

The Quanto Theory of Exchange Rates[†]

By LUKAS KREMENS AND IAN MARTIN*

We present a new identity that relates expected exchange rate appreciation to a risk-neutral covariance term, and use it to motivate a currency forecasting variable based on the prices of quanto index contracts. We show via panel regressions that the quanto forecast variable is an economically and statistically significant predictor of currency appreciation and of excess returns on currency trades. Out of sample, the quanto variable outperforms predictions based on uncovered interest parity, on purchasing power parity, and on a random walk as a forecaster of differential (dollar-neutral) currency appreciation. (JEL C53, E43, F31, F37, G12, G15)

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 equalize 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 foreign exchange literature. The UIP forecast has three appealing properties. First, it is determined by asset prices alone rather than by, say, infrequently updated and imperfectly measured macroeconomic data. Second, it has no free parameters: with no coefficients to be estimated in-sample or “calibrated,” it is perfectly suited to out-of-sample forecasting. Third, it has a straightforward interpretation as the expected exchange rate movement perceived by a risk-neutral

*Kremens: Department of Finance, London School of Economics, London, UK (email: l.kremens@lse.ac.uk); Martin: Department of Finance, London School of Economics, London, UK (email: i.w.martin@lse.ac.uk). Gita Gopinath was coeditor for this article. We thank the Systemic Risk Centre and the Paul Woolley Centre at the LSE for their support, and for providing access to data sourced from Markit under license. We are grateful to Christian Wagner, Tarek Hassan, John Campbell, Mike Chernov, Gino Cenedese, Anthony Neuberger, Dagfinn Rime, Urban Jermann, Bryn Thompson-Clarke, Adrien Verdelhan, Bernard Dumas, Pierpaolo Benigno, Alan Taylor, Daniel Ferreira, Ulf Axelson, Scott Robertson, Adrian Buss, Stig Vinther Møller, and to participants in seminars at the LSE, Imperial College, Cass Business School, LUISS, BI Business School, Boston University, and Queen Mary University of London, for their comments; and to Lerby Ergun for research assistance. Ian Martin is also grateful for support from the ERC under Starting Grant 639744.

[†]Go to <https://doi.org/10.1257/aer.20180019> to visit the article page for additional materials and author disclosure statements.

¹Some studies (e.g., Sarno, Schneider, and Wagner 2012) find that currencies with high interest rates appreciate on average, exacerbating the failure of UIP; this has become known as the forward premium puzzle. Others, such as Hassan and Mano (2017), find that exchange rates move in the direction predicted by UIP, though not by enough to offset interest rate differentials.

investor. Put differently, UIP holds if and only if the *risk-neutral* expected appreciation of a currency is equal to its *real-world* expected appreciation, the latter being the quantity relevant for forecasting exchange rate movements.

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

In this paper we propose a new predictor variable that also possesses the three appealing properties mentioned above, but which does not require that one takes the perspective of a risk-neutral investor. This alternative benchmark can be interpreted as the expected exchange rate movement that must be perceived by a risk-averse investor with log utility whose wealth is invested in the stock market. (To streamline the discussion, this description is an oversimplification and strengthening of the condition we actually need to hold for our approach to work, which is based on a general identity presented in Result 1.) This approach has been shown by Martin (2017) and Martin and Wagner (forthcoming) to be successful in forecasting returns on the stock market and on individual stocks, respectively.

It turns out that such an investor's expectations about currency returns can be inferred directly from the prices of so-called *quanto contracts*. For our purposes, the important feature of such contracts is that their prices are sensitive to the correlation between a given currency and some other asset price. Consider, for example, a quanto contract whose payoff equals the level of the S&P 500 index at time T , denominated in euros (that is, the exchange rate is fixed—in this example, at 1 euro per dollar—at initiation of the trade). The value of this contract is sensitive to the correlation between the S&P 500 index and the dollar/euro exchange rate. If the euro appreciates against the dollar at times when the index is high, and depreciates when the index is low, then this quanto contract is more valuable than a conventional, dollar-denominated, claim on the index.³

We show that the relationship between currency- i quanto forward prices and conventional forward prices on the S&P 500 index reveals the risk-neutral covariance between currency i and the index. Quantos therefore signal which currencies are risky (in that they tend to depreciate in bad times, i.e., when the S&P 500 declines) and which are hedges. It is possible, of course, that a currency is risky at one point in time and a hedge at another. Intuitively, one expects that a currency that is (currently) risky should, as compensation, have higher expected appreciation than predicted by UIP, and that hedge currencies should have lower expected appreciation.

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

³A different type of quanto contract—specifically, quanto CDS contracts—is used by Mano (2013) and Augustin, Chernov, and Song (2018) to study the relationship between currency depreciation and sovereign default.

Our framework formalizes this intuition. It also allows us to distinguish between variation in risk premia across currencies and variation over time.

It is worth emphasizing various assumptions that we do *not* make. We do not require that markets are complete (though our approach remains valid if they are). We do not assume the existence of a representative agent, nor do we assume that all economic actors are rational. The forecast in which we are interested reflects the beliefs of a rational investor, but this investor may coexist with investors with other, potentially irrational, beliefs. And we do not assume lognormality, nor do we make any other distributional assumptions: our approach allows for skewness and jumps in exchange rates. This is an important strength of our framework, given that currencies often experience crashes or jumps (as emphasized by Brunnermeier, Nagel, and Pedersen 2008; Jurek 2014; Della Corte et al. 2016; Chernov, Graveline, and Zviadadze 2018; and Farhi and Gabaix 2016; among others), and are prone to structural breaks more generally. The approach could even be used, in principle, to compute expected returns for currencies that are currently pegged but that have some probability of jumping off the peg. To the extent that skewness and jumps are empirically relevant, this fact will be embedded in the asset prices we use as forecasting variables.

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

We derive these and other theoretical results in Section I, and test them in Section II by running panel currency-forecasting regressions. The estimated coefficient on the quanto predictor variable is economically large and statistically significant. In our headline regression (20), we find t -statistics of 3.2 and 2.3, respectively, with and without currency fixed effects. (Here, as throughout the paper, we compute standard errors—and more generally the entire covariance matrix of coefficient estimates—using a nonparametric block bootstrap to account for heteroskedasticity, cross-sectional correlation across currencies, and autocorrelation in errors induced by overlapping observations.) The quanto predictor outperforms forecasting variables such as the interest rate differential, average forward discount, and the real exchange rate as a univariate forecaster of currency excess returns. On the other hand, we find that some of these variables—notably the real exchange rate and average forward discount—interact well with our quanto predictor variable, in the sense that they substantially raise R^2 above what the quanto variable achieves on its own. We interpret this fact, through the lens of the identity (6) of Result 1, as showing that these variables help to measure deviations from the log investor benchmark. We also show that the quanto predictor variable (that is, forward-looking risk-neutral covariance) predicts future realized covariance and substantially outperforms lagged realized covariance as a forecaster of exchange rates.

An important challenge is that our dataset spans a relatively short time period. If we assess the significance of joint hypothesis tests by using p -values based on the

asymptotic distributions of test statistics (with bootstrapped covariance matrices, as always), we find, in our pooled regressions, that the estimated coefficients on the quanto predictor variable and interest rate differential are consistent with the predictions of the log investor benchmark, but we can reject the hypothesis that, in addition, the intercept is zero. This rejection can be attributed to US dollar appreciation, during our sample, that was not anticipated by our model. But using asymptotic distributions of test statistics to assess p -values risks giving a false impression of precision, in view of our short sample period. In Section IIF, we bootstrap the small-sample distributions of the relevant test statistics to account for this issue. When we use the associated, more conservative, small-sample p -values, we do not reject even the most optimistic hypothesis in any of the specifications, though the individual significance of the quanto predictor becomes more marginal, with p -values ranging from 5.1 percent to 9.7 percent.

In Section III we show that the quanto variable performs well out of sample. We focus on forecasting differential returns on currencies in order to isolate the cross-sectional forecasting power of the quanto variable in a dollar-neutral way, in the spirit of Lustig, Roussanov, and Verdelhan (2011), and independent of what Hassan and Mano (2017) refer to as the dollar trade anomaly. (As noted in the preceding paragraph, the dollar strengthened against almost all other currencies over our relatively short sample, so quantos are not successful in forecasting the average performance of the dollar itself. Our findings are therefore complementary to Gourinchas and Rey 2007, who use a measure of external imbalances to forecast the appreciation of the dollar against a trade- or FDI-weighted basket of currencies.)

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. Our out-of-sample forecasts exploit the fact that our theory makes an a priori prediction for the coefficient on the quanto predictor variable. When the coefficient 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 and random walk forecasts—perfectly suited for out-of-sample forecasting. Following Meese and Rogoff (1983) and Welch and Goyal (2008), we compute mean squared errors for the differential currency forecasts made by the quanto theory and by three competitor models: UIP, which predicts currency appreciation through the interest rate differential; PPP, which uses past inflation differentials (as a proxy for expected inflation differentials) to forecast currency appreciation; and the random walk forecast. The quanto theory outperforms all three competitors. We also show that it outperforms on an alternative performance benchmark, the correct classification frontier, that has been proposed by Jordà and Taylor (2012).

I. Theory

We start with the fundamental equation of asset pricing,

$$(1) \quad E_t(M_{t+1} \tilde{R}_{t+1}) = 1,$$

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 the (real-world) expectation operator, conditional on all information available at time t , and M_{t+1} for a stochastic discount factor (SDF) that prices assets denominated in dollars. (We do not assume complete markets, so there may well be other SDFs that also price assets denominated in dollars. But all such SDFs must agree with M_{t+1} on the prices of the payoffs in which we are interested, since they are all tradable.) In equation (1), \tilde{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 $\tilde{R}_{t+1} = R_{f,t}^{\$}$; thus (1) can be rearranged as

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

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 $\tilde{R}_{t+1} = R_{f,t}^i e_{i,t+1}/e_{i,t}$ in (2) and rearranging,⁴ we find

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

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

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

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

$$(5) \quad E_t^* \left(\frac{e_{i,t+1}}{e_{i,t}} \right) = \frac{R_{f,t}^{\$}}{R_{f,t}^i}.$$

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

⁴Unlike most authors in this literature, we prefer to work with true returns, \tilde{R}_{t+1} , rather than with log returns, $\log \tilde{R}_{t+1}$, as the latter are only “an approximate measure of the rate of return to speculation,” in the words of Hansen and Hodrick (1980, p. 831).

exchange rate appreciation, the implicit assumption being made is that the covariance term can indeed be neglected.

Unfortunately, as is well known, the UIP forecast performs poorly in practice: the assumption that the covariance term is negligible in (3) (or, equivalently, that the risk-neutral expectation in (5) is close to the corresponding real-world expectation) is not valid. This is hardly surprising, given the existence of a vast literature in financial economics that emphasizes the importance of risk premia, and hence shows that the SDF M_{t+1} is highly volatile (Hansen and Jagannathan 1991). The risk adjustment term in (3) therefore cannot be neglected: expected currency appreciation depends not only on the interest rate differential, but also on the covariance between currency movements and the SDF. Moreover, it is plausible that this covariance varies both over time and across currencies. We therefore take a different approach that exploits the following observation.

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

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

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

PROOF:

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

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

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:

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

Note also that

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

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

As (3) and (6) are identities, each must hold for all currencies i in any economy that does not exhibit riskless arbitrage opportunities. Nor do they make any

assumptions about the exchange rate regime. If currency i is perfectly pegged then the covariance terms in (6) are zero, and we recover the familiar fact that countries with pegged currencies must either lose control of their monetary policy (that is, set $R_{f,t}^i = R_{f,t}^{\$}$) or restrict capital flows to prevent arbitrageurs from trading on the interest rate differential. More generally, the covariance terms should be small if a currency has a low probability of jumping off its peg.

The identity (6) generalizes (3), however, by allowing R_{t+1} to be an arbitrary return. To make the identity useful for empirical work, we want to choose a return R_{t+1} with two aims in mind. First, the residual term should be small. Second, the middle term should be easy to compute.

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

To achieve both goals simultaneously, we want to pick a return that offsets a substantial fraction of the variation⁵ in M_{t+1} , but we must do so in such a way that the risk-neutral covariance term can be measured empirically. For much of this paper, we will take R_{t+1} to be the return on the S&P 500 index. (We find similar—and internally consistent—results if R_{t+1} is set equal to the return on other stock indexes, such as the Nikkei, Euro Stoxx 50, or SMI: see Sections IB and IIA.) It is highly plausible that this return is negatively correlated with M_{t+1} , consistent with the first goal; in fact we provide conditions below under which the residual is exactly zero. We will now show that the second goal is also achieved with this choice of R_{t+1} because we can calculate the quanto-implied risk premium directly from asset prices without any further assumptions—specifically, from *quanto forward prices* (hence the name).

A. *Quantos*

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

$$(10) \quad F_t = E_t^* P_{t+1}.$$

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

⁵More precisely, all we need is to pick a return that offsets *the component of the variation in M_{t+1} that is correlated with currency movements*. But as this component will in general vary according to the currency in question, it is sensible simply to choose R_{t+1} to offset variation in M_{t+1} itself.

prices indexed by $i = 1, \dots, N$, one for each of the N currencies in our dataset. Other than in Section IB, 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

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

(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,

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

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

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

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

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

RESULT 2 (The log investor): *If we take the perspective of an investor with log utility whose wealth is fully invested in the stock index then $M_{t+1} = 1/R_{t+1}$, so that $\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$ is identically 0. The expected appreciation of currency i is then given by*

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

and the expected excess return⁶ on 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 first is the UIP prediction which, as we have seen in equation (5), equals *risk-neutral* expected currency appreciation. We will often refer to this term as the *interest rate differential* (IRD); and as above we will generally convert to net rather than gross terms by subtracting one. (We choose to refer to a high-interest-rate currency as having a *negative* interest rate differential because such a currency is forecast to depreciate by UIP.) The second is a risk adjustment term: by taking the perspective of the log investor, we have converted the general form of the residual that appears in (3) into a quantity that can be directly observed using the gap between a quanto forward and a conventional forward.⁷ Since it captures the risk premium perceived by the log investor, we refer to this term as the *quanto-implied risk premium* (QRP). Lastly, we refer to the sum of the two terms as *expected currency appreciation* (ECA = IRD + QRP).

Results 1 and 2 link expected currency returns to *risk-neutral* covariances, so deviate from the standard CAPM intuition (that risk premia are related to *true* covariances) in that they put more weight on comovement in bad states of the world. This distinction matters, given the observation of Lettau, Maggiori, and Weber (2014) that the carry trade is more correlated with the market when the market experiences negative returns. Even more important, risk-neutral covariance is directly measurable, as we have shown.⁸ In contrast, forward-looking true covariances are *not* directly observed so must be proxied somehow, typically by historical realized covariance. In Section IIC, we show that risk-neutral covariance drives out historical realized covariance as a predictor variable.

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

⁶Formally, $e_{i,t+1}/e_{i,t} - R_{f,t}^{\$}/R_{f,t}^i$ is an excess return because it is a tradable payoff whose price is 0, by (5).

⁷More generally, we can allow for the case in which the log investor chooses a portfolio $R_{p,t+1} = wR_{t+1} + (1-w)R_{f,t}^{\$}$. (The case in the text corresponds to $w = 1$.) The identity (6) then reduces to

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

We thank Scott Robertson for pointing this out to us. See footnote 12 for more discussion.

⁸While it is well known from the work of Ross (1976) and Breeden and Litzenberger (1978) that risk-neutral expectations of functions of a single asset price can typically be inferred from the price of options on that asset, Martin (2018) shows that it is in general considerably harder to infer risk-neutral expectations of functions of multiple asset prices. It is something of a coincidence that precisely the assets whose prices reveal these risk-neutral covariances are traded.

B. Alternative Benchmarks

Our choice to think from the perspective of an investor who holds the US stock market is a pragmatic one. From a purist point of view, it might seem more natural to adopt the perspective of an investor whose wealth is invested in a globally diversified portfolio;⁹ unfortunately global-wealth quantos are not traded, whereas S&P 500 quantos are. Our approach implicitly relies on an assumption that the US stock market is a tolerable proxy for global wealth. We think this assumption makes sense; it is broadly consistent with the “global financial cycle” view of Miranda-Agrippino and Rey (2015).

Nonetheless, one might wonder whether the results are similar if one uses other countries’ stock markets as proxies for global wealth.¹⁰ For, just as the forward price of the US stock index quantoed into currency i reveals the expected appreciation of currency i versus the dollar, as perceived by a log investor whose portfolio is fully invested in the US stock market, so the forward price of the currency- i stock index quantoed into dollars reveals the expected appreciation of the dollar versus currency i , as perceived by a log investor whose portfolio is fully invested in the currency- i market.

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

$$(15) \quad E_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{IRD_{i,t} + QRP_{i,t}}_{ECA_{i,t}}.$$

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

$$(16) \quad E_t \frac{1/e_{i,t+1}}{1/e_{i,t}} - 1 = \underbrace{IRD_{1/i,t} + QRP_{1/i,t}}_{ECA_{1/i,t}},$$

where we write $IRD_{1/i,t} = R_{f,t}^i/R_{f,t}^\$ - 1$, and where $QRP_{1/i,t}$ is obtained from conventional forwards and *dollar*-denominated quanto forwards on the currency- i stock market. When the left-hand side of the above equation is positive, the dollar is expected to appreciate against currency i .

In Section IIA below, we show that the two perspectives captured by (15) and (16) are broadly consistent with one another (for those currencies for which we observe the appropriate quanto forward prices). If, say, the forward price of the S&P 500 quantoed into euros implies that the euro is expected to appreciate against the dollar by 2 percent (using equation (15)), then the forward price of the Euro Stoxx 50 index quantoed into dollars typically implies that the dollar is expected to *depreciate* against the euro by about 2 percent (using equation (16)).

⁹This perspective is suggested by the analysis of Solnik (1974) and Adler and Dumas (1983), for example.

¹⁰In practice, many investors do choose to hold home-biased portfolios (French and Poterba 1991, Tesar and Werner 1995, and Warnock 2002; and see Lewis 1999 and Coeurdacier and Rey 2013 for surveys).

To be more precise, we need to take into account Siegel's "paradox" (Siegel 1972) that, by Jensen's inequality,

$$(17) \quad E_t \frac{e_{i,t+1}}{e_{i,t}} \geq \left(E_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1}.$$

(The corresponding inequality with E_t replaced by any other expectation operator also holds.) If the US and currency- i investors have the same expectations about currency appreciation then (15)–(17) imply that

$$(18) \quad \log(1 + ECA_{i,t}) \geq -\log(1 + ECA_{1/i,t}).$$

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

The difference between the two sides of (18) reflects a convexity correction whose size is determined by the amount of conditional variation in $e_{i,t+1}$:

$$\begin{aligned} \log(1 + ECA_{i,t}) - (-\log(1 + ECA_{1/i,t})) &= \log E_t \frac{e_{i,t+1}}{e_{i,t}} - \log \left[\left(E_t \frac{1/e_{i,t+1}}{1/e_{i,t}} \right)^{-1} \right] \\ &= \mathbf{c}_t(1) + \mathbf{c}_t(-1) \\ &= 2 \sum_{n \text{ even}} \frac{\kappa_{n,t}}{n!}, \end{aligned}$$

where $\mathbf{c}_t(\cdot)$ and $\kappa_{n,t}$ denote, respectively, the conditional cumulant-generating function and the n th conditional cumulant of log exchange rate appreciation at time t . In particular, $\kappa_{2,t} = \sigma_t^2$ is the conditional variance and $\kappa_{4,t}/\sigma_t^4$ the excess kurtosis of $\log e_{i,t+1}$. (For more on cumulants, see Backus, Foresi, and Telmer 2001 and Martin 2013a.)

To get a sense of the size of the convexity correction, note that if the exchange rate is lognormal then all higher cumulants are 0: $\kappa_{n,t} = 0$ for $n > 2$. Thus if exchange rate volatility, σ_t , is on the order of 10 percent, the two perspectives should disagree by about 1 percent (so in the example above, expected euro appreciation of 2 percent would be consistent with expected dollar depreciation of 1 percent). In Section IIA, we show that the convexity gap observed in our data is consistent with this calculation.

II. Empirics

We obtained forward prices and quanto forward prices on the S&P 500, together with domestic and foreign interest rates, from Markit; the maturity in each case is 24 months. The data is monthly and runs from December 2009 to October 2015 for the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), euro (EUR), British pound (GBP), Japanese yen (JPY), Korean won

(KRW), Norwegian krone (NOK), Polish zloty (PLN), and Swedish krona (SEK). As these quantos are used to forecast exchange rates over a 24-month horizon, our forecasting sample runs from December 2009 to October 2017. Markit reports consensus prices based on quotes received from a wide range of financial intermediaries. These prices are used by major OTC derivatives market makers as a means of independently verifying their book valuations and to fulfill regulatory requirements; they do not necessarily reflect transaction prices. Accounting for missing entries in our panel, we have 656 currency-month observations. (Where we do not observe a price, we treat the observation as missing. Larger periods of consecutive missing observations occur only for DKK, KRW, and PLN.)

Since the financial crisis of 2007–2009, a growing literature (including Du, Tepper, and Verdelhan 2018) has discussed the failure of covered interest parity (CIP)—the no-arbitrage relation between forward exchange rates, spot exchange rates and interest rate differentials—and established that since the financial crisis, CIP frequently does not hold if interest rates are obtained from money markets. For each maturity, we observe currency-specific discount factors directly from our Markit dataset. The implied interest rates are consistent with the observed forward prices and the absence of arbitrage. Our measure of the interest rate differentials therefore does not violate the no-arbitrage condition we require for identity (6) to hold.

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

Figure 1 plots each currency's QRP over time; for clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK). The QRP is negative for JPY and positive for all other currencies (with the partial exception of EUR, for which we observe a sign change near the end of our time period).

We plot the evolution over time of ECA (solid) and of the UIP forecast (dashed) for each of the currencies in our panel in Figure IA.6 of the online Appendix. The gap between the two lines for a given currency is that currency's QRP. Table 1 reports summary statistics of ECA. The penultimate line of the Table 1 averages the summary statistics across currencies; the last line reports summary statistics for the pooled data. Table 2 reports the same statistics for IRD and QRP.

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

Table 3 reports volatilities and correlations for the time series of individual currencies' ECA, IRD, and QRP. The table also shows three aggregated measures of volatilities and correlations. The row labeled "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 labeled "Cross section" reports cross-currency volatilities and correlations of

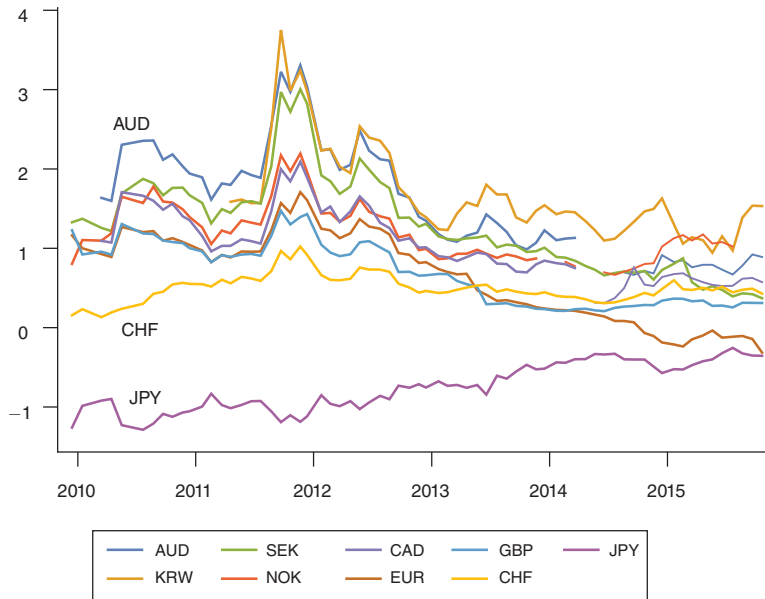


FIGURE 1. THE TIME SERIES OF QRP

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

TABLE 1—SUMMARY STATISTICS OF ECA

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

Note: This table reports annualized summary statistics (percent) of quanto-based expected currency appreciation (ECA).

time-averaged ECA, IRD, and QRP. Lastly, the row labeled “Pooled” averages on both dimensions: it reports volatilities and correlations for the pooled data.

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

TABLE 2—SUMMARY STATISTICS OF IRD AND QRP

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

Note: This table reports annualized summary statistics (percent) of UIP forecasts (IRD, panel A), and quanto-implied risk premia (QRP, panel B).

TABLE 3—VOLATILITIES AND CORRELATIONS OF ECA, IRD

	$\sigma(\text{ECA})$	$\sigma(\text{IRD})$	$\sigma(\text{QRP})$	$\rho(\text{ECA, IRD})$	$\rho(\text{ECA, QRP})$	$\rho(\text{IRD, QRP})$
AUD	0.72	1.007	0.692	0.727	-0.013	-0.696
CAD	0.526	0.353	0.441	0.558	0.748	-0.134
CHF	0.472	0.441	0.171	0.932	0.355	-0.007
DKK	0.487	0.470	0.275	0.835	0.342	-0.231
EUR	0.398	0.622	0.556	0.476	0.183	-0.777
GBP	0.350	0.223	0.389	0.137	0.822	-0.451
JPY	0.412	0.206	0.295	0.738	0.882	0.333
KRW	0.724	0.443	0.589	0.582	0.792	-0.036
NOK	0.622	0.690	0.359	0.855	0.090	-0.439
PLN	0.892	1.030	0.650	0.780	0.135	-0.514
SEK	0.656	0.905	0.616	0.733	-0.013	-0.690
Time series	0.569	0.581	0.457	0.669	0.393	-0.331
Cross section	0.786	1.242	0.751	0.817	-0.305	-0.798
Pooled	0.908	1.265	0.857	0.736	-0.026	-0.696

Notes: This table presents the standard deviations (percent) of, and correlations between, the interest rate differential (IRD), the quanto-implied risk premium (QRP), and expected currency appreciation (ECA). The row labeled *Time series* reports means of the currencies' time-series standard deviations and correlations. The row labeled *Cross section* reports cross-sectional standard deviations and correlations of time-averaged ECA, IRD, and QRP. The row labeled *Pooled* reports standard deviations and correlations of the pooled data. All quantities are expressed in annualized terms.

The correlation between IRD and QRP is negative when we pool our data ($\rho = -0.696$). Given the sign convention on IRD, this indicates that currencies with high interest rates (relative to the dollar) tend to have high risk premia; thus the predictions of the quanto theory are consistent with the carry trade literature and the findings of Lustig, Roussanov, and Verdelhan (2011). The average time-series (i.e., within-currency) correlation between IRD and QRP is more modestly negative ($\rho = -0.331$): a typical currency's risk premium tends to be higher, or less negative, at times when its interest rate is high relative to the dollar, but this tendency is fairly weak. The disparity between these two facts is accounted for by the strongly negative cross-sectional correlation between IRD and QRP ($\rho = -0.798$). If we interpret the data through the lens of Result 2, these findings suggest that the returns to the carry trade are more the result of persistent *cross-sectional* differences between currencies than of a *time-series* relationship between interest rates and risk premia. This prediction is consistent with the empirical results documented by Hassan and Mano (2017).

We see a corresponding pattern in the time-series, cross-sectional, and pooled correlations of ECA and QRP. The time-series (within-currency) correlation of the two is substantially positive ($\rho = 0.393$), while the cross-sectional correlation is negative ($\rho = -0.305$). In the time series, therefore, a rise in a given currency's QRP is associated with a rise in its expected appreciation; whereas in the cross section, currencies with relatively high QRP on average have relatively *low* expected currency appreciation on average (reflecting relatively high interest rates on average). Putting the two together, the pooled correlation is close to zero ($\rho = -0.026$). That is, Result 2 predicts that there should be no clear relationship between currency risk premia and expected currency appreciation; again, this is consistent with the findings of Hassan and Mano (2017).

These properties are illustrated graphically in Figure 2. We plot confidence ellipses centered 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 percent of its currency's observations in population. (Our interest is in the relative sizes of the ellipses: the choice of 50 percent is arbitrary.) The orientation of each ellipse illustrates the within-currency time series correlation, while the positions of the different ellipses reveal correlations across currencies. The figures refine the discussion above. QRP and IRD are negatively correlated within currency (with the exceptions of CAD, CHF, and KRW) and in the cross section. QRP and ECA are positively correlated in the time series for every currency, but exhibit negative correlation across currencies; overall, the pooled correlation between the two is close to zero.

Our empirical analysis focuses on contracts with a maturity of 24 months because these have the best data availability. But in one case—the S&P 500 index quantoed into euros—we observe a range of maturities, so can explore the term structure of QRP. We plot the time series of annualized euro-dollar QRP at horizons of 6, 12, 24, and 60 months in Figure IA.7 of the online Appendix. On average, the term structure of QRP is flat over the sample period, but QRP is slightly more volatile at shorter horizons, so that the term structure is downward-sloping when QRP spikes and upward-sloping when QRP is low.

Panel A. The relationship between QRP and IRD

Panel B. The relationship between QRP and ECA




FIGURE 2

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

A. A Consistency Check

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

Figure 3 implements (15) and (16) for the EUR/USD, JPY/USD, EUR/JPY, and EUR/CHF currency pairs. In each of the top-left, bottom-left and bottom-right panels, the solid line depicts the expected appreciation of the euro against the US dollar, yen, and Swiss franc, respectively, while the dashed line shows the expected depreciation of the three currencies against the euro (that is, we flip the sign on the “inverted” series for readability). In the top-right panel, the solid and dashed lines show the expected appreciation of the yen against the US dollar and expected depreciation of the US dollar against the yen, respectively. In every case, the two measures are strongly correlated over time and the solid line is above the dashed line, as it should be according to (18). The gaps between the measures are therefore consistent with the Jensen’s inequality correction one would expect to see if our currency forecasts measured expected currency appreciation perfectly. Moreover, given that annual exchange rate volatilities are on the order of 10 percent, the sizes of the gaps between the measures are quantitatively consistent with the Jensen’s inequality correction derived at the end of Section IB.

The EUR/CHF pair in the bottom-right panel represents a particularly interesting case study. The Swiss national bank instituted a floor on the EUR/CHF exchange rate at CHF1.20/€ in September 2011 and consequently also reduced the conditional volatility of the exchange rate. Following this, the two lines converge and the gap remains narrow, at around 0.2 percent, until January 2015 when the sudden

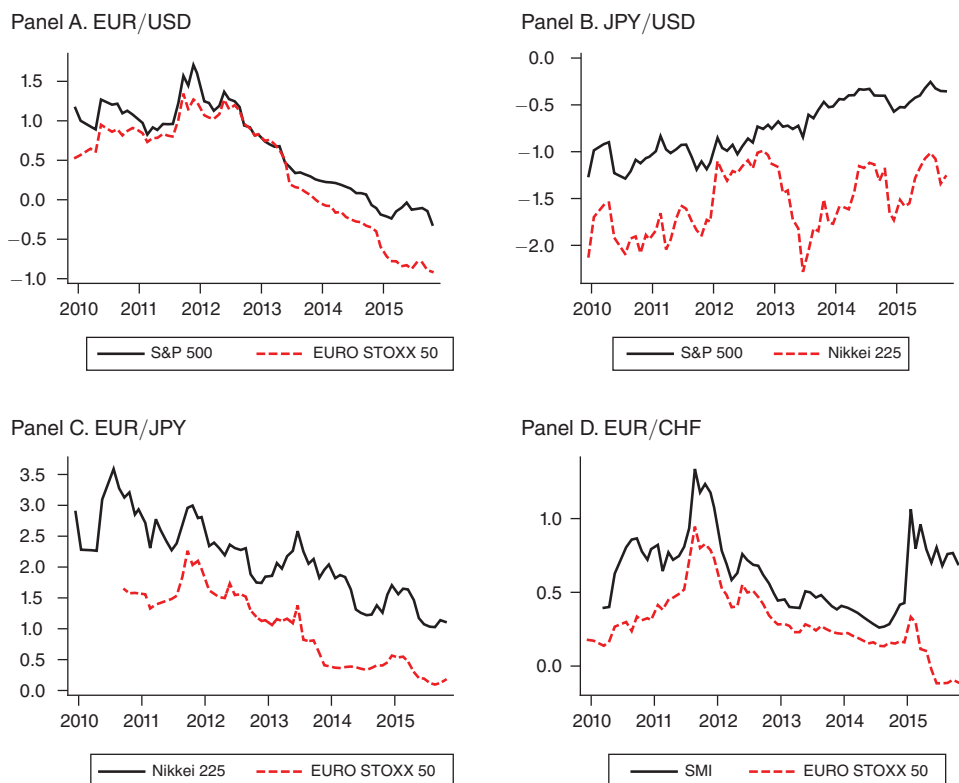


FIGURE 3

Notes: 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 line plots ECA as perceived by a log investor fully invested in the S&P 500 (panels A and B), Nikkei 225 (panel C), and SMI (panel D), respectively. The dashed 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.

removal of the floor prompted a spike in the volatility of the currency pair, visible in the figure as the point at which the two lines diverge.

B. Return Forecasting

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

To provide a sense of the data before turning to our regression results, Figures 4 and 5 represent our baseline univariate regressions graphically in the same manner as in Figure 2. Figure 4 plots realized currency excess returns (RXR) against QRP and against IRD.¹¹ 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 5 shows the corresponding results for realized currency appreciation (RCA). Panel A suggests that the within-currency correlation with the quanto predictor ECA is predominantly positive (with the exceptions of AUD and CHF), as is the cross-sectional correlation. In contrast, panel B suggests that the correlation between realized currency appreciation and interest rate differentials is close to zero both within and across currencies, consistent with the view that interest rate differentials do not help to forecast currency appreciation.

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

$$(19) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \beta QRP_{i,t} + \gamma IRD_{i,t} + \varepsilon_{i,t+1}.$$

Here (and from now on) the length of the period from t to $t + 1$ over which we measure our return realizations is 24 months, corresponding to the forecasting horizon dictated by the maturity of the quanto contracts we observe in our data.

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

$$(20) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \beta QRP_{i,t} + \varepsilon_{i,t+1},$$

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

$$(21) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \gamma IRD_{i,t} + \varepsilon_{i,t+1}.$$

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

Table 4 reports the results. We report coefficient estimates and R^2 for each regression, with and without currency fixed effects; standard errors are shown in parentheses. These standard errors are computed via a nonparametric bootstrap to account for heteroskedasticity, cross-sectional and serial correlation in our data. (The serial correlation arises due to overlapping observations: we make forecasts of 24-month excess returns at monthly intervals.) For comparison, these nonparametric standard errors exceed those obtained from a parametric residual bootstrap by a factor of

¹¹ As noted in Section I, we work with true returns as opposed to log returns. Engel (2016) points out that it may not be appropriate to view log returns as approximating true returns, as the gap between the two is of a similar order of magnitude as the risk premium itself.

Panel A. Realized currency excess return against QRP, computed from (14)

Panel B. Realized currency excess return against IRD




FIGURE 4

Notes: Realized and expected currency excess return according to (panel A) the quanto theory and (panel 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 percent of its currency’s data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse’s size reflects their volatilities. Panel A shows a dotted 45-degree line for comparison.

Panel A. Realized currency appreciation against ECA, computed from (14)

Panel B. Realized currency appreciation against IRD




FIGURE 5

Notes: Realized and expected currency excess return according to (panel A) the quanto theory and (panel 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 percent of its currency’s data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse’s size reflects their volatilities. Panel A shows a dotted 45-degree line for comparison.

around 2, and Hansen-Hodrick standard errors by a factor of around 1.3. We provide a detailed description of our bootstrap procedure and address potential small-sample concerns in Section IIF.

The estimated coefficient on the quanto-implied risk premium is positive and economically large in every specification in which it occurs. Moreover, the R^2 values are substantially higher in the two regressions (19) and (20) that feature the quanto-implied risk premium than in the regression (21) in which it does not occur. The estimate for β in our headline regression (20) is 2.604 (standard error 1.127) in the pooled regression and 4.995 (standard error 1.565) in the regression with fixed effects. The fact that these estimates are above 1 raises the possibility that beyond

TABLE 4—CURRENCY EXCESS RETURN FORECASTING REGRESSIONS

Regression	(19)	(20)	(21)
<i>Panel A. Pooled panel regressions</i>			
α (p.a.)	-0.048 (0.020)	-0.047 (0.019)	-0.030 (0.014)
β	3.394 (1.734)	2.604 (1.127)	
γ	0.769 (1.040)		-0.832 (0.651)
R^2	19.13	17.43	3.88
<i>Panel B. Panel regressions with currency fixed effects</i>			
β	5.456 (2.046)	4.995 (1.565)	
γ	0.717 (1.411)		-1.363 (1.001)
R^2	22.60	22.03	2.77

Notes: Return realizations correspond to the forecasting horizon of 24 months. The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a non-parametric block bootstrap) in parentheses, as well as R^2 (percent).

its direct importance in (6), the quanto-implied risk premium may also proxy for the second covariance term.¹² We explore this issue in Section IIE. Another noteworthy qualitative feature of our results is the consistently negative intercept, which reflects an unexpectedly strong dollar over our sample period; we discuss the statistical interpretation of this fact in Section IIF.

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

$$(22) \quad \frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta QRP_{i,t} + \gamma IRD_{i,t} + \varepsilon_{i,t+1}.$$

We do so not because we are interested in the coefficient estimates, which are mechanically related to those of regression (19), 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,

$$(23) \quad \frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta QRP_{i,t} + \varepsilon_{i,t+1},$$

¹² Another possibility is that it is more reasonable to think of a log investor as wishing to hold a levered position in the market (so $w > 1$ in the notation of footnote 7). If so, we should find a coefficient on QRP that is larger than one. We are cautious about suggesting this as an explanation, however, because a log investor would never risk bankruptcy. To match the point estimate for specification (20), we would need $w = 2.604$ or $w = 4.995$ (respectively without and with fixed effects). In the latter case, the investor would go bankrupt if the market dropped by 20 percent over the 2-year horizon.

TABLE 5—CURRENCY FORECASTING REGRESSIONS

Regression	(22)	(23)	(24)
<i>Panel A. Pooled panel regressions</i>			
α (p.a.)	−0.048 (0.020)	−0.045 (0.019)	−0.030 (0.014)
β	3.394 (1.726)	1.576 (1.172)	
γ	1.769 (1.045)		0.168 (0.651)
R^2	16.01	6.63	0.16
<i>Panel B. Panel regressions with currency fixed effects</i>			
β	5.456 (2.047)	4.352 (1.682)	
γ	1.717 (1.414)		−0.363 (1.007)
R^2	20.56	17.16	0.20

Notes: Return realizations correspond to the forecasting horizon of 24 months. The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a non-parametric block bootstrap) in parentheses, as well as R^2 (percent).

and onto interest rate differentials,

$$(24) \quad \frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \gamma \text{IRD}_{i,t} + \varepsilon_{i,t+1}.$$

As previously, we also run the three regressions (22)–(24) with fixed effects.

The regression results are shown in Table 5, which is structured similarly to Table 4. There is little evidence that the interest rate differential helps to forecast currency appreciation on its own; this is consistent with the previous set of results and with the large literature that documents the failure of UIP. In the pooled panel, the estimated γ in regression (24) is close to zero, and the R^2 is essentially zero. With fixed effects, the estimate of γ is marginally negative, providing weak evidence that currencies tend to appreciate against the dollar when their interest rate relative to the dollar is higher than its time-series mean.

More strikingly, the quanto-implied risk premium makes a very large difference in terms of R^2 , which increases by two orders of magnitude when moving from specification (24) to (22) in both the pooled regressions (0.16 percent to 16.01 percent) and the fixed effects regressions (0.20 percent to 20.56 percent). It is also interesting that when QRP is included in the regressions (with or without fixed effects) the coefficient estimate on IRD, γ , increases toward the value of 1 predicted by Result 2.

For completeness, we report the results of running regressions (20), (21), (22), and (24) separately for each currency at the 24-month horizon (and at 6- and 12-month horizons for the euro) in Table IA.5 of the online Appendix. Consistent with the previous literature (for example, Fama 1984 and Hassan and Mano 2017), the coefficient estimates are extremely noisy. A further appealing feature of Result 2 is that

it provides a justification for constraining all the coefficient on the quanto-implied risk premium to be equal across currencies, as we have done above.

C. Risk-Neutral Covariance versus True Covariance

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

$$(25) \quad E_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = R_{f,t}^{\$} \text{cov}_t \left(\frac{e_{i,t+1}}{e_{i,t}}, -\frac{1}{R_{t+1}} \right).$$

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

The fundamental challenge for a test of this prediction is that forward-looking true covariance is not directly observed. This is the major advantage of our approach: risk-neutral covariance *is* directly observed via the quanto-implied risk premium. That said, we attempt to test (25) by using lagged realized covariance, RPCL, as a proxy for true forward-looking covariance.

The results are shown in Table 6 of the Appendix. RPCL is positively related to subsequently realized currency excess returns, as suggested by (25), but it is not statistically significant in our sample, and is driven out as a predictor by risk-neutral covariance (QRP), consistent with Result 2.

In principle, this might simply indicate that lagged realized covariance is an imperfect proxy for true forward-looking covariance: perhaps the success of QRP simply reflects its superiority as a forecaster of realized covariance? Table 6 shows that risk-neutral covariance *is*, individually, a statistically significant forecaster of future realized covariance. But it is driven out when lagged realized covariance and the interest-rate differential are included in the multivariate regression (31). Moreover, the optimal covariance forecast generated by this multivariate regression is driven out by QRP in the excess-return-forecasting regression (32).

The relationship between risk-neutral covariance and true covariance is interesting in its own right. Figure 6 illustrates the empirical relationship between the covariance forecast obtained from regression (31) (our proxy for forward-looking true covariance) and forward-looking risk-neutral covariance (obtained from quanto contracts). The two are positively correlated in the cross section and in the time series, but risk-neutral covariance is generally larger (smaller) than future realized covariance for currencies with positive (negative) risk-neutral covariances. This is consistent with the observation of Lettau, Maggiori, and Weber (2014) that carry trade returns are more correlated with the market at times of negative market returns. As we will now see, it is problematic for lognormal models.

TABLE 6—REALIZED COVARIANCE REGRESSIONS

Regression	(28)	(29)	(30)	(31)	(32)
<i>Panel A. Pooled panel regression</i>					
α (p.a.)	-0.034 (0.017)	-0.047 (0.018)	-0.000 (0.001)	0.000 (0.001)	-0.047 (0.018)
β		2.798 (1.366)	0.447 (0.158)	-0.026 (0.126)	3.096 (1.639)
γ	1.307 (1.111)	-0.213 (1.193)		0.370 (0.123)	-1.103 (3.206)
δ				-0.131 (0.061)	
R^2	7.37	17.52	36.56	66.44	17.94
<i>Panel B. Panel regression with currency fixed effects</i>					
β		4.643 (2.006)	0.330 (0.168)	-0.107 (0.017)	4.988 (2.073)
γ	1.967 (1.474)	0.387 (1.384)		0.313 (0.125)	0.023 (3.300)
δ				-0.237 (0.138)	
R^2	9.14	22.27	9.43	45.69	22.03

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

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

where the summation is over daily returns on trading days s preceding t over a time frame corresponding to our forecasting horizon, h , so that $RPCL_{i,t}$ is observable at time t . We also define a realized covariance measure $RPC_{i,t}$ that is analogous to the above definition except that the summation is over trading days following t over the appropriate time-frame (so that it is not observable until time $t+h$). We test whether risk-neutral covariance forecasts realized covariance, in a univariate regression as well as in the presence of lagged realized covariance and IRD as competing predictors. Lastly, we denote by $\widehat{RPC}_{i,t}$ the optimal forecast of $RPC_{i,t}$ from regression (31) and test whether it forecasts excess returns. Now,

$$(28) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \gamma RPCL_{i,t} + \varepsilon_{i,t+1},$$

$$(29) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \beta QRP_{i,t} + \gamma RPCL_{i,t} + \varepsilon_{i,t+1},$$

$$(30) \quad RPC_{i,t} = \alpha + \beta QRP_{i,t} + \varepsilon_{i,t+1},$$

$$(31) \quad RPC_{i,t} = \alpha + \beta QRP_{i,t} + \gamma RPCL_{i,t} + \delta IRD_{i,t} + \varepsilon_{i,t+1},$$

$$(32) \quad \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^i} = \alpha + \beta QRP_{i,t} + \gamma \widehat{RPC}_{i,t} + \varepsilon_{i,t+1}.$$

Return realizations correspond to the forecasting horizon of 24 months. We report coefficient estimates for each regression, with standard errors (computed using a nonparametric block bootstrap) in brackets. See Section IIF for more details.

D. Lognormal Models

Lognormal models impose a tight connection between the covariance risk premium and the market and currency risk premium. Define the equity premium $ERP_t = \log E_t(R_{t+1}/R_{f,t}^{\$})$ and currency risk premium $CRP_{i,t} = \log E_t(\tilde{R}_{i,t+1}/R_{f,t}^{\$})$ where $\tilde{R}_{i,t+1} = R_{f,t}^i e_{i,t+1}/e_{i,t}$ is the return on the currency trade defined earlier.

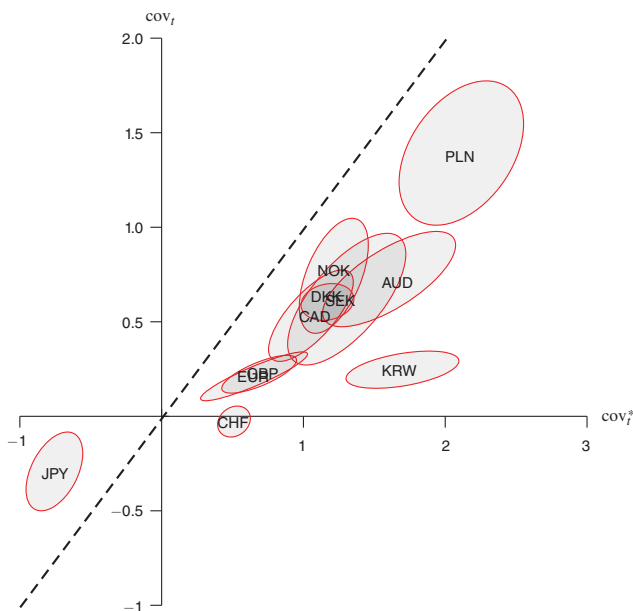


FIGURE 6

Notes: Risk-neutral and optimally predicted covariances of exchange rate movements and S&P returns. The centre of each confidence ellipse represents a currency’s average risk-neutral and realized covariance. In population, each ellipse would contain 20 percent of its currency’s data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and risk-neutral covariance for the given currency, while the ellipse’s size reflects their volatilities. We plot a dotted 45-degree line for comparison.

RESULT 3 (The Covariance Risk Premium in Lognormal Models): *Suppose that the market return, exchange rate, and SDF are conditionally jointly lognormal. Then we have*

$$(26) \quad \log \frac{\text{cov}_t(R_{t+1}, e_{i,t+1}/e_{i,t})}{\text{cov}_t^*(R_{t+1}, e_{i,t+1}/e_{i,t})} = \text{ERP}_t + \text{CRP}_{i,t},$$

or equivalently,

$$(27) \quad \text{cov}_t(r_{t+1}, \Delta e_{i,t+1}) = \text{cov}_t^*(r_{t+1}, \Delta e_{i,t+1}),$$

where $r_{t+1} = \log R_{t+1}$ and $\Delta e_{i,t+1} = \log(e_{i,t+1}/e_{i,t})$.

PROOF:

See Appendix.

Empirically, it is plausible that the right-hand side of (26) is positive for most currencies (the yen being a possible exception). But we find that the left-hand side is typically negative in our data. No lognormal model can match these patterns.

It is nonetheless an interesting exercise to see how the quanto risk premium (and the residual covariance term, which would be zero from the perspective of the log investor) behaves inside an equilibrium model. As QRP has a simple characterization in terms of risk-neutral covariance, this is an easy exercise to carry out in any equilibrium model; we suggest that it makes an interesting diagnostic for future generations of international finance models. In that spirit, we have calculated the currency risk premium, QRP, IRD, and the residual covariance term within the model of Colacito and Croce (2011).

The results are shown in Section IA.B of the online Appendix. We deviate from the symmetric baseline calibration of Colacito and Croce in order to generate a nontrivial currency risk premium. The comparative statics of their long-run risk model are such that our calibrations which yield a positive asymmetric currency risk premium generate positive risk-neutral covariance (QRP) and a positive residual. In this model, the residual covariance term therefore adds to the prediction of the quanto forecast, as opposed to offsetting it. This positive relationship between risk-neutral covariance and the residual is consistent with our finding that the slope coefficients on QRP in the predictive regressions in Section IIB are generally larger than one.

E. Beyond the Log Investor

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

To get a sense of what these assumptions may leave out, we conduct a principal components analysis on unexpected currency excess returns: that is, on the difference between realized currency excess returns and the corresponding ex ante expected returns. We calculate these unexpected excess returns in two ways. *Regression residuals* are defined as the estimated residuals $\varepsilon_{i,t+1}$ in the specification of regression (20) that includes currency fixed effects. *Theory residuals* are defined similarly, except that we impose $\alpha = 0, \beta = 1$ in (20).

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

We decompose the theory and regression residuals into their respective principal components (dropping DKK, KRW, and PLN from the panel to minimize the impact of missing observations). Table IA.2, in the online Appendix, shows the principal component loadings. The first principal component, which explains just under two thirds of the variation in residuals, can be interpreted as a level, or “dollar,” factor since it loads positively on all currencies (with the exception of GBP, in the case of the regression residuals).

Motivated by this fact, we now include an additional predictor variable, $\overline{\text{IRD}}_t$, which is calculated as the cross-sectional average of the interest rate differentials in our balanced panel of eight currencies (i.e., excluding DKK, KRW, and

TABLE 7—BEYOND THE LOG INVESTOR

Regressor	Univariate	Bivariate	3-variate	4-variate
<i>Panel regressions with currency fixed effects</i>				
QRP, β	4.995 (1.565)	5.654 (1.402)	3.799 (1.657)	3.541 (1.836)
IRD, γ				-1.059 (1.573)
$\overline{\text{IRD}}$, δ			-5.060 (1.605)	-4.266 (1.538)
RER, ζ		-0.413 (0.136)	-0.780 (0.159)	-0.804 (0.188)
R^2	22.03	35.40	43.56	44.09

Notes: This table reports the R^2 -maximizing univariate, bivariate, 3-variate, and 4-variate specifications in regressions of 24-month realized currency excess returns onto combinations of QRP, IRD, the average forward discount $\overline{\text{IRD}}$, and the real exchange rate, q . The table reports standard errors (computed using a nonparametric block bootstrap) in brackets. See Section IIE for more detail. The last line reports R^2 in percent.

PLN); Lustig, Roussanov, and Verdelhan (2014) interpret this average interest rate differential (which they refer to as the “average forward discount”) as a dollar factor and show that it helps to forecast currency returns. We also include the logarithm of the real exchange rate, which Dahlquist and Penasse (2017) have shown to be a successful forecaster of currency returns.

Table 7 reports the results of regressions of currency excess returns onto currency fixed effects and subsets of four forecasting variables: the quanto-implied risk premium (QRP), the interest rate differential (IRD), the real exchange rate (RER), and the average interest rate differential ($\overline{\text{IRD}}$). The table reports the univariate, bivariate, 3-variate, and 4-variate specifications with the highest R^2 . (Table IA.3 of the online Appendix reports the R^2 for all $2^4 - 1 = 15$ subsets of the four explanatory variables, though not—for lack of space—the estimated coefficients.) The quanto-implied risk premium features in all R^2 -maximizing regressions. The estimates of β are larger than 1 in every specification, suggesting that, over and above its relevance as a direct measure of risk-neutral covariance, the quanto-implied risk premium helps to capture the physical covariance term in (6). As we increase from one to two to three explanatory variables, R^2 increases from 22.03 percent (using QRP alone) to 35.40 percent (adding the real exchange rate) to 43.56 percent (adding the dollar factor $\overline{\text{IRD}}$). The interest rate differential itself, IRD, contributes almost no further explanatory power when it is then added as a fourth variable.

As the real exchange rate performs well, we report further results relating to it in Table IA.4 of the online Appendix.

F. Joint Hypothesis Tests and Finite-Sample Issues

We now consider the joint hypothesis tests that are suggested by Result 2. In our three main specifications (19), (20), and (22), equation (14) predicts an intercept

$\alpha = 0$, and a slope coefficient on QRP $\beta = 1$. For the excess return forecast in regression (19), it predicts that the interest rate differential should have no predictive power, i.e., $\gamma = 0$; whereas it predicts that $\gamma = 1$ in the currency-appreciation regression (22).

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

Table 8 reports p -values for tests of various hypotheses about our baseline regressions. In addition to conventional p -values calculated using the asymptotic (χ^2) distribution of the Wald test statistic, the table also reports more conservative small-sample p -values obtained from a bootstrapped test statistic distribution. We compute these small-sample p -values by constructing a small-sample distribution of the Wald test statistic for each regression. We simulate 5,000 sets of monthly data for the LHS variable under the null hypothesis of no predictability, such that the simulated data matches the monthly autocorrelation and covariance matrix of the realized, observed LHS data. We then aggregate the simulated monthly data into 24-month horizon data, like the LHS data used in our regressions (e.g., excess returns over 24 months). As we aim to measure the small-sample performance of our bootstrap routine, the simulated datasets each have the same number of data points as the observed LHS data. For each specification, we then regress the 5,000 simulated LHS data on the respective observed RHS variable(s). Where we run the regression with currency fixed effects, we use the demeaned RHS variable(s). We obtain the point estimates of the coefficients and their covariance matrix from the bootstrap routine outlined above and use the test statistics from these 5,000 regressions to construct the empirical small-sample distribution of the respective Wald statistic under the respective null hypothesis. This procedure also accounts for the potential small-sample Stambaugh bias in the p -values.

Figure IA.8, in the online Appendix, illustrates by plotting the histograms of the bootstrapped distribution of test statistics for various hypotheses on regression (22). Panels A and B show the finite-sample bootstrapped distributions of the test statistic for the hypothesis that Result 2 holds, respectively in the pooled and fixed effects regressions. The value of the test statistic in the data is indicated with an asterisk in each panel. The finite-sample and asymptotic (shown with a solid line) distributions are strikingly different. The asymptotic distribution suggests that we can reject the

TABLE 8—JOINT TESTS OF STATISTICAL SIGNIFICANCE

Regression:	(19)	(20)	(22)
<i>Panel A. Pooled panel regression</i>			
$H_0^1: \alpha = \gamma = 0, \beta = 1$	0.029/0.357		
$H_0^1: \alpha = 0, \beta = 1$		0.039/0.342	
$H_0^1: \alpha = 0, \beta = \gamma = 1$			0.030/0.340
$H_0^2: \beta = 1, \gamma = 0$	0.342/0.546		
$H_0^2: \beta = 1$		0.155/0.299	
$H_0^2: \beta = 1, \gamma = 1$			0.339/0.493
$H_0^3: \beta = 0$	0.050/0.088	0.021/0.097	0.049/0.082
<i>Panel B. Panel regression with currency fixed effects</i>			
$H_0^2: \beta = 1, \gamma = 0$	0.029/0.256		
$H_0^2: \beta = 1$		0.011/0.163	
$H_0^2: \beta = 1, \gamma = 1$			0.029/0.238
$H_0^3: \beta = 0$	0.008/0.051	0.001/0.089	0.008/0.051

Notes: Realizations for excess returns and currency appreciation correspond to the forecasting horizon of 24 months. The table reports p -values of Wald tests of various hypotheses on the regression coefficients. H_0^1 is the hypothesis suggested by Result 2: $\alpha = \gamma = 0$ and $\beta = 1$ in regression (19), $\alpha = 0$ and $\beta = 1$ in regression (20), and $\alpha = 0$ and $\beta = \gamma = 1$ in regression (22). Hypothesis H_0^2 drops the constraint that $\alpha = 0$, and therefore tests our model's ability to predict differences in currency returns but not its ability to predict the absolute level of (dollar) returns. Hypothesis H_0^3 is that QRP is not useful for forecasting. For each Wald test, we report both the asymptotic p -values obtained from the χ^2 distribution and p -values from a bootstrapped small-sample distribution (in the format asymptotic p -value/small-sample p -value).

hypothesis that Result 2 holds, but this conclusion is overturned by the finite-sample distribution. (In the pooled case, the discrepancy is largely due to the intercept, as becomes clear on comparing the asymptotic p -values for tests of hypotheses H_0^1 and H_0^2 in Table 8. The asymptotic distribution strongly penalizes the fact that the US dollar was strong over our sample period, whereas the finite-sample distribution does not.)

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

We reach similar conclusions for regressions (19) and (20): we do not reject the predictions of Result 2 in the joint Wald tests for any of the three baseline regressions using the small-sample distribution of the test statistic; and QRP remains individually significant as a predictor at the 10 percent level in all three specifications, with and without currency fixed effects, even if we take the most conservative approach to computing p -values that relies on the empirical small-sample test statistic distribution.

III. Out-of-Sample Prediction

We now test the quanto theory out of sample. Since the dollar strengthened strongly over the relatively short time period spanned by our data (as reflected in the negative intercept in our pooled panel regression (22)), we focus on forecasting differential currency appreciation. That is, we seek to predict, for example, the *relative* performance of dollar-yen versus dollar-euro.

In the previous section, we estimated the loadings on the quanto-implied risk premium, QRP, and interest rate differential, IRD, via panel regressions. These deliver the best in-sample coefficient estimates in a least-squares sense. But for an out-of-sample test we must pick the loadings a priori. Here we can exploit the distinctive feature of Result 2 that it makes specific quantitative predictions for the loadings: each should equal 1, as in the formula (14). We therefore compute out-of-sample forecasts by fixing the coefficients that appear in (22) at their theoretical values: $\alpha = 0$, $\beta = 1$, and $\gamma = 1$.

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

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

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

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

The results of this exercise are reported in Table 9. The quanto theory outperforms each of the three competitors: when the competitor model is UIP, we find that $R_{OS}^2 = 10.91$ percent; and when it is relative PPP, we find $R_{OS}^2 = 26.05$ percent. In our sample, the toughest benchmark is the random walk forecast, consistent with the findings of Rossi (2013). Nonetheless, the quanto theory easily outperforms it, with $R_{OS}^2 = 9.57$ percent.

To get a sense for whether our positive results are driven by a small subset of the currencies, Table 9 also reports the results of splitting the R^2 measure currency-by-currency: for each currency i , we define

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

TABLE 9—OUT-OF-SAMPLE FORECAST PERFORMANCE

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

Notes: The second panel reports R_{OS}^2 measures by currency. (All R_{OS}^2 measures are reported in percent.) 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.

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

Jordà and Taylor (2012) have argued that assessments of forecast performance based solely on mean squared errors may not fully reflect the economic benefits of a forecasting model. In the online Appendix, we use the approach they suggest, which essentially asks whether a predictor variable is more or less successful than competitor predictors at predicting whether a currency will appreciate or depreciate. Our approach also outperforms on their metric, both in forecasting currency excess returns and in forecasting currency appreciation.

IV. Conclusion

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

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

strength of the dollar itself on average against other currencies over our sample period: time will tell if this is a small-sample issue or something more fundamental.

We find that currencies tend to have high quanto-implied risk premia if they have high interest rates on average, relative to other currencies (a cross-sectional statement), or if they currently have unusually high interest rates (a time-series statement); and that there is more cross-sectional than time-series variation in quanto-implied risk premia. These facts explain both the existence of the carry trade and the empirical importance of persistent cross-currency asymmetries, as documented by Hassan and Mano (2017).

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

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

APPENDIX

Suppose that the SDF, X_{t+1} and Y_{t+1} are conditionally jointly lognormal, and write lower-case variables for logs of the corresponding uppercase variables. Assume further that X_{t+1} and Y_{t+1} are tradable. Then we have the following three facts:

$$\begin{aligned}\text{cov}_t(X_{t+1}, Y_{t+1}) &= E_t X_{t+1} E_t Y_{t+1} (e^{\text{cov}_t(x_{t+1}, y_{t+1})} - 1), \\ \text{cov}_t^*(X_{t+1}, Y_{t+1}) &= \text{cov}_t(X_{t+1}, Y_{t+1}) e^{\text{cov}_t(m_{t+1}, x_{t+1} + y_{t+1})}, \\ \text{cov}_t^*(X_{t+1}, Y_{t+1}) &= E_t^* X_{t+1} E_t^* Y_{t+1} (e^{\text{cov}_t^*(x_{t+1}, y_{t+1})} - 1).\end{aligned}$$

These follow by direct calculation because $\log E_t Z_{t+1} = E_t \log Z_{t+1} + \frac{1}{2} \text{var}_t \log Z_{t+1}$ for any conditionally lognormal random variable Z_{t+1} (and using the definition (4) of the risk-neutral measure to derive the second and third facts).

The first fact implies that equation (2) can be rewritten (in the lognormal case) as

$$\log E_t \frac{\tilde{R}_{t+1}}{R_{f,t}^{\$}} = -\text{cov}_t(m_{t+1}, \tilde{r}_{t+1}),$$

and in particular that $ERP_t = -\text{cov}_t(m_{t+1}, r_{t+1})$ and $CRP_{i,t} = -\text{cov}_t(m_{t+1}, \Delta e_{i,t+1})$, where ERP_t and $CRP_{i,t}$ are defined in the main text and we write $r_{t+1} = \log R_{t+1}$ and $\Delta e_{i,t+1} = \log(e_{i,t+1}/e_{i,t})$. Combined with the second fact, this gives (in the lognormal case) equation (26) in the main text:

$$\log \frac{\text{cov}_t(R_{t+1}, e_{i,t+1}/e_{i,t})}{\text{cov}_t^*(R_{t+1}, e_{i,t+1}/e_{i,t})} = ERP_t + CRP_{i,t}.$$

To see that this is equivalent to (27), exponentiate both sides and use the definitions of ERP_t and $CRP_{i,t}$, together with the first and third facts above, to conclude that

$$\frac{E_t R_{t+1} E_t e_{i,t+1}/e_{i,t} \{e^{\text{cov}(r_{t+1}, \Delta e_{i,t+1})} - 1\}}{E_t^* R_{t+1} E_t^* e_{i,t+1}/e_{i,t} \{e^{\text{cov}^*(r_{t+1}, \Delta e_{i,t+1})} - 1\}} = E_t \frac{R_{t+1}}{R_{f,t}^{\$}} E_t \frac{R_{f,t}^i e_{i,t+1}}{R_{f,t}^{\$} e_{i,t}}.$$

By the definition (4) of the risk-neutral measure, we have $E_t^* R_{t+1} = R_{f,t}^{\$}$; and similarly we have $E_t^* e_{i,t+1}/e_{i,t} = R_{f,t}^{\$}/R_{f,t}^i$ by equation (5). Equation (27) follows.

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.

REFERENCES

Adler, Michael, and Bernard Dumas. 1983. "International Portfolio Choice and Corporation Finance: A Synthesis." *Journal of Finance* 38 (3): 925–84.

Augustin, Patrick, Mikhail Chernov, and Dongho Song. 2018. "Sovereign Credit Risk and Exchange Rates: Evidence from CDS Quanto Spreads." NBER Working Paper 24506.

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

Barroso, Pedro, and Pedro Santa-Clara. 2015. "Beyond the Carry Trade: Optimal Currency Portfolios." *Journal of Financial and Quantitative Analysis* 50 (5): 1037–56.

- Breeden, Douglas T., and Robert H. Litzenberger.** 1978. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51 (4): 621–51.
- Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen.** 2008. "Carry Trades and Currency Crashes." NBER Macroeconomics Annual 2008, Vol. 23, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 313–47. Chicago: University of Chicago Press.
- Burnside, Craig.** 2011. "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Comment." *American Economic Review* 101 (7): 3456–76.
- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo.** 2011. "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial Studies* 24 (3): 853–91.
- Campbell, John Y., Karine Serfaty-de Medeiros, and Luis M. Viceira.** 2010. "Global Currency Hedging." *Journal of Finance* 65 (1): 87–121.
- Cenedese, Gino, Richard Payne, Lucio Sarno, and Giorgio Valente.** 2016. "What Do Stock Markets Tell Us About Exchange Rates?" *Review of Finance* 20 (3): 1045–80.
- Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze.** 2018. "Crash Risk in Currency Returns." *Journal of Financial and Quantitative Analysis* 53 (1): 137–70.
- Coourdacier, Nicolas, and H el ene Rey.** 2013. "Home Bias in Open Economy Financial Macroeconomics." *Journal of Economic Literature* 51 (1): 63–115.
- Colacito, Riccardo, and Mariano M. Croce.** 2011. "Risks for the Long Run and the Real Exchange Rate." *Journal of Political Economy* 119 (1): 153–82.
- Dahlquist, Magnus, and Julien Penasse.** 2017. "The Missing Risk Premium in Exchange Rates." Swedish House of Finance Research Paper 17–18.
- Della Corte, Pasquale, Lucio Sarno, Maik Schmeling, and Christian Wagner.** 2016. "Sovereign Risk and Currency Returns." Unpublished.
- Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno.** 2016. "Volatility Risk Premia and Exchange Rate Predictability." *Journal of Financial Economics* 120 (1): 21–40.
- Diebold, Francis X., and Roberto S. Mariano.** 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13 (3): 253–63.
- Du, Wenxin, Alexander Tepper, and Adrien Verdelhan.** 2018. "Deviations from Covered Interest Rate Parity." *Journal of Finance* 73 (3): 915–57.
- Engel, Charles.** 2016. "Exchange Rates, Interest Rates, and the Risk Premium." *American Economic Review* 106 (2): 436–74.
- Fama, Eugene F.** 1984. "Forward and Spot Exchange Rates." *Journal of Monetary Economics* 14 (3): 319–38.
- Farhi, Emmanuel, and Xavier Gabaix.** 2016. "Rare Disasters and Exchange Rates." *Quarterly Journal of Economics* 131 (1): 1–52.
- French, Kenneth R., and James M. Poterba.** 1991. "Investor Diversification and International Equity Markets." *American Economic Review* 81 (2): 222–26.
- Gourinchas, Pierre-Olivier, and H el ene Rey.** 2007. "International Financial Adjustment." *Journal of Political Economy* 115 (4): 665–703.
- Hansen, Lars Peter, and Robert J. Hodrick.** 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy* 88 (5): 829–53.
- Hansen, Lars Peter, and Ravi Jagannathan.** 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy* 99 (2): 225–62.
- Harvey, David, Stephen Leybourne, and Paul Newbold.** 1997. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting* 13 (2): 281–91.
- Hassan, Tarek A.** 2013. "Country Size, Currency Unions, and International Asset Returns." *Journal of Finance* 68 (6): 2269–308.
- Hassan, Tarek A., and Rui C. Mano.** 2017. "Forward and Spot Exchange Rates in a Multi-Currency World." NBER Working Paper 20294.
- Jord a,  oscar, and Alan M. Taylor.** 2012. "The Carry Trade and Fundamentals: Nothing to Fear but FEER Itself." *Journal of International Economics* 88 (1): 74–90.
- Jurek, Jakub W.** 2014. "Crash-Neutral Currency Carry Trades." *Journal of Financial Economics* 113 (3): 325–47.
- Lettau, Martin, Matteo Maggiori, and Michael Weber.** 2014. "Conditional Risk Premia in Currency Markets and Other Asset Classes." *Journal of Financial Economics* 114 (2): 197–225.
- Lewis, Karen K.** 1999. "Trying to Explain Home Bias in Equities and Consumption." *Journal of Economic Literature* 37 (2): 571–608.
- Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan.** 2011. "Common Risk Factors in Currency Markets." *Review of Financial Studies* 24 (11): 3731–77.
- Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan.** 2014. "Countercyclical Currency Risk Premia." *Journal of Financial Economics* 111 (3): 527–53.

- Lustig, Hanno, and Adrien Verdelhan.** 2007. "The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk." *American Economic Review* 97 (1): 89–117.
- Mano, Rui C.** 2013. "Exchange Rates Upon Sovereign Default." Unpublished.
- Martin, Ian W. R.** 2013a. "Consumption-Based Asset Pricing with Higher Cumulants." *Review of Economic Studies* 80 (2): 745–73.
- Martin, Ian W. R.** 2013b. "The Forward Premium Puzzle in a Two-Country World." NBER Working Paper 17564.
- Martin, Ian W. R.** 2017. "What Is the Expected Return on the Market?" *Quarterly Journal of Economics* 132 (1): 367–433.
- Martin, Ian W. R.** 2018. "Options and the Gamma Knife." *Journal of Portfolio Management* 44 (6): 47–55.
- Martin, Ian W. R., and Christian Wagner.** Forthcoming. "What is the Expected Return on a Stock?" *Journal of Finance*.
- Meese, Richard A., and Kenneth Rogoff.** 1983. "Empirical Exchange Rate Models of the Seventies: Do They Fit out of Sample?" *Journal of International Economics* 14 (1–2): 3–24.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf.** 2012. "Carry Trades and Global Foreign Exchange Volatility." *Journal of Finance* 67 (2): 681–718.
- Miranda-Agrippino, Silvia, and Hélène Rey.** 2015. "US Monetary Policy and the Global Financial Cycle." NBER Working Paper 21722.
- Patton, Andrew, Dimitris N. Politis, and Halbert White.** 2009. "Correction to 'Automatic Block-Length Selection for the Dependent Bootstrap' by D. Politis and H. White." *Econometric Reviews* 28 (4): 372–75.
- Politis, Dimitris N., and Halbert White.** 2004. "Automatic Block-Length Selection for the Dependent Bootstrap." *Econometric Reviews* 23 (1): 53–70.
- Ross, Stephen A.** 1976. "Options and Efficiency." *Quarterly Journal of Economics* 90 (1): 75–89.
- Rossi, Barbara.** 2013. "Exchange Rate Predictability." *Journal of Economic Literature* 51 (4): 1063–119.
- Sarno, Lucio, Paul Schneider, and Christian Wagner.** 2012. "Properties of Foreign Exchange Risk Premiums." *Journal of Financial Economics* 105 (2): 279–310.
- Siegel, Jeremy J.** 1972. "Risk, Interest, and Forward Exchange." *Quarterly Journal of Economics* 86 (2): 303–09.
- Solnik, Bruno H.** 1974. "An Equilibrium Model of the International Capital Market." *Journal of Economic Theory* 8 (4): 500–24.
- Tesar, Linda L., and Ingrid M. Werner.** 1995. "Home Bias and High Turnover." *Journal of International Money and Finance* 14 (4): 467–92.
- Verdelhan, Adrien.** 2010. "A Habit-Based Explanation of the Exchange Rate Risk Premium." *Journal of Finance* 65 (1): 123–46.
- Warnock, Francis E.** 2002. "Home Bias and High Turnover Reconsidered." *Journal of International Money and Finance* 21 (6): 795–805.
- Welch, Ivo, and Amit Goyal.** 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies* 21 (4): 1455–508.