

**The Quanto Theory of Exchange
Rates**

**By
Lukas Kremens
Ian Martin**

DISCUSSION PAPER NO 769

DISCUSSION PAPER SERIES

August 2017

The Quanto Theory of Exchange Rates

Lukas Kremens

Ian Martin*

August, 2017

Abstract

We present a new, theoretically motivated, forecasting variable for exchange rates that is based on the prices of quanto index contracts, and show via panel regressions that the quanto forecast variable is a statistically and economically significant predictor of currency appreciation and of excess returns on currency trades. We also test the quanto variable's ability to forecast differential currency appreciation out of sample, and find that it outperforms predictions based on uncovered interest parity, on purchasing power parity, and on a random walk.

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

*Both authors are at the London School of Economics. 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, and Dagfinn Rime 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 on, 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

¹Some studies (e.g. Sarno et al., 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.

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

appealing properties mentioned above, but which does not require that one takes the 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 that is new to this paper.) Related perspectives are adopted by Martin (2017) and by Martin and Wagner (2017) to forecast 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. The value of such a contract is sensitive to the correlation between the S&P 500 index and the dollar/euro exchange rate. If the euro is strong relative to the dollar at times when the index is high, and weak 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 and conventional forward prices on the S&P 500 index reveals the risk-neutral covariance between currency i and the index. Quantos therefore allow us to determine which currencies are risky—in that they tend to depreciate in bad times, i.e., when the stock market 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 (a distinction emphasized by Lustig et al. (2011)): according to the theory, the relative importance of the two should be revealed by the behavior of quanto prices.

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 markets are complete). We do not assume the existence of a representative agent, nor do we

³A 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.

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, 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 et al. (2008), Jurek (2014), Della Corte et al. (2016b), Chernov et al. (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.

Related to this, Burnside et al. (2011) 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. Our approach is well adapted to this view of the world: investor concerns about peso events 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 test our approach by running panel currency-forecasting regressions, and find that the quanto predictor variable is strongly significant in both statistical and economic terms. We also show that the quanto predictor variable—equivalently, risk-neutral covariance—substantially outperforms lagged realized covariance as a forecaster of exchange rates; and that it is a strongly significant predictor of future realized covariance.

We conclude by testing the out-of-sample predictive performance of the quanto variable. 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 currency 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 et al. (2011), and independent of what Hassan and Mano (2016) refer to as the dollar trade anomaly. 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 different 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 a random walk forecast. The quanto theory outperforms all three competitors.

1 Theory

We start with the fundamental equation of asset pricing,

$$E_t [M_{t+1} R_{t+1}] = 1, \quad (1)$$

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 E_t for (real-world) expectation 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 will always “think in dollars,” so M_{t+1} will always be the relevant SDF for us. 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), 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 $E_t M_{t+1} = 1/R_{f,t}^{\$}$, as can be seen by setting $R_{t+1} = R_{f,t}^{\$}$; thus (1) can be rearranged as

$$E_t [R_{t+1} - R_{f,t}^{\$}] = -R_{f,t}^{\$} \text{cov}_t [M_{t+1}, R_{t+1}]. \quad (2)$$

Consider a simple currency trade: take a dollar, convert it to foreign currency i , invest at the (gross) currency- i -riskless rate, $R_{f,t}^i$, 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 $R_{f,t}^i e_{i,t+1}/e_{i,t}$; setting $\mathbb{R}_{t+1} = R_{f,t}^i e_{i,t+1}/e_{i,t}$ in (2) and rearranging,⁴ we find that

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

From an empirical point of view, the challenging aspect of the identity (3) is the presence of the unobservable SDF M_{t+1} . It would be extremely convenient if, say, M_{t+1} were constant conditional on time- t information, for then the irritating—because hard to measure—covariance term would drop out and we would recover the prediction of uncovered interest parity (UIP) that $\mathbb{E}_t e_{i,t+1}/e_{i,t} = R_{f,t}^{\$}/R_{f,t}^i$, 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.

Equation (3) 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

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

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

$$\mathbb{E}_t^* \frac{e_{i,t+1}}{e_{i,t}} = \frac{R_{f,t}^{\$}}{R_{f,t}^i}. \quad (5)$$

Unfortunately, as is well known, the UIP prediction fares poorly empirically: the assumption that the covariance term is negligible in (3) (or, equivalently, that the risk-

⁴Unlike most authors in this literature, we prefer to work with true returns, \mathbb{R}_{t+1} , rather than with log returns, $\log \mathbb{R}_{t+1}$, as the latter are only “an approximate measure of the rate of return to speculation,” in the words of Hansen and Hodrick (1980).

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 in particular that the SDF M_{t+1} is highly volatile. In practice, therefore, the risk adjustment term in (3) cannot be neglected: expected currency appreciation depends 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 differential approach that exploits the following observation:

Result 1. *We have the identity*

$$E_t \frac{e_{i,t+1}}{e_{i,t}} = \underbrace{\frac{R_{f,t}^{\$}}{R_{f,t}^i}}_{\text{UIP forecast}} + \underbrace{\frac{1}{R_{f,t}^{\$}} \text{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right)}_{\text{quanto-implied risk premium}} - \underbrace{\text{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right)}_{\text{residual}}, \quad (6)$$

where R_{t+1} is an arbitrary gross return. The asterisk on the first covariance term in (6) indicates that it is computed using the risk-neutral probability distribution.

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

$$E_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) = \frac{1}{R_{f,t}^i}. \quad (7)$$

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:

$$\begin{aligned} \frac{1}{R_{f,t}^{\$}} \text{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right) &\stackrel{(4)}{=} E_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - R_{f,t}^{\$} E_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) \\ &\stackrel{(7)}{=} E_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - \frac{R_{f,t}^{\$}}{R_{f,t}^i}. \end{aligned} \quad (8)$$

Note also that

$$\text{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right) = E_t \left(M_{t+1} R_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) - E_t \left(\frac{e_{i,t+1}}{e_{i,t}} \right). \quad (9)$$

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

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, since it allows R_{t+1} to be an arbitrary return. To make (6) useful for empirical work, we want to choose a return R_{t+1} with two goals in mind. First, the residual term should be small. Second, the middle term (which we label the quanto-implied risk premium for reasons that will become clear in the next section) should be easy to compute.

The 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 goal 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 the remainder of the paper, we will take R_{t+1} to be the return on the S&P 500 index. 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, thereby avoiding the need to estimate it within an inevitably imperfect model.

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 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

$$F_t = E_t^* P_{t+1}. \tag{10}$$

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 $i=1, \dots, N$, that is, one for each of the N currencies in our data set. 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

$$Q_{i,t} = \frac{E_t^* e_{i,t+1} P_{t+1}}{E_t^* e_{i,t+1}}. \quad (11)$$

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

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

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,

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

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

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

Thus 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. Moreover, the risk-neutral covariance term that appears in (6) is directly revealed by the gap between quanto and conventional index forward prices. We label this term the *quanto-*

implied risk premium, since it (partially) measures the currency risk premium up to

the impact of the residual covariance term.

We still need to deal with this troublesome fi term, however. We will initially do so by taking the perspective of an unconstrained, rational investor with log utility whose wealth is fully invested in the S&P 500 index. For such an investor $M_{t+1} = 1/R_{t+1}$, so that the residual term is *exactly* zero. More generally, the quanto-implied risk premium will be informative about currency returns so long as the second covariance term in (6) can plausibly be neglected. A key advantage of our approach relative to the UIP benchmark is that it is much more reasonable to do so than to neglect the covariance term in (3).

The quanto-implied risk premium (for currency i) therefore has a simple interpretation: it is the expected excess return (on currency i) perceived by an unconstrained log investor who chooses to hold the S&P 500 index. The following result summarizes this discussion.

Result 2. *If we take the perspective of a rational investor with log utility whose wealth is fully invested in the index then*

$$E_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{\frac{R_{f,t}^{\$}}{R_{f,t}^i} - 1}_{\text{IRD}} + \underbrace{\frac{Q_{i,t} - F_t}{R_{f,t}^i P_t}}_{\text{QRP}}, \quad (14)$$

where R_{t+1} is the return on the S&P 500 index, P_t is the spot price of the index, F_t is the forward price of the index, and $Q_{i,t}$ is the quanto forward price of the index (where currency i is the quanto currency).

Rearranging, the expected excess return on speculation in currency i equals the quanto-implied risk premium:

$$E_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t}.$$

Equation (14) splits expected currency appreciation into two terms. The fi 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. (Note that a high-interest-rate currency will have a negative interest rate

diff tial, i.e. a negative UIP forecast.) 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 the *quanto forecast*, or as *expected currency appreciation* (ECA).

That said, in our analysis we will in fact allow for the presence of a non-trivial second covariance term. Throughout the paper, for each regression that we run, we report results with (as well as without) currency fi effects; these fi effects will absorb any currency-dependent but time-independent component of the covariance term. In Section 3.1, we consider various further proxies for the term that depend both on currency and time.

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 May 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 May 2017. As we show in Subsection 1.1, quantos are priced using the *joint* risk-neutral distribution of the S&P and the respective exchange rate. Consequently, they cannot be replicated using more widely available data on plain vanilla options on (separately) the S&P and the exchange rate, which each provide *marginal* distributions for the respective asset. We therefore resort directly to quanto contracts, which are traded over-the-counter. Markit reports consensus prices based on quotes received from a wide range of fi ncial 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. After accounting for missing entries in our panel (notably in DKK,

KRW, and PLN: see Figure 1) we have 610 currency-month observations.⁵

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 measure expected exchange rate appreciation if UIP held). Since the financial crisis of 2007-2009, a growing literature (including Du et al. (2016)) has discussed the failure of Covered Interest Parity—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 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. Our measure of expected currency appreciation (the quanto forecast, or ECA) is equal to the sum of IRD and QRP, as shown in equation (14) of Result 2.

Figure 1 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 quanto-implied risk premium, which varies over time and across currencies and whose magnitude is economically significant for all currencies. The quanto-implied risk premium 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). 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 the constituent parts of ECA, namely IRD and QRP. The last two lines of each panel report the statistics averaged across currencies and summary statistics of the pooled data.

The volatility of quanto-implied risk premia 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

⁵Where 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 1.

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) exhibit substantially more volatility 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 interest rate differentials and quanto-implied risk premia is negative when we pool our data ($\rho = -0.694$). 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 et al. (2011). The average time-series (i.e., within-currency) correlation between interest rate differentials and quanto-implied risk premia is more modestly negative ($\rho = -0.275$): 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 interest rate differentials and quanto-implied risk premia ($\rho = -0.802$). According to the quanto theory, therefore, 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). (These results also help to illustrate an important advantage of our approach. When, for example, Lustig et al. (2011) assess the relative importance of cross-sectional and time-series effects, they are forced to split their sample in order to estimate cross-sectional effects without using in-sample information. In contrast, our approach suggests that quanto-implied risk premia should reveal both time-series and cross-sectional dispersion in currency

risk premia in a forward-looking way.) Figure 2 makes the same point graphically by plotting each currency's quanto-implied risk premium over time; for clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK).

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 ($\rho = 0.491$), while the cross-sectional correlation is negative ($\rho = -0.321$). In the time series, therefore, a rise in a given currency's quanto-implied risk premium is associated with a rise in its expected appreciation; whereas in the cross-section, currencies with relatively high quanto-implied risk premia 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 ($\rho = -0.005$). That is, the quanto theory 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.⁶ The orientation of each ellipse illustrates the within-currency time series correlation, while the positions of the different ellipses reveal correlations across currencies. The figure refines 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 S&P-quantos with a maturity of 24 months, due to better data availability. Nonetheless, we observe quantos with a range of maturities for some currencies, including the euro, and we can therefore take a look at 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 period. However, shorter horizons are slightly more volatile resulting

⁶We are interested in the relative sizes of the ellipses, so choose 50% to make the figure readable.

in a downward sloping term structure during times when the level of QRP spikes and the smoother longer horizons lead to an upward sloping term structure following a drop in the level. We also depart from the US-based view and take the perspective of international investors by examining quantos written on foreign equity indices. We explore this in subsection 3.4.

2.1 Return forecasting

We run two sets of panel regressions in which we attempt to forecast, respectively, currency excess returns and currency appreciation. (We report the results of regressions for individual currencies in Tables 12 and 13.) The literature on exchange rate forecasting finds 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 our most aggressive prediction, as expressed in Result 2, via pooled panel regressions. We also report the results of panel regressions with currency fixed effects; by including fixed effects 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 5 and 6 represent our baseline univariate regressions graphically in the same manner as in Figure 3. Figure 5 plots realized currency excess returns (RXR) against QRP and against IRD. 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 6 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 finally run a horse race between the quanto-implied risk premium and interest

rate differential as predictors of currency excess returns ($e_{i,t+1}/e_{i,t} - R_{f,t}^{\$/i}$):

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$/i}}{R_{f,t}^{\$/i}} = \alpha + \beta \frac{Q_{f,t}^i}{P_t} + \gamma \frac{F_{f,t}^i}{R_{f,t}^i} - 1 + \varepsilon_{i,t+1}. \quad (15)$$

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

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$/i}}{R_{f,t}^{\$/i}} = \alpha + \beta \frac{Q_{f,t}^i}{P_t} + \varepsilon_{i,t+1}, \quad (16)$$

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

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$/i}}{R_{f,t}^{\$/i}} = \alpha + \gamma \frac{R_{f,t}^{\$/i}}{R_{f,t}^i} - 1 + \varepsilon_{i,t+1}. \quad (17)$$

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

Table 4 reports the results. The top panel shows coefficient estimates for each regression with and without fixed effects; standard errors (computed by block bootstrap) are shown in parentheses.⁷ The quanto-implied risk premium is positive, economically large, and strongly individually significant in every specification in which it occurs. Moreover, the R^2 values are substantially higher in the two regressions (15) and (16) that feature the quanto-implied risk premium than in the regression (17) in which it does not occur.

The bottom three panels show the p -values associated with Wald tests of various hypotheses on the regression coefficients. In two respects, the regression results are inconsistent with Result 2. The first is that the estimated coefficient β is statistically significantly larger than 1 in regressions (15) and (16) when we include currency fixed effects; and even without fixed effects, we come close to rejecting the null hypothesis $\beta = 1$ in the regression (16) at conventional significance levels. That is, the predictive power of the quanto-implied risk premium is even stronger than the theory predicts.

⁷A large fraction of our sample consists of overlapping observations. In light of the resulting autocorrelation as well as potential cross-sectional correlation, we use a block-bootstrap methodology to compute the covariance matrix of our coefficient estimates. We describe the bootstrap procedure in

more detail in Appendix A.2 and provide some evidence on its robustness to small-sample problems.

This is a relatively minor disappointment; a second, more irritating, inconsistency is that the constant α is significantly smaller than zero. For these two reasons, when we test the joint hypotheses suggested by Result 2 that $\alpha = 0$, $\beta = 1$, and $\gamma = 0$ in (15) and that $\alpha = 0$ and $\beta = 1$ in (16), we are able to reject with p -values of 0.032 and 0.018, respectively.

The significantly negative intercept α indicates that the currencies in our panel underperformed across the board, relative to the prediction of the model: that is, it reflects an unexpectedly strong dollar over our sample period. To remove this dollar effect, we also conduct a joint test of the hypotheses that $\beta = 1$ and $\gamma = 0$ in (15) (and that $\beta = 1$ in (16)) without also testing $\alpha = 0$, and thereby test whether our model forecasts *differential* currency returns in the manner implied by Result 2. At the conventional 5% or even 10% significance level, we cannot reject the hypothesis that the slope coefficients take the theory-implied values. Once we include currency effects, however, we find as previously, that the quanto-implied risk premium is an even stronger predictor of currency returns than the simplest version of the theory—which neglects the second covariance term in (6)—would imply. This suggests that beyond its direct importance in (6), the quanto-implied risk premium may also proxy for the second covariance term. We explore this possibility further in Section 3.1 below.

The bottom panel of the table reports p -values for tests of null hypotheses that the right-hand-side variables are useless, $\beta = \gamma = 0$. We are able to reject the null of no predictability with some confidence for the pooled regressions (15) and (16) (with p -values on the null of 0.057 and 0.013, respectively). In contrast, there is only weak evidence of predictability for the interest rate differential in the pooled regression (17) (p -value of 0.120). Once we include fixed effects, we can strongly reject the null of no predictability for each of the specifications.

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

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} + \gamma \frac{R_{f,t}^* - 1}{R_{f,t}^i} + \varepsilon_{i,t+1}. \quad (18)$$

The coefficient estimates in (18) are mechanically related to those of regression (15): the only difference is that the coefficient estimate on γ is exactly 1 greater in (18) than in (15). Correspondingly, we find identical p -values on the Wald tests of the joint

hypotheses implied by Result 2 in Tables 4 and 5 (up to very small deviations that can be attributed to randomness in the bootstrap). We therefore run (18) not because we are interested in the resulting coefficient estimates, 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,

$$\frac{e_{i,t+1} - 1}{e_{i,t}} = \alpha + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} + \varepsilon_{i,t+1}, \quad (19)$$

and onto interest rate differentials,

$$\frac{e_{i,t+1} - 1}{e_{i,t}} = \alpha + \gamma \left(\frac{R_{f,t}^*}{R_{f,t}^i} - 1 \right) + \varepsilon_{i,t+1}. \quad (20)$$

As previously, we also run the three regressions (18)–(20) with fixed effects.

The regression results are shown in Table 5, which is structured similarly to Table 4. There are two important conclusions that emerge from the table that are not directly implied by Table 4. First, there is little evidence that the interest rate differential helps to forecast currency appreciation on its own, consistent with the large literature that documents the failure of UIP. In the pooled panel, the estimated γ in regression (20) is close to 0 and the p -value on the null of no predictability by IRD is 0.990, and the R^2 is essentially zero. Even with fixed effects, we marginally fail to reject the null of no predictability (p -value of 0.099).

Second, the quanto-implied risk premium makes a substantial difference in terms of R^2 and is on the margin of individual significance in the pooled regressions (18) and (19) (and is strongly significant when fixed effects are included). When we include the quanto-implied risk premium in our pooled regressions instead of the interest rate differential, R^2 increases from roughly zero to 9.11%; and when the two variables are included together, R^2 increases to 17.40%. Moreover, the coefficient estimate on γ increases, in the presence of the quanto-implied risk premium, toward its theoretically predicted value of 1.

2.2 Risk-neutral covariance outperforms realized covariance

Motivated by Result 2, we have focussed our attention thus far on the *risk-neutral* covariance of currencies with stock returns, as captured by the quanto-implied risk premium. It is natural to wonder whether the empirical success of the quanto-implied risk premium merely reflects the fact that currency returns line up with *true* covariances or, equivalently, with currencies' CAPM betas. More formally, from the perspective of the log investor we can also conclude (using (3)) that

$$E_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = R_{f,t}^{\$} \text{cov}_t \left(\frac{e_{i,t+1}}{e_{i,t}}, -\frac{1}{R_{t+1}} \right). \quad (21)$$

Note that it is the real-world covariance that appears in (21). As before, this equation captures the intuition that a currency that tends to depreciate when the S&P 500 declines is risky and hence should appreciate on average, as compensation for this risk.

Various authors have explored the relationship between currency risk premia and currency betas (see, for example, Lustig and Verdelhan (2007), Campbell et al. (2010) and Burnside (2011)). Motivated by this earlier literature, and by equation (21), we define an empirical proxy for currency beta that is based on lagged realized covariance:⁸

$$\text{RPCL}_{i,t} = R_{f,t}^{\$} \frac{1}{h} \sum_{s=t-h}^t \frac{e_{i,s}}{e_{i,s-1}} \left(-\frac{1}{R_s} - \frac{1}{h} \sum_{s=t-h}^t \left(-\frac{1}{R_s} \right) \frac{e_{i,s}}{e_{i,s-1}} \right). \quad (22)$$

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

To compare the predictive performance of lagged realized covariance relative to the quanto-implied risk premium (which equals risk-neutral covariance), we regress currency excess returns onto lagged realized covariance in the panel regression

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \gamma \text{RPCL}_{i,t} + \varepsilon_{i,t+1}. \quad (23)$$

As shown in Table 6, lagged realized covariance positively predicts currency excess returns, but lacks statistical significance with a p -value of 0.141. (The table also

⁸The results are almost identical if we replace $-1/R_{s+1}$ with R_{s+1} in the definition of beta.

contains results for the corresponding regression with currency effects, which delivers similar conclusions.) We also run a horse-race with the quanto-implied risk premium,

$$e_{i,t+1} - R_{f,t}^{\$} = \alpha + \beta \frac{R_{f,t}^{i,t} - F}{R_{f,t}^i - P_t} + \gamma \text{RPCL}_{i,t} + \varepsilon_{i,t+1}, \quad (24)$$

and find that lagged realized covariance is driven out by the quanto-implied risk premium, which is individually significant.⁹ We can reject the hypothesis that $\beta = \gamma = 0$ (p -value of 0.050), but not the hypothesis that $\beta = 1$ and $\gamma = 0$ (p -value of 0.286). Moreover, R^2 more than doubles when we move from (23) to (24); this high R^2 is almost entirely due to the predictive success of the quanto-implied risk premium, since the R^2 of regression (24) is hardly any higher than in the univariate regression (16) of currency excess returns onto the quanto-implied risk premium alone.

We note that implied risk-neutral covariance and realized time-series covariance describe conceptually different objects. Similar to the variance risk premium in equity markets, the difference between expected risk-neutral and expected real-world covariance between exchange rates and equity markets can be seen as a *covariance risk premium*. We define a measure $\text{RPC}_{i,t}$ as in the definition (22) except that the summation is over daily returns on trading days *s* following *t* over the appropriate time-frame. Figure 7 shows the difference between risk-neutral covariance—measured by QRP—and realized covariance measured by RPC. Both measures express covariances over the same horizon of 24 months. With the exception of the yen, this difference is positive for all other currencies over the vast majority of the sample. A positive realized covariance premium is consistent with realized covariance rising in bad times, and the observed difference between implied and realized covariance is therefore consistent with correlation of risky currencies with equity markets rising in crisis times, while the correlations of “safe haven” assets like the Japanese yen with equity markets become more negative. Similarly, a positive covariance risk premium is consistent with volatility risk premia in equity and/or currency markets, i.e. implied volatility exceeding realised volatility for either the S&P 500, the individual currencies, or both. While implied and realized covariance may plausibly differ systematically (as argued above), our approach may outperform forecasts based on realized covariance even in the absence of a covariance

⁹The quanto-implied risk premium is even competitive with (future) realized covariance, although this is obviously not observable in real time. Results available on request.

risk premium. Implied, risk-neutral covariance is a forward-looking measure that can be obtained directly from asset prices, while measures of realized covariance that can be used for forecasting in real-time are necessarily backward-looking. We can therefore ask whether the predictive success of risk-neutral covariance can be attributed to its ability to forecast realized covariance. We run pooled regressions of realized covariance $RPC_{i,t}$ onto the quanto-implied risk premium (that is, risk-neutral covariance):

$$RPC_{i,t} = \alpha + \beta \frac{Q_{i,t} - F_t}{R_{j,t}^* P_t} + \varepsilon_{i,t+1}. \quad (25)$$

The estimated coefficient on the quanto-implied risk premium, $\beta = 0.442$, is strongly significant. We also run regressions with fixed effects, and find that the resulting estimate of β is positive and marginally significant with a p -value of 0.077. Thus the success of the quanto-implied risk premium in forecasting realized covariance is due in part to its ability to forecast persistent differences in realized covariances across currencies.

3 Beyond the log investor

The identity (6) decomposes expected currency appreciation into the UIP forecast (i.e. the interest rate differential), the quanto-implied risk premium, and a conditional covariance term ($-\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$). Thus far, we have either neglected this last term (in our pooled regressions) or allowed for a currency-dependent but time-independent component (in our fixed-eff regressions). We now explore a wider range of explanatory variables that may help to capture time variation in this third term.

3.1 Quantos and the average forward discount

We start by calculating two measures of residuals $\varepsilon_{i,t+1}$ based on the regression (16). These realized 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, $-\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$. We compute (realized) *theory residuals* by imposing

the coefficients implied by Result 2, that is, $\alpha = 0$ and $\beta = 1$; and we compute (realized) *regression residuals* using the estimated coefficients from the specification of regression (16) that includes currency fixed effects (so these regression residuals exclude any potential time-invariant cross-sectional component of the covariance term). The time series of theory and regression residuals are shown in Figure 8 for each currency.

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 factor since it loads positively on all currencies (with the exception of GBP when using regression residuals).¹⁰ It is therefore reminiscent of the ‘dollar’ factor \overline{IRD}_t constructed by Lustig et al. (2014) as the cross-sectional average of IRD. We plot \overline{IRD}_t in Figure 9; it has correlation $\rho = -0.39$ with the first principal component.

To assess the possibility that \overline{IRD}_t may capture the variation in the conditional covariance term that we have chosen to neglect thus far, we run the regression

$$e_{i,t+1} - \frac{R_{f,t}^{\$}}{R_{f,t}^f} = \alpha + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^y P_t} + \delta (\overline{IRD}_t) + \varepsilon_{i,t+1}. \quad (26)$$

The results are reported in Table 8. The estimated coefficient on the quanto-implied risk premium is, as before, individually significant whether or not we include fixed effects. The estimated coefficient on \overline{IRD}_t has the expected sign: it is negative, indicating that, all else equal, if the dollar interest rate is currently low relative to average foreign-currency interest rates, a typical foreign currency is expected to appreciate against the dollar. In other words, the dollar interest rate tends to be low (in relative terms) when the dollar is *currently* strong. This is consistent with the findings of Lustig et al. (2014). That said, \overline{IRD}_t contributes little in terms of R^2 , and the estimated coefficient δ is not significantly different from zero either with or without fixed effects. For reference, we

¹⁰The second principal component, which explains about 26% of the variation, is, loosely speaking, a yen factor.

also run the regression

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \gamma \left(\frac{R_{f,t}^{\$}}{R_{f,t}^i} - 1 \right) + \delta \text{IRD}_t + \varepsilon_{i,t+1}. \quad (27)$$

The coefficients γ and δ are not individually or jointly significant (though, of course, our sample period is considerably shorter than that of Lustig et al. (2014)). The R^2 associated with regression (27) (9.52% pooled and 9.08% with fi effects) are considerably lower than those for regression (15) in which IRD_t is replaced by the quanto-implied risk premium (22.00% and 30.31% for the pooled and fixed-eff regressions, respectively).

3.2 Quantos and the real exchange rate

Dahlquist and Penasse (2017) have shown that the logarithm of the real exchange rate, q , is a successful forecaster of currency returns: currencies with high real exchange rates depreciate on average. Table 9 reports the results of adding q as a regressor to the baseline regressions in Table 4.

Both QRP and the log real exchange rate are statistically significant in every specification in which they occur. When the log real exchange rate is added to regression (17)—which only uses IRD to forecast returns— R^2 increases to 27.11% from 8.52%.

This is lower, however, than the R^2 of 29.94% achieved by QRP in the univariate regression (16). When QRP and the real exchange rate are included in a bivariate regression, R^2 increases to 35.58%; and to 41.48% when IRD is also included.

3.3 Other proxies for the residual covariance term

Table 10 reports the results of regressions of currency excess returns onto the quanto-implied risk premium together with a range of potential proxies for the residual covariance term:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha_i + \frac{Q_{R^i}^{i,t}}{R_{f,t}^i} P^t + \gamma \text{Proxy}_{i,t} + \varepsilon_{i,t+1}.$$

We include currency fi effects to absorb any time-invariant cross-sectional component of the residual term, and constrain the coefficient on the quanto-implied risk premium to equal one in line with the identity (6). We then compare the R^2 achieved

by different proxies. For reference, the proxy-free regression on QRP alone (that is, with $\gamma = 0$) yields an R^2 of 9.21%.

In light of our baseline finding that the coefficient on the quanto-implied risk premium in the unconstrained regressions exceeds the theory-implied value of 1, the proxy reported in the table is QRP itself. The estimated coefficient γ equals 4.954 (as follows, mechanically, from our baseline regression (16)). This estimate is strongly significant, and R^2 rises to 29.94%, indicating that a substantial component of the unobserved (within-currency) residual covariance term is captured by the quanto-implied risk premium.

The next two proxies we consider are (motivated by the prior literature on interest differentials) IRD and $\overline{\text{IRD}}$. Neither is as successful as QRP in terms of individual significance or in terms of R^2 .

Further proxies are motivated by the thought that the estimated coefficients on QRP—which are significantly larger than one in our baseline regressions—may simply reflect the fact that risk aversion of the marginal investor is larger than one. From the perspective of an investor with power utility and relative risk aversion γ who holds the market, the SDF is proportional to $R_{t+1}^{-\gamma}$. If, say, γ equals two, then the residual covariance term $-\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$ is proportional to the real-world covariance whose empirical proxy is $\text{RPCL}_{i,t}$, as defined in (22). We therefore consider $\text{RPCL}_{i,t}$ as a proxy for the residual covariance term.¹¹ Similarly, we report results using $\text{RPCL}(3.5)_{i,t}$ which corresponds to the case in which $\gamma = 3.5$, and is defined by

$$\text{RPCL}(3.5)_{i,t} = R_{f,t} \left(\frac{e_{i,s+1}}{e_{i,s}} - \frac{1}{R_{s+1}^{2.5}} - \frac{1}{T-t} - \frac{1}{R_{s+1}^{2.5}} \frac{e_{i,s+1}}{e_{i,s}} \right)$$

Both measures of RPCL achieve sizeable improvements in R^2 relative to the baseline without a residual proxy, but both fall well short of the explanatory power of the quanto-implied risk premium.

The first column of the table reports results for the log real exchange rate, motivated by Dahlquist and Penasse (2017) and by the regressions discussed in the previous subsection; it does not contribute much in terms of R^2 and the estimated coefficient is not significantly different from zero (p -value = 0.100).

¹¹The coefficient of proportionality can be shown to be close to one, but we are neglecting the fact that it is time-varying.

3.4 A change of perspective

We have argued that the forward price of the US stock index quantoed into (say) euros reveals the expected appreciation of the euro versus the dollar, as perceived by a log investor whose portfolio is fully invested in the US stock market. This logic can be inverted: the forward price of a European index quantoed into dollars reveals the expected appreciation of the dollar versus the euro, as perceived by a log investor whose portfolio is fully invested in the European market.

Recall Result 2 for the expected appreciation of the euro versus the dollar,

$$E_t \frac{e_{e,t+1}}{e_{e,t}} - 1 = \frac{R^{\$}}{R_{f,t}^{\$}} - 1 + \text{QRP}_{e,t} \quad (28)$$

(To reiterate, a positive value indicates that the euro is expected to strengthen against the dollar.) The corresponding expression for the expected appreciation of the dollar versus the euro, from the perspective of a log investor who is fully invested in the Euro Stoxx 50, is

$$E_t \frac{1/e_{e,t+1}}{1/e_{e,t}} - 1 = \frac{R^e}{R_{f,t}^e} - 1 + \text{QRP}_{\$,t} \quad (29)$$

where $\text{QRP}_{\$,t}$ is obtained from conventional and dollar-denominated quanto forwards on the Euro Stoxx 50. If the right-hand side of the above equation is positive, the dollar is expected to appreciate against the euro, which would imply a fall in $e_{e,t}$.

If our approach is empirically sensible, the output of these two approaches ought to give results that are roughly consistent with one another. Thus if the forward price of the S&P 500 quantoed into euros implies that the dollar is expected to appreciate by 2% according to the right-hand side of equation (28), we would hope to find that the forward price of the Euro Stoxx 50 index quantoed into dollars implies that the euro is expected to appreciate by about -2% according to the right-hand side of equation (29).

We must however keep in mind Siegel's "paradox" (Siegel, 1972), which is the

observation that by Jensen's inequality

$$\mathbb{E}_t \frac{e_{e,t+1}}{e_{e,t}} \geq \left(\mathbb{E}_t \frac{1/e_{e,t+1}}{1/e_{e,t}} \right)^{-1} \quad \text{or, equivalently,} \quad \log \mathbb{E}_t \frac{e_{e,t+1}}{e_{e,t}} \geq -\log \mathbb{E}_t \frac{1/e_{e,t+1}}{1/e_{e,t}}.$$

This fact implies that even if (28) and (29) held perfectly, we would have (after approximating $\log(1+x) \approx x$, as is reasonable for empirically relevant values of ECA)

$$\text{ECA}_{e,t} \geq -\text{ECA}_{\$,t}.$$

The gap between the two sides of this inequality corresponds to a convexity correction whose size is determined by the amount of conditional variation in the random variable $e_{i,t+1}$. If the exchange rate is lognormal,¹² $\log(e_{i,t+1}/e_{i,t}) \sim N(\mu_t, \sigma^2)$, we have

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

Figure 10 implements the above calculations for the EUR-USD, JPY-USD, EUR-JPY, and EUR-CHF currency pairs. In the top-left, bottom-left and bottom-right panels, the solid, blue line depicts the ECA_e measure against the respective other currency, while the dashed, red line shows $-\text{ECA}_{i,t}$ against the euro (we flip the sign on the “inverted” series for readability). For the JPY-USD currency pair in the top-right panel, the lines show $\text{ECA}_{\text{¥}}$ (solid, blue) and $-\text{ECA}_{\text{¥}}$ against the yen (dashed, red), respectively. In each case, the two measures are very strongly correlated over time, and the solid, blue line is above the dashed, red line. The direction of the gaps between the measures is therefore consistent with the Jensen's inequality correction one would expect to see, if the currency forecasts from equation (14) are in fact exact measures of expected currency appreciation.

Moreover, given that annual exchange rate volatilities are on the order of 10%, the

¹²For more general distributions of log currency appreciation the correction term involves all even cumulants, since $\log \left(\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} \right) - \log \left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1} = \mathbf{c}(1) + \mathbf{c}(-1) = 2 \sum_{n \text{ even}} \frac{\kappa_n}{n!}$, where $\mathbf{c}(\cdot)$ and κ_n denote, respectively, the cumulant-generating function and the n th cumulant of log exchange rate appreciation. For an earlier treatment of cumulants in the context of exchange rates, see Backus et al. (2001).

sizes of the gaps between the two ECA measures are broadly quantitatively consistent with the Jensen’s inequality correction. Consider the EUR-CHF pair in the top-right panel as an illustrative example: 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 removal of the floor prompted a spike in the volatility of the currency pair.

Taken together, we view these results as supportive of our approach.

4 Out-of-sample prediction

We now test the predictive success of 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 significantly negative estimated intercept in our pooled panel regression (18)), 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, $(Q_{i,t} - F_i)/(R_{f,t}^i P_i)$, and interest rate differential, $R_{f,t}^s/R_{f,t}^i - 1$, 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 currency forecasts by fixing the coefficients that appear in (18) at their theoretical values: $\alpha = 0$, $\beta = 1$, $\gamma = 1$.

We compare these predictions to those of three natural competitor models that also make a priori predictions, and so do not require estimation of parameters: UIP (which predicts that currency appreciation should offset the interest rate differential, on average), a random walk (which makes the constant forecast of zero currency appreciation), and relative purchasing power parity (which predicts that currency appreciation should offset the inflation differential, on average).

To compare the forecast accuracy of the model to those of the benchmarks, we

define a dollar-neutral R^2 -measure similar to Goyal and Welch (2008):

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

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^2_{OS} versus each of the benchmarks.

The results of this exercise are reported in Table 11. Given the data of 11 currencies sampled, we form predictions on the $N(N-1)/2 = 55$ dollar-neutral currency differentials. The quanto theory outperforms each of the three competitors: when the competitor model is UIP, we find that $R^2_{OS} = 11.60\%$; and when it is relative PPP, we find $R^2_{OS} = 26.91\%$. In our sample, the toughest benchmark is the random walk (i.e., constant) forecast, consistent with the finding of Rossi (2013). Nonetheless, the quanto theory easily outperforms it, with $R^2_{OS} = 10.48\%$. Figure 11 plots the cumulative outperformance versus the competitors and shows that this accrues relatively smoothly over our sample period.

To get a sense for whether our positive results are driven by a small subset of the currencies, Table 11 also reports the results of splitting the R^2 measure currency-by-currency: for each currency i , we define

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

This quantity is positive for all i and all competitor benchmarks B : that is, 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 et al. (1997), and find that the outperformance is strongly significant.

5 Conclusion

The literature on exchange rate forecasting can be understood through the lens of an identity (3) that equates expected currency appreciation to an interest rate differential (the UIP forecast) plus a residual term that depends on the covariance between the currency in question and the stochastic discount factor. UIP holds if and only if this residual term is negligible. In reality, however, there is considerable evidence that currencies with high interest rates earn high risk premia, so the residual term is *not* negligible and UIP does not hold.

In this paper, we have derived a more general identity that opens a new line of attack for empirical work on currency forecasting. Our identity (6) features an arbitrary dollar return, R_{t+1} , that can be freely chosen to optimize the identity's usefulness for empirical work. One wants to choose R_{t+1} such that the second and third components of expected currency appreciation in (6) are each either observable or plausibly negligible. (The first component, the UIP forecast, is manifestly observable.) There is a tradeoff in doing so, because some natural choices of R_{t+1} deal with—in the sense of making observable or negligible—one of the two terms but not the other. In order to deal with both terms simultaneously, we set R_{t+1} equal to the return on the S&P 500 index. We argue that it is then much more reasonable to view the residual covariance term in (6) as small than to view the covariance term in (3) as small.

As a simple and intuitive benchmark, and to motivate our empirical work, we adopt the perspective of a rational investor with log utility who chooses to hold the S&P 500 index; then the residual covariance term in (6) is exactly zero. (In contrast, the residual in (3) is zero under the much less reasonable assumption of risk-neutrality.) The remaining terms in the identity are observable, given interest rates, S&P 500 forward prices, and S&P 500 quanto forward prices, and the fact that the predictive variables are based only on asset prices has the benefit that, in principle, we can generate currency forecasts at high frequency.

In the data, the quanto-implied risk adjustments are correlated with interest rate differentials, so the quanto theory is consistent with the presence of the carry trade. Overall, taking into account the combination of interest rate differentials and risk adjustments, the quanto forecasts indicate that the returns to high interest currencies are partially attenuated by their tendency to depreciate against low interest currencies.

We also find that our measure of currency risk premia exhibits substantially more cross-currency variation than within-currency variation, consistent with the evidence presented by Hassan and Mano (2016).

We test the theory by estimating currency-forecasting panel regressions, and find in most of our specifications that the quanto forecasting variable is significant both in statistical terms (the standard error is small) and in economic terms (the point estimate is large).

As the quanto forecasting variable (for a given currency) measures the risk-neutral covariance between that currency and the S&P 500 index, our results are related to a literature that has used lagged realized covariances as a measure of currency risk (for example Lustig and Verdelhan, 2007; Campbell et al., 2010). We show that the quanto variable successfully forecasts future realized covariances, and drives out lagged realized covariances in multivariate currency-forecasting regressions.

Beyond its implication that the quanto-implied risk premium should forecast currency excess returns, the most optimistic “log investor” version of the theory—in which we entirely neglect the residual covariance term in (6)—makes the more unusual prediction that the coefficient on the quanto-implied risk premium should take a *specific numerical value* (namely, 1); in several of our specifications, the estimated coefficient on the quanto variable is even larger than our model predicts, to the extent that we are formally able to reject this strongest version of the theory.

It follows that our assumption that the residual covariance term is negligible is too strong, so that there is scope for improving our forecasts by including proxies for the covariance term. Since we also find that the coefficient on the quanto-implied risk premium is significantly larger than one when we allow for fixed effects, we can reject the possibility that the covariance term varies by currency but not across time so that there is scope to enhance our forecasts via proxies that captures time variation in the residual covariance term, and hence in currency risk premia.

To that end, we conduct a principal component analysis of the error residuals in realized currency returns. The first principal component explains around two thirds of the variation in residuals, and loads roughly equally on all currencies. It is also reasonably strongly correlated with the average forward discount variable of Lustig et al. (2014), which is based on average interest rate differentials. We therefore add the average forward discount to our baseline specification. The quanto variable remains

significant both with and without fi effects; the average forward discount variable itself is significant in the pooled regression but not with fi effects, and in either case it only contributes a modest increase in R^2 . Other proxies, including the log real exchange rate, realized covariances, and interest rate differentials, fare little better; in fact, we find that the most successful proxy for the residual covariance term is the quanto-implied risk premium itself.

We conclude by analyzing the out-of-sample performance of our approach. If we constrain the coefficient on the quanto-implied risk premium to equal one—that is, if we adopt the perspective of the log investor—we end up with a formula for expected currency appreciation that has no free parameters. Although we are able to reject this formula in sample, the fact that it has no free parameters make it well suited to out-of-sample prediction. We test the formula’s ability to forecast differential currency appreciation out of sample and find that it outperforms three benchmark models that also have no free parameters, namely UIP, PPP, and the random walk forecast.

References

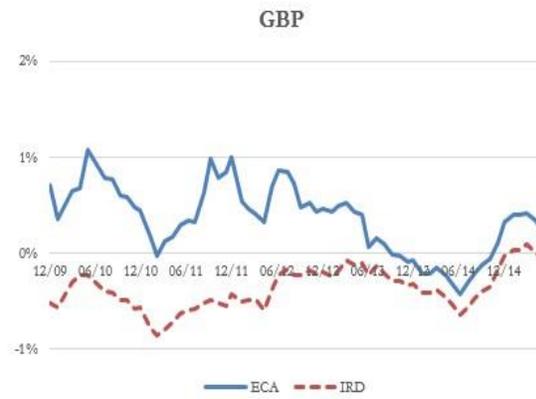
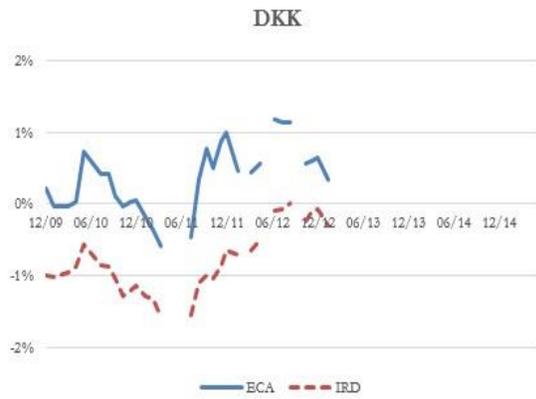
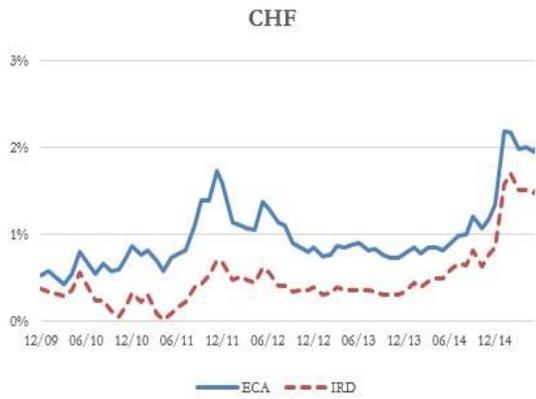
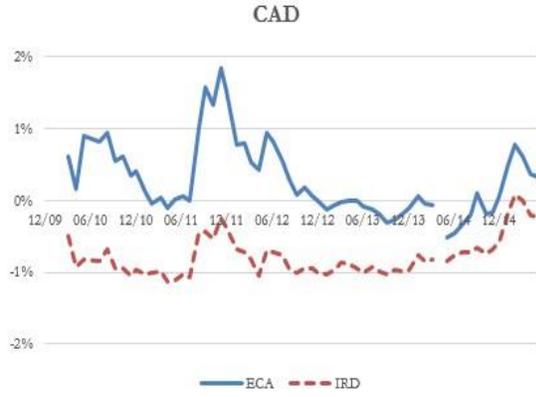
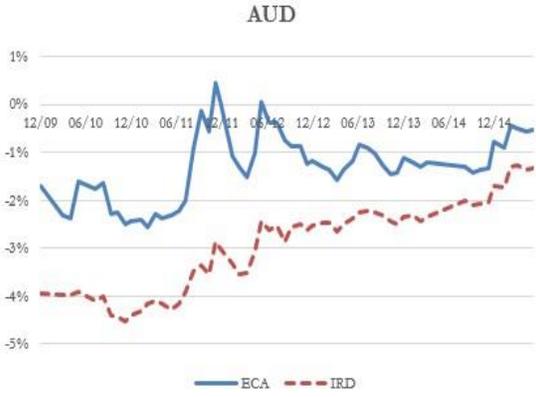
- D. K. Backus, S. Foresi, and C. I. Telmer. Affine term structure models and the forward premium anomaly. *Journal of Finance*, 56(1):279–304, 2001.
- P. Barroso and P. Santa-Clara. Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 50(5):1037–1056, 2015.
- M. K. Brunnermeier, S. Nagel, and L. H. Pedersen. Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23:313–347, 2008.
- C. Burnside. The cross section of foreign currency risk premia and consumption growth risk: Comment. *American Economic Review*, 101(7):3456–3476, 2011.
- C. Burnside, M. Eichenbaum, I. Kleshchelski, and S. Rebelo. Do peso problems explain the returns to the carry trade? *Review of Financial Studies*, 24(3):853–891, 2011.
- J. Y. Campbell, K. Serfaty-De Medeiros, and L. M. Viceira. Global currency hedging. *The Journal of Finance*, 65(1):87–121, 2010.
- M. Chernov, J. Graveline, and I. Zviadadze. Crash risk in currency returns. *Journal of Financial and Quantitative Analysis*, forthcoming, 2016.
- M. Dahlquist and J. Penasse. The missing risk premium in exchange rates. *Working Paper*, 2017.
- P. Della Corte, T. Ramadorai, and L. Sarno. Volatility risk premia and exchange rate predictability. *Journal of Financial Economics*, 20(1), 2016a.
- P. Della Corte, L. Sarno, M. Schmeling, and C. Wagner. Sovereign risk and currency returns. *Working Paper*, 2016b.
- F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(2):253–263, 1995.
- W. Du, A. Tepper, and A. Verdelhan. Deviations from covered interest parity. *Working Paper*, 2016.
- E. F. Fama. Forward and spot exchange rates. *Journal of Monetary Economics*, 14: 319–338, 1984.

- E. Farhi and X. Gabaix. Rare disasters and exchange rates. *The Quarterly Journal of Economics*, 131(1):1–52, 2016.
- A. Goyal and I. Welch. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21:1455–1508, 2008.
- L. P. Hansen and R. J. Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5):829–853, 1980.
- D. Harvey, S. Leybourne, and P. Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13:281–291, 1997.
- T. A. Hassan. Country size, currency unions, and international asset returns. *The Journal of Finance*, 68(6):2269–2308, 2013.
- T. A. Hassan and R. C. Mano. Forward and spot exchange rates in a multi-currency world. *Working Paper*, 2016.
- J. W. Jurek. Crash neutral currency carry trades. *Journal of Financial Economics*, 113(3):325–347, 2014.
- H. Lustig and A. Verdelhan. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1):89–117, 2007.
- H. N. Lustig, N. L. Roussanov, and A. Verdelhan. Common risk factors in currency markets. *Review of Financial Studies*, 24(11), 2011.
- H. N. Lustig, N. L. Roussanov, and A. Verdelhan. Countercyclical currency risk premia. *Journal of Financial Economics*, 111(3):527–553, 2014.
- R. C. Mano. Exchange rates upon sovereign default. *Working Paper*, 2013.
- I. W. R. Martin. The forward premium puzzle in a two-country world. *Working Paper*, 2013.
- I. W. R. Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433, 2017.

- I. W. R. Martin and C. Wagner. What is the expected return on a stock? *Working Paper*, 2017.
- R. Meese and K. Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24, 1983.
- L. Menkhoff, L. Sarno, M. Schmeling, and A. Schrimpf. Carry trades and global foreign exchange volatility. *Journal of Finance*, 67(2):681–718, 2012.
- A. Patton, D. N. Politis, and H. White. Correction to automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 28:372–375, 2009.
- D. N. Politis and H. White. Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23:53–70, 2004.
- B. Rossi. Exchange rate predictability. *Journal of Economic Literature*, 51(4):1063–1119, 2013.
- L. Sarno, P. Schneider, and C. Wagner. Properties of foreign exchange risk premiums. *Journal of Financial Economics*, 105(2):279–310, 2012.
- J. J. Siegel. Risk, interest rates, and the forward exchange. *Quarterly Journal of Economics*, 86:303–309, 1972.
- A. Verdelhan. A habit-based explanation of the exchange rate risk premium. *Journal of Finance*, 65(1):123–145, 2010.

A Appendix

A.1 Tables and Figures



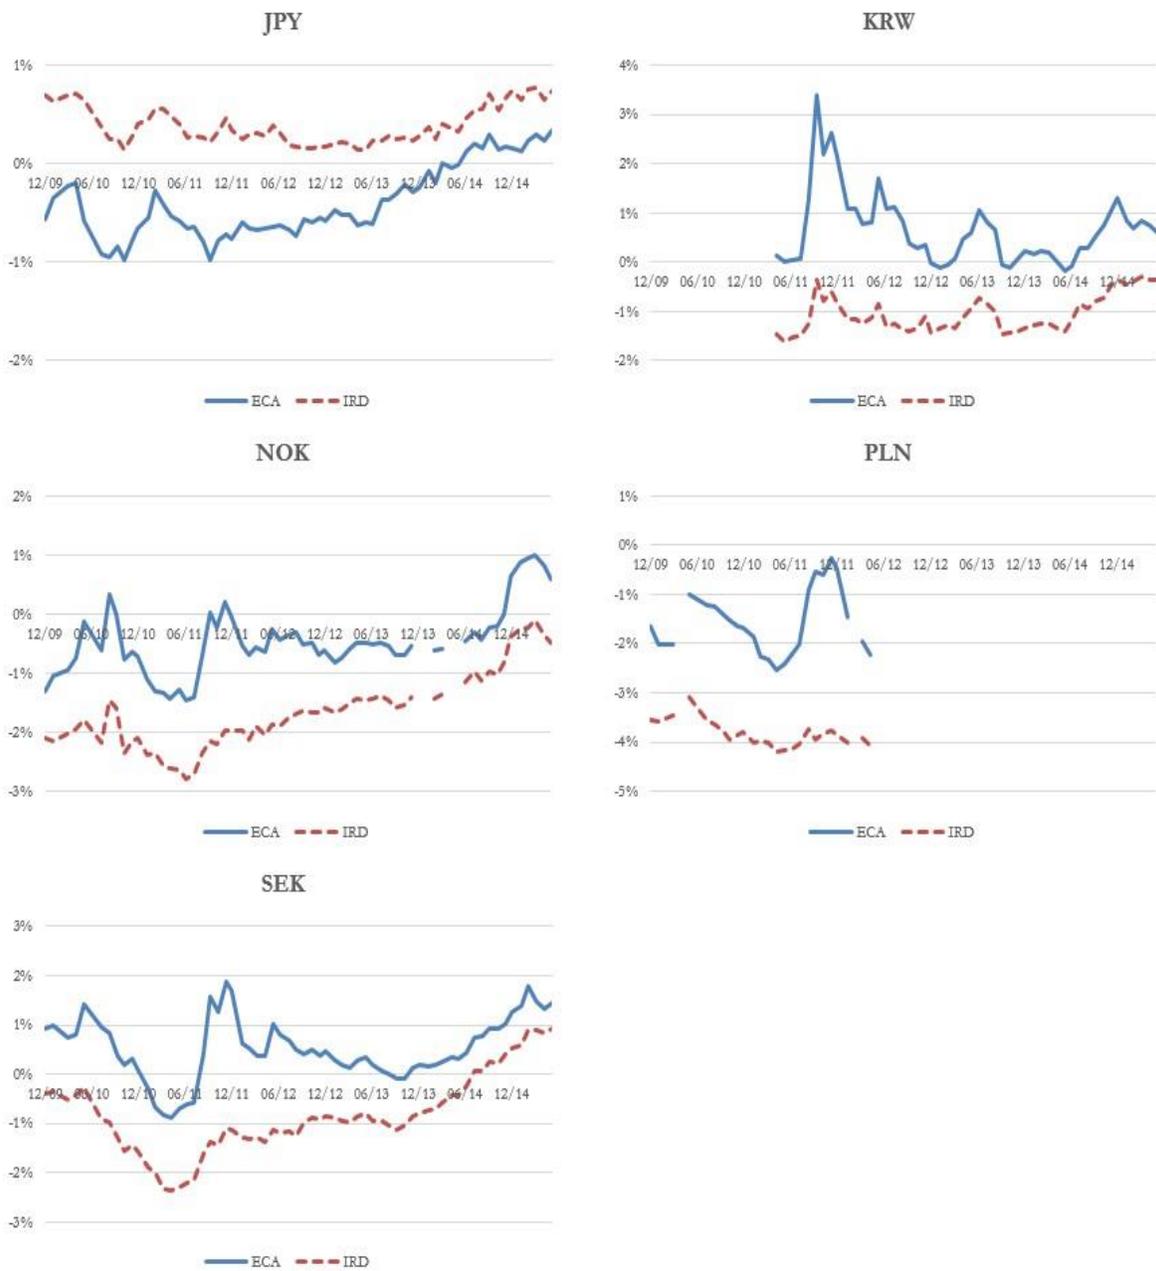


Figure 1: Time series of annualized expected currency appreciation implied by the quanto theory (ECA) and by UIP (IRD).

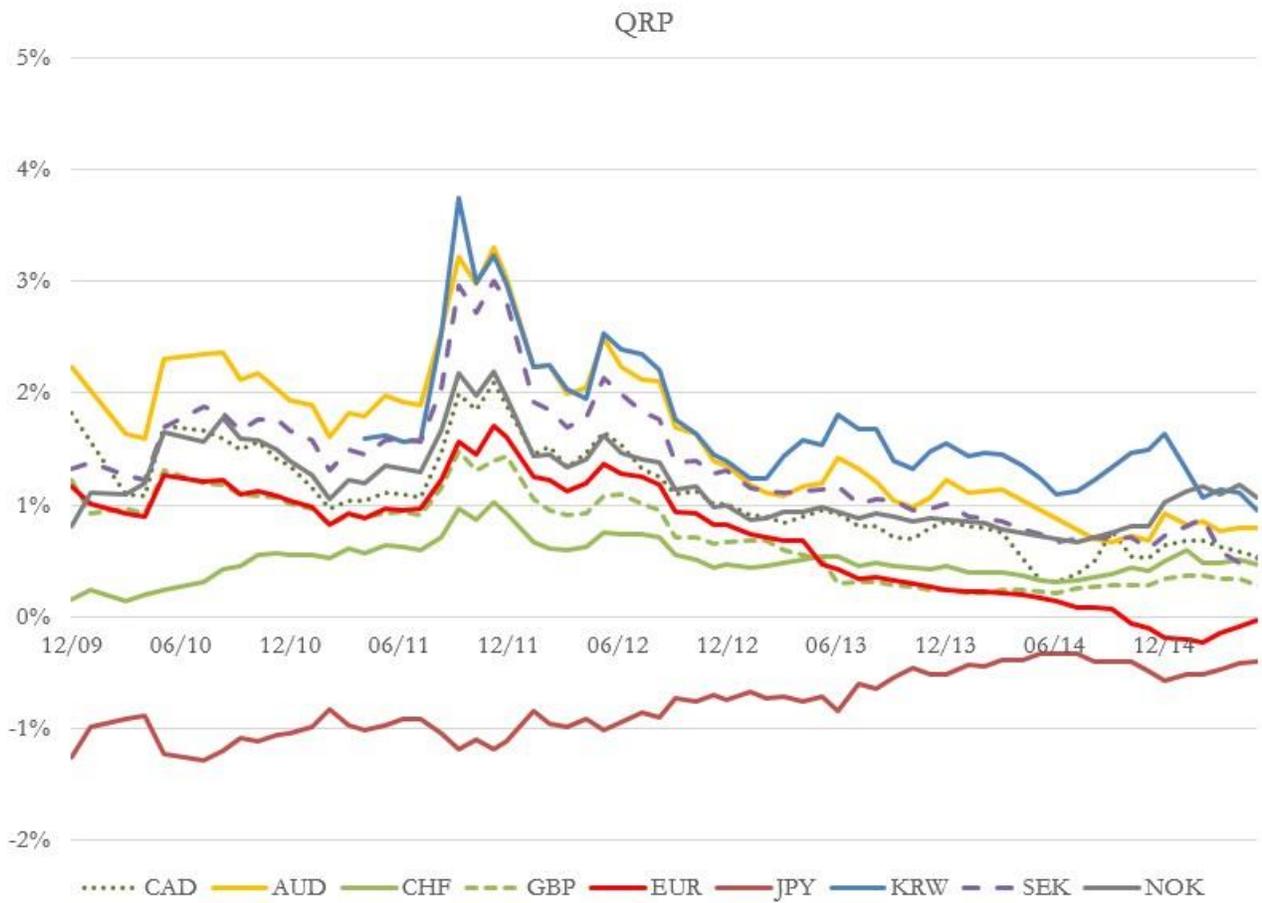


Figure 2: Currency risk premia, as measured by QRP, in the time series. For clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK).

Table 1: Summary statistics of ECA

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

ECA	Mean	Std Dev.	Skew	Kurtosis	Min	Max	Autocorr.
AUD	-1.299	0.712	0.036	-0.376	-2.550	0.450	0.851
CAD	0.304	0.539	1.028	0.604	-0.526	1.835	0.844
CHF	0.991	0.405	1.484	1.840	0.422	2.176	0.914
DKK	0.331	0.487	-0.097	-0.606	-0.587	1.172	0.762
EUR	0.577	0.411	-0.651	0.557	-0.493	1.300	0.880
GBP	0.335	0.362	-0.169	-0.651	-0.444	1.077	0.894
JPY	-0.394	0.368	0.530	-0.847	-0.978	0.346	0.943
KRW	0.662	0.743	1.641	3.317	-0.182	3.387	0.766
NOK	-0.456	0.577	0.694	0.647	-1.474	0.991	0.869
PLN	-1.523	0.723	0.838	0.223	-2.554	0.354	0.839
SEK	0.510	0.638	-0.033	-0.083	-0.907	1.885	0.872
Average	0.003	0.542	0.482	0.420	-0.934	1.361	0.858
Pooled	0.003	0.904	-0.516	0.553	-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).

IRD	Mean	Std Dev.	Skew	Kurtosis	Min	Max	Autocorr.
AUD	-2.949	0.930	-0.136	-1.094	-4.533	-1.273	0.975
CAD	-0.772	0.288	1.244	0.858	-1.133	0.077	0.842
CHF	0.484	0.359	2.028	4.147	0.013	1.690	0.935
DKK	-0.821	0.470	0.298	-0.794	-1.596	0.005	0.915
EUR	-0.128	0.587	-0.304	-0.351	-1.377	0.912	0.974
GBP	-0.374	0.218	0.025	-0.540	-0.865	0.082	0.918
JPY	0.385	0.192	0.642	-0.916	0.133	0.775	0.895
KRW	-1.053	0.379	0.684	-0.762	-1.614	-0.297	0.832
NOK	-1.661	0.639	0.641	0.153	-2.798	-0.107	0.952
PLN	-3.692	0.632	3.616	15.662	-4.215	-0.853	0.761
SEK	-0.846	0.802	0.422	0.058	-2.354	0.912	0.977
Average	-1.039	0.500	0.832	1.493	-1.849	0.175	0.907
Pooled	-1.039	1.250	-1.001	0.442	-4.533	1.690	

QRP	Mean	Std Dev.	Skew	Kurtosis	Min	Max	Autocorr.
AUD	1.650	0.679	0.466	-0.433	0.666	3.306	0.936
CAD	1.077	0.437	0.405	-0.572	0.309	2.090	0.921
CHF	0.507	0.177	0.602	1.096	0.131	1.023	0.902
DKK	1.153	0.275	0.400	0.336	0.643	1.768	0.788
EUR	0.706	0.527	-0.204	-1.134	-0.238	1.708	0.975
GBP	0.708	0.388	0.135	-1.323	0.207	1.472	0.957
JPY	-0.780	0.280	0.034	-1.230	-1.287	-0.329	0.938
KRW	1.715	0.602	1.525	2.201	0.944	3.752	0.859
NOK	1.205	0.366	0.811	0.281	0.665	2.194	0.890
PLN	2.169	0.625	0.842	0.033	1.207	3.509	0.852
SEK	1.357	0.587	0.909	0.839	0.478	3.004	0.929
Average	1.042	0.449	0.539	0.009	0.339	2.136	0.904
Pooled	1.042	0.885	-0.246	0.684	-1.287	3.752	

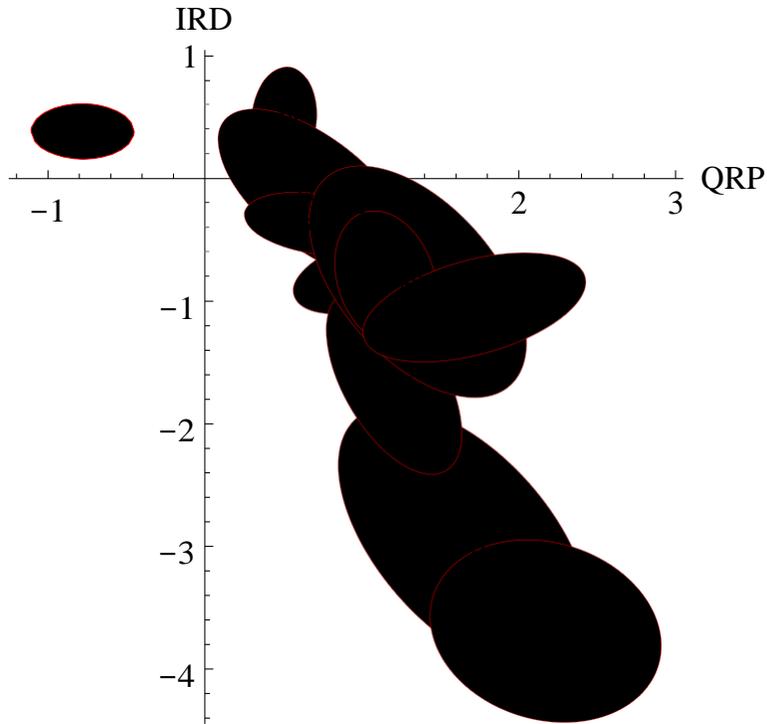
Table 3: Volatilities and correlations of ECA, IRD, and QRP

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 :

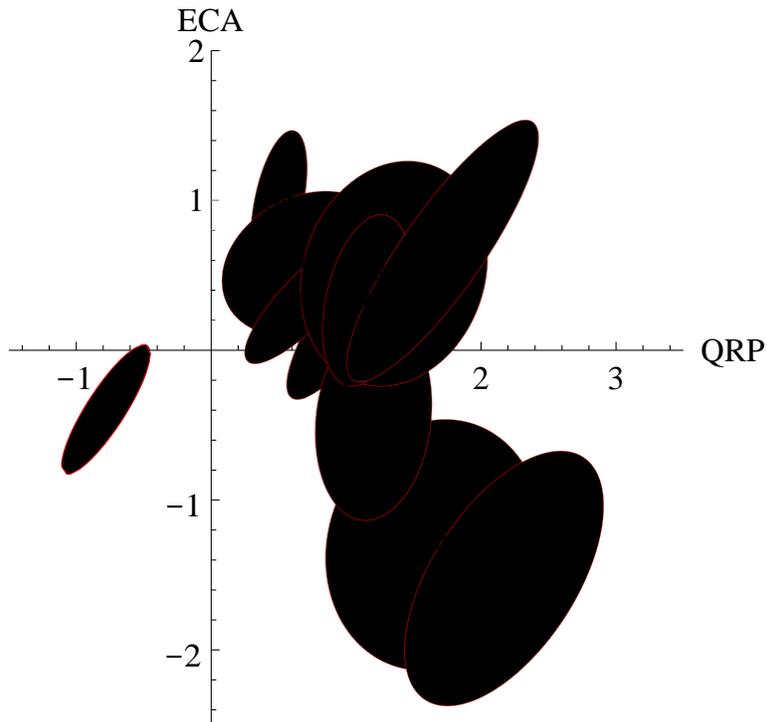
$$\begin{aligned} \text{IRD}_{i,t} &= \frac{R_{i,t}^{\$}}{D_{i,t}} - 1 \\ \text{QRP}_{i,t} &= \frac{Q_{i,t} - F_t}{R_{i,t}^{\$} P_t} \\ \text{ECA}_{i,t} &= \text{QRP}_{i,t} + \text{IRD}_{i,t}. \end{aligned}$$

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.

	$\sigma(\text{ECA})$	$\sigma(\text{IRD})$	$\sigma(\text{QRP})$	$\rho(\text{ECA}, \text{IRD})$	$\rho(\text{ECA}, \text{QRP})$	$\rho(\text{IRD}, \text{QRP})$
AUD	0.712	0.930	0.679	0.687	0.106	-0.649
CAD	0.539	0.288	0.437	0.588	0.846	0.066
CHF	0.405	0.359	0.177	0.900	0.466	0.033
DKK	0.487	0.470	0.275	0.835	0.342	-0.231
EUR	0.411	0.587	0.527	0.487	0.237	-0.733
GBP	0.362	0.218	0.388	0.179	0.834	-0.394
JPY	0.368	0.192	0.280	0.666	0.858	0.190
KRW	0.743	0.379	0.602	0.591	0.862	0.100
NOK	0.577	0.639	0.366	0.823	0.140	-0.447
PLN	0.723	0.632	0.625	0.582	0.568	-0.338
SEK	0.638	0.802	0.587	0.690	0.145	-0.616
Time-series	0.542	0.500	0.449	0.639	0.491	-0.275
Cross-section	0.819	1.298	0.779	0.823	-0.321	-0.802
Pooled	0.904	1.250	0.885	0.724	-0.005	-0.694



(a) The relationship between QRP and IRD



(b) The relationship between QRP and ECA

Figure 3: QRP plotted against IRD and against ECA. The figure plots each currency at its mean QRP and IRD (or ECA), surrounded by a confidence ellipse that would contain 50% of that currency's data points under Normality. The orientation of each ellipse reflects the time-series correlation between QRP and IRD (or ECA) for that currency, while the size reflects the volatilities of the two measures.

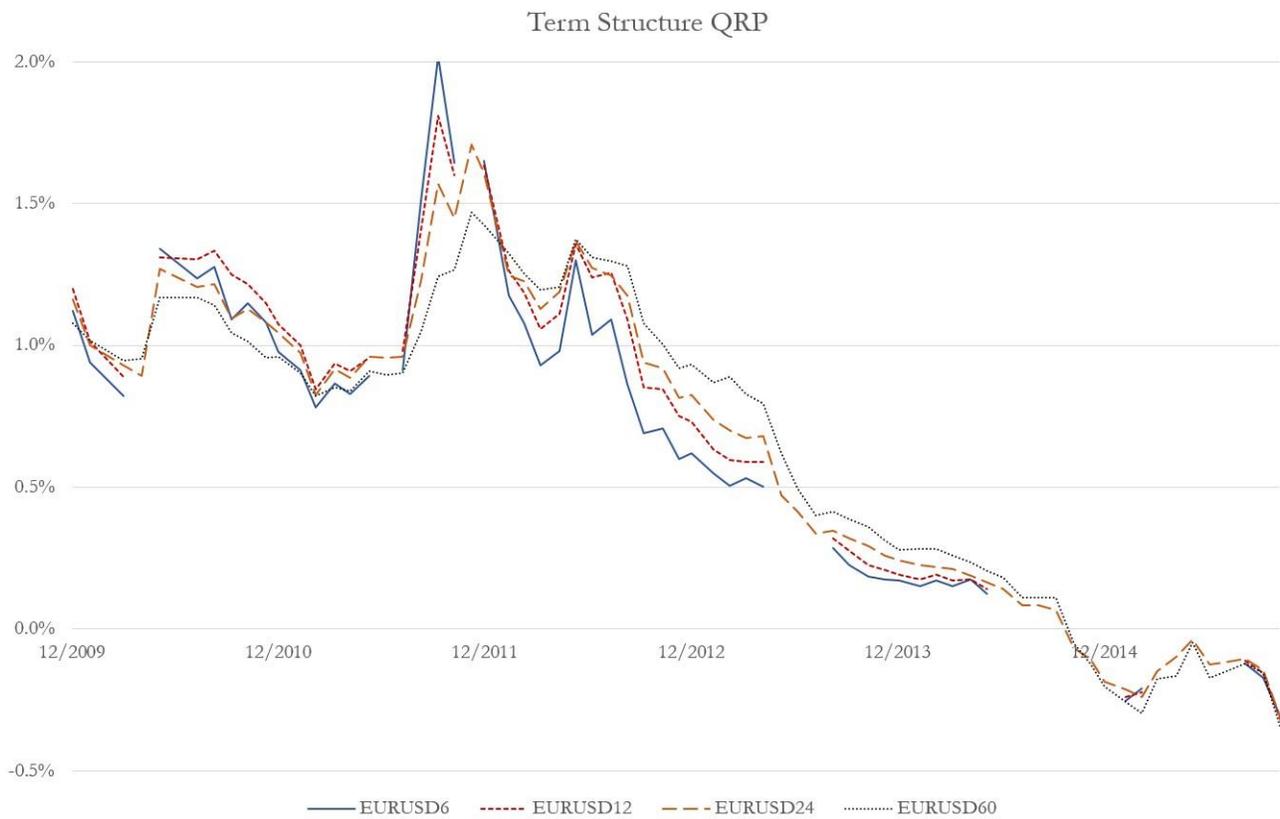
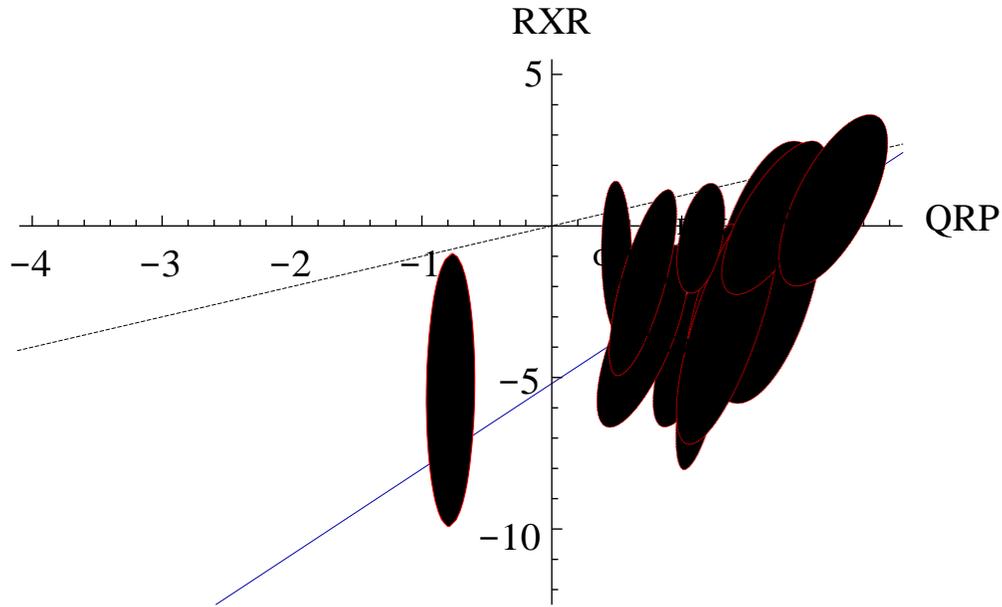
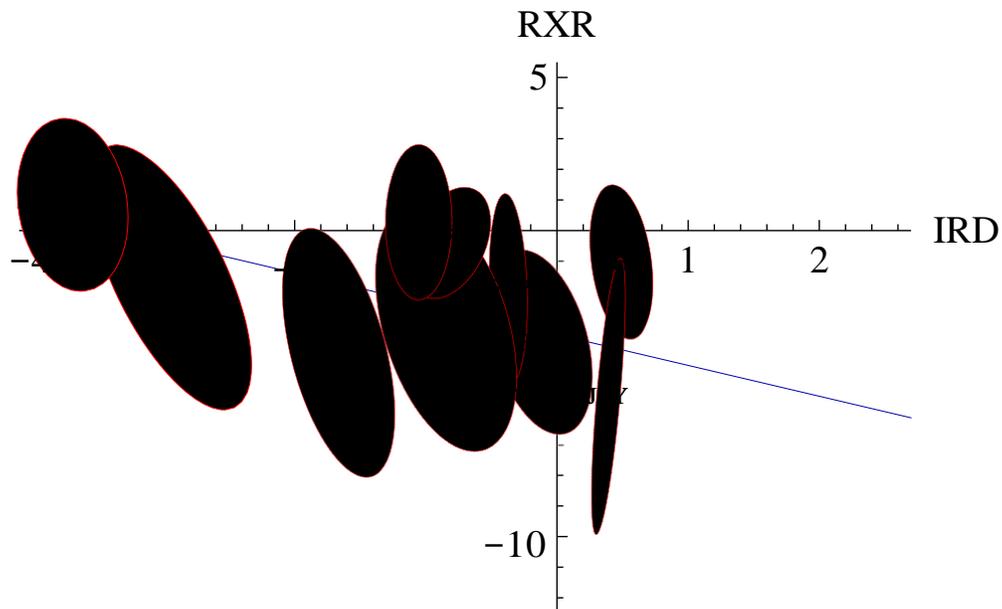


Figure 4: Term structure of the euro-dollar risk premium, as measured by QRP, in the time series for horizons of 6, 12, 24, and 60 months.

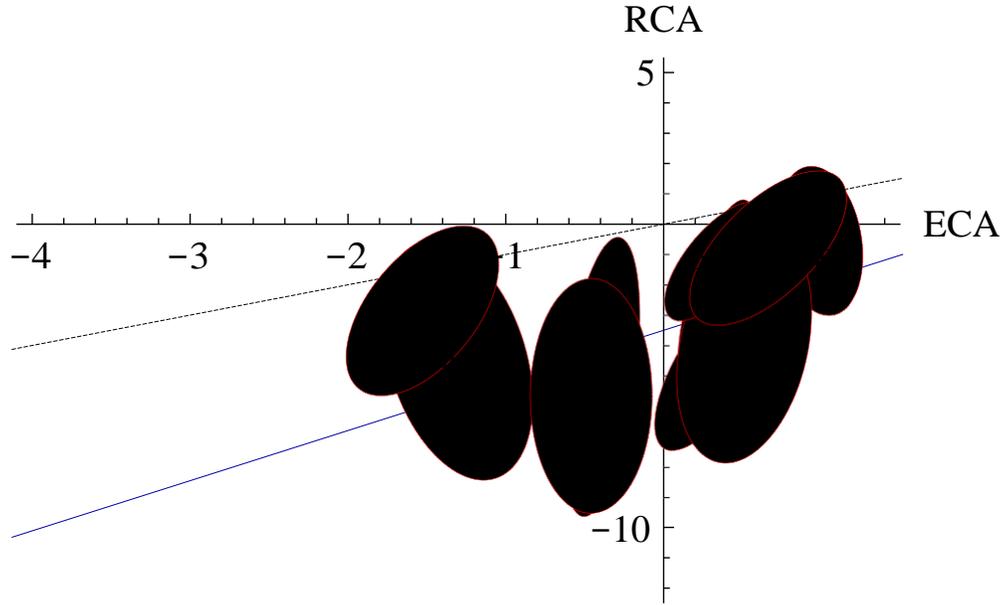


(a) Realized currency excess return against QRP, computed from (14)

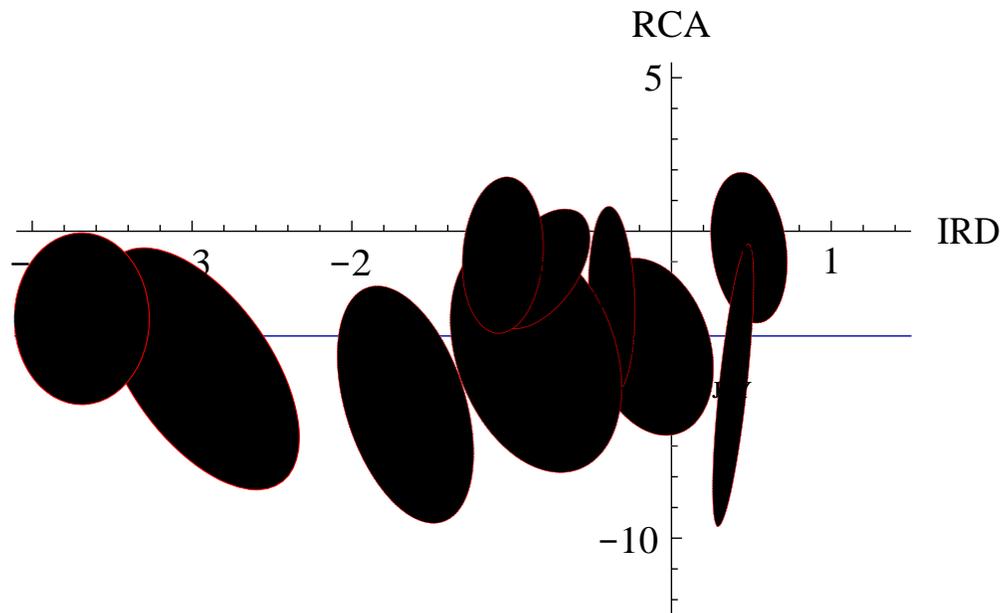


(b) Realized currency excess return against IRD

Figure 5: 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. The dotted 45° line in panel (a) indicates the prediction of the quanto theory.



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



(b) Realized currency appreciation against IRD

Figure 6: 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. The dotted 45° line in panel (a) indicates the prediction of the quanto theory.

Table 4: Currency excess return forecasting regressions

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

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$,t}}{R^{\$,t}} = a + \beta \frac{Q_{+Y}}{R^i} \frac{F^t}{F^t} - 1 + \varepsilon_{i,t+1} \quad (15)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$,t}}{R^{\$,t}} = a + \beta \frac{Q_{+Y}}{R^i} \frac{F^t}{F^t} + \varepsilon_{i,t+1} \quad (16)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$,t}}{R^{\$,t}} = a + \gamma \frac{R^{\$,t}}{R^{\$,t}} - 1 + \varepsilon_{i,t+1} \quad (17)$$

The top panel reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second panel reports R^2 (in %). The bottom three panels report p -values of Wald tests of various hypotheses on the regression coefficients. H_0^1 is the hypothesis suggested by Result 2. Hypothesis H_{β} drops the constraint that $a = 0$, 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 the right-hand side variables are not useful for forecasting.

	pooled			currency fixed effects		
Regression	(15)	(16)	(17)	(15)	(16)	(17)
a (p.a.)	-0.052 (0.019)	-0.052 (0.019)	-0.034 (0.013)			
β	3.485 (1.736)	2.821 (1.141)		5.635 (1.740)	5.954 (1.314)	
γ	0.661 (1.057)		-1.008 (0.648)	-0.633 (1.348)		-2.729 (1.053)
R^2	22.00	20.76	5.57	30.31	29.94	8.52
$H_0^1: a = \gamma = 0, \beta = 1$	0.032					
$H_{\beta}: a = 0, \beta = 1$		0.018				
$H_{\beta}: \beta = 1, \gamma = 0$	0.288			0.000		
$H_{\beta}: \beta = 1$		0.111			0.000	
$H_{\beta}: \beta = 0, \gamma = 0$	0.057			0.000		
$H_{\beta}: \beta = 0$		0.013			0.000	
$H_{\beta}: \gamma = 0$			0.120			0.010

Table 5: Currency forecasting regressions

This Table presents results from three currency forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = a + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^j P_t} + \gamma \frac{R_{f,t}^s}{R_{f,t}^i} - 1 + \varepsilon_{i,t+1} \quad (18)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = a + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^j P_t} + \varepsilon_{i,t+1} \quad (19)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = a + \gamma \frac{R_{f,t}^s}{R_{f,t}^i} - 1 + \varepsilon_{i,t+1} \quad (20)$$

The top panel reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second panel reports R^2 (in %). The bottom three panels report p -values of Wald tests of various hypotheses on the regression coefficients. H_0^1 is the hypothesis suggested by Result 2 (with the constraints $\gamma = 1$ and $\beta = 1$). Hypothesis H_0^2 drops the constraint that $a = 0$, 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 the right-hand side variables are not useful for forecasting.

	pooled			currency fixed effects		
Regression	(18)	(19)	(20)	(18)	(19)	(20)
a (p.a.)	-0.052 (0.019)	-0.051 (0.019)	-0.034 (0.013)			
β	3.485 (1.734)	1.816 (1.183)		5.635 (1.742)	5.450 (1.307)	
γ	1.661 (1.056)		-0.008 (0.656)	0.367 (1.349)		-1.729 (1.048)
R^2	17.40	9.11	0.00	26.56	26.43	3.60
$H_0^1: a = 0, \beta = \gamma = 1$	0.030					
$H_0^2: \beta = 1, \gamma = 1$	0.284			0.000		
$H_0^3: \beta = 1$		0.491			0.001	
$H_0^4: \beta = 0, \gamma = 0$	0.132			0.000		
$H_0^5: \beta = 0$		0.125			0.000	
$H_0^6: \gamma = 0$			0.990			0.099

Table 6: 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):

$$RPCL_{i,t} = R_{f,t}^S \sum_{s=t-h}^t \frac{e_{i,s}}{e_{i,s-1}} - \frac{1}{R_s} \sum_{s=t-h}^t \left(-\frac{1}{R_s} \sum_{s=t-h}^t \frac{e_{i,s}}{e_{i,s-1}} \right),$$

where the summation is over daily returns on trading days s preceding t over a time-frame corresponding to our forecasting horizon, h . We also define the realized covariance measure $RPC_{i,t}$, which is analogous to the above definition except that the summation is over trading days *following* t over the appropriate time-frame.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^S}{R_{f,t}^R} = a + \gamma RPCL_{i,t} + \varepsilon_{i,t+1} \quad (23)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^S}{R_{f,t}^R} = a + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^R} + \gamma RPCL_{i,t} + \varepsilon_{i,t+1} \quad (24)$$

$$RPC_{i,t} = a + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^R} + \varepsilon_{i,t+1} \quad (25)$$

Our data runs from December 2009 to May 2015, forming an unbalanced panel of 24-month quantos on the S&P 500 index in the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krona (DKK), Euro (EUR), British pound (GBP), Japanese yen (JPY), Korean won (KRW), Norwegian krone (NOK), Polish zloty (PLN), and Swedish krona (SEK). The top panel reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second panel reports adjusted R^2 . The bottom three panels report p -values of Wald tests of various joint hypotheses on the coefficients.

	pooled			currency fixed effects		
Regression	(23)	(24)	(25)	(23)	(24)	(25)
a (p.a.)	-0.038 (0.016)	-0.052 (0.018)	0.000 (0.002)			
β		2.855 (1.481)	0.442 (0.162)		5.222 (1.736)	0.309 (0.174)
γ	1.513 (1.028)	-0.037 (1.149)		2.543 (1.470)	0.882 (1.497)	
R^2	10.01	20.76	35.48	14.26	31.20	9.43
$H_0: a = \gamma = 0, \beta = 1$		0.0396				
$H_0: a = 0, \gamma = 1$	0.041					
$H_0: \beta = 1, \gamma = 0$		0.286			0.001	
$H_0: \gamma = 1$	0.618			0.294		
$H_0: \beta = \gamma = 0$		0.050			0.000	
$H_0: \beta = 0$			0.007			0.077
$H_0: \gamma = 0$	0.141			0.084		

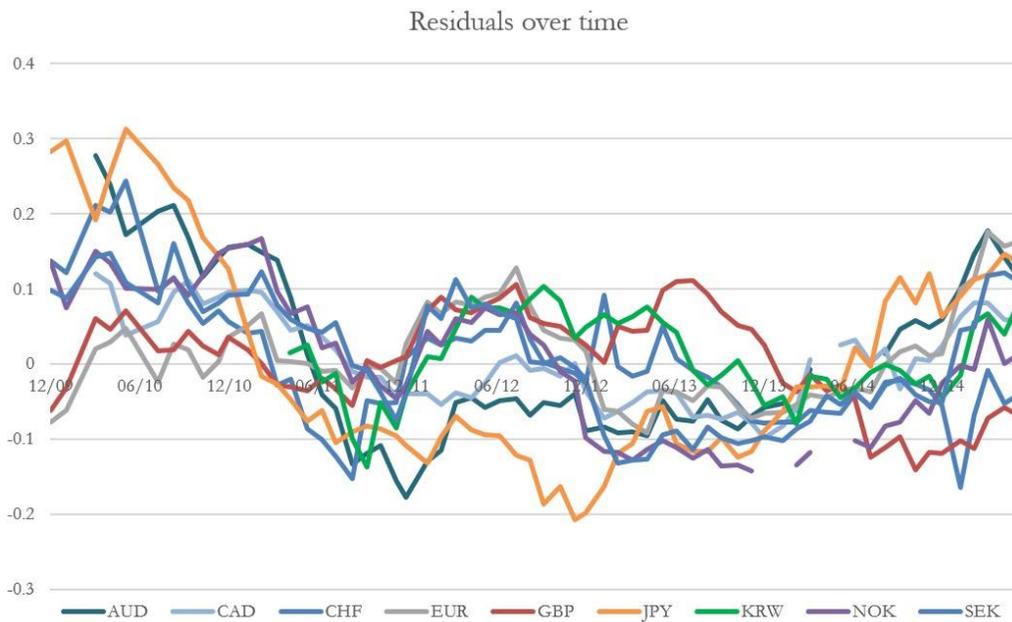
Implied vs. realized covariance



Figure 7: Covariance risk premia, as measured by the difference between implied (risk-neutral) covariance and realized covariance (QRP-RPC), in the time series. For clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK).



(a) Theory residuals



(b) Regression residuals

Figure 8: The time series of residuals by currency.

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 (16) (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).

Theory	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
AUD	0.514	0.198	0.062	-0.457	-0.241	-0.579	0.083	-0.291
CAD	0.303	-0.009	-0.103	-0.276	0.010	-0.003	-0.177	0.889
CHF	0.213	-0.082	0.672	0.286	-0.553	0.197	0.221	0.144
EUR	0.239	-0.294	-0.217	0.676	-0.078	-0.478	-0.342	0.031
GBP	0.126	-0.428	0.589	-0.153	0.600	-0.093	-0.236	-0.081
JPY	0.317	0.765	0.163	0.344	0.382	0.115	-0.107	0.000
NOK	0.480	-0.171	-0.219	-0.118	-0.161	0.599	-0.442	-0.312
SEK	0.441	-0.267	-0.255	0.145	0.314	0.131	0.730	0.005
Explained	61.77%	26.36%	6.93%	2.79%	0.96%	0.44%	0.40%	0.34%

Regression	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
AUD	0.528	-0.151	-0.008	-0.255	0.707	-0.048	0.231	-0.280
CAD	0.244	0.029	-0.180	-0.293	0.175	0.051	-0.527	0.714
CHF	0.216	0.310	0.651	0.204	0.025	-0.595	-0.008	0.193
EUR	0.120	0.263	-0.414	0.737	0.269	-0.055	-0.321	-0.153
GBP	-0.082	0.445	0.464	0.026	0.234	0.714	-0.106	-0.050
JPY	0.600	-0.479	0.237	0.304	-0.385	0.307	-0.136	-0.029
NOK	0.354	0.485	-0.158	-0.401	-0.398	-0.088	-0.288	-0.453
SEK	0.330	0.385	-0.275	0.106	-0.189	0.161	0.673	0.377
Explained	65.63%	15.18%	10.58%	3.41%	2.84%	1.40%	0.61%	0.35%

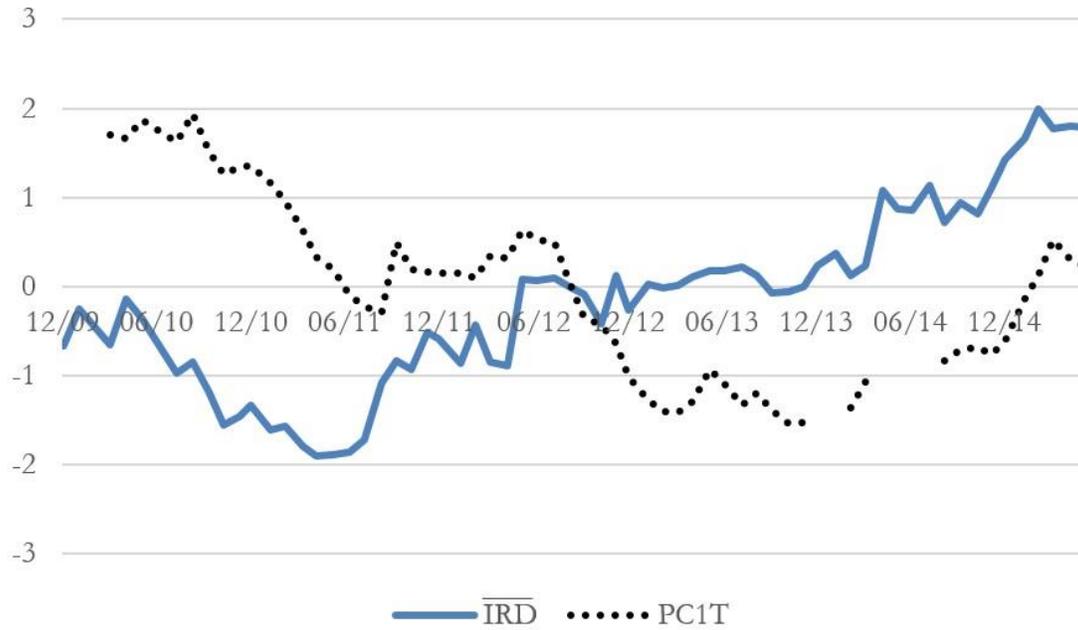


Figure 9: Time series of cross-sectional average of interest-differential (\overline{IRD}_t) and of the first principal component of theory residuals (PC1T). All series are demeaned and scaled to have unit variance.

Table 8: The connection to the average forward discount, $\overline{\text{IRD}}_t$

This Table presents results from a currency excess return forecasting regression that extends the baseline results in Table 4:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^s}{R_{f,t}^i} = a + \beta \frac{Q_{i,t}}{R_{f,t}^i} + \gamma \frac{F_{i,t}}{R_{f,t}^i} + \delta \left(\frac{R_{f,t}^s}{R_{f,t}^i} - 1 \right) + \overline{\text{IRD}}_t + \varepsilon_{i,t+1} \quad (26)$$

and, for reference, results from the regression

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^s}{R_{f,t}^i} = a + \gamma \frac{R_{f,t}^s}{R_{f,t}^i} - 1 + \delta \left(\frac{R_{f,t}^s}{R_{f,t}^i} - 1 \right) + \varepsilon_{i,t+1}. \quad (27)$$

In order to account for a time-varying dollar factor, we augment our baseline regressions (15) and (17) by including the time series of the cross-sectional averages of IRD, denoted by $\overline{\text{IRD}}_t$. To limit the impact of missing observations on the cross-sectional averages, the above regression is run on the panel of 24-month quantos on the S&P 500 index in the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), and Swedish krona (SEK) (i.e., excluding DKK, KRW, and PLN). The top panel reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second panel reports R^2 (in %). The last five panels report p -values of Wald tests of different hypotheses.

	pooled				currency fixed effects			
Regression	(26)		(27)		(26)		(27)	
a (p.a.)	-0.070	(0.029)	-0.051	(0.030)				
β	3.531	(1.617)			5.624	(1.683)		
γ	1.073	(1.042)	-0.637	(0.718)	-0.571	(1.755)	-1.912	(1.634)
δ	-5.670	(4.544)	-5.391	(6.713)	-0.220	(3.501)	-2.697	(6.613)
R^2	26.37		9.52		30.31		9.08	
$H_0: \beta = 1, \gamma = \delta = 0$	0.378				0.004			
$H_0: \beta = \gamma = \delta = 0$	0.126				0.000			
$H_0: \gamma = \delta = 0$					0.361			
$H_0: \gamma = 0$					0.375			
$H_0: \delta = 0$					0.422			
					0.323			
					0.242			
					0.683			

Table 9: 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 $q_{i,t} = \log e_{i,t} \cdot \frac{P_{i,t}}{P_{\$,t}}$, where $P_{i,t}$ and $P_{\$,t}$ are consumer price indices for country i and the US, respectively, obtained from the OECD.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = a_i + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} + \gamma \frac{R_{f,t}^{\$} - 1}{R_{f,t}^i} + \zeta q_{i,t} + \varepsilon_{i,t+1} \quad (30)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = a_i + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} + \zeta q_{i,t} + \varepsilon_{i,t+1} \quad (31)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = a_i + \gamma \frac{R_{f,t}^{\$} - 1}{R_{f,t}^i} + \zeta q_{i,t} + \varepsilon_{i,t+1} \quad (32)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = a_i + \beta \frac{Q_{i,t} - F_t}{R_{f,t}^i P_t} + \gamma \frac{R_{f,t}^{\$} - 1}{R_{f,t}^i} + \delta \left(\frac{\text{IRD}_t}{\text{IRD}_t} \right) + \zeta q_{i,t} + \varepsilon_{i,t+1} \quad (33)$$

The top panel reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second panel reports R^2 (in %). The bottom panel reports p -values for the null hypothesis that the right-hand side variables are not useful for forecasting.

currency fixed effects								
Regression	(30)		(31)		(32)		(33)	
β	4.697	(1.647)	6.127	(1.353)			3.919	(1.529)
γ	-3.087	(1.648)			-5.374	(1.230)	-1.509	(1.632)
δ							-8.710	(4.006)
ζ	-0.523	(0.207)	-0.304	(0.15)	-0.658	(0.164)	-0.716	(0.232)
R^2	41.48		35.58		27.11		45.73	
$H_0: \beta = 0$	0.004		0.000				0.010	
$H_0: \gamma = 0$	0.061				0.000		0.355	
$H_0: \delta = 0$							0.030	
$H_0: \zeta = 0$	0.011		0.039		0.000		0.002	

Table 10: Examining the residuals

This Table presents results from a currency excess return forecasting regression that extends the baseline results in Table 4:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^s}{R_{f,t}^r} = \alpha_i + \frac{Q_{i,t} - F_t}{R_{f,t}^r P_t} + \gamma \text{Proxy}_{i,t} + \varepsilon_{i,t+1}$$

Relative to our baseline regressions, we constrain the coefficient (β) on QRP to equal 1 as implied by the theory. We then include several variables to estimate the residual term: QRP, IRD, $\overline{\text{IRD}}$, RPCL as used in (23), and a variable we call RPCL(3.5) defined below.

$$\text{RPCL}(3.5)_{i,t} = R_{f,t}^s \left(\frac{e_{i,s+1}}{e_{i,s}} - \frac{1}{R_{s+1}^{2.5}} - \frac{1}{T-t} - \frac{1}{R_{s+1}^{2.5}} \frac{e_{i,s+1}}{e_{i,s}} \right),$$

The first line of the table reports coefficient estimates for each regression, with standard errors (computed using a block bootstrap) in brackets. The second line reports R^2 (in %). The third line reports p -values of tests of the hypothesis that $\gamma = 0$.

currency fixed effects							
Proxy	none	QRP	IRD	IRD	RPCL	RPCL(3.5)	q
γ		4.954 (1.317)	-2.357 (1.021)	-5.432 (5.936)	2.225 (1.403)	0.904 (0.565)	-0.264 (0.160)
R^2	9.21	29.94	15.57	14.60	20.12	20.17	13.48
$H_0: \gamma = 0$		0.000	0.021	0.360	0.113	0.110	0.100

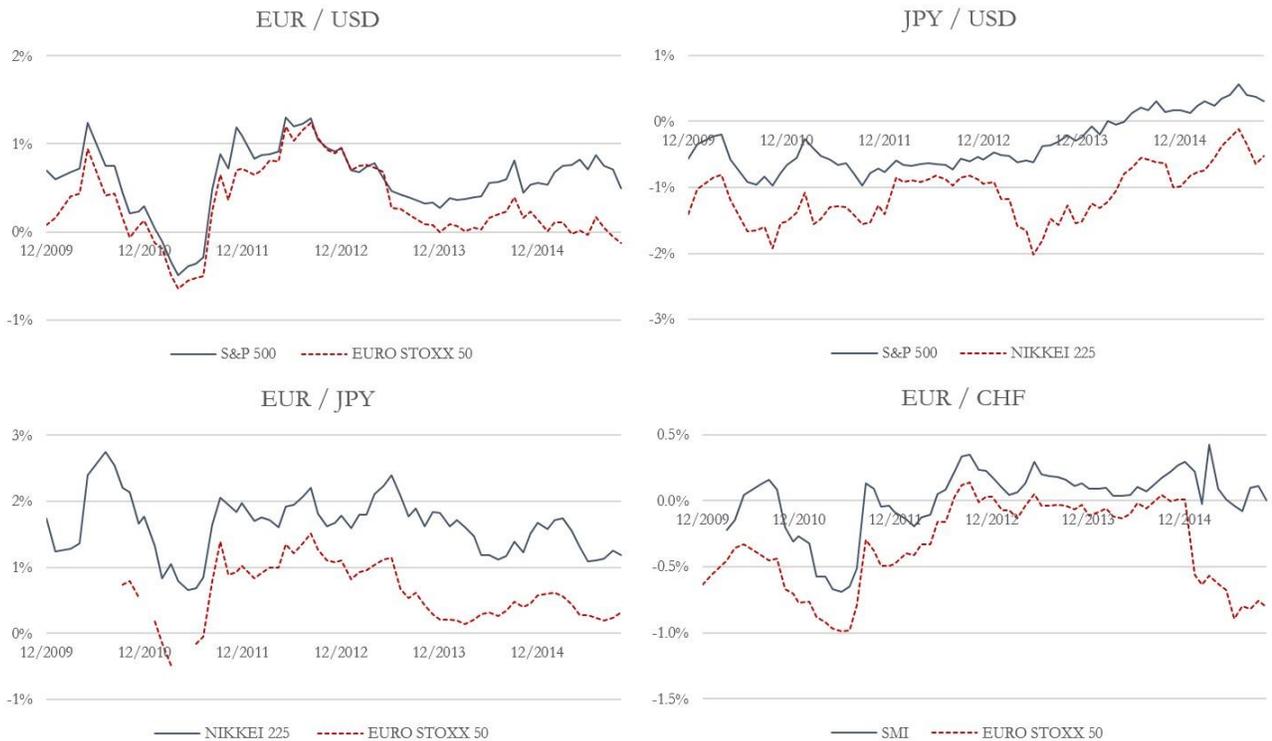


Figure 10: Expected currency appreciation over a 24-month horizon (annualized), as measured by ECA from equation (14), for the EUR-USD, JPY-USD, EUR-JPY, and EUR-CHF currency pairs. Each panel plots ECA for the respective currency pair from the two national perspectives, using quanto contracts on the 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.

Table 11: Out-of-sample forecast performance

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

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

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. We use the quanto theory and three competitor benchmarks to forecast currency appreciation as follows:

$$\begin{aligned} \text{Theory: } E_t^Q \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \frac{Q_{i,t} - F_t}{R_{f,t}^*} - 1 \\ \text{UIP: } E_t^U \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \text{IRD}_{i,t} = \frac{R_{f,t}}{R_{f,t}^i} - 1 \\ \text{Constant: } E_t^C \frac{e_{i,t+1}}{e_{i,t}} - 1 &= 0 \\ \text{PPP: } E_t^P \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \frac{\pi_t^*}{\pi_t} - 1 \end{aligned}$$

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

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

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}	11.60	10.48	26.91
$R^2_{OS,AUD}$	9.91	3.36	11.95
$R^2_{OS,CAD}$	6.56	7.26	21.43
$R^2_{OS,CHF}$	1.42	18.29	11.16
$R^2_{OS,DKK}$	10.22	7.71	23.36
$R^2_{OS,EUR}$	8.64	5.93	25.44
$R^2_{OS,GBP}$	2.67	9.89	35.04
$R^2_{OS,JPY}$	20.55	10.38	34.88
$R^2_{OS,KRW}$	22.66	17.88	35.44
$R^2_{OS,NOK}$	3.65	13.29	19.10
$R^2_{OS,PLN}$	13.33	9.00	20.10
$R^2_{OS,SEK}$	8.02	6.44	28.99
DM p -value	0.037	0.000	0.000

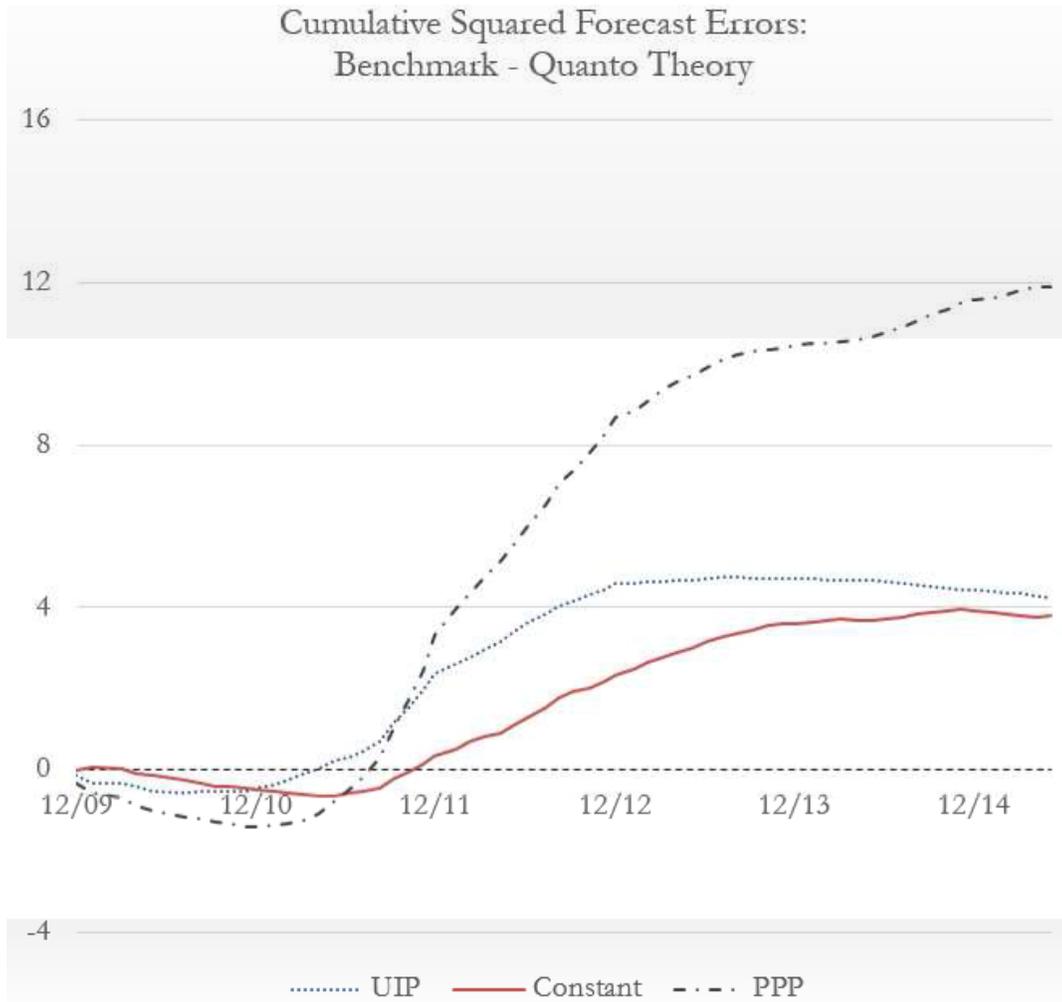


Figure 11: time series of cumulative outperformance of the quanto theory forecast over three competitors.

Table 12: Separate return forecasting regressions using quanto predictor

In addition to the pooled regressions in Tables 4 and 5, we also run regressions (16) and (18) 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. Adjusted R^2 are reported in %.

Currency	Regression (16)					Regression (18)						
	α		β		R^2	α		β		γ		R^2
AUD	-0.090	(0.069)	4.546	(3.776)	22.74	-0.153	(0.082)	0.770	(2.622)	-3.249	(2.227)	32.90
CAD	-0.107	(0.035)	6.279	(2.876)	47.04	-0.092	(0.046)	6.196	(2.882)	2.900	(2.511)	50.47
CHF	-0.002	(0.039)	-1.707	(7.074)	0.64	0.012	(0.036)	-1.506	(6.883)	-1.962	(3.369)	4.26
DKK	-0.052	(0.022)	4.125	(1.723)	17.42	-0.041	(0.021)	5.252	(1.260)	3.857	(1.671)	48.62
EUR	-0.079	(0.026)	6.099	(2.860)	51.10	-0.091	(0.023)	8.267	(2.441)	3.658	(1.847)	52.99
GBP	-0.081	(0.033)	8.731	(3.893)	54.18	-0.074	(0.035)	9.425	(3.581)	4.142	(3.112)	55.45
JPY	-0.042	(0.098)	1.603	(11.592)	0.44	-0.166	(0.084)	-1.736	(8.171)	26.588	(8.394)	53.72
KRW	-0.075	(0.034)	4.526	(1.720)	51.63	-0.083	(0.050)	4.572	(1.788)	0.265	(2.174)	52.63
NOK	-0.186	(0.059)	12.151	(4.513)	53.75	-0.201	(0.069)	10.592	(3.133)	-0.999	(2.298)	52.89
PLN	-0.090	(0.026)	4.535	(0.818)	44.68	-0.073	(0.032)	4.735	(0.863)	1.585	(0.938)	44.12
SEK	-0.135	(0.045)	7.658	(3.011)	54.83	-0.137	(0.043)	8.188	(2.381)	1.629	(1.494)	50.86
EUR (6m)	-0.040	(0.056)	3.702	(6.263)	3.17	-0.055	(0.053)	10.008	(7.198)	11.447	(8.450)	14.42
EUR (12m)	-0.071	(0.052)	6.361	(5.527)	17.98	-0.092	(0.043)	12.916	(4.771)	11.992	(4.880)	45.19

Table 13: Separate return forecasting regressions using IRD predictor

In addition to the pooled regressions in Tables 4 and 5, we also run regressions (17) and (20) 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	Regression (17)					Regression (20)				
	α		γ		R^2	α		γ		R^2
AUD	-0.151	(0.080)	-4.615	(2.530)	43.91	-0.076	(0.040)	-3.615	(2.530)	32.44
CAD	-0.020	(0.030)	2.519	(3.261)	3.30	-0.010	(0.015)	3.519	(3.261)	6.24
CHF	0.004	(0.028)	-2.986	(3.330)	8.06	0.002	(0.014)	-1.986	(3.330)	3.73
DKK	0.014	(0.023)	2.147	(2.036)	13.77	0.007	(0.012)	3.147	(2.036)	25.54
EUR	-0.040	(0.241)	-2.792	(3.229)	13.25	-0.020	(0.121)	-1.792	(3.229)	5.92
GBP	-0.032	(0.035)	3.471	(6.643)	2.70	-0.016	(0.017)	4.471	(6.643)	1.38
JPY	-0.151	(0.047)	25.109	(8.375)	51.29	-0.075	(0.023)	26.109	(8.375)	53.24
KRW	0.003	(0.035)	-0.007	(2.893)	0.00	0.001	(0.018)	0.993	(2.893)	0.97
NOK	-0.118	(0.066)	-4.716	(3.717)	24.62	-0.059	(0.033)	-3.716	(3.717)	16.86
PLN	-0.028	(0.043)	-0.997	(1.182)	2.21	-0.014	(0.021)	0.003	(1.182)	0.00
SEK	-0.057	(0.037)	-3.063	(2.839)	16.38	-0.029	(0.018)	-2.063	(2.839)	8.16
EUR (6m)	-0.007	(0.041)	2.626	(7.375)	1.23	-0.015	(0.083)	3.626	(7.375)	2.32
EUR (12m)	-0.019	(0.040)	1.869	(6.349)	1.31	-0.019	(0.040)	2.869	(6.349)	3.03

A.2 Bootstrap

We use a non-parametric bootstrap procedure to compute the covariance matrix of our coefficient estimates. A detailed background on the stationary bootstrap methodology is provided in Politis and White (2004) and Patton et al. (2009). In our 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 resulting panel is then resampled with replacement in the cross-sectional dimension by drawing blocks of uniformly distributed width (between 2 and 11). 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. To test the small-sample performance of our bootstrap procedure, we simulate 1,000 sets of excess return data under the null hypothesis of no predictability such that the simulated data matches the autocorrelation and covariance matrix of the excess returns observed in the real data. We then run regression (16) to predict the returns simulated under the null ($\alpha_0 = \beta_0 = 0$) using the real quanto data and obtain standard errors using our bootstrap procedure for each set simulated of returns. We find that even with our small sample size, the frequencies at which we reject the null of no-predictability are close to the p-values of the asymptotic χ^2 -test at the conventional test levels of 10%, 5%, and 1%. Since this test describes the forecasting performance of our quanto data in a simulated world without predictability, the appropriate rejection rates obtained from the block-bootstrap procedure in this simulation suggest that potential small-sample problems have limited impact on our inference.

Disclaimer

Markit[®] is a trade name and the property of Markit Group Limited or its affiliate (“Markit”) and is used by the London School of Economics and Political Science under license. Data provided by Markit[®]. Nothing in this publication is sponsored, endorsed, sold or promoted by Markit or its affiliates. Neither Markit nor its affiliates make any representations or warranties, express or implied, to you or any other person regarding the advisability of investing in the financial products described in this report or as to the results obtained from the use of the Markit Data. Neither Markit nor any of its affiliates have any obligation or liability in connection with the operation, marketing, trading or sale of any financial product described in this report or use of the Markit Data. Markit and its affiliates shall not be liable (whether in negligence or otherwise) to any person for any error in the Markit Data and shall not be under any obligation to advise any person of any error therein.