Forecasting BRL/USD with a Quanto Risk Premium

David Gun*　　Alex Luiz Ferreira†

Abstract

Can quanto a quanto risk premium predict BRL/USD? Using the framework of Kremens and Martin (2019) and daily data from February 2015 to March 2026, we construct an empirical proxy from the return differential between unhedged and hedged S&P 500 exposures in BRL and use its smoothed component to forecast exchange-rate movements. In sample, the proxy enters with negative and significant coefficients at the 3–12 month horizons, implying that higher covariance risk predicts subsequent BRL appreciation. Out of sample, re-estimated models fail to beat the random walk, while a theory-imposed forecast delivers positive out-of-sample R^2 at the 12 and 24 month horizons. These results show that a quanto risk premium can be recovered from public data and has forecasting value when theoretical restrictions are imposed.

JEL codes: F31; G12; G13; C53.

Keywords: exchange rates; out-of-sample forecasting; quanto risk premium; predictability; equity-FX covariance.

*Universidade de São Paulo, Departamento de Economia FEA-RP, Ribeirão Preto, 14040-900, Brazil. E-mail: davidgun@usp.br.

†Universidade de São Paulo, Departamento de Economia FEA-RP, Ribeirão Preto, 14040-900, Brazil. E-mail: alexferreira@usp.br.

1 Introduction

Can information embedded in financial assets help predict exchange-rate movements? This question lies at the center of the long-standing exchange-rate forecasting puzzle. Despite decades of research since Meese and Rogoff (1983), models based on observable macroeconomic fundamentals have rarely outperformed a random-walk benchmark in out-of-sample forecasts. One widely accepted interpretation is that exchange rates behave as forward-looking asset prices that incorporate expectations about future fundamentals as well as compensation for risk Engel and West (2005). If this view is correct, market prices that embed risk compensation should contain information about expected currency movements.

This paper investigates whether a market-based measure of joint equity–currency risk contains predictive information for the BRL/USD exchange rate. Our approach builds on the quanto-based asset-pricing framework of Kremens and Martin (2019), which identifies a pricing-relevant covariance between exchange-rate movements and foreign asset payoffs. This quanto covariance captures how currency movements interact with global asset returns in states that matter for pricing, and therefore reflects an *ex ante* variable rather than realized comovement. As such, it contains information about currency risk premia that is not fully captured by reduced-form measures of global risk and *ex post* variables, such as the observed VIX.

This variable can be directly related to the standard risk-adjusted uncovered interest parity (UIP) condition. In the UIP representation, expected currency returns depend on the covariance between exchange-rate movements and the stochastic discount factor Lustig and Verdelhan (2007). Economically, this term reflects whether the currency tends to depreciate in states that are valuable for the representative investor. For a BRL investor, the currency typically depreciates when global conditions deteriorate and marginal utility is high, so it delivers low payoffs in bad states and therefore commands a risk premium Ferreira and Moore (2015); Gaglianone and Marins (2016); Miranda-Agrippino and Rey (2020). The contribution of Kremens and Martin (2019) is to show that this UIP risk term can be decomposed into a risk-neutral covariance component—the quanto covariance—and a residual physical-risk term, thereby isolating the part of currency risk that is jointly priced with foreign asset payoffs.

Our contribution is twofold. First, we apply the quanto-based framework to BRL/USD and evaluate whether the equity–currency covariance channel carries predictive content. Second, we develop an empirical implementation for environments where quanto forward prices are not observable. In their absence, we construct a proxy based on the return differential between unhedged and quanto-hedged S&P 500 exposures in BRL and smooth it to approximate its expected component.

From the perspective of a BRL investor, joint equity–currency risk arises naturally in the payoff of foreign assets. The domestic-currency return of a foreign equity position depends on both the equity return and the exchange-rate movement, so the difference between hedged and unhedged exposures reflects the price of their joint variation. This mapping links the empirical proxy to the quanto covariance term in the theoretical decomposition. The Brazilian real provides a natural setting to study this mechanism: relative to advanced-economy currencies, it exhibits higher volatility, larger interest-rate differentials, and stronger sensitivity to global risk cycles, and tends to depreciate in periods of elevated global risk, making variation in equity–FX covariance economically large and more easily detectable.

Using data from February 2015 to March 2026, we obtain three main results. In sample, the proxy enters with negative and statistically significant coefficients, implying that higher covariance risk predicts subsequent BRL appreciation. Out of sample, re-estimated models fail to outperform the random walk, while a theory-imposed forecast delivers positive out-of-sample R^2 at the 12- and 24-month horizons. The proxy remains informative after controlling for interest-rate differentials, supporting a risk-based interpretation of BRL/USD dynamics in which expected currency movements reflect priced equity–currency covariance risk.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the theoretical framework and empirical design, including the formal definitions of the realized quanto risk premium ($\text{QRP}^{no,ov}$) and its smoothed proxy ($\widehat{\mathrm{QRP}}^{no,ov}$). Section 4 describes the data and variable construction. Section 5 reports the main predictive-regression results, including robustness and additional evidence, and Section 6 concludes.

2 Literature Review

This section reviews the exchange-rate literature most relevant to our empirical strategy, focusing on persistent out-of-sample forecasting difficulties, risk-premium interpretations of currency dynamics, and derivative-based evidence on priced covariance components. Overall, this literature suggests that, if exchange rates embed risk compensation, useful predictors should be linked to theoretically grounded covariance terms and evaluated against strict out-of-sample benchmarks.

The starting point is the seminal Meese–Rogoff finding that structural exchange-rate models rarely outperform a random walk in out-of-sample exercises Meese and Rogoff (1983). Later surveys and replications confirm that this result is remarkably persistent across samples, model classes, and forecast designs Rossi (2013). We therefore treat the random walk as the relevant benchmark in our empirical tests. This is especially challenging for Brazil, where evidence for BRL/USD indicates that forecasting performance varies with both horizon and forecast target Gaglianone and Marins (2016); Marins (2024).

Why an equity–FX covariance-risk term should matter for forecasting FX. UIP failures Fama (1984), crash-prone carry returns Brunnermeier et al. (2009), and the global comovement of currency premia Lustig et al. (2011, 2014) support a risk-based interpretation of exchange-rate dynamics. Intermediary-based models Gabaix and Maggiori (2015) further show how limited global risk-bearing capacity can amplify currency adjustments. Our paper connects to this literature by testing whether a traded proxy for equity–FX covariance contains forecasting information for BRL/USD beyond persistence and the interest-rate differential (IRD). Overall, the evidence is consistent with risk premia playing a central role in exchange-rate fluctuations.

A central empirical issue is measurement: many risk-premium proxies are latent or unstable. Derivatives offer forward-looking observables, but the mapping from risk-neutral prices to physical premia is typically model-dependent Carr and Wu (2007); Backus et al. (2001). This challenge is well illustrated by Garcia and Olivares (2001), who estimate the currency risk premium embedded in Brazilian dollar futures during the Real Plan as an unobserved state variable, extracted via a Kalman filter applied to backward-looking data. Their results confirm that the BRL risk premium is large, time-varying, and correlated with macroeconomic fundamentals—but the estimate is necessarily a filtered quantity, available only with a lag and conditional on a maintained state-space specification.

A recent strand of the literature addresses this limitation by extracting risk premia directly from the prices of derivatives traded today, rather than from historical realizations. Two contributions are especially relevant. The first is Kremens and Martin (2019), who identify a market-implied equity–FX term through the spread between standard and quanto forwards on the same underlying, bypassing latent-variable extraction altogether by reading the risk-neutral covariance directly from traded prices. Their quanto-implied risk premium (QRP) captures the component of expected currency appreciation attributable to the covariance between exchange-rate movements and stock-market returns under the risk-neutral measure. The second is Ornelas (2017), who extracts the volatility risk premium (VRP)—the gap between option-implied volatility and realized volatility—from currency options on 20 exchange rates, including 10 emerging-market currencies. Both approaches yield forward-looking, observable signals that do not require the estimation of latent state variables, and both demonstrate that derivative prices embed economically significant information about future currency returns.

These two derivative-based predictors capture distinct dimensions of currency risk. The QRP of Kremens and Martin (2019) measures the price of covariance between exchange rates and the equity market: it signals which currencies are risky (tending to depreciate when stocks fall) and which serve as hedges. The VRP of Ornelas (2017) measures the price of volatility itself: it reflects the compensation investors demand for bearing uncertainty about the magnitude of future exchange-rate movements. Despite targeting different risk channels, the two approaches share a common implication for the classic UIP puzzle. Kremens and Martin (2019) show that, once the QRP is included in forecasting regressions, the coefficient on the interest-rate differential moves toward the value of one predicted by theory. Ornelas (2017) reports an analogous finding: controlling for the VRP restores statistical support for UIP and yields the theoretically expected negative sign on the interest-rate differential. Both results suggest that UIP failed not because the underlying logic was wrong, but because a time-varying risk component was omitted from the specification.

The two approaches also share the rare achievement of outperforming the random-walk benchmark out of sample—the standard that Rossi (2013) describes as the most difficult to beat in the exchange-rate forecasting literature. Kremens and Martin (2019) demonstrate this for a panel of 11 developed and semi-developed currencies against UIP, PPP, and random-walk competitors. Ornelas (2017) reaches a similar conclusion using the VRP, finding that forecasting models based on it produce lower mean squared errors than the random walk for several currency groups.

A finding of Ornelas (2017) that is particularly relevant for our application is the role of a global VRP—an equally weighted average of the volatility risk premia across all currencies in the sample. This global measure exhibits stronger predictive power for emerging-market and Latin American currencies than individual-currency VRPs, consistent with the view that option markets in these economies are less liquid and that borrowing risk-premium information from a broader cross section reduces estimation noise. This result resonates with our strategy of using the S&P 500-based quanto spread proxy—a global equity signal—as a predictor for BRL/USD, rather than relying exclusively on local Brazilian derivatives.

Our main contribution is to provide an implementation bridge between these derivative-based insights and the Brazilian market. First, we apply the quanto-based framework to BRL/USD, an important emerging-market currency pair that is highly sensitive to global risk conditions. The relevance of the pair is underscored by the findings of Garcia and Olivares (2001), who show that the currency risk premium accounted for roughly half of expected depreciation during the managed-float period and remained an economically large component of the forward discount even after the transition to a free float in 1999. These magnitudes—far exceeding those documented for developed-market currencies—suggest that a framework capable of capturing time-varying risk premia should be especially informative for BRL/USD. Second, because OTC quanto-forward prices for Brazil are not observable in standard data sources, we build market-based proxies from the return differential between an unhedged BRL S&P 500 exposure and an otherwise identical quanto-hedged position. This realized spread proxy offers a noisy measure of the equity–FX covariance component embedded in traded returns. Third, motivated by standard asset-pricing theory, we construct alternative proxies that place greater weight on covariance in adverse market states. Since the observed proxy is backward-looking, overweighting periods of large equity drawdowns provides a simple empirical device to approximate the pricing-relevant covariance that forward-looking quanto prices would embed. Fourth, we explicitly compare flexible re-estimated forecasts with a theory-imposed specification. This positions the paper as an adaptation and empirical extension of Kremens and Martin (2019), rather than a restatement.

The literature leads to a clear empirical strategy. If the spread proxy captures priced equity–FX risk, it should enter predictive regressions with an economically meaningful sign and improve forecasting performance against random-walk and UIP benchmarks. If the relevant covariance is concentrated in adverse states, proxies that emphasize equity drawdowns should exhibit stronger predictive power. Finally, if structural restrictions help when parameters are unstable, a theory-imposed forecast should be comparatively robust out-of-sample. Sections 3–5 implement these implications by formalizing the theoretical mapping, constructing the realized spread proxy and its smoothed proxy, and evaluating forecast performance across horizons.

3 Framework and Empirical Design

This section develops a compact asset-pricing framework that clarifies the object targeted by our empirical spread proxy.

3.1 Theoretical framework

The central idea is that expected exchange-rate movements can be decomposed into an interest-rate differential and a covariance component that admits a natural quanto interpretation. Our presentation is closely related to the quanto framework of Kremens and Martin (2019), but it is written from the perspective of a domestic BRL investor. This setup is convenient for the empirical application developed below, since both returns and predictive variables are measured in Brazilian reais.

Let M_{t+1} denote the stochastic discount factor of the representative BRL investor. Let $R_{f,t}^{BR}$ and $R_{f,t}^{US}$ denote the one-period gross risk-free returns in Brazil and in the United States, respectively. Let S_t denote the spot exchange rate, expressed as BRL per U.S. dollar, so that an increase in S_t represents a depreciation of the real. Define the gross exchange-rate return as

$$\tilde{X}_{t+1} = \frac{S_{t+1}}{S_t}. \tag{1}$$

A BRL investor who purchases the U.S. risk-free asset at time t and converts the proceeds back into reais at time $t+1$ receives the gross BRL payoff $R_{f,t}^{US}\tilde{X}_{t+1}$. The Euler equation therefore implies

$$\mathbb{E}_t\left[M_{t+1}R_{f,t}^{US}\tilde{X}_{t+1}\right] = 1, \tag{2}$$

or, equivalently,

$$\mathbb{E}_t\left[M_{t+1}\tilde{X}_{t+1}\right] = \frac{1}{R_{f,t}^{US}}. \tag{3}$$

Using $\mathbb{E}_t[M_{t+1}] = 1/R_{f,t}^{BR}$ together with the covariance decomposition, expected exchange-rate returns can be written as

$$\mathbb{E}_t[\tilde{X}_{t+1}] = \frac{R_{f,t}^{BR}}{R_{f,t}^{US}} - R_{f,t}^{BR}\,\mathrm{Cov}_t(M_{t+1}, \tilde{X}_{t+1}). \tag{4}$$

Equation (4) is the risk-adjusted UIP condition: the expected exchange-rate return equals the interest-rate differential minus a pricing-kernel covariance. For BRL/USD, this covariance term is typically positive in the data: the real tends to depreciate in bad global states, precisely when the BRL investor's stochastic discount factor is high. Under this sign pattern, the covariance correction reduces expected depreciation relative to the raw Brazil–U.S. interest-rate differential, implying (4) is consistent with expected depreciation being below the observed interest differential. Equivalently, taking the interest-rate differential as given, cross-time variation in expected depreciation is explained by movements in the covariance term, rather than by UIP without risk adjustment.

To decompose this covariance term, let $\tilde{R}_{t+1}$ denote the gross BRL return on a foreign asset, written consistently with the BRL numeraire used throughout the analysis, and let Q denote the associated BRL risk-neutral measure. The covariance term in (4) can then be rewritten as the sum of a risk-neutral covariance component and a residual physical-risk term. This yields

$$\mathbb{E}_t[\tilde{X}_{t+1}] = \frac{R_{f,t}^{BR}}{R_{f,t}^{US}} + \frac{1}{R_{f,t}^{BR}} \operatorname{Cov}_t^Q(\tilde{X}_{t+1}, \tilde{R}_{t+1}) - \operatorname{Cov}_t(M_{t+1}\tilde{R}_{t+1}, \tilde{X}_{t+1}). \tag{5}$$

Equation (5) decomposes expected exchange-rate returns into three components. The first term is the BRL–USD interest-rate differential. The second is a risk-neutral covariance between exchange-rate returns and foreign returns, which is the natural quanto component of the model. The third is a residual term under the physical measure. Appendix A provides the derivation.

3.2 Empirical implementation

Because OTC quanto-forward prices are not available to us, the empirical analysis replaces the theoretical quanto premium with proxies constructed from observable traded returns.

To harmonize notation between Sections 3 and 3.1, we keep the same primitives and interpretation throughout: (i) S_t is BRL per USD, so $S_{t+1}/S_t > 1$ is BRL depreciation; (ii) $\tilde{X}_{t+1} = S_{t+1}/S_t$ is the gross FX return in the theoretical decomposition; (iii) $\tilde{R}^U_{t+1}$ and $\tilde{R}^{Q,no\,ov}_{t+1}$ are observed gross BRL returns on unhedged and hedged equity legs; and (iv) $QRP^{no\,ov}_{t+1}$ is the realized spread proxy built from those two returns. Therefore, $\tilde{X}_{t+1}$ maps into the FX component inside $\tilde{R}^U_{t+1}$, while the empirical spread proxy maps to the covariance-related term in equation (5).

Let F_t^{SPX} denote the S&P 500 futures index in USD and use S_t as defined above. The BRL value of the unhedged equity position is defined as

$$P_t^N = F_t^{SPX} S_t, \tag{6}$$

so that its one-period gross return is

$$\tilde{R}_{t+1}^U = \frac{P_{t+1}^N}{P_t^N} = \frac{F_{t+1}^{SPX}}{F_t^{SPX}} \frac{S_{t+1}}{S_t}. \tag{7}$$

Let Q_t^{TR} denote the observed S&P Quanto Total Return index in BRL. Its gross return is

$$\tilde{R}_{t+1}^{Q,TR} = \frac{Q_{t+1}^{TR}}{Q_t^{TR}}. \tag{8}$$

Since this is a total-return index, it embeds the domestic overnight remuneration. Let $\tilde{R}_{t+1}^{BR}$ denote the gross Brazilian cash return over the same period. We therefore write

$$\tilde{R}_{t+1}^{Q,TR} = \tilde{R}_{t+1}^{Q,no\,ov} \, \tilde{R}_{t+1}^{BR}, \tag{9}$$

which implies the ex-overnight quanto return

$$\tilde{R}_{t+1}^{Q,no\,ov} = \frac{\tilde{R}_{t+1}^{Q,TR}}{\tilde{R}_{t+1}^{BR}}. \tag{10}$$

This distinction is important because the theoretical decomposition isolates the covariance-related pricing term from the interest-rate differential. If the overnight component is left inside the observed quanto leg, the resulting spread proxy mixes covariance and carry. We therefore define two spreads:

$$QRP_{t+1}^{with\,ov} = 100 \left[\left(\tilde{R}_{t+1}^U - 1 \right) - \left(\tilde{R}_{t+1}^{Q,TR} - 1 \right) \right], \tag{11}$$

and

$$QRP_{t+1}^{no\,ov} = 100 \left[\left(\tilde{R}_{t+1}^U - 1 \right) - \left(\tilde{R}_{t+1}^{Q,no\,ov} - 1 \right) \right]. \tag{12}$$

The second object is the baseline empirical counterpart of the theoretical quanto premium. Both spreads are scaled by 100 so that they are expressed in percentage points, which facilitates economic interpretation.

At the monthly frequency, daily gross returns are compounded within month m:

$$\tilde{R}^{U}_{m} = \prod_{t\in m} \tilde{R}^{U}_{t}, \qquad \tilde{R}^{Q,TR}_{m} = \prod_{t\in m} \tilde{R}^{Q,TR}_{t}, \qquad \tilde{R}^{BR}_{m} = \prod_{t\in m} \tilde{R}^{BR}_{t}, \tag{13}$$

so that

$$\tilde{R}^{Q,no\,ov}_{m} = \frac{\tilde{R}^{Q,TR}_{m}}{\tilde{R}^{BR}_{m}}. \tag{14}$$

The monthly no-overnight spread proxy is then

$$QRP^{no\,ov}_{m} = 100\left[\left(\tilde{R}^{U}_{m} - 1\right) - \left(\tilde{R}^{Q,no\,ov}_{m} - 1\right)\right]. \tag{15}$$

Because $QRP^{no\,ov}_{m}$ is constructed from realized returns, it is noisy at monthly frequency. Our baseline low-frequency proxy is therefore the lagged 12-month moving average

$$\widehat{QRP}^{no\,ov}_{m} = \frac{1}{12}\sum_{j=1}^{12} QRP^{no\,ov}_{m-j}, \tag{16}$$

which uses only information available at time m.

Alternative proxy constructions, including state-dependent covariance measures, are explored in Appendix B. In the empirical results that follow, the baseline predictor is $\left(\widehat{\mathrm{QRP}}^{\mathrm{no,ov}}_{m}\right)$.

Log-investor benchmark and forecasting specification. To connect the theoretical decomposition to forecasting, we use the log-investor benchmark in Kremens and Martin (2019) as motivation for factor selection. Let S_t denote the BRL/USD spot exchange rate. Rather than forecasting the *log level* $s_{t+H} = \ln(S_{t+H})$, we specify the forecasting regression in *gross returns*:

$$\frac{S_{t+H}}{S_t} - 1 = \alpha + \beta\,\mathrm{IRD}_t + \gamma\,\widehat{\mathrm{QRP}}^{no\,ov}_{t} + u_{t+H}, \tag{17}$$

where IRD_t denotes the interest-rate differential (Brazil minus U.S.), and $\widehat{\text{QRP}}_t^{no\,ov}$ is the lagged moving-average proxy for the no-overnight quanto premium introduced in Section 3.2.

The use of gross returns—rather than log returns—is deliberate and follows from the theoretical structure of the model. As shown in Result 2 of Kremens and Martin (2019), the log-investor benchmark eliminates the unobservable residual covariance term through the identity $M_{t+1}R_{t+1} = 1$, which holds exactly only in terms of gross (arithmetic) returns. Replacing gross returns with their logarithmic approximation would reintroduce an approximation error of the same order of magnitude as the risk premium being estimated, as emphasized by Engel (2016) and discussed in Kremens and Martin (2019).

The returns specification also avoids the mechanical overlap between s_t as regressor and s_{t+H} as dependent variable that inflates fit statistics in the level specification, aligning directly with the theoretical object of interest: the *expected rate of depreciation* of the real.

We estimate the predictive regression for multiple horizons H, using heteroskedasticity and autocorrelation consistent (HAC) standard errors to account for the overlap induced by multi-period returns. We evaluate the null hypotheses

$$H_0 : \gamma = 0 \quad \text{and} \quad H_0 : \beta = 0,$$

as well as their joint significance. In addition to statistical significance, predictive performance is evaluated using in-sample R^2, out-of-sample R^2_{OS}, and forecast comparison tests based on the Clark–West and Diebold–Mariano statistics.

4 Data

Figure 1 presents the main macroeconomic variables used in the empirical part of this paper: the BRL/USD exchange rate, the Brazilian CDI rate, and the U.S. policy rate. As can be seen, the Brazilian real exhibits persistent depreciation during the sample period, with a clear level shift around the COVID-19 shock that leaves the currency structurally weaker thereafter. At the same time, Brazil maintains consistently higher interest rates than the United States, generating a sustained interest rate differential that widens again during the post-2021 tightening cycle. Together, these patterns underscore two defining features of the Brazilian case: exchange rate volatility and a persistent carry environment, both of which are central to the behavior of the realized spread proxy.

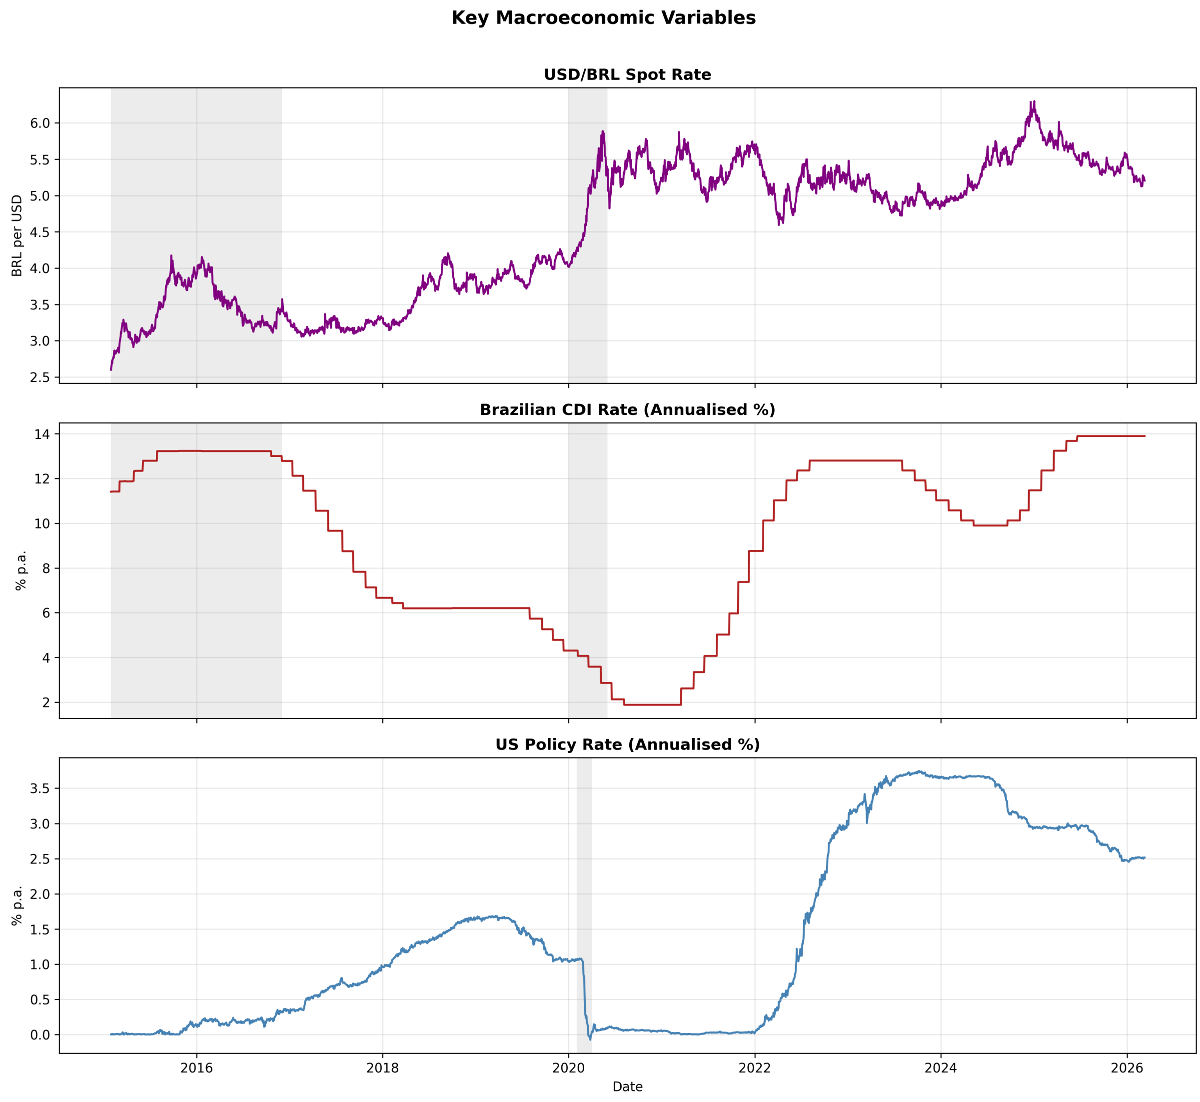

Figure 1: Key Macroeconomic Variables, 2015–2026.

Notes: Top: BRL/USD; middle: Brazilian CDI rate; bottom: U.S. policy rate. Shaded areas in the top and middle panels indicate Brazilian recession periods as dated by CODACE/FGV; in the bottom panel, shaded areas correspond to NBER-dated U.S. recessions. *Source:* Bloomberg and BCB. Own elaboration.

Table 1 reports descriptive statistics for the same series shown in Figure 1 as well as the proxies that will be discussed.

Table 1: Descriptive Statistics

Panel A: Daily Series

	R_t^Q	$\widetilde{R}_t^{BRL}$	$QRP_t^{no\,ov}$	IRD_t	S_t
Mean	0.0008	0.0009	0.0001	0.0003	4.49
Std	0.0111	0.0157	0.0111	0.0001	0.9365
IQR	0.0095	0.0165	0.0126	0.0003	1.6716

Panel B: Monthly Predictors

	ΔS_t	IRD_t	$\widehat{QRP}_t^{no\,ov}$	$\widehat{QRP}_t^{bad,6}$	$\widehat{QRP}_t^{bad,12}$
Mean	0.0051	0.0065	0.7263	0.0442	0.0631
Std	0.0452	0.0030	4.6168	0.4181	0.4832
IQR	0.0566	0.0052	5.8063	0.2905	0.3099

Notes: R_t^Q: quanto-hedged S&P 500 return; $\widetilde{R}_t^{BRL}$: unhedged S&P 500 return in BRL; $QRP_t^{no\,ov}$: realized daily no-overnight spread proxy; ΔS_t: monthly exchange-rate change; $IQR \equiv P75 - P25$. Variables with hats denote smoothed monthly proxy measures. Monthly QRP proxies are expressed in percentage points. Bad-state proxies are covariance-based measures and reported in their natural units.

Source: Bloomberg and BCB. Own elaboration.

Panel A shows that daily equity returns are volatile, with the unhedged exposure exhibiting higher dispersion than the quanto return, reflecting the additional currency component embedded in the BRL-denominated equity return.

Panel B reports the monthly predictors used in the forecasting regressions. Exchange rate changes display substantial variability over the sample, consistent with the well-known volatility of the Brazilian real. The interest rate differential is positive on average because it is defined as the Brazilian interest policy rate minus the U.S. policy rate, reflecting Brazil's structurally higher interest rates during most of the sample.

The Quanto Proxy. In this part, we explain how we obtained the spread proxy-based proxy for the quanto term using traded proxies for the unhedged BRL equity exposure and the observed S&P 500 quanto return. After removing the Brazilian overnight component embedded in the total-return quanto index, this construction yields the realized spread proxy $\mathrm{QRP}^{no\,ov}$ and its smoothed proxy $\widehat{\mathrm{QRP}}^{no\,ov}$ introduced in the previous section.

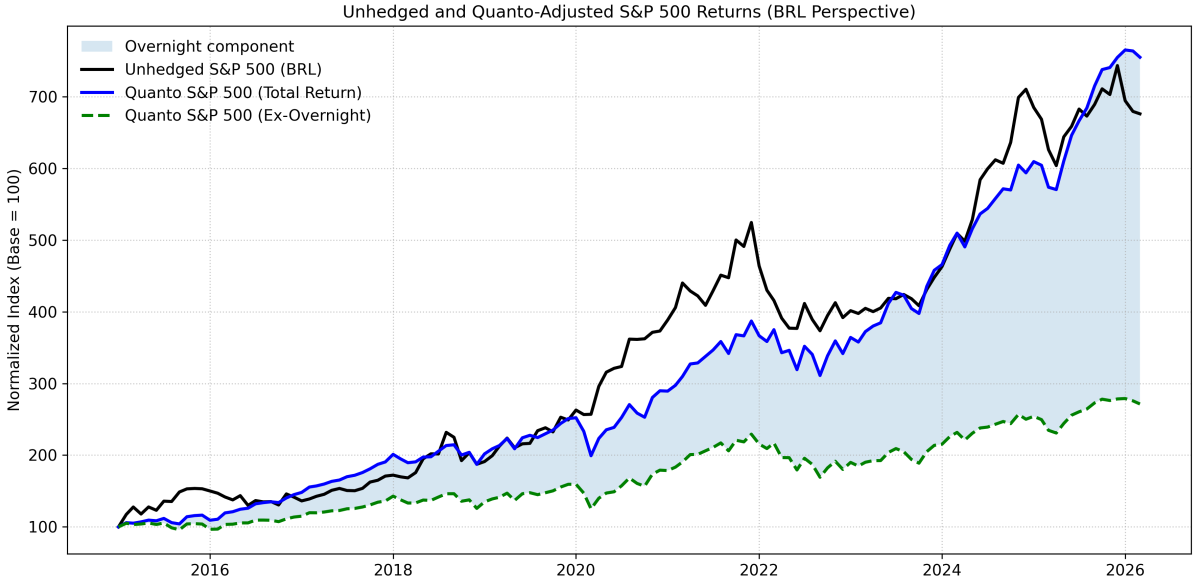

Figure 2: S&P 500 exposures from the perspective of a Brazilian investor.
Notes: Series normalized to 100 in BRL terms. *Source:* Bloomberg and BCB. Own elaboration.

Figure 2 compares the cumulative performance of three equity exposures from the perspective of a Brazilian investor. The black line reports the value of a regular S&P 500 position converted into BRL using the realized BRL/USD exchange rate. The blue line shows the S&P 500 Quanto Total Return index in BRL, which corresponds to a currency-hedged exposure that embeds the Brazilian overnight return. The dashed green line reports the corresponding quanto return after removing the overnight component, as described in Section 3.2. All series are normalized to 100 in January 2015.

The figure highlights the economic importance of the overnight component embedded in the observed quanto total-return index. The shaded area represents the cumulative contribution of this domestic cash return. Once the overnight remuneration is removed, the ex-overnight quanto series grows substantially more slowly over the sample, reflecting the fact that the domestic cash return accounts for a large share of the total performance of the quanto index.

At the same time, the regular S&P 500 exposure in BRL (black line) differs persistently from the ex-overnight quanto leg. This divergence reflects the currency component embedded in the unhedged equity return. Periods of BRL depreciation mechanically amplify the BRL value of USD-denominated equity gains, while the quanto exposure removes this exchange-rate channel.

Overall, Figure 2 illustrates two key features of the empirical construction. First, the overnight component embedded in the observed quanto total-return index is economically large and must be removed in order to isolate the pricing spread proxy implied by the theory. Second, the gap between the unhedged BRL equity return and the ex-overnight quanto leg reflects the realized counterpart of the quanto premium that motivates the proxy constructed in Section 3.2.

Dynamics of the Quanto Risk Premium. Figure 3 illustrates the monthly dynamics of the Brazilian quanto risk premium. The figure reports the realized spread proxy QRP_{no_ov} between BRL-quanto and unhedged S&P 500 exposures, together with its smoothed proxy $\widehat{QRP}_{no\_ov}$ and the interest rate differential (IRD). While the realized spread proxy is highly volatile and dominated by short-run fluctuations, the smoothed proxy isolates a more persistent component that tracks medium-run movements relevant for forecasting.

The QRP also varies systematically with macroeconomic conditions. During the domestic recession of 2015-2016, highlighted by the shaded area in Figure 3, the QRP declined, consistent with a temporary compression in perceived country risk. This period coincided with BRL appreciation and declining interest rates, culminating in a pronounced monetary easing cycle. In contrast, during the COVID-19 pandemic the BRL depreciated sharply (exceeding 4 BRL per USD), accompanied by a marked increase in the QRP.

More recently, in 2024, fiscal uncertainty surrounding the *arcabouço fiscal* and concerns about public finances were associated with a substantial BRL depreciation (above 6 BRL per USD) and a corresponding rise in both the realized and smoothed QRP. Subsequently, renewed capital inflows and BRL appreciation were accompanied by a decline in the premium.

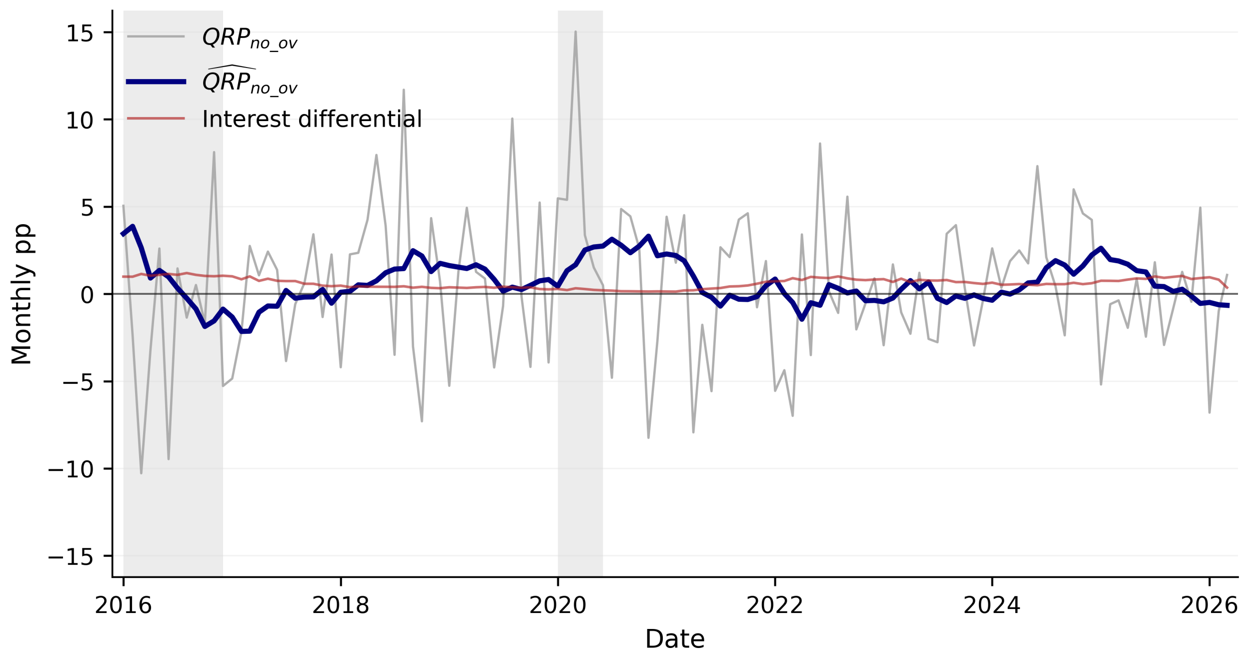

Figure 3: Monthly dynamics of the estimated quanto risk premium.

Notes: Monthly QRP_{no_ov}, $\widehat{QRP}_{no\_ov}$, and IRD (annualized percentage points). Shaded areas indicate Brazilian recession periods as dated by CODACE/FGV. *Source:* Bloomberg and BCB. Own elaboration.

Daily QRP dynamics: regimes, volatility clustering, and tail events. To characterize the realized spread proxy at the daily frequency, Figure 4 reports three complementary diagnostics based on QRP_t. Panel A plots a 60-trading-day moving average (MA60), which isolates the low-frequency component of the spread proxy. Panel B shows the corresponding 60-day rolling volatility (Vol60), highlighting time variation in uncertainty. Panel C reports a standardized measure,

$$Z_t^{60} = \frac{QRP_t - MA60_t}{STD60_t},$$

which expresses daily realizations in units of the recent conditional dispersion and facilitates comparisons of extreme movements across time.

Three patterns emerge. First, Panel A shows that the low-frequency component of QRP_t fluctuates around zero and alternates between moderately positive and negative episodes, suggesting that the spread proxy does not exhibit a persistent drift. This behavior is consistent with interpreting QRP_t as a risk-premium-like object rather than a slow-moving macroeconomic variable.

Second, Panel B displays clear volatility clustering. Periods of elevated volatility persist and are followed by calmer regimes, indicating substantial time variation in uncertainty surrounding the equity–FX relationship.

Third, Panel C provides direct evidence on tail behavior. Even after normalizing by recent volatility, the spread proxy occasionally exhibits large deviations (e.g., $|Z_t^{60}| \geq 2$), consistent with non-Gaussian tails and time-varying risk.

These features motivate the use of robust inference procedures in the predictive regressions that follow. In particular, we report heteroskedasticity and autocorrelation consistent (HAC) standard errors and complement them with bootstrap-based inference to account for potential departures from Gaussian disturbances.

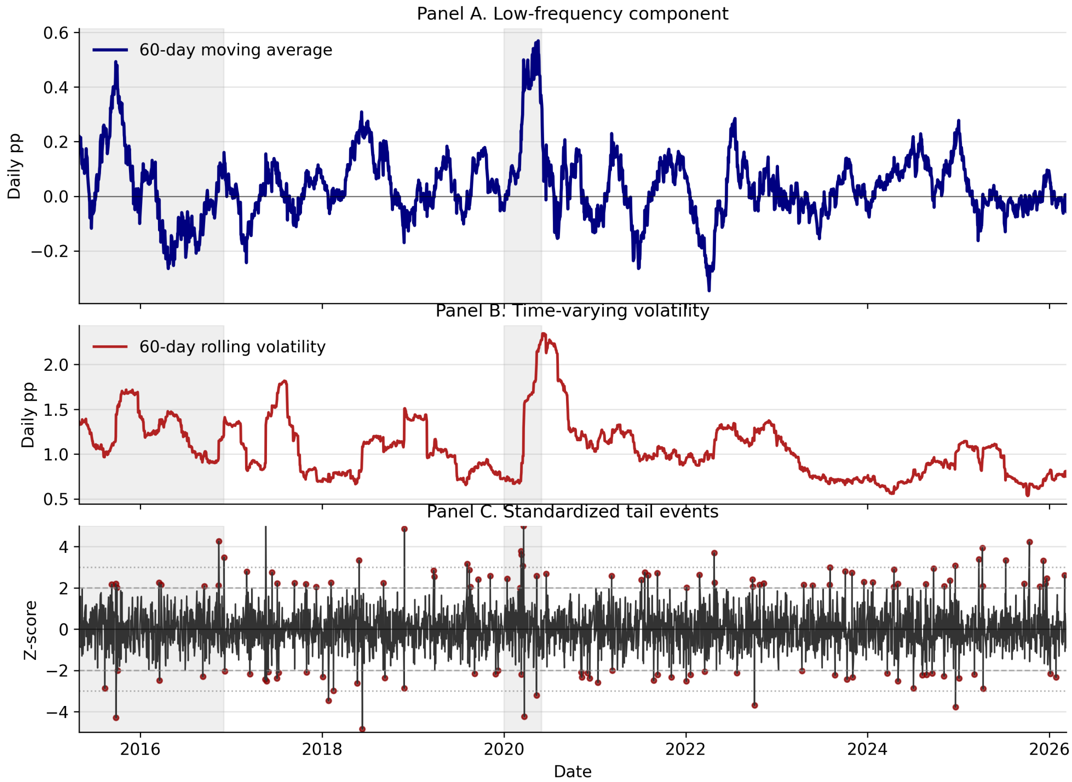

Figure 4: Daily QRP diagnostics

Notes: Panel A reports the 60-day moving average of the daily spread proxy. Panel B shows the corresponding 60-day rolling volatility. Panel C plots the standardized series Z_t^{60}, highlighting tail realizations relative to recent conditional dispersion. Shaded areas indicate Brazilian recession periods as dated by CODACE/FGV. *Source:* Bloomberg and BCB. Own elaboration.

5 Forecasting the BRL/USD rate with QRP

In this section, we investigate whether the smoothed spread proxy $\widehat{\mathrm{QRP}}_t$ contains predictive information about future BRL/USD movements. Consistent with the empirical implementation, the forecasting object is the future gross return on the BRL/USD exchange rate,

$$y_{t,H} \equiv \frac{S_{t+H}}{S_t} - 1, \tag{18}$$

where S_t denotes the BRL price of one U.S. dollar. Under this convention, $y_{t,H} > 0$ corresponds to a future depreciation of the Brazilian real against the U.S. dollar, while $y_{t,H} < 0$ corresponds to BRL appreciation. The baseline predictive regression is

$$y_{t,H} = \alpha + \beta\, \mathrm{IRD}_t + \gamma\, \widehat{\mathrm{QRP}}_t + u_{t+H}, \tag{19}$$

where IRD_t is the Brazil–U.S. interest-rate differential and $\widehat{\mathrm{QRP}}_t$ is the 12-month moving average of the no-overnight spread proxy.

Because horizons $H > 1$ imply overlapping observations, regression errors are serially correlated. We therefore report Newey–West heteroskedasticity and autocorrelation-consistent (HAC) standard errors using $H - 1$ lags. As a robustness check, we compute bootstrap p-values for the spread proxy loading using block resampling with block length tied to the forecast horizon.

Motivated by the log-investor benchmark in Kremens and Martin (2019), we also evaluate a *theory-imposed* forecast that fixes coefficients at their theoretical values. In the gross-return implementation used here, this forecast is given by

$$\widehat{y}_{t,H}^{\mathrm{THEORY}} = \mathrm{IRD}_t + \widehat{\mathrm{QRP}}_t.$$

5.1 In-sample regression results.

Tables 2 and 3 report in-sample results for $H \in \{3, 6, 12, 24\}$ months. Three findings stand out.

First, the no-overnight spread proxy enters with a negative coefficient across all horizons, but its standalone predictive content is limited. In the univariate specifications, the coefficient is generally imprecisely estimated and is not statistically significant at conventional levels, with only marginal evidence at the 12-month horizon in the BRL-depreciation regression. By contrast, once the interest-rate differential is included, the spread proxy becomes strongly significant at the 3-, 6-, and 12-month horizons in both the excess-return and exchange-rate regressions. Since the dependent variable in Table 3 is BRL depreciation, a negative coefficient implies that a higher spread proxy predicts lower subsequent depreciation, that is, a relative appreciation of the Brazilian real. Economically, this pattern is consistent with a mean-reversion interpretation: periods in which the spread proxy is elevated are followed, on average, by partial BRL recovery. At $H = 24$, however, statistical significance weakens markedly and the predictive content declines.

Second, the explanatory content of the baseline specification is economically meaningful primarily in the joint specification with the interest-rate differential. The fit is strongest at short and medium horizons, where the spread proxy and the interest-rate differential together account for a substantial share of subsequent exchange-rate variation. In the excess-return regressions, the R^2 rises from 0.264 at 3 months to 0.402 at 6 months and remains elevated at 0.356 at 12 months, before falling to 0.090 at 24 months. A similar pattern appears in the BRL-depreciation regressions, with R^2 equal to 0.202, 0.317, 0.288, and 0.065 across the four horizons. This horizon profile suggests that the baseline spread proxy captures a medium-run component of BRL/USD risk pricing that is not fully spanned by standard carry-based predictors, but whose informational content becomes too imprecise to anchor forecasting at longer horizons.

Third, alternative proxy constructions—including state-dependent covariance measures and orthogonalized variants—are examined in Appendices B–D. Those results do not alter the conclusions drawn from the baseline no-overnight proxy.

Table 2: Baseline no-overnight proxy across forecast horizons: excess returns

H	QRP only $\hat{\gamma}$	p	R^2	QRP+IRD $\hat{\gamma}$	p(QRP)	$\hat{\beta}$	p(IRD)	R^2
3	-0.610	0.410	0.0092	-2.180***	0.000	-15.168***	0.000	0.2636
6	-1.228	0.326	0.0194	-3.852***	0.000	-26.036***	0.000	0.4019
12	-2.130	0.286	0.0268	-5.956***	0.001	-37.073***	0.002	0.3560
24	-1.484	0.681	0.0069	-4.453	0.144	-25.306	0.168	0.0903

Notes: Dependent variable: future currency excess returns over horizon H. The table reports pooled regressions for the baseline no-overnight proxy across forecast horizons. "QRP only" denotes the univariate regression on $\widehat{\mathrm{QRP}}_t$, while "QRP+IRD" augments the specification with the interest-rate differential. HAC standard errors are used in the underlying regressions. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Source:* Bloomberg and BCB. Own elaboration.

Table 3: Baseline no-overnight proxy across forecast horizons: BRL depreciation

H	QRP only $\hat{\gamma}$	p	R^2	QRP+IRD $\hat{\gamma}$	p(QRP)	$\hat{\beta}$	p(IRD)	R^2
3	-0.900	0.172	0.0220	-2.162***	0.000	-12.188***	0.000	0.2015
6	-1.719	0.101	0.0451	-3.748***	0.000	-20.140***	0.000	0.3168
12	-2.820*	0.078	0.0619	-5.585***	0.000	-26.796**	0.011	0.2881
24	-2.406	0.408	0.0269	-4.053*	0.089	-14.035	0.337	0.0651

Notes: Dependent variable: BRL depreciation against the U.S. dollar over horizon H, defined so that negative values correspond to BRL appreciation. The table reports pooled regressions for the baseline no-overnight proxy across forecast horizons. "QRP only" denotes the univariate regression on $\widehat{\mathrm{QRP}}_t$, while "QRP+IRD" augments the specification with the interest-rate differential. HAC standard errors are used in the underlying regressions. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Source:* Bloomberg and BCB. Own elaboration.

Moreover, the Wald tests reinforce the horizon pattern in the baseline in-sample regressions (see Appendix C). At the 3, 6, and 12 month horizons, both the spread proxy and the interest-rate differential are individually and jointly significant in the specification, for both excess returns and BRL depreciation. At the 24-month horizon, however, the joint predictive content weakens sharply and can no longer be rejected in either specification. In excess-return regressions, neither regressor remains individually significant, while in BRL-depreciation regressions only the QRP proxy is marginally significant on its own. Overall, the statistical evidence indicates that the predictive channel is concentrated at short and medium horizons rather than at the longest horizon.

5.2 Out-of-Sample Forecast Performance

We evaluate real-time forecasting performance using an expanding-window scheme. At each forecast origin, the estimation sample is restricted to information available prior to the target observation, thereby preserving a strict real-time information set.

Our baseline benchmark is the random walk without drift. Since the forecasting target is the net BRL depreciation return, the random-walk forecast implies $E_t[y_{t,H}] = 0$. We compare this benchmark with three alternatives: an interest-rate model (UIP), a recursively re-estimated predictive regression (MAIN), and the theory-imposed forecast (THEORY).

Table 4 reports the results for the baseline no-overnight spread proxy. The recursively estimated MAIN specification fails to outperform the random walk at any horizon. Its out-of-sample R^2 is negative throughout and deteriorates sharply as the horizon lengthens, suggesting that repeated coefficient re-estimation in a short and noisy sample does not recover a stable forecasting signal from the spread proxy.

By contrast, the theory-imposed forecast is more stable. Although it remains slightly below the random walk at short horizons, it turns positive at longer horizons, reaching $R^2_{OS} = 0.4\%$ at $H = 12$ and $R^2_{OS} = 6.9\%$ at $H = 24$. These gains are economically meaningful: our R^2_{OS} of 6.9% at the two-year horizon is remarkably consistent in magnitude with the 9.6% out-of-sample R^2 reported by Kremens and Martin (2019) for a broad panel of developed-market currencies. Generating a positive out-of-sample R^2 against the driftless random walk is already a rare achievement in the exchange-rate forecasting literature Rossi (2013), and the fact that the Brazilian real—an inherently more volatile emerging-market currency—yields gains of comparable order to those found for G10 pairs provides striking external validity for the quanto theory. The evidence therefore favors the theoretical structure of the quanto forecast over unrestricted recursive slope estimation, confirming the core insight of Kremens and Martin (2019): that imposing theory-implied coefficients disciplines the forecast and circumvents the parameter estimation error that plagues conventional approaches in short, noisy samples.

Diebold–Mariano tests reported in Appendix F support a similar interpretation. They provide no evidence that the theory-imposed forecast outperforms the random walk at short horizons, but become more favorable as the horizon lengthens, with the strongest support at 24 months. Taken together, the out-of-sample evidence points to limited but non-negligible predictive gains from theory-based discipline at medium and long horizons.

Table 4: Out-of-sample forecasting performance for BRL depreciation relative to the random-walk benchmark

	$H = 3$	$H = 6$	$H = 12$	$H = 24$
N_{OOS}	60	57	51	39
MSFE (RW)	0.0037	0.0061	0.0102	0.0098
R^2_{OS} (UIP)	−0.121	−0.149	−0.600	−2.133
R^2_{OS} (MAIN)	−0.092	−0.166	−0.646	−3.294
R^2_{OS} (THEORY)	−0.026	−0.012	0.004	0.069

Notes: Expanding-window out-of-sample forecasts for net BRL depreciation returns, relative to the random-walk-without-drift benchmark. Positive R^2_{OS} values indicate lower MSFE than the benchmark, computed as $R^2_{OS} = 1 - \frac{\mathrm{MSFE}_{model}}{\mathrm{MSFE}_{RW}}$; MSFE denotes mean squared forecast error. *Source:* Bloomberg and BCB. Own elaboration.

5.3 Controlling for global risk: VIX

We augment the baseline regression with lagged $\log(\mathrm{VIX}_{t-1})$ and, in an additional specification, interact it with the spread proxy. All specifications are estimated for future gross BRL depreciation, ensuring consistency across the empirical analysis.

Figure 5 plots the standardized QRP proxy together with lagged log VIX. The two series display little systematic comovement over time, with a correlation close to zero, suggesting that the spread proxy is largely orthogonal to global volatility conditions.

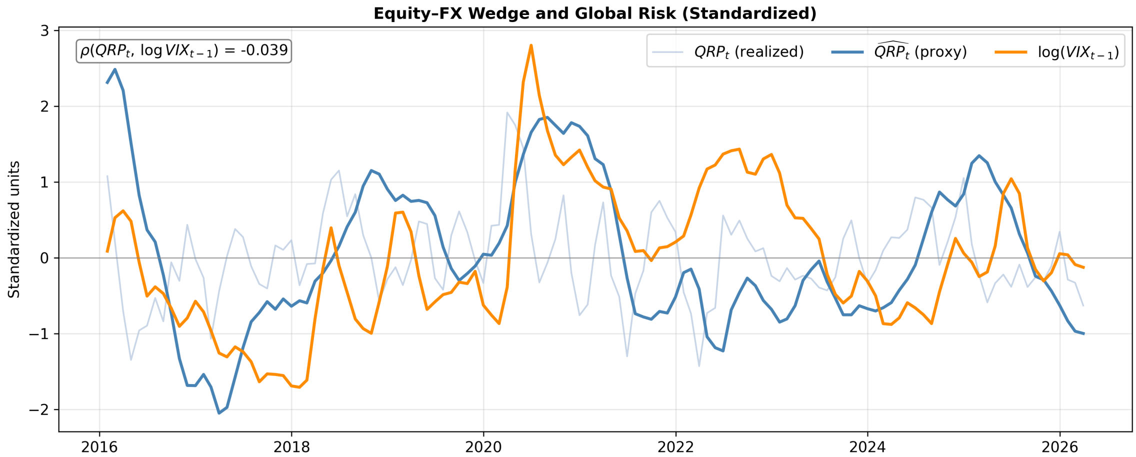

Figure 5: QRP proxy and VIX (standardized).

Notes: Standardized monthly QRP proxy and log VIX. *Source:* Bloomberg, BCB and Yahoo Finance. Own elaboration.

Table 5 reports the corresponding regressions. At short horizons ($H = 3, 6$), the spread proxy remains strongly significant after controlling for lagged VIX, while VIX itself is insignificant or only marginally significant. At longer horizons ($H = 12, 24$), the spread proxy loses statistical significance and lagged VIX becomes significant. The interaction between the spread proxy and VIX is never statistically significant.

These results are informative for interpretation. If the spread proxy were merely a relabeling of global risk, its predictive content should disappear uniformly once VIX is included. This is not supported by the data. Instead, the spread proxy retains predictive power at short and medium horizons, where the baseline signal is strongest, while global volatility conditions become more relevant at longer horizons.

The negative coefficients on both the spread proxy and VIX indicate that higher risk premia or global volatility predict lower subsequent BRL depreciation (i.e., relative appreciation), consistent with a risk-compensation mechanism in currency markets.

Table 5: **Robustness to global uncertainty: controlling for VIX**

Horizon	$\widehat{\text{QRP}}_t$	$\log(\text{VIX}_{t-1})$
3	-1.883^{***} (0.007)	-0.026 (0.297)
6	-3.075^{***} (0.002)	-0.061^{*} (0.057)
12	-3.820 (0.103)	-0.169^{***} (0.003)
24	-0.685 (0.765)	-0.271^{**} (0.044)

Notes: Slope coefficients with bootstrap p-values for the spread proxy and lagged VIX in the additive specification. An extended specification that includes a $\widehat{\text{QRP}}_t \times \log(\text{VIX}_{t-1})$ interaction term was also estimated; the interaction coefficient is statistically insignificant at all horizons (smallest bootstrap p-value $= 0.219$ at $H = 24$). Source: Bloomberg, BCB and Yahoo Finance. Own elaboration.

Additional checks. To keep the main text focused, additional exercises are reported in the appendix. Appendices B–D document alternative proxy constructions and orthogonalization diagnostics. Appendix E presents Clark–West forecast-comparison tests for the main forecast proxy, Appendix F reports Diebold–Mariano tests comparing the theory-imposed forecast against the benchmark models, and Appendix G evaluates robustness to alternative smoothing windows for the main proxy.

Taken together, these supplemental results reinforce the central message of the paper: the baseline QRP proxy is informative in sample, and the most robust out-of-sample gains arise when the quanto structure is used to discipline the forecast rather than when coefficients are repeatedly re-estimated.

6 Conclusion

This paper studied whether equity–FX risk pricing helps explain and forecast BRL depreciation against the U.S. dollar from the perspective of a Brazilian investor. Following Kremens and Martin (2019), it constructed market-based proxies from the return differential between unhedged and hedged S&P 500 exposures in BRL and evaluated their predictive content using publicly available data.

The evidence yields three main conclusions. First, the baseline no-overnight proxy is a strong in-sample predictor at the 3, 6, and 12 month horizons, but loses precision at 24 months. Second, alternative constructions that emphasize adverse equity-market states can extend the signal to longer horizons, but these results are less robust. Third, out-of-sample performance depends critically on how theoretical structure is imposed: recursively re-estimated models fail to outperform the random walk, while the theory-imposed forecast delivers gains from 6 months onward, particularly at 12 and 24 months. This pattern mirrors the main result of Kremens and Martin (2019): in short and noisy samples, unrestricted estimation leads to instability, whereas imposing theoretical restrictions improves predictive accuracy.

The proxy remains informative at short horizons even after controlling for lagged VIX, while global volatility becomes more relevant only at longer horizons, indicating that the predictive content is not driven by broad global risk conditions. Instead, it reflects a component linked to equity–currency comovement that is not fully spanned by standard global risk factors. According to the theory, this component corresponds to *ex-ante* pricing of an idiosyncratic quanto covariance risk: currencies that depreciate when global risk is high must offer compensation in equilibrium, and the proxy isolates this channel using traded indices.

Overall, the results support the quanto framework in an emerging-market setting. Even without access to OTC quanto prices or cross-currency panels, the pricing-relevant covariance component can be approximated using observable data and generates economically meaningful forecasting gains when theoretical restrictions are imposed. Future work could explore derivative-based proxies, extend the analysis to other currencies, and investigate whether the predictability concentrates in specific macro regimes such as commodity cycles or episodes of elevated domestic political uncertainty. For practitioners, the proxy seems to offer a tractable, market-based input for currency overlay strategies and for central banks monitoring market-implied currency risk premia as a complement to standard external vulnerability indicators.

Acknowledgements

The authors used artificial-intelligence tools exclusively for language editing and drafting assistance. All research questions, theoretical development, empirical analysis, and conclusions are entirely the authors' own.

References

Backus, D. K., Foresi, S., and Telmer, C. I. (2001). Affine term structure models and the forward premium anomaly. *Journal of Finance*, 56(1):279–304.

Brunnermeier, M. K., Nagel, S., and Pedersen, L. H. (2009). Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23(1):313–347.

Carr, P. and Wu, L. (2007). Stochastic skew in currency options. *Journal of Financial Economics*, 86(1):213–247.

Engel, C. (2016). Exchange rates, interest rates, and the risk premium. *American Economic Review*, 106(2):436–474.

Engel, C. and West, K. D. (2005). Exchange rates and fundamentals. *Journal of Political Economy*, 113(3):485–517.

Fama, E. F. (1984). Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3):319–338.

Ferreira, A. L. and Moore, M. J. (2015). Carry trade e risco cambial: um conto de dois fatores. *Revista Brasileira de Economia*, 69(4):429–449.

Gabaix, X. and Maggiori, M. (2015). International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420.

Gaglianone, W. P. and Marins, J. T. M. (2016). Evaluation of exchange rate point and density forecasts: an application to brazil. Technical Report 446, Banco Central do Brasil. Brasília.

Garcia, M. and Olivares, G. (2001). O prêmio de risco da taxa de câmbio no brasil durante o plano real. *Revista Brasileira de Economia (RBE)*, 55(2):151–182.

Kremens, L. and Martin, I. (2019). The quanto theory of exchange rates. *American Economic Review*, 109(3):810–843.

Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common risk factors in currency markets. *Review of Financial Studies*, 24(11):3731–3777.

Lustig, H., Roussanov, N., and Verdelhan, A. (2014). Countercyclical currency risk premia. *Journal of Financial Economics*, 111(3):527–553.

Lustig, H. and Verdelhan, A. (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1):89–117.

Marins, J. T. M. (2024). Predictability of exchange rate densities for emerging economies in the short run. Technical Report 588, Banco Central do Brasil.

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1–2):3–24.

Miranda-Agrippino, S. and Rey, H. (2020). Us monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776.

Ornelas, J. R. H. (2017). Expected currency returns and volatility risk premia. Working Paper Series 454, Banco Central do Brasil.

Rossi, B. (2013). Exchange rate predictability. *Journal of Economic Literature*, 51(4):1063–1119.

Appendix

A Derivation of the Quanto Decomposition

This appendix derives equation (5). We retain the perspective of a domestic BRL investor and use the BRL risk-neutral measure associated with the domestic numeraire.

Let $\pi(s)$ denote the physical probability of state s at time $t+1$. Define the BRL risk-neutral probability by

$$q(s) = \pi(s) M_{t+1}(s) R_{f,t}^{BR}. \tag{20}$$

Under this measure, any payoff Z_{t+1} satisfies

$$\mathbb{E}_t^Q[Z_{t+1}] = R_{f,t}^{BR} \mathbb{E}_t[M_{t+1} Z_{t+1}]. \tag{21}$$

We normalize the foreign return $\tilde{R}_{t+1}$ so that its expected gross return under the BRL risk-neutral measure equals the domestic risk-free rate,

$$\mathbb{E}_t^Q[\tilde{R}_{t+1}] = R_{f,t}^{BR}. \tag{22}$$

Using (21), this implies

$$\mathbb{E}_t[M_{t+1} \tilde{R}_{t+1}] = 1. \tag{23}$$

Next, consider the joint BRL payoff $\tilde{X}_{t+1} \tilde{R}_{t+1}$. Its Euler equation is

$$\mathbb{E}_t[M_{t+1} \tilde{X}_{t+1} \tilde{R}_{t+1}] = 1. \tag{24}$$

Applying (21) to this payoff yields

$$\mathbb{E}_t^Q[\tilde{X}_{t+1} \tilde{R}_{t+1}] = R_{f,t}^{BR}. \tag{25}$$

Expanding the left-hand side,

$$\mathbb{E}_t^Q[\tilde{X}_{t+1}]\,\mathbb{E}_t^Q[\tilde{R}_{t+1}] + \mathrm{Cov}_t^Q(\tilde{X}_{t+1}, \tilde{R}_{t+1}) = R_{f,t}^{BR}. \tag{26}$$

Substituting (22) into (26) gives

$$R_{f,t}^{BR}\mathbb{E}_t^Q[\tilde{X}_{t+1}] + \mathrm{Cov}_t^Q(\tilde{X}_{t+1}, \tilde{R}_{t+1}) = R_{f,t}^{BR}. \tag{27}$$

Now apply (21) to $\tilde{X}_{t+1}$. Using equation (3) from the main text,

$$\mathbb{E}_t^Q[\tilde{X}_{t+1}] = R_{f,t}^{BR}\mathbb{E}_t[M_{t+1}\tilde{X}_{t+1}] = \frac{R_{f,t}^{BR}}{R_{f,t}^{US}}. \tag{28}$$

Substituting (28) into (27) yields

$$1 = \frac{R_{f,t}^{BR}}{R_{f,t}^{US}} + \frac{1}{R_{f,t}^{BR}}\,\mathrm{Cov}_t^Q(\tilde{X}_{t+1}, \tilde{R}_{t+1}). \tag{29}$$

We next derive the physical-measure counterpart. Using (24), the covariance decomposition, and (23),

$$1 = \mathbb{E}_t[\tilde{X}_{t+1}]\,\mathbb{E}_t[M_{t+1}\tilde{R}_{t+1}] + \mathrm{Cov}_t(M_{t+1}\tilde{R}_{t+1}, \tilde{X}_{t+1}) = \mathbb{E}_t[\tilde{X}_{t+1}] + \mathrm{Cov}_t(M_{t+1}\tilde{R}_{t+1}, \tilde{X}_{t+1}). \tag{30}$$

Hence,

$$\mathbb{E}_t[\tilde{X}_{t+1}] = 1 - \mathrm{Cov}_t(M_{t+1}\tilde{R}_{t+1}, \tilde{X}_{t+1}). \tag{31}$$

Finally, substituting (29) into (31) gives

$$\mathbb{E}_t[\tilde{X}_{t+1}] = \frac{R_{f,t}^{BR}}{R_{f,t}^{US}} + \frac{1}{R_{f,t}^{BR}}\,\mathrm{Cov}_t^Q(\tilde{X}_{t+1}, \tilde{R}_{t+1}) - \mathrm{Cov}_t(M_{t+1}\tilde{R}_{t+1}, \tilde{X}_{t+1}), \tag{32}$$

which is equation (5) in the main text.

B Construction of the bad-state covariance proxies

This appendix provides details on the construction of the state-dependent covariance proxies introduced in Section 3.2.

Let

$$r_t^{SPX} = \frac{F_t^{SPX}}{F_{t-1}^{SPX}} - 1, \qquad r_t^{FX} = \frac{S_{t-1}}{S_t} - 1, \tag{33}$$

where $r_t^{FX} > 0$ denotes BRL appreciation. The simple realized covariance at the daily frequency is

$$rcov_t = r_t^{SPX} r_t^{FX}. \tag{34}$$

To place more weight on adverse equity states, we define a weighted covariance over a rolling window of H months:

$$badcovW_t^{(H)} = \frac{\sum_{j=0}^{L_H-1} \max\bigl(-r_{t-j}^{SPX}, 0\bigr)^p \; r_{t-j}^{SPX} r_{t-j}^{FX}}{\sum_{j=0}^{L_H-1} \max\bigl(-r_{t-j}^{SPX}, 0\bigr)^p}, \tag{35}$$

where L_H is the number of trading days in the H-month window and $p > 0$ controls how aggressively the weighting scheme emphasizes large equity drawdowns.

The monthly state variable is then obtained by sampling the rolling daily estimator at the last trading day of month m:

$$QRP_m^{badstate,H} = 10{,}000 \cdot badcovW_{t_m}^{(H)}, \qquad H \in \{3, 6, 12\}, \tag{36}$$

where t_m denotes the last trading day of month m. The scaling factor 10,000 is used only for readability.

The sign convention in this appendix is intentionally tied to BRL appreciation inside the covariance object, because the proxy is meant to capture whether Brazilian investors are exposed to joint equity–currency losses in adverse global states. In the forecasting regressions generated by the notebook, however, the target variable is future gross BRL depreciation. As a result, a negative coefficient on $QRP_m^{badstate,H}$ is economically consistent with subsequent BRL recovery.

The main text keeps the no-overnight spread proxy as the baseline empirical counterpart of the theoretical quanto premium. The bad-state covariance measures are interpreted as complementary reduced-form proxies designed to emphasize the part of equity–FX comovement that is most likely to be priced in bad states of nature.

C Wald tests for the baseline pooled specifications

Table 6 reports Wald tests for the pooled baseline horse-race specifications that include both the smoothed no-overnight quanto proxy and the interest-rate differential. The results confirm a clear horizon pattern. At the 3, 6, and 12 month horizons, both regressors are individually significant and jointly significant in the excess-return and BRL-depreciation specifications. At the 24 month horizon, however, the joint null can no longer be rejected in either specification. In the excess-return regression, neither regressor remains individually significant. In the BRL-depreciation regression, the quanto proxy remains only marginally significant on its own, while the interest-rate differential is no longer significant. Overall, the Wald evidence indicates that the predictive content of the baseline specification is concentrated at short and medium horizons and weakens substantially at the longest horizon.

Table 6: Wald tests for the pooled baseline QRP+IRD specifications

H	QRP = 0		IRD = 0		Joint = 0	
	Wald	p-value	Wald	p-value	Wald	p-value
Panel A. Excess returns (Eq. (17))						
3	20.644	0.0000***	20.767	0.0000***	30.381	0.0000***
6	15.523	0.0001***	21.168	0.0000***	23.741	0.0000***
12	11.234	0.0008***	9.958	0.0016***	12.852	0.0016***
24	2.140	0.1435	1.903	0.1677	3.357	0.1866
Panel B. BRL depreciation (Eq. (19))						
3	20.452	0.0000***	13.491	0.0002***	25.253	0.0000***
6	15.589	0.0001***	13.341	0.0003***	18.855	0.0001***
12	12.258	0.0005***	6.453	0.0111**	12.624	0.0018***
24	2.885	0.0894*	0.923	0.3367	3.601	0.1652

Notes: This table reports Wald tests for the pooled baseline regressions that include both the smoothed no-overnight quanto proxy and the interest-rate differential (IRD). Panel A corresponds to the excess-return specification in Eq. (17), and Panel B to the BRL-depreciation specification in Eq. (19). For each horizon H, the table reports tests of the null hypotheses that the coefficient on QRP is zero, that the coefficient on IRD is zero, and that both coefficients are jointly zero. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Source:* Bloomberg and BCB. Own elaboration.

D Proxy variants and orthogonalization diagnostics

This appendix compares alternative proxy constructions used in the forecasting exercise. All statistics are based on regressions for future gross BRL depreciation. Four results emerge.

First, the with-overnight spread proxy delivers slightly higher in-sample fit than the no-overnight version at short and medium horizons ($H = 3, 6, 12$). The no-overnight specification remains the preferred baseline on economic grounds.

Second, normalization weakens predictive performance. The normalized proxy consistently exhibits lower R^2 and smaller statistical significance relative to the raw spread proxy.

Third, recursive orthogonalization performs poorly at all horizons. The orthogonalized proxy is statistically insignificant and yields low explanatory power, indicating that orthogonalization removes an empirically relevant component of the signal.

Fourth, bad-state covariance proxies dominate at longer horizons. At $H = 12$ and $H = 24$, these specifications deliver the highest R^2 in the comparison block, substantially outperforming the baseline spread proxy.

This long-horizon evidence should be interpreted with caution. While HAC-based inference suggests strong statistical significance, block-bootstrap p-values are materially larger, especially at $H = 24$. This pattern indicates that the bad-state results are economically strong but less stable once small-sample dependence is taken into account. Coefficient magnitudes are not directly comparable across proxies due to differences in scaling. Table 7 reports the full set of results.

Table 7: Proxy variants comparison

H	Proxy variant	γ	p_{HAC}	p_{boot}	R^2	N
Horizon $H = 3$						
3	With overnight spread proxy (MA-12)	-2.2177	0.000	0.000	0.2131	120
3	No overnight spread proxy (MA-12)	-2.1619	0.000	0.000	0.2015	120
3	Main proxy normalized (QRP_no_ov)	-0.0669	0.010	0.011	0.1153	109
3	Bad-state covariance 12m	-0.0660	0.050	0.049	0.1478	107
3	Bad-state covariance 3m	-0.0714	0.104	0.185	0.1046	116
3	Bad-state covariance 6m	-0.0586	0.115	0.180	0.0994	113
3	Main proxy orthogonalized	-0.9191	0.842	0.846	0.0995	120
Horizon $H = 6$						
6	With overnight spread proxy (MA-12)	-3.9368	0.000	0.000	0.3456	117
6	No overnight spread proxy (MA-12)	-3.7484	0.000	0.000	0.3168	117
6	Main proxy normalized (QRP_no_ov)	-0.1356	0.002	0.002	0.2151	106
6	Bad-state covariance 12m	-0.1411	0.002	0.002	0.3221	104
6	Bad-state covariance 6m	-0.1487	0.002	0.040	0.2504	110
6	Bad-state covariance 3m	-0.1512	0.020	0.149	0.2007	113
6	Main proxy orthogonalized	-1.2490	0.883	0.864	0.1430	117
Horizon $H = 12$						
12	With overnight spread proxy (MA-12)	-5.9963	0.000	0.000	0.3300	111
12	Bad-state covariance 12m	-0.2390	0.000	0.001	0.3862	98
12	Bad-state covariance 6m	-0.3254	0.000	0.000	0.4255	104
12	Bad-state covariance 3m	-0.3627	0.000	0.003	0.3426	107
12	No overnight spread proxy (MA-12)	-5.5854	0.000	0.002	0.2881	111
12	Main proxy normalized (QRP_no_ov)	-0.2088	0.007	0.019	0.1742	100
12	Main proxy orthogonalized	4.1074	0.848	0.833	0.0977	111
Horizon $H = 24$						
24	Bad-state covariance 12m	-0.3302	0.000	0.134	0.4198	86
24	Bad-state covariance 6m	-0.3848	0.001	0.166	0.3358	92
24	Bad-state covariance 3m	-0.4575	0.005	0.175	0.2912	95
24	With overnight spread proxy (MA-12)	-4.9453	0.007	0.057	0.1001	99
24	No overnight spread proxy (MA-12)	-4.0525	0.089	0.186	0.0651	99
24	Main proxy normalized (QRP_no_ov)	0.1746	0.305	0.442	0.0265	88
24	Main proxy orthogonalized	18.1414	0.649	0.626	0.0186	99

Source: Bloomberg and BCB. Own elaboration.

E Clark–West tests (main forecast proxy)

This appendix reports Clark–West forecast-comparison tests using the main smoothed forecasting proxy adopted in the paper, namely the no-overnight spread proxy. Small one-sided p-values favor the recursively estimated MAIN model.

Table 8 shows only limited evidence in favor of MAIN. Against the random-walk benchmark, the test statistic is positive at the 3- and 6-month horizons, but the corresponding one-sided p-values remain above conventional 5% significance levels. Against UIP, the results are similarly weak, with only marginal evidence at short horizons and no support at medium or long horizons. At 12 and 24 months, the Clark–West statistics do not support the recursively estimated MAIN specification.

Overall, the Clark–West evidence does not provide robust support for the recursively estimated forecasting model when the baseline no-overnight proxy is used. This result is consistent with the broader out-of-sample message of the paper: the main forecasting gains do not come from unrestricted recursive estimation.

Table 8: Clark–West tests with the main forecast proxy

H	N	$t_{\text{vs RW}}$	$p_{\text{vs RW}}$	$t_{\text{vs UIP}}$	$p_{\text{vs UIP}}$
3	60	1.3627	0.0865	1.4272	0.0768
6	57	1.2317	0.1090	0.7481	0.2272
12	51	-0.1797	0.5713	0.5172	0.3025
24	39	-1.2755	0.8989	-2.3730	0.9912

Notes: One-sided Clark–West tests; small p-values favor MAIN. The proxy used in this block is the main smoothed no-overnight spread proxy adopted in the baseline specification. *Source:* Bloomberg and BCB. Own elaboration.

F Diebold–Mariano tests for the theory-imposed forecast

This appendix reports Diebold–Mariano comparisons between the theory-imposed forecast and the random-walk benchmark, using the baseline no-overnight spread proxy in the out-of-sample exercise. Positive statistics indicate lower MSFE for the theory-imposed forecast.

Table 9 shows a horizon-dependent pattern consistent with the main text. At short horizons, the Diebold–Mariano statistics provide no evidence that the theory-imposed forecast improves upon the random walk. The evidence becomes more favorable as the horizon lengthens, but remains weak at 12 months. At 24 months, the statistic turns more supportive, with marginal significance under the one-sided test. Overall, the formal test results are consistent with the main out-of-sample finding: theory-based discipline does not generate large forecasting gains, but it becomes more valuable at longer horizons.

Table 9: Diebold–Mariano tests: theory-imposed forecast versus random walk

H	N	DM(T vs RW)	p_1	p_2
3	60	-0.6745	0.7500	0.5000
6	57	-0.3165	0.6242	0.7517
12	51	0.0965	0.4615	0.9231
24	39	1.3643	0.0862	0.1725

Notes: Positive Diebold–Mariano statistics indicate lower MSFE for the theory-imposed forecast. p_1 denotes the one-sided p-value and p_2 the two-sided p-value. *Source:* Bloomberg and BCB. Own elaboration.

G Robustness to smoothing windows

This appendix evaluates the sensitivity of the proxy for quanto covariance risk to the choice of smoothing window. We consider moving averages of 6, 12, and 18 months applied to the no-overnight proxy.

The results show that the 12-month specification strikes the best balance across horizons. Shorter windows (MA-6) produce a noisier series and weaken the predictive relationship, particularly at short horizons. Longer windows (MA-18) smooth the series further but do not deliver systematic improvements relative to the baseline.

Overall, the evidence indicates that additional smoothing beyond 12 months yields limited gains, while shorter windows fail to adequately filter high-frequency variation. For this reason, the 12-month moving average is used as the baseline specification.

Table 10: Robustness to smoothing windows (main proxy)

H	MA window	γ	p_{boot}	R^2
Horizon $H = 3$				
3	MA(6)	-0.428	0.316	0.080
3	MA(12)	-2.162	0.000	0.202
3	MA(18)	-2.613	0.001	0.136
Horizon $H = 6$				
6	MA(6)	-0.950	0.094	0.161
6	MA(12)	-3.748	0.000	0.317
6	MA(18)	-4.511	0.001	0.229
Horizon $H = 12$				
12	MA(6)	-2.773	0.008	0.259
12	MA(12)	-5.585	0.002	0.288
12	MA(18)	-7.852	0.004	0.221
Horizon $H = 24$				
24	MA(6)	-2.829	0.179	0.112
24	MA(12)	-4.053	0.186	0.065
24	MA(18)	-5.421	0.308	0.043

Notes: Slope coefficients with bootstrap p-values for alternative moving-average windows applied to the no-overnight spread proxy. The dependent variable is future gross BRL depreciation. *Source:* Bloomberg and BCB. Own elaboration.