Brazilian Interest Rates: Term Premium, Fiscal Dominance, and the Real Curve

Marcos A. Mollica

September 2026

A Brazilian DI rate move is routinely read as a shift in the market's Selic call, which is often wrong: it can instead reflect a change in the compensation investors demand for Brazilian duration risk. This paper separates the two using sixteen years of DI futures data (2010–2026) and Adrian, Crump and Moench's (2013) no-arbitrage framework, proposing term premium composition as a marker for fiscal dominance risk versus credible policy delivery, tested on two historical episodes and applied to 2025–26. Properly controlling for US term premia means testing for cointegration, not assuming it: both curves prove genuinely cointegrated with their US counterparts, yielding a "basal" premium and a direct measure of Brazil's own share on top of it. The identical machinery, applied to Brazil's real (NTN-B) curve, recovers a real term premium and breakeven decomposition against the same control. A companion paper turns this machinery toward monetary policy, estimating Brazil's neutral rate three ways.

1 Introduction

A DI rate is a sum of two economically distinct objects: the average policy rate the market expects Copom to deliver over the life of the contract, and the compensation the market demands to bet on that path rather than hold cash. Writing $y_t^{(n)}$ for the n-period DI rate observed at time t and r_t for the short policy rate,

$$y_t^{(n)} = \underbrace{\frac{1}{n}\sum_{i=0}^{n-1} E_t[r_{t+i}]}_{\text{expected average Selic path}} + \underbrace{\phi_t^{(n)}}_{\text{term premium}}$$

The identity is definitional; the decomposition is not. Most curve-watchers read a rate move as the first term moving, because the DI rate itself does not report the split. This paper estimates it directly, on the full DI futures history available (2010–2026), and asks what that split reveals about three specific episodes in Brazilian macro-financial history and about the relationship between Brazilian and US duration risk compensation.

The contribution is sixfold. First, a statistical characterization shows the dominant level, slope, and curvature factors are stable in shape but not in concentration across four structurally distinct regimes. Second, a no-arbitrage decomposition recovers the term premium at every maturity of interest. Third, the factor composition of a given episode's term premium is proposed as a structurally motivated candidate marker for a specific and consequential question: is the market pricing fiscal dominance risk, or a credible central bank delivering an inflation-targeting mandate? These are observationally similar in the rate level alone – both can coincide with high or rising rates – but structurally distinct in which curve factor carries the premium. This paper states the diagnostic and its evidentiary scope before testing it, tests it on two historical episodes, and applies it to the current

2025–26 episode strictly on the conditional basis that scope allows. Fourth, this paper documents in full a specific internal-consistency requirement in this modeling class – that the shocks used to recover risk prices must be computed with the same corrected persistence parameter used everywhere else in the model (Appendix A.3.2) – and shows that enforcing it is what lets the diagnostic in the third contribution tell two economically distinct episodes apart at all. Fifth, the identical machinery is applied to Brazil's real (NTN-B) curve – a real term premium, a check of whether that premium reflects priced risk or technical dislocation, a joint nominal-real decomposition of the breakeven into expected inflation and an inflation risk premium, and a global control isolating the Brazil-specific share of each. Sixth, controlling for the US term premium is done properly rather than by assumption: cointegration is tested for, not assumed away, on both curves, confirmed at every maturity examined, and used to build a basal premium – the level Brazil's term premium would sit at given only current global conditions – which in turn decomposes every observed Brazilian yield into an expected-policy-path, a global, and a Brazil-specific component in one place, rather than leaving the global share implicit in a differenced regression's low R^2.

On generalizability. Everything in this paper is built and validated on Brazil because Brazil is where the underlying research programme has been developed most extensively, not because the tooling is Brazil-specific. The curve construction, the factor decomposition, the no-arbitrage term premium estimation, the diagnostic of Section 6, and the cointegration-tested approach to a global control in Sections 7 and 8.5 all depend on nothing particular to Brazilian institutions – they require only a liquid local short-rate futures curve and a published central bank policy rate, which numerous other markets in the mandate's core and secondary universe (Mexico, Colombia, Chile, Poland, the Czech Republic, among others) also offer.

The remainder of the paper is organized as follows. Section 2 situates the paper in the term structure and fiscal dominance literatures. Section 3 details data construction and estimation. Section 4 presents the factor-stability result. Section 5 presents the term premium decomposition. Section 6 develops the fiscal-dominance-versus-credibility diagnostic, states its evidentiary limits, and applies it to 2015–16, 2021–22, and live to 2025–26. Section 7 controls for the US term premium. Section 8 turns the identical machinery toward Brazil's real (NTN-B) curve, decomposing the real term premium and the breakeven, and controlling for the US real premium. Section 9 concludes. A companion paper picks up exactly where this one's machinery leaves off, turning it toward monetary policy directly.

2 Related Literature

No-arbitrage term structure decomposition. The regression-based approach used throughout this paper follows Adrian, Crump and Moench (2013), who show that risk prices in a Gaussian affine term structure model can be recovered from a sequence of ordinary least squares regressions rather than maximum likelihood, at a substantial computational advantage relative to the Kim and Wright (2005) three-factor arbitrage-free framework, which remains a standard alternative reference point, particularly in policy applications at the Federal Reserve, whose published decomposition of the US Treasury curve is used directly in Section 7. Duffee (2002) is the origin of the "essentially affine" price-of-risk specification, $\lambda_t = \lambda_0 + \lambda_1 X_t$, that this entire framework relies on. Duffee (2011) is the direct caveat on trusting it too far: almost half of the variation in bond risk premia, by his estimate, is not detectable from the cross-section of yields alone, meaning a meaningful share of what actually prices risk in this class of model may not be spanned by the PCA factors this paper extracts from the curve itself. This is not a defect specific to any one implementation; it is a standing caveat on the entire approach.

Small-sample bias in near-unit-root VARs. The persistence of the state vector governing the term structure is well known to be biased downward in finite samples when the true process is close to a unit root. Bauer, Rudebusch and Wu (2012) develop a bootstrap correction for this bias in exactly this class of model; it is applied throughout this paper, and the appendix (A.3.2) states the specific internal-consistency requirement this correction imposes on every downstream quantity computed from the corrected persistence matrix.

Fiscal dominance and term premium. Sargent and Wallace (1981) establish the theoretical mechanism by which an unsustainable fiscal path can force monetary policy to accommodate rather than lead, undermining the credibility on which conventional policy transmission relies. Blanchard (2004) applies this mechanism directly to Brazil, arguing that under sufficiently adverse fiscal-debt dynamics, a higher policy rate can perversely raise default risk and depreciate the currency rather than tighten conditions. This paper's diagnostic in Section 6 operationalizes exactly this distinction – fiscal dominance risk versus credible policy delivery – as a factor-composition question, and grounds the 2021–22 side of the comparison in Brazil's own 2021 central bank independence reform, discussed there in detail.

The expectations hypothesis and its violations. A large literature predates the affine no-arbitrage approach used here and motivates why an expectations-only reading of the curve is unreliable in the first place. Campbell and Shiller (1991) show that regressions of future short-rate changes on the current yield spread reject the expectations hypothesis across maturities and markets, with slope coefficients running the wrong sign at short horizons – the canonical empirical motivation for treating the term premium as a genuinely separate, time-varying object rather than a constant or negligible residual. Fama and Bliss (1987) reach a related conclusion from forward rates directly, showing that the forward-spot spread has predictive power for future spot rates beyond what the expectations hypothesis allows, power that is naturally interpreted as compensation for risk rather than information about the future path.

Real yields and inflation compensation. Extracting expected inflation and an inflation risk premium from a breakeven requires separating nominal yields into a real component and expected inflation rather than simply subtracting one fitted curve from another. Abrahams, Adrian, Crump and Moench (2016) develop a joint affine model of nominal and real (inflation-indexed) yields that avoids the well-documented liquidity and specification problems of a simple nominal-minus-real breakeven, anchoring the decomposition to a survey or realized inflation series rather than leaving expected inflation unidentified. Section 8.4 applies this joint framework to the Brazilian nominal (DI) and real (NTN-B) curves, anchored to realized IPCA in the absence of a usable long-run survey series.

Brazil's real term premium and TIPS liquidity. Vicente (2021) and Kubudi and Vicente (2018) apply term-structure decompositions to Brazilian NTN-B and DI data specifically, finding a real and inflation risk premium that is small, smoothly increasing with maturity, and mostly positive outside crisis episodes – the benchmark Section 8.4's joint estimation is checked against. Campbell, Sunderam and Viceira (2016) document that nominal bonds can switch from behaving as an inflation bet to behaving as a deflation hedge during acute growth scares, which shows up as a sharply negative inflation risk premium – the pattern this paper's own estimation reproduces around the 2020–21 pandemic. Andreasen, Christensen and Riddell (2021) document a persistent TIPS liquidity premium that complicates a naive real yield reading in the US context; the same joint framework used for their US decomposition is used for the US real-curve control in Section 8.5.

3 Data and Estimation

Data construction, curve interpolation, and the complete step-by-step ACM estimation procedure – including the implementation choices that materially affect the result and are not obvious from the model's equations alone – are documented in full in the Appendix, which allows independent reproduction from the underlying DI futures data alone. State variables $X_t \in \mathbb{R}^3$ are the level, slope, and curvature factors extracted by principal components analysis on standardized DI yield levels (1 to 110 months, monthly, 2010–2026), with the extracted factor scores further rescaled to unit variance component-wise. The factors follow a Gaussian VAR(1) under the physical measure, bias-corrected via Bauer-Rudebusch-Wu bootstrap. Risk prices are recovered via the three-step ACM regression procedure and used to construct the risk-neutral measure and the affine bond-pricing recursion. The term premium is the gap between the resulting model-implied yield and the expectations-hypothesis forecast of the average future short rate. The same state vector and VAR dynamics are reused without modification in a companion paper on monetary policy: the expected-path forecast that defines $\phi_t^{(n)}$ here is, by construction, also the model's own forecast of where Selic is going, and the VAR's unconditional mean is, by construction, the level the model expects the policy rate to settle at absent further shocks – both used directly there rather than requiring a separate estimation.

4 Result 1: Factor Shape Is Stable, Factor Concentration Is Not

Principal component analysis is run separately on the daily changes of the DI curve within each of four regimes: the calm 2010–13 period, the 2015–16 political and fiscal crisis, the 2021–22 post-pandemic hiking cycle, and the current 2025–26 episode.

Across all four regimes the shape of the loadings is remarkably stable: the slope factor's zero-crossing sits consistently between two-and-a-half and three years; the curvature factor's trough sits consistently between two and two-and-a-half years. What moves with the regime is not shape but concentration – the level factor's share of variance rises from 83.6% in the calm 2010–13 period to 90–94% in every subsequent stressed or fast-moving episode, monotonically across the three later regimes. Quieter markets leave more room for idiosyncratic, vertex-specific variation; stressed or fast-moving markets reprice the whole curve on a single dominant theme. This stability in shape is part of the justification for pooling sixteen years into a single estimation, and the state vector built from these factors is what every subsequent result in this paper is built on – including, in Section 6, the observation that which factor carries a given episode's term premium is informative even when the level-concentration statistic alone is not.

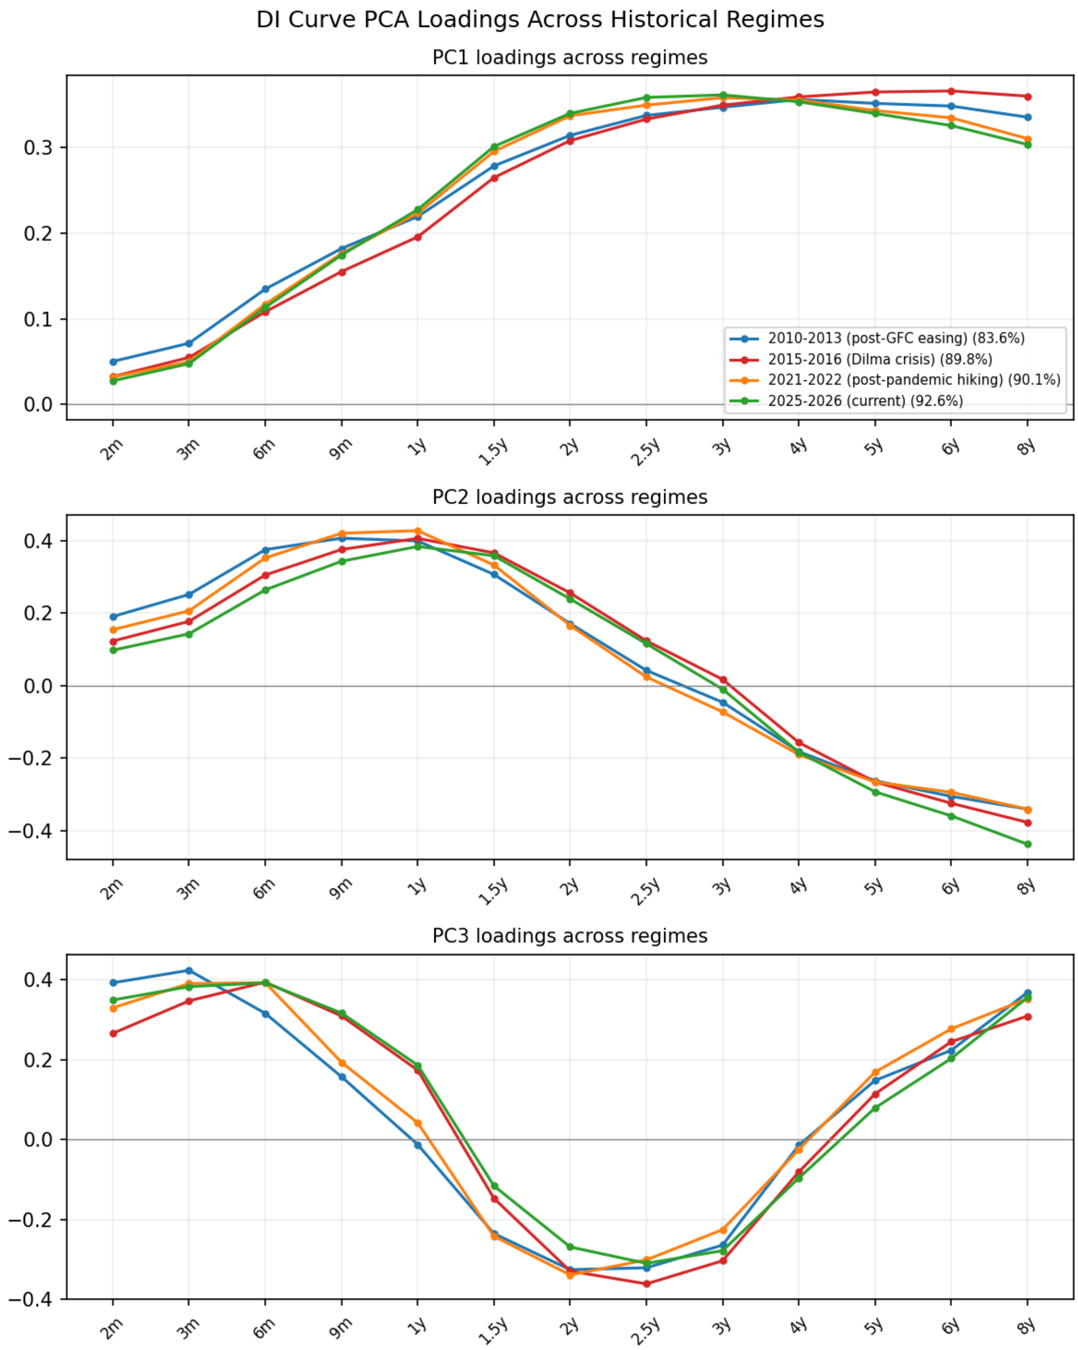

Figure 1: DI curve PCA loadings across the four regimes: level (top), slope (middle), curvature (bottom). Shape holds across regimes; concentration in the level factor rises monotonically from the calm period onward, per Table 1.

Regime	PC1 (level)	PC2 (slope)	PC3 (curvature)	Cumulative
2010–13 (calm)	83.6%	12.0%	1.9%	97.5%
2015–16 (Dilma crisis)	89.8%	7.6%	1.2%	98.6%
2021–22 (hiking cycle)	90.1%	7.1%	1.3%	98.5%
2025–26 (current)	92.6%	5.9%	1.0%	99.5%

Table 1: Variance share by regime, first three principal components.

5 Result 2: The Term Premium

The 6-, 36-, and 60-month term premium is estimated across the full sample, alongside the actual DI yield levels, with regime windows highlighted. The maturity structure itself is informative before any single episode is examined: the 6-month term premium is an order of magnitude smaller than the 36- and 60-month readings throughout the sample (sixteen-year average 0.03pp against 0.55pp and 0.82pp respectively), consistent with there being little time at short horizons for priced risk to diverge materially from the expected path – a basic sanity check the model passes by construction of the underlying no-arbitrage recursion, not by assumption.

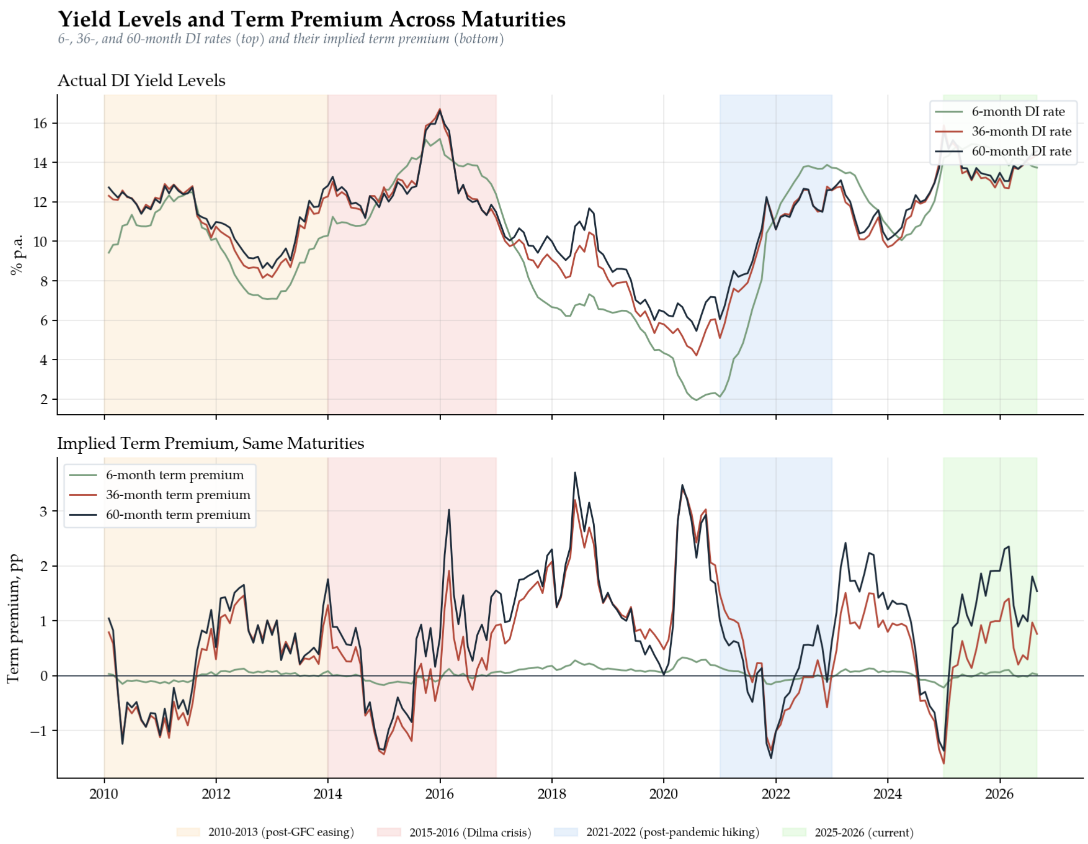

Figure 2: Actual DI yield levels (top) and implied term premium (bottom), 6-, 36-, and 60-month, across the full sample. The 6-month premium (green) is essentially flat throughout; the 36- and 60-month premia (red, black) carry the episode-level variation the rest of this section reads.

Three episodes anchor the reading. 2015–16 shows the highest term premium of the sample, peaking at $\phi^{(36)} \approx 1.9$pp, $\phi^{(60)} \approx 3.0$pp in February 2016, at the peak of the Dilma-era fiscal and political crisis. 2021–22, developed at length in Section 6, shows the term premium turning negative despite one of the fastest hiking cycles in Brazilian history. The current episode (2025–26) shows the term premium elevated and rising, reaching levels not seen since 2015–16.

A single-month reading in either episode is noisy – the estimation is a point estimate from a finite sample, and this paper reads each episode on a multi-month basis (a sustained climb or a sustained trough across several consecutive months) rather than treating any individual month's value as a standalone finding. The 2021–22 window shows eight consecutive months of negative median term premium (October 2021 through May 2022), and the 2015–16 window shows a sustained climb through the second half of 2015 into the February 2016 peak, not an isolated spike. That persistence of sign, not any single reading, is what the rest of this paper treats as evidence.

6 A Proposed Diagnostic for Regime Type: Fiscal Dominance versus Credible Tightening

6.1 The puzzle, stated plainly

Two of this paper's three headline episodes involve Selic rising. In 2015–16, hiking from roughly 11% to 14.25% coincided with the highest term premium in the sample. In 2021–22, hiking from 2% to 13.75% – a larger, faster move – coincided with the term premium turning negative. A reader could reasonably ask whether these are simply two draws from a noisy process, or whether hiking has no stable relationship with the term premium at all. This section proposes a specific, falsifiable answer to that question, states before looking at the data what evidence would support or undermine it, and only then runs the test.

6.2 The hypothesis

Aggressive hiking can raise the term premium through a well-documented channel: higher rates raise debt-service costs on domestically-held, often floating-rate public debt, and if the market doubts the fiscal path can absorb that cost, the hiking itself becomes evidence of fiscal stress rather than evidence of a central bank in control – the Sargent-Wallace (1981) / Blanchard (2004) fiscal dominance mechanism cited in Section 2. Aggressive hiking can just as easily lower the term premium for the opposite reason: if the institution doing the hiking is credible, the hiking itself resolves uncertainty rather than creating it, and a credible, front-loaded cycle can compress the premium for bearing a narrower, technical uncertainty even as it lifts realized yields. The two mechanisms are observationally similar in the rate level alone and require a different kind of evidence to tell apart.

6.3 Why level and curvature are the proposed markers

Fiscal dominance risk is a risk to the policy anchor itself – if the market doubts the central bank's ability to deliver its mandate at all, that doubt applies to every point on the curve simultaneously, with no reason to concentrate at a particular maturity. Structurally, this is a level phenomenon – the whole curve repricing together, the exact signature Result 1 already documents intensifying in every stressed episode. Credible tightening is different in kind: if the institution's ability to deliver is not in question, the only remaining uncertainty is technical – how high the terminal rate goes and how quickly the cycle turns – and that uncertainty concentrates precisely where the peak-and-pivot of the cycle gets priced, the belly of the curve, the region a level-slope-curvature decomposition assigns to the curvature factor by construction. This gives a testable, structural prediction: a fiscal-dominance episode should show up level-driven; a credible-tightening episode should show up curvature-driven. This prediction is consistent with, though not directly tested by, an existing literature on what these two factors separately capture – the level-slope-curvature decomposition originates with Litterman and Scheinkman (1991); Diebold and Li (2006) show curvature carries medium-term dynamics a level factor alone misses; Mönch (2012) finds curvature innovations specifically predictive of the curve's own future evolution in a way consistent with markets pricing a communicated policy stance; and Wright (2011), together with the emerging-market decomposition literature that follows it – including Ceballos, Naudon and Romero's (2015) application of this same ACM-style approach to Chile – repeatedly finds local term premia responding to inflation uncertainty and country-specific credibility conditions, distinct from the expected-path component. None of these findings, to this author's knowledge, combines level-versus-curvature factor composition into a diagnostic for fiscal dominance specifically; that combination is this section's own proposal.

6.4 What this test can and cannot show

Before any data are examined episode-by-episode, it is worth being precise about the evidentiary status of what follows, since the result is easy to overread once it is in hand. Sections 6.5 and 6.6 report whether the structurally motivated prediction above matches the factor composition of Brazil's two clearest historical instances of the two regime types – 2015–16 and 2021–22. That is a real, non-trivial check: the prediction was stated before the data were examined, and it could have failed. It is not, and cannot be, a validated classifier from two data points. Two episodes, both from the same country, cannot distinguish "this mechanism is real" from "this mechanism happens to fit these two cases." A skeptic's reasonable next question – what would this diagnostic say about a case where the true regime is known but is neither of these two, or about a fiscal-dominance episode outside Brazil – cannot be answered from the evidence in this paper. The honest label for the result that follows is a structurally motivated, suggested interpretation, not a robust or validated test. The only extension that would actually upgrade its evidentiary status is applying the identical decomposition to fiscal-dominance and credible-tightening episodes in other emerging markets, where the mechanism proposed in Section 6.2 is not specific to Brazilian institutions; that cross-country test is flagged in Section 9 as the paper's most important direction for future work, precisely because nothing short of it would turn this from a suggested reading into a tested one. Section 6.7 applies the diagnostic to the current episode strictly on this conditional basis, after – not before – these limits are on the record.

6.5 Test: decomposing the term premium by factor

Both the model-implied yield and the expectations-hypothesis yield are affine in the state vector X_t, so their difference – the term premium – can be written $\phi_t^{(n)} = c_0 + c_1' X_t$ for maturity-specific constants derivable directly from the affine recursion coefficients. This decomposition, applied separately to each episode, matches the prediction of Section 6.3. 2015–16 is level-driven: the first principal component is elevated well above its historical mean throughout the episode, consistent with the whole curve repricing together on a single fiscal-credibility theme. 2021–22 is curvature-driven instead: the level factor sits close to zero throughout this episode, while the curvature factor swings sharply – from roughly $+2.3$ to -1.6 in standardized units – as the curve's belly reprices during the front-loaded hiking cycle. The curvature channel's own contribution to the term premium reached -0.72pp at its most negative point in this window.

6.6 An institutional anchor for the 2021–22 reading

The curvature-driven signature of 2021–22 has a specific, dateable institutional counterpart. On 24 February 2021, Brazil enacted Complementary Law 179, granting the Central Bank formal statutory autonomy for the first time – fixed four-year terms for the president and directors, explicitly decoupled from the presidential election cycle, after thirty years of legislative debate. The 2021–22 hiking cycle began essentially at the moment this law took effect. This is close to as clean a natural experiment as macro-financial data offers for the mechanism in Section 6.2: the same action – aggressive tightening – undertaken by an institution that had just gained the specific legal protection the fiscal dominance literature identifies as the missing ingredient in 2015–16. This is corroborating context for the reading in Section 6.5, not independent statistical evidence for it; it does not change the $n = 2$ limitation stated in Section 6.4.

6.7 Applying the diagnostic live: what does 2025–26 look like, on this reading?

What follows should be read as the logical implication of the suggested interpretation developed above, applied to a case where the true regime is not yet known, subject in full to the limitations stated in Section 6.4 – not as a demonstrated finding. On the factor-composition reading alone, the current episode looks more like 2015–16 than 2021–22: the level factor has been consistently elevated and rising through mid-2026, not the near-zero readings that characterized the credible-tightening episode. Taken alongside the timing already noted elsewhere – Brazil headed into one of its closest presidential races in a generation, with an October 2026 runoff at a documented dead heat – if the diagnostic proposed in this section holds beyond the two cases it was built on, this composition would be more consistent with the market pricing systemic, whole-curve uncertainty about the policy regime than a narrower, technical uncertainty about the path to a terminal rate. That is a real conditional statement worth having on record given the timing, and it is exactly as strong as the evidence in Section 6.4 allows it to be – a single additional data point read through an interpretation calibrated on two prior cases, not a third confirmation of a validated pattern. It should be weighted accordingly until the cross-country test proposed in Section 9 exists.

7 Controlling for the US Term Premium

A natural objection to Section 6 is that some of the observed co-movement could reflect global duration risk being repriced everywhere at once, rather than anything Brazil-specific. Using the New York Fed's own published monthly ACM decomposition of the US 2-year and 10-year term premia – the identical methodology of this paper, applied by the Fed to the US curve – each Brazilian term premium is regressed on both.

A methodological point worth stating explicitly, since it changes the headline number. Both the Brazilian and US term premium series are highly persistent (monthly autocorrelation 0.91 and 0.97 respectively) – close enough to a unit root that a regression run in raw levels risks the classic Granger-Newbold spurious-regression problem: two independently trending series can show substantial, statistically "significant" co-movement in levels purely as an artifact of shared persistence, unrelated to any genuine relationship. Run in levels, the regression reports R^2 of 19.9% (36-month) and 16.7% (60-month) – a result this paper does not use on its own, for reasons made precise below.

That precision matters, because "both series look like unit roots" is not by itself a license to difference: two unit-root series can also be genuinely cointegrated, in which case a stable long-run relationship exists and differencing away the levels would discard it rather than correct for it. This is directly testable. An Engle-Granger test (Engle and Granger, 1987) – ADF (Dickey and Fuller, 1979) applied to the residual of the levels regression above – rejects a unit root in that residual decisively (36-month: $p = 0.002$; 60-month: $p = 0.0004$), and a KPSS test (Kwiatkowski, Phillips, Schmidt and Shin, 1992) on the same residual fails to reject stationarity, agreeing in the same direction.[1] The two series are cointegrated: a genuine long-run equilibrium relationship exists between Brazilian and US term premia, not merely a shared trend.

The genuinely informative specification, given cointegration, is an error-correction model rather than a first-differenced regression alone: $\Delta\phi_t^{Br} = \alpha(\phi_{t-1}^{Br} - \beta'\phi_{t-1}^{US} - c) + \gamma'\Delta\phi_t^{US} + \varepsilon_t$, with the error-correction term taken directly from the Engle-Granger first stage above. Estimated with Newey-West standard errors:

Maturity	α (speed of adjustment)	t	R^2
36-month	-0.126	-3.97	7.1%
60-month	-0.153	-4.82	8.3%

Table 2: Error-correction estimates. The short-run coefficients on $\Delta\phi_{2y}^{US}$ and $\Delta\phi_{10y}^{US}$ are not reported in this row: neither is statistically distinguishable from zero at either maturity ($|t| < 0.5$ throughout).

Two things are true at once, and neither one is the whole story on its own. The error-correction term is strongly significant at both maturities: when Brazil's term premium drifts away from its long-run relationship with the US premium, it adjusts back at a rate of 13–15% per month, confirming a genuine, stable equilibrium link between the two series over the full sixteen-year sample. At the same time, the short-run coefficients – how much of this month's US term premium *change* shows up in Brazil's premium change

[1] The two tests are used together because their null hypotheses are opposite: the ADF test's null is that the series *has* a unit root, so rejecting it is evidence of stationarity; the KPSS test's null is that the series *is* stationary, so failing to reject it is evidence of the same thing, from the other direction. A single test's conclusion could reflect that test's own assumptions or known weaknesses (ADF's low power in small samples; KPSS's sensitivity to the bandwidth choice) rather than a genuine property of the data. Agreement between two tests built to fail in different ways is meaningfully stronger evidence than either alone, and is the standard used throughout this paper wherever a cointegration or stationarity claim is made.

this same month – are indistinguishable from zero, exactly consistent with the pure first-differenced regression's own finding. The month-to-month co-movement really is close to nothing; what the differenced-only specification missed is that a real long-run tether exists underneath that near-zero short-run link. For the Section 6 diagnostic specifically, the relevant object is the short-run one: what moves Brazil's term premium episode by episode, month to month, is overwhelmingly local, exactly as a differenced-only reading concluded – the long-run cointegrating relationship is a genuine, separate finding about where Brazil's premium is anchored over long horizons, not a reason to revisit that conclusion.

The long-run relationship does have a direct, practical use of its own, distinct from anything above: a "basal" premium, the level Brazil's term premium would sit at if it were exactly on its estimated long-run relationship with the US – what the market is pricing for Brazilian duration risk once genuinely global conditions are accounted for, with no Brazil-specific factor pushing it away in either direction. This is the Engle-Granger first-stage fitted value, evaluated at the US's current level, and it answers a question market participants ask often in practice: given where US term premia sit today, where should Brazil's converge to, absent its own idiosyncratic risk? As of the most recent observation (July 2026, US 2Y/10Y term premia at 0.10pp/0.77pp), the basal 36-month Brazilian premium is 0.36pp against an actual reading of 0.97pp – a 0.61pp idiosyncratic premium priced in today, beyond what the US relationship alone implies. At 60 months the gap is wider: a 0.74pp basal level against an actual 1.81pp, a 1.06pp idiosyncratic premium. Both readings say the same thing from a different angle than Section 6's own diagnostic: not only is Brazil's premium moving on local news month to month, its current *level* is running meaningfully above what a stable relationship with US duration pricing alone would justify – consistent with, and a second independent confirmation of, the elevated-and-restrictive reading this paper's other sections reach by different routes. This basal-level calculation is inherently a snapshot, not a fixed constant: it moves whenever the US term premium does, and is reported here as of the sample's final observation rather than as a permanent benchmark.

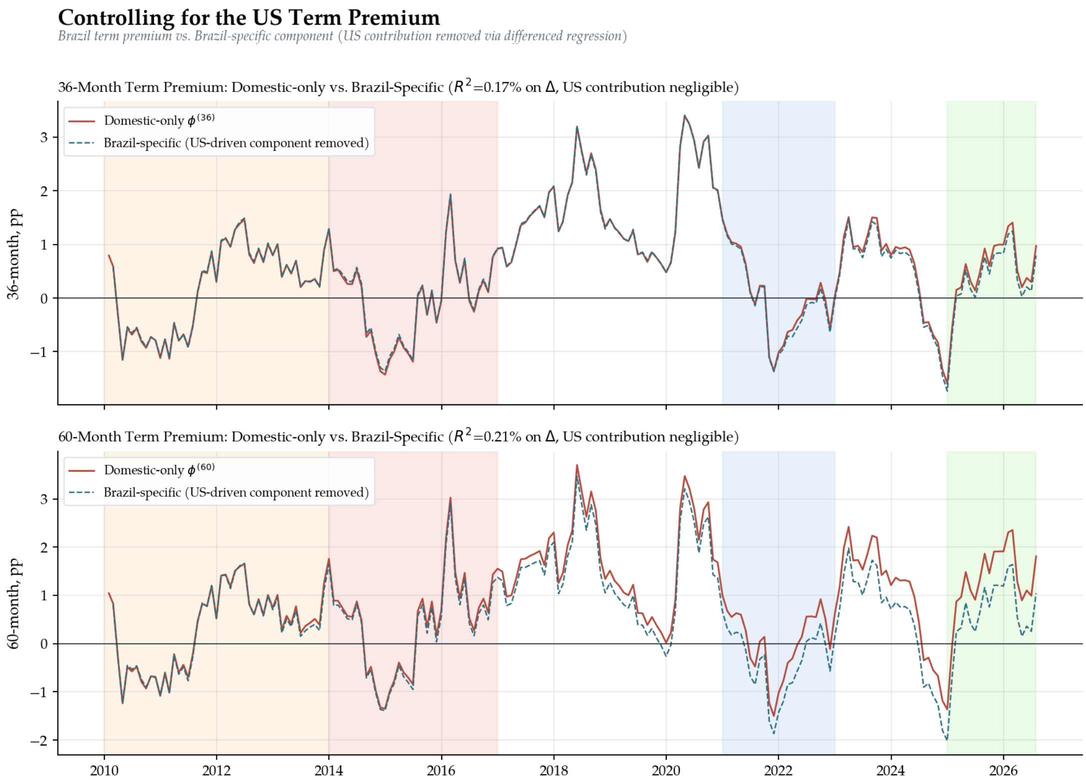

Figure 3: Brazilian term premium, domestic-only versus Brazil-specific component (US 2y/10y contribution removed via the correctly-specified differenced regression), 36- and 60-month. The two lines are visually close to indistinguishable throughout the sample – the correct visual signature of a near-zero R^2 – and a useful sanity check on the construction itself: had the two series diverged substantially despite the tiny R^2, that divergence would itself have signaled an error in how the residual was built.

8 Real Yields and Brazil's Fiscal Risk Premium

Sections 4–7 apply the affine decomposition to Brazil's nominal (DI) curve. The identical machinery – PCA state extraction, Bauer-Rudebusch-Wu-corrected VAR, ACM risk-price recovery – applies equally to Brazil's real curve, built from NTN-B inflation-linked bonds, and this section reports that companion estimation in full: a real term premium, a test of whether that premium reflects priced risk or technical market dislocation, a joint decomposition of the breakeven into expected inflation and an inflation risk premium, and a global control isolating the Brazil-specific share of the result. This is not a restatement of the narrower, self-contained real-curve neutral-rate calculation in the companion monetary-policy paper. This section is the fuller, ANBIMA-validated real curve construction; the two share a data source, not a pipeline.

8.1 Constructing the real curve

Brazil's official real-yield curve (ANBIMA's ETTJ) is refit daily but its public data feed retains only roughly the last thirteen months – not enough history for a term premium model that needs to see multiple interest-rate cycles. Individual NTN-B bond tickers each have deep history, but no single bond spans the full sample: each covers only its own issuance-to-maturity window. The construction used here instead pulls every individual NTN-B bond's own yield history – 33 bonds in total, spanning November 2005 through August 2026 – and fits a cross-sectional Nelson-Siegel-Svensson curve across whichever bonds were outstanding on each trading day. This is mechanically what ANBIMA's own methodology does; the difference is that it is reproducible here and extendable to any maturity, including ones ANBIMA's own vertex grid does not cover.

The construction fits cleanly: 5,166 of 5,166 trading days fit successfully, with a mean pricing error of 7 basis points against the actual observed bond yields. It was cross-checked against a separate, independently-built real curve (a constant-maturity 1Y/2Y/5Y/10Y series built directly from ANBIMA's official daily prices), and the two curves agree to within a median 1–4bp with no systematic bias – reassuring, given the two were built by entirely different methods.

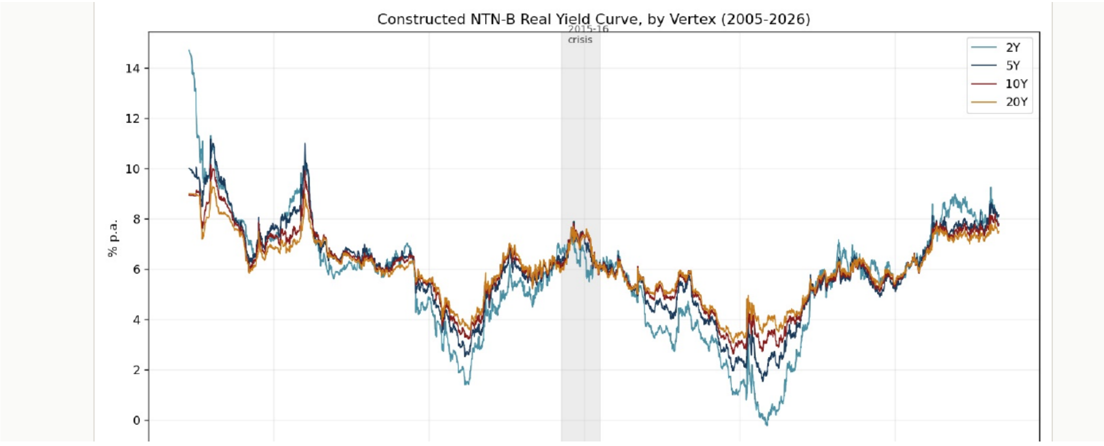

Figure 4: Real yield curve constructed from cross-sectional Nelson-Siegel-Svensson fits across all outstanding NTN-B bonds, 2005–2026. Every major Brazilian rate cycle is visible without being told where to look: the pre-2008 high-real-rate regime, the 2011–13 easing trough, the 2015–16 fiscal crisis, the pandemic collapse toward zero, and the current 2024–26 rise.

8.2 The real term premium

Applying the ACM/BRW decomposition to the constructed real curve, targeting the 2Y, 5Y, 10Y and 20Y vertices, produces a term premium series spanning the full 2005–2026 sample. The Q-measure dynamics are stable throughout (largest eigenvalue modulus 0.996, no explosion), and the model refits the actual observed curve tightly (mean error near zero, standard deviation 0.09–0.17pp across vertices).

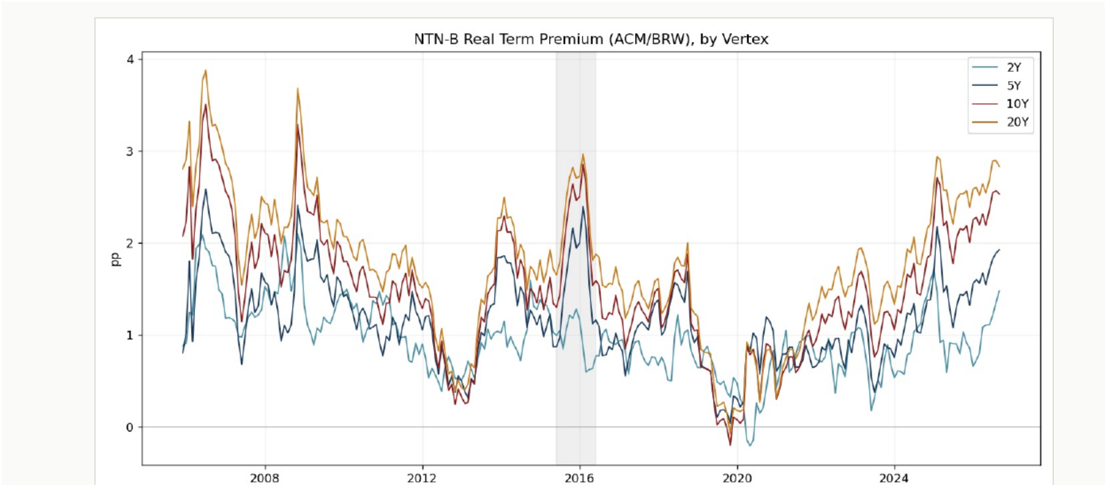

Figure 5: Real term premium, by vertex, 2005–2026. Shaded region marks the 2015–16 fiscal crisis.

The premium behaves exactly as the theory predicts: it rises with maturity (longer bonds carry more of the risk being priced), it spikes clearly during the 2015–16 crisis, compresses toward – and briefly through – zero during the pandemic's emergency easing, and has been rising steadily since 2022.

Vertex	Real yield	Expected path	Term premium
2Y	7.76%	6.50%	1.48pp
5Y	8.16%	5.88%	1.92pp
10Y	7.78%	5.25%	2.53pp
20Y	7.60%	4.77%	2.83pp

Table 3: Real yield, expected path, and term premium by vertex, as of end-August 2026.

At the 20-year point, the current term premium (2.83pp) is now approaching the 2015–16 crisis peak (2.97pp, January 2016) – the worst fiscal-risk episode in Brazil's modern rates history, at least on the raw, uncontrolled comparison. Whether that comparison survives once global term-premium co-movement is accounted for is qualified in Section 8.5.

8.3 Fiscal risk or technical dislocation?

A reasonable objection to reading an elevated real term premium as fiscal risk is that it could instead reflect a technical dislocation – buyer saturation, shrinking multimarket fund flows, regulatory changes – rather than a coherent repricing of risk. This is at least indirectly checkable with the model already in hand: if the elevated premium were mainly technical, the yield curve should get *messier* exactly when the premium spikes, since a technical dislocation shows up as harder-to-fit, noisier cross-sectional pricing rather than a smooth repricing of a common factor.

It doesn't. The cross-sectional NSS fitting error – also used as the liquidity proxy in the breakeven decomposition below – sits at essentially normal levels during both stress episodes: 0.071pp average during the 2015–16 crisis and 0.068pp during the current episode, against a 0.073pp full-sample average. Even the single worst month of each episode (January 2016, January 2025) shows fitting error well below the 90th percentile of the whole 2005–2026 sample.

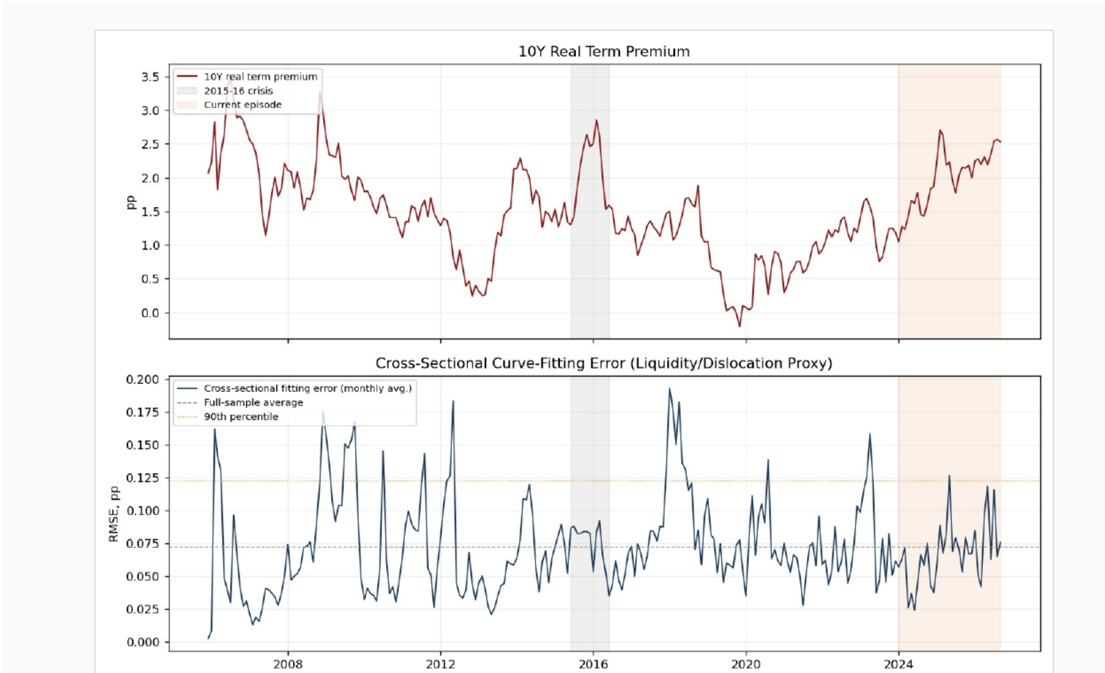

Figure 6: Top: 10Y real term premium. Bottom: cross-sectional NSS fitting error (RMSE against observed bond yields), a direct proxy for curve dislocation. The premium spikes in both shaded episodes; the fitting error does not.

This is not a clean, direct test – it doesn't use actual flow data (multimarket fund AUM, foreign NTN-B holdings, primary dealer positioning), which would be the natural way to settle the question properly and is flagged as a next step rather than attempted here. An affine term structure model also can't cleanly separate "fundamental risk" from "technical" as line items – anything priced and persistent, including a structural liquidity or balance-sheet constraint if one exists, folds into the same term premium this section reports. What the evidence above does support is narrower but still useful: whatever is driving the elevated premium, it is not showing up as the kind of cross-sectional pricing breakdown a temporary technical dislocation would produce. The curve is expensive and coherent, not dislocated.

8.4 Decomposing the breakeven

The companion theoretical case behind this empirical extension argues that the breakeven (nominal minus real yield) should widen through the same channel as the real premium: in the same fiscal-stress states that raise the real premium, expected inflation also tends to be higher, and that covariance with the discount factor should show up as its own, separate inflation risk premium. Testing this requires a nominal curve built with the same rigor as the real one, and – as an internal methodological correction below shows – requires estimating the nominal and real curves *jointly*, not separately.

The breakeven is decomposed using the joint real-nominal (RealACM) framework of Abrahams, Adrian, Crump and Moench (2016), fitting both curves *together* so that expected inflation and the inflation risk premium share a consistent factor structure rather than being backed out from two independently-estimated models, anchored to realized IPCA – the same tool used for the US decomposition in Section 8.5. This joint approach is what produces an inflation risk premium with the properties the Brazil-specific literature documents (small, smoothly increasing with maturity, mostly positive outside crisis episodes; Vicente, 2021; Kubudi & Vicente, BCB Working Paper 452): fitting each curve independently and subtracting the two afterward does not share this factor structure, and produces a premium swinging from -2.9pp to $+2.7$pp at 2Y and negative on average – properties with no comparable precedent or clean economic story.

Brazil's DI futures – the instrument underlying the nominal curve – are only meaningfully liquid out to roughly 9–10 years; there is essentially no tradable nominal curve beyond that. The real curve, built from actual NTN-B bonds with maturities out to 2060, has no such constraint. This is a genuine data limitation, not a modeling choice: the breakeven decomposition below is therefore reported at 2Y, 5Y and 10Y only. The 20Y real term premium reported above stands on its own, with no nominal or breakeven counterpart available from this data source.

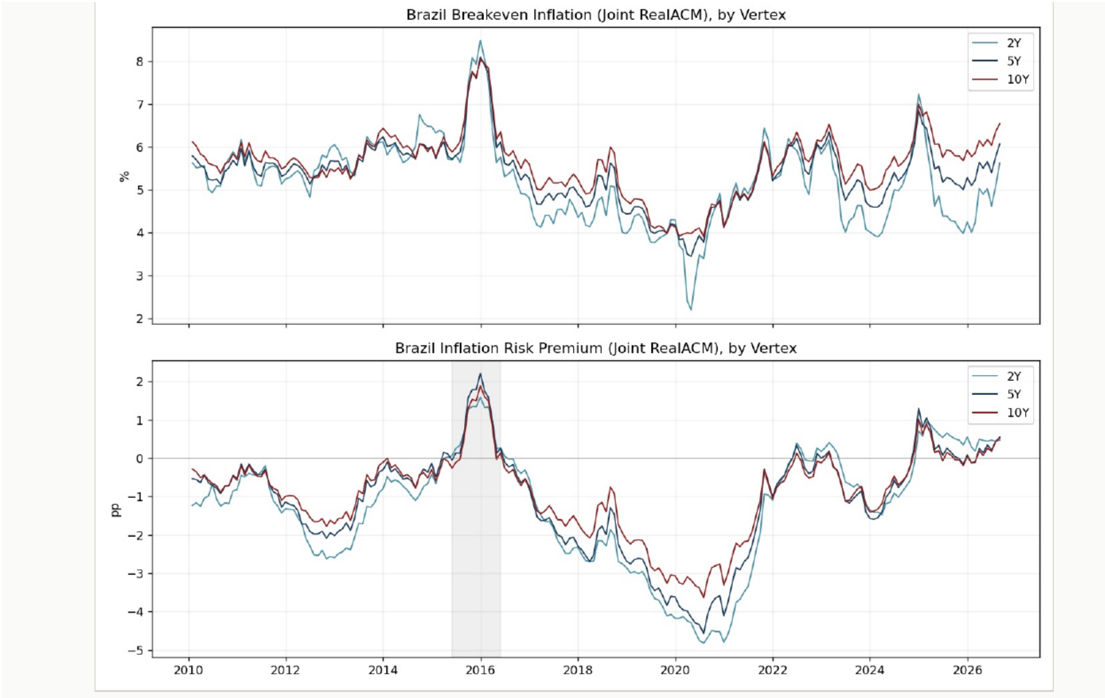

Figure 7: Top: breakeven inflation (joint RealACM). Bottom: inflation risk premium, by vertex, 2010–2026. Shaded region marks the 2015–16 fiscal crisis.

Vertex	Nominal TP	Real TP	Breakeven	Inflation risk premium
2Y	1.09pp	0.62pp	5.63%	+0.47pp
5Y	2.63pp	2.09pp	6.07%	+0.55pp
10Y	3.70pp	3.18pp	6.54%	+0.52pp

Table 4: Nominal and real term premium, breakeven inflation, and the implied inflation risk premium by vertex, end-August 2026.

With the joint model, the inflation risk premium is well-behaved and economically legible across the full 2010–2026 sample. Two episodes anchor the read. The **2015–16 fiscal crisis** shows the premium rising cleanly across all three vertices, peaking in December 2015 at +1.6pp (2Y), +2.2pp (5Y), +1.9pp (10Y) – the same crisis-sensitivity the real premium showed, now visible on the inflation side too. The **2020–21 pandemic** shows the premium falling to its deepest point in the whole sample, −4.6pp at 5Y in July 2020. This is not a contradiction; it matches a pattern already documented in the international literature (Campbell, Sunderam and Viceira, 2016): during acute deflation scares, nominal bonds stop behaving as an inflation bet and start behaving as a deflation hedge, which shows up as a sharply negative inflation risk premium. Brazil's Selic fell to 2% during this period with real fears of a demand collapse – the mechanism fits. The current reading – a small, positive premium of roughly 0.5pp across all three vertices – sits far below the 2015–16 crisis peak and close to the calm-period norm: unlike the real term premium, the inflation side of this story is not currently signaling elevated fiscal-inflation stress.

8.5 The global control: does Brazil's story survive?

Both results above are Brazil-only estimates. A fair objection is that global term premia move together, so a rise in the Brazilian real premium could partly reflect a global risk-appetite shift rather than anything Brazil-specific – the identical discipline Section 7 already applies to the nominal term premium belongs here too.

The US real term premium was estimated with the same joint Treasury-TIPS decomposition of Abrahams, Adrian, Crump and Moench (2016) – the natural real-curve counterpart to the nominal ACM model used throughout this paper – fed with generic Treasury and TIPS yields (2Y–30Y) and a VIX-based liquidity control, following Andreasen, Christensen and Riddell (2021).

A methodological correction: levels versus differences. Brazil's real premium regressed on the US real premium in levels reports an implausibly strong relationship ($R^2 = 0.55$ at 10Y). Unit root tests (ADF, KPSS) on both series told a clear story before the regression should have been trusted: the US real term premium cannot be distinguished from a unit-root process by either test, and Brazil's own series is borderline. Two persistent, independently-trending series will produce an inflated R^2 in levels even with no genuine short-run linkage – precisely the problem Section 7's own US control ran into on the nominal side.

That persistence is not, on its own, a license to simply difference and move on: an Engle-Granger test – ADF applied to the residual of the levels regression above, using the US real term premium built from Bloomberg-sourced Treasury and TIPS generic yields (2004–2026) with a VIX-based liquidity control, following Andreasen, Christensen and Riddell (2021) – rejects a unit root in that residual at both 5Y ($p = 0.016$) and 10Y ($p = 0.002$), with KPSS agreeing in the same direction at both maturities. Brazil's and the US's real term premia are cointegrated, not merely co-trending, and the result is stable under an independent robustness check: rebuilding the US series with a cruder curve-fitting-error liquidity proxy instead of VIX gives cointegration statistics and error-correction coefficients within a few thousandths of the VIX-based figures at both maturities. The genuinely informative specification, given cointegration, is an error-correction model rather than the levels regression above – estimated identically to Section 7's nominal case: $\Delta\phi_t^{Br,\text{real}} = \alpha(\phi_{t-1}^{Br,\text{real}} - \beta\phi_{t-1}^{US,\text{real}} - c) + \gamma\Delta\phi_t^{US,\text{real}} + \varepsilon_t$, with the error-correction term taken directly from the Engle-Granger first stage above. Estimated with Newey-West standard errors:

Maturity	α (speed of adjustment)	t
5Y	-0.065	-2.50
10Y	-0.139	-4.09

Table 5: Error-correction estimates, Brazil vs. US real term premium. The short-run coefficient on $\Delta\phi^{US,\text{real}}$ is not statistically significant at either maturity ($t = 1.25$ at 5Y, $t = 1.94$ at 10Y).

The pattern is the same one Section 7 documents for the nominal control: a genuine, statistically decisive long-run relationship between Brazilian and US real term premia, coexisting with short-run co-movement that is weak-to-absent. As with the nominal case, this paper uses the short-run object as the operative one going forward: the section immediately below switches to first differences precisely because the short-run coefficient here is indistinguishable from zero, exactly the result a differenced regression alone would report. That choice is not a rejection of the cointegration finding – it reflects what this paper's own diagnostics need. Section 6's use of Brazil's term premium reads episode-level, month-to-month movement, which is what the (insignificant) short-run coefficient speaks to; the long-run cointegrating relationship documented here is a genuine, separate finding about where Brazil's real premium is anchored over long horizons, reported for completeness rather than folded into anything downstream.

The same basal-premium calculation introduced for the nominal case in Section 7 applies here, with one honest difference worth flagging rather than glossing over: the estimated long-run slope on the US real term premium is small and imprecisely signed (5Y: -0.13; 10Y: $+0.03$), against a slope near or above 1 on the nominal side. The cointegrating relationship is real – the residual is stationary at both maturities, exactly as reported above – but it is closer to "Brazil's real premium reverts to a roughly stable historical anchor, largely independent of where the US real premium currently sits" than to "Brazil's real premium scales with the US's." The basal level is therefore far less sensitive to where the US real premium happens to be on any given day, and should be read as a stable long-run anchor rather than a live, US-level-tracking benchmark. With that caveat stated: as of the most recent observation, the basal 5Y Brazilian real premium is 1.01pp against an actual reading of 1.86pp (a 0.85pp idiosyncratic premium), and at the longest common maturity the basal level is 1.73pp against an actual 2.62pp (a 0.90pp idiosyncratic premium) – both in the same direction and of similar magnitude to the nominal curve's own basal-premium gap, a third independent confirmation, via yet another route, that Brazil's currently elevated term premium is substantially more than a global phenomenon passing through.

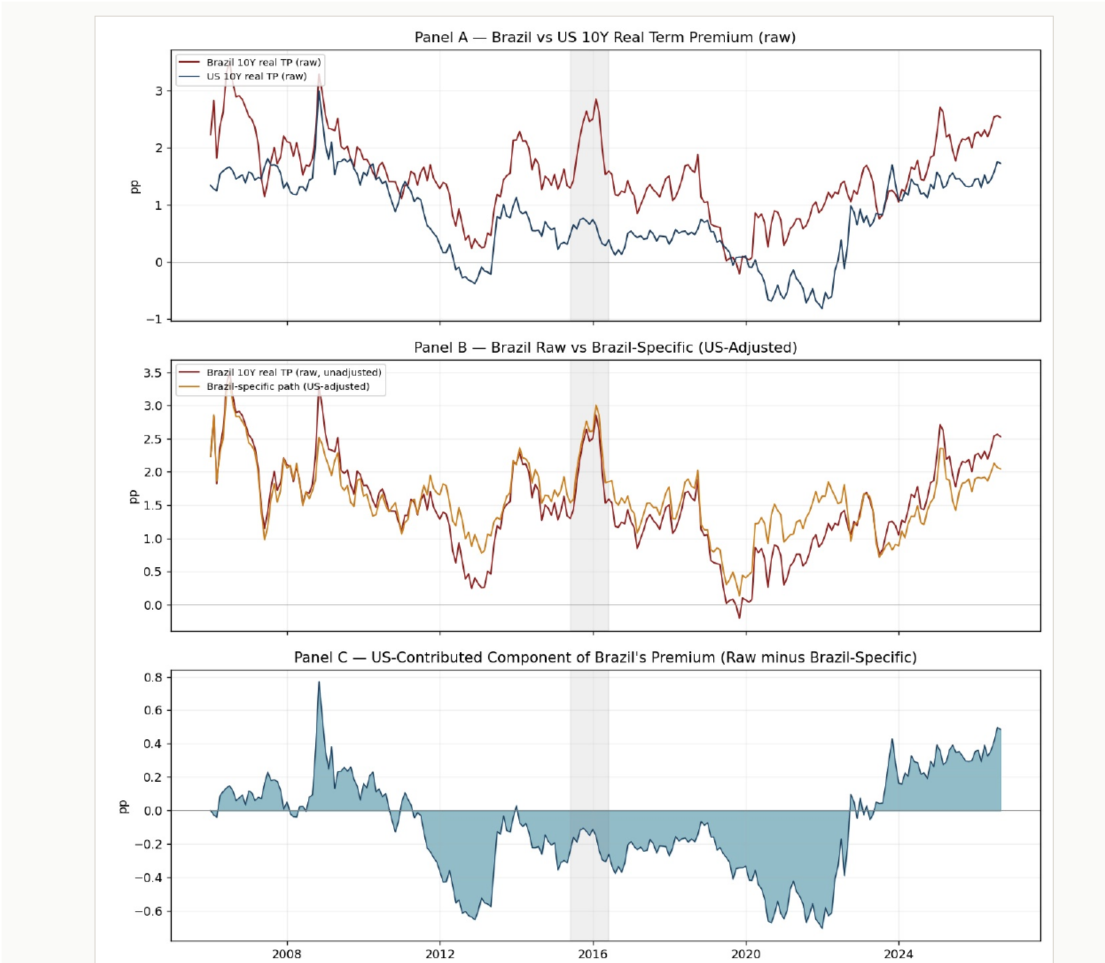

Figure 8: Panel A: Brazil vs. US 10Y real term premium, raw. Panel B: Brazil's raw premium against the Brazil-specific path (the same path with the US-explained component removed). Panel C: the difference between them – the US-contributed component of Brazil's premium, positive when the US is pushing Brazil's raw reading up, negative when it's pulling it down. Shaded region marks the 2015–16 fiscal crisis.

Vertex	β on US TP	t-stat	R^2
5Y real TP	0.24	3.0	0.03
10Y real TP	0.44	5.4	0.11

Table 6: Brazil real term premium regressed on the US real term premium, first differences.

The relationship is real – both betas remain statistically significant – but far smaller than the levels regression implied: R^2 falls from 0.55 to 0.11 at 10Y, roughly a fifth of the earlier estimate. Most of what looked like tight global co-movement was two series sharing a slow common drift over 2005–2026, not a strong contemporaneous link. Panel C makes the point precisely: during the 2015–16 crisis, the US contribution was actually slightly *negative* (averaging -0.20pp) – that spike was almost entirely Brazil's own. Today's reading is different in kind, not just degree: the US contribution is positive and rising, around $+0.49$pp as of August 2026, meaning a genuine share of the current elevated premium is global co-movement, not fiscal-risk-specific.

	Raw 10Y real TP	Brazil-specific path
2015–16 crisis peak	2.97pp	3.01pp
Current (Aug 2026)	2.53pp	2.04pp

Table 7: Raw versus Brazil-specific 10Y real term premium at the two episodes' respective peaks.

With the properly specified control, the current episode's Brazil-specific component (2.04pp) sits meaningfully below the 2015–16 peak (3.01pp) – but not nearly as far below as the flawed levels-based estimate would have suggested (0.34pp vs. 1.44pp of the gap). The honest summary: the current real term premium is elevated for both global and Brazil-specific reasons; the Brazil-specific part is real and substantial, and it is running somewhat, but not dramatically, below the worst acknowledged fiscal crisis in the modern sample.

The inflation risk premium's relationship to the US survives the same scrutiny less comfortably, and the reason is worth stating before the numbers: unlike a term premium, which compensates for duration risk that is plausibly priced through a common channel – global capital moving across borders in search of duration exposure, arbitraging away persistent cross-country gaps – an inflation risk premium compensates for uncertainty about a specific country's own inflation process. Brazilian and US inflation are driven by different central banks, different fiscal regimes, and largely different shocks; there is no comparable arbitrage mechanism that would force the two countries' IRP to share a long-run trend. The data does not contradict this prior. An Engle-Granger test, run with the same properly-sourced US IRP series used for the real term premium above, finds no cointegration at either maturity (5Y: ADF $p = 0.17$; 10Y: ADF $p = 0.24$) – a materially weaker result than the real term premium's own robust finding, and one that gets weaker rather than stronger under the more careful of the two liquidity constructions tried. With cointegration inconclusive, the differenced specification is what this section falls back on, exactly as Sections 7 and 8.5's real-premium control do when cointegration is confirmed and an error-correction model is warranted – the difference here is that no such long-run relationship is established in the first place, so differencing is not a second-best approximation to something real, it is simply the right test for a relationship that may not exist at all. That test itself does not support one: Brazil's IRP loads *negatively* on the US IRP at 5Y, weakly significant and consistent under both liquidity constructions tried ($\beta = -0.16$, $t \approx -2.1$ under either one), and the relationship is statistically indistinguishable from zero at 10Y under either one ($\beta = -0.09$ to -0.15, $t = -0.4$ to -1.1). A negative point estimate that is not robust across maturities, sits close to zero, and has no clear economic mechanism behind it in either direction is best read as no meaningful relationship, not as a weak positive or negative one – the same conclusion the cointegration test above already pointed to, from a different angle.

Putting the pieces together: a full decomposition of Brazil's current yields. Every observed Brazilian yield can now be split into three economically distinct pieces: what Copom is expected to do, what global duration pricing alone would charge for Brazilian risk, and what is specifically Brazilian on top of that.

$$y_t^{(n)} = \underbrace{EH_t^{(n)}}_{\text{expected policy path}} + \underbrace{\phi_t^{(n),\text{basal}}}_{\text{US-consistent premium}} + \underbrace{\text{idiosyncratic}_t^{(n)}}_{\text{Brazil-specific premium}}$$

	Observed yield	Expected path (EH)	Term premium: Total	Basal	Idiosyncratic	Yield if idiosyncratic premium were zero
Nominal, 36mo	14.12	13.14	0.97	0.36	0.61	13.51
Nominal, 60mo	14.35	12.54	1.81	0.74	1.06	13.29
Real, 5Y	8.00	6.14	1.85	1.01	0.85	7.15
Real, ∼9Y	7.96	5.34	2.62	1.73	0.90	7.07

Table 8: Decomposing Brazil's current yields into what Copom is expected to do, what global pricing alone would charge, and what is Brazil-specific (% p.a., 31 July 2026). **Basal** is the term premium implied by the estimated long-run relationship between Brazil's and the US's term premium (Tables 2 and 5 above), evaluated at the US's current level – the premium Brazil would carry if it moved only with global duration pricing, with no country-specific risk priced in. **Idiosyncratic** is the residual: observed premium minus basal, specifically Brazilian by construction, and the two always sum to the observed total premium. The final column, EH plus basal, is the counterfactual yield if that idiosyncratic premium went to zero – not a forecast, but where the curve would sit today with Copom's expected path and global pricing unchanged and none of Brazil's own risk priced in.

Two things stand out reading across rather than down. First, the idiosyncratic premium is not concentrated at one end of either curve – it runs from roughly 0.6pp at the nominal short end to just over 1pp at the nominal belly, and sits in a narrower 0.85–0.90pp band on the real curve, whose weak long-run sensitivity to the US level (documented above) makes its own basal figure the less dynamically informative of the two. Second, the gap between the observed and fully-corrected yield is not small relative to the moves this market actually trades: at the nominal 60-month point alone it is worth just over a full percentage point, comparable in size to a genuine Copom hiking or cutting cycle rather than a rounding-level curiosity. The size of the current alarm is closer to the raw numbers than an overcorrected control would suggest – but the correction itself is the more important result. Brazil carries a real, substantial, mostly idiosyncratic risk premium at both 5Y and 10Y, currently running somewhat below – not dramatically below – its 2015–16 crisis level. The inflation side of the story, once properly estimated with a joint nominal-real model, is coherent on its own terms (small and positive today, spiking with genuine Brazil-specific stress, collapsing during global deflation scares); its relationship to the US premium, unlike the real term premium's, does not survive scrutiny as a meaningful one in either direction, consistent with inflation risk being a largely country-specific object with no clear cross-border arbitrage channel to tether it to. The 20Y point remains real-only by construction, not by omission: neither Brazil's DI/NTN-F curve nor the US TIPS/Treasury pair used here has genuine nominal depth that far out, so no breakeven or global control is buildable there from data that exists. What remains open is sourcing a genuine liquidity proxy for the Brazilian nominal leg, to replace the current curve-fitting-error stand-in.

9 Conclusion and Further Work

This paper separates Brazilian DI yields into an expected-policy-path component and a term premium, using a no-arbitrage regression framework applied uniformly across sixteen years and four structurally distinct regimes. It proposes that the term premium's factor composition – level-driven or curvature-driven – is a structurally motivated candidate marker for whether the market is pricing fiscal dominance risk or a credible, institutionally independent central bank, states the evidentiary limits of that proposal before testing it, and tests it on two historical episodes with an unusually clean institutional anchor (Brazil's 2021 central bank autonomy law). The paper is explicit that two episodes support a suggested interpretation, not a validated test, and applies the same conditional reading to the current episode, where the composition is level-driven and therefore closer in kind to 2015–16 than to 2021–22. This distinction depends on a specific internal-consistency requirement in the estimation, documented in enough detail (Appendix A.3.2) to be checked in other applications of this modeling class. Controlling for the US term premium properly means testing for cointegration rather than assuming a shared trend must be differenced away, and doing so changes the finding rather than just qualifying it: both the nominal and real Brazil-US relationships are genuinely cointegrated, month-to-month co-movement is close to nothing on both curves, and the resulting basal premium – where Brazil's term premium would sit given only current global conditions – puts a number on how much of the current elevated reading is specifically Brazilian rather than passing through from abroad, across every maturity and curve examined. The identical machinery is then applied to Brazil's real (NTN-B) curve, recovering a real term premium and, via a joint nominal-real decomposition, an inflation risk premium – the real term premium checked against, and surviving, the same cointegration-tested global control, while the inflation risk premium's own relationship to the US does not survive scrutiny as a meaningful one in either direction, consistent with inflation risk being a substantially more country-specific object than duration risk. A companion paper reuses this same state vector and VAR dynamics, without re-estimation, to turn the decomposition toward monetary policy directly.

Three extensions are natural next steps. First and most important, the fiscal-dominance-versus-credibility diagnostic of Section 6 is, as stated there explicitly, a suggested interpretation calibrated on two within-country episodes, not a validated test; the only extension that would actually upgrade its evidentiary status is applying the identical decomposition to fiscal-dominance and credible-tightening episodes in other emerging markets, where the underlying mechanism is not specific to Brazilian institutions. Second, a walk-forward, out-of-sample evaluation of the rates-side relative-value application (outlined at the level of the general framework only, without disclosing specific signal construction) is the natural empirical complement to the in-sample results presented throughout. Third, Section 8.3's fiscal-risk-versus-technical-dislocation check uses cross-sectional curve-fitting error as an indirect liquidity proxy; the natural next step, flagged there explicitly, is testing it directly against flow data – multimarket fund AUM, foreign NTN-B ownership, primary dealer positioning – once that data is sourced.

References

1. Abrahams, M., Adrian, T., Crump, R. K., & Moench, E. (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics*, 84, 182–200.

2. Adrian, T., Crump, R. K., & Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1), 110–138.

3. Andreasen, M. M., Christensen, J. H. E., & Riddell, S. (2021). The TIPS liquidity premium. *Review of Finance*, 25(6), 1639–1675.

4. Bauer, M. D., Rudebusch, G. D., & Wu, J. C. (2012). Correcting estimation bias in dynamic term structure models. *Journal of Business & Economic Statistics*, 30(3), 454–467.

5. Blanchard, O. (2004). Fiscal dominance and inflation targeting: Lessons from Brazil. *NBER Working Paper No. 10389*.

6. Campbell, J. Y., Sunderam, A., & Viceira, L. M. (2016). Inflation bets or deflation hedges? The changing risks of nominal bonds. *Critical Finance Review*, 6(2), 263–301.

7. Campbell, J. Y., & Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, 58(3), 495–514.

8. Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366a), 427–431.

9. Ceballos, L., Naudon, A., & Romero, D. (2015). Nominal term structure and term premia: Evidence from Chile. Working paper, cited via Banco Central de Chile / IMF "Challenges for Central Banking" volume; exact original venue to be confirmed before submission.

10. Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2), 337–364.

11. Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance*, 57(1), 405–443.

12. Duffee, G. R. (2011). Information in (and not in) the term structure.

Review of Financial Studies, 24(9), 2895–2934.

13. Engle, R. F., & Granger, C. W. J. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55(2), 251–276.

14. Fama, E. F., & Bliss, R. R. (1987). The information in long-maturity forward rates. *American Economic Review*, 77(4), 680–692.

15. Kim, D. H., & Wright, J. H. (2005). An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. *Federal Reserve Board Finance and Economics Discussion Series* 2005-33.

16. Kubudi, D., & Vicente, J. (2018). Inflation risk premium in Brazil: A term structure approach. *Banco Central do Brasil Working Paper Series* No. 452.

17. Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3), 159–178.

18. Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of Fixed Income*, 1(1), 54–61.

19. Mönch, E. (2012). Term structure surprises: The predictive content of curvature, level, and slope. *Journal of Applied Econometrics*, 27(4), 574–619.

20. Sargent, T. J., & Wallace, N. (1981). Some unpleasant monetarist arithmetic. *Federal Reserve Bank of Minneapolis Quarterly Review*, 5(3), 1–17.

21. Vicente, J. (2021). Sovereign risk and the term structure of interest rates: Notes on the Brazilian case. *Banco Central do Brasil*, mimeo; exact origi-

nal venue to be confirmed before submission.

22. Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, 101(4), 1514–1534.

Appendix: Detailed Estimation Procedure

This appendix gives the complete, step-by-step estimation procedure behind Result 2 and the US-control exercise, in enough detail to reproduce every figure in this paper from the underlying DI futures data alone. It is intended as a companion to the methodology section, not a restatement of it: the main text gives the model; this appendix gives every implementation choice, including several that materially affect the result and are not obvious from the model's equations alone.

A.1 From Bloomberg contracts to a fitted curve

Raw data. Daily settlement prices for all listed DI1 futures contracts (Bloomberg tickers of the form ODmyy Comdty, e.g. ODF10 for the January 2010 contract), 2010–2026, plus the B3 trading holiday calendar.

Expiry and business-day counting. Each contract's ticker encodes its expiry month via a standard futures month code (F=Jan, G=Feb, H=Mar, ..., Z=Dec) and a two-digit year. DI1 contracts expire on the first business day of the expiry month. Days-to-expiry d_u for a given trade date is the count of B3 business days between that date and expiry, using the official holiday calendar. Only observations with $d_u > 0$ are retained.

Curve interpolation. For each trade date, the set of (days-to-expiry, rate) pairs across all live contracts is interpolated to a target maturity via flat-forward interpolation on log discount factors:

$$P(d_u) = (1 + y(d_u))^{-d_u/252}, \qquad \log P(d_{uf}) = \text{interp}(d_{uf}, d_u, \log P(d_u))$$

with linear interpolation performed in $(d_u, \log P)$ space, then converted back to an annualized rate.

Two panels are built from this common interpolation routine: a **daily panel**

(13 standard tenors, used for Layer 1/Result 1), and a **monthly panel** (112 vertices at every 21-business-day multiple, resampled to month-end, used for Layer 2/the ACM estimation, and reused without modification by the companion monetary-policy paper).

A.2 Layer 1: Regime-conditional PCA (Result 1)

For each of the four regime windows, principal components are extracted from the daily changes of the 13-tenor panel, $\Sigma w_k = \lambda_k w_k$, with the sign of each loading vector fixed by convention (level loadings normalized positive; slope loadings normalized positive at the short end; curvature loadings normalized negative at the belly) purely for visual comparability across regimes. This layer is fully independent of Layer 2: it uses no VAR, no risk pricing, and no affine recursion.

A.3 Layer 2: The ACM estimation, step by step

A.3.1 State variables. The monthly panel of 112 maturities, standardized column-wise, gives raw factor scores $X_t^{\text{raw}} = Y_z w_{1:3}$, rescaled to unit variance component-wise before proceeding: $X_t = X_t^{\text{raw}} / \text{std}(X_t^{\text{raw}})$. This rescaling is necessary but not sufficient on its own; the input yields are standardized as well.

A.3.2 VAR(1) and the Bauer-Rudebusch-Wu correction. $X_{t+1} = \mu + \Phi X_t + v_{t+1}$ is first estimated by OLS, giving $\hat{\mu}_{OLS}, \hat{\Phi}_{OLS}$. A 2,000-replication bootstrap gives the estimated small-sample bias in $\hat{\Phi}_{OLS}$; the correction $\Phi = \hat{\Phi}_{OLS} - \delta \cdot \text{bias}$ is applied with $\delta = 1$. The innovations v_t used downstream in the excess-return regressions (A.3.4) *must* be computed using the same, bias-corrected Φ used for the risk-neutral dynamics (A.3.5), the physical-measure forecast (A.3.6), and the companion paper's steady-state and forward-iteration calculations – not the uncorrected $\hat{\Phi}_{OLS}$:

$$v_t = X_t - (\mu + \Phi X_{t-1})$$

Using $\hat{\Phi}_{OLS}$-based residuals while using the corrected Φ elsewhere introduces an internal inconsistency between the shocks entering the risk-price estimation and the dynamics those risk prices are subsequently used to price. This does not show up as an estimation error in any individual step; it shows up only in the aggregate shape and magnitude of the recovered term premium and, by the identical logic, in any steady-state or forward-path quantity computed from μ and Φ in the companion monetary-policy paper.

A.3.3 Short-rate equation. $y_t^{(1)} = \delta_0 + \delta_1' X_t + \varepsilon_t$, estimated by OLS on the shortest available monthly maturity.

A.3.4 Excess-return regressions. For each of 16 maturities n (2, 3, 6, 9, 12, 15, 18, 24, 30, 36, 48, 60, 72, 84, 96, and 108 months), the one-month holding-period log excess return $rx_{t+1}^{(n)} = \log P_{t+1}^{(n-1)} - \log P_t^{(n)} - \frac{1}{12}\log(1 + y_t^{(1)})$ is regressed on the contemporaneous VAR innovation and lagged state, giving $a \in \mathbb{R}^{16}$, $B \in \mathbb{R}^{16\times 3}$, $C \in \mathbb{R}^{16\times 3}$.

A.3.5 Risk-price recovery. Stacking the no-arbitrage restriction across all 16 maturities and solving by pseudo-inverse least squares: $\lambda_0 = (B\Sigma_v)^+(a + \text{conv})$, $\lambda_1 = (B\Sigma_v)^+ C$. Omitting the Σ_v factor is not equivalent except in the degenerate case $\Sigma_v = I$; only the Σ_v-inclusive version is consistent with the stated pricing kernel.

A.3.6 Risk-neutral dynamics and the affine recursion. $\mu^Q = \mu - \Sigma_v \lambda_0$, $\Phi^Q = \Phi - \Sigma_v \lambda_1$; log bond prices $p_t^{(n)} = A_n + B_n' X_t$ via the standard recursion, with $A_1 = -\delta_0/12$, $B_1 = -\delta_1/12$. No eigenvalue clamp is applied to Φ^Q: with the internal-consistency requirement of A.3.2 satisfied, Φ^Q's eigenvalues are stationary without any additional constraint.

A.3.7 Expectations-hypothesis yield, the term premium, and the policy-path forecast. $y_t^{(n),EH} = \frac{1}{n}\sum_{i=0}^{n-1} E_t[\delta_0 + \delta_1' X_{t+i}]$, with $E_t[X_{t+i}] = \Phi^i X_t + \sum_{j=0}^{i-1} \Phi^j \mu$, computed by direct iteration of the bias-corrected VAR. The term premium is $\phi_t^{(n)} = y_t^{(n),\text{model}} - y_t^{(n),EH}$. This same forward iteration, and the VAR's own unconditional mean $\bar{r} = \delta_0 + \delta_1' \mu$, are reused directly, without re-estimation, by the companion monetary-policy paper's predicted-path and steady-state calculations.

A.4 The US term premium control

The NY Fed's published monthly ACM decomposition of the US Treasury curve supplies $\phi_t^{(2y),US}$ and $\phi_t^{(10y),US}$. Each Brazilian term premium is regressed on both, in first differences (Section 7), by OLS with Newey-West HAC standard errors.

A.5 The real curve

Section 8's real term premium, breakeven, and US control use the fuller of two real-curve constructions this research programme has built: a cross-sectional Nelson-Siegel-Svensson fit across all outstanding NTN-B bonds (33 bonds, 2005–2026), duration-based interpolation validated against an independent ANBIMA curve to single-digit basis point agreement, and – for the breakeven only – the joint Abrahams-Adrian-Crump-Moench (2016) nominal-real framework anchored to realized IPCA. A companion monetary-policy paper uses a narrower, self-contained real-curve construction of its own, for its own real neutral-rate calculation specifically. The two share a data source, not a pipeline.

A.6 A validation checklist

1. Mean excess returns by maturity should be small and positive, increasing with maturity. Large, negative, roughly maturity-independent val-

ues indicate a risk-free-leg scaling error (A.3.4).

2. Φ^Q's eigenvalue moduli should be under 1 without needing a clamp. If a clamp is required, check the internal-consistency requirement of A.3.2 first.

3. The fitted model yield should track the actual observed curve closely (mean absolute error well under 0.5pp) – a large, systematic gap indicates a specification error upstream of the term premium calculation itself.

4. The 2015–16 and 2021–22 episodes should show clearly opposite-signed term premium behavior of comparable magnitude. A model that reproduces one episode but collapses the other toward zero is displaying a diagnosable asymmetry traceable to A.3.2.