Online Appendix to
"International Correlation Risk"

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

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

Abstract

The Online Appendix contains additional results not included in the main paper. In Section OA-1, we show that our FX correlation dispersion measure FXC is very robust to different choices regarding its construction. Section OA-2 shows that our cross-sectional asset pricing tests are robust to using the non-traded ΔFXC factor instead of the traded HML^C factor and to using different sample periods and sample currencies. In Section OA-3, we confirm that our cross-sectional results with respect to the FX correlation risk premiums are robust to alternative construction methods. In Section OA-4, we discuss the effect of jumps on implied FX correlation measures. Section OA-5 presents summary statistics for FX variances and FX variance risk premiums. Section OA-6 shows that sorting on exposure to the FX correlation risk factor is not subsumed by exposure to an FX variance risk factor. Finally, Section OA-7 explores the spanning properties of FX correlation risk.

OA-1. Alternative construction of FXC

In the paper, we document a strong cross-sectional association between the average level and the cyclicality of conditional FX correlation: high average correlation FX pairs become even more correlated during bad times, whereas the correlation of low correlation FX pairs falls. This empirical result motivates the use of our FX correlation dispersion measure FXC, which is defined as the difference in average conditional FX correlation between the top and the bottom deciles of FX pairs sorted on their conditional correlation.

[Insert Figure OA-1 here.]

Instead of using deciles, we can alternatively construct the measure as the difference in correlations between the top and bottom quintiles or quartiles ($FXC^{Quintile}$ and $FXC^{Quartile}$, respectively). Increasing the size of the top and bottom groups reduces the correlation spread and, thus, the average level of the dispersion measure. In Panel A of Figure OA-1, we plot the original measure, along with the two alternative measures. Both alternative measures are very highly correlated with the original measure, both in levels (0.99 and 0.98, respectively, for $FXC^{Quintile}$ and $FXC^{Quartile}$) and in first differences (0.95 and 0.94, respectively). Not surprisingly, the portfolio results are very robust to using the alternative measures: Panels B and C present three G10 currency portfolios sorted on the alternative ΔFXC betas for various sample periods. The results for both alternative measures are qualitatively the same as those for the original measure. In particular, for all subsamples, excess currency returns are decreasing in the ΔFXC betas and the spread between the high and the low ΔFXC beta portfolio is substantial and statistically significant.

OA-2. Cross-sectional asset pricing

In Table 10 of the main paper, we present estimates of the market price of FX correlation risk using various test assets. Table OA-1 contains the first-stage regression estimates for test asset sets (3) and (4) that are omitted in Table 10.

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

[Insert Table OA-1 here.]

Instead of using HML^C returns, a traded factor, we can also perform our asset pricing tests using the non-traded factor ΔFXC. Figure OA-2 illustrates the performance of the ΔFXC factor by plotting the predicted annualized excess returns for various test assets (G10 currencies in Panel A, currency portfolios using all currencies and developed country currencies, in Panels B and C, respectively) against the actual annualized mean excess returns. Compared with the results when using the traded factor HML^C, the second-stage regression R^2s are lower, but still high: 0.79, 0.78 and 0.58, respectively. Table OA-2 compares the estimates for the market prices of risk of two factors, HML^C returns and FXC innovations, for the samples of all currencies and developed country currencies, using eight test assets (four currency portfolios sorted on FX correlation betas and four currency portfolios sorted on nominal interest rates) in each case. We consider three sample periods for each specification: January 1996 to December 2013, January 1984 to December 2013, and January 1996 to July 2007. Our results are fairly robust across alternative sample periods and factor specifications. In particular, the market price of risk for the non-traded correlation factor ΔFXC is always negative and often significant. Compared to our benchmark sample period (January 1996 to December 2013), the results for the January 1984 to December 2013 sample are slightly weaker, while the price of risk estimates when we exclude the financial crisis (January 1996 to July 2007) are strongly significant. Overall, FXC innovations perform reasonably well in pricing the cross section of currency returns, albeit somewhat worse than our traded factor, HML^C returns.

[Insert Figure OA-2 and Table OA-2 here.]

OA-3. Alternative definitions for FX correlation risk premiums

In the paper, we show that average FX correlation risk premiums and average FX correlations are negatively related in the cross section of FX pairs. We measure FX correlation risk premiums using information available at time t: we calculate the conditional risk-neutral correlations using currency option prices observed at time t, while our proxy for conditional correlations under the physical measure is the average realized correlation over the three-month period ending at t, sampled daily. Instead of using a three-month window, we can instead proxy for the conditional FX correlations under the physical measure using an one-month window of past daily exchange rates (i.e., daily data from $t-1$ to t), or using an one-month window of future daily exchange rates (i.e., daily data from t to $t+1$). Figure OA-3 provides scatter plots of average FX correlation risk premiums against the average FX correlations constructed using each of those two alternative measures of physical measure conditional FX correlations: Panel A refers to measures constructed using data over the month ending at t, whereas Panel B refers measures constructed using data between t and $t+1$. As we can see, the negative cross-sectional association between average FX correlation risk premiums and average FX correlations is robust to alternative measures of conditional FX correlation under the physical measure.

[Insert Figure OA-3 here.]

OA-4. The effect of jumps on implied FX correlation

To construct the measures of model-free implied exchange rate moments, we follow the methodology of Britten-Jones and Neuberger (2000). They impose no-arbitrage conditions and show that the risk-neutral expected return variance of an underlying asset is fully specified by a continuum of call and put options on that asset, provided that the price of the underlying asset is a diffusion process. Given recent empirical evidence of priced jump risk in exchange rates (see, e.g., Chernov, Graveline, and Zviadadze (2016)), it is natural to ask whether that methodology remains valid when the underlying price process includes jumps. In the following, we address this question in two ways. First, we show that the Britten-Jones and Neuberger (2000) methodology is valid even in the presence of jumps, provided that the higher order moments of the jump distribution are not very large. Second, we also consider the approach of Martin (2016), who derives a measure of risk-neutral expected variance that is robust to jumps, and we show that the methodologies of Britten-Jones and Neuberger (2000) and Martin (2016) produce virtually identical implied FX correlation and FX correlation risk premium measures.

We first consider the case of the underlying price process being a diffusion. The proof follows Britten-Jones and Neuberger (2000), except for the fact that instead of the risk-neutral measure $\mathbb{Q}$, we make use of the forward measure $\mathbb{Q}_T$. We denote by F_t the price at time t of a forward contract on the underlying asset with maturity T. We assume that no-arbitrage holds and, hence, there exists a forward measure $\mathbb{Q}_T$ such that

$$\frac{dF_t}{F_t} = \sqrt{V_t} dW_t^{\mathbb{Q}_T},$$

where $W^{\mathbb{Q}_T}$ is a standard Brownian motion under the T-forward measure $\mathbb{Q}_T$ and V_t is the instantaneous variance of the return of the forward contract. Applying Itô's lemma, we have

$$d\ln F_t = -\frac{1}{2} V_t dt + \sqrt{V_t} dW_t^{\mathbb{Q}_T}.$$

Integrating over time and taking expectations, we get

$$E_0^{\mathbb{Q}_T}\left(\int_0^T V_t dt\right) = 2\left(\ln F_0 - E_0^{\mathbb{Q}_T}(\ln F_T)\right),$$

where the LHS is equivalent to $E_0^{\mathbb{Q}_T}\left(\int_0^T \left(\frac{dF_t}{F_t}\right)^2\right)$.

Consider an option on the same underlying asset that expires at time T and has strike price K, and denote its price by $C(K,T)$. Then,

$$\int_0^\infty \frac{C(K,T) - \max(0, F_0 - K)}{K^2} dK = \ln F_0 - E_0^{\mathbb{Q}_T}(\ln F_T). \tag{OA-4.1}$$

Integration by parts of the LHS yields

$$\begin{aligned}\int_0^\infty \frac{C(K,T) - \max(0, F_0 - K)}{K^2} dK &= -\left.\frac{C(K,T) - \max(0, F_0 - K)}{K}\right|_0^\infty \\ &\quad + \int_0^\infty \frac{\partial_K C(K,T) + 1_{F_0 > K}}{K} dK\end{aligned}$$

where $\partial_K C(K,T)$ is the partial derivative of the option price with respect to the strike price K. Note that the first term of the RHS of the equation disappears assuming that the density of the forward price distribution is bounded. Since

$$\partial_K C(K,T) = E_0^{\mathbb{Q}_T}\left(\frac{\partial \max(F_T - K, 0)}{\partial K}\right) = -E_0^{\mathbb{Q}_T}(1_{F_T > K}),$$

we get that the second term of the RHS is equal to

$$\int_0^\infty \frac{\partial_K C(K,T) + 1_{F_0 > K}}{K} dK = \ln F_0 - E_0^{\mathbb{Q}_T}(\ln F_T).$$

Hence, we have shown that

$$E_0^{\mathbb{Q}_T}\left(\int_0^T \left(\frac{dF_t}{F_t}\right)^2\right) = 2\int_0^\infty \frac{C(K,T) - \max(F_0 - K, 0)}{K^2} dK. \tag{OA-4.2}$$

We use equation (OA-4.2) to calculate the implied FX variances used to construct our implied FX correlation measures.

Now, we can consider the impact of jumps. Under some regularity conditions, the forward price (which is a martingale) can be decomposed into two components: a purely continuous martingale and a discontinuous martingale.[1] In particular, let the forward price process now be

$$\frac{dF_t}{F_t} = \sqrt{V_t} dW_t^{\mathbb{Q}_T} + J_t dN_t - \mu_t \lambda_t dt,$$

[1] See for example Protter (1990), page 8. This is sometimes taken as a defining property of martingales; see, for example, Jacod and Shiryaev (1987).

where N_t is a pure jump process with time-varying intensity λ_t, J_t is the jump size with instantaneous mean μ_t, and the jump component is independent of the diffusion component. Applying Itô's lemma, we have

$$d\ln F_t = -\frac{1}{2}V_t dt + \sqrt{V_t} dW^{\mathbb{Q}_T} + \ln(1 + J_t)dN_t - \mu_t \lambda_t dt.$$

Using the approximation $\ln(1 + J_t) \approx J_t - \frac{1}{2}J_t^2$, we have

$$E_0^{\mathbb{Q}_T}\left(\int_0^T d\ln F_t\right) \approx -\frac{1}{2}E_0^{\mathbb{Q}_T}\left(\int_0^T \left(V_t + \lambda_t J_t^2\right) dt\right),$$

and

$$E_0^{\mathbb{Q}_T}\left(\int_0^T \left(V_t + \lambda_t J_t^2\right) dt\right) \approx 2\left(\ln F_0 - E_0^{\mathbb{Q}_T}\left(\ln F_T\right)\right).$$

Since the LHS of this equation is the expected return variance, we have that

$$E_0^{\mathbb{Q}_T}\left(\int_0^T \left(\frac{dF_t}{F_t}\right)^2\right) \approx 2\left(\ln F_0 - E_0^{\mathbb{Q}_T}\left(\ln F_T\right)\right). \quad \text{(OA-4.3)}$$

Note that equation (OA-4.1) holds even in the presence of jumps because we do not make any use of the return process itself. Hence, if we combine equation (OA-4.1) with equation (OA-4.3), we note that the Britten-Jones and Neuberger (2000) methodology is valid even in the presence of jumps.

One might be worried that equation (OA-4.3) does not hold exactly; the approximation error arises from the higher order terms in the expansion of $\ln(1 + J_t)$. One way to address this issue is to follow the approach of Martin (2016), who considers "simple variance swaps", which are robust to jumps in the underlying price process. To that end, we first compute implied FX variances using the Martin (2016) approach and use them to calculate implied FX correlations and FX correlation risk premiums. Then, we compare the series of implied FX correlations and FX correlation risk premiums obtained using the Martin (2016) approach with those obtained using the Britten-Jones and Neuberger (2000) approach.

Table OA-3 presents the results of that comparison. We first consider implied FX correlations. For each FX pair, we report the average difference in implied FX correlations between the two methodologies; note that, overall, the average differences are extremely small, ranging from −0.004 to 0.003. To formally assess whether those average differences are statistically different from zero, we perform an F-test of equal implied FX correlation averages between the two methodologies for each FX pair and report the corresponding p-value. We find that we cannot reject the null of zero difference at any conventional significance level for any of the 36 FX pairs. For each FX pair, we also calculate (but not report) the correlation between the two implied FX correlation time series: the two methods generate almost perfectly correlated series, with the lowest correlation of the 36 being 0.9994. Finally, we perform the same exercise for FX correlation risk premiums. The two methodologies produce almost perfectly correlated CRP series (the lowest correlation of the 36 is 0.991), and, as the last column of Table OA-3 shows, we cannot reject the null of no difference in average CRP between the two methodologies for any FX pair.

[Insert Table OA-3 here.]

The aforementioned comparison suggests that our methodology is robust to the presence of jumps. However, it might be the case that, due to the rarity of jumps, using all the sample days for our analysis obfuscates the potential impact of jumps. To address that point, we also perform the F-test of equal average implied FX correlations and equal average FX correlation risk premiums only in the sub-sample of "jump days", defined as days characterized by jumps in exchange rates or in exchange rate volatility. In particular, we use the "jump days" identified in Chernov, Graveline, and Zviadadze (2016): they focus on four exchange rates (AUD/USD, CHF/USD, GBP/USD, and JPY/USD; henceforth, the "CGZ set") and estimate jumps both in the corresponding excess FX returns (essentially, spot exchange rates) and the conditional variances of excess FX returns (essentially, conditional FX variances) using data on spot exchange rates and currency options.[2]

[2]We thank the authors for providing us with the data.

Since Chernov, Graveline, and Zviadadze (2016) consider four USD exchange rates, we first focus on the 6 exchange rate pairs generated by those four USD exchange rates. For each FX pair, we consider three subsamples: the first subsample consists of all days in which either of the two exchange rates in the corresponding FX pair jumps (denoted by "FX jumps"), the second subsample of all days in which the volatility of either of the two exchange rates in the corresponding FX pair jumps ("Vol jumps"), and the third subsample of all days in which either of two exchange rates or either of two exchange rate volatilities associated with the corresponding FX pair jumps ("Any jumps"). We present the results in Table OA-4. In panel A, for each FX pair and each subsample, we report the number of observations in the subsample (number of "jump days" N) and the p-value of the F-test for equality in the average implied FX correlation between the two methodologies. We find that we cannot reject the null hypothesis of equal average implied FX correlations between the two methodologies for any FX pair and any subsample. In Panel B, we repeat the same exercise for FX correlation risk premiums and find that the null hypothesis of no difference in averages between methodologies cannot be rejected for any FX pair and any subsample.

[Insert Table OA-4 here.]

Finally, we construct three further subsamples as follows: the first subsample consists of all days in which any exchange rate in the CGZ set jumps, the second subsample consists of all days in which the volatility of any exchange rate in the CGZ set jumps, and the third sample consists of all days in which any exchange rate or any exchange rate volatility in the CGZ set jumps. The first subsample (denoted by "FX jumps") has $N = 37$ daily observations, the second subsample ("Vol jumps") has $N = 133$ observations and the third subsample ("Any jumps") has $N = 160$ observations. This sample construction methodology allows us to consider all 36 exchange rate pairs in our benchmark set. Thus, for each of the 36 FX pairs and for each of three subsamples, we calculate the subsample average of the implied correlation and of the CRP using the methodologies of Britten-Jones and Neuberger (2000) and Martin (2016), and perform an F-test for the equality of the corresponding averages between methodologies. We report the p-values of the F-statistics in Table OA-5. Again, the null hypotheses of equal average implied FX correlations and equal average FX correlation risk premiums between methodologies cannot be rejected for any FX pair and any subsample.[3]

[Insert Table OA-5 here.]

We conclude that our findings on implied FX correlations and FX correlation risk premiums appear to be robust to the presence of jumps in exchange rates or exchange rate volatility.

OA-5. FX variances and FX variance risk premiums

The summary statistics for the realized and implied FX variance of the G10 exchange rates, as well as for their variance risk premiums, are provided in Table OA-6. All measures are expressed in squared percentage points.

Despite the evidence for significant variance risk premiums in equity markets (see, for example, Bollerslev, Tauchen, and Zhou (2009)), the results for currency markets are mixed. On the one hand, the variance risk premiums for a number of exchange rates are either not statistically significant (AUD, CHF, NZD) or are marginally significant (NOK), while, on the other hand, the variance risk premiums for the CAD and the SEK are significant at the 5% level and the variance risk premiums for the EUR, GBP and JPY are significant at the 1% level. For all exchange rates with a statistically significant average variance risk premium, that average is positive. The average variance risk premium across all exchange rates is 0.54 (in squared monthly percentage points), which is smaller by a factor of more than ten compared with the equity variance risk premium, and by a factor of 4.5 compared to the Treasury variance risk

[3] For each FX pair and for each subsample, we again calculate the correlation between the implied FX correlation series obtained using the two methodologies, and find that the lowest correlation out of the 108 (36 FX pairs, with 3 subsamples each) is 0.9989. The corresponding number for the FX CRP series is 0.9983.

premium.[4] Furthermore, there is an almost even split between exchange rates that have left- and right-skewed variance risk premium distributions. The two most negatively skewed variance risk premium distributions are those of the AUD and the NZD exchange rates against the USD; notably, the AUD and the NZD are typically used as investment currencies in the carry trade. On the other hand, the EUR, the JPY and the GBP exchange rates against the USD have the most positively skewed variance risk premium distributions.

[Insert Table OA-6 here.]

OA-6. Double sorts on correlation and variance

Menkhoff, Sarno, Schmeling, and Schrimpf (2012) show that FX variance is priced in currency markets. In our model, the cross-sectional average of all FX conditional variances is

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

Therefore, FXV innovations satisfy

$$\Delta FXV_{t+1} - E_t\left(\Delta FXV_{t+1}\right) = -\frac{1}{N}\sum_i \left(\sqrt{\gamma^i} - \sqrt{\gamma^0}\right)^2 \xi^w \sqrt{z^w_t} u^w_{t+1} - 2\kappa\xi\sqrt{z_t} u^g_{t+1}.$$

FXV innovations have two components: a component that is perfectly negatively correlated with the first global shock u^w and a component that is perfectly negatively correlated with the second global shock u^g. The first component generates positive correlation between FXV innovations and FXC innovations: a negative u^w shock raises the global pricing factor z^w, increasing both FXV and FXC. On the other hand, the second component generates negative correlation between FXV innovations and FXC innovations: negative u^g shocks raise the local pricing factor z and, thus, increase the conditional variance of all FX pairs (and therefore FXV) due to country-specific risk, but they reduce the importance of heterogeneity in global risk exposure, tightening the cross section of FX correlation and, thus, decreasing FXC. As a result, the correlation between FXV innovations and FXC innovations depends on the relative strength of the two components.

In the data, the correlation between FXC and FCV innovations is 0.3, suggesting that the effect of the first component dominates. To show that sorting on exposure to the FX correlation factor is not subsumed by exposure to the FX variance factor, we double-sort currencies on exposure to those two factors. The small number of available currencies restricts us to a two-by-two sort, for a total of four portfolios: high and low exposure to the FX correlation factor and high and low exposure to the FX variance factor, with exposure being measured by the return beta with respect to FXC and FXV innovations, respectively. FXV innovations, denoted by ΔFXV, are defined as the first differences in FXV levels. We perform this exercise for two sets of currencies: the G10 sample and the full sample. The results are presented in Table OA-7.

Panel A presents the results sorting on a single dimension, ΔFXC beta or ΔFXV beta, forming two portfolios for each dimension. The top row of Panel A shows that the HML^C portfolio (high ΔFXC beta currencies minus low ΔFXC beta currencies, regardless their ΔFXV beta) yields a negative and statistically significant average excess return, in line with the results presented in the main paper. The HML^V portfolio, which is obtained by investing in high ΔFXV beta currencies and shorting low ΔFXV beta currencies, also yields negative average excess returns, which is consistent with the positive correlation between ΔFXC and ΔFXV. However, the HML^V returns are not statistically significant.

Panels B and C present the portfolio results when double-sorting on ΔFXC beta and ΔFXV beta, for the G10 currencies and the full set of currencies, respectively. Conditional on ΔFXV exposure, exposure on ΔFXC is negatively associated with returns for both sets of currencies and for both levels of ΔFXV exposure (high or low); however, the

[4] The corresponding numbers are calculated using the methodology reported in Choi, Mueller, and Vedolin (2016) for a slightly longer sample. The corresponding values are 6 and 2.4 for the equity and Treasury variance risk premiums, respectively.

return difference between high ΔFXC beta and low ΔFXC beta portfolios is statistically significant only conditioning on high ΔFXV exposure. On the other hand, conditioning on ΔFXC exposure, FX variance risk does not seem to be associated with currency returns, as none of the return differentials between high ΔFXV beta portfolios and low ΔFXV beta portfolios is statistically significant at the 5% level. Furthermore, most of the return differentials have a positive, rather than a negative, sign.

In summary, Table OA-7 confirms that our findings with respect to FX correlation risk are robust to accounting for exposure to FX variance risk.

[Insert Table OA-7 here.]

OA-7. Spanning

To explore the spanning properties of FX correlation risk, we regress ΔFXC, the FX correlation risk factor, on: all nine G10 exchange rate changes, all nine G10 currency excess returns, the changes in the implied FX correlations of all 36 G10 FX pairs, and the changes in the FX correlation risk premiums of all 36 G10 FX pairs. We report the regression R^2s and adjusted R^2s in the first four rows of Table OA-8. We find that changes in FX correlation dispersion are almost unspanned by FX changes and currency excess returns: FX changes or currency excess returns can only explain about 12% of the variation in ΔFXC. In comparison, ΔFXC appears to be well-spanned by changes in FX correlation and changes in FX correlation risk premiums: changes in implied FX correlation are able to explain more than three times as much (40%) as FX changes or currency excess returns, and changes FX correlation risk premiums almost four times as much (46%). The results for adjusted R^2s, which take into account the number of regressors, yield similar results: 7% for exchange rate changes and currency excess returns, compared to 24% for changes in implied FX correlation and 32% for changes in FX correlation risk premiums.

[Insert Table OA-8 here.]

For robustness, we also regress ΔFXC on the principal components (PCs) of exchange rate changes, currency excess returns, changes in implied FX correlations and changes in FX correlation risk premiums. We consider the first 6 PCs of exchange rate changes and currency excess returns, as they capture about 95% of their variation, and the first 6 PCs of changes in implied FX correlations and changes in FX correlation risk premiums. Since the first 6 PCs capture only about 72% (70%) of the variation in changes in implied FX correlations (changes in FX correlation risk premiums), we also consider 20 (21) PCs, which capture about 95% of the variation of changes in implied FX correlations (changes in FX correlation risk premiums). The results are reported in the last six rows of Table OA-8 and are consistent with our previous findings.

References

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

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

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

Choi, H., Mueller, P., Vedolin, A., 2016. Bond variance risk premiums, working Paper, London School of Economics.

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

Jacod, J., Shiryaev, A., 1987. Limit theorems for stochastic processes, springer-Verlag, Berlin.

Martin, I., 2016. What is the expected return on the market? forthcoming, Quarterly Journal of Economics.

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

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55 (3), 703–708.

Protter, P., 1990. Stochastic integration and differential equations: A new approach, springer, Berlin.

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

Table OA-1. Estimating the price of FX correlation risk: extended sample.
The table reports the results for the estimation of the market price of FX correlation risk. Panel A reports the estimates of the test asset factor betas and their Newey and West (1987) standard errors (in parentheses) in the first-stage regressions. The test assets used in the first-stage regressions include four currency portfolios sorted on exposure to the correlation risk factor ΔFXC (denoted by Pf^{C}) and four currency portfolios sorted on nominal interest rates (denoted by Pf^{F}); currency portfolios are constructed using either the full set of currencies or the set of developed country currencies, as described in Section 3 of the paper. Panel B reports the Fama and MacBeth (1973) estimates of factor prices and their standard errors (in parentheses). Shanken (1992)-corrected standard errors are reported in brackets. Monthly data from January 1996 through December 2013.

Panel A: Factor betas

	All countries				Developed countries			
	α	DOL	HML^C	R^2	α	DOL	HML^C	R^2
Pf1^{C}	0.07	1.17	-0.34	0.89	0.00	1.11	-0.52	0.88
	(0.06)	(0.03)	(0.03)		(0.08)	(0.04)	(0.04)	
Pf2^{C}	-0.02	1.05	-0.11	0.88	0.04	1.15	-0.06	0.74
	(0.06)	(0.04)	(0.03)		(0.10)	(0.06)	(0.05)	
Pf3^{C}	0.02	0.87	0.14	0.79	-0.06	1.15	0.17	0.80
	(0.06)	(0.04)	(0.03)		(0.09)	(0.06)	(0.04)	
Pf4^{C}	0.09	0.76	0.31	0.75	0.11	0.96	0.53	0.75
	(0.06)	(0.04)	(0.03)		(0.09)	(0.05)	(0.04)	
Pf1^{F}	-0.14	0.90	0.18	0.80	0.02	0.96	0.39	0.61
	(0.06)	(0.04)	(0.04)		(0.12)	(0.06)	(0.06)	
Pf2^{F}	-0.01	0.93	0.00	0.89	-0.03	1.01	0.07	0.75
	(0.05)	(0.03)	(0.02)		(0.07)	(0.07)	(0.05)	
Pf3^{F}	0.07	0.99	-0.10	0.88	-0.05	1.12	-0.15	0.80
	(0.05)	(0.04)	(0.03)		(0.08)	(0.05)	(0.05)	
Pf4^{F}	0.24	1.14	-0.07	0.79	0.15	1.31	-0.17	0.77
	(0.09)	(0.05)	(0.04)		(0.11)	(0.07)	(0.08)	

Panel B: Factor prices

All countries			Developed countries		
DOL	HML^C	R^2	DOL	HML^C	R^2
0.15	-0.67	0.81	0.13	-0.51	0.90
(0.14)	(0.22)		(0.15)	(0.17)	
[0.14]	[0.23]		[0.15]	[0.18]	

Table OA-2. The price of FX correlation risk.

The table reports estimates for the price of FX correlation risk. We consider two alternative FX correlation risk factors, HML^C returns, a traded factor (Panel A), and ΔFXC, a non-traded factor (Panel B). We focus on either the set of all currencies or the set of developed country currencies; in each case, the test assets are four currency portfolios sorted on exposure to FX correlation risk and four currency portfolios sorted on nominal interest rates. Standard errors are reported in parentheses. Observations are monthly, for three alternative periods: January 1996-December 2013, January 1984-December 2013, and January 1996-July 2007. Regression R^2s are also provided.

Panel A: Factor prices using HML^C

Sample	Countries	DOL		HML^C		R^2
January 1996–December 2013	All	0.15	(0.14)	-0.67	(0.22)	0.81
	Developed	0.13	(0.15)	-0.51	(0.17)	0.90
January 1984–December 2013	All	0.26	(0.11)	-0.27	(0.17)	0.85
	Developed	0.25	(0.12)	-0.34	(0.15)	0.94
January 1996–July 2007	All	0.15	(0.14)	-0.87	(0.28)	0.70
	Developed	0.09	(0.16)	-0.73	(0.19)	0.78

Panel B: Factor prices using ΔFXC

Sample	Countries	DOL		ΔFXC		R^2
January 1996–December 2013	All	0.17	(0.14)	-0.13	(0.04)	0.78
	Developed	0.14	(0.15)	-0.08	(0.04)	0.58
January 1984–December 2013	All	0.27	(0.11)	-0.10	(0.04)	0.89
	Developed	0.27	(0.12)	-0.12	(0.05)	0.96
January 1996–July 2007	All	0.16	(0.14)	-0.12	(0.03)	0.76
	Developed	0.14	(0.16)	-0.10	(0.03)	0.54

Table OA-3. Differences in implied FX correlation and FX correlation risk premiums.
The table presents time-series averages of the differences in estimates of implied FX correlations (Diff. in IC) and FX correlation risk premiums (Diff. in CRP) between the methodologies of Britten-Jones and Neuberger (2000) and Martin (2016), for each of the 36 exchange rate pairs in the sample. The table also reports the p-value of the corresponding F-statistic that tests for equality in average implied FX correlations or in average FX correlation risk premiums between the two methodologies. Monthly data from January 1996 to December 2013 (options data for EUR start in January 1999).

FX pair		Diff. in IC	p-value	Diff. in CRP	p-value
AUD	CAD	0.002	0.86	0.002	0.84
AUD	CHF	0.000	0.98	0.000	0.99
AUD	EUR	0.002	0.84	0.002	0.87
AUD	GBP	0.002	0.64	0.002	0.80
AUD	JPY	-0.003	0.50	-0.003	0.85
AUD	NOK	0.000	0.95	0.000	0.99
AUD	NZD	0.001	0.96	0.001	0.99
AUD	SEK	0.000	0.95	0.000	0.94
CAD	CHF	-0.001	0.89	-0.001	1.00
CAD	EUR	0.000	0.87	0.000	0.98
CAD	GBP	0.000	0.97	0.000	0.98
CAD	JPY	-0.003	0.47	-0.003	0.56
CAD	NOK	0.000	0.91	0.000	0.97
CAD	NZD	0.001	0.99	0.001	0.97
CAD	SEK	0.000	0.71	0.000	0.77
CHF	EUR	0.000	0.89	0.000	0.97
CHF	GBP	0.000	0.92	0.000	0.95
CHF	JPY	-0.001	0.56	-0.001	0.93
CHF	NOK	0.000	0.73	0.000	0.95
CHF	NZD	0.000	0.93	0.000	0.91
CHF	SEK	0.001	0.68	0.001	0.85
EUR	GBP	0.001	0.72	0.001	0.91
EUR	JPY	0.000	0.74	0.000	0.98
EUR	NOK	0.000	0.85	0.000	0.94
EUR	NZD	0.003	0.79	0.003	0.98
EUR	SEK	0.000	0.99	0.000	0.93
GBP	JPY	-0.002	0.46	-0.002	0.84
GBP	NOK	0.000	0.99	0.000	0.94
GBP	NZD	0.003	0.60	0.003	0.95
GBP	SEK	0.000	0.87	0.000	0.97
JPY	NOK	-0.003	0.12	-0.003	0.83
JPY	NZD	-0.003	0.49	-0.003	0.75
JPY	SEK	-0.004	0.16	-0.004	0.67
NOK	NZD	0.000	0.93	0.000	0.94
NOK	SEK	0.000	0.97	0.000	1.00
NZD	SEK	0.000	0.94	0.000	0.97

Table OA-4. Differences in implied correlation and correlation risk premiums between methodologies: FX pair jump days. The table reports the p-value of the F-statistic that tests for equality in average implied FX correlations (Panel A) or in average FX correlation risk premiums (Panel B) between the methodologies of Britten-Jones and Neuberger (2000) and Martin (2016), for each of the six exchange rate pairs in the sample. In each panel, we consider three subsamples: all days in which either of the two exchange rates in the corresponding FX pair jumps ("FX jumps"), all days in which the volatility of either of the two exchange rates in the corresponding FX pair jumps ("Vol jumps"), and all days in which either exchange rate or either exchange rate volatility associated with the corresponding FX pair jumps ("Any jumps"). In each case, the number of jump days (N) is also reported. Data from January 1996 to December 2013.

Panel A: Implied correlation

FX pair		FX jumps	Vol jumps	Any jumps
AUD	CHF	0.98	0.98	0.94
N		15	67	77
AUD	GBP	0.97	0.91	0.88
N		14	74	84
AUD	JPY	0.99	0.90	0.98
N		29	69	90
CHF	GBP	1.00	0.99	0.93
N		6	77	83
CHF	JPY	0.94	0.93	0.82
N		22	81	99
GBP	JPY	0.97	0.81	0.82
N		22	88	106

Panel B: Correlation risk premium

FX pair		FX jumps	Vol jumps	Any jumps
AUD	CHF	1.00	0.99	0.82
N		15	67	77
AUD	GBP	0.96	0.85	0.94
N		14	74	84
AUD	JPY	0.97	0.96	0.99
N		29	69	90
CHF	GBP	1.00	0.97	0.99
N		6	77	83
CHF	JPY	1.00	0.99	0.94
N		22	81	99
GBP	JPY	0.98	0.94	0.93
N		22	88	106

Table OA-5. Differences in implied correlation and correlation risk premiums between methodologies: All jump days. The table reports the p-value of the F-statistic that tests for equality in average implied FX correlations (Columns 2-4) or in average FX correlation risk premiums (Columns 5-7) between the methodologies of Britten-Jones and Neuberger (2000) and Martin (2016), for each of the 36 exchange rate pairs in the sample. Data from January 1996 to December 2013 (data for EUR start in January 1999). We consider three subsamples: all days in which any exchange rate in the Chernov, Graveline, and Zviadadze (2016) (CGZ) set jumps ("FX jumps", sample size $N = 37$ days), all days in which the volatility of any exchange rate in the CGZ set jumps ("Vol jumps", sample size $N = 133$ days), and all days in which any exchange rate or any exchange rate volatility in the CGZ set jumps ("Any jumps", sample size $N = 160$ days).

		Implied correlation			Correlation risk premium		
FX pair		FX jumps	Vol jumps	Any jumps	FX jumps	Vol jumps	Any jumps
AUD	CAD	0.98	0.95	0.95	0.98	0.95	0.95
AUD	CHF	0.98	0.97	0.97	0.99	0.98	0.98
AUD	EUR	0.97	0.89	0.89	0.98	0.85	0.87
AUD	GBP	0.99	0.87	0.89	0.98	0.82	0.84
AUD	JPY	0.99	0.85	0.84	0.97	0.92	0.92
AUD	NOK	1.00	1.00	1.00	1.00	0.89	0.91
AUD	NZD	1.00	0.99	0.99	1.00	0.97	0.98
AUD	SEK	1.00	1.00	1.00	0.99	0.96	0.96
CAD	CHF	1.00	0.95	0.95	0.98	0.95	0.95
CAD	EUR	1.00	0.98	0.98	1.00	1.00	1.00
CAD	GBP	1.00	1.00	1.00	1.00	0.99	0.99
CAD	JPY	0.99	0.70	0.73	0.98	0.89	0.88
CAD	NOK	1.00	0.99	1.00	1.00	0.99	0.99
CAD	NZD	1.00	1.00	1.00	1.00	1.00	1.00
CAD	SEK	0.98	0.92	0.92	0.96	0.85	0.86
CHF	EUR	0.99	0.96	0.97	0.98	0.96	0.96
CHF	GBP	1.00	0.98	0.98	0.99	0.98	0.98
CHF	JPY	0.95	0.90	0.89	1.00	0.99	0.99
CHF	NOK	0.99	0.90	0.90	0.99	0.97	0.98
CHF	NZD	0.99	0.97	0.97	1.00	1.00	1.00
CHF	SEK	0.99	0.88	0.88	0.99	0.94	0.93
EUR	GBP	0.96	0.96	0.94	0.99	0.89	0.89
EUR	JPY	0.99	0.90	0.91	0.97	0.96	0.96
EUR	NOK	1.00	0.93	0.93	1.00	1.00	1.00
EUR	NZD	0.97	0.90	0.90	0.98	0.97	0.96
EUR	SEK	0.99	0.93	0.93	0.99	0.98	0.98
GBP	JPY	0.96	0.79	0.79	0.97	0.94	0.93
GBP	NOK	0.99	0.99	0.99	0.99	1.00	0.99
GBP	NZD	0.97	0.86	0.87	0.95	0.97	0.98
GBP	SEK	0.98	0.95	0.95	0.98	0.97	0.96
JPY	NOK	0.91	0.68	0.67	0.88	0.89	0.89
JPY	NZD	0.96	0.84	0.83	0.98	0.93	0.92
JPY	SEK	0.91	0.67	0.65	0.97	0.92	0.90
NOK	NZD	1.00	1.00	0.99	1.00	0.99	0.99
NOK	SEK	0.99	0.96	0.97	0.99	0.99	0.98
NZD	SEK	1.00	1.00	1.00	0.99	0.98	0.98

Table OA-6. Summary statistics for variances and variance risk premiums.
The table reports summary statistics for realized and implied FX variances (Panels A and B, respectively) and FX variance risk premiums (Panel C). The FX variance risk premium is defined as the difference between the implied and realized FX variance. Realized FX variances are calculated using past daily log exchange rate changes over a three month window. Implied FX variances are calculated from daily option prices on the underlying exchange rates. FX variances and variance risk premiums are monthly and expressed in squared percentage points. Monthly data from January 1996 to December 2013 (options data for EUR start in January 1999).

Panel A: Realized variance									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	14.52	6.55	10.07	8.34	7.05	10.44	11.97	15.20	12.02
StDev	24.57	8.57	7.28	5.70	7.67	7.86	11.53	16.60	12.01
Skewness	6.51	4.48	3.88	3.49	4.69	3.05	3.78	4.48	3.97
Kurtosis	49.68	28.15	20.78	19.48	26.91	15.17	19.44	27.16	21.11
AC(1)	0.84	0.91	0.84	0.92	0.93	0.86	0.93	0.89	0.94
Panel B: Implied variance									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	13.32	6.97	10.49	10.54	7.71	11.94	12.73	14.84	12.97
StDev	14.63	7.65	5.62	7.87	7.33	8.59	10.05	11.97	10.52
Skewness	6.74	4.26	3.25	3.91	5.22	3.95	4.17	4.15	3.93
Kurtosis	65.86	26.34	19.89	24.87	37.06	28.87	26.29	29.47	22.95
AC(1)	0.69	0.86	0.73	0.77	0.80	0.68	0.80	0.74	0.82
Panel C: Variance risk premium									
	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Mean	-1.19	0.42	0.43	1.72	0.66	1.50	0.77	-0.36	0.95
StDev	14.74	3.16	5.35	4.58	3.77	5.81	5.79	9.09	6.47
t-stat	-1.19	1.97	1.17	5.04	2.57	3.78	1.94	-0.58	2.16
Skewness	-8.68	-1.32	-1.96	5.47	3.05	1.90	-0.19	-5.66	0.18
Kurtosis	84.82	34.11	20.07	54.18	45.93	32.81	19.86	47.94	30.79
AC(1)	0.60	0.45	0.58	0.45	0.46	0.33	0.57	0.52	0.54

Table OA-7. Portfolios sorted on correlation and variance exposure.
The table reports the properties of the returns of currency portfolios sorted on exposure to the FX correlation factor ΔFXC and the FX variance factor ΔFXV. For Panel A, we sort currencies into two portfolios based on their exposure to ΔFXC or ΔFXV; HML^C (HML^V) denotes a long-short portfolio that invests in the high ΔFXC (ΔFXV) beta currencies and shorts the low ΔFXC (ΔFXV) beta currencies. For Panels B and C, we double sort currencies into portfolios on their ΔFXC beta and their ΔFXV beta; Panel B reports the results for the G10 currencies, whereas Panel C reports the results for the full set of currencies. In the table, we report annualized USD excess returns, expressed in percentage terms. Monthly data from January 1996 through December 2013. Standard deviations are in parentheses, t-statistics are in brackets.

Panel A: HML^C and HML^V

	G10 countries			All countries		
	Mean	Std	t-stat	Mean	Std	t-stat
HML^C	−4.23	(6.21)	[−2.89]	−3.03	(6.11)	[−2.11]
HML^V	−1.76	(6.32)	[−1.18]	−0.95	(5.27)	[−0.76]

Panel B: G10 countries

		ΔFXC		
		Low beta	High beta	Difference
	Low beta	2.26	0.83	−1.43
	Std	(10.34)	(9.65)	(5.17)
	t-stat.	[0.93]	[0.36]	[−1.17]
ΔFXV	High beta	5.08	1.26	−3.82
	Std	(11.54)	(9.42)	(6.72)
	t-stat.	[1.87]	[0.57]	[−2.41]
	Difference	2.82	0.43	
	Std	(7.32)	(4.74)	
	t-stat.	[1.63]	[0.39]	

Panel C: All countries

		ΔFXC		
		Low beta	High beta	Difference
	Low beta	2.00	1.10	−0.90
	Std	(9.78)	(8.17)	(6.80)
	t-stat.	[0.87]	[0.57]	[−0.56]
ΔFXV	High beta	2.83	−0.86	−3.69
	Std	(9.49)	(5.86)	(7.14)
	t-stat.	[1.27]	[−0.62]	[−2.19]
	Difference	0.83	−1.96	
	Std	(5.45)	(6.39)	
	t-stat.	[0.65]	[−1.30]	

Table OA-8. Spanning regressions for ΔFXC.
The table reports the R^2 and adjusted R^2 of contemporaneous linear regressions of ΔFXC, the innovations of the FX correlation dispersion measure FXC, on spot exchange rate changes (Δs), currency excess returns (rx), changes in implied FX correlations (ΔIC) and changes in FX correlation risk premiums (ΔCRP); the results are reported in the first four rows of the table. In addition, we extract principal components and regress ΔFXC on the first principal components (PCs) of exchange rate changes, currency excess returns, changes in implied FX correlations and changes in FX correlation risk premiums; we use the first 6 PCs for exchange rate changes and currency excess returns and the first 6 or 20 (6 or 21) PCs for changes in FX implied correlations (FX correlation risk premiums). The results are reported in the last six rows of the table. The first 6 PCs of exchange rate changes and currency excess returns capture roughly 95% of the variation in exchange rate changes and currency excess returns, respectively. The first 6 and 20 (6 and 21) PCs of the changes in implied FX correlations (FX correlation risk premiums) capture roughly 72% and 95% (70% and 95%) of their overall variation, respectively. Monthly data from January 1999 through December 2013.

Type of variables	Spanning set	R^2	Adj. R^2
Δs	G10 (9 series)	0.119	0.072
rx	G10 (9 series)	0.120	0.073
ΔIC	G10 (36 series)	0.398	0.240
ΔCRP	G10 (36 series)	0.460	0.320
Δs	G10: 6 PCs	0.098	0.067
rx	G10: 6 PCs	0.099	0.068
ΔIC	G10: 6 PCs	0.215	0.187
ΔIC	G10: 20 PCs	0.324	0.236
ΔCRP	G10: 6 PCs	0.250	0.223
ΔCRP	G10: 21 PCs	0.383	0.298

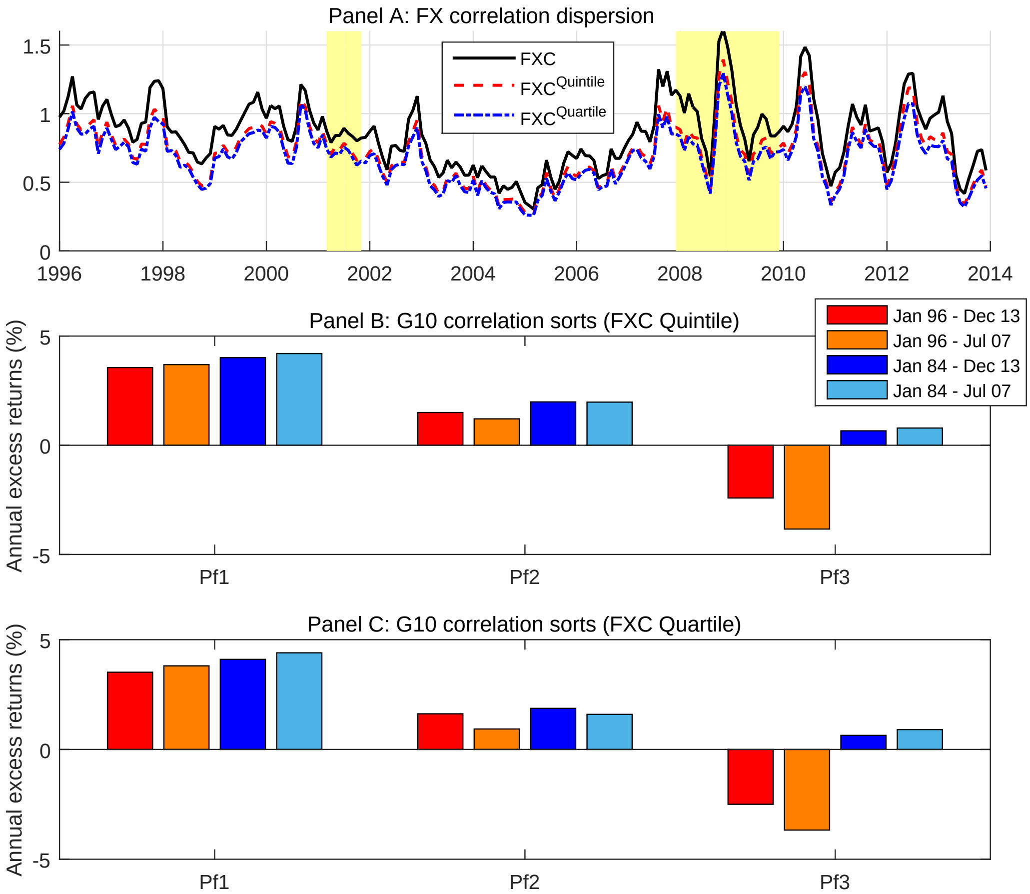

Fig. OA-1. Alternative FX correlation dispersion measures and currency portfolio excess returns.
Panel A plots the time series of three FX correlation dispersion measures, denoted by FXC, $FXC^{Quintile}$ and $FXC^{Quartile}$. The main measure, FXC, is calculated as the difference between the average conditional correlation of high and low correlation FX pairs; high (low) correlation FX pairs are defined as the FX pairs in the highest (lowest) decile across all 36 G10 FX pairs sorted on conditional FX correlations. The deciles are rebalanced every month. The two alternative measures, $FXC^{Quintile}$ and $FXC^{Quartile}$, are computed in the same fashion as FXC, but instead of using deciles, we use quintiles for $FXC^{Quintile}$ and quartiles for $FXC^{Quartile}$. Panels B and C present the average excess returns of currency portfolios sorted at time t on exposure to FX correlation risk at the end of period $t-1$, for four sample periods: January 1996-December 2013, January 1996-July 2007, January 1984-December 2013, and January 1984-July 2007. In each case, currencies are sorted into three portfolios. Exposure to FX correlation risk is measured by regressing currency excess returns on the innovations of a given alternative FX correlation dispersion measure ($FXC^{Quintile}$ for Panel B, $FXC^{Quartile}$ for Panel C) over the preceding 36 months. Portfolio 1 (Pf1) contains the currencies with the lowest pre-sort beta, beta whereas Portfolio 3 (Pf3) contains the currencies with the highest pre-sort beta.

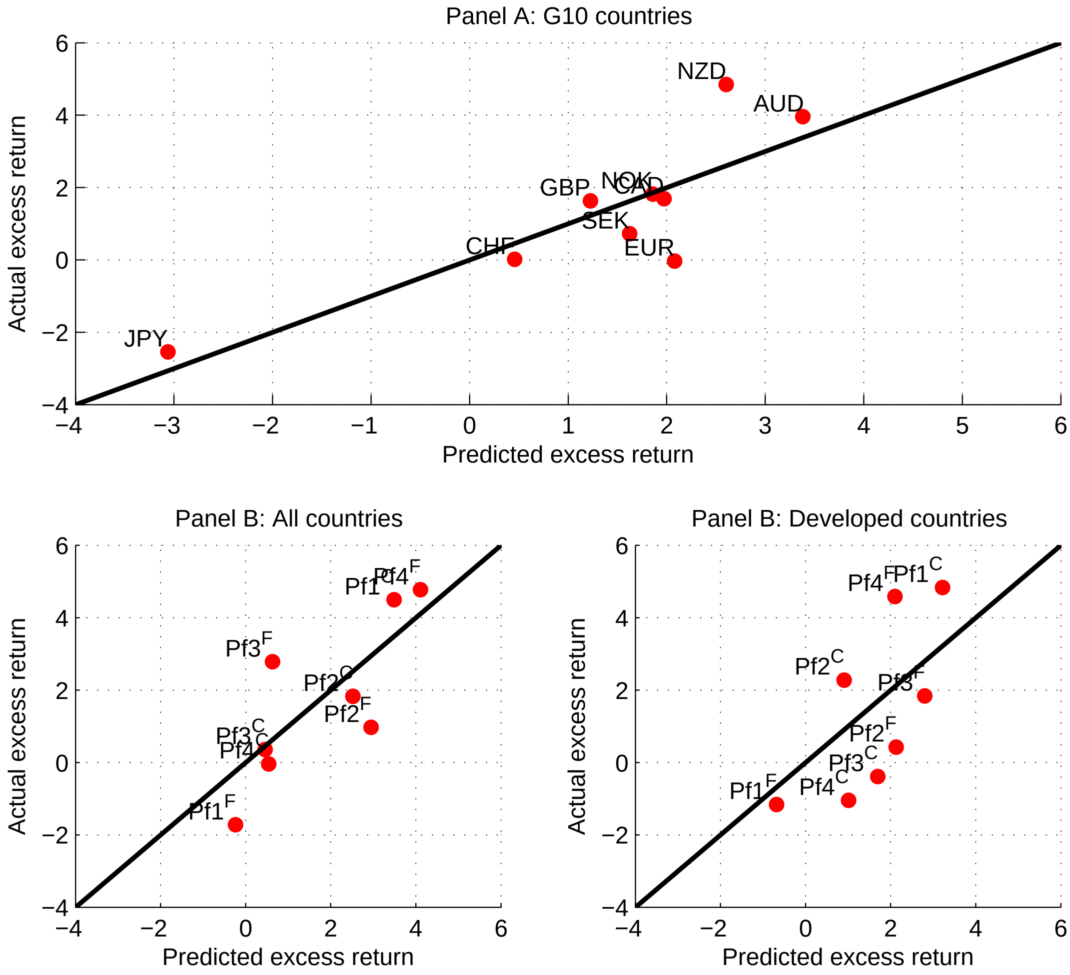

Fig. OA-2. Model performance with various test assets (model with ΔFXC).

The figure plots the realized annualized mean excess returns against the predicted annualized excess returns for various test assets using a linear pricing model that includes the dollar factor *DOL* and the FX correlation factor ΔFXC. Excess returns are reported in percentage points. Panel A displays the results for the nine G10 currencies. Panels B and C display the results for four currency portfolios sorted on ΔFXC betas (Pf1^{C} to Pf4^{C}) and four currency portfolios sorted on nominal interest rates (Pf1^{F} to Pf4^{F}), constructed using all currencies and developed country currencies, respectively. Observations are monthly, from January 1996 to December 2013.

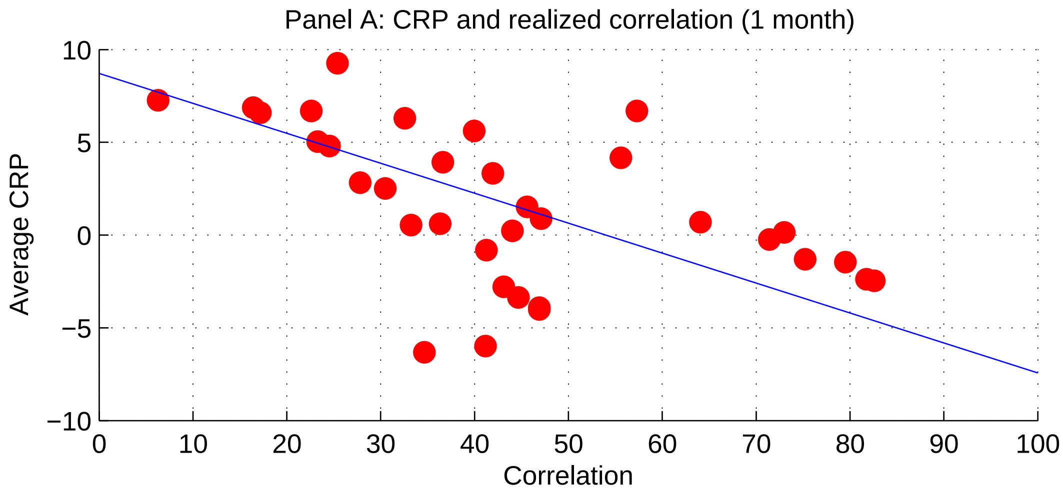

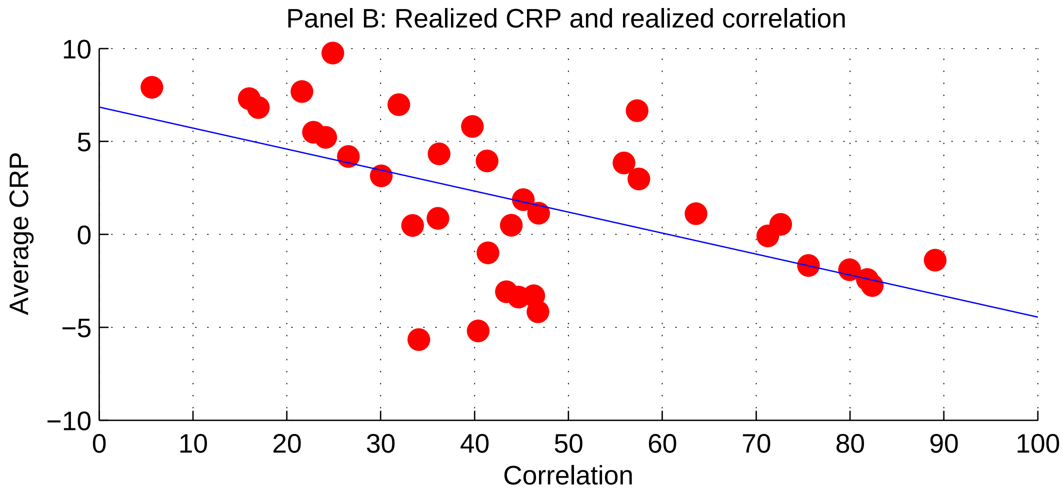

Fig. OA-3. G10 realized correlation risk premiums and realized correlations.
The figure plots the average correlation risk premiums for all 36 G10 FX pairs against their average realized correlations. In Panel A, conditional FX correlations at time t are proxied by the realized daily FX correlation over a past window of a month, i.e. over the period $[t-1, t]$. In Panel B, conditional FX correlations at time t are proxied by the realized daily FX correlation over a future window of a month, i.e., over the period $[t, t+1]$. Correlation risk premiums and correlations are expressed in percentage terms. Data are monthly, from January 1996 to December 2013 (options data for EUR start in January 1999).