The Quanto Theory of Exchange Rates

Lukas Kremens Ian Martin*

January, 2018

Abstract

We derive a new identity that relates expected exchange rate appreciation to a risk-neutral covariance term, and use it to motivate a currency forecasting variable based on the prices of quanto index contracts. We show via panel regressions that the quanto forecast variable is an economically and statistically significant predictor of currency appreciation and of excess returns on currency trades. Out of sample, the quanto variable outperforms predictions based on uncovered interest parity, on purchasing power parity, and on a random walk as a forecaster of differential (dollar-neutral) currency appreciation.

JEL codes: G12, G15, F31, F37, F47.

*Kremens: l.kremens@lse.ac.uk, LSE. Martin: i.w.martin@lse.ac.uk, LSE. We thank the Systemic Risk Centre at the LSE for their support, and for providing access to data sourced from Markit under license. We are grateful to Christian Wagner, Tarek Hassan, John Campbell, Mike Chernov, Gino Cenedese, Anthony Neuberger, Dagfinn Rime, Urban Jermann, Bryn Thompson-Clarke, Adrien Verdelhan, Bernard Dumas, Pierpaolo Benigno, Alan Taylor, and to participants in seminars at the LSE, Imperial College, Cass Business School, LUISS, and BI Business School, for their comments; and to Lerby Ergun for research assistance.

It is notoriously hard to forecast movements in exchange rates. A large part of the literature is organized around the principle of uncovered interest parity (UIP), which predicts that expected exchange rate movements offset interest rate differentials and therefore equalise expected returns across currencies. Unfortunately many authors, starting from Hansen and Hodrick (1980) and Fama (1984), have shown that this prediction fails: returns have historically been larger on high interest rate currencies than on low interest rate currencies.[1]

Given its empirical failings, it is worth reflecting on why UIP represents such an enduring benchmark in the FX literature. The UIP forecast has (at least) three appealing properties. First, UIP forecasts are determined by asset prices alone rather than by, say, infrequently updated and imperfectly measured macroeconomic data. Second, the UIP forecast has no free parameters; with no coefficients to be estimated in-sample or "calibrated," it is perfectly suited to out-of-sample forecasting. Third, the UIP forecast has a straightforward interpretation: it is the expected exchange rate movement that must be perceived by a risk-neutral investor. Put differently, UIP holds if and only if the *risk-neutral* expected appreciation of a currency is equal to its *real-world* expected appreciation, the latter being the quantity relevant for forecasting exchange rate movements.

There is, however, no reason to expect that the real-world and risk-neutral expectations should be similar. On the contrary, the modern literature in financial economics has documented that large and time-varying risk premia are pervasive across asset classes, so that risk-neutral and real-world distributions are very different from one another: in other words, the perspective of a risk-neutral investor is not useful from the point of view of forecasting. Thus, while UIP has been a useful organizing principle for the empirical literature on exchange rates, its predictive failure is no surprise.[2]

In this paper we propose a new predictor variable that also possesses the three appealing properties mentioned above, but which does not require that one takes the

[1]Some studies (e.g. Sarno, Schneider and Wagner, 2012) find that currencies with high interest rates appreciate on average, exacerbating the failure of UIP; this has become known as the forward premium puzzle. Others, such as Hassan and Mano (2016), find that exchange rates move in the direction predicted by UIP, though not by enough to offset interest rate differentials.

[2]Various authors have fleshed out this point in the context of equilibrium models: see for example Verdelhan (2010), Hassan (2013), and Martin (2013*b*). On the empirical side, authors including Menkhoff et al. (2012), Barroso and Santa-Clara (2015) and Della Corte, Ramadorai and Sarno (2016) have argued that it is necessary to look beyond interest rate differentials to explain the variation in currency returns.

perspective of a risk-neutral investor. This alternative benchmark can be interpreted as the expected exchange rate movement that must be perceived by a risk-averse investor with log utility whose wealth is invested in the stock market. (To streamline the discussion, this description is an oversimplification and strengthening of the condition we actually need to hold for our approach to work, which is based on a general identity presented in Result 1.) This approach has been shown by Martin (2017) and Martin and Wagner (2017) to be successful in forecasting returns on the stock market and on individual stocks, respectively.

It turns out that such an investor's expectations about currency returns can be inferred directly from the prices of so-called *quanto contracts*. Consider, for example, a quanto contract whose payoff equals the level of the S&P 500 index at time T, denominated in euros (that is, the exchange rate is fixed—in this example, at 1 euro per dollar—at initiation of the trade). The value of this contract is sensitive to the correlation between the S&P 500 index and the dollar/euro exchange rate. If the euro appreciates against the dollar at times when the index is high, and depreciates when the index is low, then this quanto contract is more valuable than a conventional, dollar-denominated, claim on the index.[3] We show that the relationship between currency-i quanto forward prices and conventional forward prices on the S&P 500 index reveals the risk-neutral covariance between currency i and the index. Quantos therefore signal which currencies are risky—in that they tend to depreciate in bad times, i.e., when the S&P 500 declines—and which are hedges; it is possible, of course, that a currency is risky at one point in time and a hedge at another. Intuitively, one expects that a currency that is (currently) risky should, as compensation, have higher expected appreciation than predicted by UIP, and that hedge currencies should have lower expected appreciation. Our framework formalizes this intuition. It also allows us to distinguish between variation in risk premia across currencies and variation over time.

It is worth emphasizing various assumptions that we do *not* make. We do not require that markets are complete (though our approach remains valid if they are). We do not assume the existence of a representative agent, nor do we assume that all economic actors are rational: the forecast in which we are interested reflects the beliefs of a rational investor, but this investor may coexist with investors with other,

[3] A different type of quanto contract—specifically, quanto CDS contracts—is used by Mano (2013) to estimate risk-neutral expectations of currency depreciation conditional on sovereign default.

potentially irrational, beliefs. And we do not assume lognormality, nor do we make any other distributional assumptions: our approach allows for skewness and jumps in exchange rates. This is an important strength of our framework, given that currencies often experience crashes or jumps (as emphasized by Brunnermeier, Nagel and Pedersen (2008), Jurek (2014), Della Corte et al. (2016), Chernov, Graveline and Zviadadze (2016) and Farhi and Gabaix (2016), among others), and are prone to structural breaks more generally. The approach could even be used, in principle, to compute expected returns for currencies that are currently pegged but that have some probability of jumping off the peg. To the extent that skewness and jumps are empirically relevant, this fact will be embedded in the asset prices we use as forecasting variables.

Our approach is therefore well adapted to the view of the world put forward by Burnside et al. (2011), who argue that the attractive properties of carry trade strategies in currency markets may reflect the possibility of peso events in which the SDF takes extremely large values. Investor concerns about such events, if present, should be reflected in the forward-looking asset prices that we exploit, and thus our quanto predictor variable should forecast high appreciation for currencies vulnerable to peso events even if no such events turn out to happen in sample.

We derive these and other theoretical results in Section 1, and test them in Section 2 by running panel currency-forecasting regressions. The estimated coefficient on the quanto predictor variable is economically large and statistically significant: in our headline regression (23), we find t-statistics of 3.2 and 2.3 (respectively with and without currency fixed effects; and with standard errors computed using a nonparametric block bootstrap). The quanto predictor outperforms forecasting variables such as the interest rate differential, average forward discount, and the real exchange rate as a univariate forecaster of currency excess returns. On the other hand, we find that some of these variables—notably the real exchange rate and average forward discount—interact well with our quanto predictor variable, in the sense that they substantially raise R^2 above what the quanto variable achieves on its own. We interpret this fact, through the lens of the identity (6) of Result 1, as showing that these variables help to measure deviations from the log investor benchmark. We also show that the quanto predictor variable—that is, forward-looking risk-neutral covariance—predicts future realized covariance and substantially outperforms lagged realized covariance as a forecaster of exchange rates.

In Section 2.5, we consider various joint hypotheses motivated by Result 2 on the coefficient estimates in the panel regressions of Section 2. To do so, we conduct Wald tests with (as always) covariance matrices computed using a nonparametric block bootstrap, and with p-values calculated in two ways: (i) using, as is conventional, the asymptotic distribution of the Wald test statistic and (ii) using a bootstrapped small-sample distribution. The latter approach directly accounts for the fact that our dataset spans a relatively short time period. Using the asymptotic p-values, we find, in our pooled regressions, that the estimated coefficients on the quanto predictor variable and interest rate differential are consistent with the predictions of Result 2, but we can reject the hypothesis that, in addition, the intercept is zero; this rejection can be attributed to US dollar appreciation, over our sample period, that was not anticipated by our model. In the regressions with currency fixed effects, we can reject the prediction of Result 2 because the coefficient on the quanto variable is even larger than expected. When we use the more conservative small-sample p-values, however, we do not reject even the most optimistic hypothesis in any of the specifications, though the individual significance of the quanto predictor becomes more marginal (with p-values ranging from 5.1% to 9.7%).

In Section 3 we show that the quanto variable performs well out of sample. In a recent survey of the literature, Rossi (2013) emphasizes that the exchange-rate forecasting literature has struggled to overturn the frustrating fact, originally documented by Meese and Rogoff (1983), that it is hard even to outperform a random walk forecast out of sample. Since our data span a relatively short period (from 2009 to 2017) over which the dollar strengthened against almost all the other currencies in our dataset, we focus on forecasting differential returns on currencies. This allows us to isolate the cross-sectional forecasting power of the quanto variable in a dollar-neutral way, in the spirit of Lustig, Roussanov and Verdelhan (2011), and independent of what Hassan and Mano (2016) refer to as the dollar trade anomaly. (Our findings are therefore complementary to Gourinchas and Rey (2007), who use a measure of external imbalances to forecast the appreciation of the dollar itself against a trade- or FDI-weighted basket of currencies.)

Our out-of-sample forecasts exploit the fact that the theory makes an a priori prediction for the coefficient on the predictor variable. When the coefficient on the quanto predictor is fixed at the level implied by the theory, we end up with a forecast of currency appreciation that has no free parameters, and which is therefore—like the UIP forecast—perfectly suited for out-of-sample forecasting. Following Meese and Rogoff (1983) and Goyal and Welch (2008), we compute mean squared error for the differential currency forecasts made by the quanto theory and by three competitor models: UIP, which predicts currency appreciation through the interest rate differential; PPP, which uses past inflation differentials (as a proxy for expected inflation differentials) to forecast currency appreciation; and the random walk forecast. The quanto theory outperforms all three competitors. We also show that it outperforms on an alternative performance benchmark, the correct classification frontier, that has been proposed by Jordà and Taylor (2012).

1 Theory

We start with the fundamental equation of asset pricing,

$$\mathbb{E}_t \left(M_{t+1} \widetilde{R}_{t+1} \right) = 1, \tag{1}$$

since this will allow us to introduce some notation. Today is time t; we are interested in assets with payoffs at time $t+1$. We write $\mathbb{E}_t$ for the (real-world) expectation operator, conditional on all information available at time t, and M_{t+1} for a stochastic discount factor (SDF) that prices assets denominated in dollars. (We do not assume complete markets, so there may well be other SDFs that also price assets denominated in dollars. But all such SDFs must agree with M_{t+1} on the prices of the payoffs in which we are interested, since they are all tradable.) In equation (1), $\widetilde{R}_{t+1}$ is the gross return on some arbitrary dollar-denominated asset or trading strategy. If we write $R^{\$}_{f,t}$ for the gross one-period dollar interest rate, then the equation implies that $\mathbb{E}_t\, M_{t+1} = 1/R^{\$}_{f,t}$, as can be seen by setting $\widetilde{R}_{t+1} = R^{\$}_{f,t}$; thus (1) can be rearranged as

$$\mathbb{E}_t\, \widetilde{R}_{t+1} - R^{\$}_{f,t} = -R^{\$}_{f,t} \operatorname{cov}_t \left(M_{t+1}, \widetilde{R}_{t+1} \right). \tag{2}$$

Consider a simple currency trade: take a dollar, convert it to foreign currency i, invest at the (gross) currency-i riskless rate, $R^{i}_{f,t}$, for one period, and then convert back to dollars. We write $e_{i,t}$ for the price in dollars at time t of a unit of currency i, so that the gross return on the currency trade is $R^{i}_{f,t} e_{i,t+1}/e_{i,t}$; setting $\widetilde{R}_{t+1} = R^{i}_{f,t} e_{i,t+1}/e_{i,t}$ in

(2) and rearranging,[4] we find that

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} = \underbrace{\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{UIP forecast}} \underbrace{- R^{\$}_{f,t} \operatorname{cov}_t \left(M_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right)}_{\text{residual}}. \tag{3}$$

This (well known) identity can also be expressed using the risk-neutral expectation $\mathbb{E}^*_t$, in terms of which the time t price of any payoff, X_{t+1}, received at time $t+1$ is

$$\text{time } t \text{ price of a claim to } X_{t+1} = \frac{1}{R^{\$}_{f,t}} \mathbb{E}^*_t X_{t+1} = \mathbb{E}_t \left(M_{t+1} X_{t+1} \right). \tag{4}$$

The first equality is the defining property of the risk-neutral probability distribution. The second equality (which can be thought of as a dictionary for translating between risk-neutral and SDF notation) can be used to rewrite (3) as

$$\mathbb{E}^*_t \left(\frac{e_{i,t+1}}{e_{i,t}} \right) = \frac{R^{\$}_{f,t}}{R^{i}_{f,t}}. \tag{5}$$

From an empirical point of view, the challenging aspect of the identity (3) is the presence of the unobservable SDF M_{t+1}. If M_{t+1} were constant conditional on time t information then the covariance term would drop out and we would recover the UIP prediction that $\mathbb{E}_t \, e_{i,t+1}/e_{i,t} = R^{\$}_{f,t}/R^{i}_{f,t}$, according to which high-interest-rate currencies are expected to depreciate. Thus, if the UIP forecast is used to predict exchange rate appreciation, the implicit assumption being made is that the covariance term can indeed be neglected.

Unfortunately, as is well known, the UIP forecast performs poorly in practice: the assumption that the covariance term is negligible in (3) (or, equivalently, that the risk-neutral expectation in (5) is close to the corresponding real-world expectation) is not valid. This is hardly surprising, given the existence of a vast literature in financial economics that emphasizes the importance of risk premia, and hence shows that the SDF M_{t+1} is highly volatile (Hansen and Jagannathan, 1991). The risk adjustment term in (3) therefore cannot be neglected: expected currency appreciation depends

[4] Unlike most authors in this literature, we prefer to work with true returns, $\widetilde{R}_{t+1}$, rather than with log returns, $\log \widetilde{R}_{t+1}$, as the latter are only "an approximate measure of the rate of return to speculation," in the words of Hansen and Hodrick (1980). We elaborate on this point in Section 2.

not only on the interest rate differential, but also on the covariance between currency movements and the SDF. Moreover, it is plausible that this covariance varies both over time and across currencies. We therefore take a different approach that exploits the following observation:

Result 1. *Let R_{t+1} be an arbitrary gross return. We have the identity*

$$\mathbb{E}_t \, \frac{e_{i,t+1}}{e_{i,t}} = \underbrace{\frac{R_{f,t}^{\$}}{R_{f,t}^{i}}}_{\textit{UIP forecast}} + \underbrace{\frac{1}{R_{f,t}^{\$}} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right)}_{\textit{quanto-implied risk premium}} \underbrace{- \operatorname{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right)}_{\textit{residual}}. \tag{6}$$

The asterisk on the first covariance term in (6) *indicates that it is computed using the risk-neutral probability distribution.*

Proof. Setting $\widetilde{R}_{t+1} = R_{f,t}^{i} e_{i,t+1}/e_{i,t}$ in (1) and rearranging, we have

$$\mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) = \frac{1}{R_{f,t}^{i}}. \tag{7}$$

We can use (4) and (7) to expand the risk-neutral covariance term that appears in the identity (6) and express it in terms of the SDF:

$$\begin{aligned} \frac{1}{R_{f,t}^{\$}} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right) &\overset{(4)}{=} \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - R_{f,t}^{\$} \, \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) \\ &\overset{(7)}{=} \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - \frac{R_{f,t}^{\$}}{R_{f,t}^{i}}. \end{aligned} \tag{8}$$

Note also that

$$\operatorname{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right) = \mathbb{E}_t \left(M_{t+1} R_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) - \mathbb{E}_t \left(\frac{e_{i,t+1}}{e_{i,t}} \right). \tag{9}$$

Subtracting (9) from (8) and rearranging, we have the result. □

As (3) and (6) are identities, each must hold for all currencies i in any economy that does not permit riskless arbitrage opportunities. The identity (6) generalizes (3), however, by allowing R_{t+1} to be an arbitrary return. To make the identity useful for empirical work, we want to choose a return R_{t+1} with two aims in mind. First, the residual term should be small. Second, the middle term should be easy to compute.

These two goals are in tension. If we set $R_{t+1} = R^{\$}_{f,t}$, for example, then (6) reduces to (3), which achieves the second of the goals but not the first. Conversely, one might imagine setting R_{t+1} equal to the return on an elaborate portfolio exposed to multiple risk factors and constructed in such a way as to minimise the volatility of $M_{t+1}R_{t+1}$: this would achieve the first but not necessarily the second, as will become clear in the next section.

To achieve both goals simultaneously, we want to pick a return that offsets a substantial fraction of the variation in M_{t+1}; but we must do so in such a way that the risk-neutral covariance term can be measured empirically. For much of this paper, we will take R_{t+1} to be the return on the S&P 500 index. (We find similar—and internally consistent—results if R_{t+1} is set equal to the return on other stock indexes, such as the Nikkei, Euro Stoxx 50, or SMI: see Sections 1.2 and 2.1.) It is highly plausible that this return is negatively correlated with M_{t+1}, as dictated by the first goal; in fact we provide conditions below under which the residual is exactly zero. We will now show that the second goal is also achieved with this choice of R_{t+1} because we can calculate the quanto-implied risk premium directly from asset prices without any further assumptions—specifically, from *quanto forward prices* (hence the name).

1.1 Quantos

An investor who is bullish about the S&P 500 index might choose to go long a *forward contract* at time t, for settlement at time $t+1$. If so, he commits to pay F_t at time $t+1$ in exchange for the level of the index, P_{t+1}. The dollar payoff on the investor's long forward contract is therefore $P_{t+1} - F_t$ at time $t+1$. Market convention is to choose F_t to make the market value of the contract equal to zero, so that no money needs to change hands initially. This requirement implies that

$$F_t = \mathbb{E}^*_t \, P_{t+1}. \tag{10}$$

A *quanto forward contract* is closely related. The key difference is that the quanto forward commits the investor to pay $Q_{i,t}$ *units of currency* i at time $t+1$, in exchange for P_{t+1} *units of currency* i. (At each time t, there are N different quanto prices indexed by $i = 1, \ldots, N$, one for each of the N currencies in our data set. Other than in Section 1.2, the underlying asset is always the S&P 500 index, whatever the currency.)

The payoff on a long position in a quanto forward contract is therefore $P_{t+1} - Q_{i,t}$ units of currency i at time $t+1$; this is equivalent to a time $t+1$ dollar payoff of $e_{i,t+1}(P_{t+1} - Q_{i,t})$. As with a conventional forward contract, the market convention is to choose the quanto forward price, $Q_{i,t}$, in such a way that the contract has zero value at initiation. It must therefore satisfy

$$Q_{i,t} = \frac{\mathbb{E}_t^* \, e_{i,t+1} P_{t+1}}{\mathbb{E}_t^* \, e_{i,t+1}}. \tag{11}$$

(We converted to dollars because $\mathbb{E}_t^*$ is the risk-neutral expectations operator that prices *dollar* payoffs.) Combining equations (5) and (11), the quanto forward price can be written

$$Q_{i,t} = \frac{R_{f,t}^i}{R_{f,t}^\$} \, \mathbb{E}_t^* \, \frac{e_{i,t+1} P_{t+1}}{e_{i,t}},$$

which implies, using (5) and (10), that the gap between the quanto and conventional forward prices captures the conditional risk-neutral covariance between the exchange rate and stock index,

$$Q_{i,t} - F_t = \frac{R_{f,t}^i}{R_{f,t}^\$} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, P_{t+1} \right). \tag{12}$$

We will make the simplifying assumption that dividends earned on the index between time t and time $t+1$ are known at time t and paid at time $t+1$. It then follows from (12) that

$$\frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} = \frac{1}{R_{f,t}^\$} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right), \tag{13}$$

so the quanto forward and conventional forward prices are equal if and only if currency i is uncorrelated with the stock index under the risk-neutral measure. This allows us to measure the risk-neutral covariance term that appears in (6) directly from the gap between quanto and conventional index forward prices (which, as noted, we will refer to as the quanto-implied risk premium).

We still have to deal with the final covariance term in the identity (6). The next result exhibits a case in which this covariance term is exactly zero.

Result 2 (The log investor)**.** *If we take the perspective of an investor with log utility whose wealth is fully invested in the stock index then* $M_{t+1} = 1/R_{t+1}$*, so that*

$\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$ *is identically zero. The expected appreciation of currency* i *is then given by*

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{\frac{R^{\$}_{f,t}}{R^i_{f,t}} - 1}_{IRD_{i,t}} + \underbrace{\frac{Q_{i,t} - F_t}{R^i_{f,t} P_t}}_{QRP_{i,t}}, \tag{14}$$

and the expected excess return[5] *on currency* i *equals the quanto-implied risk premium:*

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^i_{f,t}} = \frac{Q_{i,t} - F_t}{R^i_{f,t} P_t}.$$

Equation (14) splits expected currency appreciation into two terms. The first is the UIP prediction which, as we have seen in equation (5), equals *risk-neutral* expected currency appreciation. We will often refer to this term as the *interest rate differential* (IRD); and as above we will generally convert to net rather than gross terms by subtracting 1. (We choose to refer to a high-interest-rate currency as having a *negative* interest rate differential because such a currency is forecast to depreciate by UIP.) The second is a risk adjustment term: by taking the perspective of the log investor, we have converted the general form of the residual that appears in (3) into a quantity that can be directly observed using the gap between a quanto forward and a conventional forward. Since it captures the risk premium perceived by the log investor, we refer to this term as the *quanto-implied risk premium* (QRP). Lastly, we refer to the sum of the two terms as *expected currency appreciation* (ECA = IRD + QRP).

Results 1 and 2 link expected currency returns to *risk-neutral* covariances, so deviate from the standard CAPM intuition (that risk premia are related to *true* covariances) in that they put more weight on comovement in bad states of the world. This distinction matters, given the observation of Lettau, Maggiori and Weber (2014) that the carry trade is more correlated with the market when the market experiences negative returns. Even more important, risk-neutral covariance is directly measurable, as we have shown. In contrast, forward-looking true covariances are *not* directly observed so must be proxied somehow, typically by historical realized covariance. In Section 2.3, we show that risk-neutral covariance drives out historical realized covariance as a predictor variable.

[5]Formally, $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^i_{f,t}$ is an excess return because it is a tradable payoff whose price is zero, by (5).

The assumption that underpins Result 2, $M_{t+1} = 1/R_{t+1}$, implies that

$$\mathbb{E}_t \frac{R^{\$}_{i,t+1}}{R_{t+1}} = 1, \tag{15}$$

where $R^{\$}_{i,t+1}$ is the US dollar return on the currency-i stock market. This is reminiscent of the *uncovered equity parity* condition proposed by Hau and Rey (2006); while they do not provide a formal definition, a natural characterization of the condition (analogous to the UIP condition, along the lines of Cappiello and De Santis (2007)) is that expected stock market returns are equated in dollars:

$$\mathbb{E}_t R^{\$}_{i,t+1} = \mathbb{E}_t R_{t+1}. \tag{16}$$

The conditions (15) and (16) are similar but distinct. For the sake of argument, if stock market returns and exchange rates are jointly lognormal then (15) implies that the expected returns are not in general equated: instead,

$$\mathbb{E}_t R^{\$}_{i,t+1} = \mathbb{E}_t R_{t+1} \times \exp\left\{\text{cov}_t\left(r^{\$}_{i,t+1}, r_{t+1}\right) - \text{var}_t\, r_{t+1}\right\} ,$$

where lower-case letters indicate log returns. Thus a sufficiently volatile (and correlated) foreign stock market may earn a higher expected dollar return than the US stock market according to (15) but not according to (16).

Lastly, we emphasize that while Result 2 represents a useful benchmark and is the jumping-off point for our empirical work, in our analysis below we will also allow for the presence of the final covariance term in the identity (6). Throughout the paper, we do so in a simple way by reporting regression results with (and without) currency fixed effects, to account for any currency-dependent but time-independent component of the covariance term. In Section 2.4, we consider further proxies that depend both on currency and time.

1.2 Alternative benchmarks

Our choice to think from the perspective of an investor who holds the US stock market is a pragmatic one. From a purist point of view, it might seem more natural to adopt the perspective of an investor whose wealth is invested in a globally diversified portfolio;

unfortunately global-wealth quantos are not traded, whereas S&P 500 quantos are. Our approach implicitly relies on an assumption that the US stock market is a tolerable proxy for global wealth. We think this assumption makes sense; it is broadly consistent with the 'global financial cycle' view of Miranda-Agrippino and Rey (2015).

Nonetheless, one might wonder whether the results are similar if one uses other countries' stock markets as proxies for global wealth.[6] For, just as the forward price of the US stock index quantoed into currency i reveals the expected appreciation of currency i versus the dollar, as perceived by a log investor whose portfolio is fully invested in the US stock market, so the forward price of the currency-i stock index quantoed into dollars reveals the expected appreciation of the dollar versus currency i, as perceived by a log investor whose portfolio is fully invested in the currency-i market.

Recall Result 2 for the expected appreciation of currency i versus the dollar,

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{\text{IRD}_{i,t} + \text{QRP}_{i,t}}_{\text{ECA}_{i,t}}. \tag{17}$$

(To reiterate, a positive value indicates that currency i is expected to strengthen against the dollar.) The corresponding expression for the expected appreciation of the dollar versus currency i, from the perspective of a log investor whose wealth is fully invested in the currency-i stock market, is

$$\mathbb{E}_t^i \frac{1/e_{i,t+1}}{1/e_{i,t}} - 1 = \underbrace{\text{IRD}_{1/i,t} + \text{QRP}_{1/i,t}}_{\text{ECA}_{1/i,t}}, \tag{18}$$

where (with a slight abuse of notation) we write $\text{IRD}_{1/i,t} = R^i_{f,t}/R^\$_{f,t} - 1$, and where $\text{QRP}_{1/i,t}$ is obtained from conventional and *dollar*-denominated quanto forwards on the currency-i stock market. When the left-hand side of the above equation is positive, the dollar is expected to appreciate against currency i.

In Section 2.1 below, we show that the two perspectives captured by (17) and (18) are broadly consistent with one another (for those currencies for which we observe the appropriate quanto forward prices). If, say, the forward price of the S&P 500 quantoed into euros implies that the euro is expected to appreciate against the dollar by 2%

[6] In practice, many investors do choose to hold home-biased portfolios (French and Poterba (1991), Tesar and Werner (1995), and Warnock (2002); and see Lewis (1999) and Coeurdacier and Rey (2013) for surveys).

(using equation (17)), then the forward price of the Euro Stoxx 50 index quantoed into dollars typically implies that the dollar is expected to *de*preciate against the euro by about 2% (using equation (18)). To be more precise, we need to take into account Siegel's "paradox" (Siegel, 1972) that, by Jensen's inequality,

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} \geq \left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1} . \tag{19}$$

(The corresponding inequality with $\mathbb{E}_t$ replaced by any other expectation operator also holds.) If the US and currency-i investors have the same expectations about currency appreciation then (17)–(19) imply that

$$\log\left(1 + \mathrm{ECA}_{i,t}\right) \geq -\log\left(1 + \mathrm{ECA}_{1/i,t}\right) . \tag{20}$$

In practice $\log(1 + \mathrm{ECA}) \approx \mathrm{ECA}$, so the above inequality is essentially equivalent to $\mathrm{ECA}_{i,t} \geq -\mathrm{ECA}_{1/i,t}$: thus (continuing the example) if the euro is expected to appreciate by 2% against the dollar, then the dollar should be expected to depreciate against the euro by at most 2%.

The difference between the two sides of (20) reflects a convexity correction whose size is determined by the amount of conditional variation in $e_{i,t+1}$. If the exchange rate is lognormal, $\log(e_{i,t+1}/e_{i,t}) \sim N(\mu_t, \sigma_t^2)$, then by a straightforward calculation[7]

$$\begin{aligned} \log\left(1 + \mathrm{ECA}_{i,t}\right) - \left(-\log\left(1 + \mathrm{ECA}_{1/i,t}\right)\right) &= \log \mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \log \left[\left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1} \right] \\ &= \sigma_t^2 . \end{aligned} \tag{21}$$

Thus if exchange rate volatility is on the order of 10%, the two perspectives should disagree by about 1% (so in the example above, expected euro appreciation of 2% would be consistent with expected dollar depreciation of 1%). In Section 2.1, we show that the convexity gap observed in our data is indeed consistent with (21).

[7] More generally the correction involves all even cumulants: $\log \mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \log \left[\left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1} \right] = \boldsymbol{c}(1) + \boldsymbol{c}(-1) = 2 \sum_{n \text{ even}} \kappa_n / n!$, where $\boldsymbol{c}(\cdot)$ and κ_n denote, respectively, the cumulant-generating function and the nth cumulant of log exchange rate appreciation; see Martin (2013*a*). In particular, $\kappa_2 = \sigma_t^2$ and κ_4/σ_t^4 is the excess kurtosis of $\log e_{i,t+1}$. If the log exchange rate is normally distributed, as in (21), all cumulants above the second are zero. For an early treatment of cumulants in the context of exchange rates see Backus, Foresi and Telmer (2001).

2 Empirics

We obtained forward prices and quanto forward prices on the S&P 500, together with domestic and foreign interest rates, from Markit; the maturity in each case is 24 months. The data is monthly and runs from December 2009 to October 2015 for the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), Euro (EUR), British pound (GBP), Japanese yen (JPY), Korean won (KRW), Norwegian krone (NOK), Polish zloty (PLN), and Swedish krona (SEK). Since these quantos are used to forecast exchange rates over a 24-month horizon, our forecasting sample runs from December 2009 to October 2017. Markit reports consensus prices based on quotes received from a wide range of financial intermediaries. These prices are used by major OTC derivatives market makers as a means of independently verifying their book valuations and to fulfil regulatory requirements; they do not necessarily reflect transaction prices. Accounting for missing entries in our panel, we have 656 currency-month observations.[8]

Since the financial crisis of 2007-2009, a growing literature (including Du, Tepper and Verdelhan (2016)) has discussed the failure of covered interest parity (CIP)—the no-arbitrage relation between forward exchange rates, spot exchange rates and interest rate differentials—and established that since the financial crisis, CIP frequently does not hold if interest rates are obtained from money markets. For each maturity, we observe currency-specific discount factors directly from our Markit data set. The implied interest rates are consistent with the observed forward prices and the absence of arbitrage. Our measure of the interest rate differentials therefore does not violate the no-arbitrage condition we require for identity (6) to hold.

The two building blocks of our empirical analysis are the currencies' quanto-implied risk premia (QRP, which measure the risk-neutral covariances between each currency and the S&P 500 index, as shown in equation (13)), and their interest rate differentials vis-à-vis the US dollar (IRD, which would equal expected exchange rate appreciation if UIP held). Our measure of expected currency appreciation (the quanto forecast, or ECA) is equal to the sum of IRD and QRP, as in equation (14).

Figure 1 plots each currency's QRP over time; for clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK). The QRP

[8]Where we do not observe a price, we treat the observation as missing. Larger periods of consecutive missing observations occur only for DKK, KRW, and PLN and are shown as gaps in Figure 2.

is negative for JPY and positive for all other currencies (with the partial exception of EUR, for which we observe a sign change in QRP near the end of our time period).

Figure 2 shows the evolution over time of ECA (solid) and of the UIP forecast (dashed) for each of the currencies in our panel. The gap between the two lines for a given currency is that currency's QRP. Table 1 reports summary statistics of ECA. The penultimate line of the table averages the summary statistics across currencies; the last line reports summary statistics for the pooled data. Table 2 reports the same statistics for IRD and QRP.

The volatility of QRP is similar to that of interest rate differentials, both currency-by-currency and in the panel. There is considerably more variability in IRD and QRP when we pool the data than there is in the time series of a typical currency: this reflects substantial dispersion in IRD and QRP across currencies that is captured in the pooled measure but not in the average time series.

Table 3 reports volatilities and correlations for the time series of individual currencies' ECA, IRD, and QRP. The table also shows three aggregated measures of volatilities and correlations. The row labelled "Time series" reports time-series volatilities and correlations for a typical currency, calculated by averaging time-series volatilities and correlations across currencies. Conversely, the row labelled "Cross section" reports cross-currency volatilities and correlations of time-averaged ECA, IRD, and QRP. Lastly, the row labelled "Pooled" averages on both dimensions: it reports volatilities and correlations for the pooled data.

All three variables (ECA, IRD, and QRP) are more volatile in the cross section than in the time series. This is particularly true of interest rate differentials, which exhibit far more dispersion across currencies than over time.

The correlation between IRD and QRP is negative when we pool our data ($\rho = -0.696$). Given the sign convention on IRD, this indicates that currencies with high interest rates (relative to the dollar) tend to have high risk premia; thus the predictions of the quanto theory are consistent with the carry trade literature and the findings of Lustig, Roussanov and Verdelhan (2011). The average time-series (i.e., within-currency) correlation between IRD and QRP is more modestly negative ($\rho = -0.331$): a typical currency's risk premium tends to be higher, or less negative, at times when its interest rate is high relative to the dollar, but this tendency is fairly weak. The disparity between these two facts is accounted for by the strongly negative cross-sectional correlation between IRD and QRP ($\rho = -0.798$). If we interpret the data through the lens of Result 2, these findings suggest that the returns to the carry trade are more the result of persistent *cross-sectional* differences between currencies than of a *time-series* relationship between interest rates and risk premia. This prediction is consistent with the empirical results documented by Hassan and Mano (2016).

We see a corresponding pattern in the time-series, cross-sectional, and pooled correlations of ECA and QRP. The time-series (within-currency) correlation of the two is substantially positive ($\rho = 0.393$), while the cross-sectional correlation is negative ($\rho = -0.305$). In the time series, therefore, a rise in a given currency's QRP is associated with a rise in its expected appreciation; whereas in the cross-section, currencies with relatively high QRP on average have relatively *low* expected currency appreciation on average (reflecting relatively high interest rates on average). Putting the two together, the pooled correlation is close to zero ($\rho = -0.026$). That is, Result 2 predicts that there should be no clear relationship between currency risk premia and expected currency appreciation; again, this is consistent with the findings of Hassan and Mano (2016).

These properties are illustrated graphically in Figure 3. We plot confidence ellipses centred on the means of QRP and IRD in panel (a), and of QRP and ECA in panel (b), for each currency. The sizes of the ellipses reflect the volatilities of IRD and QRP (or ECA): under joint normality, each ellipse would contain 50% of its currency's observations in population.[9] The orientation of each ellipse illustrates the within-currency time series correlation, while the positions of the different ellipses reveal correlations across currencies. The figures refine the discussion above. QRP and IRD are negatively correlated within currency (with the exceptions of CAD, CHF, and KRW) and in the cross-section. QRP and ECA are positively correlated in the time series for every currency, but exhibit negative correlation across currencies; overall, the pooled correlation between the two is close to zero.

Our empirical analysis focuses on contracts with a maturity of 24 months, since these have the best data availability. But in the case of the S&P 500 index quantoed into euros, we observe a range of maturities so can explore the term structure of QRP. Figure 4 plots the time series of annualized euro-dollar QRP for horizons of 6, 12, 24, and 60 months. On average, the term structure of QRP is flat over the sample

[9]Our interest is in the relative sizes of the ellipses, so the choice of 50% is arbitrary.

period. However, shorter horizons are slightly more volatile, so that the term structure is downward-sloping when QRP spikes and upward-sloping when QRP is low.

2.1 A consistency check

Our data also includes quanto forward prices of certain other stock indexes, notably the Nikkei, Euro Stoxx 50, and SMI. We can use this data to explore the predictions of Section 1.2, which provides a consistency check on our empirical strategy.

Figure 5 implements (17) and (18) for the EUR-USD, JPY-USD, EUR-JPY, and EUR-CHF currency pairs. In each of the top-left, bottom-left and bottom-right panels, the solid line depicts the expected appreciation of the euro against the US dollar, yen, and Swiss franc, respectively, while the dashed line shows the expected *de*preciation of the three currencies against the euro (that is, we flip the sign on the "inverted" series for readability). In the top-right panel, the solid and dashed lines show the expected appreciation of the yen against the US dollar and expected depreciation of the US dollar against the yen, respectively. In every case, the two measures are strongly correlated over time and the solid line is above the dashed line, as they should be according to (20). The gaps between the measures are therefore consistent with the Jensen's inequality correction one would expect to see if our currency forecasts measured expected currency appreciation perfectly. Moreover, given that annual exchange rate volatilities are on the order of 10%, the sizes of the gaps between the measures are quantitatively consistent with the Jensen's inequality correction derived in equation (21).

The EUR-CHF pair in the bottom-right panel represents a particularly interesting case study. The Swiss national bank instituted a floor on the EUR-CHF exchange rate at CHF1.20/€ in September 2011 and consequently also reduced the conditional volatility of the exchange rate. Following this, the two lines converge and the gap stays very narrow at around 0.2% up until January 2015, when the sudden removal of the floor prompted a spike in the volatility of the currency pair, visible in the figure as the point at which the two lines diverge.

2.2 Return forecasting

We run two sets of panel regressions in which we attempt to forecast, respectively, currency excess returns and currency appreciation. The literature on exchange rate forecasting has found it substantially more difficult to forecast pure currency appreciation than currency excess returns, so the second set of regressions should be considered more empirically challenging. In each case, we test the prediction of Result 2 via pooled panel regressions. We also report the results of panel regressions with currency fixed effects; by doing so, we allow for the more general possibility that there is a currency-dependent—but time-independent—component in the second covariance term that appears in the identity (6).

To provide a sense of the data before turning to our regression results, Figures 6 and 7 represent our baseline univariate regressions graphically in the same manner as in Figure 3. Figure 6 plots realized currency excess returns (RXR) against QRP and against IRD.[10] Excess returns are strongly positively correlated with QRP both within currency and in the cross-section, suggesting strong predictability with a positive sign. The correlation of RXR with IRD is negative in the cross-section but close to zero, on average, within currency.

Figure 7 shows the corresponding results for realized currency appreciation (RCA). Panel (a) suggests that the within-currency correlation with the quanto predictor ECA is predominantly positive (with the exceptions of AUD and CHF), as is the cross-sectional correlation. In contrast, panel (b) suggests that the correlation between realized currency appreciation and interest rate differentials is close to zero both within and across currencies, consistent with the view that interest rate differentials do not help to forecast currency appreciation.

We first run a horse race between the quanto-implied risk premium and interest rate differential as predictors of currency excess returns:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta \, \mathrm{QRP}_{i,t} + \gamma \, \mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{22}$$

We also run two univariate regressions. The first of these,

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta \, \mathrm{QRP}_{i,t} + \varepsilon_{i,t+1}, \tag{23}$$

[10] As noted in Section 1, we work with true returns as opposed to log returns. Engel (2016) points out that it may not be appropriate to view log returns as approximating true returns, since the gap between the two is a similar order of magnitude as the risk premium itself.

is suggested by Result 2. The second uses interest rate differentials to forecast currency excess returns, as a benchmark:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma \, \mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{24}$$

We also run all three regressions with currency fixed effects α_i in place of the shared intercept α.

Table 4 reports the results. We report coefficient estimates and R^2 for each regression, with and without currency fixed effects; standard errors are shown in parentheses. These standard errors are computed via a nonparametric bootstrap to account for the cross-sectional and serial correlation structure in our data. (For comparison, these nonparametric standard errors exceed those obtained from a parametric residual bootstrap by up to a factor of 2.) We provide a detailed description of our bootstrap procedure and address potential small-sample concerns in Section 2.5.

The estimated coefficient on the quanto-implied risk premium is positive and economically large in every specification in which it occurs. Moreover, the R^2 values are substantially higher in the two regressions (22) and (23) that feature the quanto-implied risk premium than in the regression (24) in which it does not occur. The estimate for β in our headline regression (23) is 2.604 (standard error 1.127) in the pooled regression and 4.995 (standard error 1.565) in the regression with fixed effects. The fact that these estimates are above 1 raises the possibility that beyond its direct importance in (6), the quanto-implied risk premium may also proxy for the second covariance term. We explore this issue in Section 2.4. Another noteworthy qualitative feature of our results is the consistently negative intercept, which reflects an unexpectedly strong dollar over our sample period; we discuss the statistical interpretation of this fact in Section 2.5.

Following Fama (1984), we can also test how the theory fares at predicting currency appreciation $(e_{i,t+1}/e_{i,t} - 1)$. To do so, we run the regression

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta \, \mathrm{QRP}_{i,t} + \gamma \, \mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{25}$$

We do so not because we are interested in the coefficient estimates, which are mechanically related to those of regression (22), but because we are interested in the R^2.

To explore the relative importance of the quanto-implied risk premium and interest rate differentials for forecasting currency appreciation, we run univariate regressions of currency appreciation onto the quanto-implied risk premium,

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta \, \text{QRP}_{i,t} + \varepsilon_{i,t+1}, \tag{26}$$

and onto interest rate differentials,

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \gamma \, \text{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{27}$$

As previously, we also run the three regressions (25)–(27) with fixed effects.

The regression results are shown in Table 5, which is structured similarly to Table 4. There is little evidence that the interest rate differential helps to forecast currency appreciation on its own; this is consistent with the previous set of results and with the large literature that documents the failure of UIP. In the pooled panel, the estimated γ in regression (27) is close to 0, and the R^2 is essentially zero. With fixed effects, the estimate of γ is marginally negative, providing weak evidence that currencies tend to appreciate against the dollar when their interest rate relative to the dollar is higher than its time-series mean.

More strikingly, the quanto-implied risk premium makes a very large difference in terms of R^2, which increases by two orders of magnitude when moving from specification (27) to (25) in both the pooled regressions (0.16% to 16.01%) and the fixed-effects regressions (0.20% to 20.56%). It is also interesting to note that when QRP is included in the regressions (with or without fixed effects) the coefficient estimate on IRD, γ, increases toward the value of 1 predicted by Result 2.

For completeness, Table 6 reports the results of running regressions (23), (24), (25), and (27) separately for each currency at the 24-month horizon, and at 6- and 12-month horizons for the euro. Consistent with the previous literature (for example Fama (1984) and Hassan and Mano (2016)), the coefficient estimates are extremely noisy. A further appealing feature of Result 2 is that it provides a justification for constraining all the coefficient on the quanto-implied risk premium to be equal across currencies, as we have done above.

2.3 Risk-neutral covariance vs. true covariance

We have emphasized the importance of risk-neutral covariances of currencies with stock returns, as captured by quanto-implied risk premia, and below we will show that risk-neutral covariance performs well empirically. But it is natural to wonder whether this empirical success merely reflects the fact that currency returns line up with *true* covariances, as studied by Lustig and Verdelhan (2007), Campbell, Medeiros and Viceira (2010), Burnside (2011) and Cenedese et al. (2016), among others. More formally, from the perspective of the log investor we can conclude, from (3), that

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = R^{\$}_{f,t} \operatorname{cov}_t \left(\frac{e_{i,t+1}}{e_{i,t}}, -\frac{1}{R_{t+1}} \right). \tag{28}$$

Note that it is the true, not the risk-neutral, covariance that appears in this equation.

The fundamental challenge for a test of this prediction is that forward-looking true covariance is not directly observed. This is the major advantage of our approach: risk-neutral covariance *is* directly observed via the quanto-implied risk premium. That said, we attempt to test (28) by using lagged realized covariance, RPCL, as a proxy for true forward-looking covariance. The results are shown in Table 8 of the Appendix. RPCL is positively related to subsequently realized currency excess returns, as suggested by (28), but it is driven out as a predictor by risk-neutral covariance (QRP), consistent with Result 2. We also find that risk-neutral covariance is a statistically significant forecaster of future realized covariance.

2.4 Beyond the log investor

The identity (6) expresses expected currency appreciation as the sum of IRD, QRP, and a covariance term, $-\operatorname{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$. Thus far, we have either assumed that this term is constant across currencies and over time (so is captured by the constant in our pooled regressions) or that it has a currency-dependent but time-independent component (so is captured by fixed effects).

To get a sense of what these assumptions may leave out, we conduct a principal components analysis on unexpected currency excess returns: that is, on the difference between realized currency excess returns and the corresponding ex ante expected returns. We calculate these unexpected excess returns in two ways. *Regression residuals* are defined as the estimated residuals $\varepsilon_{i,t+1}$ in the specification of regression (23) that includes currency fixed effects. *Theory residuals* are defined similarly, except that we impose $\alpha = 0$, $\beta = 1$ in (23).

These residuals reflect both the ex ante residual from the identity (6) and the ex post realizations of unexpected currency returns. The identity implies that the predictable component of the realized residuals—if there is one—reveals the covariance term, $-\operatorname{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$.

We decompose the theory and regression residuals into their respective principal components (dropping DKK, KRW, and PLN from the panel to minimize the impact of missing observations). Table 7 shows the principal component loadings. The first principal component, which explains just under two thirds of the variation in residuals, can be interpreted as a level, or 'dollar,' factor since it loads positively on all currencies (with the exception of GBP when in the case of the regression residuals).

Motivated by this fact, we now include an additional predictor variable, $\overline{IRD}_t = (1/N_t)\sum_i \mathrm{IRD}_{i,t}$, where N_t is the number of currencies in our dataset at time t; Lustig, Roussanov and Verdelhan (2014) interpret this average interest rate differential (which they refer to as the 'average forward discount') as a dollar factor and show that it helps to forecast currency returns. We also include the logarithm of the real exchange rate, which Dahlquist and Penasse (2017) have shown to be a successful forecaster of currency returns.

Table 9 reports the results of regressions of currency excess returns onto currency fixed effects and subsets of four forecasting variables: the quanto-implied risk premium (QRP), the interest rate differential (IRD), the real exchange rate (RER), and the average interest rate differential ($\overline{\mathrm{IRD}}$). The table reports the univariate, bivariate, 3-variate, and 4-variate specifications with the highest R^2. (Table 10 reports the R^2 for all $2^4-1 = 15$ subsets of the four explanatory variables, though not—for lack of space—the estimated coefficients.) The quanto-implied risk premium features in all R^2-maximizing regressions. The estimates of β are larger than 1 in every specification, suggesting that, over and above its relevance as a direct measure of risk-neutral covariance, the quanto-implied risk premium helps to capture the physical covariance term in (6). As we increase from one to two to three explanatory variables, R^2 increases from 22.03% (using QRP alone) to 35.25% (adding the real exchange rate) to 43.28% (adding the dollar factor $\overline{\mathrm{IRD}}$). The interest rate differential itself, IRD, contributes almost no further explanatory power when it is then added as a fourth variable. Since the real exchange rate performs well, Table 11 reports further results that incorporate it as a regressor.

2.5 Joint hypothesis tests and finite-sample issues

We now consider the joint hypothesis tests that are suggested by Result 2. In our three main specifications (22), (23), and (25), equation (14) predicts an intercept $\alpha = 0$, and a slope coefficient on QRP $\beta = 1$. For the excess return forecast in regression (22), it predicts that the interest rate differential should have no predictive power, i.e. $\gamma = 0$; whereas it predicts that $\gamma = 1$ in the currency-appreciation regression (25).

Here, as elsewhere, we use a nonparametric bootstrap procedure to compute the covariance matrix of coefficient estimates. A detailed exposition of the bootstrap methodology is provided in Politis and White (2004) and Patton, Politis and White (2009). In the bootstrap procedure, we resample the data by drawing with replacement blocks of 24 time-series observations from the panel while ensuring that this time-series resampling is synchronized in the cross-section. The length of the time-series blocks is chosen to equal the forecasting horizon of 24 months. The resulting panel is then resampled with replacement in the cross-sectional dimension by drawing blocks of uniformly distributed width (between 2 and 11, the latter being the width of the full cross-section). Since currencies which are adjacent in the panel are more likely to be included together in any given one of these cross-sectional blocks, we permute the cross-section of our panel randomly before each resampling. We then compute the point estimates of the coefficients from the two-dimensionally resampled panel and repeat this procedure 100,000 times. The standard errors are then computed as the standard deviations of the respective coefficients across the 100,000 bootstrap repetitions.

Table 12 reports p-values for tests of various hypotheses about our baseline regressions. In addition to conventional p-values calculated using the asymptotic (chi-squared) distribution of the Wald test statistic, the table also reports more conservative small-sample p-values obtained from a bootstrapped test statistic distribution. We compute these small-sample p-values by constructing a small-sample distribution of the Wald test-statistic for each regression: We simulate 5,000 sets of monthly data for the LHS variable under the null hypothesis of no predictability, such that the simulated data matches the monthly autocorrelation and covariance matrix of the realized, observed LHS data. We then aggregate the simulated monthly data into 24-month horizon data, like the LHS data used in our regressions (e.g. excess returns over 24 months). As we aim to measure the small-sample performance of our bootstrap routine, the simulated data sets each have the same number of data points as the observed LHS data. For each specification, we then regress the 5,000 simulated LHS data on the respective observed RHS variable(s). Where we run the regression with currency fixed effects, we use the demeaned RHS variable(s). We obtain the point estimates of the coefficients and their covariance matrix from the bootstrap routine outlined above and use the test statistics from these 5,000 regressions to construct the empirical small-sample distribution of the respective Wald statistic under the respective null hypothesis.

Figure 8 illustrates by plotting the histograms of the bootstrapped distribution of test statistics for various hypotheses on regression (25). Panels a and b show the finite-sample bootstrapped distributions of the test statistic for the hypothesis that Result 2 holds, respectively in the pooled and fixed-effects regressions. The value of the test statistic in the data is indicated with an asterisk in each panel. The finite-sample and asymptotic (shown with a solid line) distributions are strikingly different: the asymptotic distribution suggests that we can reject the hypothesis that Result 2 holds, but this conclusion is overturned by the finite-sample distribution. (In the pooled case, the discrepancy is largely due to the intercept, as becomes clear on comparing the asymptotic p-values for tests of hypotheses H_0^1 and H_0^2 in Table 12: the asymptotic distribution penalizes the fact that the US dollar was strong over our sample period, whereas the finite-sample distribution does not.)

In contrast, the asymptotic and finite-sample distributions tell more or less the same story in panels c and d, which show the corresponding results for tests (without and with fixed effects) of the hypothesis H_0^3 that $\beta = 0$, i.e. that QRP is not useful in forecasting currency appreciation. While the small-sample distributions of the test statistics exhibit fatter tails than the asymptotic χ^2 distribution, the discrepancy between the two is small by comparison with panels a and b, and even using the finite-sample distribution we can reject the hypothesis with some confidence (with p-values of 0.082 and 0.051 in the pooled and fixed-effects cases, respectively).

We reach similar conclusions for regressions (22) and (23): we do not reject the predictions of Result 2 in the joint Wald tests for any of the three baseline regressions using the small-sample distribution of the test statistic; and QRP remains individually significant as a predictor at the 10% level in all three specifications, with and without currency fixed effects, even if we take the most conservative approach to computing p-values that relies on the empirical small-sample test statistic distribution.

3 Out-of-sample prediction

We now test the quanto theory out of sample. Since the dollar strengthened strongly over the relatively short time period spanned by our data (as reflected in the negative intercept in our pooled panel regression (25)), we focus on forecasting differential currency appreciation: that is, we seek to predict, for example, the *relative* performance of dollar-yen versus dollar-euro.

In the previous section, we estimated the loadings on the quanto-implied risk premium, QRP, and interest rate differential, IRD, via panel regressions. These deliver the best in-sample coefficient estimates in a least-squares sense. But for an out-of-sample test we must pick the loadings a priori. Here we can exploit the distinctive feature of Result 2 that it makes specific quantitative predictions for the loadings: each should equal 1, as in the formula (14). We therefore compute out-of-sample forecasts by fixing the coefficients that appear in (25) at their theoretical values: $\alpha = 0$, $\beta = 1$, $\gamma = 1$.

We compare these predictions to those of three competitor models: UIP (which predicts that currency appreciation should offset the interest rate differential, on average), a random walk without drift (which makes the constant forecast of zero currency appreciation, and which is described in the survey of Rossi (2013) as "the toughest benchmark to beat"), and relative purchasing power parity (which predicts that currency appreciation should offset the inflation differential, on average). These models are natural competitors because, like our approach, they make a priori predictions without requiring estimation of parameters, and so avoid in-sample/out-of-sample issues.

3.1 Mean squared errors

To compare the forecast accuracy of the model to those of the benchmarks, we define a dollar-neutral R^2-measure similar to that of Goyal and Welch (2008):

$$R_{OS}^2 = 1 - \frac{\sum_i \sum_j \sum_t (\varepsilon_{i,t+1}^Q - \varepsilon_{j,t+1}^Q)^2}{\sum_i \sum_j \sum_t (\varepsilon_{i,t+1}^B - \varepsilon_{j,t+1}^B)^2},$$

where $\varepsilon_{i,t+1}^Q$ and $\varepsilon_{i,t+1}^B$ denote forecast errors (for currency i against the dollar) of the quanto theory and the benchmark, respectively, so our measure compares the accuracy of differential forecasts of currencies i and j against the dollar. We hope to find that the quanto theory has lower mean squared error than each of the competitor models, that is, we hope to find positive R_{OS}^2 versus each of the benchmarks.

The results of this exercise are reported in Table 13. The quanto theory outperforms each of the three competitors: when the competitor model is UIP, we find that $R_{OS}^2 = 10.91\%$; and when it is relative PPP, we find $R_{OS}^2 = 26.05\%$. In our sample, the toughest benchmark is the random walk forecast, consistent with the findings of Rossi (2013). Nonetheless, the quanto theory easily outperforms it, with $R_{OS}^2 = 9.57\%$.

To get a sense for whether our positive results are driven by a small subset of the currencies, Table 13 also reports the results of splitting the R^2 measure currency-by-currency: for each currency i, we define

$$R_{OS,i}^2 = 1 - \frac{\sum_j \sum_t (\varepsilon_{i,t+1}^Q - \varepsilon_{j,t+1}^Q)^2}{\sum_j \sum_t (\varepsilon_{i,t+1}^B - \varepsilon_{j,t+1}^B)^2}.$$

This quantity is positive for all i and all competitor benchmarks B, indicating that the quanto theory outperforms all three benchmarks for all 11 currencies. We run Diebold–Mariano tests (Diebold and Mariano, 1995) of the null hypothesis that the quanto theory and competitor models perform equally well for all currencies, using a small-sample adjustment proposed by Harvey, Leybourne and Newbold (1997), and find that the outperformance is strongly significant.

3.2 Binary forecast accuracy

Jordà and Taylor (2012) have argued that assessments of forecast performance based solely on mean squared errors may not fully reflect the economic benefits of a forecasting model. In this section, we follow the approach of Jordà and Taylor by computing a *correct classification frontier* (CCF) to assess the performance of our forecasts in correctly predicting the direction of currency movements.

Denote by $f^Q_{i,j,t} = \text{QRP}_{i,t} - \text{QRP}_{j,t}$ and $f^B_{i,j,t}$ the forecasts obtained, respectively, from the quanto variable and a competitor benchmark for currency pair (i, j) at time t. Similarly, $r_{i,t} = e_{i,t+1}/e_{i,t} - R^{\$}_t/R^i_t$ denotes the realized excess return of the currency i against the dollar, and $r_{i,j,t} = r_{i,t} - r_{j,t}$ represents the dollar-neutral return in currency pair (i, j). We calculate the *true positive* (TP) and *true negative* (TN) rates for each forecasting model as a function of a threshold, c:

$$TP(c) = P(f^Q_{i,j,t} > c \mid r_{i,j,t} > 0), \tag{29}$$

$$TN(c) = P(f^Q_{i,j,t} < c \mid r_{i,j,t} < 0) \tag{30}$$

The two rates describe, respectively, the fractions of ex post positive long and short returns that were correctly identified ex ante as profitable long and short positions by the forecasting model.

For the same 55 dollar-neutral currency pairs used above, we find that $TP(0) = 0.50$, $TN(0) = 0.64$, with a weighted average correct classification of 0.57 for the quanto forecast. Since binary accuracy does not reflect the magnitudes of returns from the signal, we follow Jordà and Taylor (2012) and denote the ex post maximum gain from long positions by $L = \sum_{r_{i,j,t}>0} r_{i,j,t}$, and similarly for the ex post maximum returns from short positions $S = -\sum_{r_{i,j,t}<0} r_{i,j,t}$. We compute the return-weighted true positive (TP*) and true negative (TN*) rates as

$$TP^*(c) = \frac{\sum_{f^Q_{i,j,t}>c|r_{i,j,t}>0} r_{i,j,t}}{L}, \tag{31}$$

$$TN^*(c) = \frac{-\sum_{f^Q_{i,j,t}<c|r_{i,j,t}<0} r_{i,j,t}}{S} \tag{32}$$

When we weight the forecasts by the realization of the excess return they would have earned, the true rates increase to $TP^*(0) = 0.58$, $TN^*(0) = 0.67$, with a weighted average of 0.63. Both rates increase relative to the equally-weighted classifications, which implies that the direction of excess return realizations is more likely to have been predicted by the quanto variable when these realizations are large.

The CCF (and analogously CCF*) is defined as the set of pairs $\{TP(c), TN(c)\}$

for all possible values of c between $-\infty$ and ∞. Varying the threshold level, c, trades off true positives against true negatives by shifting the direction of the forecast. For instance, for $c = \infty$, the true negative rate is maximized at $TN = 1$, at the cost of $TP = 0$. Since $TN(c)$ and $TP(c)$ must lie between 0 and 1, we can plot the resulting CCF in the unit square, and compute the *area under the CCF* (AUC). Intuitively, the AUC can be interpreted as the probability that the forecast for a randomly chosen positive return realization will be higher than that for a randomly chosen negative return realization. Under the UIP forecast the excess return on any currency is 0, so the CCF is the diagonal with slope -1 in the unit square and, accordingly, $AUC = 0.5$.

We benchmark the quanto forecast against the driftless random walk model considered above (which forecasts the currency excess return as being equal to the interest rate differential). Figure 9 shows the resulting CCFs. The quanto forecast outperforms the random walk model for equally-weighted and return-weighted classifications. For the quanto forecast, $AUCQ = 0.60$ and $AUCQ^* = 0.70$, while the random walk model achieves $AUCRW = 0.55$ and $AUCRW^* = 0.60$. Both forecasts correctly identify large returns more often than small returns, as the CCF* (red) lies above the CCF (blue) in both cases.

We also reverse the conditioning in the true positive and true negative rates, to calculate how likely a forecast is to signal the correct direction of trade, and denote these by $PT(c)$ and $NT(c)$, respectively:

$$PT(c) = P(r_{i,j,t} > c \mid f^{Q}_{i,j,t} > 0), \tag{33}$$

$$NT(c) = P(r_{i,j,t} < c \mid f^{Q}_{i,j,t} < 0) \tag{34}$$

We find $PT(0) = 0.60$, $NT(0) = 0.54$, $PT^*(0) = 0.65$, and $NT^*(0) = 0.63$. Plotting the resulting CCFs, Figure 10 shows that, again, the quanto variable outperforms the random walk forecast with AUC-measures of $AUCQ = 0.60$, $AUCQ^* = 0.71$, against the random walk model with $AUCRW = 0.55$ and $AUCRW^* = 0.60$.

4 Conclusion

UIP forecasts that high interest rate currencies should depreciate on average; it reflects the expected currency appreciation that a genuinely risk-neutral investor would perceive in equilibrium. Unsurprisingly—given that the financial economics literature has repeatedly documented the importance of risk premia—the UIP forecast performs extremely poorly in practice.

We have proposed an alternative forecast, the quanto-implied risk premium, that can be interpreted as the expected excess return on a currency perceived by an investor with log utility whose wealth is fully invested in the stock market. Like the UIP forecast, the quanto forecast has no free parameters and can be computed directly from asset prices. Unlike the UIP forecast, the quanto forecast performs well empirically both in and out of sample.

Empirically, we find that currencies tend to have high quanto-implied risk premia if they have high interest rates on average, relative to other currencies (a cross-sectional statement), or if they currently have unusually high interest rates (a time-series statement); and there is more cross-sectional than time-series variation in quanto-implied risk premia. These facts explain both the existence of the carry trade and the empirical importance of persistent cross-currency asymmetries, as documented by Hassan and Mano (2016).

The interpretation of the quanto-implied risk premium as revealing the log investor's expectation of currency excess returns is a special case of the identity (6), which decomposes expected currency appreciation into the interest rate differential (the UIP term), risk-neutral covariance (the quanto-implied risk premium), and a real-world covariance term that, we argue, is likely to be small—and in particular, smaller than the corresponding covariance term in the well-known identity (3). In the log investor case, this real-world covariance term is exactly zero, a fact we use to provide intuition and to motivate our out-of-sample analysis. But we also allow for deviations from the log investor benchmark—that is, for a nontrivial real-world covariance term—by running regressions including currency fixed effects, realized covariance, interest rate differentials, the average forward discount of Lustig, Roussanov and Verdelhan (2014), and the real exchange rate, as in Dahlquist and Penasse (2017), in addition to the quanto-implied risk premium itself. The quanto-implied risk premium is the best performing univariate predictor, and features in every R^2-maximizing multivariate specification.

Let us note, finally, that although we have argued that quanto-implied risk premia should (in theory) and do (in practice) predict currency excess returns, we have said nothing about why a particular currency should have a high or low quanto-implied risk premium at a given time. Analogously, the CAPM predicts that assets' betas should forecast their returns but has nothing to say about why a given asset has a high or low beta. Connecting quanto-implied risk premia to macroeconomic fundamentals is an interesting topic for future research.

References

Backus, D. K., S. Foresi, and C. I. Telmer. 2001. "Affine Term Structure Models and the Forward Premium Anomaly." *Journal of Finance*, 56(1): 279–304.

Barroso, P., and P. Santa-Clara. 2015. "Beyond the Carry Trade: Optimal Currency Portfolios." *Journal of Financial and Quantitative Analysis*, 50(5): 1037–1056.

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

Burnside, C. 2011. "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Comment." *American Economic Review*, 101(7): 3456–3476.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo. 2011. "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial Studies*, 24(3): 853–891.

Campbell, J. Y., K. Serfaty-De Medeiros, and L. M. Viceira. 2010. "Global Currency Hedging." *The Journal of Finance*, 65(1): 87–121.

Cappiello, L., and R. A. De Santis. 2007. "The Uncovered Return Parity Condition." *ECB Working Paper Series 812*.

Cenedese, G., R. Payne, L. Sarno, and G. Valente. 2016. "What do Stock Markets Tell Us About Exchange Rates?" *Review of Finance*, 20(3): 1045–1080.

Chernov, M., J. Graveline, and I. Zviadadze. 2016. "Crash Risk in Currency Returns." *Journal of Financial and Quantitative Analysis*, forthcoming.

Coeurdacier, N., and H. Rey. 2013. "Home Bias in Open Economy Financial Macroeconomics." *Journal of Economic Literature*, 51(1): 63–115.

Dahlquist, M., and J. Penasse. 2017. "The Missing Risk Premium in Exchange Rates." *Working Paper.*

Della Corte, P., L. Sarno, M. Schmeling, and C. Wagner. 2016. "Sovereign Risk and Currency Returns." *Working Paper.*

Della Corte, P., T. Ramadorai, and L. Sarno. 2016. "Volatility Risk Premia and Exchange Rate Predictability." *Journal of Financial Economics*, 20(1).

Diebold, F. X., and R. S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics*, 13(2): 253–263.

Du, W., A. Tepper, and A. Verdelhan. 2016. "Deviations from Covered Interest Parity." *Working Paper.*

Engel, C. 2016. "Exchange Rates, Interest Rates, and the Risk Premium." *American Economic Review*, 106(2): 436–474.

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

Farhi, E., and X. Gabaix. 2016. "Rare Disasters and Exchange Rates." *The Quarterly Journal of Economics*, 131(1): 1–52.

French, K. R., and J. M. Poterba. 1991. "Investor Diversification and International Equity Markets." *American Economic Review*, 81(2): 222–226.

Gourinchas, P. O., and H. Rey. 2007. "International Financial Adjustment." *Journal of Political Economy*, 115(4): 665–703.

Goyal, A., and I. Welch. 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21: 1455–1508.

Hansen, L. P., and R. Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 99(2): 225–262.

Hansen, L. P., and R. J. Hodrick. 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy*, 88(5): 829–853.

Harvey, D., S. Leybourne, and P. Newbold. 1997. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting*, 13: 281–291.

Hassan, T. A. 2013. "Country Size, Currency Unions, and International Asset Returns." *The Journal of Finance*, 68(6): 2269–2308.

Hassan, T. A., and R. C. Mano. 2016. "Forward and Spot Exchange Rates in a Multi-Currency World." *Working Paper.*

Hau, H., and H. Rey. 2006. "Exchange Rates, Equity Prices, and Capital Flows." *Review of Financial Studies*, 19(1): 273–317.

Jordà, O., and A. M. Taylor. 2012. "The Carry Trade and Fundamentals: Nothing to Fear but FEER itself." *Journal of International Economics*, 88: 74–90.

Jurek, J. W. 2014. "Crash Neutral Currency Carry Trades." *Journal of Financial Economics*, 113(3): 325–347.

Lettau, M., M. Maggiori, and M. Weber. 2014. "Conditional Risk Premia in Currency Markets and Other Asset Classes." *Journal of Financial Economics*, 114(2): 197–225.

Lewis, K. K. 1999. "Trying to Explain Home Bias in Equities and Consumption." *Journal of Economic Literature*, 37(2): 571–608.

Lustig, H., and A. Verdelhan. 2007. "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk." *American Economic Review*, 97(1): 89–117.

Lustig, H. N., N. L. Roussanov, and A. Verdelhan. 2011. "Common Risk Factors in Currency Markets." *Review of Financial Studies*, 24(11).

Lustig, H. N., N. L. Roussanov, and A. Verdelhan. 2014. "Countercyclical Currency Risk Premia." *Journal of Financial Economics*, 111(3): 527–553.

Mano, R. C. 2013. "Exchange Rates Upon Sovereign Default." *Working Paper.*

Martin, I. W. R. 2013*a*. "Consumption-Based Asset Pricing with Higher Cumulants." *Review of Economic Studies*, 80(2): 745–773.

Martin, I. W. R. 2013*b*. “The Forward Premium Puzzle in a Two-Country World.” *Working Paper.*

Martin, I. W. R. 2017. “What is the Expected Return on the Market?” *Quarterly Journal of Economics*, 132(1): 367–433.

Martin, I. W. R., and C. Wagner. 2017. “What is the Expected Return on a Stock?” *Working Paper.*

Meese, R., and K. Rogoff. 1983. “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics*, 14(1-2): 3–24.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012. “Carry Trades and Global Foreign Exchange Volatility.” *Journal of Finance*, 67(2): 681–718.

Miranda-Agrippino, S., and H. Rey. 2015. “World Asset Markets and the Global Financial Cycle.” *NBER Working Paper No. 21722.*

Patton, A., D. N. Politis, and H. White. 2009. “Correction to Automatic Block-Length Selection for the dependent Bootstrap.” *Econometric Reviews*, 28: 372–375.

Politis, D. N., and H. White. 2004. “Automatic Block-Length Selection for the dependent Bootstrap.” *Econometric Reviews*, 23: 53–70.

Rossi, B. 2013. “Exchange Rate Predictability.” *Journal of Economic Literature*, 51(4): 1063–1119.

Sarno, L., P. Schneider, and C. Wagner. 2012. “Properties of Foreign Exchange Risk Premiums.” *Journal of Financial Economics*, 105(2): 279–310.

Siegel, J. J. 1972. “Risk, interest rates, and the forward exchange.” *Quarterly Journal of Economics*, 86: 303–309.

Tesar, L. L., and I. M. Werner. 1995. “Home Bias and High Turnover.” *Journal of International Money and Finance*, 14(4): 467–492.

Verdelhan, A. 2010. “A Habit-Based Explanation of the Exchange Rate Risk Premium.” *Journal of Finance*, 65(1): 123–145.

Warnock, F. E. 2002. “Home Bias and High Turnover Reconsidered.” *Journal of International Money and Finance*, 21(6): 795–805.

A Appendix

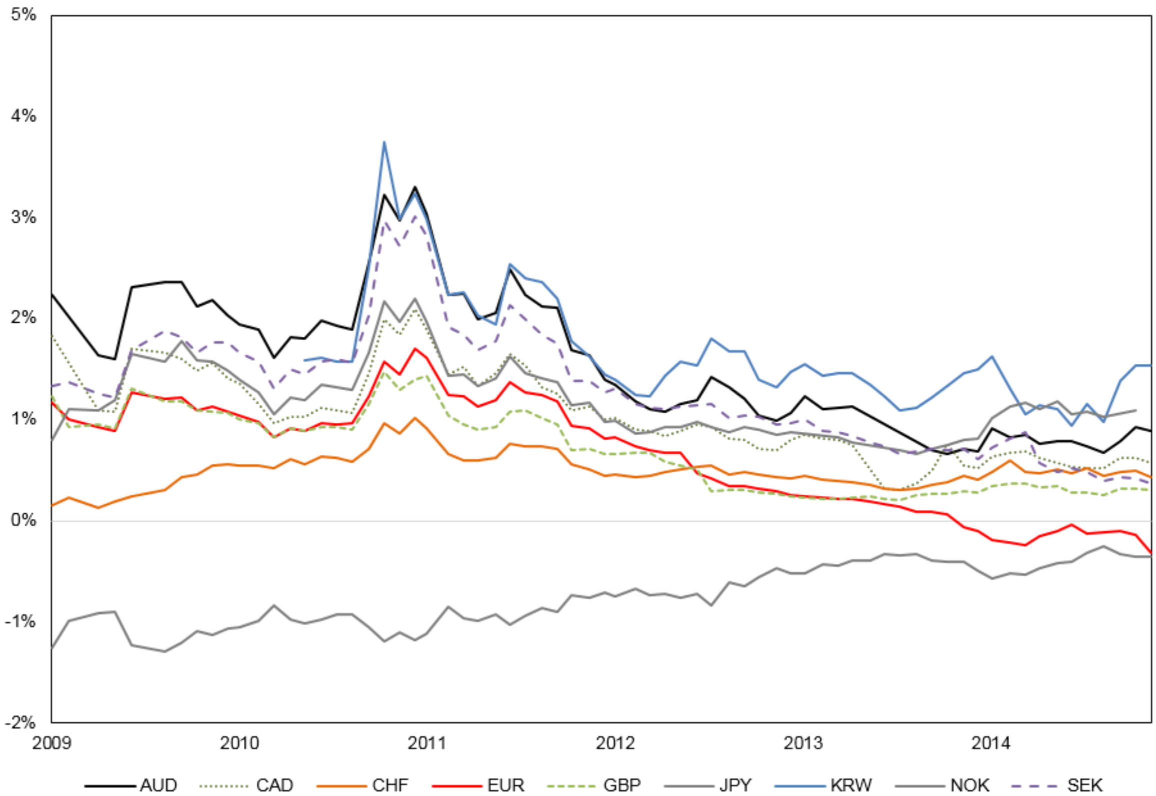

Figure 1: The time series of QRP. The figure drops two currencies (PLN and DKK) for which we have highly incomplete time series.

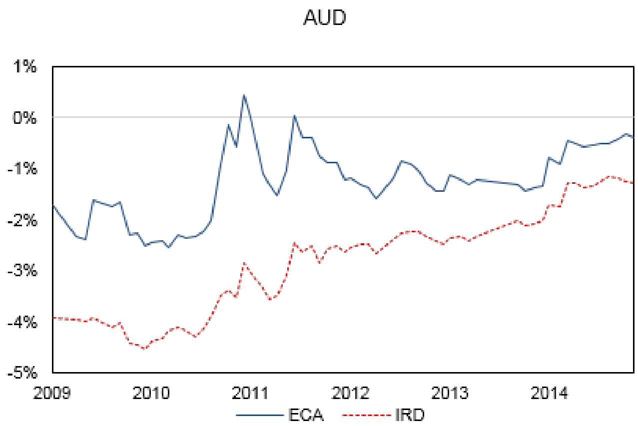
AUD
1%
0%
-1%
-2%
-3%
-4%
-5%
2009
2010
2011
2012
2013
2014
ECA
IRD

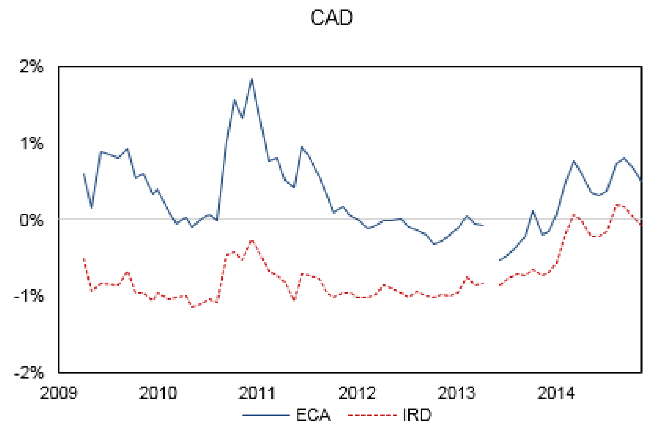
CAD
2%
1%
0%
-1%
-2%
2009
2010
2011
2012
2013
2014
ECA
IRD

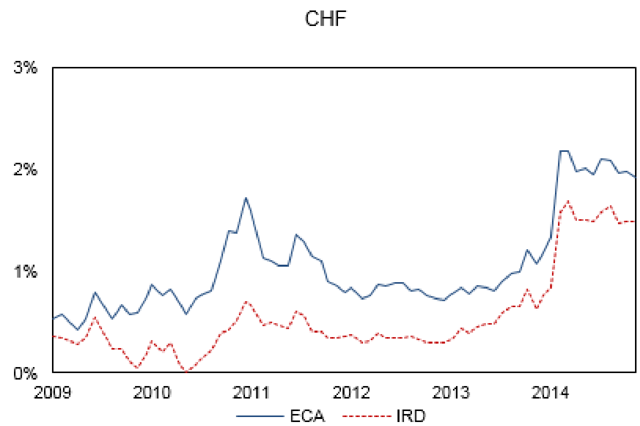
CHF
3%
2%
1%
0%
2009
2010
2011
2012
2013
2014
ECA
IRD

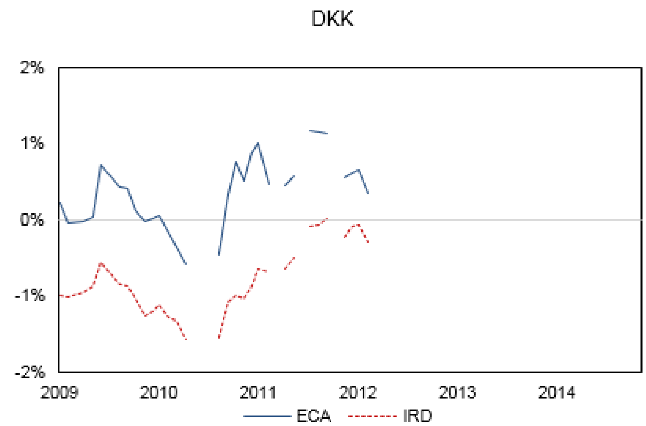
DKK
2%
1%
0%
-1%
-2%
2009
2010
2011
2012
2013
2014
ECA
IRD

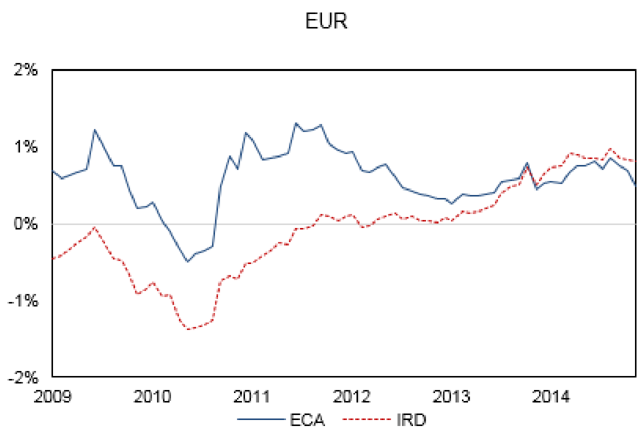
EUR
2%
1%
0%
-1%
-2%
2009
2010
2011
2012
2013
2014
ECA
IRD

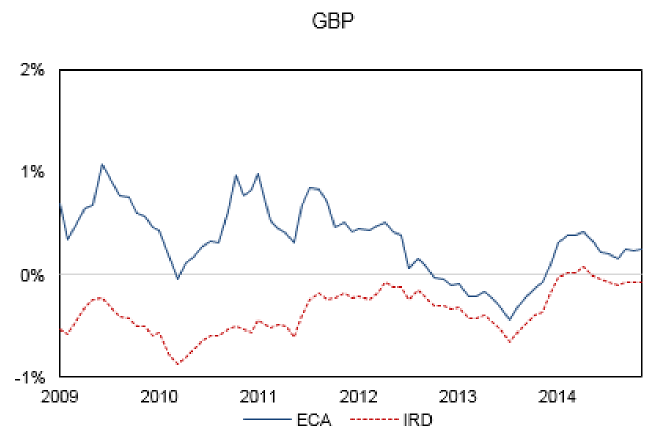
GBP
2%
1%
0%
-1%
2009
2010
2011
2012
2013
2014
ECA
IRD

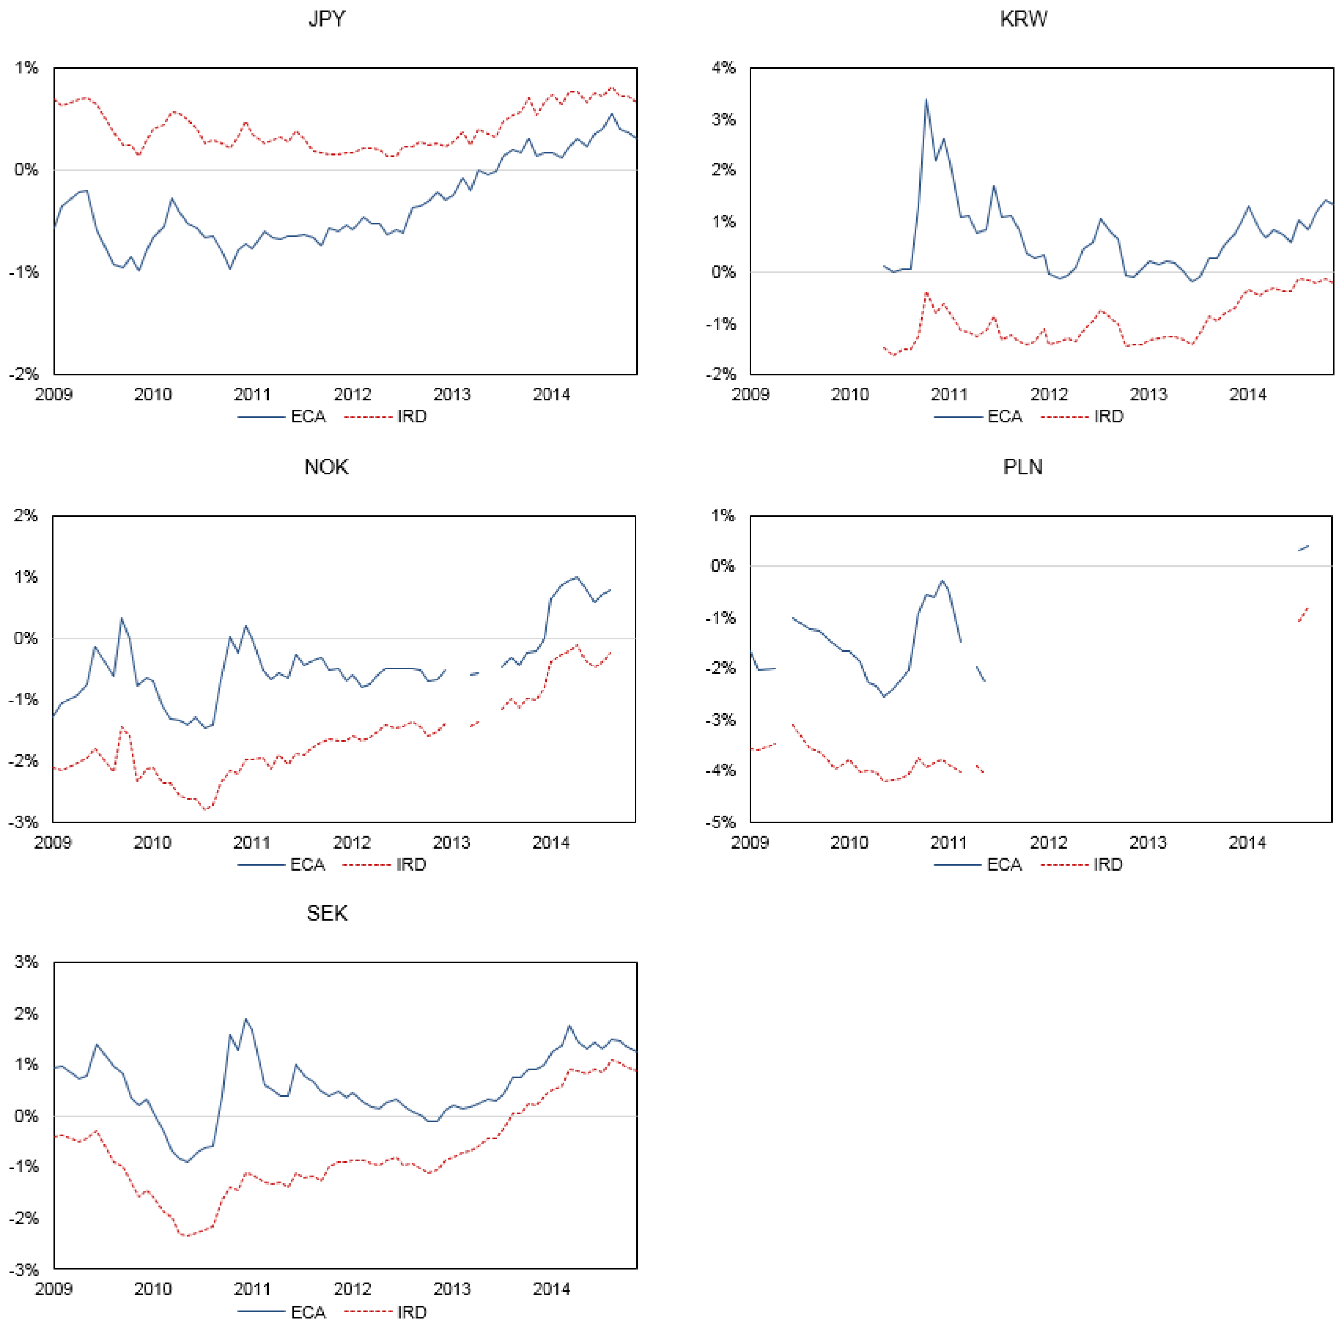

Figure 2: Time series of annualized expected currency appreciation implied by the quanto theory (ECA) and by UIP (IRD).

Table 1: Summary statistics of ECA

This table reports annualized summary statistics (in %) of quanto-based expected currency appreciation (ECA).

	Mean	**Std Dev.**	**Skew**	**Kurtosis**	**Min**	**Max**	**Autocorr.**
			Expected currency appreciation, ECA				
AUD	−1.231	0.723	−0.114	−0.577	−2.550	0.450	0.864
CAD	0.327	0.526	0.909	0.494	−0.526	1.835	0.845
CHF	1.064	0.472	1.147	0.210	0.422	2.176	0.934
DKK	0.331	0.487	−0.097	−0.606	−0.587	1.172	0.762
EUR	0.587	0.398	−0.725	0.799	−0.493	1.300	0.877
GBP	0.326	0.350	−0.103	−0.517	−0.444	1.077	0.894
JPY	−0.337	0.412	0.484	−0.989	−0.978	0.555	0.953
KRW	0.706	0.724	1.455	2.922	−0.182	3.387	0.770
NOK	−0.398	0.622	0.624	0.040	−1.474	0.991	0.877
PLN	−1.340	0.892	0.759	−0.479	−2.554	0.436	0.881
SEK	0.574	0.656	−0.143	−0.340	−0.907	1.885	0.885
Average	0.056	0.569	0.382	0.087	−0.934	1.388	0.867
Pooled	0.056	0.908	−0.500	0.630	−2.554	3.387	

Table 2: Summary statistics of IRD and QRP

This table reports annualized summary statistics (in %) of UIP forecasts (IRD, top panel), and quanto-implied risk premia (QRP, bottom).

	Mean	**Std Dev.**	**Skew**	**Kurtosis**	**Min**	**Max**	**Autocorr.**
				Interest rate differential, IRD			
AUD	−2.815	1.007	−0.104	−1.081	−4.533	−1.168	0.979
CAD	−0.712	0.353	1.121	0.204	−1.133	0.195	0.890
CHF	0.560	0.441	1.501	1.137	0.013	1.690	0.953
DKK	−0.821	0.470	0.298	−0.794	−1.596	0.005	0.915
EUR	−0.056	0.622	−0.282	−0.509	−1.377	0.983	0.977
GBP	−0.352	0.223	−0.098	−0.745	−0.865	0.082	0.925
JPY	0.410	0.206	0.476	−1.229	0.133	0.809	0.909
KRW	−0.973	0.443	0.587	−1.017	−1.614	−0.116	0.877
NOK	−1.596	0.690	0.587	−0.286	−2.798	−0.107	0.955
PLN	−3.422	1.030	2.010	2.733	−4.215	−0.806	0.967
SEK	−0.715	0.905	0.430	−0.421	−2.354	1.105	0.981
Average	−0.954	0.581	0.593	−0.183	−1.849	0.243	0.939
Pooled	−0.954	1.265	−0.952	0.657	−4.533	1.690	
				Quanto-implied risk premium, QRP			
AUD	1.584	0.692	0.546	−0.454	0.666	3.306	0.941
CAD	1.039	0.441	0.509	−0.572	0.309	2.090	0.926
CHF	0.504	0.171	0.663	1.405	0.131	1.023	0.900
DKK	1.153	0.275	0.400	0.336	0.643	1.768	0.788
EUR	0.643	0.556	−0.104	−1.274	−0.315	1.708	0.978
GBP	0.678	0.389	0.270	−1.318	0.207	1.472	0.959
JPY	−0.746	0.295	−0.033	−1.287	−1.287	−0.255	0.945
KRW	1.679	0.589	1.605	2.582	0.944	3.752	0.859
NOK	1.198	0.359	0.876	0.462	0.665	2.194	0.890
PLN	2.083	0.650	0.814	0.026	1.194	3.509	0.868
SEK	1.289	0.616	0.801	0.620	0.371	3.004	0.938
Average	1.009	0.457	0.577	0.048	0.321	2.143	0.908
Pooled	1.009	0.857	−0.107	0.658	−1.287	3.752	

Table 3: Volatilities and correlations of ECA, IRD, and QRP

This Table presents the standard deviations (in %) of, and correlations between, the interest rate differential (IRD), the quanto-implied risk premium (QRP), and expected currency appreciation (ECA), calculated from (14) for each currency i:

$$\mathrm{IRD}_{i,t} = \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} - 1$$
$$\mathrm{QRP}_{i,t} = \frac{Q_{i,t} - F_t}{R^{i}_{f,t} P_t}$$
$$\mathrm{ECA}_{i,t} = \mathrm{QRP}_{i,t} + \mathrm{IRD}_{i,t}.$$

The row labelled "Time series" reports means of the currencies' time-series standard deviations and correlations. The row labelled "Cross section" reports cross-sectional standard deviations and correlations of time-averaged ECA, IRD, and QRP. The row labelled "Pooled" reports standard deviations and correlations of the pooled data. All quantities are expressed in annualized terms.

	$\sigma(ECA)$	$\sigma(IRD)$	$\sigma(QRP)$	$\rho(ECA, IRD)$	$\rho(ECA, QRP)$	$\rho(IRD, QRP)$
AUD	0.723	1.007	0.692	0.727	−0.013	−0.696
CAD	0.526	0.353	0.441	0.558	0.748	−0.134
CHF	0.472	0.441	0.171	0.932	0.355	−0.007
DKK	0.487	0.470	0.275	0.835	0.342	−0.231
EUR	0.398	0.622	0.556	0.476	0.183	−0.777
GBP	0.350	0.223	0.389	0.137	0.822	−0.451
JPY	0.412	0.206	0.295	0.738	0.882	0.333
KRW	0.724	0.443	0.589	0.582	0.792	−0.036
NOK	0.622	0.690	0.359	0.855	0.090	−0.439
PLN	0.892	1.030	0.650	0.780	0.135	−0.514
SEK	0.656	0.905	0.616	0.733	−0.013	−0.690
Time-series	0.569	0.581	0.457	0.669	0.393	−0.331
Cross-section	0.786	1.242	0.751	0.817	−0.305	−0.798
Pooled	0.908	1.265	0.857	0.736	−0.026	−0.696

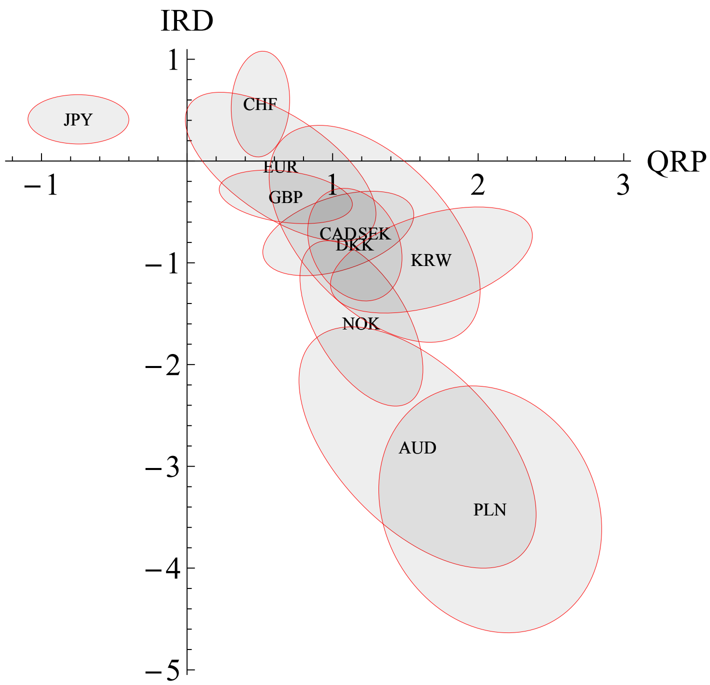

(a) The relationship between QRP and IRD

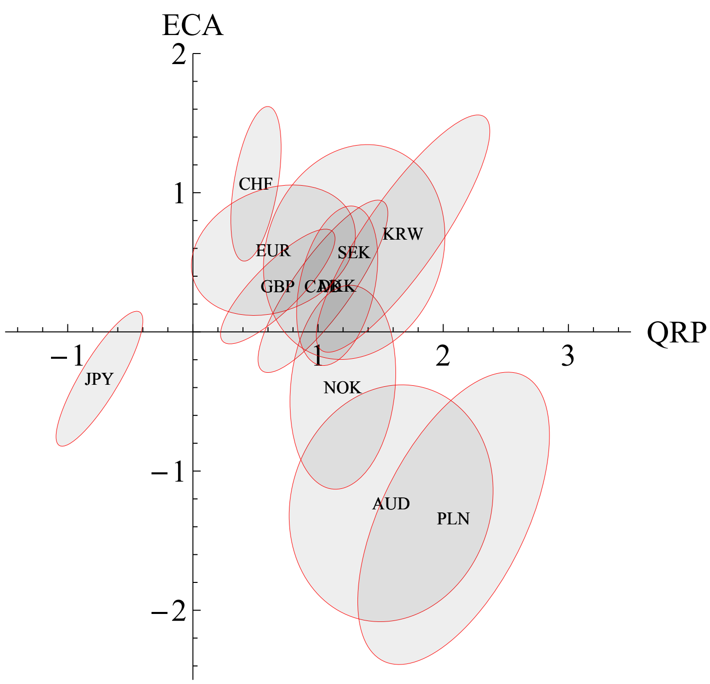

(b) The relationship between QRP and ECA

Figure 3: For each currency, the figures plot mean QRP and IRD (or ECA) surrounded by a confidence ellipse whose orientation reflects the time-series correlation between QRP and IRD (or ECA), and whose size reflects their volatilities. The location and orientation of the ellipses in panel (a) indicate that high interest rates are associated with high quanto-implied risk premia in the cross section and in the time series.

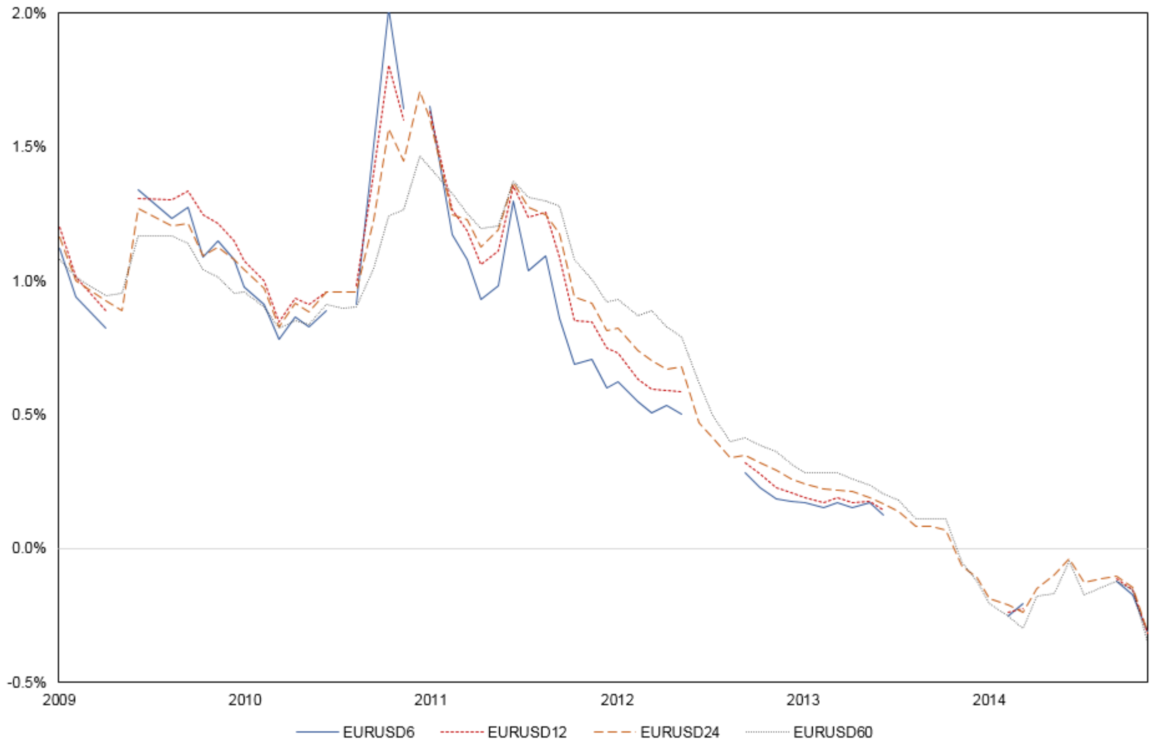

Figure 4: Term structure of the euro-dollar risk premium, as measured by QRP, in the time series for horizons of 6, 12, 24, and 60 months.

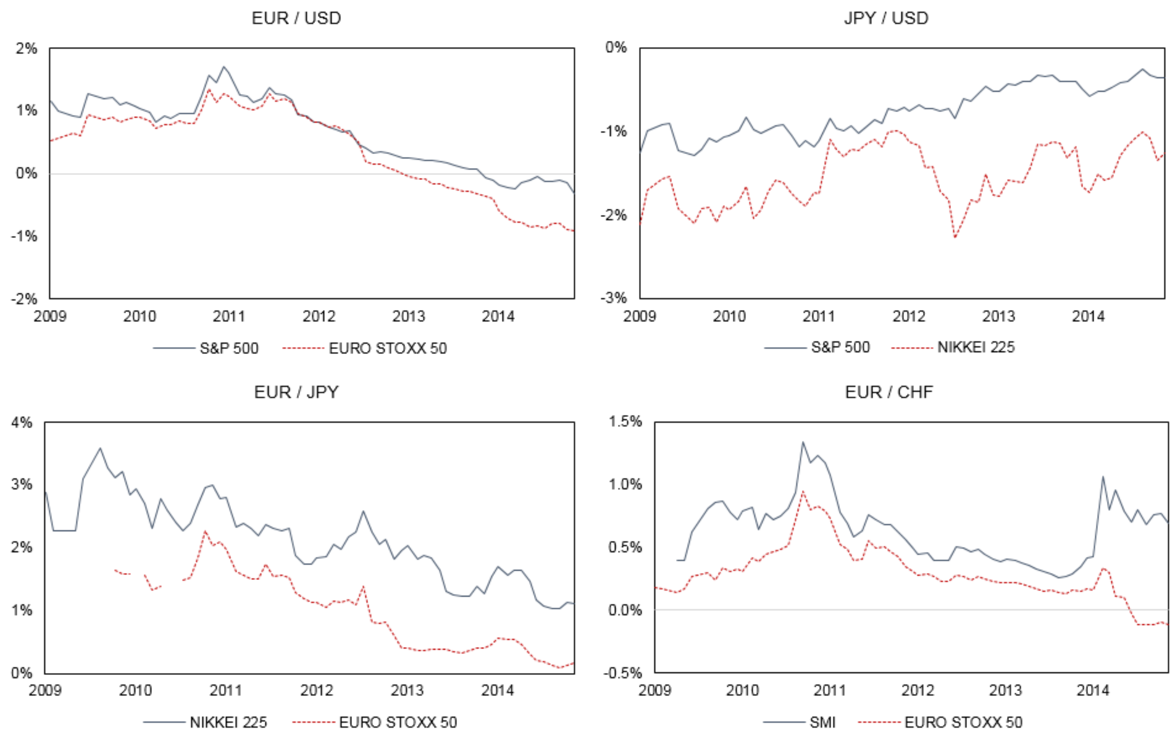

Figure 5: Expected currency appreciation over a 24-month horizon (annualized), as measured by ECA from equation (14), for the EUR-USD, JPY-USD, EUR-JPY, and EUR-CHF currency pairs. Each panel plots ECA for the respective currency pair from the two national perspectives, using quanto contracts on the respective domestic index denominated in the respective foreign currency. The solid blue line plots ECA as perceived by a log investor fully invested in the S&P (top two panels), Nikkei (bottom left panel), and SMI (bottom right panel), respectively. The dashed red line plots the negative of ECA for the same currency pair (inverting the exchange rate) from the perspective of a log investor fully invested in the respective foreign equity index.

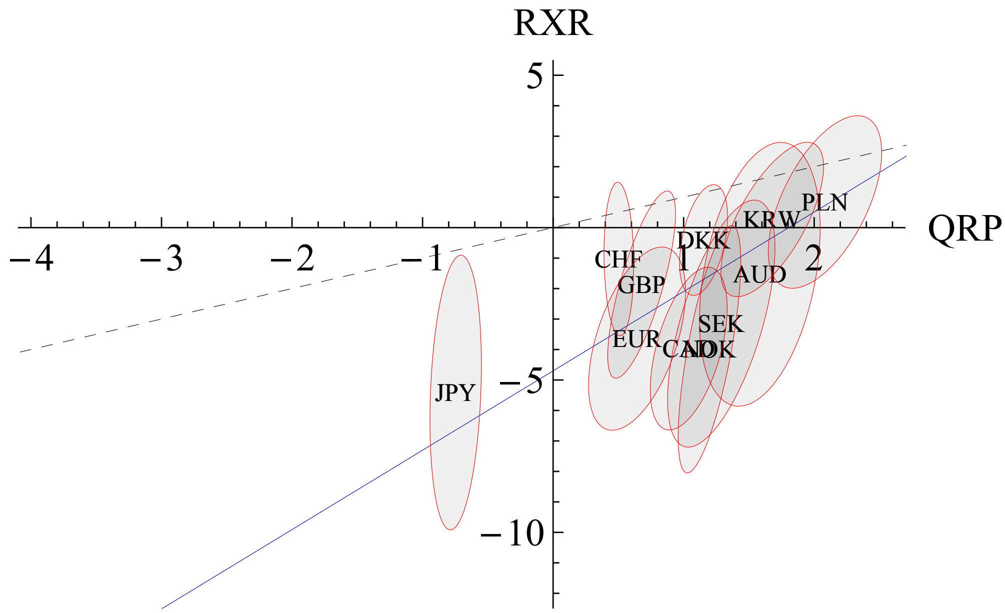

(a) Realized currency excess return against QRP, computed from (14)

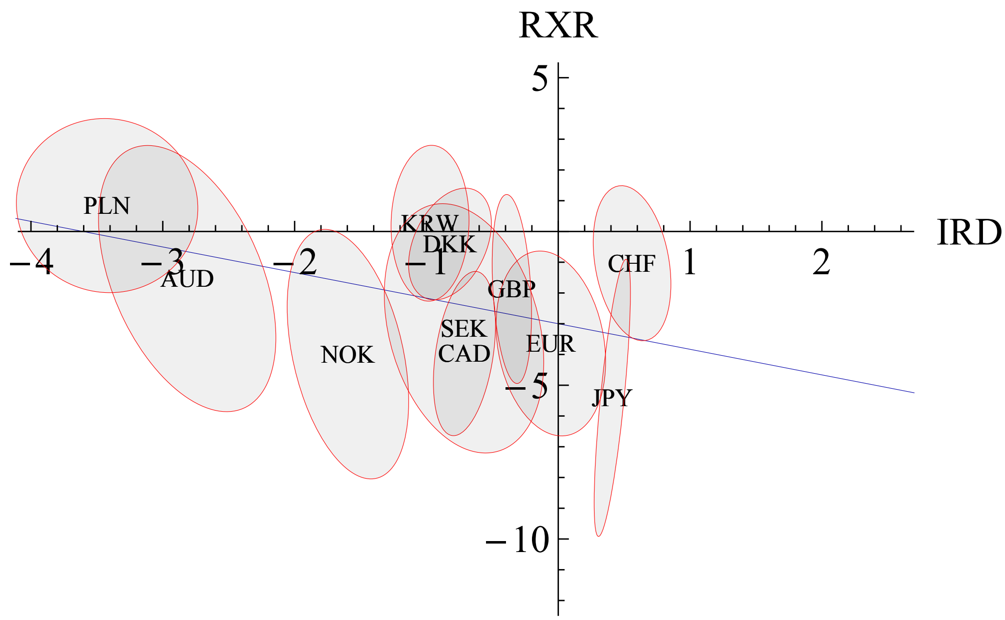

(b) Realized currency excess return against IRD

Figure 6: Realized and expected currency excess return according to (a) the quanto theory and (b) UIP. The centre of each confidence ellipse represents a currency's mean expected and realized currency excess return. In population, each ellipse would contain 20% of its currency's data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse's size reflects their volatilities. Panel (a) shows a dotted 45° line for comparison.

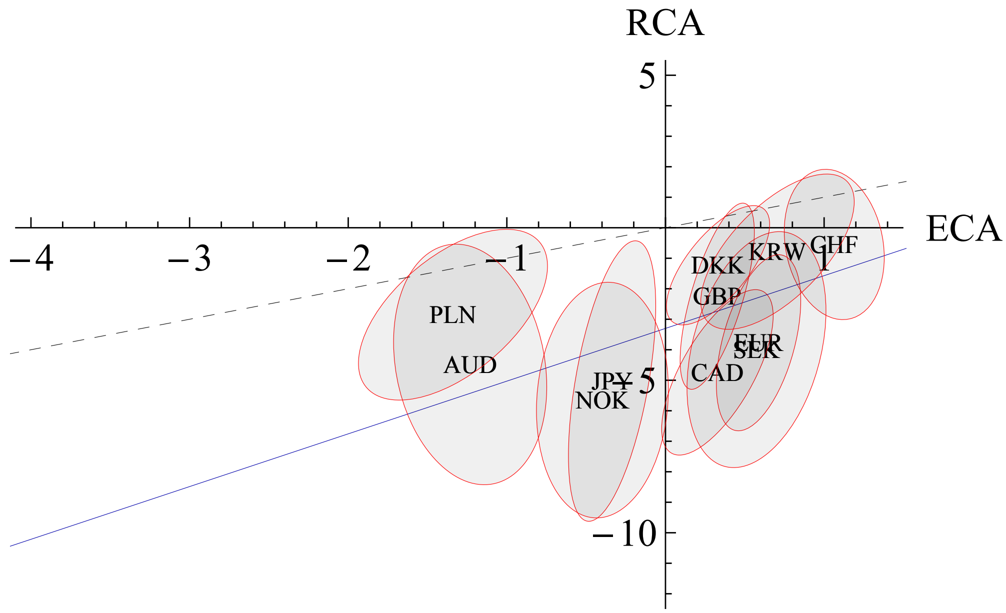

(a) Realized currency appreciation against ECA, computed from (14)

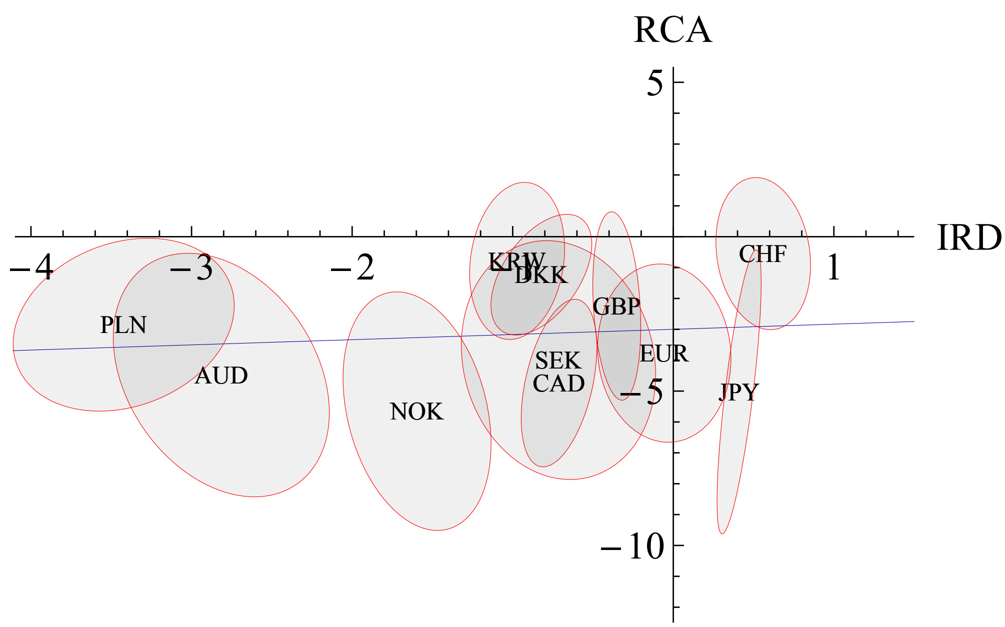

(b) Realized currency appreciation against IRD

Figure 7: Realized and expected currency appreciation according to (a) the quanto theory and (b) UIP. The centre of each confidence ellipse represents a currency's mean expected and realized currency appreciation. In population, each ellipse would contain 20% of its currency's data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse's size reflects their volatilities. Panel (a) shows a dotted 45° line for comparison.

Table 4: Currency excess return forecasting regressions

This Table presents results from three currency excess return forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \quad (22)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \quad (23)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \quad (24)$$

The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, as well as R^2 (in %).

Panel A: Pooled panel regressions			
Regression	(22)	(23)	(24)
α (p.a.)	−0.048	−0.047	−0.030
	(0.020)	(0.019)	(0.014)
β	3.394	2.604	
	(1.734)	(1.127)	
γ	0.769		−0.832
	(1.040)		(0.651)
R^2	19.13	17.43	3.88
Panel B: Panel regressions with currency fixed effects			
β	5.456	4.995	
	(2.046)	(1.565)	
γ	0.717		−1.363
	(1.411)		(1.001)
R^2	22.60	22.03	2.77

Table 5: Currency forecasting regressions

This Table presents results from three currency forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{25}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \tag{26}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{27}$$

The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, as well as R^2 (in %).

Panel A: Pooled panel regressions			
Regression	(25)	(26)	(27)
α (p.a.)	−0.048	−0.045	−0.030
	(0.020)	(0.019)	(0.014)
β	3.394	1.576	
	(1.726)	(1.172)	
γ	1.769		0.168
	(1.045)		(0.651)
R^2	16.01	6.63	0.16
Panel B: Panel regressions with currency fixed effects			
β	5.456	4.352	
	(2.047)	(1.682)	
γ	1.717		−0.363
	(1.414)		(1.007)
R^2	20.56	17.16	0.20

Table 6: Separate return forecasting regressions using QRP and IRD predictors

This table reports the results of running regressions (23), (24), (25), and (27) separately for each currency at the 24-month horizon, and at 6- and 12-month horizons for the euro. We report the OLS estimates along with Hansen–Hodrick standard errors. R^2 are reported in %.

Currency	AUD	CAD	CHF	DKK	EUR	EUR	EUR	GBP	JPY	KRW	NOK	PLN	SEK
Horizon	24m	24m	24m	24m	6m	12m	24m	24m	24m	24m	24m	24m	24m
Panel A: Regression (23): $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^{i}_{f,t} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.062	−0.085	−0.003	−0.052	−0.040	−0.071	−0.060	−0.086	−0.012	−0.068	−0.180	−0.065	−0.106
	(0.071)	(0.042)	(0.038)	(0.022)	(0.056)	(0.052)	(0.030)	(0.031)	(0.090)	(0.034)	(0.061)	(0.026)	(0.048)
β	3.258	4.754	−1.657	4.125	3.702	6.361	4.148	9.217	4.750	4.227	11.860	3.580	5.930
	(3.991)	(3.546)	(6.903)	(1.723)	(6.263)	(5.527)	(3.367)	(3.791)	(10.959)	(1.757)	(4.698)	(0.956)	(3.316)
R^2	12.15	25.39	0.60	17.42	3.17	17.98	25.93	57.48	4.06	46.59	49.96	33.01	38.00
Panel B: Regression (24): $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^{i}_{f,t} = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.091	−0.006	0.001	0.014	−0.015	−0.019	−0.034	−0.043	−0.152	0.007	−0.091	0.005	−0.042
	(0.084)	(0.030)	(0.027)	(0.023)	(0.083)	(0.040)	(0.025)	(0.034)	(0.046)	(0.034)	(0.065)	(0.045)	(0.035)
γ	−2.859	4.135	−2.246	2.147	2.626	1.869	−1.439	−5.564	25.539	0.312	−3.310	−0.118	−1.765
	(2.743)	(3.543)	(3.067)	(2.036)	(7.375)	(6.349)	(3.255)	(6.779)	(8.318)	(3.011)	(3.698)	(1.211)	(2.730)
R^2	19.82	12.30	7.33	13.77	1.23	1.31	3.90	6.93	57.26	0.14	14.39	0.09	7.28
Panel C: Regression (25): $e_{i,t+1}/e_{i,t} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.093	−0.055	0.010	−0.041	−0.055	−0.092	−0.078	−0.082	−0.165	−0.063	−0.185	−0.041	−0.117
	(0.087)	(0.044)	(0.035)	(0.021)	(0.053)	(0.043)	(0.027)	(0.033)	(0.079)	(0.046)	(0.070)	(0.032)	(0.043)
β	0.698	5.291	−1.698	5.252	10.008	12.916	7.321	9.760	−1.348	4.241	11.230	4.736	7.895
	(3.130)	(2.984)	(6.621)	(1.260)	(7.198)	(4.771)	(2.895)	(3.519)	(7.485)	(1.719)	(3.491)	(0.848)	(2.552)
γ	−1.525	6.019	−1.250	3.857	11.447	11.992	4.651	3.094	27.182	1.514	0.253	2.419	2.938
	(2.429)	(2.637)	(3.050)	(1.671)	(8.450)	(4.880)	(2.175)	(3.124)	(8.344)	(2.149)	(2.402)	(1.003)	(1.683)
R^2	9.79	46.74	3.04	48.62	14.42	45.19	33.51	57.29	59.41	48.22	46.61	45.28	39.00
Panel D: Regression (27): $e_{i,t+1}/e_{i,t} - 1 = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.091	−0.006	0.001	0.014	−0.007	−0.019	−0.034	−0.043	−0.152	0.007	−0.091	0.005	−0.042
	(0.084)	(0.030)	(0.027)	(0.023)	(0.041)	(0.040)	(0.025)	(0.034)	(0.046)	(0.034)	(0.065)	(0.045)	(0.035)
γ	−1.859	5.135	−1.246	3.147	3.626	2.869	−0.439	−4.564	26.539	1.312	−2.310	0.882	−0.765
	(2.743)	(3.543)	(3.067)	(2.036)	(7.375)	(6.349)	(3.255)	(6.779)	(8.318)	(3.011)	(3.698)	(1.211)	(2.730)
R^2	9.47	17.78	2.38	25.54	2.32	3.03	0.38	4.77	59.13	2.48	7.57	4.79	1.45

Table 7: Principal components analysis of residuals

This table reports the loadings on the principal components of realized residuals obtained from the quanto theory (top panel) and the fixed-effects specification of regression (23) (bottom panel). In order to limit the impact of missing observations, the residuals are only obtained for the balanced panel of currencies (excluding DKK, KRW, and PLN).

	PC1	**PC2**	**PC3**	**PC4**	**PC5**	**PC6**	**PC7**	**PC8**
				Theory residuals				
AUD	0.520	0.160	0.108	−0.443	−0.273	0.235	0.578	−0.183
CAD	0.311	−0.015	−0.107	−0.257	−0.090	0.458	−0.490	0.606
CHF	0.194	−0.124	0.644	0.344	−0.534	−0.270	−0.067	0.228
EUR	0.243	−0.265	−0.308	0.688	−0.119	0.490	0.127	−0.179
GBP	0.083	−0.471	0.579	−0.104	0.552	0.296	−0.046	−0.176
JPY	0.353	0.741	0.200	0.325	0.397	0.009	−0.145	−0.055
NOK	0.472	−0.194	−0.190	−0.147	−0.099	−0.334	−0.527	−0.532
SEK	0.427	−0.283	−0.238	0.093	0.382	−0.472	0.324	0.446
Explained	61.26%	26.49%	7.26%	2.80%	0.93%	0.53%	0.39%	0.34%
				Regression residuals				
AUD	0.532	0.138	0.019	−0.261	0.665	−0.025	−0.368	−0.227
CAD	0.276	−0.057	−0.175	−0.271	0.248	0.057	0.657	0.566
CHF	0.177	−0.243	0.662	0.273	0.070	−0.594	0.052	0.193
EUR	0.178	−0.291	−0.430	0.732	0.248	−0.004	0.205	−0.244
GBP	−0.086	−0.440	0.489	0.024	0.195	0.714	0.073	−0.082
JPY	0.558	0.539	0.243	0.289	−0.372	0.303	0.154	−0.050
NOK	0.369	−0.451	−0.060	−0.399	−0.409	−0.148	0.229	−0.506
SEK	0.351	−0.384	−0.209	0.068	−0.295	0.144	−0.555	0.516
Explained	65.70%	16.33%	10.65%	3.10%	2.12%	1.20%	0.54%	0.34%

Table 8: Realized covariance regressions

This Table presents results of regressions using the lagged realized covariance of exchange rate movements with the negative reciprocal of the S&P 500 return (RPCL) as a proxy for the currency beta:

$$\mathrm{RPCL}_{i,t} = R^{\$}_{f,t}\left(\sum_{t-h}^{t}\left[\frac{e_{i,s}}{e_{i,s-1}}\left(-\frac{1}{R_s}\right)\right] - \frac{1}{h}\sum_{t-h}^{t}\left(-\frac{1}{R_s}\right)\sum_{t-h}^{t}\frac{e_{i,s}}{e_{i,s-1}}\right),$$

where the summation is over daily returns on trading days s preceding t over a time-frame corresponding to our forecasting horizon, h, so that $\mathrm{RPCL}_{i,t}$ is observable at time t.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma\,\mathrm{RPCL}_{i,t} + \varepsilon_{i,t+1} \quad (35)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{RPCL}_{i,t} + \varepsilon_{i,t+1} \quad (36)$$

We also define a realized covariance measure $\mathrm{RPC}_{i,t}$ that is analogous to the above definition except that the summation is over trading days *following* t over the appropriate time-frame (so that it is not observable until time $t+h$). We test whether risk-neutral covariance forecasts realized covariance by running the following regression.

$$\mathrm{RPC}_{i,t} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \quad (37)$$

We report coefficient estimates for each regression, with standard errors (computed using a nonparametric block bootstrap) in brackets. See Section 2.5 for more detail.

Panel A: Pooled panel regression			
Regression	(35)	(36)	(37)
α (p.a.)	−0.034	−0.047	−0.000
	(0.017)	(0.018)	(0.001)
β		2.798	0.447
		(1.366)	(0.158)
γ	1.307	−0.213	
	(1.111)	(1.193)	
R^2	7.37	17.52	36.56
Panel B: Panel regression with currency fixed effects			
β		4.643	0.330
		(2.006)	(0.168)
γ	1.967	0.387	
	(1.474)	(1.384)	
R^2	9.14	22.27	9.43

Table 9: Beyond the log investor

This table reports the R^2-maximizing univariate, bivariate, 3-variate, and 4-variate specifications in regressions of realized currency excess returns onto combinations of QRP, IRD, the average forward discount $\overline{\text{IRD}}$, and the real exchange rate, q. The table reports standard errors (computed using a nonparametric block bootstrap) in brackets. See Section 2.5 for more detail. The last line reports R^2 (in %).

Panel regressions with currency fixed effects				
Regressor	univariate	bivariate	3-variate	4-variate
QRP, β	4.995	5.646	3.865	3.570
	(1.565)	(1.376)	(1.617)	(1.809)
IRD, γ				-1.096
				(1.460)
$\overline{\text{IRD}}$, δ			-9.763	-8.251
			(3.325)	(3.045)
RER, ζ		-0.408	-0.767	-0.795
		(0.134)	(0.162)	(0.190)
R^2	22.03	35.25	43.28	43.88

Table 10: R^2 of different variable combinations

This table reports the R^2 (in %) from currency excess return forecasting regressions (with currency fixed effects) using all possible univariate, bivariate, 3-variate and 4-variate combinations of the quanto-implied risk premium (QRP), the interest rate differential (IRD), the average interest rate differential ($\overline{\text{IRD}}$), and the real exchange rate (RER).

	univariate	bivariate	3-variate	4-variate
QRP	22.03			
RER	7.88			
IRD	2.77			
$\overline{\text{IRD}}$	1.99			
QRP, RER		35.25		
$\overline{\text{IRD}}$, RER		33.70		
IRD, RER		27.73		
QRP, $\overline{\text{IRD}}$		22.73		
QRP, IRD		22.60		
IRD, $\overline{\text{IRD}}$		2.79		
QRP, $\overline{\text{IRD}}$, RER			43.28	
QRP, IRD, RER			39.54	
IRD, $\overline{\text{IRD}}$, RER			36.40	
QRP, IRD, $\overline{\text{IRD}}$			22.79	
QRP, IRD, $\overline{\text{IRD}}$, RER				43.88

Table 11: Quantos and the real exchange rate

This Table presents results from currency excess return forecasting regressions that extend the baseline results in Table 4 by adding the log real exchange rate to the regressors on the right-hand side. Following Dahlquist and Penasse (2017), we compute the log real exchange rate as $\mathrm{RER}_{i,t} = \log\left(e_{i,t}\frac{P_{i,t}}{P_{\$,t}}\right)$, where $P_{i,t}$ and $P_{\$,t}$ are consumer price indices for country i and the US, respectively, obtained from the OECD.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{38}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \beta\,\mathrm{QRP}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{39}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \gamma\,\mathrm{IRD}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{40}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{41}$$

The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, see Section 2.5 for more detail.

Panel regressions with currency fixed effects				
Regression	(38)	(39)	(40)	(41)
QRP, β	4.316	5.646		
	(1.821)	(1.376)		
IRD, γ	-2.557		-4.726	
	(1.533)		(1.233)	
RER, ζ	-0.604	-0.408	-0.715	-0.311
	(0.192)	(0.134)	(0.190)	(0.161)
R^2	39.54	35.25	27.73	7.88

Table 12: Joint tests of statistical significance

This Table presents results from three currency forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\text{QRP}_{i,t} + \gamma\,\text{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{22}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\text{QRP}_{i,t} + \varepsilon_{i,t+1} \tag{23}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\text{QRP}_{i,t} + \gamma\,\text{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{25}$$

The Table reports p-values of Wald tests of various hypotheses on the regression coefficients. H_0^1 is the hypothesis suggested by Result 2: $\alpha = \gamma = 0$ and $\beta = 1$ in regression (22), $\alpha = 0$ and $\beta = 1$ in regression (23), and $\alpha = 0$ and $\beta = \gamma = 1$ in regression (25). Hypothesis H_0^2 drops the constraint that $\alpha = 0$, and therefore tests our model's ability to predict differences in currency returns but not its ability to predict the absolute level of (dollar) returns. Hypothesis H_0^3 is that QRP is not useful for forecasting. For each Wald test, we report both the asymptotic p-values obtained from the χ^2 distribution and p-values from a bootstrapped small-sample distribution (in the format asymptotic p-value / small-sample p-value).

Panel A: Pooled panel regression			
Regression	(22)	(23)	(25)
H_0^1: $\alpha = \gamma = 0, \beta = 1$	0.029 / 0.357		
H_0^1: $\alpha = 0, \beta = 1$		0.039 / 0.342	
H_0^1: $\alpha = 0, \beta = \gamma = 1$			0.030 / 0.340
H_0^2: $\beta = 1, \gamma = 0$	0.342 / 0.546		
H_0^2: $\beta = 1$		0.155 / 0.299	
H_0^2: $\beta = 1, \gamma = 1$			0.339 / 0.493
H_0^3: $\beta = 0$	0.050 / 0.088	0.021 / 0.097	0.049 / 0.082
Panel B: Panel regression with currency fixed effects			
H_0^2: $\beta = 1, \gamma = 0$	0.029 / 0.256		
H_0^2: $\beta = 1$		0.011 / 0.163	
H_0^2: $\beta = 1, \gamma = 1$			0.029 / 0.238
H_0^3: $\beta = 0$	0.008 / 0.051	0.001 / 0.089	0.008 / 0.051

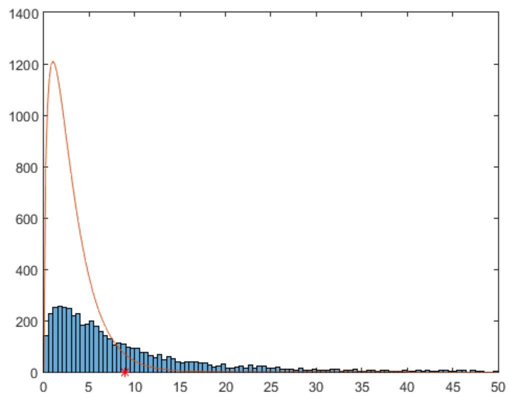

(a) Pooled, H_0^1

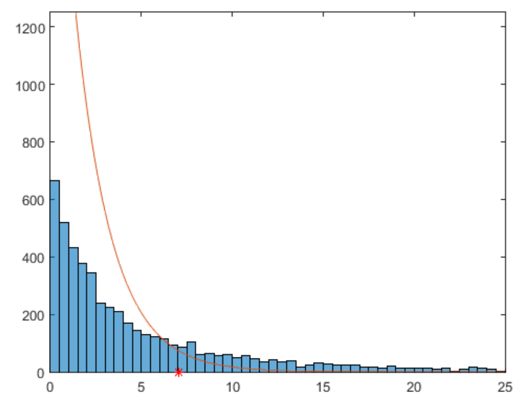

(b) Fixed effects, H_0^2

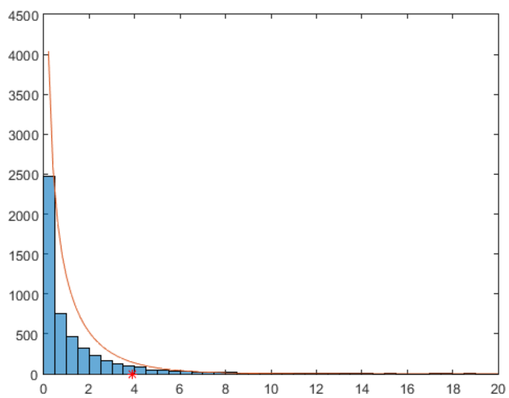

(c) Pooled, H_0^3

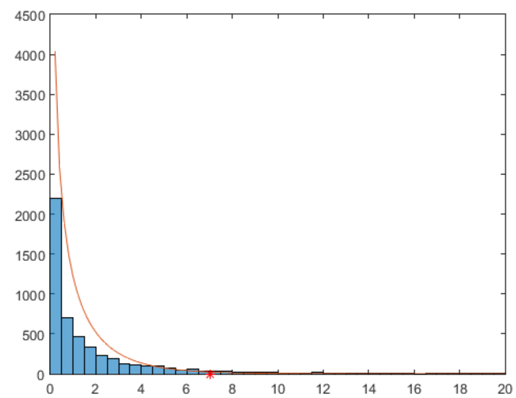

(d) Fixed effects, H_0^3

Figure 8: Histogram of the small-sample distributions of the test statistics for various hypotheses on regression (25). The asymptotic distribution is shown as a solid line. Asterisks indicate the test statistics for the original sample.

Table 13: Out-of-sample forecast performance

We define a dollar-neutral out-of-sample R^2 similar to Goyal and Welch (2008):

$$R^2_{OS} = 1 - \frac{\sum_i \sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_i \sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2},$$

where $\varepsilon^Q_{i,t+1}$ and $\varepsilon^B_{i,t+1}$ denote forecast errors (for currency i against the dollar) of the quanto theory and the benchmark, respectively. We use the quanto theory and three competitor benchmarks to forecast currency appreciation as follows:

$$\text{Theory: } \mathbb{E}^Q_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \text{QRP}_{i,t} + \text{IRD}_{i,t}$$

$$\text{UIP: } \mathbb{E}^U_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \text{IRD}_{i,t}$$

$$\text{Constant: } \mathbb{E}^C_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = 0$$

$$\text{PPP: } \mathbb{E}^P_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \left(\frac{\pi^\$_t}{\pi^i_t}\right)^2 - 1$$

We also report results for the following decomposition of R^2_{OS}, which focusses on dollar-neutral forecast performance for currency i:

$$R^2_{OS,i} = 1 - \frac{\sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2}.$$

The second panel reports R^2_{OS} measures by currency. (All R^2_{OS} measures are reported in %.) The last line of the table reports p-values for a small-sample Diebold–Mariano test of the null hypothesis that the quanto theory and competitor model perform equally well for all currencies.

Benchmark	**IRD**	**Constant**	**PPP**
R^2_{OS}	10.91	9.57	26.05
$R^2_{OS,AUD}$	9.71	0.93	11.42
$R^2_{OS,CAD}$	6.24	6.55	21.31
$R^2_{OS,CHF}$	1.40	16.37	11.43
$R^2_{OS,DKK}$	10.22	7.71	23.36
$R^2_{OS,EUR}$	7.65	5.36	24.56
$R^2_{OS,GBP}$	2.98	9.74	32.35
$R^2_{OS,JPY}$	19.21	9.59	33.74
$R^2_{OS,KRW}$	21.98	17.09	34.71
$R^2_{OS,NOK}$	3.43	12.86	18.97
$R^2_{OS,PLN}$	13.25	8.32	19.62
$R^2_{OS,SEK}$	7.68	5.88	28.22
DM p-value	0.039	0.000	0.000

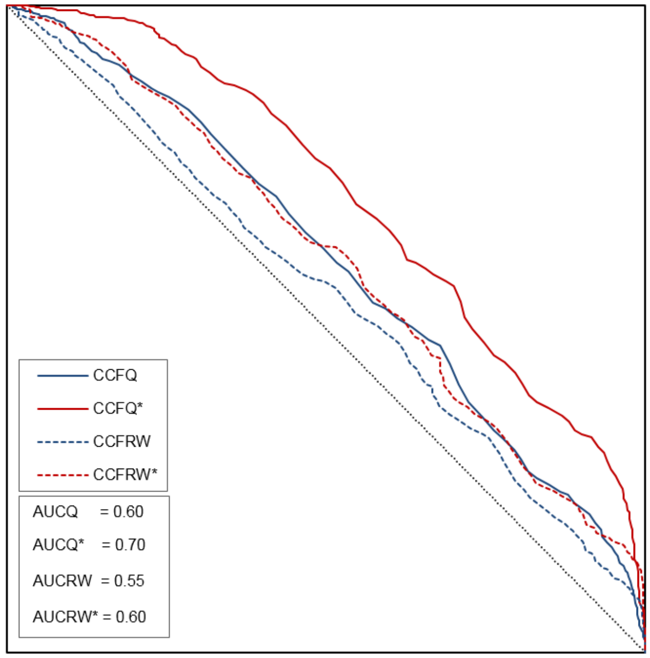

Figure 9: Correct classification frontier (CCF) and AUC statistics for the quanto forecast, and a competitor excess return forecast under which exchange rates follow a random walk.

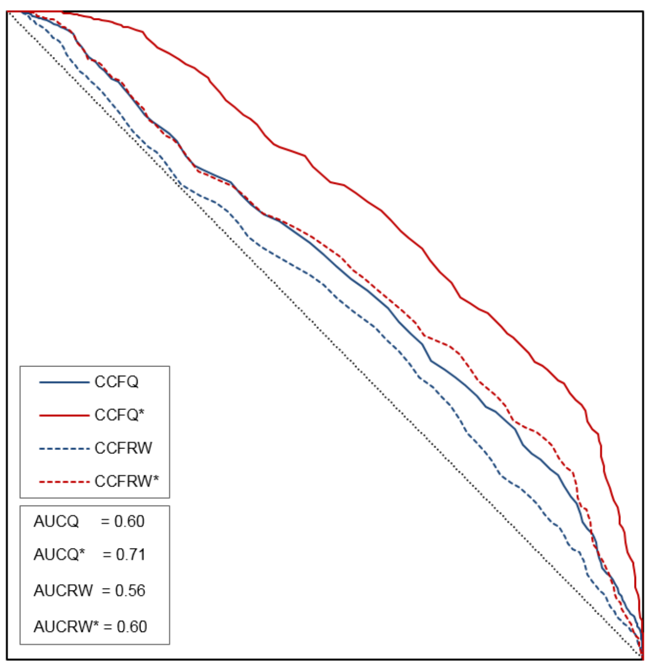

Figure 10: Reverse-conditioned correct classification frontier (CCF) and AUC statistics for the quanto forecast, and a competitor excess return forecast under which exchange rates follow a random walk.

Disclaimer