International Correlation Risk*

Philippe Mueller
London School of Economics[†]

Andreas Stathopoulos
University of Washington[‡]

Andrea Vedolin
London School of Economics[§]

Abstract

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

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

First Version: December 2011

This Version: December 2015

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

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

[‡]Foster School of Business, 4277 NE Stevens Way, Seattle, WA 98195, USA, Email: astath@uw.edu

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

The importance of correlation in financial markets has spawned a large literature documenting the properties of asset return correlations. Correlation risk is priced in equity markets, arguably due to the deterioration of investors' investment opportunities that results from a reduction in diversification benefits when asset return correlations increase. Yet, the existing literature has largely ignored the foreign exchange (FX) market. In this paper, we explore the empirical properties of conditional FX correlations both under the physical and under the risk-neutral measure and propose a reduced-form no-arbitrage model that is consistent with our empirical findings and illustrates the effects of spanned and unspanned currency risk on FX correlations.

We start by documenting the empirical properties of conditional FX correlations. We consider exchange rates against the U.S. dollar (USD) and show that there exists substantial cross-sectional heterogeneity in the average conditional correlation of FX pairs, suggesting the existence of ex ante heterogeneity in exchange rates. Furthermore, using several business cycle proxies, we find that FX pairs with high average correlation become more correlated in adverse economic times, whereas FX pairs with low average correlation become even less correlated in those states. As a result, the cross-sectional dispersion of FX correlations widens in bad states and tightens in good states of the world. Consider, for example, the exchange rates of three currencies against the USD: the Japanese yen (JPY), a low interest rate currency, and the Australian and the New Zealand dollar (AUD and NZD), two high interest rate currencies. The average correlation between the JPY and either the AUD or the NZD exchange rate is fairly low (0.16 and 0.15, respectively) and procyclical. On the other hand, the average correlation between the exchange rates of the two high interest rate currencies is 0.76 and countercyclical.

We exploit the cyclical properties of FX correlation by defining an FX correlation dispersion measure, FXC, using the conditional correlations of the G10 exchange rates. To construct our measure, we sort FX pairs into deciles based on their conditional FX correlation and subtract the average conditional FX correlation of the bottom decile from the average conditional FX correlation of the top decile. We then verify that the resulting dispersion measure is strongly countercyclical, being positively correlated with the market variables that are associated with bad states. Finally, we sort currencies into portfolios and find that currencies with low FXC betas have high average excess returns, whereas currencies with high FXC betas yield low excess returns. For our benchmark sample of G10 currencies, the return to HML^C, a currency portfolio that goes short the high FXC beta currencies and invests in the low FXC beta currencies, generates a highly significant average annual excess return of 6.4% with a Sharpe ratio of 0.82. We then estimate the price of FX correlation risk using a two-factor model that includes the dollar factor from Lustig, Roussanov, and Verdelhan (2011) and the return to HML^C, our traded FX correlation factor. Using different sets of test assets and estimation periods, we find that our estimates of the price of FX correlation risk range from -51 to -67 basis points (bps) per month.

We conclude our empirical investigation by using currency option prices in order to extract conditional FX correlation dynamics under the risk-neutral measure. We calculate FX correlation risk premia, defined as the difference between conditional FX correlations under the risk-neutral and the physical measures, and we find a strongly negative cross-sectional relationship between average FX correlations and average correlation risk premia: FX pairs characterized by low average correlations tend to exhibit high correlation risk premia (i.e., they are on average more correlated under the risk-neutral measure than under the physical measure), whereas FX pairs that are highly correlated on average have low correlation risk premia. Thus, the cross-sectional dispersion of FX correlations is on average lower under the risk-neutral measure than under the physical measure. We also show that there is a very strong negative time series relationship between FX correlations and FX correlation risk premia for almost all FX pairs. As regards cyclicality, FX pairs with high average correlation risk premia have countercyclical correlation risk premia, whereas pairs with low correlation risk premia have procyclical premia. Thus, bad states amplify the magnitude of FX correlation risk premia, 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 physical and risk-neutral measure FX correlation dynamics. Under the physical measure, the negative association between FXC betas and currency returns suggests that U.S. investors require an FX risk premium for being exposed to states in which the cross section of FX correlations *widens*. However, FX options data are priced in a way that suggests that U.S. investors worry about states in which the cross section of FX correlations *tightens*. The key to addressing this apparent conundrum is that, in our model, FX correlation risk is not spanned by exchange rates: the marginal utility of U.S. investors is exposed to shocks that affect conditional FX correlations, but not exchange rates themselves.

In the model, each country's SDF is exposed to two global shocks, as well as a single country-specific shock. Importantly, countries have heterogeneous loadings on the first global shock, but identical loadings on the second global shock. This implies that exchange rates are exposed only to the first global shock, as the second global shock cancels out and does not affect exchange rates at all. As a result, the steady-state cross-sectional distribution of conditional FX correlations is determined by the cross section of exposures to the first global shock: on average, FX pairs that correspond to foreign countries with similar exposure to the first global shock (called similar FX pairs) are more correlated 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 across 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 returns.

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: 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. Both shocks are priced, but not symmetrically: U.S. investors price the second shock more severely than the first, so they attach a high price to states characterized by high relative values of the local pricing factor. Since those are exactly the states in which the cross-sectional dispersion of FX correlation tightens, our model is able to match the cross sectional properties of average correlation risk premia.

On the other hand, only the first global shock is priced in currency markets, as this is the *only* global shock to which exchange rates are exposed: exchange rates do not span FX correlation risk, as they are unaffected by the second global shock. This lack of spanning allows our model to generate a negative relationship between FXC betas and currency returns: investing in exchange rates requires compensation for the only global shock that exchange rates are exposed to, the first global shock. Since negative realizations of that shock lead to a widening of the cross section of FX correlations, investors require high returns for negative FXC beta currencies, in line with our cross-sectional empirical findings.

A simulated version of our model generates realized FX correlations, implied FX correlations and FX correlation risk premia 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.

Related literature: This paper is part of the literature addressing the risk–return relationship in 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. Importantly, they 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 factor is crucial for explaining the behavior of FX correlations under both the physical and the risk-neutral measure.

Recent international finance models that address the cross section of currency risk premia by assuming ex ante heterogeneity across countries include Hassan (2013), Martin (2013), Tran (2013), Backus, Gavazzoni, Telmer, and Zin (2013), Ready, Roussanov, and Ward (2013), and Colacito, Croce, Gavazzoni, and Ready (2014), the latter of which extends Colacito and Croce (2013). In all models, high (low) interest rate currencies are risky (hedges) because they depreciate (appreciate) in bad global states; this is because high interest rate countries are those with low exposure to global risk: small countries, countries with smooth non-traded output, countries with more procyclical monetary policy, commodity producers, or countries with less exposure to global long-run endowment shocks, depending on the model. Gabaix and Maggiori (2015) take a different approach and consider a model of imperfect financial markets in which the FX carry trade constitutes compensation for exposure to shocks in the risk-bearing capacity of financiers.

In recent empirical work, Lustig and Verdelhan (2007), Menkhoff, Sarno, Schmeling, and Schrimpf (2012), Mancini, Ranaldo, and Wrampelmeyer (2013), Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan (2015), Lettau, Maggiori, and Weber (2014) and Dobrynskaya (2014) argue that the carry trade can be explained as compensation for exposure to consumption risk, global FX volatility risk, FX liquidity risk, disaster risk, or downside market risk, respectively. Cenedese, Sarno, and Tsiakas (2014) find that a high cross-sectional average of currency excess return variance predicts carry trade losses.

The rest of the paper is organized as follows. Section 1 describes the data. Section 2 contains our empirical findings regarding the cross section and time series properties of FX correlations, as well as the pricing of correlation risk in currency markets. The stylized facts concerning the FX correlation risk premia are presented in Section 3. Section 4 presents our no-arbitrage model, and Section 5 concludes. The Appendix contains details on the construction of the realized and implied correlation measures, robustness checks and details on the model. Additional results and robustness checks are deferred to an Online Appendix.

1 Data

We start by describing the data. We calculate physical measure (realized) FX correlations using daily spot exchange rates. To construct measures of implied FX correlations we use daily FX options data. Our benchmark sample period starts in January 1996 and ends in December 2013, and is dictated by the availability of the options data.

Spot and forward rates: To calculate physical measure FX moments, we use daily spot exchange rates from WM/Reuters obtained through Datastream. We also collect one-month forward rates from WM/Reuters in order to calculate forward discounts. The spot and forward rates are fixed at 4 p.m. UK time, which is standard in the FX market.

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.[1] We use the U.S. dollar as the base currency, so superscript i always denotes the foreign currency. Our benchmark sample comprises the nine G10 foreign currencies (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD, SEK) from January 1996 to December 2013. Before the introduction of the EUR in 1999 we use the German Mark (DEM) in its place. In the Appendix, we present results for an extended sample of developed and emerging market currencies. Monthly log excess returns from holding the foreign currency i are computed as $rx_{t+1}^i = f_t^i - s_{t+1}^i$.

Table 1 presents the properties of the G10 exchange rates during the benchmark sample. In line with the literature on the FX carry trade, we find that currencies with high nominal interest rates tend to achieve higher average dollar excess returns: the NZD, AUD and NOK are characterized by high nominal interest rates, as well as

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

high average currency excess returns, while the reverse is true for the JPY, CHF and EUR/DEM.

[Insert Table 1 here.]

For robustness checks, we extend the cross section of currencies to a set of developed and emerging market currencies. The developed country sample, apart from the G10 currencies, includes the currencies of Austria, Belgium, Denmark, Finland, France, Greece, Italy, Ireland, Netherlands, Portugal and Spain. The full sample includes all the developed country currencies, along with the currencies of Czech Republic, Hungary, India, Indonesia, Kuwait, Malaysia, Mexico, Philippines, Poland, Singapore, South Africa, South Korea, Taiwan and Thailand.[2]

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

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

2 Exchange rate correlations

In this section, we first document a negative relationship between the unconditional average and the cyclicality of conditional FX correlations and, based on this observation, we construct a novel FX correlation risk factor, FXC, which reflects the cross-sectional dispersion of FX correlations. Sorting currencies into portfolios based on exposure to the FXC factor reveals a significant return spread between the high and low correlation risk portfolio. Consistent with this finding, we show that FX correlation risk has a negative price in currency markets.

2.1 Properties of exchange rate correlations

We use daily spot exchange rates to calculate conditional FX correlations under the physical measure. In particular, we proxy the conditional 1-month FX correlation of each FX pair at time t with its corresponding realized correlation over a rolling 3-month window of past daily observations. Appendix A provides the details.

The first two columns of Table 2 present the time-series mean and standard deviation of the conditional FX correlation of each of the 36 G10 FX pairs. The average conditional correlation is positive for all 36 FX pairs, indicating all pairs of dollar exchange rates exhibit a positive comovement on average. The cross-sectional average of the conditional correlation means is 0.454 but there is substantial cross-sectional variation in the average conditional FX correlation: the averages range from almost zero (CAD/JPY with 0.054, 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 0.888).[3] This variation suggests considerable ex ante heterogeneity across exchange rate pairs which is manifested as fixed effects in average FX correlations. Furthermore, conditional FX correlations exhibit non-trivial variability across time: the cross-sectional average of the standard deviation of conditional FX correlations is 0.23, ranging from 0.09 (EUR/NOK

[3]Beginning September 2011, the Swiss National Bank imposed a cap of 1.2 CHF to the 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.

pair) to 0.34 (AUD/JPY pair), suggesting non-trivial swings in the degree of exchange rate comovement across time for all FX pairs.

[Insert Table 2 here.]

Given the time variation in conditional FX correlations, we then explore whether that time variation is cyclical and, if so, whether there is cross-sectional heterogeneity in its properties. To that end, we consider the comovement of conditional FX correlations with market variables that are well-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 VIX (VIX). $GVol$ is constructed as in Lustig, Roussanov, and Verdelhan (2011). GFI is constructed based on the method proposed by Hu, Pan, and Wang (2013) but calculated using an international sample of government bond securities as in Malkhozov, Mueller, Vedolin, and Venter (2015). TED is from FRED and is the spread between the three month USD LIBOR and the three month Treasury Bill rate. VIX is backed out from options on the S&P 500 stock index and available from the CBOE. TED and VIX are U.S. specific measures but are often used as global indicators. $GVol$ and GFI are calculated using international data in local currencies. 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 here.]

As seen in the table, we find substantial cross-sectional heterogeneity regarding the cyclicality properties of conditional FX correlations. Consider for example the FX pairs GBP/JPY and NOK/SEK: the former is a pair with low average conditional correlation (0.22), while the latter is an FX pair with high conditional correlation on average (0.80).

Interestingly, the cyclicality measures of the GBP/JPY pair are all negative, ranging from -0.43 (VIX) to -0.35 ($GVol$), indicating lower FX comovement during periods when the market variables are elevated, which are typically considered bad states of the world. On the other hand, the cyclicality measures of the NOK/SEK pair are all positive, ranging from 0.03 (GFI) to 0.16 ($GVol$), suggesting higher comovement of the two exchange rates in bad states.

To determine whether there is a cross-sectional pattern in the cyclicality properties of FX correlation, we plot each cyclicality measure of the 36 FX pairs against their average conditional correlation; Panels A to D in Figure 1 present the plots for the four cyclicality measures. Each panel also presents the line of best fit from a cross-sectional regression of average conditional FX correlations on the FX cyclicality measure featured in the panel. We report the details of the four cross-sectional regressions in Panel A of Table 4: for each regression, we document the point estimate of the slope coefficient, its asymptotic t-statistic, and the 95% bootstrapped confidence interval (2.5 and 97.5 bootstrap percentiles). The asymptotic t-statistic is calculated using White (1980) standard errors that adjust for cross-sectional heteroskedasticity. The bootstrapped confidence interval allows to adjust for potential small sample biases. 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 relationship between average conditional FX correlation and FX correlation cyclicality. Indeed, we can see that the FX pairs with the highest average realized correlations exhibit either non-cyclical or slightly countercyclical correlations. In the other extreme, the FX pairs with the lowest realized correlations are characterized by strongly procyclical FX correlations.

[Insert Figure 1 and Table 4 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. Thus, the difference in conditional correlations between high correlation FX pairs and low correlation FX pairs is also countercyclical, increasing during crises and declining during booms. This can easily be seen by constructing an FX correlation dispersion measure as follows: each period t, we sort all FX pairs according to their conditional correlation, defined as the realized correlation over the past three months. We then calculate the average conditional correlation for the top and bottom decile (which consists of four pairs each) and take the difference of the two values as our dispersion measure at time t, FXC_t. Due to the time variation in conditional FX correlations, there is turnover in both the top and bottom deciles; in order to abstract from composition effects, we also compute an alternative dispersion measure (FXC^{UNC}) by using the top and bottom deciles of FX pairs based on the unconditional realized correlations.

We plot the time series of the level of the two FX correlation dispersion measures in Panel A of Figure 2. The correlation between the two series is 0.87, indicating that the two measures are very similar.[4] Indeed, during the financial crisis the two measures are almost perfectly correlated, as there is almost no turnover in the extreme deciles of FX pair conditional correlation. In Panel B, we plot the (standardized) market variables we used to measure the cyclicality of correlations. Table 5 reports the unconditional correlations between our two FX correlation dispersion measures and the market variables. All correlations are significantly positive, confirming our previous empirical findings.

[Insert Figure 2 and Table 5 here.]

2.2 Correlation risk and the cross section of currency returns

We can now explore the pricing of FX correlation risk in the cross section of currency returns. To do so, we sort currencies into portfolios based on their exposure to our dispersion measure FXC. We measure exposure to FX correlation dispersion by the currency return beta with respect to innovations in the FX correlation dispersion measure FXC; innovations for the period t to $t+1$ are defined as the average of first differences in conditional FX correlation for the FX pairs that belong to the top and bottom decile in

[4] The Online Appendix presents additional results for alternative construction methods. Overall, we find that results using the alternative measures remain qualitatively the same.

period t.[5] Our currency portfolios are rebalanced monthly: each month t we calculate rolling betas using 36 monthly observations and, hence, portfolios are formed using only information available at time t.

We sort the nine G10 currencies into three portfolios; the first portfolio ($Pf1^C$) contains the currencies with the low FXC betas while the last portfolio ($Pf3^C$) contains the high FXC 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 three FXC-beta-sorted currency portfolios using the G10 currencies. Notably, there is an inverse relationship between exposure to FXC and average portfolio returns: average portfolio returns are monotonically decreasing in the FXC beta. 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 6.4% with an associated Sharpe ratio of 0.82 and a t-statistic 3.5.

[Insert Table 6 here.]

Our finding of a strongly negative return for HML^C is robust to different sample periods. Our benchmark sample period starts in January 1996 in line with the availability of the options data. However, the construction of the risk factor FXC does not rely on implied correlations but only on realized correlations that are calculated using daily changes in log exchange rates. Hence, the time span for calculating the factor can be extended further back. In particular, we consider the subperiods from January 1984 to December 2013, from January 1984 to July 2007 and from January 1996 to July 2007. Consistent with our results in the benchmark period, we find an inverse relationship between exposure to the FX correlation factor FXC and average portfolio returns in each of the three additional sample periods. Although the return differences across portfolios somewhat attenuate when the sample period is extended back to 1984 (Panels B and C), shorting the HML^C portfolio still yields highly significant annualized average excess returns of between 3.5% and 3.7%, respectively. On the other hand, ending the sample

[5] Innovations in FXC are not the first differences in the level of the factor, as the composition of the deciles changes over time. On the other hand, since the FX pairs used to calculate FXC^{UNC} are fixed, innovations in FXC^{UNC} can be simply defined as first differences in the level of the factor.

in July 2007, in the beginning of the financial crisis, results in an increase of the average excess return of the HML^C portfolio to more than 7% and of its Sharpe ratio to 1.1. Overall, our results are very robust to different sample periods and do not appear to be driven by the recent financial crisis.

[Insert Figure 3 here.]

We can also extend the currency sample and consider a cross section that includes other developed country currencies (the developed country sample) and one that includes the entirety of the developed sample and also some emerging currencies (the full sample).[6] For each of the two extended samples, we construct four FXC-beta-sorted portfolios. Figure 3 presents the average excess return of HML^C for the three sets of currencies (G10, all countries and developed countries) and four subperiods. We find that there is a consistently negative relationship between average portfolio excess returns and exposure to correlation risk for all currency samples and sample periods. Average HML^C excess returns are significant at the 5% level for all currency and period samples with the exception of the samples starting in 1984 for the full (developed and emerging) set of currencies. For example, for the benchmark subperiod from January 1996 to December 2013 the average annualized return of HML^C in the developed country cross section is 5.5% (with a t-statistic of 2.4) and its Sharpe ratio is 0.59. For the full cross section of currencies, HML^C yields 4.0% on average (with a t-statistic of 2.0) and a Sharpe ratio of 0.46.

2.3 The price of correlation risk

Given the significant excess returns to the HML^C portfolio, it is natural to test whether correlation risk is priced in the cross section of currencies. 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 returns and shown by Lustig, Roussanov, and Verdelhan (2011) to act as a level factor in currency markets,

[6] See Section 1 for a full list of the currencies in the respective samples.

and the second factor is HML^C, the return difference between the high and the low correlation beta portfolio for the sample of G10 currencies.

In particular, we consider the following model:

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

where rx^i denotes the excess return in levels (i.e., corrected for the Jensen term). To estimate the factor prices λ we follow the two-stage procedure of Fama and MacBeth (1973): first, we run a time series regression of returns on the factors and then we run a cross-sectional regression of average portfolio returns on factor betas. We do not include a constant in the cross-sectional regression of the second stage.[7]

Panel A in Table 7 reports the first stage regression results. We consider 15 test assets: the three currency portfolios sorted on exposure to FXC, three currency portfolios sorted on forward discounts (called"carry portfolios" and denoted by $Pf1^F$, $Pf2^F$ and $Pf3^F$) and nine individual G10 exchange rates. As expected, the HML^C betas of the correlation portfolios are monotonically increasing, while the HML^C betas of the carry portfolios are monotonically decreasing in forward discounts. 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.

Panel B presents the second-stage results for various sets of test assets. For set (1), we estimate the market price of risk using only the three correlation-sorted ($Pf1^C$ to $Pf3^C$) and the three carry ($Pf1^F$ to $Pf3^F$) portfolios. For set (2), we also add 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 for each second-stage regression. We find a significantly negative price of correlation risk of -58bps or -54bps per month for sets (1) and (2), respectively. Those estimates are not significantly different from the average HML^C excess return of -54bps per month. The second stage R^2 are very high for both sets of test assets (0.99 and 0.93, respectively).

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

[Insert Table 7 and Figure 4 here.]

For robustness, we extend our sample to developed countries and developed and emerging countries, labelled sets (3) and (4), respectively. For each of the two additional sets, we use eight test assets: four correlation and four carry portfolios. The second stage results are also provided in Panel B of Table 7 and the estimates are in line with our benchmark results: the price of correlation risk is significantly negative, ranging between -51bps and -67bps per month for sets (3) and (4), respectively. The R^2 are again high: 0.90 for developed currencies and 0.81 for the full set of currencies.[8]

Figure 4 illustrates the performance of our two-factor model by plotting the predicted annualized excess returns for the test assets against their actual counterparts: Panels A, B and C refer to the test assets and prices of risk of sets (2), (3) and (4), respectively. For all three sets, the deviations from the model are small, as expected by the high cross-sectional R^2s in the second-stage regressions.

We have shown that our traded correlation risk factor HML^C acts as a slope factor regarding the pricing of currency risk. Thus, a natural question that arises regards the relationship between HML^C and the Lustig, Roussanov, and Verdelhan (2011) carry trade factor HML^{FX}, which reflects the returns to a portfolio that invests in high interest rate currencies and shorts low interest rate currencies. Empirically, the two factors are strongly negatively correlated, suggesting that they capture similar sources of systematic risk. We defer a more detailed exploration of the relationship between HML^C and HML^{FX} to Section 4.4, which discusses the two factors in the context of our proposed no-arbitrage model.[9]

[8]To save space we defer the first stage regression results for the additional sets of currencies to the Online Appendix.

[9]In the Online Appendix, we also consider the relationship between our FX correlation risk factor and the FX volatility risk factor of Menkhoff, Sarno, Schmeling, and Schrimpf (2012).

3 Exchange rate correlation risk premia

In this section, we document the cross-sectional and time series properties of correlation risk premia (CRP). Then, we explore the links between FX correlation risk premia and FX correlations.

3.1 The cross sectional properties of correlation risk premia

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

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

We only consider one-month premia, i.e., $T = t + 1$ for a monthly frequency.[10]

To calculate the risk-neutral (implied) conditional FX correlations we follow the literature on model-free measures of implied volatility and covariance using daily FX option prices. The details of the calculations are outlined in Appendix B. Given the availability of FX options, we calculate correlation risk premia for the nine G10 currencies during the sample period from 1996 to 2013 for a total of 216 monthly observations. For the EUR, the options data starts in 1999 for a total of 181 observations.

Columns (3) and (4) of Table 2 present the time-series mean and standard deviation, respectively, of the implied conditional FX correlations for each of the 36 G10 FX pairs. The cross-sectional mean of average implied FX correlation is 0.48, slightly higher than its physical measure counterpart (0.45). However, the cross-sectional range of average implied FX correlations is lower than that of physical measure ones: the lowest average implied conditional FX correlation is 0.14 (CAD/JPY pair) and the highest is 0.88 (CHF/EUR pair), whereas the average realized correlations are 0.05 and 0.89 for the same pairs. Thus, the heterogeneity of FX pairs regarding their average conditional

[10]Variance risk premia are defined analogously as the difference between the risk-neutral and objective measures of FX variance. A brief discussion of their summary statistics is deferred to the Online Appendix.

correlation is lower under the risk-neutral measure than under the physical measure. However, the volatility of conditional implied FX correlations has the same order of magnitude as the volatility of conditional FX correlations under the physical measure, with standard deviations ranging from 0.07 to 0.34 and the cross-sectional average being 0.19 (for the realized correlations, the range is from 0.09 to 0.34 and the average is slightly higher at 0.23).

The last five columns of Table 2 present the descriptive statistics for correlation risk premia. From left to right, we report the time-series mean and standard deviation of the correlation risk premia for each exchange rate pair, followed by the asymptotic t-statistic and the bootstrapped 95% confidence interval of the CRP mean. Average FX correlation risk premia can be substantial and exhibit considerable cross-sectional heterogeneity, with their size and sign varying greatly across FX pairs. In particular, average correlation risk premia range from -0.069 (CAD/SEK) to 0.099 (JPY/NOK), with the cross-sectional mean being 0.016, indicating that risk-neutral FX correlations are on average slightly higher than realized FX correlations. Roughly two thirds of the average CRP are positive and one third are negative; overall, three quarters of all average premia are significant at the 5% level according to either the asymptotic or the bootstrapped distributions. The average of the bottom quartile of correlation risk premia is -0.04, whereas the top quartile average is 0.07.[11] Furthermore, correlation risk premia are very volatile: despite the fact that CRP are much smaller than either physical measure 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 the cross-sectional average equal to 0.14), suggesting that there is substantial time variation in the disparity between the physical and the risk neutral measure.

To explore whether average correlation risk premia exhibit a cross-sectional pattern, we plot the average CRP of all G10 exchange rate pairs against their average realized correlations. Figure 5 presents the plot, along with the line of best fit. The cross-sectional correlation between average FX correlation risk premia and average FX realized correlations is -0.55. For example, the AUD/JPY pair, characterized by a very low average

[11] In terms of order of magnitude this is up to over 40% of the correlation risk premium reported for the equity market by Driessen, Maenhout, and Vilkov (2009).

realized FX correlation (0.16), has a positive and highly significant average CRP of 0.083. On the other hand, 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 premia on the average realized correlations yields a statistically significant slope coefficient of -0.144.[12] The strongly negative cross-sectional relationship between average realized FX correlations and average FX correlation risk premia is what generates the lower cross-sectional range of average correlations under the risk-neutral measure, as opposed to the average correlations under the physical measure that we discussed earlier.

[Insert Figure 5 here.]

3.2 The time series properties of correlation risk premia

We now turn to the time series properties of conditional implied FX correlations and FX correlation risk premia. The first four columns of Table 8 provide summary statistics on the time-series correlations between physical measure and risk-neutral measure conditional FX correlations: for each FX pair, we report the unconditional correlation coefficient between the two time series, as well as its asymptotic t-statistic and the 95% bootstrapped confidence interval. Physical measure and implied conditional FX correlations exhibit substantial comovement for all FX pairs, with unconditional correlations between the two ranging from 0.70 to 0.92, all statistically significant. Therefore, implied conditional FX correlations appear to track observed FX correlations very well.

[Insert Table 8 here.]

The last four columns of Table 8 report descriptive statistics on the unconditional correlation between physical measure correlations and CRP. We find that the correlation averaging -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

[12] Its asymptotic t-statistic, calculated with White (1980) standard errors, is -5.80 and the bootstrapped 95% confidence interval is $[-0.154, -0.076]$.

between the physical measure and the risk neutral measure. This relationship is pervasive and robust: 35 of the 36 unconditional correlations are negative, ranging from -0.760 to -0.102, with all but one of those being statistically different from zero.

Finally, to assess the cyclicality of correlation risk premia, we construct CRP cyclicality measures. We define those in a fashion similar to FX correlation cyclicality measures: they are unconditional correlations between correlation risk premia and the four market variables we used before. The last four columns of Table 3 provide the four CRP cyclicality measures for the G10 FX pairs. Panels A to D of Figure 6 then plot those measures against average correlation risk premia. We find a positive cross-sectional association: FX pairs with high average CRP have countercyclical correlation risk premia, whereas pairs with low average CRP have procyclical premia. The regression results in Panel B of Table 4 suggest that this positive cross-sectional association is statistically significant for all four cyclicality measures.

[Insert Figure 6 here.]

In sum, FX pairs with high average physical measure correlations or low average correlation risk premia exhibit countercyclical correlations and procyclical correlation risk premia, whereas FX pairs with low average correlations or high average correlation risk premia have procyclical correlations and countercyclical correlation risk premia. Thus, just as the cross-sectional dispersion in conditional FX correlations is countercyclical, so is the dispersion in correlation risk premia: in bad times, the premia of FX pairs with high average CRP increase and the premia of FX pairs with low average CRP decline.

4 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 cross-country comovement in the pricing of local risk. This assumption allows our model to match the joint empirical properties of FX correlations and FX correlation risk premia.

4.1 Model setup

In order to illustrate the basic economic mechanisms in operation, we first focus on a simple version of our full model, which will be referred to as the benchmark model.

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

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

where z and z^w are the local and the global pricing factors, respectively. Both pricing factors are common to all countries. Notably, countries are ex ante heterogeneous only with regard to their exposure γ to the first global shock u^w; all other SDF parameters are identical across countries. As we will see, differences in γ capture the exchange rate fixed effect that is manifested, inter alia, in the cross-sectional differences in average FX correlations discussed in the empirical section. In our model, global risk exposure γ is exogenous. Richer models that endogenize unconditional cross-sectional differences in global risk exposure include Hassan (2013), Martin (2013), Tran (2013), Backus, Gavazzoni, Telmer, and Zin (2013), Ready, Roussanov, and Ward (2013), and Colacito, Croce, Gavazzoni, and Ready (2014).

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 $\sqrt{\delta z_t}$. On the other hand, countries have heterogeneous exposure to the first global shock u^w, and its price in country i is $\sqrt{\gamma^i z_t^w}$.

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

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

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

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

It also features countercyclical pricing of risk. To ensure that all pricing factors are strictly positive, we further assume that the Feller conditions $2\lambda\bar{z} > \xi^2$ and $2\lambda^w\bar{z}^w > (\xi^w)^2$ are satisfied. All parameters except α, χ and φ are strictly positive. All the shocks in our model are i.i.d. standard normal.

Lastly, we assume that inflation is constant, normalized to zero for all countries, so real interest rates and exchange rates coincide with their nominal counterparts. We relax this assumption in the full version of our model.

4.2 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 units of domestic currency, in real terms). As a result of financial market completeness, real exchange rate changes equal the SDF differential between the two countries,

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

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

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

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

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

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

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

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

$$cov_t\left(\Delta q_{t+1}^{i}, \Delta q_{t+1}^{j}\right) = \kappa z_t + D^{i,j} z_t^w,$$

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

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

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

The first component of conditional FX covariance is due to common exposure to the domestic local shock, as the two exchange rates are mechanically correlated through their relationship 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. On average, foreign countries with similar exposure to the global shock u^w (i.e., that satisfy $D^{i,j} > 0$) have exchange rates that covary more than the exchange rates of countries that have dissimilar exposure to global FX risk. Furthermore, the effect of fluctuations in z^w on conditional FX covariance depends on the type of the FX pair concerned. In particular, 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: similar FX pairs have higher unconditional correlations than dissimilar ones. As a result, an increase in the global pricing factor z^w increases the cross-sectional dispersion of conditional FX correlations, as it raises the correlation of exchange rates with high average correlation (similar FX pairs) and decreases the correlation of exchange rates with low average correlation (dissimilar FX pairs). In the limit, as $z^w \to \infty$ similar exchange rates become perfectly positively correlated and dissimilar exchange rates become perfectly negatively correlated.

On the other hand, an increase of the local pricing factor z increases both the FX variance and the FX covariance of all exchange rate pairs, the latter due to the domestic currency effect. When $z \to \infty$ the correlation of all FX pairs converges to $\frac{1}{2}$. This happens because all cross-sectional differences in global risk exposure become second-order and what ultimately drives FX comovement is the domestic currency effect. In particular, the limit behavior of log exchange rate changes is described by

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

so 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 towards $\frac{1}{2}$. Thus, when the local pricing factor increases the conditional correlation of similar exchange rates (which have high unconditional correlations) declines, whereas the conditional correlation of dissimilar exchange rates (with low unconditional correlations) increases, leading to a tightening of the cross section of conditional FX correlations.

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

[Insert Figure 7 here.]

We first consider the impact of the global pricing factor z^w; the left panels of Figure 7 present the results. In particular, Panels A, C and E plot 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 to the dissimilar exchange rate pair (1,3) and Panel E plots the difference in the conditional FX correlations of the two FX pairs. Panel A shows that the conditional correlation of the similar FX pair is always increasing in z^w, as the similarity of the two exchange rates to global risk exposure increases their comovement when global fluctuations become more highly priced. Exactly the opposite occurs for the dissimilar exchange rate pair: as seen in Panel C, an increase in z^w always reduces their conditional correlation. Taken together, these results imply that the disparity in conditional FX correlations is increasing in z^w, as illustrated in Panel E.

We now turn to the effects of the local pricing factor z. The results are presented in the right panels of Figure 7; Panels B, D and F 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. As seen in Panel B, the conditional correlation between the two similar exchange rates is decreasing in z regardless of the value of z^w. This is due to the domestic currency effect, which pushes the correlation of the two exchange rates towards $\frac{1}{2}$ when z takes large values. On the other hand, this effect induces a negative relationship between the value of the local pricing factor and the conditional FX correlation of the dissimilar pair, as seen in Panel D. As a result, the difference in the two FX correlations is decreasing in z regardless of the value of z^w: Panel F shows that as z increases all foreign currencies appreciate and depreciate together against the domestic currency.

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, that implies that changes in FXC reflect both u^w shocks (with a positive sign) and u^g shocks (with a negative sign). Empirically, we have seen that 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 risk u^w, rather than to the second global shock u^g.

4.3 Correlation risk and the cross section of FX returns

Recall that the USD excess return for investing in the currency of country i is given by:

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

so FX 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

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

FX risk premia have two components: a part that compensates domestic investors for the fact that investing in a foreign currency essentially entails shorting the 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 premia is solely due to heterogeneity in exposure to u^w, i.e. heterogeneity in γ. In particular, 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 in order to hold currency i. Conversely, currencies of countries with high exposure to u^w have a negative compensation 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. We have seen that fluctuations in FXC reflect both innovations in the global pricing factor z^w (which are scaled multiples of global shock u^w) and innovations in the the local pricing factor z^w (scaled multiples of global shock u^g). Importantly, both kinds of innovations are globally priced and they have opposite effects on FXC, so it is not trivial to establish whether a positive loading of an asset return on FXC should be associated with a positive or a negative risk premium: assets should earn a negative premium for a positive loading on FXC that arises from exposure to u^w, and a positive premium for a positive loading that arises from exposure to u^g. However, there is no ambiguity 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:

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

Thus, countries i with a higher SDF exposure γ^i to global risk u^w than the domestic country have USD exchange rates with a positive conditional loading on ΔFXC; conversely, the USD exchange rate of countries with $\gamma^i < \gamma^0$ has a negative loading on ΔFXC. Given the negative cross sectional relationship between γ and currency risk premia, that implies a negative risk premium for high FXC beta exchange rates and a positive premium for low FXC beta exchange rates, in line with our empirical finding of a negative price of FXC.

4.4 Correlation risk and the carry trade factor

In Section 2.3 we show empirically that the traded correlation risk factor HML^C prices the cross section of currencies. At this stage we can discuss how correlation risk relates in the model to the well-known Lustig, Roussanov, and Verdelhan (2011) carry trade factor, a portfolio that invests in high interest rate currencies and shorts low interest rate currencies.

The real interest rate of country i is given by

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

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 risk have a relatively low (high) interest rate, due to a higher (lower) precautionary savings motive. As a result, high interest rate currencies are associated with low γs and, thus, high risk premia.

The excess return to the carry trade portfolio HML^{FX} is defined as

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

with high interest rate (low γ, according to the model) currencies in set H and low interest rate (high γ) currencies in set L. Provided that currency portfolios contain enough currencies so that the local shocks cancel out, the return innovations of the HML^{FX} portfolio are perfectly positively correlated with the global shock u^w:

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

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

On the other hand, fluctuations in the correlation factor FXC capture both kinds of global innovations, 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} will always have better pricing abilities than FXC in the cross section of currency returns. To get a cleaner measure of u^w innovations, we should consider FX excess return differentials, which are only exposed 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:

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

with high-ΔFXC-loading (i.e. high γ) currencies in set H and low-ΔFXC-loading (low γ) currencies in set L. 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:

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

Therefore, HML^C returns are perfectly negatively correlated with HML^{FX} returns for a large enough set of currencies: they both reflect u^w shocks and, thus, should have the same explanatory power for the cross section of FX returns: high γ currencies, which hedge u^w risk, have a low interest rate and a high ΔFXC beta and low currency risk premia, whereas low γ (high interest rate, low ΔFXC beta) have high currency risk premia. Figure 8 plots HML^C and HML^{FX} from January 1996 to December 2013. Indeed, the correlation between the two time series is -0.66, suggesting a very strong negative association between the two factors that is consistent with our model predictions.

[Insert Figure 8 here.]

4.5 The properties of correlation risk premia

We now turn to correlation risk premia and show that our empirical findings can be rationalized in our model if the domestic agent prices fluctuations in the local pricing factor sufficiently more strongly than fluctuations in the global pricing factor.

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

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

so the drift adjustment is positive and equal to $\xi^w\sqrt{\gamma^0} z_t^w$. We can rewrite equation (1) as a square root process,

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

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

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

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. Notably, the drift adjustment of the two factors depends crucially on the volatility parameters ξ and ξ^w, which determine the sensitivity of the pricing factors to shocks u^g and u^w respectively, and the exposure parameters δ and γ^0, which regulate the pricing of shocks u^g and u^w respectively. The higher ξ is compared to ξ^w, and the higher δ is relative to γ^0, the higher the relative drift adjustment of the local pricing factor over the global pricing factor, as the shocks of the former are more highly priced compared to the shocks of the latter.

Note that for the local pricing factor we have

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

under the risk-neutral measure, compared to

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

under the physical measure. Given the higher steady-state and higher persistence of the local pricing factor under the risk-neutral measure, the wedge $E^{\mathbb{Q}}_t(z_{t+s}) - E_t(z^{\mathbb{P}}_{t+s})$ is always positive and increasing in z_t.[13] Exactly the same is true for the global pricing factor z^w.

Thus, the magnitude of FX correlation risk premia is determined by the disparity between the risk-neutral measure and the physical measure behavior of z and z^w. The expression for FX correlation risk premia is derived in Appendix C. Of particular rele-

[13] In particular, the wedge is an affine function of z_t, 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 for $x \in (0,1)$.

vance is the case in which the domestic agent prices fluctuations in the local pricing factor more heavily than fluctuations in the global pricing factor, i.e., when $\xi\sqrt{\delta} >> \xi^w\sqrt{\gamma^0}$. In that case, the FX correlation under the risk-neutral measure is akin to the physical measure FX correlation with a large upwards adjustment for the local pricing factor and a much smaller upwards adjustment for the global pricing factor, so the relative importance of the global pricing factor is smaller under the risk-neutral measure than under the physical measure. This has implications for both the cross section and the time series of FX correlation risk premia.

We can start with the cross-sectional implications. As discussed, the global pricing factor is relatively less important under the risk-neutral measure than under the physical measure. As a result, the risk-neutral FX correlation is always lower than the physical correlation for similar FX pairs, and higher than the physical correlation for dissimilar FX pairs. This implies that similar FX pairs have negative average CRP and dissimilar FX pairs have positive average CRP, which generates a negative cross-sectional relationship between average FX correlations and average CRP, in line with the empirical findings presented in Figure 5. The left side panels of Figure 7 provide a useful visualization: keeping z^w constant, higher values for z correspond to lower conditional correlation for similar FX pairs and higher conditional correlation for dissimilar FX pairs. Indeed, the high z curves are always below the low z curves for similar FX pairs (Panel A) and always above the low z curves for dissimilar FX pairs (Panel C), implying negative CRP for similar FX pairs and positive CRP for dissimilar FX pairs.

We can now turn to the time series of correlation risk premia. The key feature of the model is that risk-neutral FX correlations are less sensitive to the value of z^w than their physical measure counterparts; this is because conditional FX correlation is typically a less steep function of z^w for higher values of z. This implies that in states in which z^w is high, which are the states identified as bad by our four empirical business cycle proxies, the increase (decrease) of the correlation of similar (dissimilar) FX pairs that occurs in the physical measure is attenuated in the risk-neutral measure. This amplifies the discrepancy between risk-neutral and physical measure FX correlation for all FX pairs: the (negative) correlation risk premia for similar FX pairs decrease further and the (positive) premia for dissimilar FX pairs increase further. Since the correlation of similar FX pairs is increasing in z^w and the correlation of dissimilar FX pairs is decreasing in z^w, the aforementioned behavior of CRP implies a negative time-series correlation between FX correlation and CRP levels for all FX pairs and a widening of the cross-sectional dispersion of CRP when z^w is elevated, consistent with our empirical findings. Graphically, consider the left panels of Figure 7 and focus on the distance between the low z and high z curves: as the global pricing factor z^w increases, the distance between the two curves widens for both similar and dissimilar FX pairs, since conditional FX correlation is a flatter function of z^w for higher z values. This widening characterizes the behavior of correlation risk premia in the region of the state state in which the economy spends most of its time, although it starts to reverse for very high values of z^w.

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

4.6 Quantitative performance

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

In this section, we consider the full version of our model, which allows for non-trivial exogenous inflation processes with unpriced innovations. In particular, the local pricing factor of country i, z^i, satisfies

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

so it is driven by both the global shock u^g and the local shock u^i. If we assume that $\rho = 0$ and that all local pricing factors have the same initial value, then all local pricing factors are identical and we retrieve our benchmark model. On the other hand, if we assume that $\rho = 1$, then we retrieve the model in Lustig, Roussanov, and Verdelhan (2014), which features independent local pricing factors.

The inflation process for country i is given by

$$\pi^i_{t+1} = \bar{\pi} + \zeta z^w_t + \sqrt{\sigma}\eta^i_{t+1}.$$

Inflation is exposed to i.i.d. innovations η^i. Since inflation shocks are unpriced, there is no inflation risk premium.

We provide a more detailed description of the full model and its calibration and smilation in Appendix D. The values of our calibrated parameters are reported in Table 9. To illustrate the importance of comovement in the local pricing factors, we consider the two polar values of ρ: $\rho = 0$, as in our benchmark model, and $\rho = 1$, as in Lustig, Roussanov, and Verdelhan (2014).

Table 10 reports moments for inflation, real and nominal interest rates and real and nominal exchange rates for the U.S. and the foreign countries for our model ($\rho = 0$). All moments are well matched, although the model slightly overshoots inflation volatility and, as a result, nominal interest rate volatility. The model exhibits an annualized SDF volatility of about 60%, in line with empirical bounds. Finally, the model generates a strong carry trade effect, with the return on the FX carry portfolio having an average excess return of 2.62%. It is worth mentioning that 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 -0.45, 0.02, and 0.55, respectively, with average excess returns monotonically increasing in the HML^{FX} portfolio betas. This mirrors the real data where betas are increasing in the interest rate differentials and excess returns are increasing in the portfolio betas.

[Insert Tables 9 and 10 here.]

Table 11 presents the FX correlation moments both our model ($\rho = 0$) and for the model with independent local pricing factors ($\rho = 1$). We first discuss our model. As seen in Table 11, the model generates a non-trivial cross-sectional spread in average physical and implied FX correlation coefficients, in line with empirical evidence, and is able to closely match the cross-sectional average of mean FX correlations. The only weakness of the model regards the magnitude of correlation risk premia: the model-implied correlation risk premia are much lower (in absolute terms) than their empirical counterparts. Notably though, the model is able to successfully generate both positive and negative correlation risk premia, as in the data.

[Insert Table 11 here.]

For the model to generate higher correlation risk premia, δ, countries' exposure to the second global stock u^g, has to be higher: a higher exposure to the second global shock would increase the relative pricing of local pricing factor fluctuations by the domestic agent and, thus, would raise the pricing importance of states characterized by high values of z. Although the value of δ does not affect FX properties, as exposure to the second global shock is identical across countries and, thus, cancels out from exchange rates, higher values of δ would increase SDF volatility and lower interest rates by strengthening the precautionary savings motive. Furthermore, if the local pricing factor is not exactly identical across countries (i.e., $\rho > 0$), then exchange rate volatility is increasing in δ.

The model is able to replicate the almost perfect positive cross-sectional relationship between average realized and average implied correlations and, crucially, the strongly negative cross-sectional relationship between average realized correlations and average CRP. In particular, in our 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, as happens empirically. Regarding the time series, the model generates a perfect time-series correlation between realized and implied correlation and a negative time series correlation between realized correlation and CRP for all FX pairs. Empirically, realized and implied correlations are very highly correlated for all FX pairs (unconditional correlations range from 0.70 to 0.92) and almost all FX pairs display a negative time series association between realized correlation and CRP, so the time series properties of the model are fully in line with the data.

Finally, we consider the asset pricing implications of the model. In the simulated data, the annualized average excess return for the currency portfolio that is long currencies with high exposure to FXC innovations and short currencies with a low exposure is -1.21%, suggesting a negative price for exposure to FX correlation risk. As discussed in the previous section, our FXC factor is not unrelated to the Lustig, Roussanov, and Verdelhan (2011) HML^{FX} factor: indeed, the HML^{FX} factor is priced in the cross-section of FXC-beta-sorted currency portfolio returns, with the low, medium, and high FXC exposure portfolio having an HML^{FX} beta of 0.28, 0.03, and -0.19, respectively. This is in line with the data where the HML^{FX} betas of the correlation sorted portfolios are also monotonically decreasing. Furthermore, there is a negative cross-sectional relationship between interest rates and FXC betas: the low, medium and high FXC loading portfolio has an average interest rate differential (against the domestic country) of 0.65%, 0.04% and -0.51%, respectively, which is again in line with the data.

We now turn to the model with independent local pricing factors. In that model, foreign pricing factors are driven by foreign idiosyncratic shocks, unpriced by the domestic investor. As a result, the domestic investor only prices fluctuations in the domestic pricing factor z^0 and in the global pricing factor z^w, but not fluctuations in the foreign pricing factors. This greatly affects the properties of FX correlation risk premia. Recall that when local pricing factors are identical, pricing states in which the local pricing factor has a high value delivers negative correlation risk premia for similar FX pairs and positive premia for dissimilar FX pairs, consistent with the empirical evidence. However, in the absence of comovement in local pricing factors, the risk adjustment for z^0 tends to generate positive correlation risk premia for all FX pairs; a discussion is provided in Appendix D. Indeed, in our simulation average correlation risk premia are positive for all currencies, ranging from 0.01% to 0.04%. The small magnitude of FX correlation risk premia is due to the low value of parameter κ, which regulates (along with ξ) the pricing of z^0 fluctuations when local pricing factor shocks are country-specific.

Furthermore, the model generates an almost perfect positive cross-sectional association between average FX correlation and average FX CRP, at odds with empirical evidence. This is because the pricing of z^0 fluctuations is not strong enough to dominate the pricing effects of z^w fluctuations in our parametrization. Lastly, the model fails to match the empirical time series properties of CRP: the time series of simulated realized correlations and CRP are uncorrelated for all FX pairs, at odds with the strong negative correlation that characterizes their empirical counterparts. Increasing the value of local risk loading κ would provide some limited improvement of the cross-sectional properties of simulated average FX correlation risk premia, but would also change the properties of a number of variables, including FX correlations, in counterfactual ways; for example, it would reduce the relative importance of heterogeneous country exposure to global risk, generating a very tight cross section of FX correlations—all average FX correlations would be around 0.5 due to the domestic currency effect.

5 Conclusion

We document a negative cross-sectional relationship between average FX correlations and correlation cyclicality, implying that FX pairs that are highly correlated on average become even more correlated in bad times while pairs characterized by low average correlations become even less correlated. Thus, FX correlations become more cross-sectionally dispersed in adverse economic states.

We capture the countercyclicality of cross-sectional dispersion in conditional FX correlations by constructing the FX correlation factor FXC, 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 FXC beta currency portfolios is economically and statistically large (6.4% annually) and that the price of FX correlation risk is almost -7% per year. In short, we find that investors want to be compensated for investing in currencies that perform badly during periods of increased cross-sectional dispersion in conditional FX correlations.

Defining the FX correlation risk premium as the difference between the FX correlation under the risk-neutral and physical probability measure, we find a strong negative relationship between FX correlations and FX correlation risk premia both in the cross section and in the time series. In the cross section, FX pairs with high average correlation exhibit low (or negative) average correlation risk premia, while the opposite is true for FX pairs with low average correlations. Furthermore, the cross sectional dispersion of FX correlation risk premia increases in bad economic states.

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

References

Backus, D. K., F. Gavazzoni, C. Telmer, and S. E. Zin (2013): "Monetary Policy and the Uncovered Interest Rate Parity Puzzle," Working Paper, New York University.

Bollerslev, T. G., G. Tauchen, and H. Zhou (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.

Brandt, M. W., and F. X. Diebold (2006): "A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations," *Journal of Business*, 79, 61–73.

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

Cenedese, G., L. Sarno, and I. Tsiakas (2014): "Foreign Exchange Risk and the Predictability of Carry Trade Returns," *Journal of Banking and Finance*, 42, 302–313.

Colacito, R., and M. Croce (2013): "International Asset Pricing with Recursive Preferences," *Journal of Finance*, 68, 2651–2686.

Colacito, R., M. M. Croce, F. Gavazzoni, and R. C. Ready (2014): "Currency Risk Factors in a Recursive Multi-Country Economy," Working Paper, UNC.

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

Dobrynskaya, V. (2014): "Downside Market Risk of Carry Trades," *Review of Finance*, 24, 1–29.

Driessen, J., P. Maenhout, and G. Vilkov (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., and J. MacBeth (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81, 607–636.

Farhi, E., S. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2015): "Crash Risk in Currency Markets," Working Paper, New York University.

Fenn, D. J., S. D. Howison, M. McDonald, S. Williams, and N. F. Johnson (2009): "The Mirage Of Triangular Arbitrage In The Spot Foreign Exchange Market," *International Journal of Theoretical and Applied Finance*, 12, 1105–1123.

Gabaix, X., and M. Maggiori (2015): "International Liquidity and Exchange Rate Dynamics," *forthcoming, Quarterly Journal of Economics*.

Garman, M. B., and S. W. Kohlhagen (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., J. Pan, and J. Wang (2013): "Noise as Information for Illiquidity," *Journal of Finance*, 68, 2341–2382.

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

Kozhan, R., and W. W. Tham (2012): "Execution Risk in High-Frequency Arbitrage," *Management Science*, 58, 2131–2149.

Lettau, M., M. Maggiori, and M. Weber (2014): "Conditional Risk Premia in Currency Markets and Other Asset Classes," *Journal of Financial Economics*, 114, 197–225.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24, 3731–3777.

——— (2014): "Countercylical Currency Risk Premia," *Journal of Financial Economics*, 111, 527–553.

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

Malkhozov, A., P. Mueller, A. Vedolin, and G. Venter (2015): "International Illiquidity," Working Paper, London School of Economics.

Mancini, L., A. Ranaldo, and J. Wrampelmeyer (2013): "Liquidity in the Foreign Exchange Market: Measurement, Commonality, and Risk Premiums," *Journal of Finance*, 68, 1805–1841.

Martin, I. (2013): "The Forward Premium Puzzle in a Two-Country World," Working Paper, Stanford University.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012): "Carry Trades and Global Foreign Exchange Volatility," *Journal of Finance*, 67, 681–718.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

Ready, R., N. Roussanov, and C. Ward (2013): "Commodity Trade and the Carry Trade: a Tale of Two Countries," Working Paper, University of Rochester.

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," Working Paper, Olin Business School.

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.

Wystup, U. (2006): *FX Options and Structured Products*. Chichester: John Wiley and Sons.

Appendix A Realized variances and correlations

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

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

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

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

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

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

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

Appendix B Implied variances and correlations

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

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

where S_t is the underlying spot exchange rate and $\mathrm{P}(K,T)$ and $\mathrm{C}(K,T)$ are the respective put and call prices with maturity date T and strike K. In practice, the number of traded options for any underlying asset is finite; hence the available strike price series is a finite sequence. Calculating the model-free implied variance involves the entire cross section of option prices: for each maturity T, all five strikes are taken into account. These are quoted in terms of the option delta. In addition, we use daily spot rates and one-month Eurocurrency (LIBOR) rates from Datastream. Following the conventions in the FX market we use the use the Garman and Kohlhagen (1983) valuation formula to extract the relevant strike prices and to calculate the corresponding option prices.[14]

To approximate the integral in equation (A-1), we adopt a trapezoidal integration scheme over the range of strike prices covered by our dataset. Jiang and Tian (2005) report two types of implementation errors: (i) truncation errors due to the non-availability of an infinite range of strike prices; and (ii) discretization errors that arise due to the unavailability of a

[14] See, e.g., Wystup (2006) for the specifics of FX options conventions.

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.[15] For the construction we require all cross rates for three currencies, S_t^i, S_t^j, and S_t^{ij}. The absence of triangular arbitrage then implies that:[16] $S_t^{ij} = S_t^i / S_t^j$. Taking logs, we derive the following equation:

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

Finally, taking variances yields:

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

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

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

Using the standard replication arguments, we find that:

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

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

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

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

[16] 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 no effect on calculated quantities.

Appendix C Correlation risk premia in the model

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

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

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

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

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

$$\begin{aligned} CRP_t^{i,j} \;=\; & \frac{\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}} z_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)\left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}} z_t^w\right)}{\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}} z_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}} z_t^w\right)}\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}} z_t\right) + \left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)^2\left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}} z_t^w\right)}} \\ & - \frac{\kappa\left(A + B z_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)\left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)\left(A^w + B^w z_t^w\right)}{\sqrt{2\kappa\left(A + B z_t\right) + \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2\left(A^w + B^w z_t^w\right)}\sqrt{2\kappa\left(A^{\mathbb{Q}} + B^{\mathbb{Q}} z_t\right) + \left(\sqrt{\gamma^j} - \sqrt{\gamma^0}\right)^2\left(A^w + B^w z_t^w\right)}}. \end{aligned}$$

where 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 are positive constants, defined as follows. For the local pricing factor we have

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

under the risk-neutral measure and

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

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 $X_s = A_s, B_s, A_s^{\mathbb{Q}}, B_s^{\mathbb{Q}}, A_s^w, B_s^w, A_s^{w,\mathbb{Q}}$ and $B_s^{w,\mathbb{Q}}$, we respectively define $X = A, B, A^{\mathbb{Q}}, B^{\mathbb{Q}}, A^w, B^w, A^{w,\mathbb{Q}}$ and $B^{w,\mathbb{Q}}$ as

$$X \equiv \sum_{s=0}^{T-t-1} X_s.$$

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 fluctuations in the local pricing factor are priced more heavily than fluctuations in the global pricing factor by the domestic agent, i.e., when $\xi\sqrt{\delta} >> \xi^w\sqrt{\gamma^0}$, then it holds that

$$\left(A^{\mathbb{Q}} + B^{\mathbb{Q}} z_t\right) - \left(A + B z_t\right) >> \left(A^{w,\mathbb{Q}} + B^{w,\mathbb{Q}} z_t^w\right) - \left(A^w + B^w z_t^w\right),$$

which, as discussed in the main text, suggests a negative cross-sectional relationship between average FX correlations and average FX correlation risk premia.

Appendix D The full model

The full model allows for imperfect comovement across local pricing factors and introduces non-trivial inflation dynamics.

If $\rho > 0$, the local pricing factors have different realizations due to the independence of the local shocks. As a result, countries 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:

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

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

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

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

The nominal stochastic discount factor of country i, $m^{i,\$}$ is

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

so nominal exchange rate changes satisfy

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

Inflation differentials add unpriced volatility to exchange rates due to the idiosyncratic nature of the inflation shocks η. In particular, the conditional variance of nominal log exchange rate changes is given by

$$var_t\left(\Delta s_{t+1}^i\right) = var_t\left(\Delta q_{t+1}^i\right) + var_t\left(\pi_{t+1}^i - \pi_{t+1}^0\right) = var_t\left(\Delta q_{t+1}^i\right) + 2\sigma.$$

Furthermore, domestic inflation shocks enhance the domestic currency effect as regards FX comovement:

$$cov_t\left(\Delta s_{t+1}^i, \Delta s_{t+1}^j\right) = cov_t\left(\Delta q_{t+1}^i, \Delta q_{t+1}^j\right) + \sigma.$$

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

Consider the case of independent local pricing factors ($\rho = 1$). In that case, innovations in the domestic pricing factor z^0 are still priced by the domestic agent, but innovations in the foreign pricing factors are not. Therefore, to understand risk adjustment in the risk-neutral measure we only need to consider the impact of z^w and z^0 on conditional FX correlations. We illustrate the effect of z^w and z^0 in Figure 9. As in Figure 7, we study a world of three foreign countries: countries 1 and 2 are less exposed to the first global shock u^w than the domestic country, while country 3 is more exposed.

[Insert Figure 9 here.]

The left panels of Figure 9 depict the conditional FX correlations as a function of the global pricing factor z^w for different values of the domestic pricing factor ($z^0 = 0.2\bar{z}$, $\bar{z}$ and $5\bar{z}$, depicted with circles, solid lines and squares, respectively), holding all the foreign pricing factors equal to 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 benchmark model: as z^w increases, similarities and dissimilarities in exposure to global risk get amplified. However, the position of the curves corresponding to different constant values of z^0 suggests that conditional FX correlation is not a monotonic function of the domestic pricing factor.

The right panels of Figure 9 present conditional FX correlations as a function of z^0 for different values of z^w, assuming that all foreign pricing factors are equal to their steady-state values, and confirm that the relationship between z^0 and the conditional FX correlation is not monotonic. For small values of z^0, conditional FX correlation is higher than its steady-steady value for both similar and dissimilar FX pairs. This is because the local pricing factor of each country regulates its exposure to the second global shock u^g, which now affects exchange rates. For small values of z^0, all FX pairs are similar regarding their exposure to u^g, as the loading of all foreign countries is lower than the domestic loading. As the value of z^0 increases, conditional FX correlation decreases, since 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, increases in z^0 increase the similarity of all FX pairs regarding their exposure to u^g, as the domestic loading becomes much higher than all foreign loadings. As a result, all FX pairs become more correlated.

Crucially, in the absence of comovement in local pricing factors across countries the model's ability to match the negative cross-sectional relationship between average FX correlations and average FX correlation risk premia is severely hindered. Recall that in the benchmark model the desired cross-sectional pattern is achieved by pricing states of the world characterized by high values of the common local pricing factor z. However, Figure 9 shows that when only z^0 fluctuations are priced, pricing states in which the domestic pricing factor z^0 has a high value does not generate the desired cross-sectional pattern. Indeed, conditional FX correlations are increasing in z^0 (in the region above $\bar{z}$) for both similar and dissimilar FX pairs, suggesting that the aforementioned risk adjustment tends to generate positive correlation risk premia for all FX pairs. Graphically, in Figure 9 the curve that corresponds to $z^0 = 5\bar{z}$ is above the curve for $z^0 = \bar{z}$ in both Panel A (similar FX pairs) and Panel C (dissimilar FX pairs).

We can also show that shutting down exposure to the second global shock u^g (by setting $\delta = 0$) in the model with independent local pricing factors does not help with addressing the cross section of FX correlation risk premia. In that case, increases in z^0 always raise the correlation of all exchange rates, implying that the pricing of high z^0 states generates positive average correlation risk premia for all FX pairs.

Our model has $15 + (I + 1)$ parameters: five common SDF parameters (α, χ, ϕ, κ and δ), $I + 1$ heterogeneous parameters (the loading γ for each country), seven common pricing factor parameters—four for the local pricing factor (λ, $\bar{z}$, ξ and ρ) and three for the global pricing factor (λ^w, $\bar{z^w}$ and ξ^w) —and, finally, three common inflation parameters ($\bar{\pi}$, ζ and σ). In our calibration, we follow Lustig, Roussanov, and Verdelhan (2011, 2014) and reduce the set of parameters by imposing the constraint that the loadings γ^i are equally spaced across the foreign countries. In particular, 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 $\gamma^{i} = \gamma^{min} + \frac{i-1}{I-1}(\gamma^{max} - \gamma^{min})$. Thus, there are 18 parameters of interest in total.

We set the values of all parameters (except α, ξ, ξ^{w} and ρ) equal to the corresponding values in Lustig, Roussanov, and Verdelhan (2014). Notably, the calibration in Lustig, Roussanov, and Verdelhan (2014) targets specific interest rate, inflation and exchange rate moments, but does not involve any moments related to FX correlations or FX correlation risk premia. We depart from that calibration as regards the values of α, ξ and ξ^{w} because the Lustig, Roussanov, and Verdelhan (2014) parametrization delivers too high real interest rate means and too low real interest rate volatilities compared to the corresponding empirical values in our sample. To match those moments, we target the mean and variance of the U.S. real interest rate and estimate the three aforementioned parameters using GMM: we leave α 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. In sum, we do not use any of the moments related to either FX correlation or FX correlation risk premia for our calibration. The values of our calibrated parameters are reported in Table 9.[17]

To illustrate the importance of comovement in the local pricing factors, we consider the two polar values of ρ: $\rho = 0$, in which case local pricing factors are identical, as in our benchmark model, and $\rho = 1$, in which case the local pricing factors are independent across countries. In the latter case, our model is identical to the model in Lustig, Roussanov, and Verdelhan (2014). We consider a world with ten countries (I=9 foreign countries and the U.S.) and simulate the model for each of the two polar cases. Each simulation runs for 55,000 monthly periods, and is initialized at the steady-state values $\bar{z}$ and $\bar{z}^{w}$; to reduce the effect of initial conditions, we discard the first 5,000 observations. 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.

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

Appendix E Tables

Table 1
Summary statistics for G10 currencies

This table reports summary statistics for the G10 currencies, namely the mean excess return, standard deviation, corresponding Sharpe ratio, skewness, kurtosis, and forward discount $f_t - s_t$. All returns are excess returns in USD, annualized and expressed in percent. Monthly data from January 1996 through December 2013. Before 1999 we use the DEM instead of the EUR.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	3.01	1.12	-0.39	-0.46	1.37	-2.74	1.17	3.73	0.22
StDev	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

Table 2
Summary statistics for correlations and correlation risk premia

This table reports means and standard deviations for realized and implied correlations, as well as correlation risk premia for all FX pairs. Correlation risk premia (CRP) are defined as the difference between the implied and realized correlations. Realized correlations (RC) are calculated using past daily log exchange rate changes over a three month window. Implied correlations (IC) are calculated from daily option prices on the underlying exchange rates. The last two columns show bootstrapped 95% confidence interval (using the 2.5 and 97.5 percentiles). Monthly data from January 1996 to December 2013 (options data for EUR start in January 1999).

		RC		IC		CRP				
FX Pair		Mean	StDev	Mean	StDev	Mean	StDev	t-stat	2.5%	97.5%
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AUD	CAD	0.471	0.25	0.430	0.27	-0.041	0.15	-4.07	-0.060	-0.023
AUD	CHF	0.357	0.27	0.405	0.20	0.048	0.15	4.73	0.028	0.068
AUD	EUR	0.450	0.28	0.544	0.16	0.019	0.09	2.81	0.006	0.031
AUD	GBP	0.422	0.24	0.453	0.19	0.031	0.12	3.86	0.014	0.046
AUD	JPY	0.155	0.34	0.238	0.26	0.083	0.16	7.58	0.062	0.103
AUD	NOK	0.467	0.26	0.431	0.29	-0.036	0.20	-2.64	-0.064	-0.010
AUD	NZD	0.755	0.16	0.739	0.15	-0.016	0.08	-2.97	-0.026	-0.005
AUD	SEK	0.474	0.25	0.480	0.20	0.005	0.13	0.61	-0.012	0.022
CAD	CHF	0.233	0.28	0.283	0.21	0.050	0.15	4.94	0.031	0.070
CAD	EUR	0.307	0.30	0.405	0.19	0.024	0.13	2.45	0.005	0.044
CAD	GBP	0.281	0.27	0.307	0.23	0.025	0.15	2.34	0.004	0.044
CAD	JPY	0.054	0.26	0.136	0.19	0.082	0.16	7.33	0.060	0.104
CAD	NOK	0.340	0.28	0.341	0.28	-0.002	0.18	-0.17	-0.028	0.022
CAD	NZD	0.413	0.23	0.352	0.34	-0.061	0.22	-4.19	-0.092	-0.035
CAD	SEK	0.352	0.26	0.287	0.29	-0.069	0.17	-5.96	-0.094	-0.047
CHF	EUR	0.888	0.13	0.875	0.12	-0.010	0.08	-1.69	-0.020	0.002
CHF	GBP	0.580	0.19	0.605	0.15	0.025	0.11	3.32	0.010	0.039
CHF	JPY	0.405	0.26	0.456	0.18	0.051	0.14	5.15	0.032	0.070
CHF	NOK	0.726	0.16	0.731	0.12	0.006	0.11	0.73	-0.009	0.021
CHF	NZD	0.358	0.23	0.370	0.20	0.012	0.16	1.06	-0.010	0.033
CHF	SEK	0.707	0.16	0.712	0.13	0.004	0.10	0.58	-0.010	0.017
EUR	GBP	0.644	0.15	0.683	0.10	0.003	0.08	0.54	-0.009	0.015
EUR	JPY	0.324	0.27	0.364	0.20	0.067	0.15	5.84	0.046	0.089
EUR	NOK	0.825	0.09	0.798	0.07	-0.025	0.06	-5.20	-0.035	-0.016
EUR	NZD	0.440	0.23	0.501	0.17	0.005	0.12	0.55	-0.013	0.022
EUR	SEK	0.816	0.11	0.817	0.08	-0.022	0.06	-4.64	-0.031	-0.012
GBP	JPY	0.217	0.26	0.293	0.19	0.076	0.15	7.29	0.056	0.095
GBP	NOK	0.577	0.16	0.638	0.12	0.059	0.16	5.39	0.038	0.080
GBP	NZD	0.415	0.23	0.404	0.22	-0.011	0.14	-1.15	-0.029	0.006
GBP	SEK	0.560	0.16	0.598	0.13	0.037	0.13	4.26	0.021	0.053
JPY	NOK	0.248	0.26	0.347	0.21	0.099	0.16	9.22	0.079	0.119
JPY	NZD	0.146	0.32	0.233	0.24	0.087	0.18	7.09	0.063	0.111
JPY	SEK	0.241	0.27	0.294	0.20	0.052	0.16	4.95	0.033	0.072
NOK	NZD	0.449	0.22	0.413	0.27	-0.036	0.20	-2.65	-0.064	-0.011
NOK	SEK	0.796	0.10	0.780	0.11	-0.016	0.08	-2.93	-0.026	-0.006
NZD	SEK	0.439	0.23	0.403	0.27	-0.036	0.18	-2.89	-0.060	-0.013

Table 3
Cyclicality of realized correlations and correlation risk premia

This table reports unconditional correlations of realized correlations (columns (1) to (4)) and correlation risk premia (columns (5) to (8)) with the global equity volatility measure used in Lustig, Roussanov, and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin, and Venter (2015) (GFI), the CBOE VIX (VIX), and the TED spread (TED), respectively. Unconditional correlations are calculated using monthly data from January 1996 through December 2013.

		RC cyclicality				CRP cyclicality			
		$GVol$	GFI	TED	VIX	$GVol$	GFI	TED	VIX
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AUD	CAD	0.174	-0.016	-0.081	0.168	-0.090	-0.203	-0.029	-0.180
AUD	CHF	-0.110	-0.342	-0.241	-0.180	0.068	0.116	0.024	0.062
AUD	EUR	0.100	-0.217	-0.079	0.008	0.040	0.007	-0.076	0.060
AUD	GBP	0.016	-0.207	-0.047	-0.102	0.004	0.062	-0.070	0.053
AUD	JPY	-0.328	-0.488	-0.365	-0.395	0.077	0.162	0.110	0.082
AUD	NOK	0.143	-0.145	-0.037	0.089	-0.096	-0.113	-0.328	-0.116
AUD	NZD	0.298	-0.125	0.014	0.287	-0.107	0.036	-0.016	-0.138
AUD	SEK	0.121	-0.161	-0.084	0.050	-0.141	-0.017	-0.115	-0.125
CAD	CHF	-0.099	-0.251	-0.223	-0.164	0.120	0.099	0.167	0.103
CAD	EUR	0.070	-0.133	-0.106	-0.009	-0.056	-0.014	0.076	-0.031
CAD	GBP	0.042	-0.060	-0.021	-0.041	0.090	-0.156	-0.150	0.066
CAD	JPY	-0.284	-0.405	-0.322	-0.383	0.050	0.097	0.065	0.063
CAD	NOK	0.102	-0.065	-0.063	0.053	-0.038	-0.151	-0.132	-0.043
CAD	NZD	0.166	-0.005	-0.060	0.174	0.084	-0.321	-0.182	-0.018
CAD	SEK	0.134	-0.025	-0.066	0.069	-0.078	-0.091	-0.187	-0.028
CHF	EUR	-0.221	-0.107	-0.030	-0.250	0.330	0.122	0.178	0.308
CHF	GBP	-0.159	-0.323	-0.256	-0.265	0.069	0.114	0.113	0.087
CHF	JPY	-0.146	-0.063	-0.028	-0.223	0.069	0.114	0.002	0.133
CHF	NOK	-0.269	-0.045	-0.130	-0.276	0.103	-0.019	0.098	0.130
CHF	NZD	-0.106	-0.241	-0.256	-0.114	0.142	-0.026	-0.031	0.084
CHF	SEK	-0.186	-0.221	-0.013	-0.265	0.037	-0.050	0.059	0.025
EUR	GBP	0.105	-0.155	-0.137	-0.018	-0.216	-0.137	-0.043	-0.184
EUR	JPY	-0.281	-0.178	-0.215	-0.301	0.173	0.228	0.190	0.208
EUR	NOK	-0.064	0.137	0.026	-0.056	-0.063	-0.062	0.032	-0.042
EUR	NZD	0.135	-0.106	-0.057	0.104	-0.002	-0.111	-0.205	-0.022
EUR	SEK	0.077	-0.169	0.077	-0.025	-0.177	-0.107	0.058	-0.186
GBP	JPY	-0.353	-0.412	-0.368	-0.433	0.158	0.213	0.149	0.166
GBP	NOK	0.026	-0.041	-0.118	-0.041	-0.038	-0.010	0.058	0.017
GBP	NZD	0.059	-0.099	0.000	-0.007	0.001	-0.196	-0.227	0.006
GBP	SEK	0.097	-0.163	-0.065	0.006	-0.211	0.013	-0.028	-0.128
JPY	NOK	-0.340	-0.219	-0.303	-0.354	0.199	0.212	0.262	0.226
JPY	NZD	-0.327	-0.361	-0.352	-0.317	0.064	0.077	0.129	0.008
JPY	SEK	-0.343	-0.314	-0.224	-0.399	0.224	0.256	0.121	0.253
NOK	NZD	0.163	-0.059	-0.028	0.161	-0.062	-0.179	-0.301	-0.101
NOK	SEK	0.156	0.030	0.141	0.144	-0.086	-0.022	-0.105	-0.047
NZD	SEK	0.171	-0.065	-0.054	0.144	-0.118	-0.154	-0.284	-0.154

Table 4
Cross-sectional cyclicality regressions

Panel A presents the output of the cross-sectional regressions of average realized correlations on each of the four FX correlation cyclicality measures. Panel B presents the output of the cross-sectional regressions of average correlation risk premia on each of the four CRP cyclicality measures. Each panel reports the slope coefficients, their t-statistic, their bootstrapped 95% confidence interval, as well as the regression R^2. Each FX correlation cyclicality measure (CRP cyclicality measure) is defined as the unconditional correlation of conditional FX correlation (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 (2015) (GFI), the CBOE VIX (VIX), and the TED spread (TED). The cyclicality measures are calculated using monthly data from January 1996 through December 2013 and are reported in Table 3. The t-statistics in parentheses are calculated using White (1980) standard errors.

	Slope	t-stat	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					
	Slope	t-stat	2.5%	97.5%	R^2
$GVol$	0.163	(2.63)	0.006	0.198	0.22
GFI	0.248	(8.98)	0.108	0.282	0.63
TED	0.203	(6.57)	0.073	0.262	0.48
VIX	0.198	(3.73)	0.064	0.232	0.33

Table 5
Correlations of FX correlation dispersion measures and market variables

This table reports the correlation coefficients between the FX correlation dispersion measures FXC and FXC^{UNC}, the global equity volatility measure used in Lustig, Roussanov, and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin, and Venter (2015) (GFI), the CBOE VIX (VIX), and the TED spread (TED). Monthly data from January 1996 through December 2013.

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

Table 6
FXC-beta-sorted portfolios

The table reports summary statistics for three G10 currency portfolios sorted on exposure to the correlation risk factor FXC. Exposure is measured by regressing currency excess returns on innovations in the correlation risk factor over the preceding 36 months. Portfolio 1 ($\mathrm{Pf1}^{\mathrm{C}}$) contains the three currencies with the lowest pre-sort FXC betas whereas Portfolio 3 ($\mathrm{Pf3}^{\mathrm{C}}$) contains the three currencies with the highest pre-sort FXC betas. HML^{C}, denotes the long-short portfolio that invests in the high correlation beta currencies ($\mathrm{Pf1}^{\mathrm{C}}$) and shorts the low correlation beta currencies ($\mathrm{Pf1}^{\mathrm{C}}$). Monthly data: for Panel A from January 1996 through December 2013, for Panel B from January 1984 through December 2013, for Panel C from January 1984 through July 2007 and for Panel D from January 1996 through July 2007.

	Panel A: January 1996–December 2013			
	$\mathrm{Pf1}^{\mathrm{C}}$	$\mathrm{Pf2}^{\mathrm{C}}$	$\mathrm{Pf3}^{\mathrm{C}}$	HML^{C}
Mean	4.04	0.99	-2.38	-6.42
StDev	10.26	9.11	7.86	7.83
t-stat	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 1984–December 2013			
	$\mathrm{Pf1}^{\mathrm{C}}$	$\mathrm{Pf2}^{\mathrm{C}}$	$\mathrm{Pf3}^{\mathrm{C}}$	HML^{C}
Mean	4.37	1.58	0.65	-3.72
StDev	9.62	9.44	8.87	8.37
t-stat	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 C: January 1984–July 2007			
	$\mathrm{Pf1}^{\mathrm{C}}$	$\mathrm{Pf2}^{\mathrm{C}}$	$\mathrm{Pf3}^{\mathrm{C}}$	HML^{C}
Mean	4.36	1.61	0.91	-3.45
StDev	8.00	9.05	9.00	8.02
t-stat	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
	Panel D: January 1996–July 2007			
	$\mathrm{Pf1}^{\mathrm{C}}$	$\mathrm{Pf2}^{\mathrm{C}}$	$\mathrm{Pf3}^{\mathrm{C}}$	HML^{C}
Mean	3.84	0.74	-3.51	-7.35
StDev	7.34	8.07	7.56	6.68
t-stat	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

Table 7
Estimating the price of correlation risk

This table reports the results for the estimation of the market price of correlation risk. Panel A reports factor betas and Newey and West (1987) standard errors (in parentheses) for the first stage regressions for various test assets. The test assets include three correlation portfolios (Pf^{C}) sorted based on exposure to the correlation risk factor FXC from Table 6, three interest rate sorted portfolios (Pf^{F}) that are constructed by sorting on the interest rate differential as well as the nine individual G10 currencies. Panel B reports the Fama and MacBeth (1973) factor prices and standard errors (in parentheses). Shanken (1992)-corrected standard errors are reported in brackets. Overall, we consider four sets of test assets. Set (1) considers only the three correlation and three carry portfolios from Panel A, while for set (2) we add the individual nine G10 currencies as test assets. For sets (3) and (4) we construct four correlation and interest rate sorted portfolios each but with an extended set of currencies to match those used in Lustig, Roussanov, and Verdelhan (2011) (further details are given in Section 1). Set (3) includes developed currencies only while the portfolios in set (4) are constructed using the all currencies. The first stage beta estimates for sets (3) and (4) are provided in the Online Appendix. Monthly data from January 1996 through December 2013.

Panel A: Factor betas

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

Panel B: Factor prices

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

Table 8
Time series correlations for RC, IC, and CRP

This table reports the time series correlations between realized correlations (RC) and implied correlations (IC), and between realized correlations and correlation risk premia (CRP) for all FX pairs. In addition to the means, we report t-statistics and 95% bootstrapped confidence intervals. Correlation risk premia (CRP) are defined as the difference between the implied and realized correlations. Implied correlations (IC) are calculated from daily option prices on the underlying exchange rates. Realized correlations (RC) are calculated using past daily log exchange rate changes over a three month window. Monthly data from January 1996 to December 2013 (options data for EUR start in January 1999).

		Correlation RC/IC				Correlation RC/CRP			
FX Pair		Mean	t-stat	2.5%	97.5%	Mean	t-stat	2.5%	97.5%
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AUD	CAD	0.843	22.88	0.800	0.875	-0.102	-1.49	-0.243	0.046
AUD	CHF	0.844	22.97	0.805	0.877	-0.695	-14.15	-0.756	-0.627
AUD	EUR	0.923	32.01	0.901	0.941	-0.714	-13.59	-0.782	-0.638
AUD	GBP	0.876	26.54	0.844	0.905	-0.656	-12.71	-0.732	-0.566
AUD	JPY	0.892	28.89	0.855	0.922	-0.695	-14.13	-0.764	-0.610
AUD	NOK	0.744	16.28	0.679	0.807	-0.213	-3.18	-0.317	-0.091
AUD	NZD	0.872	26.01	0.833	0.906	-0.457	-7.52	-0.646	-0.212
AUD	SEK	0.870	25.82	0.840	0.902	-0.618	-11.49	-0.723	-0.490
CAD	CHF	0.856	24.22	0.827	0.885	-0.684	-13.73	-0.756	-0.594
CAD	EUR	0.864	22.86	0.822	0.899	-0.702	-13.17	-0.785	-0.602
CAD	GBP	0.825	21.34	0.776	0.869	-0.518	-8.86	-0.640	-0.371
CAD	JPY	0.777	18.03	0.708	0.829	-0.680	-13.57	-0.737	-0.622
CAD	NOK	0.780	18.26	0.723	0.838	-0.316	-4.88	-0.465	-0.168
CAD	NZD	0.784	18.48	0.730	0.838	0.161	2.39	0.011	0.308
CAD	SEK	0.813	20.44	0.766	0.856	-0.137	-2.02	-0.241	-0.024
CHF	EUR	0.846	21.21	0.717	0.946	-0.603	-10.09	-0.743	-0.278
CHF	GBP	0.816	20.63	0.757	0.862	-0.640	-12.17	-0.715	-0.554
CHF	JPY	0.835	22.19	0.788	0.874	-0.733	-15.76	-0.785	-0.665
CHF	NOK	0.725	15.42	0.632	0.816	-0.671	-13.23	-0.763	-0.525
CHF	NZD	0.724	15.35	0.661	0.783	-0.532	-9.19	-0.619	-0.428
CHF	SEK	0.757	16.94	0.668	0.832	-0.560	-9.88	-0.683	-0.386
EUR	GBP	0.774	16.33	0.707	0.837	-0.592	-9.80	-0.697	-0.463
EUR	JPY	0.858	22.29	0.811	0.898	-0.760	-15.60	-0.813	-0.704
EUR	NOK	0.704	13.23	0.628	0.776	-0.632	-10.87	-0.773	-0.379
EUR	NZD	0.770	16.12	0.703	0.830	-0.467	-7.04	-0.597	-0.329
EUR	SEK	0.721	13.89	0.659	0.786	-0.549	-8.76	-0.697	-0.326
GBP	JPY	0.824	21.30	0.770	0.867	-0.713	-14.87	-0.778	-0.634
GBP	NOK	0.332	5.14	0.149	0.477	-0.710	-14.77	-0.771	-0.644
GBP	NZD	0.812	20.32	0.773	0.852	-0.350	-5.47	-0.498	-0.199
GBP	SEK	0.644	12.31	0.575	0.717	-0.615	-11.41	-0.747	-0.462
JPY	NOK	0.795	19.15	0.743	0.837	-0.572	-10.21	-0.657	-0.473
JPY	NZD	0.831	21.83	0.777	0.875	-0.680	-13.55	-0.746	-0.603
JPY	SEK	0.825	21.34	0.775	0.865	-0.699	-14.29	-0.762	-0.627
NOK	NZD	0.699	14.29	0.630	0.764	-0.157	-2.32	-0.267	-0.051
NOK	SEK	0.701	14.36	0.643	0.761	-0.347	-5.42	-0.521	-0.148
NZD	SEK	0.750	16.58	0.684	0.805	-0.158	-2.34	-0.253	-0.053

Table 9
Parameter values

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

SDF PARAMETERS							
α	χ	ϕ	κ	δ	γ^0	γ^{min}	γ^{max}
0.0067	0.89	0.06	0.04	2.78	0.36	0.22	0.49

PRICING FACTOR PARAMETERS					
λ	$\bar{z}$	ξ	λ^w	$\bar{z^w}$	ξ^w
0.09	0.0077	0.0322	0.01	0.0209	0.0132

INFLATION PARAMETERS		
$\bar{\pi}$	ζ	σ
-0.0031	0.25	0.0037^2

Table 10
Simulated moments: Interest rates, inflation, and exchange rates

This table reports empirical moments (first column) and simulated moments (second column) for the full model with $\rho = 0$ (identical local pricing factors). The first panel reports annualized means and standard deviations of real U.S. interest rates and cross-sectional averages of mean and standard deviation of foreign interest rates. The second panel reports cross-sectional averages of exchange rate volatility and autocorrelation. The third panel reports the average and standard deviation of U.S. inflation and cross-sectional averages of mean and standard deviation of foreign inflation. The fourth panel reports annualized means and standard deviations of nominal U.S. interest rates and cross-sectional averages of mean and standard deviation of foreign interest rates. The fifth panel reports cross-sectional averages of the volatility and autocorrelation of nominal exchange rates. The sixth panel reports the average excess return on the HML^{FX} factor and the last panel reports the cross-sectional average of the standard deviation of the real and nominal log SDF.

Moment	Data	Model
$E\left(r^{U.S.}\right)$	0.28%	0.24%
$Std\left(r^{U.S.}\right)$	1.35%	1.34%
$E_{\text{cross}}\left(E\left(r^{FGN}\right)\right)$	1.15%	0.30%
$E_{\text{cross}}\left(Std\left(r^{FGN}\right)\right)$	1.19%	1.35%
$E_{\text{cross}}\left(Std\left(\Delta q_{t+1}\right)\right)$	10.82%	9.41%
$E_{\text{cross}}\left(AC\left(\Delta q_{t+1}\right)\right)$	-0.01	0.00
$E\left(\pi^{U.S.}\right)$	2.32%	2.34%
$Std\left(\pi^{U.S.}\right)$	1.27%	1.70%
$E_{\text{cross}}\left(E\left(\pi^{FGN}\right)\right)$	1.56%	2.35%
$E_{\text{cross}}\left(Std\left(\pi^{FGN}\right)\right)$	1.12%	1.69%
$E\left(r^{\$,U.S.}\right)$	2.60%	2.58%
$Std\left(r^{\$,U.S.}\right)$	0.62%	1.36%
$E_{\text{cross}}\left(E\left(r^{\$,FGN}\right)\right)$	2.70%	2.64%
$E_{\text{cross}}\left(Std\left(r^{\$,FGN}\right)\right)$	0.44%	1.37%
$E_{\text{cross}}\left(Std\left(\Delta s_{t+1}\right)\right)$	10.76%	9.59%
$E_{\text{cross}}\left(AC\left(\Delta s_{t+1}\right)\right)$	0.01	0.00
$E\left(rx_{t+1}^{HML}\right)$	5.41%	2.62%
$E_{\text{cross}}\left(Std(m_{t+1})\right)$	-	0.59
$E_{\text{cross}}\left(Std(m_{t+1}^{\$})\right)$	-	0.60

Table 11
Simulated moments: FX correlations and risk premia

This table reports empirical moments (first column) and simulated moments for the full model with $\rho = 0$ (identical local pricing factors) and for the full model with $\rho = 1$ (independent local pricing factors) (second and third column, respectively). All moments refer to nominal exchange rates. The first panel reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average realized FX correlations, respectively. The second panel reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average implied FX correlations. The third panel reports the cross-sectional mean and the 2.5 and 97.5 percentiles of average FX CRP. The fourth panel reports the cross-sectional correlation between average realized and average implied FX correlation and the cross-sectional correlation between average realized correlation and average CRP. The fifth panel reports the cross-sectional average of the correlation between realized and implied FX correlation and the cross-sectional mean as well as the 2.5 and 97.5 percentiles of the correlation between realized correlation and CRP.

MOMENT	DATA	MODEL	
		IDENTICAL z	INDEPENDENT z
$2.5\%_{\text{cross}}\,(E(RC))$	0.13	0.01	0.29
$E_{\text{cross}}\,(E(RC))$	0.45	0.39	0.39
$97.5\%_{\text{cross}}\,(E(RC))$	0.83	0.65	0.45
$2.5\%_{\text{cross}}\,(E(IC))$	0.22	0.02	0.29
$E_{\text{cross}}\,(E(IC))$	0.48	0.39	0.39
$97.5\%_{\text{cross}}\,(E(IC))$	0.82	0.64	0.45
$2.5\%_{\text{cross}}\,(CRP)$	-6.24%	-0.15%	0.01%
$E_{\text{cross}}\,(CRP)$	1.58%	0.13%	0.03%
$97.5\%_{\text{cross}}\,(CRP)$	8.85%	0.49%	0.04%
$corr_{\text{cross}}\,(E(RC), E(IC))$	0.98	1.00	1.00
$corr_{\text{cross}}\,(E(RC), E(CRP))$	-0.55	-0.99	0.95
$E_{\text{cross}}\,(corr(RC, IC))$	0.79	1.00	1.00
$2.5\%_{\text{cross}}\,(corr(RC, CRP))$	-0.74	-0.83	-0.08
$E_{\text{cross}}\,(corr(RC, CRP))$	-0.52	-0.68	-0.04
$97.5\%_{\text{cross}}\,(corr(RC, CRP))$	-0.07	-0.35	-0.02

Appendix F Figures

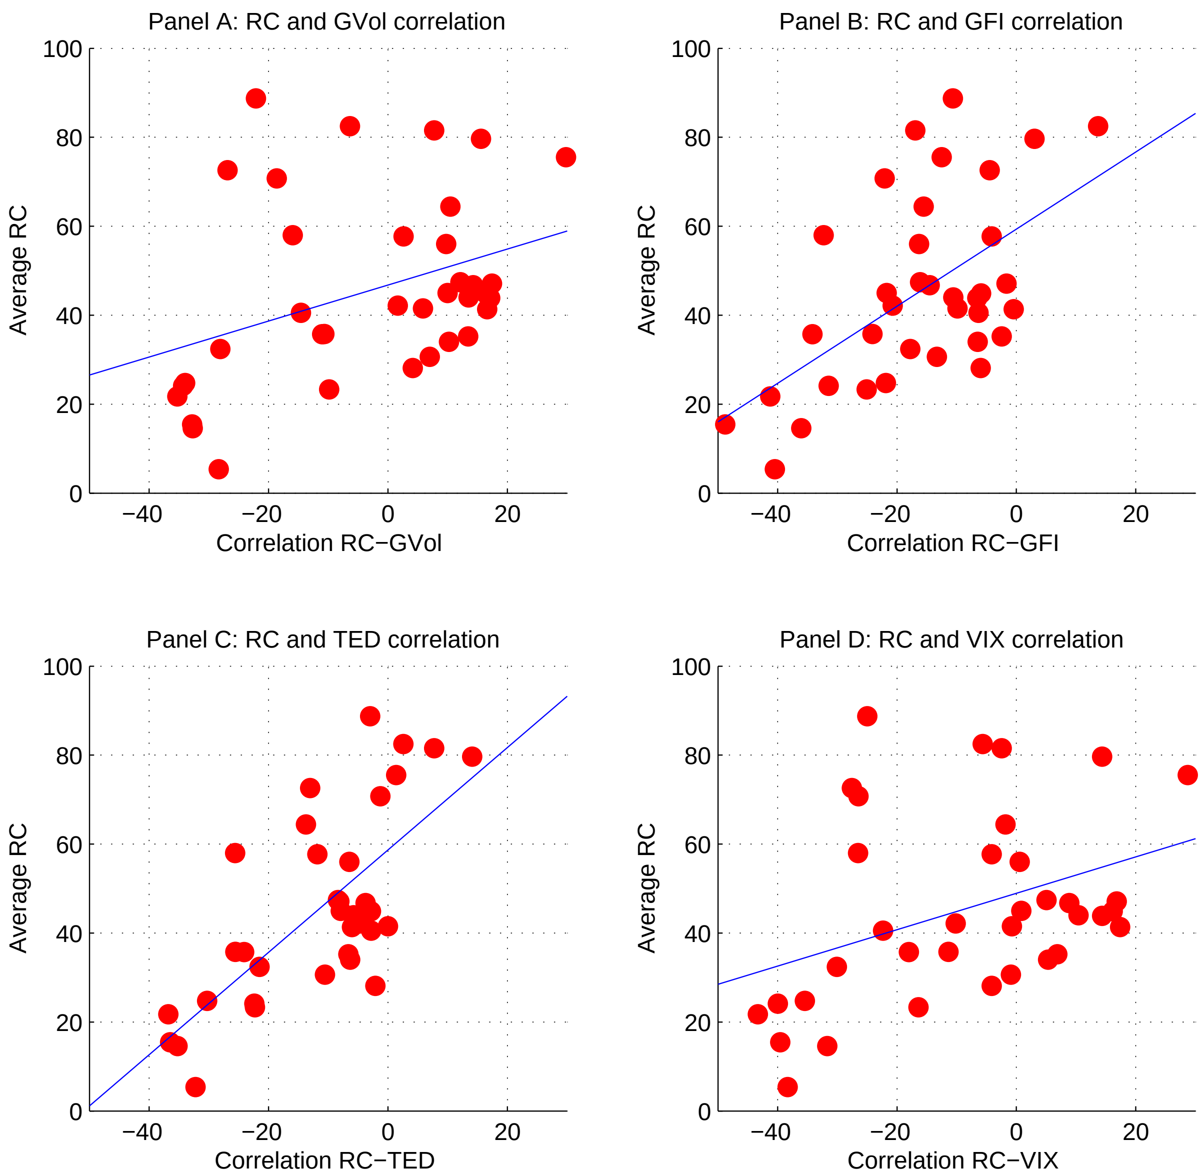

Figure 1. Realized correlations and cyclicality

This figure illustrates the relationship between measures of cyclicality of FX correlations and average realized correlations. Cyclicality is measured by the correlation between the realized correlation time series for a FX pair and a business cycle proxy. The proxies 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 (2015), the TED spread (TED, Panel C), and the CBOE VIX (VIX, Panel D). Monthly data from January 1996 to December 2013.

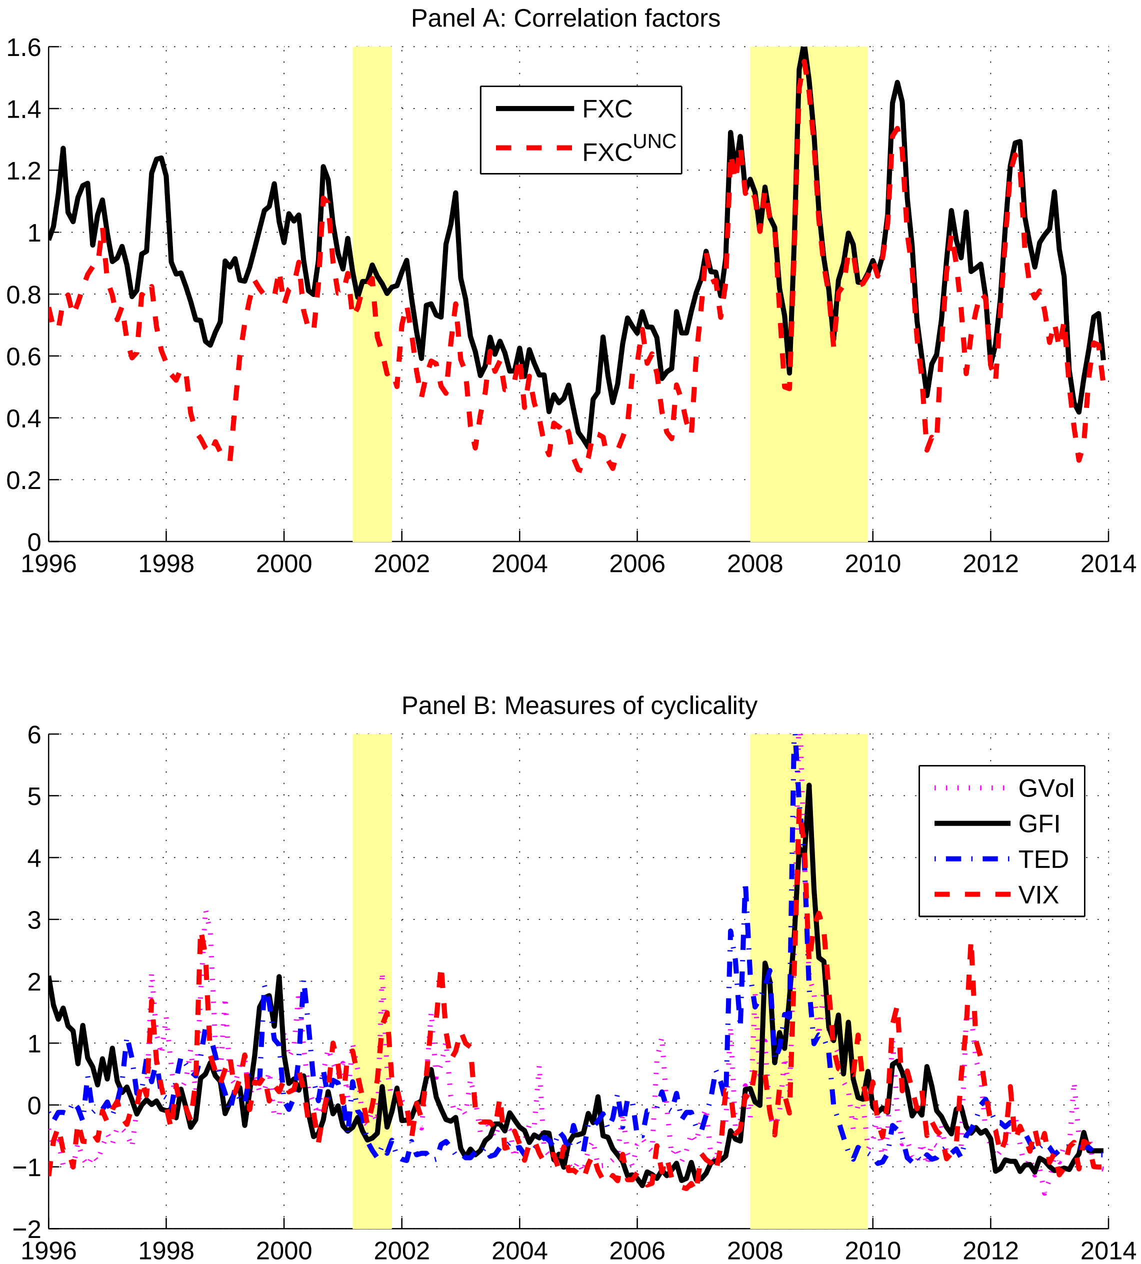

Figure 2. FX correlation dispersion measures and market variables

Panel A plots the FX correlation dispersion measure FXC calculated as the difference between the average high and low correlation FX pairs (solid line). The two groups consist of the highest and lowest decile of realized correlations across all 36 G10 FX pairs. The deciles are rebalanced every month. FXC is calculated for the period from January 1996 to December 2013. The alternative dispersion measure FXC^{UNC} is calculated as the difference of correlations between the decile of high correlation pairs and the decile of low correlation pairs measured over the whole sample period. Panel B plots the global equity volatility measure used in Lustig, Roussanov, and Verdelhan (2011) ($GVol$), the global funding illiquidity measure of Malkhozov, Mueller, Vedolin, and Venter (2015) (GFI), the CBOE VIX (VIX), and the TED spread (TED). All series are standardized to have zero mean and a standard deviation of one. The shaded areas depict NBER recessions.

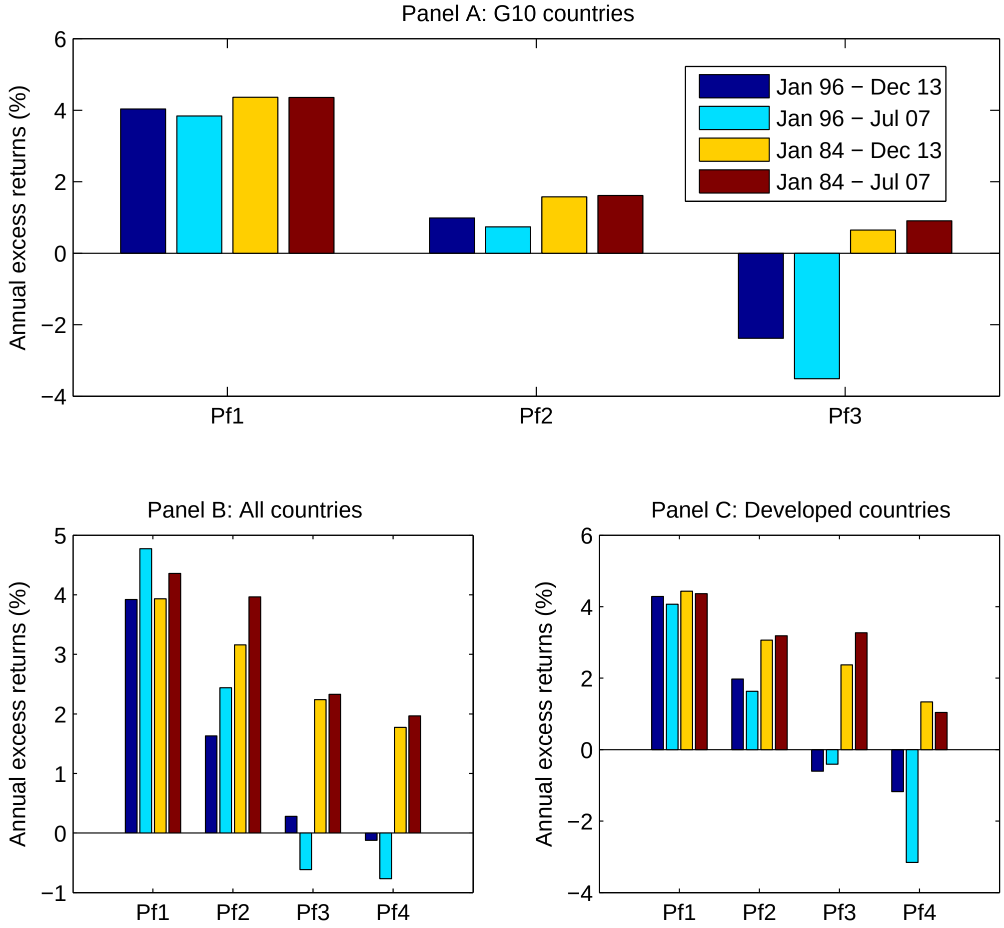

Figure 3. Portfolios sorted on exposure to correlation risk

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

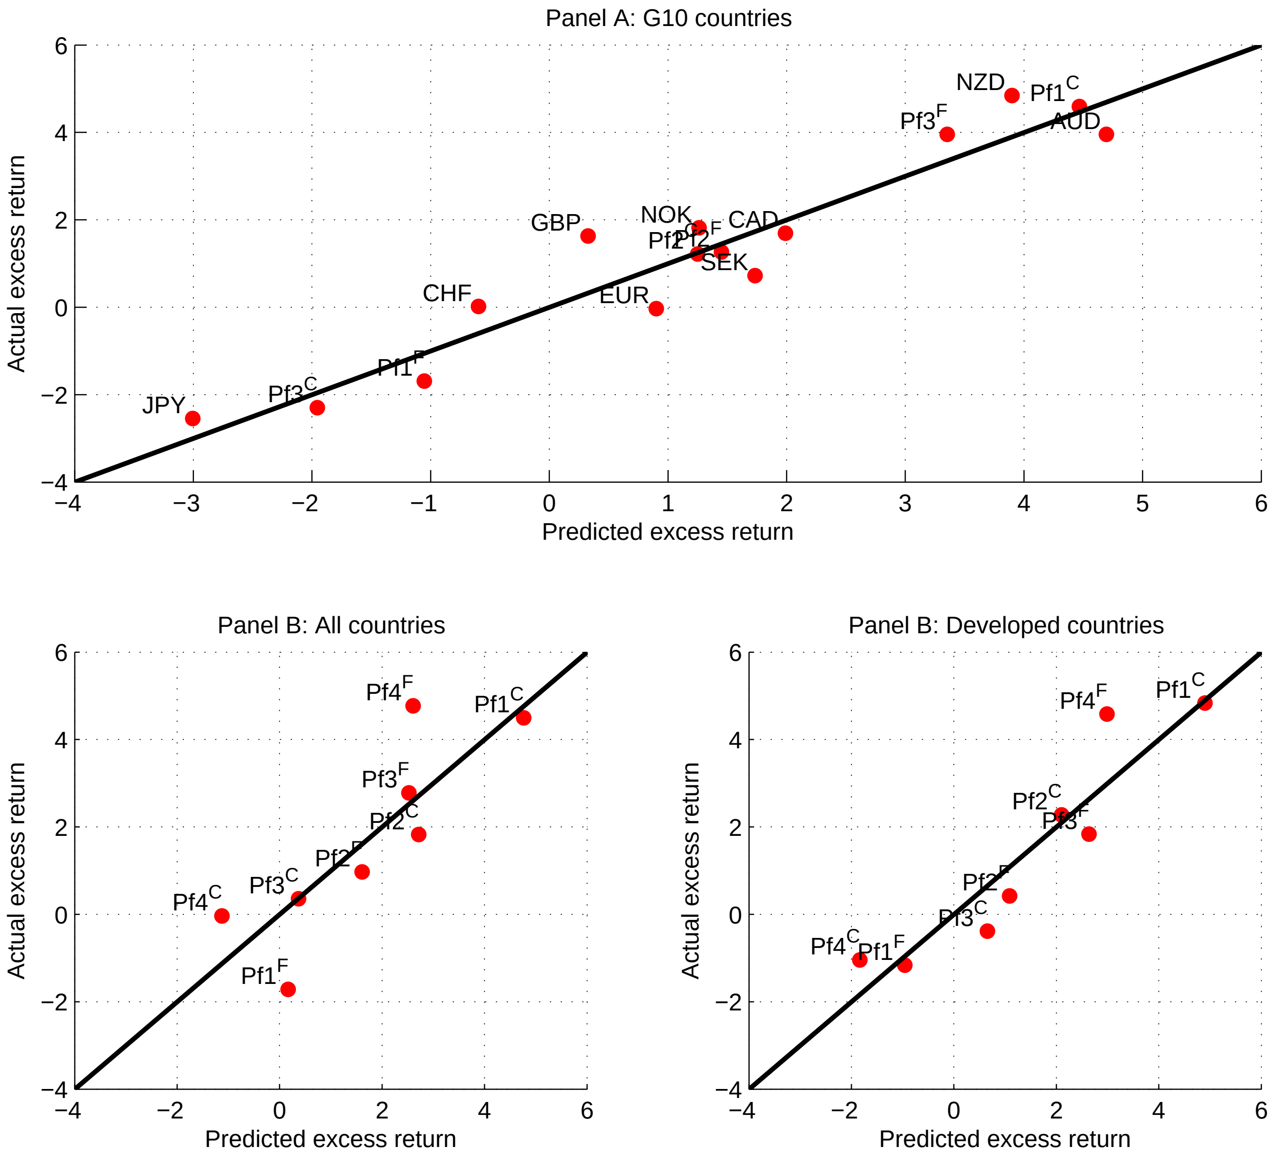

Figure 4. Model performance with various test assets

The figure plots the actual annualized mean excess returns in percent versus the predicted excess returns for various test assets using a linear pricing model that includes the dollar factor DOL and the HML^C correlation factor. Panel A displays the results for the nine G10 currencies and the three interest rate (Pf1^{F} to Pf3^{F}) and correlation (Pf1^{C} to Pf3^{C}) portfolios for the G10 currencies. Panels B and C display the results for the four interest rate (Pf1^{F} to Pf4^{F}) and correlation (Pf1^{C} to Pf4^{C}) portfolios constructed using all or only developed currencies, respectively. Monthly data from January 1996 to December 2013.

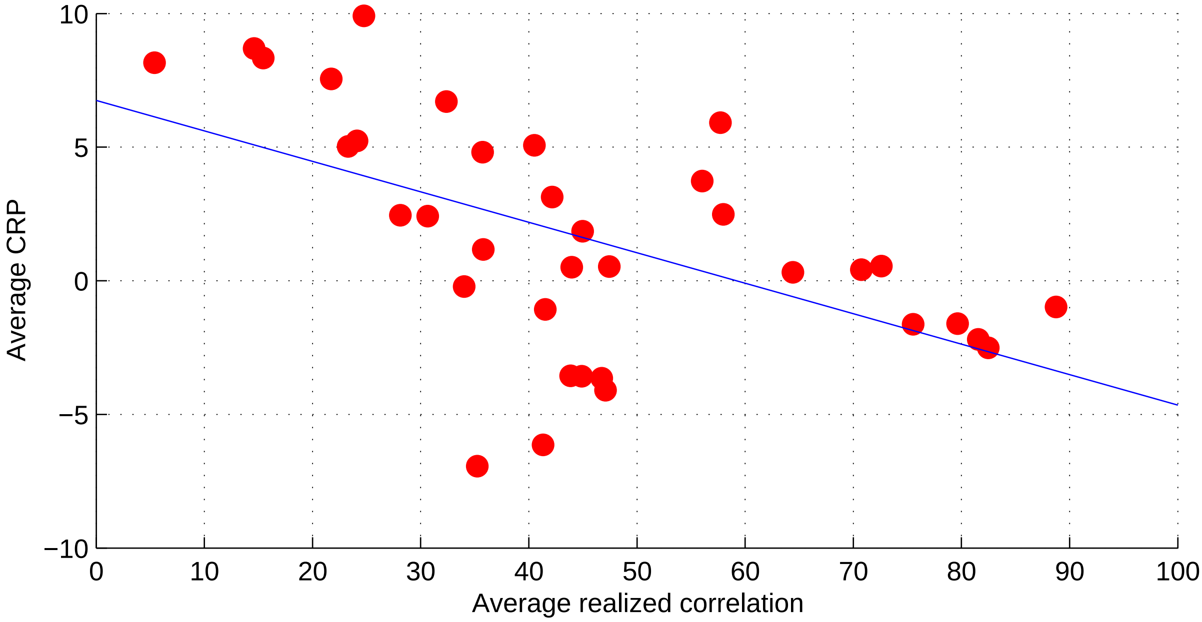

Figure 5. G10 realized correlation and correlation risk premia

This figure plots the average correlation risk premia for all 36 G10 exchange rate pairs against their average realized correlations. Correlation risk premia and correlations are expressed in percentage points. Monthly data from January 1996 (EUR since January 1999) to December 2013.

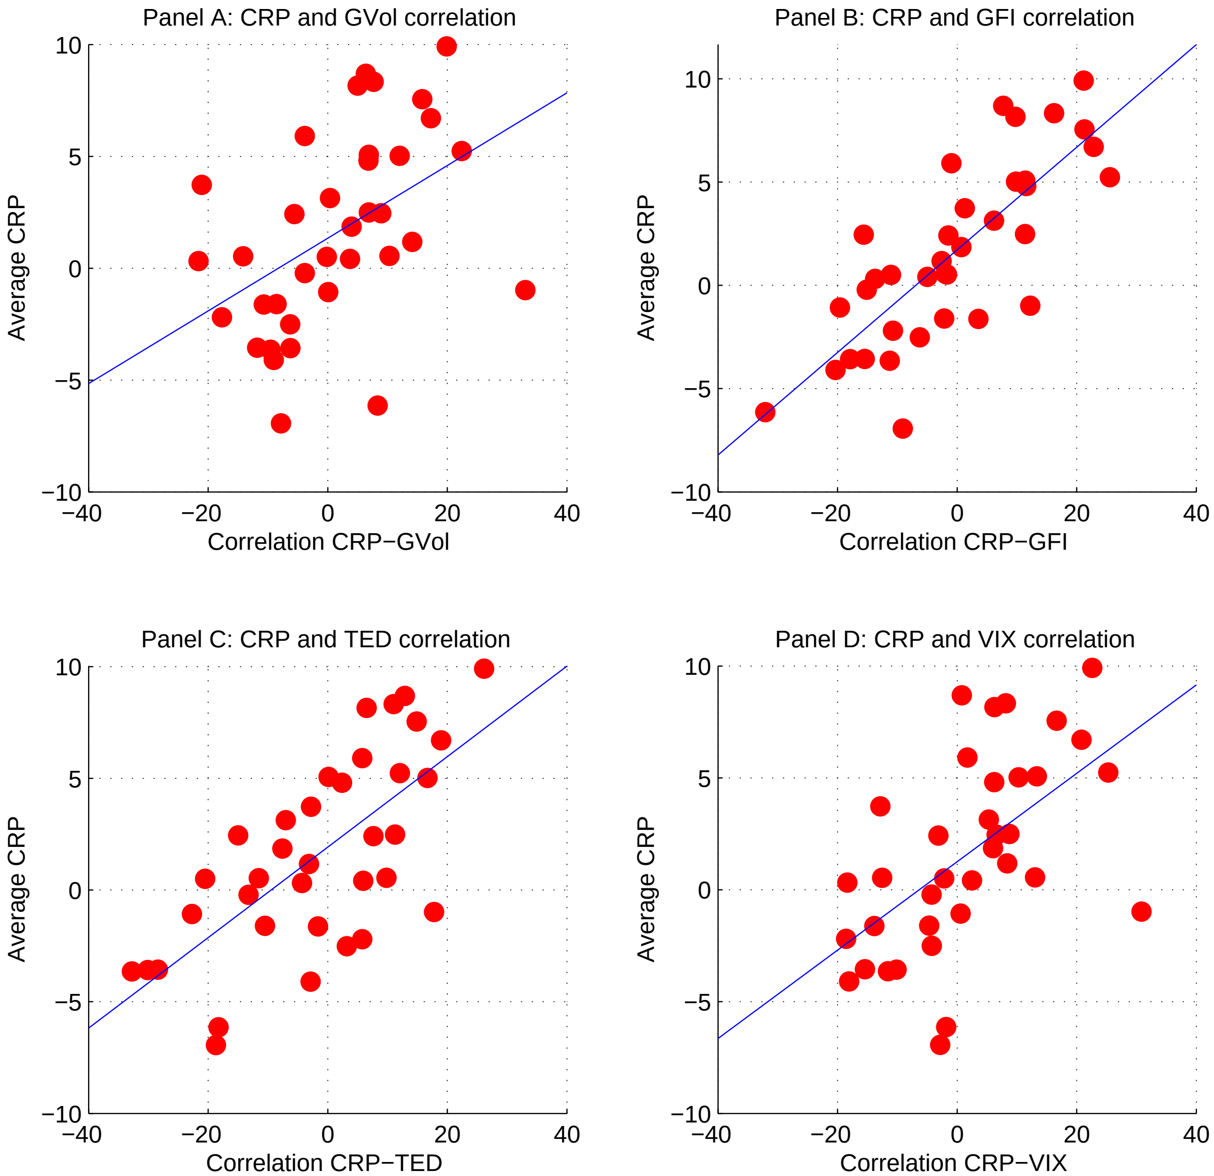

Figure 6. Correlation risk premia cyclicality and CRP

This figure illustrates the relationship between measures of cyclicality of correlation risk premia and average correlation risk premia. Cyclicality is measured by the correlation between the realized correlation (or correlation risk premia) time series for a FX pair and a business cycle proxy. The proxies 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 (2015), the TED spread (TED, Panel C), and the CBOE VIX (VIX, Panel D). Monthly data from January 1996 to December 2013.

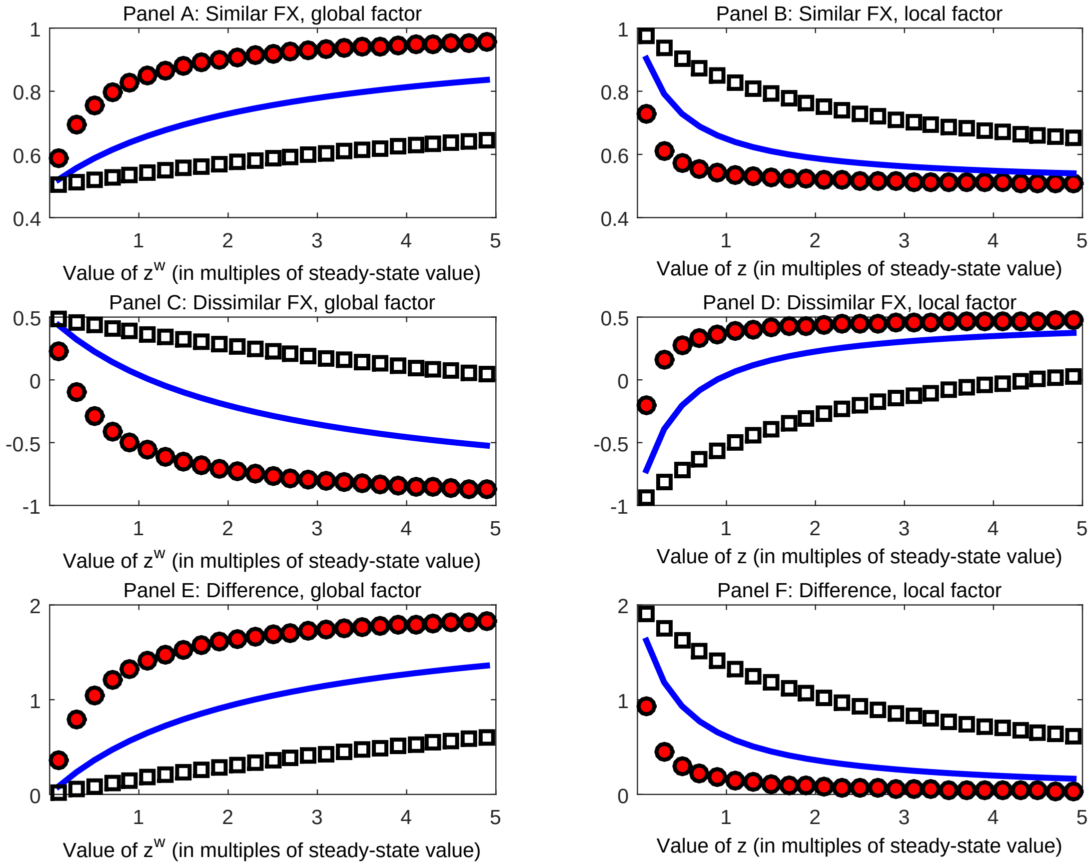

Figure 7. Model-implied FX correlations: identical local pricing factors

The figure displays the properties of conditional real FX correlation in the model with identical local pricing factors. Panels A, C and E plot conditional FX correlation as a function of the global pricing factor z^w, holding the local pricing factor z constant: Panel A refers to the conditional FX correlation of the similar FX pair (1,2), Panel C refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel E refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the local pricing factor z being equal to 0.2, 1, and 5 times its steady-state value $\bar{z}$, respectively. Panels B, D and F plot conditional FX correlation as a function of the local pricing factor z, holding the global pricing factor z^w constant: Panel B refers to the conditional FX correlation of the similar FX pair (1,2), Panel D refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel F refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the global pricing factor z^w being equal to 0.2, 1, and 5 times its steady-state value $\bar{z}^w$, respectively. To plot the figures, we set the model parameters equal to their calibrated values in Table 9.

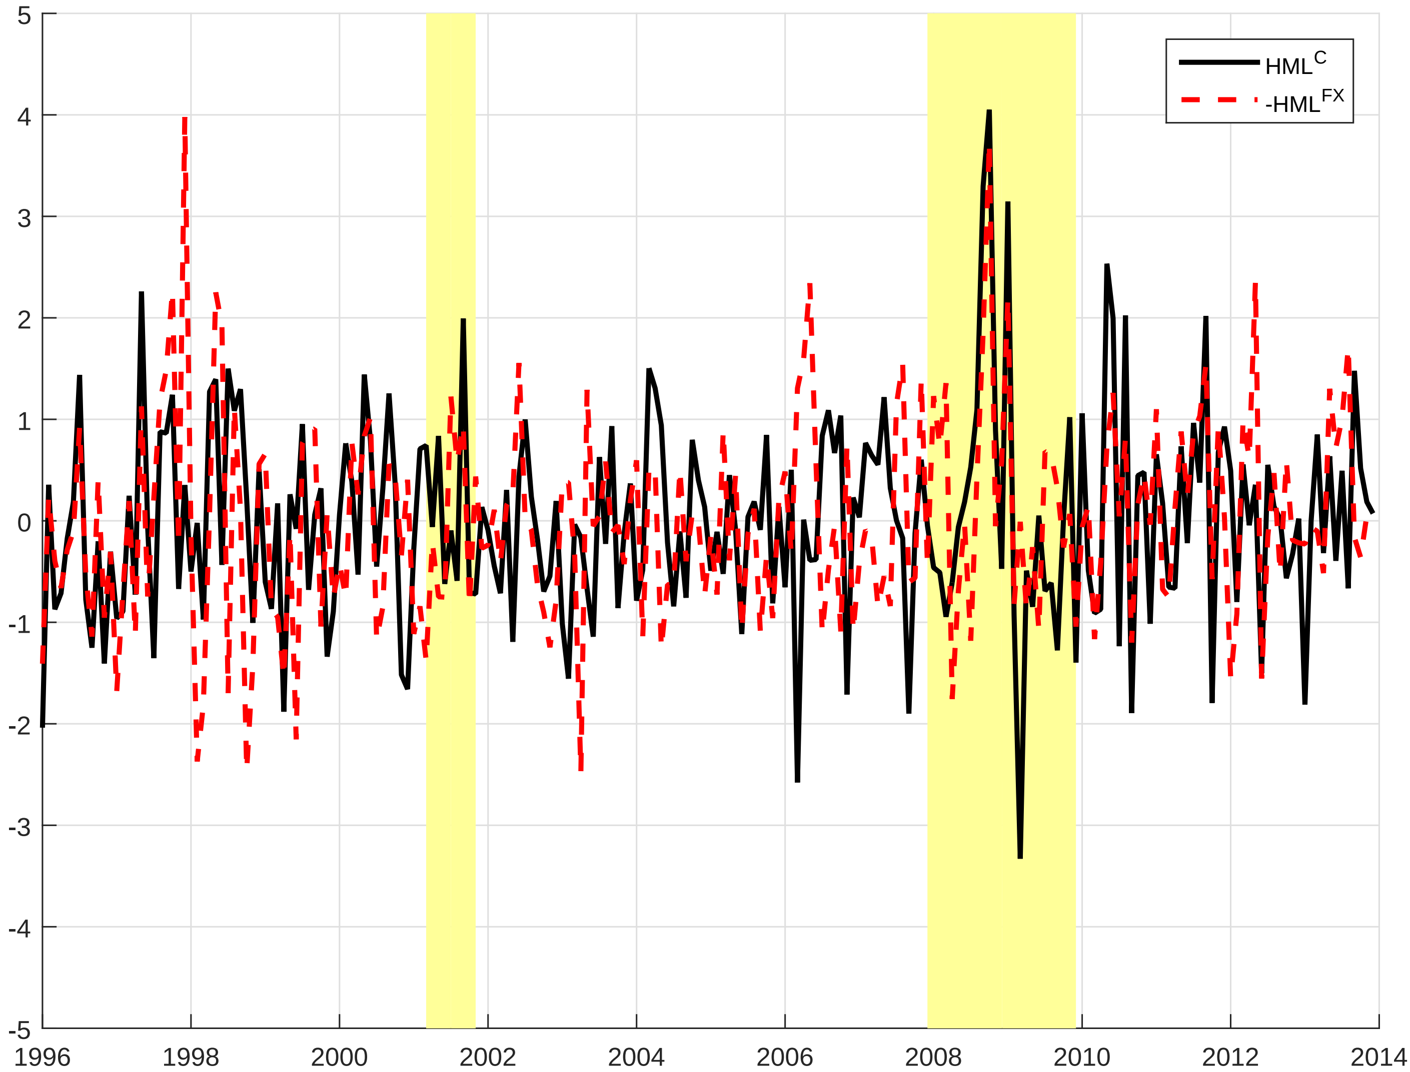

Figure 8. FX correlation and FX carry factors

The figure plots the return on HML^C, the long-short portfolio that invests in the high correlation beta currencies ($\mathrm{Pf1}^{\mathrm{C}}$) and shorts the low correlation beta currencies ($\mathrm{Pf1}^{\mathrm{C}}$) (solid line) and the (negative of) the FX carry factor HML^{FX} from Lustig, Roussanov, and Verdelhan (2011) (dashed line). All series are standardized to have zero mean and a standard deviation of one. The shaded areas depict NBER recessions. Monthly data from January 1996 to December 2013.

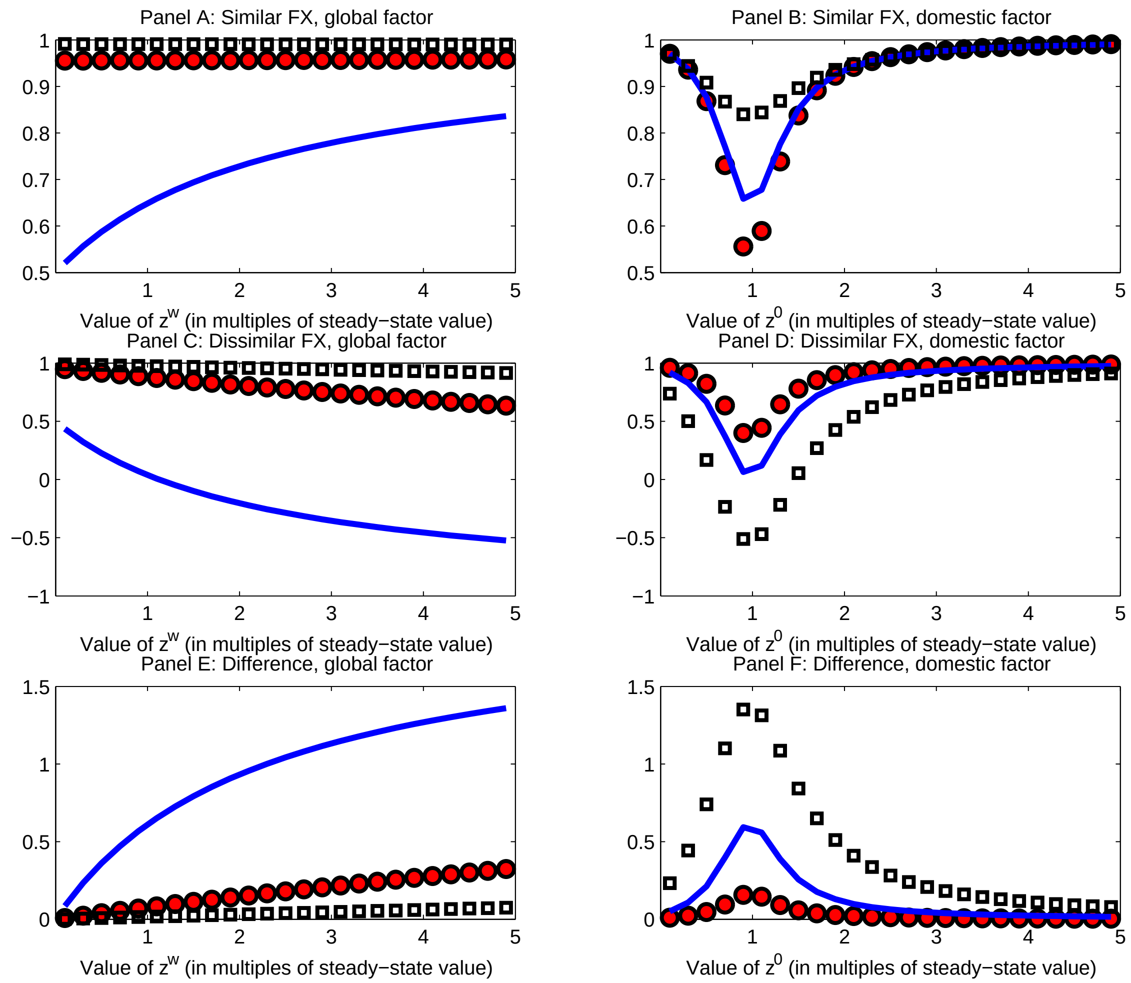

Figure 9. Model-implied FX correlations: independent local pricing factors

The figure displays the properties of conditional real FX correlation in the model with independent local pricing factors. Panels A, C and E plot conditional FX correlation as a function of the global pricing factor z^w, holding all the local pricing factors constant: Panel A refers to the conditional FX correlation of the similar FX pair (1,2), Panel C refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel E refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the domestic pricing factor z^0 being equal to 0.2, 1, and 5 times its steady-state value $\bar{z}$, respectively, and all the foreign pricing factors being equal to their common steady-state value $\bar{z}$. Panels B, D and F plot conditional FX correlation as a function of the domestic pricing factor z^0, holding the global pricing factor z^w constant: Panel B refers to the conditional FX correlation of the similar FX pair (1,2), Panel D refers to the conditional FX correlation of the dissimilar FX pair (1,3) and Panel F refers to difference in the conditional FX correlation between the two pairs. In each panel, the circles, solid line and squares plot the conditional FX correlation conditional on the global pricing factor z^w being equal to 0.2, 1, and 5 times its steady-state value $\bar{z}^w$, respectively, and all the foreign pricing factors being equal to their steady-state value $\bar{z}$. To plot the figures, we set the model parameters equal to their calibrated values in Table 9.