairs as z^0 increases and thus generates negative CRP for similar FX pairs and positive CRP for dissimilar FX pairs. Lastly, the model with $\rho = 0$ fails to match the empirical time series properties of FX correlation risk premiums. On average, the time series of simulated physical FX correlations and FX CRP are almost uncorrelated, at odds with the strongly negative correlation that characterizes their empirical counterparts.

To explore the behavior of FX correlation risk premiums for intermediate values of ρ, Panel B of Fig. 11 plots the correlation coefficient of average FX correlations and average CRP for $\rho = \{0, 0.05, ..., 0.95, 1\}$. As the value of ρ increases, and thus the local pricing factors become more correlated across countries, the cross-sectional correlation between average FX correlations and average FX correlation risk premiums tends to decline. We find that high values of ρ are needed for this correlation to become statistically significant, as the cross-sectional correlation is negative and significant at the 5% level only for ρ values of 0.65 and higher. Taken together, Panels A and B of Fig. 11 show that only very high values of ρ can jointly satisfy the physical and the risk-neutral measure properties of FX correlations.

A weakness of our benchmark model, which imposes the polar condition of $\rho = 1$, is that the cross-sectional rank of interest rates (nominal and real) is fixed across time, as cross-sectional interest rate disparity is generated only by the fixed parameter γ. In reality, the cross-sectional ranking of interest rates is time-varying, so this feature of the model is not realistic and precludes matching salient empirical findings, such as the dollar carry trade discussed in Lustig, Roussanov and Verdelhan (2014). However, we can show that a very small relaxation of the assumption of identical local pricing factors allows the model to generate realistic cross-sectional properties of interest rates without compromising the desirable features of the benchmark model for FX correlations.

Consider the average interest rate differential between the foreign countries and the domestic country (AFD, average forward discount):

$$AFD_t = \frac{1}{I}\sum_{i=1}^{I} r_t^i - r_t^0 = \left(\chi - \frac{1}{2}\kappa - \frac{1}{2}\delta\right)\left(\frac{1}{I}\sum_{i=1}^{I} z_t^i - z_t^0\right) + \frac{1}{2}\left(\gamma^0 - \frac{1}{I}\sum_{i=1}^{I}\gamma^i\right) z_t^w. \tag{25}$$

Eq. (25) is valid for both nominal and real interest rate differentials. If the local pricing factor is identical across countries ($\rho = 1$), then the first term drops out and the AFD solely reflects fluctuations in the global pricing factor z^w, never changing sign. However, if the local pricing factors differ across countries ($0 \le \rho < 1$), then the AFD can change sign across time, as it reflects fluctuations in both z^w and the local pricing factors. In the special, and empirically plausible (if the domestic country is the United States) case that the domestic SDF loading on global risk u^w is close to the average foreign loading ($\gamma^0 \simeq \frac{1}{I}\sum_{i=1}^{I}\gamma^i$), the sign of the AFD each period is determined by the sign of the local pricing factor differential. Assuming that the precautionary savings motive dominates the intertemporal smoothing motive ($\chi < \frac{1}{2}\kappa + \frac{1}{2}\delta$) and that the number of foreign countries I is large enough so that the average of the foreign local pricing factors is always close to their common steady state value $\bar{z}$,

$$\frac{1}{I}\sum_{i=1}^{I} z_t^i \to \bar{z}, \tag{26}$$

then the AFD is positive (negative) when the domestic local pricing factor z^0 is higher (lower) than its steady state value. In that case, a domestic investor engaging in the dollar carry trade, i.e., investing in foreign currencies when $AFD > 0$ and shorting them when $AFD < 0$, takes (insures) FX risk when the domestic pricing factor z^0 is transitorily high (low).

To show that our model can address the salient cross-sectional properties of interest rates, we simulate the model setting $\rho = 0.999$, keeping all other parameters at their Table 8 values. In simulated data, this ρ value implies an average cross-sectional correlation of 0.999 for the local pricing factors. The simulated moments are presented in the last column of Table 10. Overall, we find that the model with $\rho = 0.999$ preserves the key FX correlation features of the benchmark model. As regards the dollar carry trade, its empirical annualized return for the G10 currencies from January 1996 to December 2013 is 5.26% using the nominal AFD and 3.48% using the real AFD. In the model, the two strategies are identical, yielding an annualized return of 1.82%, so the model undershoots both empirical returns. However, the model is able to almost perfectly match the turnover of interest rate-sorted currency portfolios, as it generates a monthly turnover of 0.049, virtually identical to the empirical turnover of 0.047 observed in the G10 sample from January 1996 to December 2013.

6. Conclusion

We show that FX correlations become more cross-sectionally dispersed in adverse economic states, and we construct an FX correlation dispersion measure, denoted by FXC and defined as the difference between the conditional correlation of the most and least conditionally correlated FX pairs. We then sort currencies into portfolios based on their exposure to FXC innovations and show that the spread between high and low ΔFXC beta currency portfolios is economically and statistically significant (6.42% annually), suggesting that investors want to be compensated for investing in currencies that perform badly during periods of increased cross-sectional dispersion in conditional FX correlations. Then, defining the FX correlation risk premium as the difference between the FX correlation under the risk-neutral and the physical probability measures, we find a strongly negative cross-sectional association between average FX correlations and average FX correlation risk premiums. FX pairs with high average correlation exhibit low (or negative) average correlation risk premiums, and the opposite is true for FX pairs with low average correlations.

We rationalize our empirical findings with a no-arbitrage model of exchange rates that is able to jointly match the salient properties of FX correlations under both the physical and the risk-neutral measure. Our findings suggest that a possible avenue for richer no-arbitrage models that feature endogenously determined stochastic discount factors and aim to explain the dynamics of FX correlation is the incorporation of unspanned risk. In that class of models, any shock that affects countries' SDF identically (and thus does not enter exchange rates) and causes the cross section of FX correlation to tighten, has the potential to address the apparent inconsistency between the behavior of FX correlations under the physical measure and under the risk-neutral measure. That said, we stress that unspanned risk is not the only possible avenue to be considered. It is possible that alternative economic mechanisms, including market segmentation or other frictions in financial markets, can also play a role in addressing our empirical findings.

Appendix A. Realized FX moments

We use daily spot exchange rates to calculate measures of realized FX moments. $\Delta s_t^i = \ln\left(S_t^i\right) - \ln(S_{t-1}^i)$ denotes the daily log change for exchange rate i. The annualized realized FX variance observed at t is then calculated as

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

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.

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

$$\mathrm{RCov}_t^{i,j} = \frac{252}{K} \sum_{k=0}^{K-1} \Delta s_{t-k}^i \Delta s_{t-k}^j. \tag{A.2}$$

Finally, the realized FX correlation is defined as the ratio of corresponding realized FX covariance and the product of the respective FX standard deviations:

$$\mathrm{RC}_t^{i,j} = \mathrm{RCov}_t^{i,j} / \sqrt{\mathrm{RV}_t^i}\sqrt{\mathrm{RV}_t^j}. \tag{A.3}$$

Appendix B. Price of FX correlation risk

We consider the two-factor model,

$$\mathrm{E}[rx^i] = \beta_i^{DOL}\lambda^{DOL} + \beta_i^{HML^C}\lambda^{HML^C}, \tag{B.1}$$

where rx^i denotes the excess return in levels (i.e., corrected for the Jensen term). To estimate the factor prices λ^{DOL} and λ^{HML^C}, we follow the two-stage procedure of Fama and MacBeth (1973). First, we run a time series regression of excess returns on the factors. Second, we run a cross-sectional regression of average excess returns on factor betas. We do not include a constant in the cross-sectional regression of the second stage, as the dollar factor DOL essentially acts a constant, as discussed in Lustig, Roussanov and Verdelhan (2011).

The first part of Table B1 reports the first-stage regression results. We consider 15 test assets: three currency portfolios sorted on exposure to ΔFXC ($\mathrm{Pf1^C}$, $\mathrm{Pf2^C}$, and $\mathrm{Pf3^C}$), three currency portfolios sorted on forward discounts (called carry portfolios and denoted by $\mathrm{Pf1^F}$, $\mathrm{Pf2^F}$, and $\mathrm{Pf3^F}$) and nine individual G10 exchange rates. As expected, the HML^C betas of the ΔFXC beta-sorted portfolios are monotonically increasing. Notably, the HML^C betas of the carry portfolios are monotonically decreasing, with low (high) interest rate currencies having a positive (negative) HML^C beta. Finally, the HML^C betas for the individual G10 currencies are highly negatively correlated with their average excess returns over the sample period, with the correlation coefficient being −0.92.

[Insert Table B1 near here.]

The second part of Table B1 presents the second-stage results for various sets of test assets. Set (1) contains only the three ΔFXC beta-sorted portfolios (Pf1^C to Pf3^C) and the three carry portfolios (Pf1^F to Pf3^F), and Set (2) contains the test assets of Set (1) along with the nine individual G10 currencies. For both sets, we report the point estimates of the prices of risk, along with their standard errors (in parentheses) and Shanken (1992)–corrected standard errors (in brackets). We also report the R^2 of each second-stage regression. Correlation risk has a significantly negative price, with λ^{HML^C} being −0.58% (−0.54%) per month for Set (1) [Set (2)]. Those estimates are not significantly different from the average HML^C return of −0.54% per month. The second-stage R^2 is very high for both regressions (0.99 and 0.93, respectively).

For robustness, we consider additional developed and emerging country currencies. Set (3) of test assets contains four ΔFXC beta-sorted and four forward discount-sorted portfolios, using all developed country currencies. Set (4) contains four ΔFXC beta-sorted and four forward discount-sorted portfolios, using the full set of currencies. To conserve space, we discuss the first-stage regression results for the test assets in Sets (3) and (4) in the Online Appendix, and provide only the second-stage results in Table B1. We find that the λ^{HML^C} estimates are in line with our benchmark results. The price of correlation risk is estimated at −0.51% and −0.67% per month in Sets (3) and (4), respectively, with both estimates being statistically significant at the 5% level. The regression R^2 is 0.90 for Set (3) and 0.81 for Set (4). For robustness, the Online Appendix contains price of risk estimates using FXC innovations, a non-traded factor, in lieu of HML^C returns, a traded factor. We find that FXC innovations also have a negative price in the cross section of currency returns.

Our traded correlation risk factor HML^C acts as a slope factor regarding the pricing of currency risk, so a natural question that arises regards the relation between HML^C and the Lustig, Roussanov and Verdelhan (2011) carry trade factor HML^{FX}. The carry trade factor reflects the returns to a portfolio that invests in high interest rate currencies and shorts low interest rate currencies, and has also been shown to act as a slope factor. Using monthly data from January 1996 to December 2013, we find that the two factors are strongly negatively correlated, with the correlation coefficient between the two time series being −0.66. This result suggests that the two factors capture similar sources of risk.

The highly negative association between HML^{FX} and HML^C is fully consistent with our proposed no-arbitrage model. In the model, the excess return to the carry trade portfolio is defined as

$$HML^{FX}_{t+1} = \frac{1}{N}\sum_{i\in HF} rx^i_{t+1} - \frac{1}{N}\sum_{i\in LF} rx^i_{t+1}, \tag{B.2}$$

with high interest rate (low γ, according to the model) currencies in set HF and low interest rate (high γ) currencies in set LF. Provided that currency portfolios contain enough currencies so that the local shocks average zero, HML^{FX} innovations are perfectly positively correlated with the global shock u^w:

$$HML^{FX}_{t+1} - E_t\left(HML^{FX}_{t+1}\right) = \frac{1}{N}\left(\sum_{i\in LF}\sqrt{\gamma^i} - \sum_{i\in HF}\sqrt{\gamma^i}\right)\sqrt{z^w_t}\,u^w_{t+1}. \tag{B.3}$$

Thus, HML^{FX} returns capture exposure to the global shock u^w, which is the only global shock priced in currency markets.

In contrast, FXC innovations capture both kinds of global shocks, u^w and u^g, so they provide a very noisy measure of the part of FX correlation risk that is priced in foreign exchange markets. It follows that HML^{FX} always has better pricing ability than ΔFXC in the cross section of currency returns. To get a cleaner measure of u^w innovations, we can consider FX return differentials, which are exposed only to u^w shocks. In particular, consider portfolio HML^C, which is long currencies with high ΔFXC loading and short currencies with low ΔFXC loading. Its return is

$$HML^C_{t+1} = \frac{1}{N}\sum_{i\in HC} rx^i_{t+1} - \frac{1}{N}\sum_{i\in LC} rx^i_{t+1}, \tag{B.4}$$

with high-ΔFXC-loading (i.e., high γ) currencies in set HC and low-ΔFXC-loading (low γ) currencies in set LC. Provided that the long and the short positions of the portfolio contain enough currencies so that the local shocks cancel out, the return innovations of the HML^C portfolio are perfectly negatively correlated with the global shock u^w:

$$HML^C_{t+1} - E_t\left(HML^C_{t+1}\right) = \frac{1}{N}\left(\sum_{i\in LC}\sqrt{\gamma^i} - \sum_{i\in HC}\sqrt{\gamma^i}\right)\sqrt{z^w_t}\,u^w_{t+1}. \tag{B.5}$$

Therefore, HML^C return innovations are perfectly negatively correlated with HML^{FX} return innovations, as they both reflect u^w shocks and, thus, should have the same explanatory power for the cross section of FX returns. In short, high γ currencies, which hedge u^w risk, have low interest rates, high HML^C betas, low HML^{FX} betas and low risk premiums. Conversely, low γ (i.e. high interest rate, low HML^C beta, high HML^{FX} beta) currencies have high risk premiums.

In the Online Appendix, we discuss the relation between our FX correlation risk factor and another known FX risk factor, the FX volatility factor of Menkhoff, Sarno, Schmeling and Schrimpf (2012).

Appendix C. Implied FX moments

We follow Demeterfi, Derman, Kamal and Zou (1999) and Britten-Jones and Neuberger (2000) to obtain a model-free measure of implied volatility. They show that if 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:

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

where S_t is the underlying spot exchange rate, $\mathrm{P}(K,T)$ and $\mathrm{C}(K,T)$ are the respective put and call option prices with maturity date T and strike price K, and r is the continuously compounded interest rate of the quote currency. In practice, the number of traded options for any underlying asset is finite, so the available strike price series is a finite sequence. Calculating the model-free implied variance involves the entire cross section of option prices, so, 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 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. See, e.g., Wystup (2006) for the specifics of FX options conventions.

To approximate the integral in Eq. (C.1), we adopt a trapezoidal integration scheme over the range of strike prices in our data set. Jiang and Tian (2005) report two types of implementation errors: truncation errors due to the unavailability of an infinite range of strike prices, and discretization errors that arise due to the lack of availability 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. Brandt and Diebold (2006) use the same approach to construct realized covariances of exchange rates from range-based volatility estimators. Our construction methodology relies on state prices being sufficiently similar for the different agents (countries). For the construction, we require all cross rates for three currencies, S_t^i, S_t^j, and S_t^{ij}, i.e., the two exchange rates against the domestic (base) currency and the exchange rate between the two foreign currencies. The absence of triangular arbitrage then implies that $S_t^{ij} = S_t^i / S_t^j$. Taking logs, we derive the equation

$$\ln\left(\frac{S_T^{ij}}{S_t^{ij}}\right) = \ln\left(\frac{S_T^i}{S_t^i}\right) - \ln\left(\frac{S_T^j}{S_t^j}\right). \tag{C.2}$$

Finally, taking variances yields

$$\int_t^T \left(\sigma_u^{ij}\right)^2 du = \int_t^T \left(\sigma_u^i\right)^2 du + \int_t^T \left(\sigma_u^j\right)^2 du - 2\int_t^T \gamma_u^{i,j} du, \tag{C.3}$$

where $\gamma_t^{i,j}$ denotes the covariance of returns between domestic currency FX pairs i and j. Solving for the covariance term, we obtain

$$\int_t^T \gamma_u^{i,j} du = \frac{1}{2}\int_t^T \left(\sigma_u^i\right)^2 du + \frac{1}{2}\int_t^T \left(\sigma_u^j\right)^2 ds - \frac{1}{2}\int_t^T \left(\sigma_u^{ij}\right)^2 du. \tag{C.4}$$

Using the standard replication arguments, we find that

$$\begin{aligned} \mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \gamma_u^{i,j}\, du\right) = \; & e^{r(T-t)}\Big(\int_t^{S_t^i} \frac{1}{K^2}\mathrm{P}^i(K,T)\, dK + \int_{S_t^i}^{\infty} \frac{1}{K^2}\mathrm{C}^i(K,T)\, dK \\ & + \int_t^{S_t^j} \frac{1}{K^2}\mathrm{P}^j(K,T)\, dK + \int_{S_t^j}^{\infty} \frac{1}{K^2}\mathrm{C}^j(K,T)\, dK \\ & - \int_t^{S_t^{ij}} \frac{1}{K^2}\mathrm{P}^{ij}(K,T)\, dK - \int_{S_t^{ij}}^{\infty} \frac{1}{K^2}\mathrm{C}^{ij}(K,T)\, dK\Big). \end{aligned} \tag{C.5}$$

The model-free implied correlation can then be calculated using expression (C.5) and the model-free implied variance expression (C.1):

$$\mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \rho_u^{i,j} du\right) \equiv \frac{\mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \gamma_u^{i,j} ds\right)}{\sqrt{\mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \left(\sigma_u^i\right)^2 du\right)}\sqrt{\mathrm{E}_t^{\mathbb{Q}}\left(\int_t^T \left(\sigma_u^j\right)^2 du\right)}}. \tag{C.6}$$

Our methodology relies on the absence of triangular arbitrage in currency markets. Recent 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, 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 a negligible effect on our calculations.

There is recent empirical evidence suggesting that jump risk could be present in the FX market, see, e.g., Chernov, Graveline and Zviadadze (2016), Jurek (2014), and Farhi, Fraiberger, Gabaix, Ranciere and Verdelhan (2015). Britten-Jones and Neuberger (2000) show that the risk-neutral expected integrated return variance is fully specified by a continuum of call and put options, provided that the price of the underlying asset is a diffusion process. However, in the Online Appendix, we show that our analysis is robust to the presence of jumps.

Appendix D. FX correlation risk premiums in the model

For period $[t, T]$, the expected variance of the changes in the log exchange rate i is given by

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

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

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

For the local pricing factor, we have

$$E_t^{\mathbb{Q}}(z_{t+s}) = \left(1 - (1 - \lambda^{\mathbb{Q}})^s\right) \bar{z}^{\mathbb{Q}} + (1 - \lambda^{\mathbb{Q}})^s z_t \equiv A_s^{\mathbb{Q}} + B_s^{\mathbb{Q}} z_t \tag{D.3}$$

under the risk-neutral measure and

$$E_t(z_{t+s}) = (1 - (1 - \lambda)^s)\, \bar{z} + (1 - \lambda)^s z_t \equiv A_s + B_s z_t \tag{D.4}$$

under the physical measure, with $A_s^{\mathbb{Q}} > A_s$ and $B_s^{\mathbb{Q}} > B_s$ for all $s > 0$. A similar notation can be used for the global pricing factor z^w. For $X_s = \{A_s, B_s, A_s^{\mathbb{Q}}, B_s^{\mathbb{Q}}, A_s^w, B_s^w, A_s^{w,\mathbb{Q}}, B_s^{w,\mathbb{Q}}\}$, we respectively define $X = \{A, B, A^{\mathbb{Q}}, B^{\mathbb{Q}}, A^w, B^w, A^{w,\mathbb{Q}}, B^{w,\mathbb{Q}}\}$ as $X \equiv \sum_{s=0}^{T-t-1} X_s$.

The expected FX correlation is defined as the ratio of the corresponding expected FX covariance over the product of the square root of the two FX variances, as in Section 3. Thus, the FX correlation risk premium can be written as

$$CRP_t^{i,j} = \frac{\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t\right) + D^{i,j}\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w\right)}{\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w\right)}\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t\right) + \left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)^2\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w\right)}} \tag{D.5}$$

$$- \frac{\kappa\left(A + Bz_t\right) + D^{i,j}\left(A^w + B^w z_t^w\right)}{\sqrt{2\kappa\left(A + Bz_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2\left(A^w + B^w z_t^w\right)}\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t\right) + \left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)^2\left(A^w + B^w z_t^w\right)}}. \tag{D.6}$$

The magnitude of the correlation risk premium depends on the difference between the risk-neutral measure parameters $A^{\mathbb{Q}}$, $B^{\mathbb{Q}}$, $A^{w,\mathbb{Q}}$, and $B^{w,\mathbb{Q}}$ and the physical measure parameters A, B, A^w, and B^w. When the domestic agent prices fluctuations in the local pricing factor more heavily than fluctuations in the global pricing factor, i.e., when $\xi\sqrt{\delta} >> \xi^w\sqrt{\gamma^0}$, then

$$\left(A^{\mathbb{Q}} + B^{\mathbb{Q}}z_t\right) - \left(A + Bz_t\right) >> \left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}}z_t^w\right) - \left(A^w + B^w z_t^w\right), \tag{D.7}$$

implying that the risk adjustment for the local pricing factor z is quantitatively larger than the risk adjustment for the global pricing factor z^w as regards FX correlation. The implications of such risk adjustment for the cross-sectional and time series properties of FX correlation risk premiums are discussed in Section 4.

As regards the nesting model, the law of motion for the global pricing factor z^w under the risk-neutral measure is identical to its risk-neutral measure law of motion in the model with identical pricing factors, given in Eq. (17), and the law of motion of the domestic local pricing factor z^0 is

$$\Delta z_{t+1}^0 = \lambda^{0,\mathbb{Q}}\left(\bar{z}^{0,\mathbb{Q}} - z_t^0\right) - \xi\sqrt{z_t^0}\left(\sqrt{\rho}u_{t+1}^{g,\mathbb{Q}} + \sqrt{1-\rho}u_{t+1}^{0,\mathbb{Q}}\right), \tag{D.8}$$

where $\bar{z}^{0,\mathbb{Q}} \equiv \frac{\lambda}{\lambda^{0,\mathbb{Q}}}\bar{z}$ and $\lambda^{0,\mathbb{Q}} = \lambda - \xi\left(\sqrt{\rho}\sqrt{\delta} + \sqrt{1-\rho}\sqrt{\kappa}\right)$, as both components of the innovations in z^0 are priced by the domestic investor. For the foreign local pricing factors z^i with $i = 1, ..., I$, the risk-neutral measure law of motion is

$$\Delta z_{t+1}^i = \lambda\left(\bar{z} - z_t^i\right) + \xi\sqrt{\rho}\sqrt{\delta}\sqrt{z_t^i}\sqrt{z_t^0} - \xi\sqrt{z_t^i}\left(\sqrt{\rho}u_{t+1}^{g,\mathbb{Q}} + \sqrt{1-\rho}u_{t+1}^{i,\mathbb{Q}}\right), \tag{D.9}$$

as the domestic investor prices only the global component $\sqrt{\rho}u^g$ of the foreign local pricing factor innovations, but not their local component $\sqrt{1-\rho}u^i$.

Appendix E. Model calibration and simulation

Excluding ρ, the nesting model has $14 + (I + 1)$ parameters in total: five common SDF parameters (α, χ, ϕ, κ, and δ), $I + 1$ heterogeneous parameters (the loading γ^i for each country), six common 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 three common inflation parameters ($\bar{\pi}$, ζ, and σ).

To calibrate our benchmark model, we impose $\rho = 1$ and then largely follow Lustig, Roussanov, and Verdelhan (2011, 2014). We reduce the set of parameters by imposing the constraint that the loadings γ^i are equally spaced across the foreign countries. We assume that the first foreign country has loading γ^{min}, the last foreign country has loading γ^{max}, and each intermediate foreign country $i = 2, ..., I - 1$ has loading $\gamma^i = \gamma^{min} + \frac{i-1}{I-1}(\gamma^{max} - \gamma^{min})$. To generate a large effect of the local pricing factor, in line with our model, we set $\delta = 40$ and $\lambda = 0.25$. The latter value ensures that the local pricing factor z is stationary under both the physical and the risk-neutral measure. Furthermore, we set γ^{min} to 0.20 [instead of 0.18, as in the Lustig, Roussanov and Verdelhan (2014) calibration], to achieve a more realistic cross-sectional dispersion in interest rates and FX correlations. In unreported results, using $\gamma^{min} = 0.18$ does not affect our results substantially. All the other parameters, with the exception of χ, ξ, ξ^w, and $\bar{\pi}$, are set equal to the corresponding values in Lustig, Roussanov and Verdelhan (2014), which are set to target specific interest rate, inflation, and exchange rate moments, but no moments related to FX correlations or FX correlation risk premiums. Finally, we set χ, ξ, ξ^w, and $\bar{\pi}$ using a Generalized Method of Moments approach as follows. We target three moments: the cross-sectional average of the time series mean and variance of the real interest rates of the ten countries, and the cross-sectional average of the time series mean of the inflation rates of the ten countries. In the estimation, we leave $\bar{\pi}$ unconstrained but constrain the ratio of $\frac{\xi}{\xi^w}$ to equal 2.43, which is the parameter ratio in the Lustig, Roussanov and Verdelhan (2014) calibration. The values of our calibrated parameters are reported in Table 8. Regarding the calibration data, we proxy interest rate differentials against the USD by the corresponding forward discounts, and the nominal USD interest rate by the Fama-French one-month Treasury bill rate. Inflation in each country is calculated using the corresponding Consumer Price Index, and real interest rates are calculated as the difference between nominal interest rates and inflation rates.

Finally, we simulate the model for different values of ρ. We consider two types of simulations: small sample and large sample. For a given value of ρ, a small sample simulation consists of one thousand simulation paths of 5,216 monthly observations each, initialized at the steady state values $\bar{z}$ and $\bar{z}^w$. To reduce the effect of initial conditions, we discard the first five thousand observations, so we are left with 216 observations for each path, allowing us to study the small-sample properties of the moments of interest. For a given value of ρ, a large-sample simulation consists of a single path of 55 thousand monthly observations, initialized at the steady state values $\bar{z}$ and $\bar{z}^w$. Again, we discard the first five thousand observations and calculate moments using the last 50 thousand observations. For both kinds of simulations, conditional FX moments (realized and implied) are 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. At each period, conditional expectations are calculated using averages across one hundred simulations, with the exception of the benchmark model ($\rho = 1$), in which case we use closed-form expressions for the conditional expectations.

References

Backus, D.K., Gavazzoni, F., Telmer, C., Zin, S.E., 2013. Monetary policy and the uncovered interest rate parity puzzle. Unpublished working paper, New York University, NY.

Bollerslev, T.G., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.

Brandt, M.W., Diebold, F.X., 2006. A no-arbitrage approach to range-based estimation of return covariances and correlations. Journal of Business 79, 61–73.

Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. Journal of Finance 55, 839–866.

Chernov, M., Graveline, J., Zviadadze, I., 2016. Crash risk in currency returns. forthcoming, Journal of Financial and Quantitative Analysis .

Colacito, R., Croce, M.M., 2013. International asset pricing with recursive preferences. Journal of Finance 68, 2651–2686.

Colacito, R., Croce, M.M., Gavazzoni, F., Ready, R.C., 2015. Currency risk factors in a recursive multi-country economy. Unpublished working paper, University of North Carolina Kenan-Flagler Business School, Chapel Hill, NC.

Demeterfi, K., Derman, E., Kamal, M., Zou, J., 1999. A guide to volatility and variance swaps. Journal of Derivatives 6, 9–32.

Driessen, J., Maenhout, P., Vilkov, G., 2009. The price of correlation risk: Evidence from equity options. Journal of Finance 64, 1377–1406.

Fama, E.F., 1984. Forward and spot exchange rates. Journal of Monetary Economics 14, 319–338.

Fama, E.F., MacBeth, J., 1973. Risk, return, and equilibrium: Empirical tests. Journal of Political Economy 81, 607–636.

Farhi, E., Fraiberger, S., Gabaix, X., Ranciere, R., Verdelhan, A., 2015. Crash risk in currency markets. Unpublished working paper, New York University, New York, NY.

Fenn, D.J., Howison, S.D., McDonald, M., Williams, S., Johnson, N.F., 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., Kohlhagen, S.W., 1983. Foreign currency option values. Journal of International Money and Finance 2, 231–237.

Hassan, T., 2013. Country size, currency unions, and international asset returns. Journal of Finance 68, 2269–2308.

Hu, G.X., Pan, J., Wang, J., 2013. Noise as information for illiquidity. Journal of Finance 68, 2341–2382.

Jiang, G., Tian, Y., 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., Tham, W.W., 2012. Execution risk in high-frequency arbitrage. Management Science 58, 2131–2149.

Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. Review of Financial Studies 24, 3731–3777.

Lustig, H., Roussanov, N., Verdelhan, A., 2014. Countercylical currency risk premia. Journal of Financial Economics 111, 527–553.

Malkhozov, A., Mueller, P., Vedolin, A., Venter, G., 2016. International illiquidity. Unpublished working paper, London School of Economics, London, UK.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012. Carry trades and global foreign exchange volatility. Journal of Finance 67, 681–718.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Ready, R., Roussanov, N., Ward, C., 2016. Commodity trade and the carry trade: a tale of two countries. forthcoming, Journal of Finance .

Shanken, J., 1992. On the estimation of beta pricing models. Review of Financial Studies 5, 1–34.

Tran, N.K., 2013. Growth risk of nontraded industries and asset pricing. Unpublished working paper, Olin Business School at Washington Univesity, St. Louis, MO.

Verdelhan, A., 2015. The share of systematic risk in bilateral exchange rates. forthcoming, Journal of Finance .

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.

Wystup, U., 2006. FX Options and Structured Products. John Wiley and Sons, Chichester, England.

Table 1
Summary statistics: individual currencies.
The table reports summary statistics for individual currencies. For each foreign currency i we report the mean, standard deviation, Sharpe ratio, skewness, and kurtosis of US dollar excess returns $f_t^i - s_{t+1}^i$ and the mean forward discount $f_t^i - s_t^i$. Excess returns are annualized and expressed in percentage points. Panel A reports monthly data from January 1996 through December 2013; Panel B, monthly data from January 1984 through December 2013. In both panels, before January 1999, we use the German mark (DEM) in the place of the euro. AUD = Australian dollar; CAD = Canadian dollar; CHF = Swiss franc; GBP = Pound sterling; JPY = Japanese yen; NOK = Norwegian krone; NZD = New Zealand dollar; SEK = Swedish krona.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Panel A: January 1996–December 2013									
Mean	3.01	1.12	-0.39	-0.46	1.37	-2.74	1.17	3.73	0.22
Standard deviation	12.78	8.50	10.91	10.25	8.50	10.78	11.15	13.09	11.22
Sharpe ratio	0.24	0.13	-0.04	-0.05	0.16	-0.25	0.11	0.29	0.02
Skewness	-0.60	-0.60	0.13	-0.15	-0.50	0.48	-0.36	-0.37	-0.08
Kurtosis	5.29	7.26	4.40	3.80	4.73	5.22	4.10	4.85	3.61
$f_t - s_t$	2.12	-0.04	-2.00	-0.60	0.91	-3.01	0.98	2.70	-0.10
Panel B: January 1984–December 2013									
Mean	2.96	1.15	1.21	1.60	2.43	0.14	2.99	4.88	2.34
Standard deviation	12.08	7.15	11.93	11.14	10.37	11.38	11.05	13.25	11.36
Sharpe ratio	0.24	0.16	0.10	0.14	0.23	0.01	0.27	0.37	0.21
Skewness	-0.72	-0.65	0.00	-0.21	-0.23	0.32	-0.48	-1.01	-0.46
Kurtosis	5.62	8.90	3.56	3.43	5.36	4.26	4.20	9.41	4.44
$f_t - s_t$	3.12	0.77	-1.83	-0.61	1.89	-2.64	2.23	4.15	1.60

Table 2

Summary statistics: foreign exchange (FX) correlations and FX correlation risk premiums.

The table reports means and standard deviations for realized and implied FX correlations (RC and IC, respectively), as well as FX correlation risk premiums (CRPs), for all FX pairs. Correlation risk premiums are defined as the difference between the implied and realized correlations. Realized correlations are calculated using past daily log exchange rate changes over a three month window. Implied correlations are calculated from daily option prices on the underlying exchange rates. The last two columns report the bootstrapped 95% confidence interval (using the 2.5 and 97.5 percentiles). We use monthly data from January 1996 to December 2013 [options data for the euro (EUR) start in January 1999]. AUD = Australian dollar; CAD = Canadian dollar; CHF = Swiss franc; GBP = Pound sterling; JPY = Japanese yen; NOK = Norwegian krone; NZD = New Zealand dollar; SEK = Swedish krona.

	RC		IC		CRP				
FX pair	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation	*t*-statistic	2.5%	97.5%
AUDCAD	0.471	0.25	0.430	0.27	-0.041	0.15	-4.07	-0.060	-0.023
AUDCHF	0.357	0.27	0.405	0.20	0.048	0.15	4.73	0.028	0.068
AUDEUR	0.450	0.28	0.544	0.16	0.019	0.09	2.81	0.006	0.031
AUDGBP	0.422	0.24	0.453	0.19	0.031	0.12	3.86	0.014	0.046
AUDJPY	0.155	0.34	0.238	0.26	0.083	0.16	7.58	0.062	0.103
AUDNOK	0.467	0.26	0.431	0.29	-0.036	0.20	-2.64	-0.064	-0.010
AUDNZD	0.755	0.16	0.739	0.15	-0.016	0.08	-2.97	-0.026	-0.005
AUDSEK	0.474	0.25	0.480	0.20	0.005	0.13	0.61	-0.012	0.022
CADCHF	0.233	0.28	0.283	0.21	0.050	0.15	4.94	0.031	0.070
CADEUR	0.307	0.30	0.405	0.19	0.024	0.13	2.45	0.005	0.044
CADGBP	0.281	0.27	0.307	0.23	0.025	0.15	2.34	0.004	0.044
CADJPY	0.054	0.26	0.136	0.19	0.082	0.16	7.33	0.060	0.104
CADNOK	0.340	0.28	0.341	0.28	-0.002	0.18	-0.17	-0.028	0.022
CADNZD	0.413	0.23	0.352	0.34	-0.061	0.22	-4.19	-0.092	-0.035
CADSEK	0.352	0.26	0.287	0.29	-0.069	0.17	-5.96	-0.094	-0.047
CHFEUR	0.888	0.13	0.875	0.12	-0.010	0.08	-1.69	-0.020	0.002
CHFGBP	0.580	0.19	0.605	0.15	0.025	0.11	3.32	0.010	0.039
CHFJPY	0.405	0.26	0.456	0.18	0.051	0.14	5.15	0.032	0.070
CHFNOK	0.726	0.16	0.731	0.12	0.006	0.11	0.73	-0.009	0.021
CHFNZD	0.358	0.23	0.370	0.20	0.012	0.16	1.06	-0.010	0.033
CHFSEK	0.707	0.16	0.712	0.13	0.004	0.10	0.58	-0.010	0.017
EURGBP	0.644	0.15	0.683	0.10	0.003	0.08	0.54	-0.009	0.015
EURJPY	0.324	0.27	0.364	0.20	0.067	0.15	5.84	0.046	0.089
EURNOK	0.825	0.09	0.798	0.07	-0.025	0.06	-5.20	-0.035	-0.016
EURNZD	0.440	0.23	0.501	0.17	0.005	0.12	0.55	-0.013	0.022
EURSEK	0.816	0.11	0.817	0.08	-0.022	0.06	-4.64	-0.031	-0.012
GBPJPY	0.217	0.26	0.293	0.19	0.076	0.15	7.29	0.056	0.095
GBPNOK	0.577	0.16	0.638	0.12	0.059	0.16	5.39	0.038	0.080
GBPNZD	0.415	0.23	0.404	0.22	-0.011	0.14	-1.15	-0.029	0.006
GBPSEK	0.560	0.16	0.598	0.13	0.037	0.13	4.26	0.021	0.053
JPYNOK	0.248	0.26	0.347	0.21	0.099	0.16	9.22	0.079	0.119
JPYNZD	0.146	0.32	0.233	0.24	0.087	0.18	7.09	0.063	0.111
JPYSEK	0.241	0.27	0.294	0.20	0.052	0.16	4.95	0.033	0.072
NOKNZD	0.449	0.22	0.413	0.27	-0.036	0.20	-2.65	-0.064	-0.011
NOKSEK	0.796	0.10	0.780	0.11	-0.016	0.08	-2.93	-0.026	-0.006
NZDSEK	0.439	0.23	0.403	0.27	-0.036	0.18	-2.89	-0.060	-0.013

Table 3
Cyclicality of realized foreign exchange (FX) correlations and FX correlation risk premiums (CRPs).
The table reports the unconditional correlation of realized correlations (RC cyclicality) and correlation risk premiums (CRP cyclicality) with four market variables: the global equity volatility measure used in Lustig, Roussanov and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin and Venter (2016) (GFI), the TED spread (TED), and the Chicago Board Options Exchange Volatility Index (VIX). Unconditional correlations are calculated using monthly data from January 1996 through December 2013 [options data for the euro (EUR) start in January 1999]. AUD = Australian dollar; CAD = Canadian dollar; CHF = Swiss franc; GBP = Pound sterling; JPY = Japanese yen; NOK = Norwegian krone; NZD = New Zealand dollar; SEK = Swedish krona.

	RC cyclicality				CRP cyclicality			
FX pair	$GVol$	GFI	TED	VIX	$GVol$	GFI	TED	VIX
AUDCAD	0.174	-0.016	-0.081	0.168	-0.090	-0.203	-0.029	-0.180
AUDCHF	-0.110	-0.342	-0.241	-0.180	0.068	0.116	0.024	0.062
AUDEUR	0.100	-0.217	-0.079	0.008	0.040	0.007	-0.076	0.060
AUDGBP	0.016	-0.207	-0.047	-0.102	0.004	0.062	-0.070	0.053
AUDJPY	-0.328	-0.488	-0.365	-0.395	0.077	0.162	0.110	0.082
AUDNOK	0.143	-0.145	-0.037	0.089	-0.096	-0.113	-0.328	-0.116
AUDNZD	0.298	-0.125	0.014	0.287	-0.107	0.036	-0.016	-0.138
AUDSEK	0.121	-0.161	-0.084	0.050	-0.141	-0.017	-0.115	-0.125
CADCHF	-0.099	-0.251	-0.223	-0.164	0.120	0.099	0.167	0.103
CADEUR	0.070	-0.133	-0.106	-0.009	-0.056	-0.014	0.076	-0.031
CADGBP	0.042	-0.060	-0.021	-0.041	0.090	-0.156	-0.150	0.066
CADJPY	-0.284	-0.405	-0.322	-0.383	0.050	0.097	0.065	0.063
CADNOK	0.102	-0.065	-0.063	0.053	-0.038	-0.151	-0.132	-0.043
CADNZD	0.166	-0.005	-0.060	0.174	0.084	-0.321	-0.182	-0.018
CADSEK	0.134	-0.025	-0.066	0.069	-0.078	-0.091	-0.187	-0.028
CHFEUR	-0.221	-0.107	-0.030	-0.250	0.330	0.122	0.178	0.308
CHFGBP	-0.159	-0.323	-0.256	-0.265	0.069	0.114	0.113	0.087
CHFJPY	-0.146	-0.063	-0.028	-0.223	0.069	0.114	0.002	0.133
CHFNOK	-0.269	-0.045	-0.130	-0.276	0.103	-0.019	0.098	0.130
CHFNZD	-0.106	-0.241	-0.256	-0.114	0.142	-0.026	-0.031	0.084
CHFSEK	-0.186	-0.221	-0.013	-0.265	0.037	-0.050	0.059	0.025
EURGBP	0.105	-0.155	-0.137	-0.018	-0.216	-0.137	-0.043	-0.184
EURJPY	-0.281	-0.178	-0.215	-0.301	0.173	0.228	0.190	0.208
EURNOK	-0.064	0.137	0.026	-0.056	-0.063	-0.062	0.032	-0.042
EURNZD	0.135	-0.106	-0.057	0.104	-0.002	-0.111	-0.205	-0.022
EURSEK	0.077	-0.169	0.077	-0.025	-0.177	-0.107	0.058	-0.186
GBPJPY	-0.353	-0.412	-0.368	-0.433	0.158	0.213	0.149	0.166
GBPNOK	0.026	-0.041	-0.118	-0.041	-0.038	-0.010	0.058	0.017
GBPNZD	0.059	-0.099	0.000	-0.007	0.001	-0.196	-0.227	0.006
GBPSEK	0.097	-0.163	-0.065	0.006	-0.211	0.013	-0.028	-0.128
JPYNOK	-0.340	-0.219	-0.303	-0.354	0.199	0.212	0.262	0.226
JPYNZD	-0.327	-0.361	-0.352	-0.317	0.064	0.077	0.129	0.008
JPYSEK	-0.343	-0.314	-0.224	-0.399	0.224	0.256	0.121	0.253
NOKNZD	0.163	-0.059	-0.028	0.161	-0.062	-0.179	-0.301	-0.101
NOKSEK	0.156	0.030	0.141	0.144	-0.086	-0.022	-0.105	-0.047
NZDSEK	0.171	-0.065	-0.054	0.144	-0.118	-0.154	-0.284	-0.154

Table 4
Cross-sectional foreign exchange (FX) cyclicality regressions.
Panel A presents the output of cross-sectional regressions of average realized FX correlations on each of the four FX correlation cyclicality measures. Panel B presents the output of cross-sectional regressions of average FX correlation risk premiums (CRPs) on each of the four FX CRP cyclicality measures. Each panel reports the regression slope coefficients, their t-statistics, their bootstrapped 95% confidence intervals, and the regression R^2s. For Panel A (Panel B) results, each FX correlation cyclicality measure (FX CRP cyclicality measure) is defined as the unconditional correlation of realized FX correlation (FX CRP) with a given market variable. The market variables are the global equity volatility measure used in Lustig, Roussanov and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin and Venter (2016) (GFI), the TED spread (TED), and the Chicago Board Options Exchange Volatility Index (VIX). The cyclicality measures are calculated using monthly data from January 1996 through December 2013 (options data for the euro start in January 1999) and are reported in Table 3. The t-statistics (in parentheses) are calculated using White (1980) standard errors.

Measure	Slope	t-statistic	2.5%	97.5%	R^2
Panel A: Average RC and RC cyclicality					
$GVol$	0.404	(2.45)	0.064	1.000	0.14
GFI	0.867	(5.14)	0.176	1.054	0.32
TED	1.151	(7.31)	0.348	1.638	0.50
VIX	0.409	(2.66)	0.148	0.892	0.15
Panel B: Average CRP and CRP cyclicality					
$GVol$	0.166	(2.66)	0.007	0.199	0.22
GFI	0.249	(9.00)	0.108	0.284	0.63
TED	0.203	(6.61)	0.073	0.263	0.48
VIX	0.201	(3.80)	0.065	0.233	0.34

Table 5
Unconditional correlation of foreign exchange (FX) correlation dispersion measures and market variables.
The table reports the correlation coefficients between the FX correlation dispersion measures FXC and FXC^{UNC} and four market variables: the global equity volatility measure used in Lustig, Roussanov and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin and Venter (2016) (GFI), the TED spread (TED), and the Chicago Board Options Exchange Volatility Index (VIX). Panel A refers to monthly data from January 1996 through December 2013; Panel B, to monthly data from January 1984 through December 2013. In both panels, bootstrap standard errors are in parentheses.

Measure	FXC^{UNC}	$GVol$	GFI	TED	VIX
Panel A: January 1996–December 2013					
FXC	0.86	0.35	0.48	0.42	0.45
	(0.02)	(0.08)	(0.06)	(0.07)	(0.07)
FXC^{UNC}		0.26	0.44	0.41	0.39
		(0.10)	(0.07)	(0.07)	(0.08)
$GVol$			0.53	0.59	0.81
			(0.08)	(0.08)	(0.04)
GFI				0.57	0.61
				(0.07)	(0.07)
TED					0.43
					(0.09)
Panel B: January 1984–December 2013					
FXC	0.89	0.22	0.32	0.26	0.21
	(0.01)	(0.06)	(0.05)	(0.05)	(0.07)
FXC^{UNC}		0.21	0.33	0.28	0.19
		(0.06)	(0.05)	(0.05)	(0.07)
$GVol$			0.12	0.41	0.79
			(0.07)	(0.08)	(0.03)
GFI				0.61	0.18
				(0.04)	(0.08)
TED					0.41
					(0.09)

Table 6
ΔFXC beta-sorted currency portfolios.
The table reports summary statistics for the excess returns of three currency portfolios sorted on exposure to ΔFXC, the innovations to the foreign exchange (FX) correlation dispersion measure FXC. Portfolio 1 (Pf1^{C}) contains the three currencies with the lowest pre-sort ΔFXC betas, and Portfolio 3 (Pf3^{C}) contains the three currencies with the highest pre-sort ΔFXC betas. HML^C denotes the portfolio that has along position in the high correlation beta currencies (Pf3^{C}) and a short position in the low correlation beta currencies (Pf1^{C}). Panel A refers to monthly data from January 1996 through December 2013; Panel B, from January 1996 through July 2007; Panel C, from January 1984 through December 2013; Panel D, from January 1984 through July 2007.

Moment	Pf1^{C}	Pf2^{C}	Pf3^{C}	HML^C
Panel A: January 1996–December 2013				
Mean	4.04	0.99	-2.38	-6.42
Standard deviation	10.26	9.11	7.86	7.83
t-statistic	1.67	0.46	-1.28	-3.47
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: January 1996–July 2007				
Mean	3.84	0.74	-3.51	-7.35
Standard deviation	7.34	8.07	7.56	6.68
t-statistic	1.78	0.31	-1.58	-3.74
Skewness	0.17	0.49	0.11	-0.01
Kurtosis	3.35	3.10	2.76	2.92
Sharpe ratio	0.52	0.09	-0.46	-1.10
Panel C: January 1984–December 2013				
Mean	4.37	1.58	0.65	-3.72
Standard deviation	9.62	9.44	8.87	8.37
t-statistic	2.48	0.92	0.40	-2.43
Skewness	-0.43	-0.24	-0.26	0.06
Kurtosis	6.09	3.73	3.96	3.71
Sharpe ratio	0.45	0.17	0.07	-0.44
Panel D: January 1984–July 2007				
Mean	4.36	1.61	0.91	-3.45
Standard deviation	8.00	9.05	9.00	8.02
t-statistic	2.64	0.87	0.49	-2.09
Skewness	0.18	-0.22	-0.28	-0.19
Kurtosis	3.81	3.79	4.04	3.13
Sharpe ratio	0.54	0.18	0.10	-0.43

Table 7
Time series correlations of foreign exchange (FX) correlations and FX correlation risk premiums (CRPs).
The table reports the time series correlations between realized FX correlations (RC) and implied FX correlations (IC) and between realized FX correlations and FX correlation risk premiums (CRPs), for all FX pairs. In addition to the correlation estimates, we report their t-statistics and 95% bootstrapped confidence intervals. FX correlation risk premiums are defined as the difference between the implied and realized FX correlations. Realized FX correlations are calculated using past daily log exchange rate changes over a three month window. Implied FX correlations are calculated from daily option prices on the underlying exchange rates. We use monthly data from January 1996 to December 2013 [options data for the euro (EUR) start in January 1999]. AUD = Australian dollar; CAD = Canadian dollar; CHF = Swiss franc; GBP = Pound sterling; JPY = Japanese yen; NOK = Norwegian krone; NZD = New Zealand dollar; SEK = Swedish krona.

	Correlation RC/IC				Correlation RC/CRP			
FX pair	Mean	t-statistic	2.5%	97.5%	Mean	t-statistic	2.5%	97.5%
AUDCAD	0.843	22.88	0.800	0.875	-0.102	-1.49	-0.243	0.046
AUDCHF	0.844	22.97	0.805	0.877	-0.695	-14.15	-0.756	-0.627
AUDEUR	0.923	32.09	0.901	0.941	-0.714	-13.63	-0.782	-0.638
AUDGBP	0.876	26.54	0.844	0.905	-0.656	-12.71	-0.732	-0.566
AUDJPY	0.892	28.89	0.855	0.922	-0.695	-14.13	-0.764	-0.610
AUDNOK	0.744	16.09	0.679	0.807	-0.213	-3.15	-0.317	-0.091
AUDNZD	0.872	26.01	0.833	0.906	-0.457	-7.52	-0.646	-0.212
AUDSEK	0.870	25.82	0.840	0.902	-0.618	-11.49	-0.723	-0.490
CADCHF	0.856	24.22	0.827	0.885	-0.684	-13.73	-0.756	-0.594
CADEUR	0.864	22.93	0.822	0.899	-0.702	-13.21	-0.785	-0.602
CADGBP	0.825	21.24	0.776	0.869	-0.518	-8.82	-0.640	-0.371
CADJPY	0.777	18.03	0.708	0.829	-0.680	-13.57	-0.737	-0.622
CADNOK	0.780	18.18	0.723	0.838	-0.316	-4.85	-0.465	-0.168
CADNZD	0.784	18.48	0.730	0.838	0.161	2.39	0.011	0.308
CADSEK	0.813	20.34	0.766	0.856	-0.137	-2.01	-0.241	-0.024
CHFEUR	0.846	21.27	0.717	0.946	-0.603	-10.12	-0.743	-0.278
CHFGBP	0.816	20.63	0.757	0.862	-0.640	-12.17	-0.715	-0.554
CHFJPY	0.835	22.19	0.788	0.874	-0.733	-15.76	-0.785	-0.665
CHFNOK	0.725	15.42	0.632	0.816	-0.671	-13.23	-0.763	-0.525
CHFNZD	0.724	15.35	0.661	0.783	-0.532	-9.19	-0.619	-0.428
CHFSEK	0.757	16.94	0.668	0.832	-0.560	-9.88	-0.683	-0.386
EURGBP	0.774	16.38	0.707	0.837	-0.592	-9.82	-0.697	-0.463
EURJPY	0.858	22.35	0.811	0.898	-0.760	-15.65	-0.813	-0.704
EURNOK	0.704	13.27	0.628	0.776	-0.632	-10.90	-0.773	-0.379
EURNZD	0.770	16.17	0.703	0.830	-0.467	-7.06	-0.597	-0.329
EURSEK	0.721	13.93	0.659	0.786	-0.549	-8.78	-0.697	-0.326
GBPJPY	0.824	21.30	0.770	0.867	-0.713	-14.87	-0.778	-0.634
GBPNOK	0.282	4.30	0.077	0.448	-0.711	-14.79	-0.767	-0.647
GBPNZD	0.812	20.32	0.773	0.852	-0.350	-5.47	-0.498	-0.199
GBPSEK	0.644	12.31	0.575	0.717	-0.615	-11.41	-0.747	-0.462
JPYNOK	0.795	19.15	0.743	0.837	-0.572	-10.21	-0.657	-0.473
JPYNZD	0.831	21.83	0.777	0.875	-0.680	-13.55	-0.746	-0.603
JPYSEK	0.825	21.34	0.775	0.865	-0.699	-14.29	-0.762	-0.627
NOKNZD	0.699	14.29	0.630	0.764	-0.157	-2.32	-0.267	-0.051
NOKSEK	0.701	14.36	0.643	0.761	-0.347	-5.42	-0.521	-0.148
NZDSEK	0.750	16.58	0.684	0.805	-0.158	-2.34	-0.253	-0.053

Table 8
Parameter values.
The table reports the calibrated parameter values used for the model simulations. All countries share the same parameter values except for γ. As regards γ, γ^0 is the parameter for the domestic country, and the values for the foreign γ^i, $i = 1, ..., 9$, are equally spaced on the interval $[\gamma^{min}, \gamma^{max}]$.

Parameter symbol	Parameter value
Panel A: Stochastic discount factor parameters	
α	0.0076
χ	19.4551
ϕ	0.06
κ	0.04
δ	40
γ^0	0.36
γ^{min}	0.20
γ^{max}	0.49
Panel B: Pricing factor parameters	
λ	0.25
$\bar{z}$	0.0077
ξ	0.0393
λ^w	0.01
$\bar{z}^w$	0.0209
ξ^w	0.0162
Panel C: Inflation parameters	
$\bar{\pi}$	-0.0039
ζ	0.25
σ	0.0037^2

Table 9
Simulated moments (benchmark model): interest rates, inflation, and exchange rates.
The table reports empirical moments and simulated moments for the model with identical local pricing factors (benchmark model). For each empirical moment, the table reports the value of the moment in the sample and the moment bootstrap standard error (in parentheses). Bootstrapping involves one thousand block bootstrap samples of 216 monthly observations each, with a block length of three observations. For each simulated moment, the table reports the point estimate and the standard error (in parentheses). The former is the moment average across one thousand simulations, and the latter is the moment standard deviation across those simulations. Panel A reports the annualized mean and standard deviation of the US real interest rate and the cross-sectional average of the mean and standard deviation of foreign real interest rates. Panel B reports the cross-sectional average of real exchange rate volatility and autocorrelation. Panel C reports the annualized mean and standard deviation of US inflation and the cross-sectional average of the mean and standard deviation of foreign inflation. Panel D reports the annualized mean and standard deviation of the US nominal interest rate and the cross-sectional average of the mean and standard deviation of foreign nominal interest rates. Panel E reports the cross-sectional average of nominal exchange rate volatility and autocorrelation.

Moment	Data	Benchmark model
Panel A: Real interest rates		
$E\left(r^{US}\right)$	0.28%	0.74%
	(0.46%)	(1.96%)
$Std\left(r^{US}\right)$	1.35%	1.08%
	(0.13%)	(0.17%)
$E_{\text{cross}}\left(E\left(r^{FGN}\right)\right)$	1.15%	0.94%
	(0.19%)	(1.85%)
$E_{\text{cross}}\left(Std\left(r^{FGN}\right)\right)$	1.19%	1.08%
	(0.03%)	(0.17%)
Panel B: Exchange rates		
$E_{\text{cross}}\left(Std\left(\Delta q_{t+1}\right)\right)$	10.82%	9.52%
	(0.59%)	(0.73%)
$E_{\text{cross}}\left(AC\left(\Delta q_{t+1}\right)\right)$	-0.01	0.00
	(0.05)	(0.04)
Panel C: Inflation		
$E\left(\pi^{US}\right)$	2.32%	1.83%
	(0.33%)	(3.86%)
$Std\left(\pi^{US}\right)$	1.27%	1.59%
	(0.14%)	(0.29%)
$E_{\text{cross}}\left(E\left(\pi^{FGN}\right)\right)$	1.56%	1.85%
	(0.17%)	(3.84%)
$E_{\text{cross}}\left(Std\left(\pi^{FGN}\right)\right)$	1.12%	1.59%
	(0.04%)	(0.28%)
Panel D: Nominal interest rates		
$E\left(r^{NOM,US}\right)$	2.60%	2.58%
	(0.25%)	(2.09%)
$Std\left(r^{NOM,US}\right)$	0.62%	1.11%
	(0.02%)	(0.20%)
$E_{\text{cross}}\left(E\left(r^{NOM,FGN}\right)\right)$	2.70%	2.77%
	(0.15%)	(2.20%)
$E_{\text{cross}}\left(Std\left(r^{NOM,FGN}\right)\right)$	0.44%	1.13%
	(0.02%)	(0.21%)
Panel E: Exchange rates		
$E_{\text{cross}}\left(Std\left(\Delta s_{t+1}\right)\right)$	10.76%	9.69%
	(0.62%)	(0.72%)
$E_{\text{cross}}\left(AC\left(\Delta s_{t+1}\right)\right)$	0.01	0.00
	(0.06)	(0.04)

Table 10
Simulated moments: foreign exchange (FX) correlations and FX correlation risk premiums (CRPs).
The table reports empirical moments and simulated moments for the model with $\rho = 1$, $\rho = 0$, and $\rho = 0.999$. All moments refer to nominal exchange rates. For each empirical moment, the table reports the value of the moment in the sample and the moment bootstrap standard error (in parentheses). Bootstrapping involves one thousand block bootstrap samples of 216 monthly observations each, with a block length of three observations. For each simulated moment, the table reports the point estimate and the standard error (in parentheses). The former is the moment average across one thousand simulations, and the latter is the moment standard deviation across those simulations. Panel A reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average realized FX correlations. Panel B reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average implied FX correlations. Panel C reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average FX CRP. Panel D reports the cross-sectional correlation between average realized and average implied FX correlation and the cross-sectional correlation between average realized FX correlation and average FX CRP. Panel E reports the cross-sectional average of the correlation between realized and implied FX correlation and the cross-sectional average of the correlation between realized FX correlation and FX CRP.

		Model		
Moment	Data	$\rho = 1$	$\rho = 0$	$\rho = 0.999$
Panel A: Realized correlation				
$2.5\%_{\text{cross}}\,(E(RC))$	0.09	0.01	0.30	0.06
	(0.03)	(0.17)	(0.04)	(0.15)
$E_{\text{cross}}\,(E(RC))$	0.45	0.39	0.40	0.40
	(0.02)	(0.04)	(0.03)	(0.04)
$97.5\%_{\text{cross}}\,(E(RC))$	0.86	0.66	0.49	0.64
	(0.01)	(0.06)	(0.03)	(0.05)
Panel B: Implied correlation				
$2.5\%_{\text{cross}}\,(E(IC))$	0.17	0.03	0.30	0.09
	(0.02)	(0.16)	(0.04)	(0.15)
$E_{\text{cross}}\,(E(IC))$	0.48	0.40	0.40	0.41
	(0.01)	(0.04)	(0.03)	(0.04)
$97.5\%_{\text{cross}}\,(E(IC))$	0.85	0.65	0.49	0.63
	(0.01)	(0.05)	(0.03)	(0.05)
Panel C: Correlation risk premiums				
$2.5\%_{\text{cross}}\,(CRP)$	-6.62%	-0.89%	0.00%	-0.71%
	(1.41%)	(0.18%)	(0.03%)	(0.16%)
$E_{\text{cross}}\,(CRP)$	1.58%	0.71%	0.04%	0.56%
	(0.57%)	(0.20%)	(0.02%)	(0.16%)
$97.5\%_{\text{cross}}\,(CRP)$	9.43%	2.75%	0.08%	2.23%
	(1.20%)	(0.55%)	(0.03%)	(0.48%)
Panel D: Cross-sectional correlations				
$corr_{\text{cross}}\,(E(RC), E(IC))$	0.98	1.00	1.00	1.00
	(0.01)	(0.00)	(0.00)	(0.00)
$corr_{\text{cross}}\,(E(RC), E(CRP))$	-0.55	-0.99	0.00	-0.99
	(0.10)	(0.01)	(0.22)	(0.00)
Panel E: Cross-sectional averages				
$E_{\text{cross}}\,(corr(RC, IC))$	0.79	1.00	1.00	1.00
	(0.02)	(0.00)	(0.00)	(0.00)
$E_{\text{cross}}\,(corr(RC, CRP))$	-0.52	-0.77	-0.02	-0.80
	(0.03)	(0.13)	(0.03)	(0.10)

Table B1
Estimates of the price of correlation risk.
The table reports the results for the estimation of the market price of correlation risk. The first part of the table reports factor betas and Newey and West (1987) standard errors (in parentheses) for the first stage regressions for various test assets. The test assets are three currency portfolios (Pf^C) sorted on exposure to the correlation risk factor ΔFXC, three currency portfolios (Pf^F) sorted on interest rate differentials, and nine individual currencies. The second part of the table reports the Fama and MacBeth (1973) factor prices and standard errors (in parentheses). Shanken (1992)–corrected standard errors are reported in brackets. We consider four sets of test assets. Set (1) contains only the three ΔFXC beta-sorted and the three interest rate-sorted portfolios from Panel A, and Set (2) also includes the nine individual G10 currencies. Set (3) contains four ΔFXC beta-sorted and four interest rate-sorted currency portfolios, using all developed country currencies. Set (4) contains four ΔFXC beta-sorted and four interest rate-sorted currency portfolios, using the full set of currencies. The first-stage beta estimates for Sets (3) and (4) are provided in the Online Appendix. We use monthly data from January 1996 through December 2013. Regression R^2s are also provided. AUD = Australian dollar; CAD = Canadian dollar; CHF = Swiss franc; GBP = Pound sterling; JPY = Japanese yen; NOK = Norwegian krone; NZD = New Zealand dollar; SEK = Swedish krona.

Factor betas				
Asset	α	DOL	HML^C	R^2
$Pf1^C$	-0.01	1.03	-0.52	0.40
	(0.07)	(0.05)	(0.03)	
$Pf2^C$	-0.02	1.11	0.00	0.10
	(0.09)	(0.06)	(0.04)	
$Pf3^C$	-0.03	1.03	0.48	-0.20
	(0.07)	(0.05)	(0.03)	
$Pf1^F$	-0.06	0.98	0.33	-0.12
	(0.10)	(0.06)	(0.06)	
$Pf2^F$	-0.03	1.03	-0.05	0.12
	(0.08)	(0.04)	(0.04)	
$Pf3^F$	0.03	1.16	-0.32	0.30
	(0.09)	(0.07)	(0.06)	
AUD	-0.09	1.20	-0.52	0.39
	(0.13)	(0.08)	(0.08)	
CAD	-0.04	0.66	-0.19	0.17
	(0.11)	(0.07)	(0.07)	
CHF	0.04	1.24	0.31	-0.05
	(0.14)	(0.08)	(0.07)	
EUR	-0.09	1.22	0.07	0.08
	(0.11)	(0.07)	(0.05)	
GBP	0.10	0.75	0.08	0.03
	(0.13)	(0.09)	(0.06)	
JPY	0.04	0.63	0.57	-0.25
	(0.22)	(0.12)	(0.10)	
NOK	0.03	1.24	0.02	0.11
	(0.13)	(0.09)	(0.08)	
NZD	0.06	1.27	-0.39	0.32
	(0.15)	(0.08)	(0.11)	
SEK	-0.10	1.29	-0.05	0.14
	(0.11)	(0.07)	(0.06)	

Factor prices			
Set	λ^{DOL}	λ^{HML^C}	R^2
Set (1)	0.09	-0.58	0.99
	(0.15)	(0.15)	
	[0.15]	[0.15]	
Set (2)	0.09	-0.54	0.93
	(0.15)	(0.20)	
	[0.15]	[0.20]	
Set (3)	0.13	-0.51	0.90
	(0.15)	(0.17)	
	[0.15]	[0.18]	
Set (4)	0.15	-0.67	0.81
	(0.14)	(0.22)	
	[0.14]	[0.23]	

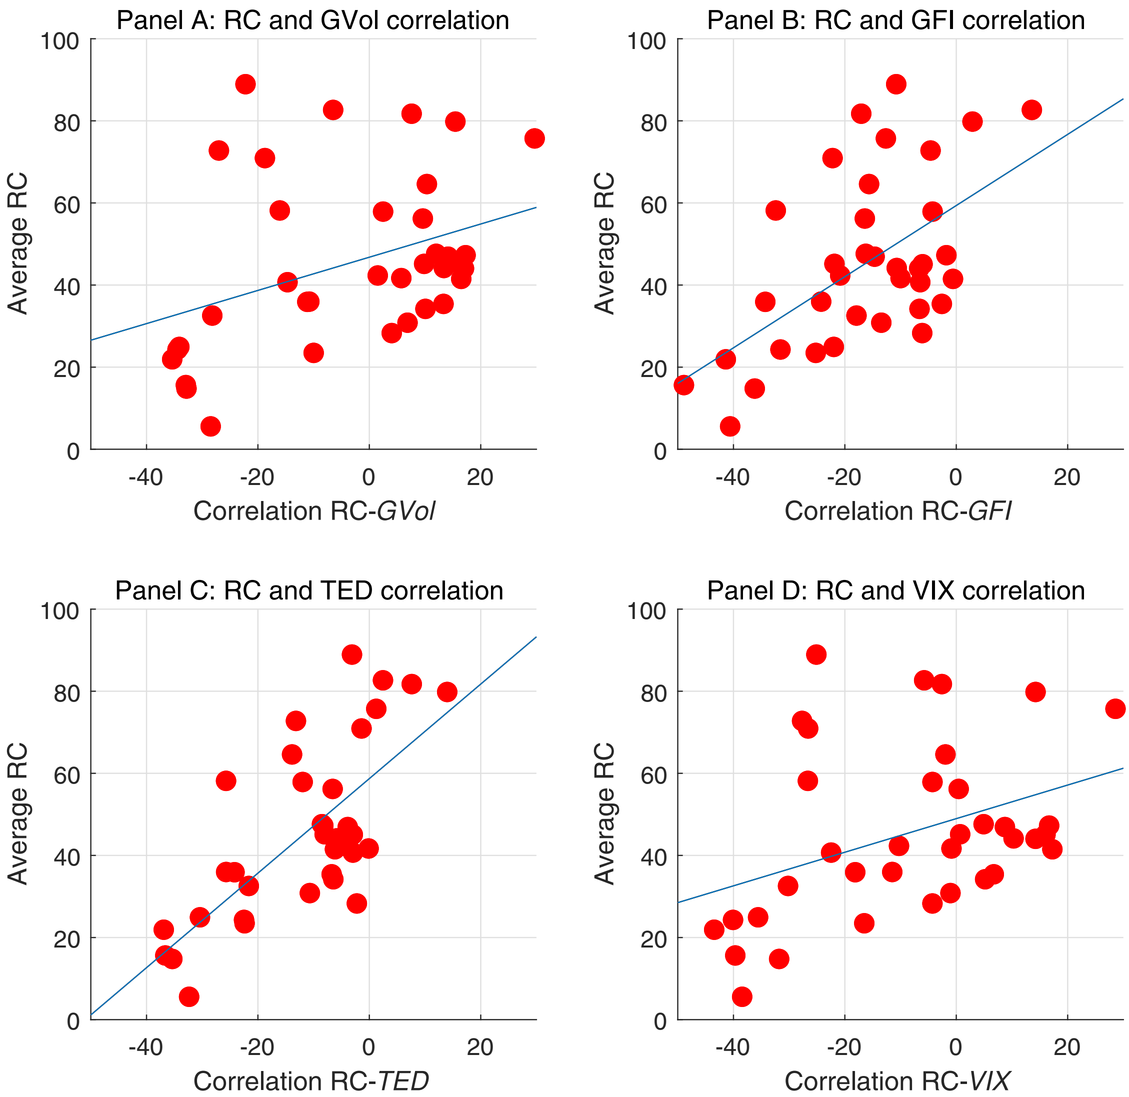

Fig. 1. Average realized foreign exchange (FX) correlations and FX correlation cyclicality. The figure illustrates the association between average realized FX correlations and measures FX correlation cyclicality. For each FX pair, FX correlation cyclicality is measured by the unconditional correlation between the realized FX correlation of the pair and a market variable that acts as a business cycle proxy. The market variables considered are the global equity volatility measure from Lustig, Roussanov and Verdelhan (2011) ($GVol$, Panel A), the global funding illiquidity measure (GFI, Panel B) from Malkhozov, Mueller, Vedolin and Venter (2016), the TED spread (TED, Panel C), and the Chicago Board Options Exchange Volatility Index (VIX, Panel D). We use monthly data from January 1996 to December 2013. In each panel, the line of best fit is also shown. RC = realized correlation.

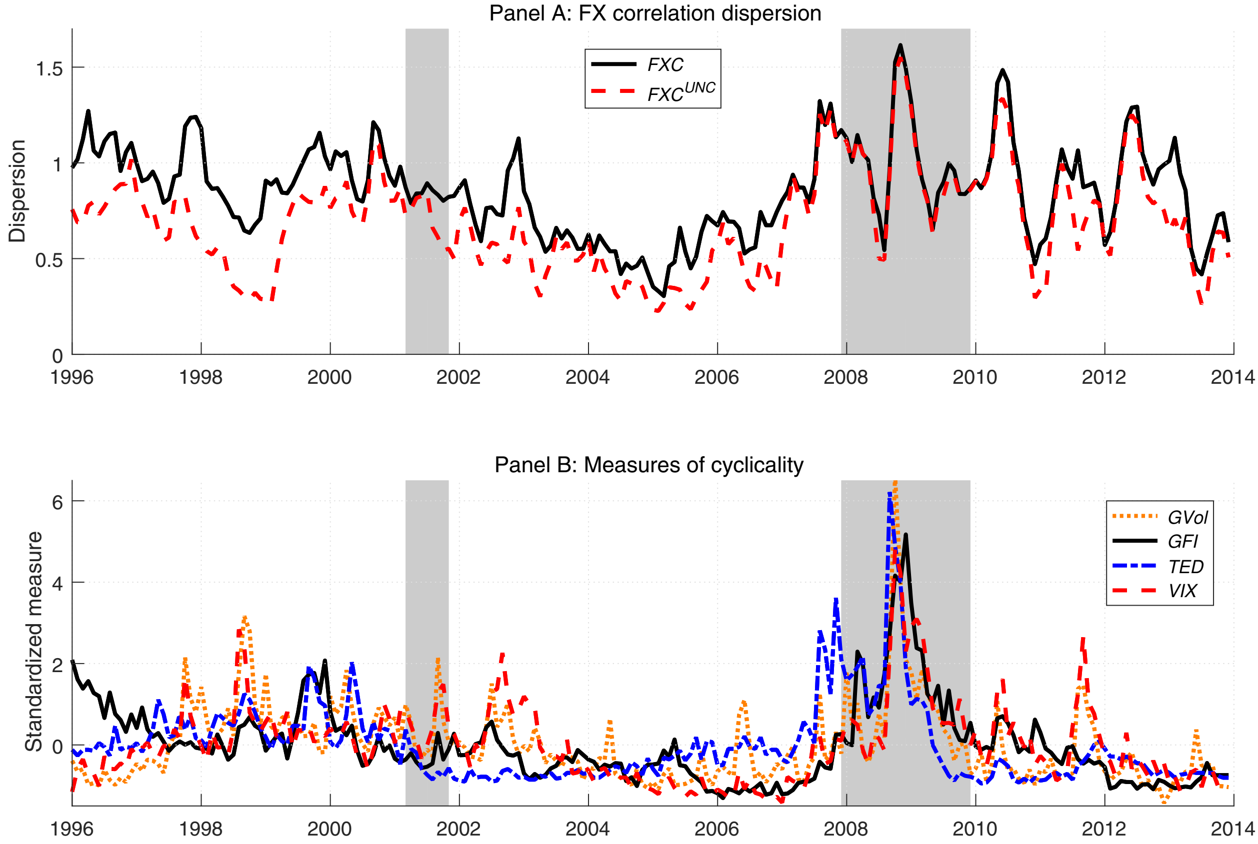

Fig. 2. Foreign exchange (FX) correlation dispersion measures and market variables. Panel A plots the time series of the two FX correlation dispersion measures, FXC and FXC^{UNC}, from January 1996 to December 2013. FXC (solid line) is calculated as the difference between the average FX correlation of high- and low-correlation FX pairs. The two groups of FX pairs consist of the highest and lowest deciles of realized FX correlations across all 36 FX pairs in our sample, with the deciles being rebalanced every month. FXC^{UNC} (dashed line) is calculated as the difference in average correlations between the decile of high average correlation FX pairs and the decile of low average correlation FX pairs. Panel B plots the time series of the global equity volatility measure used in Lustig, Roussanov and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin and Venter (2016) (GFI), the TED spread (TED), and the Chicago Board Options Exchange Volatility Index (VIX), from January 1996 to December 2013. All series in Panel B are standardized to have zero mean and a standard deviation of one. In both panels, the shaded areas correspond to National Bureau of Economic Research recessions.

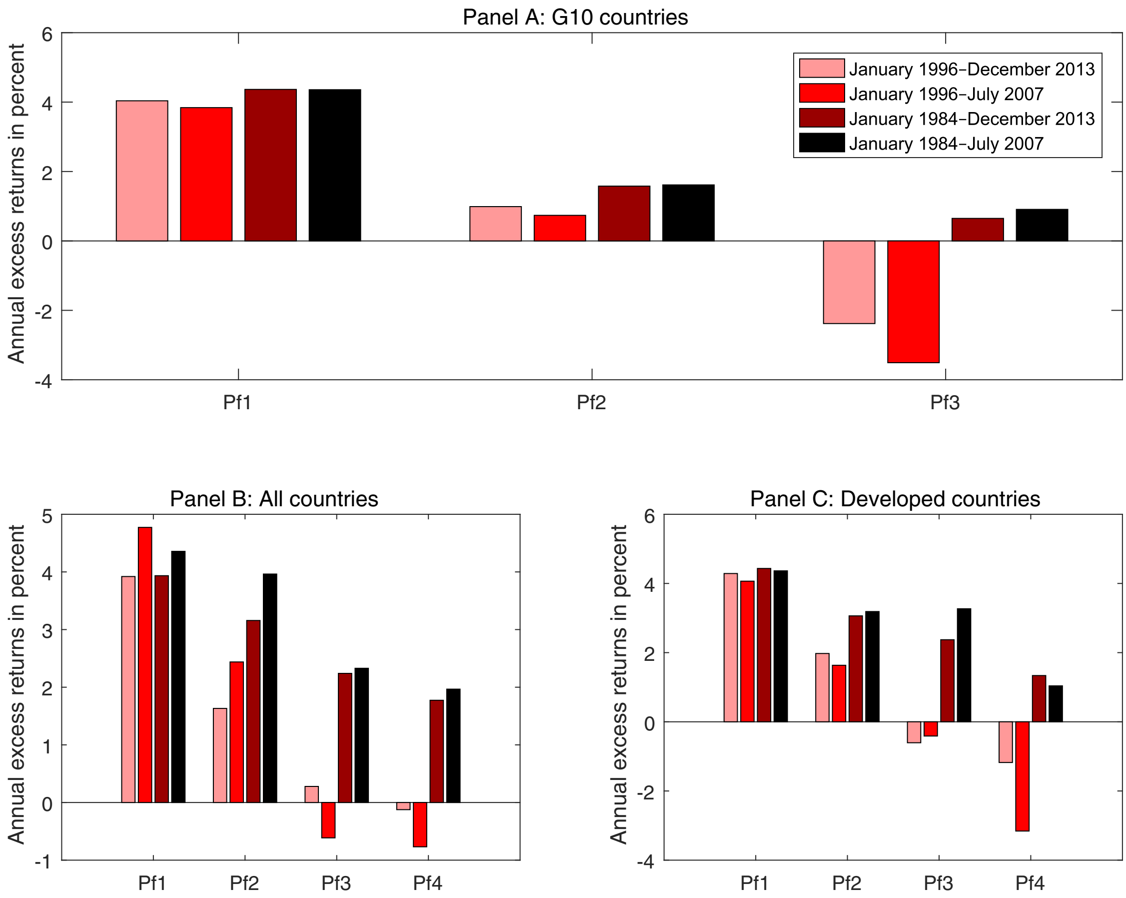

Fig. 3. Currency portfolios sorted on exposure to the foreign exchange (FX) correlation factor ΔFXC. The figure displays annualized average excess returns of currency portfolios, for different currency and period samples. Currencies are sorted into portfolios at time t based on their exposure to ΔFXC at the end of period $t-1$. Exposure is measured by regressing currency excess returns on the FX correlation risk factor ΔFXC over the preceding 36 months. Panel A presents the portfolio excess returns for the Group of Ten (G10) set of currencies (three ΔFXC beta-sorted currency portfolios). Panels B and C present the portfolio excess returns for the currencies in the developed country set and in the full country set, respectively (four ΔFXC beta-sorted currency portfolios for each set). In each panel, Portfolio 1 (Pf1) contains the currencies with the lowest pre-sort ΔFXC betas, and Portfolio 3 or 4 (Pf3 or Pf4), depending on the set of currencies, contains the currencies with the highest pre-sort ΔFXC betas. In each panel, average annualized portfolio excess returns are reported for four sample periods: January 1996–December 2013, January 1996–July 2007, January 1984–December 2013, and January 1984–July 2007.

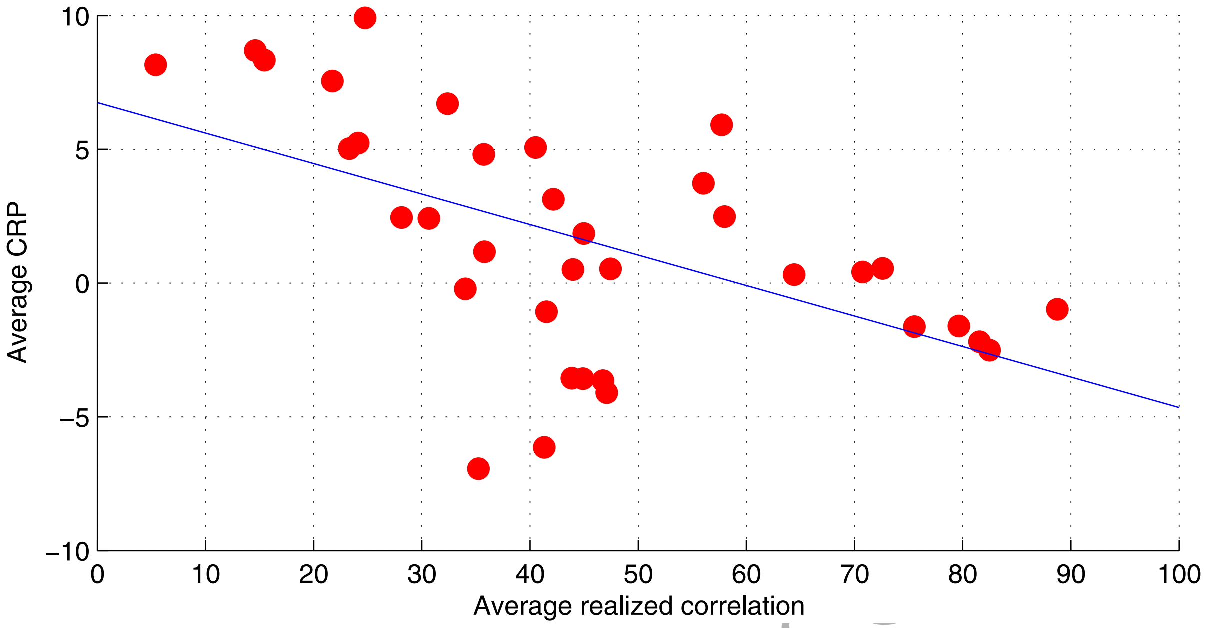

Fig. 4. Average realized foreign exchange (FX) correlations and average FX correlation risk premiums (CRPs). The figure plots the average FX correlation risk premiums for all 36 exchange rate pairs in our sample against the corresponding average realized FX correlations. Average FX correlation risk premiums and average realized FX correlations are expressed in percentage points. We use monthly data from January 1996 to December 2013 (options data for the euro start in January 1999). The line of best fit is also shown.

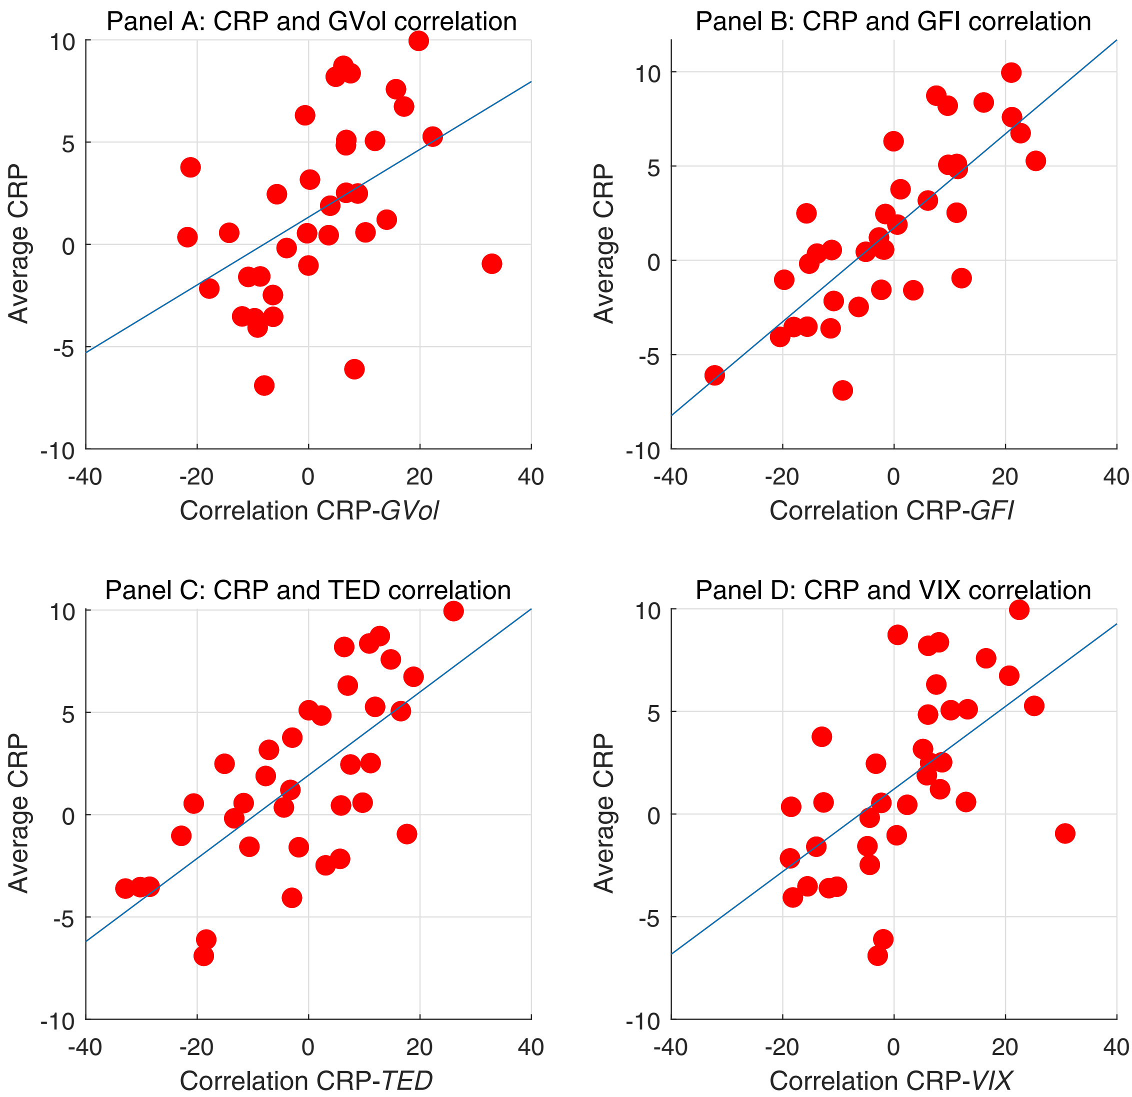

Fig. 5. Average foreign exchange (FX) correlation risk premiums (CRPs) and FX CRP cyclicality. The figure illustrates the association between average FX correlation risk premiums and measures FX correlation risk premium cyclicality. For each FX pair, FX correlation risk premium cyclicality is measured by the unconditional correlation between the FX correlation risk premium of the pair and a market variable that acts as a business cycle proxy. The market variables considered are the global equity volatility measure from Lustig, Roussanov and Verdelhan (2011) ($GVol$, Panel A), the global funding illiquidity measure (GFI, Panel B) from Malkhozov, Mueller, Vedolin and Venter (2016), the TED spread (TED, Panel C), and the Chicago Board Options Exchange Volatility Index (VIX, Panel D). We use monthly data from January 1996 to December 2013 (options data for the euro start in January 1999). In each panel, the line of best fit is also shown.

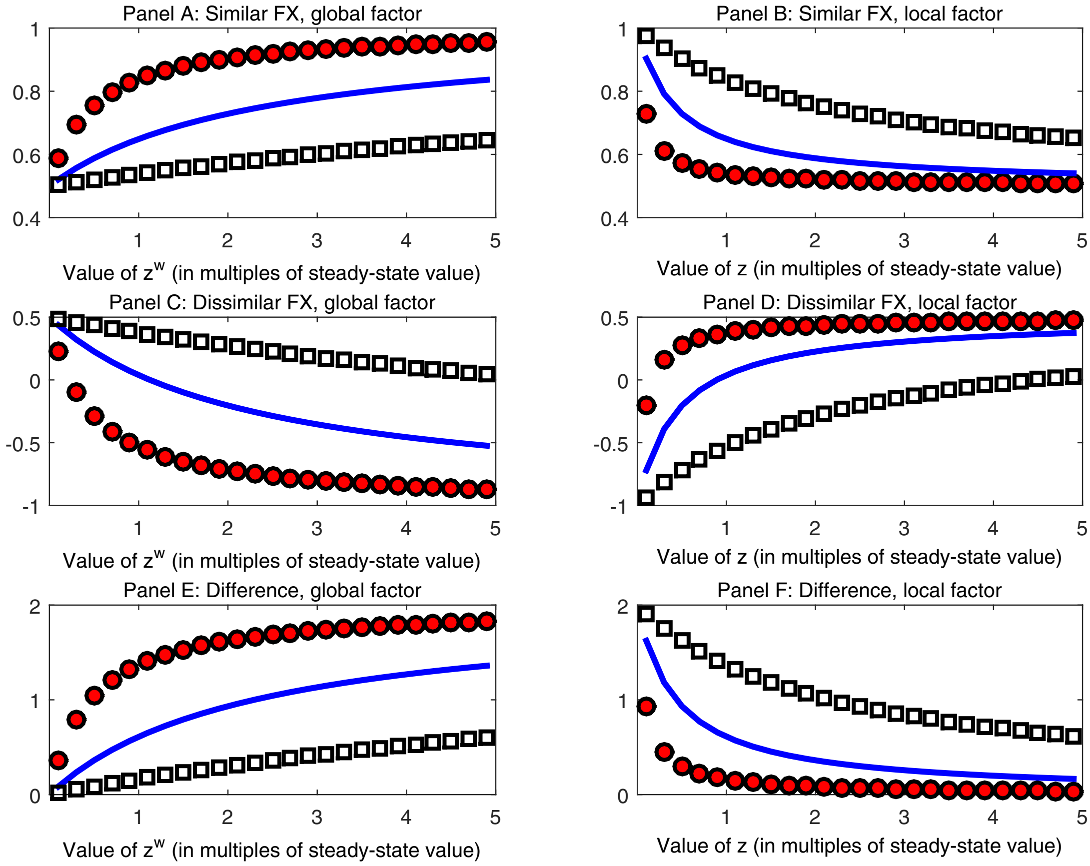

Fig. 6. Model-implied foreign exchange (FX) correlations. The figure displays the properties of conditional real FX correlation in the model with identical local pricing factors. Panels A, C, and E plot the conditional FX correlation as a function of the global pricing factor z^w, holding the local pricing 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 the difference in conditional FX correlation between the two pairs. In each panel, the circles, solid line, and squares plot the conditional FX correlation, assuming that the local pricing factor z is equal to 0.2, 1, and 5 times its steady state value $\bar{z}$, respectively. Panels B, D, and F plot the 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 the difference in conditional FX correlation between the two pairs. In each panel, the circles, solid line, and squares plot the conditional FX correlation assuming that the global pricing factor z^w is equal to 0.2, 1, and 5 times its steady state value $\bar{z}^w$, respectively. To plot the figures, we set the model parameters equal to their calibrated values in Table 8. To ensure symmetry, we set the values of the country exposures to global FX risk such that the condition $D^{1,2} = -D^{1,3} > 0$ is satisfied. We achieve that by imposing $\gamma^1 = \gamma^{min}$ and $\gamma^3 = \gamma^{max}$, and setting γ^2 so that the symmetry condition holds.

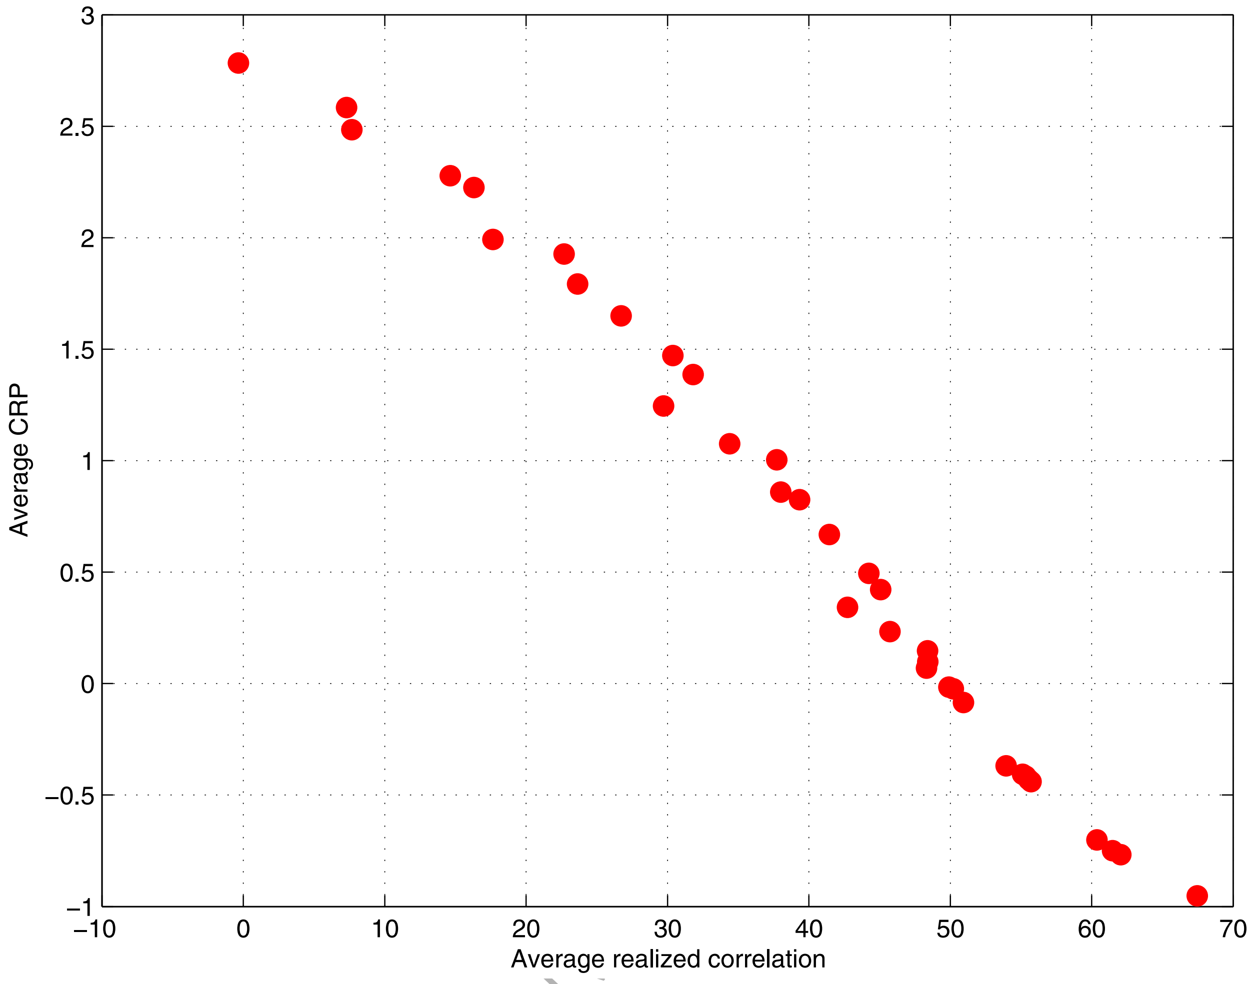

Fig. 7. Model-implied average realized foreign exchange (FX) correlations and average FX correlation risk premiums (CRPs). The figure plots the average FX correlation risk premiums for all 36 exchange rate pairs against the corresponding average realized FX correlations using simulated data for the model with identical local pricing factors ($\rho = 1$). The parameter values are reported in Table 8, and the simulation details can be found in Appendix E. Average FX correlation risk premiums and average realized FX correlations are expressed in percentage points.

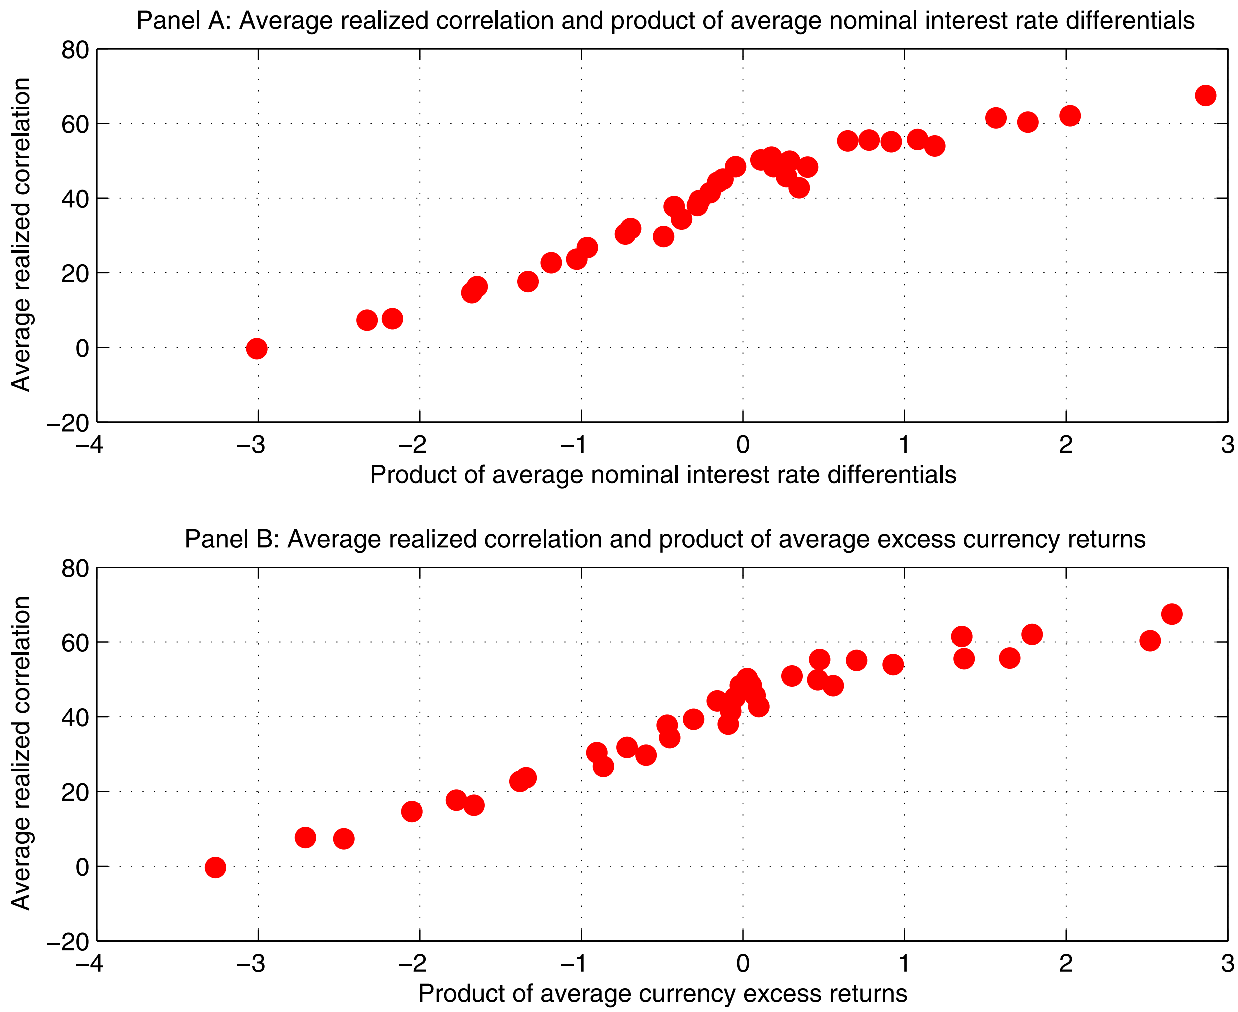

Fig. 8. Model-implied average realized foreign exchange (FX) correlations and products of average nominal interest rate differentials and average currency excess returns. The figure plots the average realized FX correlations for all 36 exchange rate pairs against the corresponding product of average nominal interest rate differentials (Panel A) or the product of average currency excess returns (Panel B) for the model with identical local pricing factors ($\rho = 1$). The parameter values are reported in Table 8, and the simulation details can be found in Appendix E. Average realized FX correlations are expressed in percentage points. Products of nominal interest rate differentials and currency excess returns are expressed in squared percentage points. Nominal interest rate differentials and currency excess returns are annualized.

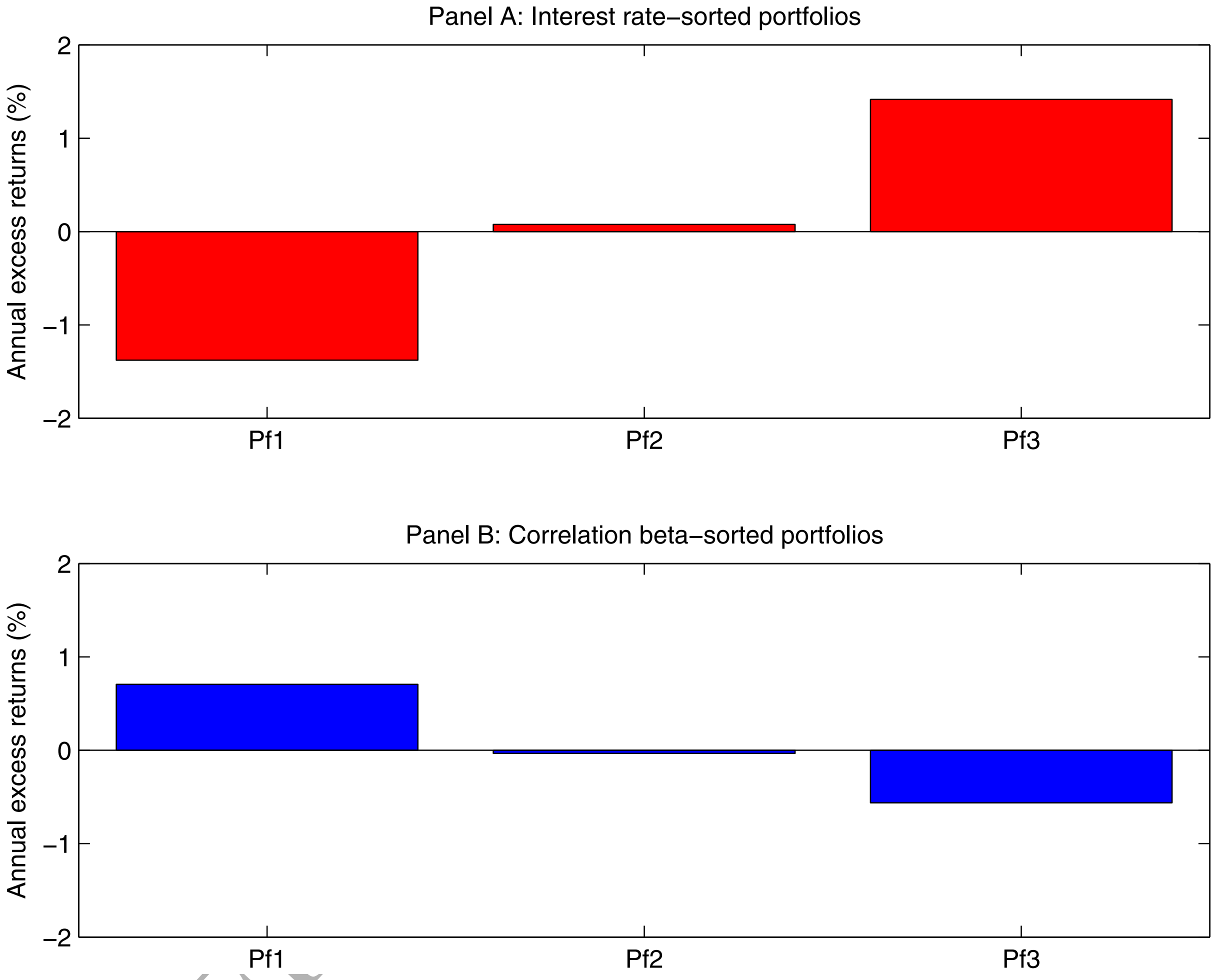

Fig. 9. Model-implied currency portfolio excess returns. The figure displays average annualized portfolio excess returns for interest rate-sorted (Panel A) and ΔFXC beta-sorted (Panel B) currency portfolios using simulated data for the model with identical local pricing factors ($\rho = 1$). For Panel A, currencies are sorted into portfolios according to their nominal interest rate, with monthly rebalancing. Portfolio 1 (Pf1) contains low interest rate currencies, and Portfolio 3(Pf3) contains high interest rate currencies. For Panel B, currencies are sorted into portfolios on their exposure to ΔFXC at the end of period $t - 1$, with monthly rebalancing. Exposure is measured by regressing currency excess returns on the correlation risk factor ΔFXC over the preceding 36 months. Portfolio 1 (Pf1) contains the currencies with the lowest pre-sort ΔFXC betas, and Portfolio 3 (Pf3) contains the currencies with the highest pre-sort ΔFXC betas. The parameter values are reported in Table 8, and the simulation details can be found in Appendix E.

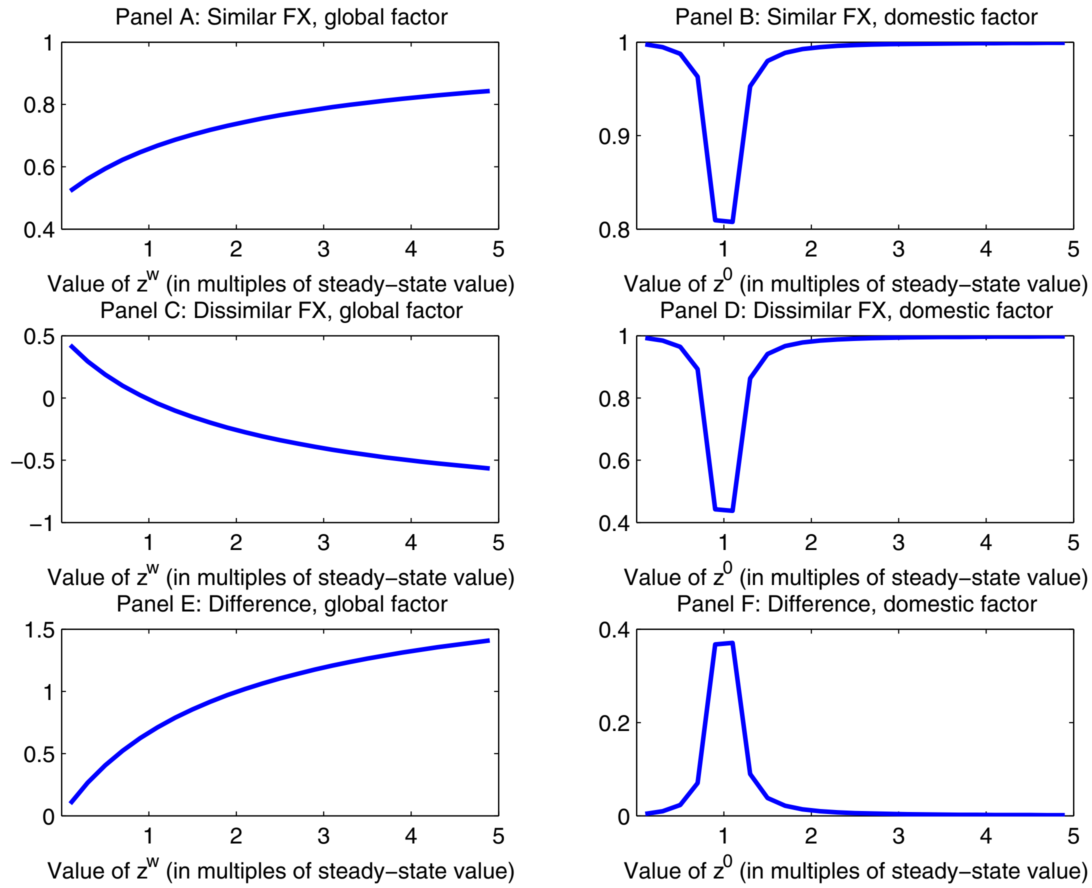

Fig. 10. Model-implied foreign exchange (FX) correlations: independent local pricing factors. The figure displays the properties of conditional real FX correlation in the model with independent local pricing factors ($\rho = 0$). Panels A, C, and E plot the conditional FX correlation as a function of the global pricing factor z^w, holding all the local pricing factors constant at their common steady state level $\bar{z}$. 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 the difference in conditional FX correlation between the two pairs. Panels B, D, and F plot the conditional FX correlation as a function of the domestic pricing factor z^0, holding the global pricing factor z^w constant at its steady state level $\bar{z}^w$ and all the foreign local pricing factors constant at their common steady state value $\bar{z}$. 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 the difference in conditional FX correlation between the two pairs. To plot the figures, we set the model parameters equal to their calibrated values in Table 8. To ensure symmetry, we set the values of the country exposures to global FX risk such that the condition $D^{1,2} = -D^{1,3} > 0$ is satisfied. We achieve that by imposing $\gamma^1 = \gamma^{min}$ and $\gamma^3 = \gamma^{max}$, and setting γ^2 so that the symmetry condition holds.

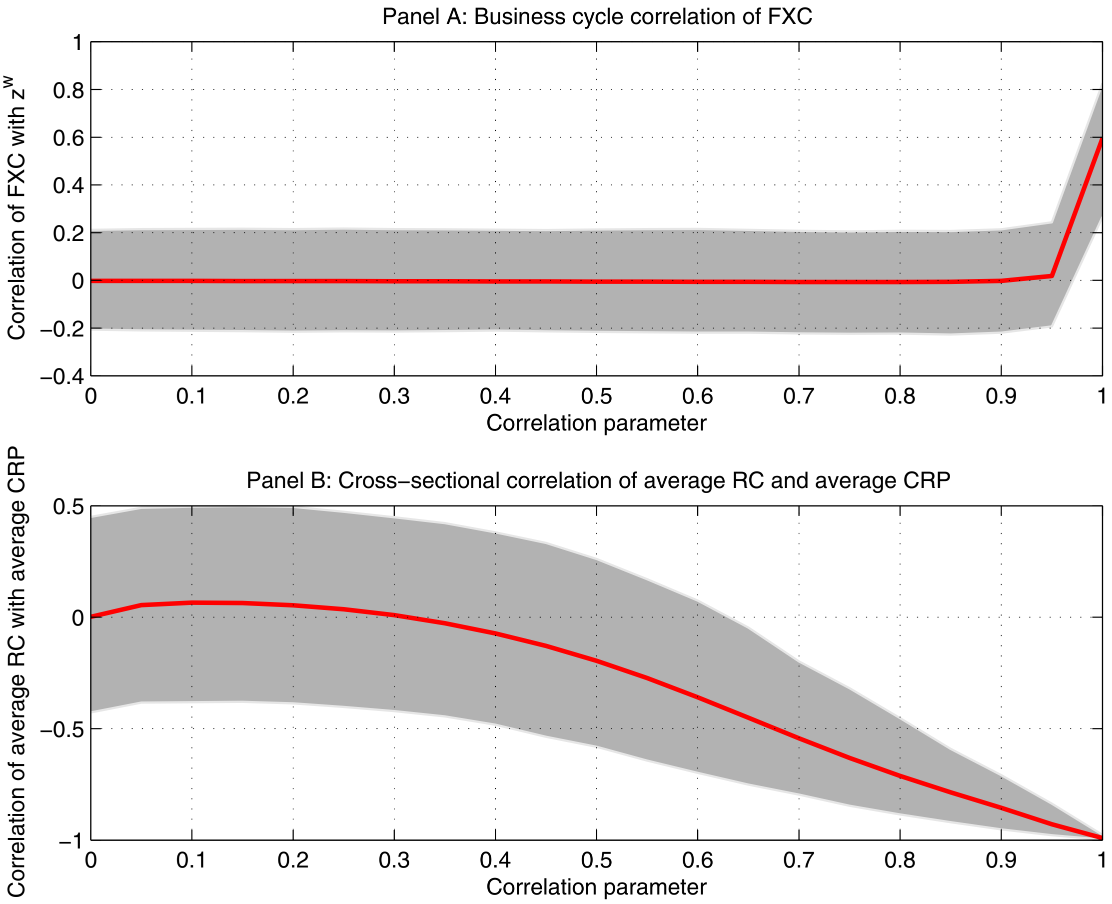

Fig. 11. Model-implied correlations as function of parameter ρ. The figure presents the point estimates (solid line) and the 95% confidence intervals (shaded area) of correlations of interest in simulated data for different values of the correlation parameter ρ. A value of $\rho = 0$ corresponds to the model with independent local pricing factors, and a value of $\rho = 1$ corresponds to the benchmark model with identical local pricing factors. We consider 21 values of ρ, ranging from $\rho = 0$ to $\rho = 1$, in increments of 0.05. Panel A presents the correlation between FXC, the measure of cross-sectional dispersion in conditional FX correlation, and z^w, the global pricing factor. Panel B presents the cross-sectional correlation between average FX correlations and average FX correlation risk premiums (CRPs) across FX pairs. With the exception of parameter ρ, the parameter values are reported in Table 8. The simulation details can be found in Appendix E. RC = realized correlation.