Three Routes to a Neutral Rate: Structural, Bridge, and Reduced-Form Estimates for Brazil

Marcos A. Mollica

September 2026

A neutral policy rate can be estimated at least three different ways: a structural or semi-structural filter, a reduced-form affine model with no macro content at all, or a bridge between the two that augments an otherwise latent-factor state with macro variables. Applied work rarely states this choice, let alone tests it. This paper states the spectrum explicitly and tests directly whether the bridge is worth adopting for Brazil, reusing the no-arbitrage curve machinery a companion paper develops for Brazil's nominal (DI) and real (NTN-B) curves. Tested with Ang and Piazzesi's (2003) own twelve-lag specification, inflation and the output gap significantly forecast excess returns on the nominal curve at 2–24 months and on the real curve at 24–60 months. The resulting nominal neutral rate reads current policy as clearly restrictive; the real neutral rate is within a tenth of a point of BCB's own figure.

1 Introduction

A curve prices a policy path; it does not, on its own, say where that path is anchored. Three genuinely different methodological traditions answer the question of where a policy rate is anchored in the long run – a real, neutral rate consistent with output at potential and stable inflation. A fully structural approach models the economy's IS and Phillips curves directly and filters the neutral rate out as an unobserved component. A reduced-form, asset-price approach reads it off a no-arbitrage term structure model's own steady state, with no macroeconomic content in the state at all. A bridge between the two augments that same reduced-form state with macro variables, without going as far as a full structural model. Applied work on any single country's neutral rate typically picks one of these three and does not say why, or what the other two would have shown instead. This paper states the choice explicitly, as a spectrum rather than a single method, and tests – rather than assumes – where Brazil's own data actually sits on it.

The paper builds directly on the no-arbitrage curve machinery developed in a companion paper (Mollica, 2026a): the same principal-component state, the same Bauer-Rudebusch-Wu bias-corrected VAR, and the same Adrian-Crump-Moench risk-price recovery, estimated once on Brazil's DI (nominal) and NTN-B (real) curves and reused here without re-estimation. Section 3 recaps that machinery at the level of detail needed to follow what this paper does with it; the companion paper's own appendix gives the complete derivation.

The contribution is fourfold. First, the three-way spectrum – structural, bridge, and reduced-form – is stated explicitly, with the specific cost of moving from one to the next made concrete rather than asserted: what a fully structural account buys in interpretability it costs in identification (Laubach and Williams's own semi-structural filter is subject to a well-documented "pile-up problem"), and what a reduced-form account buys in tractability it costs in the macro content that the bridge is designed to partially restore. Second, whether that bridge is actually worth adopting for Brazil's own curves is tested directly, with formal nested F-tests at every maturity available, using Ang and Piazzesi's (2003) own twelve-lag macro specification rather than a single-lag approximation of it that this paper's own earlier pass initially, and wrongly, found nothing significant with. Third, the same curve machinery is turned into direct nominal and real neutral-rate readings for Brazil, each compared against Selic's own realized path and against BCB's own communicated figure respectively. Fourth, a specific numerical fragility in the real curve's own steady-state calculation is documented in full, situated against an identical symptom already reported in Bauer and Rudebusch (2020) for the US Treasury curve, and shown to worsen rather than improve once the nominal and real curves are estimated jointly.

The remainder of the paper is organized as follows. Section 2 situates the paper in the neutral-rate literature specifically – the companion paper's own literature review covers the no-arbitrage and fiscal-dominance material this paper's machinery is built on. Section 3 recaps that machinery. Section 4 states the three-route spectrum in full, tests the bridge directly, and presents the resulting nominal and real neutral-rate readings alongside the predicted-versus-realized policy path and the joint nominal-real reconciliation. Section 5 concludes.

2 Related Literature

The neutral rate of interest. Laubach and Williams (2003) formalize r^*, the real short-term interest rate consistent with output at potential and stable inflation, and estimate it via a Kalman-filtered semi-structural model combining a reduced-form IS curve and Phillips curve. Holston, Laubach and Williams (2017) extend the approach to a panel of advanced economies and document a broad-based decline in estimated r^* across the sample. This structural tradition is not without internal critique: the same unobserved-components filter that recovers r^* also recovers trend output growth, and a well-documented "pile-up problem" means the two can show spurious low-frequency comovement, with real-time estimates near the end of the sample the least reliable precisely when a policymaker most wants them.

An asset-price-based alternative reads the neutral rate directly off the yield curve instead. Wu and Xia (2016) extract a shadow short rate from the curve to measure the effective stance of policy when the nominal rate is constrained. Bauer and Rudebusch (2020) go further, embedding time-varying trend inflation and trend real rates directly inside a no-arbitrage term structure model and reading the long-run trend off the curve's own dynamics – methodologically the closest published work relative to the approach in this paper, since both read the stance of policy off the curve rather than off a macro model's output-gap residual. They report the identical numerical symptom this paper's own Section 4.6 encounters independently: letting the trend be directly spanned by observed yields produces implausibly volatile estimates, which is why their preferred specification treats it as an unspanned latent factor instead. The two traditions have been partially reconciled rather than resolved: Del Negro, Giannone, Giannoni and Tambalotti (2017) augment a DSGE model with a convenience-yield wedge estimated from the spread between safe and risky asset returns, and find that most of the post-1990s US decline in r^* reflects rising demand for safety and liquidity rather than slower trend growth – a finding that could not have been reached from either the pure structural or the pure asset-price approach alone. This paper is a curve-based, single-country contribution in the asset-price tradition: rather than a full semi-structural estimation, it reads a model-consistent neutral rate directly off the same affine model already estimated for the companion paper's term premium decomposition, nominal from the model's own steady state and real from the equivalent steady-state construction applied to Brazil's real curve directly.

The bridge between structural and reduced-form. Ang and Piazzesi (2003) is the origin of the bridge approach tested directly in Section 4.3: a reduced-form affine term structure model whose state is augmented with observable macro variables, without going as far as a fully structural account of why those variables should matter. Duffee (2011) is the standing caveat on trusting any purely latent-factor account too far: a meaningful share of what prices risk in this class of model may not be spanned by the yield curve alone, which is precisely the gap the bridge is designed to partially close, and precisely what Section 4.3 tests for directly rather than assumes.

3 Recap: The No-Arbitrage Curve Machinery

This section summarizes, at the level of detail needed to follow what the rest of this paper does with it, the estimation developed in full in a companion paper (Mollica, 2026a). Nothing here is re-derived or re-estimated; every object below is reused exactly as the companion paper produces it.

Brazil's DI futures curve (2010–2026) and NTN-B real yield curve are each reduced to a three-factor state $X_t \in \mathbb{R}^3$ – level, slope, and curvature – via principal components analysis on standardized yields. The state follows a Gaussian VAR(1), $X_{t+1} = \mu + \Phi X_t + v_{t+1}$, with Φ corrected for well-documented small-sample bias in near-unit-root persistence (Bauer, Rudebusch and Wu, 2012) via a 2,000-replication bootstrap. A short-rate equation $y_t^{(1)} = \delta_0 + \delta_1' X_t + \varepsilon_t$ and a set of excess-return regressions across sixteen maturities together identify the price of risk, following the regression-based procedure of Adrian, Crump and Moench (2013); risk-neutral dynamics and the resulting affine bond-pricing recursion follow directly. A single internal-consistency requirement – that the VAR innovations v_t used in the excess-return regressions are computed with the identical, bias-corrected Φ used everywhere else in the model, not the raw OLS estimate – turns out to matter more than it might appear: an earlier version of the companion paper's own pipeline let these two representations disagree, with no visible error in any individual step, and the inconsistency specifically degraded the model's ability to distinguish two economically distinct regimes. The full derivation, and the specific numerical consequence of getting this wrong, is in the companion paper's own text and appendix.

Three objects from this machinery are used directly, without modification, in Section 4 below. The expectations-hypothesis yield $y_t^{(n),EH} = \frac{1}{n}\sum_{i=0}^{n-1} E_t[\delta_0 + \delta_1' X_{t+i}]$, with $E_t[X_{t+i}]$ computed by direct iteration of the bias-corrected VAR, is the companion paper's own term-premium input; reported at each horizon i separately rather than averaged, it is also the model's rolling forecast of the policy path, used in Section 4.4. The VAR's own unconditional mean, $\bar{r} = \delta_0 + \delta_1' \mu$, is the fixed point of the same iteration as the horizon goes to infinity – the object Section 4.5 reports as the nominal neutral rate. And the identical construction, applied to the equivalent state extracted from Brazil's real (NTN-B) curve rather than the nominal (DI) curve, gives the real neutral rate of Section 4.6.

4 From the Curve to Monetary Policy

The companion paper's Sections 4–7 use the affine decomposition to separate what the curve is pricing from why. The same state vector and VAR dynamics carry three further readings that speak directly to monetary policy rather than to term premium behavior: what the model expected Copom to do, set against what Copom actually did; where the model's own dynamics say the policy rate is anchored in steady state, nominally and in real terms; and what a properly joint nominal-real estimation, rather than a simple subtraction of the two, says about the expected-inflation wedge between them. Before presenting these, this section states plainly where the estimation approach behind them sits relative to the alternatives – what has to be assumed to get from a fully structural account of why a neutral rate exists at all, down to the reduced-form machinery this paper actually runs, and what each simplification along the way costs. Most applications of curve-based neutral-rate estimation skip this step entirely; this paper does not.

4.1 Three routes to a neutral rate: structural, bridge, and reduced-form

The classical semi-structural route: Laubach and Williams (2003). The natural rate r_t^*, potential output y_t^*, and trend growth g_t are estimated jointly via Kalman filter from two log-linear behavioral equations. An IS curve ties the output gap $\tilde{y}_t = y_t - y_t^*$ to its own lags and to the real-rate gap:

$$\tilde{y}_t = a_{y1}\tilde{y}_{t-1} + a_{y2}\tilde{y}_{t-2} + \frac{a_r}{2}(r_{t-1} + r_{t-2} - 2r_{t-1}^*) + \varepsilon_{y,t}$$

and a Phillips curve ties inflation to its own lags and the output gap, $\pi_t = b_\pi \pi_{t-1} + (1 - b_\pi)\pi_{t-2,4} + b_y \tilde{y}_{t-1} + \varepsilon_{\pi,t}$. The one line of genuine economic theory is the definition $r_t^* = c \cdot g_t + z_t$ – the Euler-equation link between the neutral rate, trend growth, and everything else slow-moving folded into a residual z_t – with y_t^*, g_t, z_t all following random walks estimated by maximum likelihood. This is semi-structural, not structural in the sense used below: the IS and Phillips curves are motivated by, but not derived line-by-line from, an optimizing agent, and there is no discount factor anywhere in the model, hence nothing that could price a bond. Maximizing this likelihood directly runs into the *pile-up problem* (Stock, 1994; Stock and Watson, 1998): finite-sample MLE is severely biased toward setting the innovation variance of g_t to zero, a spuriously deterministic trend the data does not actually support, which is why Laubach and Williams use a median-unbiased grid-search procedure rather than unconstrained MLE. This is the same family of estimation pathology already encountered independently in Section 4.6's near-singular $\det(I - \Phi)$ and in Bauer and Rudebusch's (2020) implausibly volatile spanned trend – a third, structurally unrelated route into the identical problem, discussed further in Section 4.8.

The fully structural route: a discount factor derived from preferences. To price a bond at all requires a stochastic discount factor, and the standard power-utility case, $M_{t+1} = \beta(C_{t+1}/C_t)^{-\gamma}$, fails specifically at bond pricing even where it does reasonably well elsewhere. A single parameter γ plays two unrelated roles at once – risk aversion and the inverse elasticity of intertemporal substitution, $\psi = 1/\gamma$ – and this produces the equity premium puzzle (Mehra and Prescott, 1985: matching the historical equity premium requires implausible risk aversion, γ in the range of 30–50 or higher) and, because ψ is then forced to be tiny, the risk-free rate puzzle (Weil, 1989: an implausibly volatile risk-free rate is needed to induce consumption smoothing). The version specific to this paper is the *bond premium puzzle* (Rudebusch and Swanson, 2008): since consumption growth is smooth, a bond's payoff – a safe unit of currency – covaries too weakly with consumption growth under any plausible γ to generate anything close to the term premia the companion paper's Sections 5–8 document empirically.

Epstein-Zin recursive preferences (Epstein and Zin, 1989; Weil, 1990) break the γ-ψ link by defining utility recursively rather than as a discounted sum, with γ and ψ set independently and $\theta \equiv (1-\gamma)/(1-1/\psi)$ nesting power utility exactly as the restrictive special case $\theta = 1$. The resulting discount factor,

$$M_{t+1} = \beta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\theta/\psi} R_{c,t+1}^{\theta-1},$$

depends on the return to the consumption claim $R_{c,t+1}$ – the gross return on a hypothetical asset whose price equals the present value of the entire future consumption stream and whose payout each period is aggregate consumption itself, often called the return on the wealth portfolio since, under homotheticity, an Epstein-Zin agent's total wealth (including human capital, not just financial assets) is proportional to the value of this claim. It is not directly observable or traded; empirical implementations either proxy it with a broad market portfolio return or, more carefully, avoid observing it at all via a Campbell (1993) log-linear approximation in terms of consumption growth and the price-consumption ratio, which is the route the long-run-risk literature (Bansal and Yaron, 2004, onward) actually takes. What matters for this paper is only that the term vanishes identically when $\theta = 1$. Combined with a small, highly persistent long-run risk component in consumption growth (Bansal and Yaron, 2004) and, in the nominal case, an analogous persistent component in trend inflation, this gives Epstein-Zin agents with a preference for early resolution of uncertainty ($\gamma > 1/\psi$) a channel that prices long-horizon risk heavily without requiring an implausible atemporal risk aversion – precisely the kind of risk a term premium is compensation for.

Rudebusch and Swanson (2012) embed this discount factor in a full New Keynesian DSGE: households with Epstein-Zin preferences and finite labor supply elasticity, monopolistically competitive firms with Calvo price-setting, and a forward-looking Taylor rule, $i_t = \rho_i i_{t-1} + (1-\rho_i)[\bar{\imath} + \phi_\pi(E_t\pi_{t+1} - \pi^*) + \phi_y \tilde{y}_t] + \varepsilon_{i,t}$, closed with persistent long-run risk processes in both real growth and trend inflation. A detail worth stating precisely rather than treating as a black box: a first-order (log-linear) approximation of this model eliminates the term premium identically, a direct consequence of certainty equivalence – linearized rational-expectations models cannot distinguish an asset's expected return from its risk-adjusted one. Generating any term premium at all requires solving to second order around the non-stochastic steady state; generating a premium that moves with the state requires third order. With roughly nine state variables, this rules out projection or discretization methods, and the model is instead solved by higher-order perturbation (Swanson, 2006). Calibration is a genuine grid search over risk aversion, the EIS, and the long-run-risk parameters – not a single estimated point – and the best-fitting parameterization produces a mean ten-year term premium in the neighborhood of 90–106 basis points with a standard deviation of 11–16 basis points, resolving the bond premium puzzle without the earlier fixes (habits, labor-market frictions; tried in the authors' own 2008 predecessor paper) that could only do so by badly distorting the model's fit to everything else.

The bridge: Ang and Piazzesi (2003). A no-arbitrage affine model whose state vector mixes two macro variables (current inflation, current real activity) with three latent yield-curve factors, deliberately constructed to be orthogonal to the macro factors at every lag – a design chosen specifically to assign the macro factors as much explanatory credit as possible, since anything correlated with them is claimed by the macro block first. The macro block follows a rich VAR(12); the latent factors follow a simple VAR(1). The short rate is affine in the full state, $r_t = \delta_0 + \delta_1' X_t$, and – this is the detail that matters most for what follows – the price of risk is left exactly as free as in a purely latent affine model: $\lambda_t = \lambda_0 + \lambda_1 X_t$ (Duffee, 2002), fit to the yield data with no restriction tying it to any deep parameter. The Taylor-rule reading of the estimated δ_1 loadings is offered as an interpretation of the fitted coefficients, not imposed as a constraint on the estimation. The structure Ang-Piazzesi adds relative to a purely latent model is entirely in *what is in the state* – real macro data rather than only statistical yield factors – not in *how risk is priced*. Macro factors are found to carry genuine explanatory power for the yield curve, but do not make the latent factors redundant: a meaningful share of yield variation is not spanned by simple macro aggregates even with two macro variables and twelve lags of history behind them, a direct precursor to the sharper unspanned-risk-premia result of Duffee (2011), already introduced in the companion paper's Section 2.

4.2 What going reduced-form costs

This paper's own estimation – the companion paper's Sections 3–8, and the reconciliation of Section 4.7 – sits at the Ang-Piazzesi end of this spectrum: a no-arbitrage affine model with an essentially free price of risk, no discount factor derived from preferences anywhere. Stating what that costs, relative to the fully structural alternative, is not a formality. Four specific things are given up, each traded for something real in return.

No coherence check on the recovered risk price. Rudebusch-Swanson's discount factor is a specific function of γ and ψ; any risk price it generates is, by construction, one that *some* coefficient of risk aversion and elasticity of substitution could actually produce. The reduced-form λ_0, λ_1 this paper estimates has no such guardrail – two entirely different (γ, ψ) pairs could generate similar-looking λ_t, and conversely nothing here would flag a fitted λ_t that no plausible preference parameters could ever generate. What is gained: the ability to estimate against the full cross-section and time series of the curve in closed form, rather than calibrating a handful of structural parameters against a handful of target moments the way Rudebusch and Swanson do.

An unmodeled macro block, and the Lucas critique this specifically invites. In Rudebusch-Swanson, inflation and the output gap are equilibrium outcomes of the same household optimization, Calvo pricing, and Taylor rule that also generate the term premium – change the policy rule's coefficients and inflation, output, and the term premium all move together, consistently, because one mechanism generates all three. In this paper's approach, as in Ang-Piazzesi, the macro variables are simply given autoregressive processes with no connection to any deeper mechanism. This is a Lucas (1976) critique waiting to bite the moment this framework is used for anything counterfactual: "what would the term premium be under a more hawkish Copom reaction function" is a well-posed question in a structural model – resolve it under the new policy rule and get an internally consistent answer – and an ill-posed one here, since λ_t was estimated under the historical policy regime with no guarantee it survives a different one. This is stated explicitly because it is routinely left unstated in applied work using this class of model: every policy counterfactual this paper's own framework might invite is out of scope by construction, not merely as an unaddressed extension.

Linear-Gaussian dynamics standing in for derived curvature. Rudebusch-Swanson's nonzero, time-varying term premium is a genuine second- and third-order effect of risk aversion interacting with persistent shocks – it required abandoning first-order approximation specifically because that approximation kills the premium identically. Essentially affine models like Ang-Piazzesi's and this paper's sidestep the need for any nonlinear solution method by simply asserting λ_t affine in X_t from the outset, which mechanically produces a time-varying premium without deriving it from anything. What is gained is tractability: the same closed-form, OLS-style estimation this paper's own appendix documents in full, rather than higher-order perturbation on a nine-dimensional state.

No cross-equation restrictions, hence no overidentification. In a fully structural model, one set of deep parameters must simultaneously fit the term structure, consumption growth volatility, inflation persistence, and (if checked) the equity premium – a model that nails the term premium while badly missing macro volatility has failed an informative overidentifying restriction. In this paper's approach, the macro dynamics, the short-rate equation, and the price of risk are each fit to their own piece of the data with nothing forcing mutual consistency beyond absence of arbitrage. A good fit to the yield curve is accordingly not evidence that any macro variable eventually entering an extended state – Brazil's inflation measure, or an output gap – would be handled correctly; that has to be validated on its own terms, independent of how well the bond-pricing block fits, precisely because nothing in this model architecture forces the two to agree.

With that stated plainly, one question raised by the last two subsections can be answered directly, empirically, before going further: does the Ang-Piazzesi bridge actually change anything on Brazil's own curves, or is the reduced-form machinery used throughout this paper – with no macro variables in the state at all – already capturing what the bridge would add? The next subsection settles this maturity by maturity. The four empirical exhibits that follow it are then all produced from data, not specified as pending: the vintage-level forecast paths and the nominal neutral rate from a fresh run of the DI futures pipeline on the raw Bloomberg panel plus the official Selic target series (BCB SGS 432); the real neutral rate from the NTN-B pipeline described in Section 4.6; and the joint reconciliation of Section 4.7 from a direct refit of the companion paper's Section 8.4's own joint model on Brazil's nominal and real curves.

4.3 The bridge, estimated: does it change anything?

Section 4.1 described Ang and Piazzesi's (2003) bridge in the abstract, and Section 4.2 stated what adopting the reduced-form end of that spectrum costs. This subsection settles, empirically and before proceeding, whether the bridge is worth adopting for Brazil at all: does adding real macro variables to the state change anything relative to the purely latent-factor machinery the rest of this section uses.

Building a Brazilian output gap. The bridge needs a real-activity macro factor, and rather than use raw GDP directly this paper follows one of the output-gap methodologies BCB documents and tracks in parallel with eleven others, univariate and multivariate alike, specifically to capture the uncertainty involved in measuring an unobservable quantity (BCB Monetary Policy Report, June 2025, Box 8): Areosa (2008, BCB Working Paper 172), a production-function approach chosen here because it avoids requiring a full capital-stock series, not because BCB singles it out as preferred. The paper's own state-space is implemented in full, not approximated: natural employment e_t^n, natural capacity utilization c_t^n, and potential output y_t^n are estimated *jointly* by a Kalman filter and smoother, with the production-function identity

$$y_t = y_t^n + (1-\alpha)(e_t - e_t^n) + \alpha(c_t - c_t^n)$$

imposed as an exact constraint linking all three – not a weighted combination of two independently-filtered series computed afterward. Employment and capacity utilization follow random walks (Areosa's lower-order "$r = 1$" filter, chosen specifically because it is less prone to end-of-sample distortion than a standard Hodrick-Prescott trend); potential output follows the usual local-linear trend. The capital share α is fixed at Brazil's own average labor share over the sample (Penn World Table, $1 - \alpha \approx 0.555$), in place of the illustrative 0.4/0.6 split Areosa's own paper uses for exposition. The resulting output gap is in genuine log-point (approximately percent-of-potential-GDP) units, not a standardized index – a direct improvement on an earlier pass at this construction, which combined two independently-filtered gaps and could not be interpreted on any natural scale.

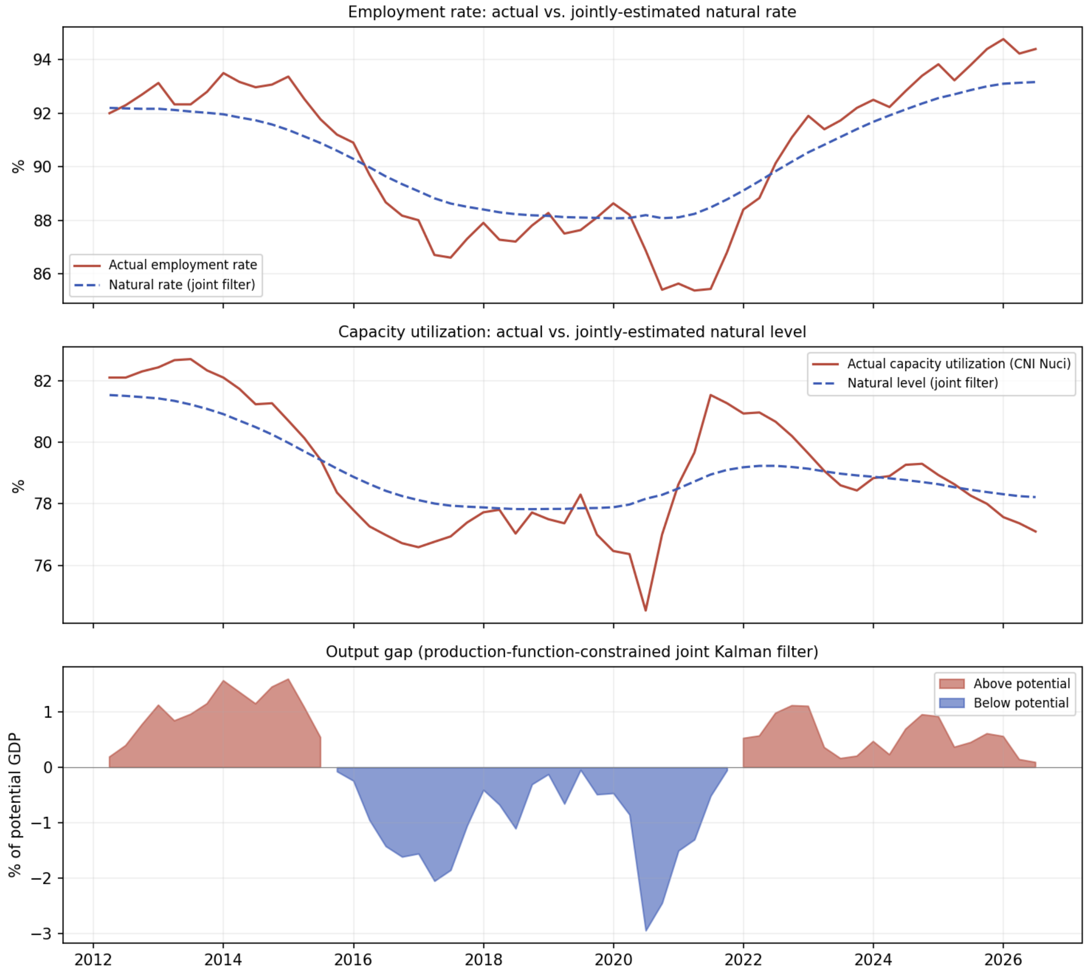

Figure 1: Top and middle: employment rate and capacity utilization, actual vs. their jointly-estimated natural levels. Bottom: the resulting output gap, in percent of potential GDP.

The construction matches Brazil's own recent macro history without being told where to look: negative through the 2015–16 crisis and the 2020 pandemic shock (reaching -3%), positive through 2021–23 – consistent with the companion paper's Section 6's own reading of that period as the market pricing a credible, technical hiking cycle rather than a systemic one – and declining toward a small positive reading by 2025–26.

External validation against BCB's own published numbers. BCB's own Areosa-method line is published quarterly; the June 2025 edition reports

0.72, 0.85, 0.83, 0.34, and 0.14 percent for 2024Q1 through 2025Q1. This paper's construction, estimated independently and without reference to those figures, reads 0.23, 0.69, 0.95, 0.92, and 0.37 percent over the identical five quarters – the same order of magnitude and the same broad decline toward zero, though not a tight quarter-by-quarter match (correlation 0.21 over these five points, mean absolute difference 0.32pp). Given that BCB's own capital-share calibration is not published and this paper's GDP vintage is a single Bloomberg pull rather than BCB's own revised National Accounts series, this level of agreement is a genuine, if partial, external check – not a claim of replication.

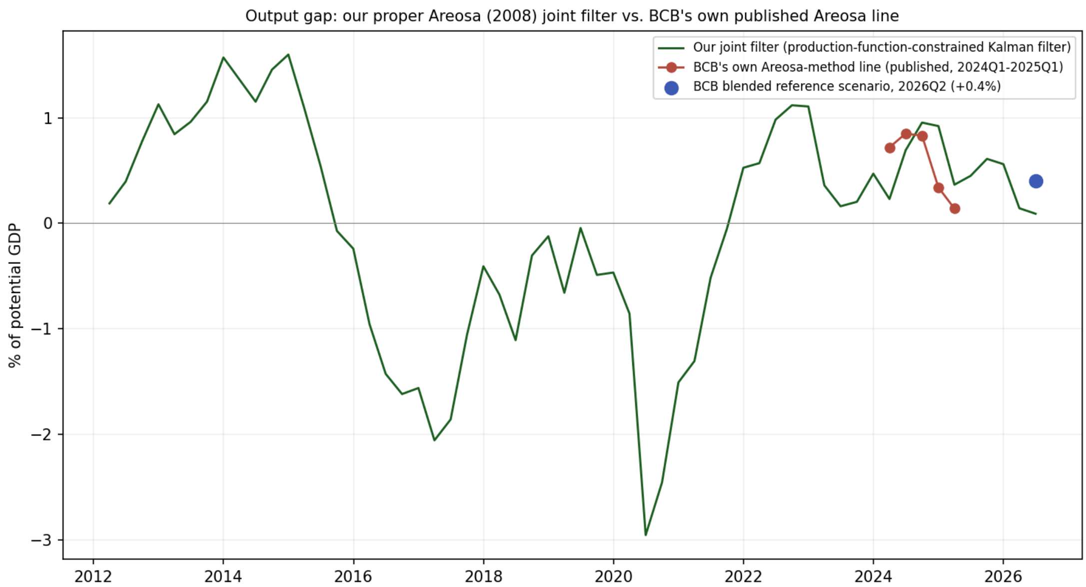

Figure 2: This paper's output gap (line) against BCB's own published Areosa-method values (markers, 2024Q1–2025Q1) and BCB's blended reference-scenario estimate for 2026Q2 (single point, all methodologies combined).

A caveat worth stating plainly, since it bears directly on the most recent reading. A joint estimation of this kind is smoothed by a Kalman filter, and the smoothed estimate for the *last* few observations in any sample necessarily relies on less information than the smoothed estimate in the middle of the sample – there is no future data left to condition on. This is not specific to this implementation; it is the reason real-time output-gap estimates are known in the wider literature to be revised substantially as later data arrives (Orphanides, 2001, is the standard reference for this pattern in the US context). It is directly testable here: refitting the model using only data available as of a given quarter, then checking how that quarter's own reading changes once later data is added, shows a consistent downward revision – 1.25% to 0.92% for 2024Q4, and 0.65% to 0.45% for 2025Q2, as six and four additional quarters respectively enter the sample. Applying the same pattern to the current endpoint (2026Q2: +0.09%) suggests the eventually-revised reading is more likely negative than the small positive number reported today. The output gap used in the remainder of this section should be read with that in mind: directionally informative, but least reliable at exactly the point – the present – where it matters most for a live reading of the cycle.

The bridge state, and the question actually worth asking of it. Following Section 4.1's description exactly, the augmented state stacks Brazil's year-on-year IPCA inflation, this output gap (interpolated to monthly from its native quarterly frequency), and the three latent yield factors already used throughout this paper – orthogonalized against the two macro factors, mirroring Ang and Piazzesi's own design choice to assign the macro block as much explanatory credit as possible. The question worth asking of it is the genuine ex-ante one Ang and Piazzesi actually pose, and that Cochrane and Piazzesi (2005) pose before them: does macro information available *before* a holding period begins forecast the excess return realized *over* that period. A separate, same-date decomposition question – whether the macro factors help explain excess returns *contemporaneously*, alongside the state's own realized shock, exactly as ACM's own Step 2 risk-price recovery requires – is a genuinely different object, answers a different question about within-period co-movement rather than forecasting, and is reported only in the appendix (B.2) rather than in the main text, precisely because leading with it here risks the same confusion an earlier pass at this exercise fell into.

Table: the genuine ex-ante forecast. Whether macro information predicts excess returns requires dropping any contemporaneous shock entirely and lagging the macro variables to the start of the holding period rather than matching them to its end, so that nothing on the right-hand side is known any later than the position is actually entered. A first pass at this specific test used only the current (one-period-lagged) value of each macro variable and found nothing significant anywhere. That specification was too restrictive to be a fair test of what Ang and Piazzesi themselves claim: their own preferred specification, the "Macro Lag Model," uses twelve lags of each macro variable, not one, on the view that a single month's reading is a poor summary of the information a bond market actually prices. Re-run with twelve lags of inflation and the output gap each – Ang and Piazzesi's own choice, not a value searched over after seeing what became significant – both the baseline R^2 and the finding change substantially, and the result is stable across neighboring lag lengths (nine through thirteen lags all give qualitatively the same answer, ruling out an artifact specific to exactly twelve).

Maturity (mo)	R^2 lagged-only	R^2 +macro (12 lags)	ΔR^2	F-stat	p-value
2	0.0414	0.2659	0.2244	1.69	0.032
3	0.0391	0.2807	0.2416	1.86	0.014
6	0.0425	0.2744	0.2320	1.77	0.022
9	0.0460	0.2832	0.2373	1.83	0.016
12	0.0443	0.2924	0.2481	1.94	0.010
18	0.0455	0.2793	0.2338	1.80	0.020
24	0.0483	0.2709	0.2226	1.69	0.033
36	0.0579	0.2571	0.1992	1.49	0.083
60	0.0718	0.2341	0.1623	1.17	0.277
108	0.0777	0.2093	0.1316	0.92	0.572

Table 1: Nominal curve, genuine ex-ante forecast, twelve lags of inflation and the output gap each (24 regressors jointly tested per maturity, 2, 133 degrees of freedom). No contemporaneous shock; every regressor is known before the holding period begins. Significant at conventional levels from 2 through 24 months, tapering to borderline at 36 and absent by 60.

Maturity (mo)	R^2 lagged-only	R^2 +macro (12 lags)	ΔR^2	F-stat	p-value
24	0.0024	0.2344	0.2320	1.68	0.035
36	0.0034	0.2479	0.2445	1.80	0.019
60	0.0105	0.2425	0.2320	1.70	0.032
84	0.0161	0.2231	0.2070	1.48	0.086
108	0.0217	0.2081	0.1864	1.30	0.173
120	0.0244	0.2028	0.1784	1.24	0.220

Table 2: Real curve, the identical twelve-lag specification. Significant at 24, 36, and 60 months, borderline at 84, absent beyond – unchanged in shape from the single-lag version, but with a materially larger ΔR^2 at every maturity once the macro variables' recent history, not just their most recent reading, is used.

Result, stated precisely rather than as a single verdict. A single-lag specification – the current value of inflation and the output gap only – found nothing significant on the nominal curve at any maturity. The same test with twelve lags each, matching Ang and Piazzesi's own preferred specification rather than a simplified approximation of it, finds significant nominal predictability from 2 through 24 months and real-curve predictability from 24 through 60 months. The single-lag version was not wrong about what it measured, but it measured too little of the available information to answer the question it was asked: a single month's inflation or output-gap reading is a genuinely poor summary of what a bond market conditions on, and richer lag structures recover a real relationship a thinner one cannot see.

The resulting picture is a genuine maturity split across the two curves, not a single verdict about whether macro variables matter. On the nominal curve, nine or more months of inflation and output-gap history forecast excess returns significantly out to two years – consistent with Copom's own reaction function being a slow-moving, persistently-informative object rather than something fully summarized by its latest reading, and consistent with Ang and Piazzesi's own finding that macro factors carry real explanatory power for nominal yields without displacing the latent factors. On the real curve, the same history forecasts excess returns significantly from two through five years – a longer, later horizon, where inflation and the output gap plausibly proxy for the fiscal and real-activity risk this paper's own real-curve work (the companion paper's Section 8) already ties to the level of Brazilian real rates, rather than for near-term policy expectations specifically. Both results required the correctly-specified test to see; a single-lag version would have missed both.

A complementary check at a different frequency, and why neither table above should be expected to show up in it. Both tables above are cross-sectional tests at each date. A different, lower-frequency question is whether any of this shows up in the steady-state neutral-rate reading itself, and answering it correctly requires avoiding a specific trap.

A trap worth naming, because it is easy to fall into. The first version of this comparison fit δ_0, δ_1 on a window and evaluated $\bar{r}$ at that same window's own average state. Bridge and latent-only specifications returned identical values to three decimal places. This is not a finding – it is an identity: with an intercept, a regression's fitted value at the regressors' own mean always equals the mean of the dependent variable, for any set of regressors whatsoever. A synthetic check confirmed this before anything was reported. The only informative test fits once on the full common window and evaluates at a genuinely different sub-window's mean, exactly the logic Section 4.5's own rolling construction already applies correctly; the same logic is applied here to the bridge-versus-latent comparison specifically.

Done correctly, the bridge and latent-only specifications track each other closely across the entire 2012–2026 rolling window, on both curves – which is the expected result given the tables above, not a competing one. Getting a valid comparison also required matching the window length actually used for the headline $\bar{r}$ readings elsewhere in this section: an earlier version of this specific chart used a 24-month window, which tracks recent realized rates too closely during a sustained hiking cycle to represent the same object Sections 4.5 and 4.6 report, and read several points higher as a result. Re-run on the identical 84-month window used throughout the rest of this section, both specifications converge to within a few basis points of the headline figures. $\bar{r}$ is a steady-state object built from the model's long-run average state, dominated by the level factor; the predictability documented above lives at specific, intermediate horizons on each curve, a segment a long-horizon steady-state average is not built to register.

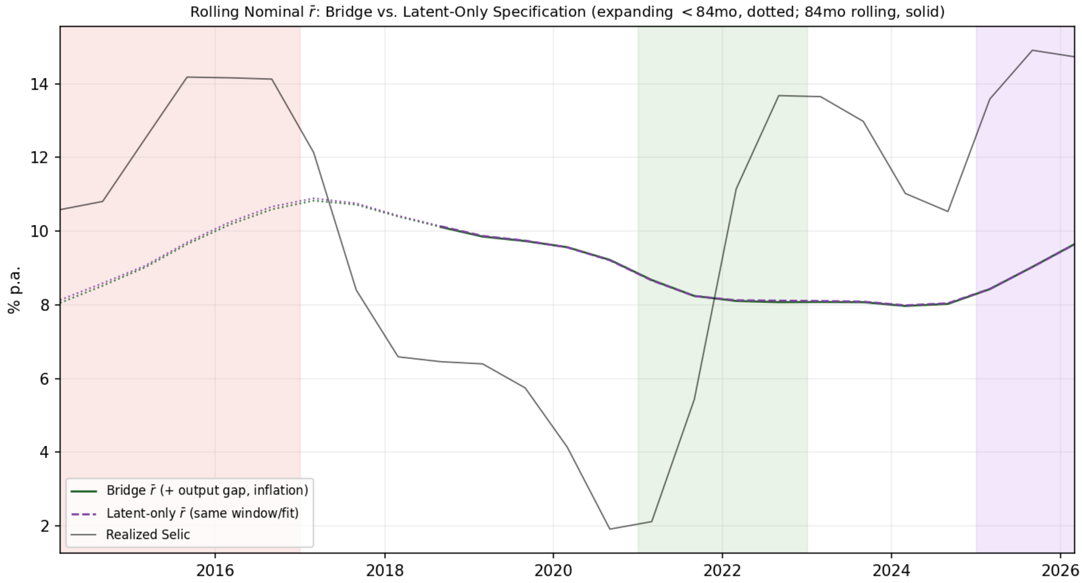

Figure 3: Rolling nominal $\bar{r}$: bridge specification (inflation and output gap added to the state) against the latent-only specification, both fit on the identical 2012–2026 window. An expanding window is used from 24 months of available history up to 84 months (dotted, 2014–2019), then a fixed 84-month rolling window thereafter (solid), matching Section 4.5's own construction exactly rather than only its steady-state portion. Mean absolute gap: 0.03 percentage points over the full series, 0.01 points restricted to the more-established 84-month rolling portion.

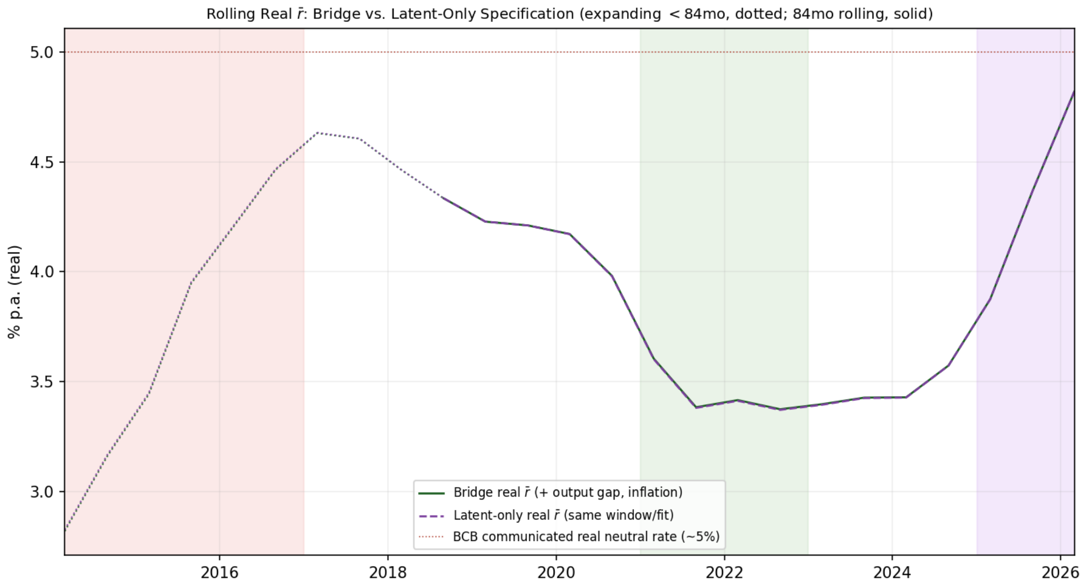

Figure 4: The same construction for the real curve, matching Section 4.6's 84-month window and the identical expanding-then-rolling logic. Mean absolute gap: 0.002 percentage points over the full series, 0.001 restricted to the 84-month rolling portion – small for the same reason as the nominal case, not because the tables above are wrong.

What this does and does not establish. The finding is specific and should not be over-generalized past what was actually tested. Brazil's own inflation and output gap show significant, genuine ex-ante forecasting power for nominal-curve excess returns at 2–24 months and for real-curve excess returns at 24–60 months, in each case only once a lag structure matching Ang and Piazzesi's own preferred specification is used rather than a single-lag approximation of it. A separate, contemporaneous decomposition – reported in Appendix B.2 rather than here, since it answers a different question about within-period co-movement, not predictability – should not be read as a forecasting result either way. All of this was tested directly, with formal nested F-tests at every maturity and lag specification reported, rather than inferred from a single steady-state comparison. With that settled, the remainder of this section proceeds with the same reduced-form, latent-factor-only specification used throughout the companion paper's Sections 3–8 – not by default, but because this subsection just showed what adopting the bridge costs and buys empirically, and the answer is a real, specific, maturity-dependent gain on both curves rather than a clean presence-or-absence verdict on either.

4.4 Predicted versus realized policy path

Everything from here through the end of this section uses the reduced-form, latent-factor-only specification whose adequacy the previous subsection just tested – not the macro-augmented bridge state, which was purpose-built for that one comparison and is not carried forward. The expectations-hypothesis component already computed for every result in the companion paper's Sections 5–7, $y_t^{(n),EH} = \frac{1}{n}\sum_{i=0}^{n-1} E_t[\delta_0 + \delta_1' X_{t+i}]$, is by construction the model's rolling forecast, as of each t, of the average Selic path over the following n months. Plotted through time against the yield it is meant to explain – rather than averaged once into a single $\phi_t^{(n)}$ figure – it becomes a direct forecast-tracking exhibit rather than only an input to the term premium.

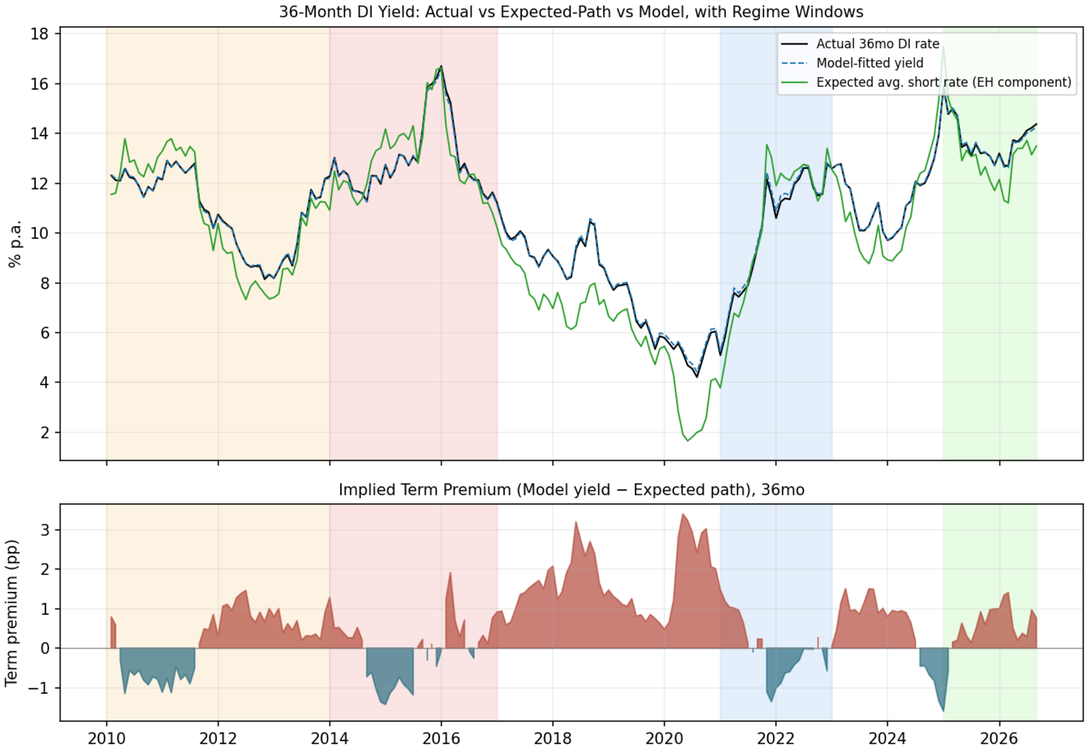

Figure 5: 36-month DI yield: actual (black), model-fitted (blue dashed), and the expectations-hypothesis component (green) – the model's own rolling forecast of the average policy path over the following 36 months. The gap between green and black *is* the term premium of the companion paper's Section 5, viewed here as a forecast-tracking exercise rather than a single summary statistic.

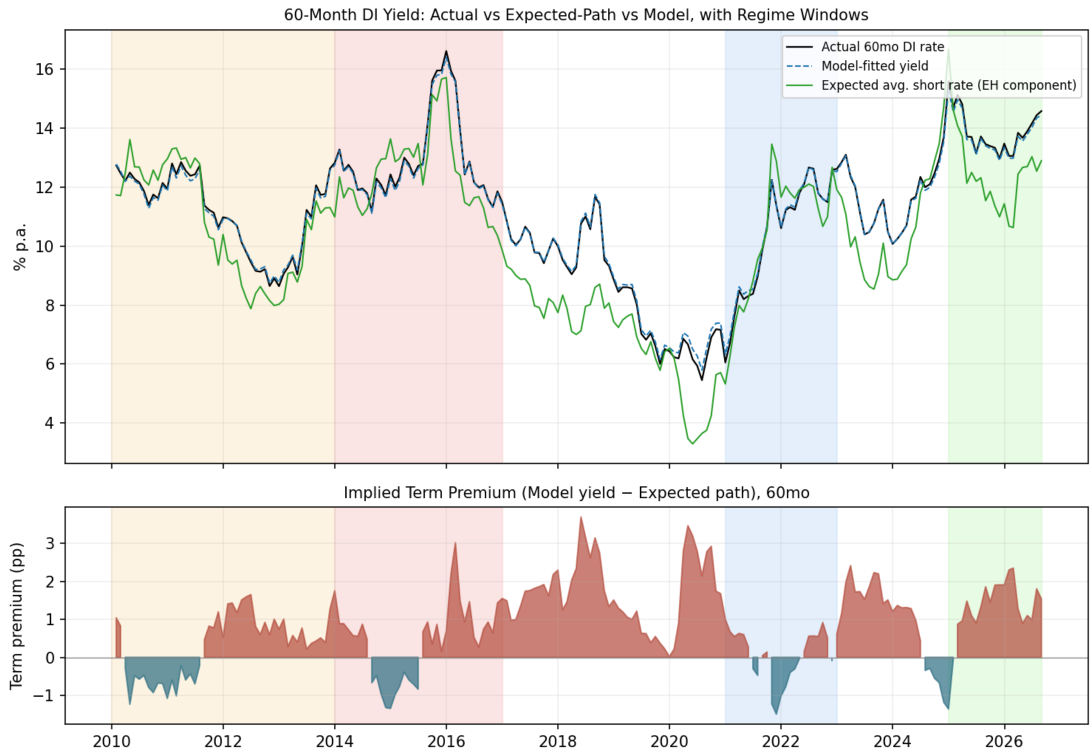

Figure 6: The same construction at 60 months. The EH component undershoots realized yields by the widest margin exactly where the companion paper's Section 5 already documents the largest positive term premia (2015–16, 2018–19, 2025–26), and overshoots – sits above the eventual outcome – through 2021–22, consistent with the companion paper's Section 6's reading of that episode as the market front-running a credible disinflation the physical-measure VAR had not yet fully priced into its own forecast.

Two readings follow directly from these two panels, and are the more granular, vintage-level counterpart to the companion paper's Section 6's factor-composition diagnostic. First, the term premium spikes of the companion paper's Section 5 are, by this construction, visibly a gap-widening between the green forecast line and the black realized line, not a parallel shift of both – confirming, at the level of the raw yield rather than the extracted factors, that those episodes involve the market pricing something beyond its own average rate expectation, not merely revising that expectation upward. Second, the sign of the gap is informative on its own: a green line running persistently *below* black (2015–16, and the 2025–26 episode developing now) says the model's physical-measure dynamics were, at the time, underpredicting how high or how sustained the tightening would be – while a green line running *above* black through 2021–22 says the opposite, that the VAR's own backward-looking dynamics were slower to price in the credibility of that cycle than the market's forward curve was. This is a useful check on the VAR itself, independent of the term premium application: a model whose physical-measure forecast is persistently biased in one direction within a regime is telling the modeler something about the estimation window, not only about term premium risk.

A finer-grained version of the same exercise – the full forecast path $E_t[r_{t+i}]$ from a small number of individual vintage dates, against realized Selic (BCB SGS 432) rather than the DI curve's own short end – isolates genuine forecast surprises at the moment they occurred, rather than reading them off the smoothed rolling series above.

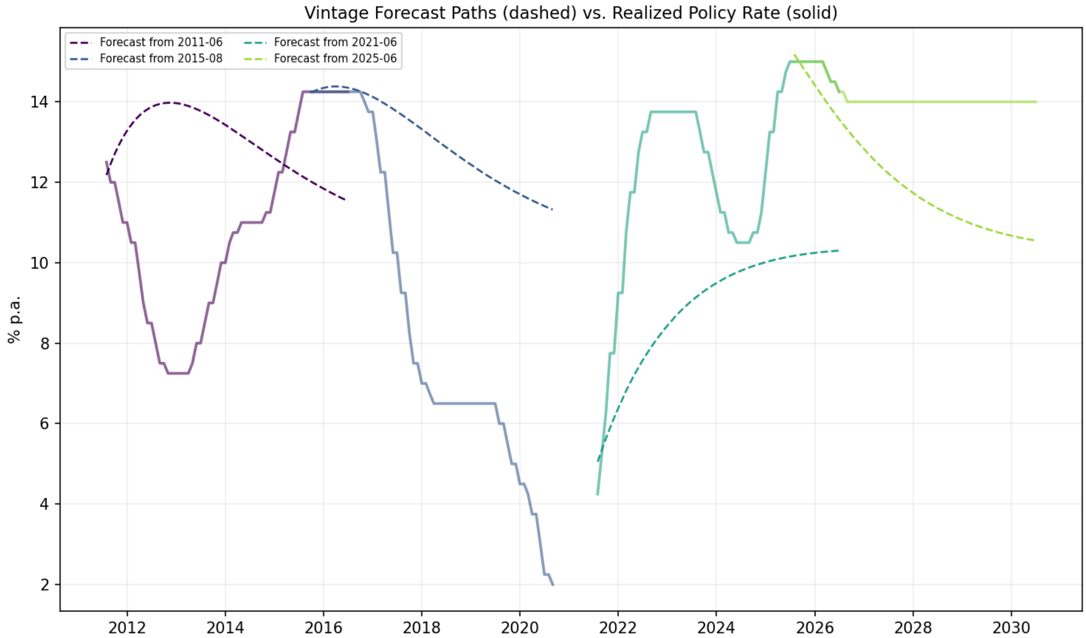

Figure 7: Vintage forecast paths (dashed), $E_t[r_{t+i}]$ iterated forward from four representative dates, against the Selic path that subsequently materialized (solid). Each color is one vintage: the model's forecast as of that date, and what actually happened afterward.

The four vintages tell four different stories, and each matches an episode already discussed elsewhere in this paper. From **2011-06**, the model expected a further gradual decline; Selic instead round-tripped back up into the 2015–16 crisis, a genuine surprise the forecast path entirely misses. From **2015-08**, near the crisis peak, the model expected a slow grind down from the 14% area; realized Selic fell considerably faster, through the 2017–20 easing cycle to 2%, again undershooting the forecast persistently – consistent with the credible, front-loaded disinflation this paper's the companion paper's Section 6 diagnostic reads as a hallmark of the 2021–22-style regime, arriving here a few years early in the data. From **2021-06**, at the start of the post-pandemic hiking cycle, the model's own physical-measure dynamics badly underestimated how far and how fast Copom would move – the forecast path tops out near 10% while Selic actually reached nearly 14%, the single largest miss of the four vintages, and a direct, vintage-level confirmation of the companion paper's Section 6.7's reading of that cycle as one the backward-looking VAR had not yet learned to expect. From **2025-06**, the most recent vintage, the model expects a gradual easing from the 15% peak toward roughly 10.5% by 2030; Selic has already eased slightly faster than that path over the few months of data available so far, though it is too early in this vintage's life to call it a repeat of 2015–16's persistent undershoot rather than noise.

The general pattern across all four vintages is one-sided in a specific way: this VAR's physical-measure forecasts are more often surprised by moves being larger and faster than expected than the reverse, in both directions (the 2011 and 2015 vintages were surprised by faster easing and re-hiking respectively; the 2021 vintage was surprised by faster hiking). A backward-looking linear model built to smooth over regime transitions will structurally lag behind sharp ones; that is a property of the model worth keeping in mind when reading the term premium estimates of the companion paper's Sections 5–8, which are conditioned on this same VAR's forecast path being the correct baseline against which risk compensation is measured.

4.5 The model-implied nominal neutral rate

The VAR's unconditional mean under the physical measure, μ, together with the short-rate equation $y_t^{(1)} = \delta_0 + \delta_1' X_t$ estimated in the companion paper's Section 3, gives a model-consistent steady-state short rate:

$$\bar{r} = \delta_0 + \delta_1' \mu$$

This is the nominal short rate the model's own dynamics say Selic is being pulled toward, absent further shocks – a curve-implied companion to the semi-structural r^* estimates of Laubach and Williams (2003) and Holston, Laubach and Williams (2017), and methodologically closer in spirit to the shadow-rate approach of Wu and Xia (2016), since both read the anchor off the term structure's own dynamics rather than off a macro model's output-gap residual.

As with the real curve in Section 4.6, the textbook closed-form μ is not what is actually used: $\bar{r}$ is computed from the fitted short-rate equation evaluated at each estimation window's own average state, for the same numerical-robustness reason developed in full there. Unlike the real curve, this substitution turns out to matter only modestly for the nominal side – the raw VAR intercept and the average-state definition agree to within about a quarter of a percentage point at the full-sample level, since the nominal curve's persistence, while still close to the unit circle in places (rolling eigenvalue moduli range 0.96–1.07 across the sample), is on average further from a genuine unit root than the real curve's level factor. μ is regime-dependent in exactly the sense the companion paper's Section 4 documents for the factor loadings, so $\bar{r}$ is estimated on an 84-month rolling window (stepped every 6 months) once enough history exists; before that, from 24 months of history at the start of the sample, the window expands rather than waiting for a full 84 months to accumulate, so the series can begin close to the start of the sample rather than only in 2017. The two regimes are visually distinguished in Figure 8 (dashed for the expanding-window portion, solid once the window is full) precisely because the earlier estimates rest on less history and should be read with correspondingly more caution.

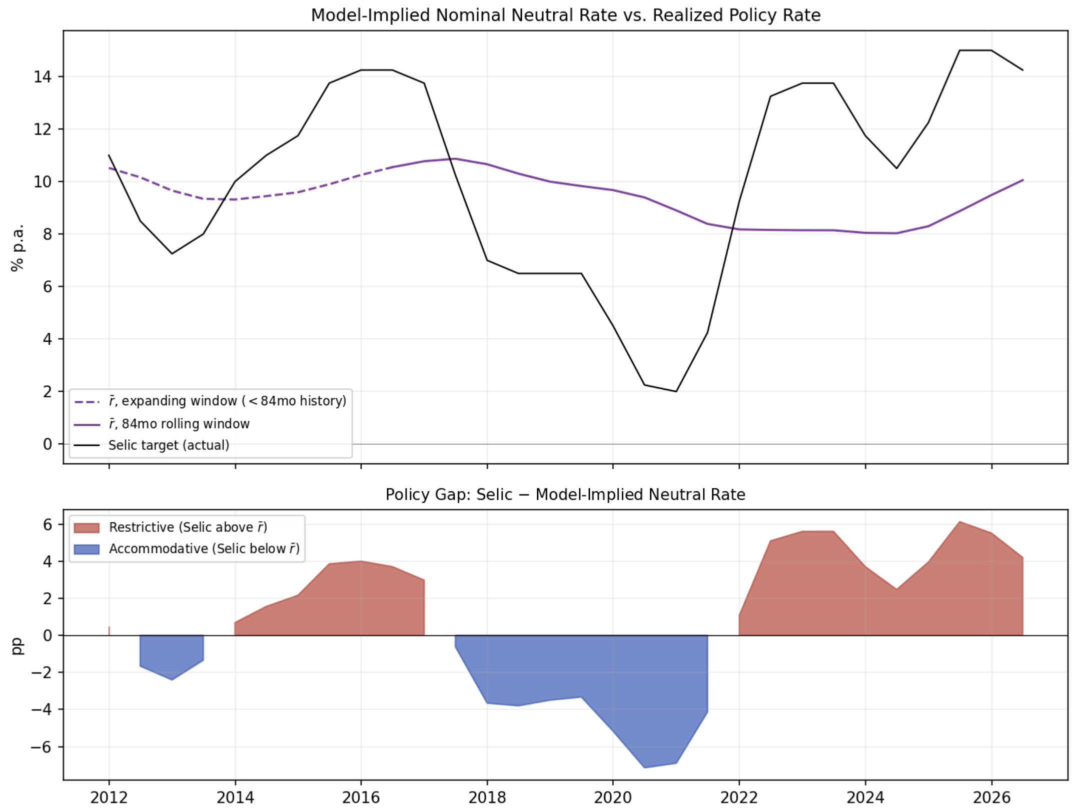

Figure 8: Top: model-implied nominal neutral rate (purple – dashed while the window is still expanding, solid once it reaches the full 84 months) against the realized Selic target (black, BCB SGS 432). Bottom: the policy gap, Selic minus $\bar{r}$ – positive (red) reads restrictive, negative (blue) reads accommodative.

Result. The rolling nominal $\bar{r}$ opens the sample near 10.5% in late 2011, drifts down through 9.3–10.9% across 2012–17, declines further to a trough near 8.0–8.1% through 2022–24, then rises again to 10.05% by mid-2026 – read the way a stance indicator is normally read, Selic above the rolling line signals restrictive policy on the model's own terms, Selic below signals accommodative, and the 2025–26 configuration (Selic at 14–15%, $\bar{r}$ at roughly 10%) reads as clearly restrictive on the model's own terms, consistent with the elevated and rising term premium the companion paper's Section 5 already documents for the same window. The bottom panel makes the same read directly rather than requiring a visual comparison of two lines: Brazil moves through four distinct policy-stance regimes over the sample – mildly accommodative in 2012–13, moderately restrictive into the 2015–16 crisis (peaking near +4pp, the same episode the companion paper's Section 6 reads as level-driven and fiscal in origin), sharply and persistently accommodative through the 2017–21 easing cycle and pandemic response (troughing beyond −7pp in 2020), and restrictive again from 2022 onward, twice approaching +6pp – once in 2023 and again in the current 2025–26 episode. The gap's magnitude and persistence – not any single month's reading, for the same reason argued in the companion paper's Section 5 – is the object of interest; a rolling window with substantial month-to-month overlap by construction produces a smooth series, and should not be read as implying the neutral rate itself moves smoothly in reality.

4.6 The model-implied real neutral rate

A nominal neutral rate is of limited use on its own; the comparison the literature actually cares about is the real rate against a proxy for potential growth or against other economies' own r^* estimates. Rather than deflating $\bar{r}$ by realized IPCA directly – which conflates the neutral rate with whatever inflation happened to print, backward-looking and noisy at the point-in-time level this exercise needs – the real neutral rate is estimated directly off Brazil's own real (NTN-B) curve, using the identical PCA/VAR/BRW/short-rate-equation construction of the companion paper's Section 3, applied to real rather than nominal yields:

$$\bar{r}^{\,\text{real}} = \delta_0^{\text{real}} + (\delta_1^{\text{real}})'\overline{X^{\text{real}}}$$

This is a self-contained estimation on the real curve's own state dynamics – it does not require the nominal DI side, or an inflation anchor, at all, since it never has to separate a breakeven from a real yield in the first place. It is a narrower exercise than the full joint nominal-real (Abrahams-Adrian-Crump-Moench) decomposition used in this author's companion NTN-B work to recover breakeven inflation and the inflation risk premium; reconciling the two – checking that $\bar{r}^{\,\text{real}}$ here plus a model-consistent expected-inflation path recovers $\bar{r}$ from Section 4.5 – is a natural next step once the nominal pipeline is run, but is not needed to answer the real-rate question on its own.

Data and curve construction. The real curve is built from the full history of outstanding NTN-B bonds (34 bonds, 2005–2026), fitting a cross-sectional Nelson-Siegel curve – $y(\tau) = \beta_0 + \beta_1 \frac{1-e^{-\tau/\lambda}}{\tau/\lambda} + \beta_2 \left(\frac{1-e^{-\tau/\lambda}}{\tau/\lambda} - e^{-\tau/\lambda} \right)$ – separately at each month-end, then evaluating the fitted curve at nine target maturities (1, 2, 3, 4, 5, 7, 10, 15, and 20 years) to build a monthly panel, 2010–2026 (200 observations). Time-to-maturity is used for τ, computed directly from each bond's known maturity date – a standard and defensible convention for a Nelson-Siegel fit in its own right. The fit is tight throughout – a mean absolute cross-sectional pricing error of roughly 6 basis points in a typical month, rising to a still-small 20–25bp in a handful of recent, thinner months.

An estimation subtlety the level factor's persistence forces. PCA on this real curve recovers a level factor carrying 97.0% of variance – even more concentrated than the nominal curve's level factor in the companion paper's Table 1 – with a slope and curvature factor carrying 2.5% and 0.4%. Fitting the bias-corrected VAR(1) to this state vector, the level factor's persistence comes back with an eigenvalue modulus of approximately 1.01: a near unit root. This matters for the steady-state calculation specifically, not just as an abstract stationarity concern: the textbook closed-form unconditional mean of a VAR(1), $(I - \Phi)^{-1}\mu$, requires $(I - \Phi)$ to be safely invertible, and here it is nearly singular ($\det(I - \Phi) \approx -4 \times 10^{-4}$), so that formula is numerically unreliable – small, statistically insignificant differences in the estimated μ translate into large, sign-flipping swings in the implied steady state. This is itself a substantive finding rather than a nuisance to route around silently: over this sixteen-year sample, the real curve's level factor is too persistent for a formal infinite-horizon anchor to be statistically identified, consistent with a slow-moving structural or fiscal-risk-premium component in the level of Brazilian real rates rather than a cleanly mean-reverting process. $\bar{r}^{\text{real}}$ is therefore defined as the fitted short-rate equation evaluated at the model's own average state over the estimation window, $\overline{X}^{\text{real}}$ – robust to the near-singular inversion, and a direct answer to the practical question of what level the model's own dynamics have been centered on, without extrapolating to an ill-identified infinite horizon.

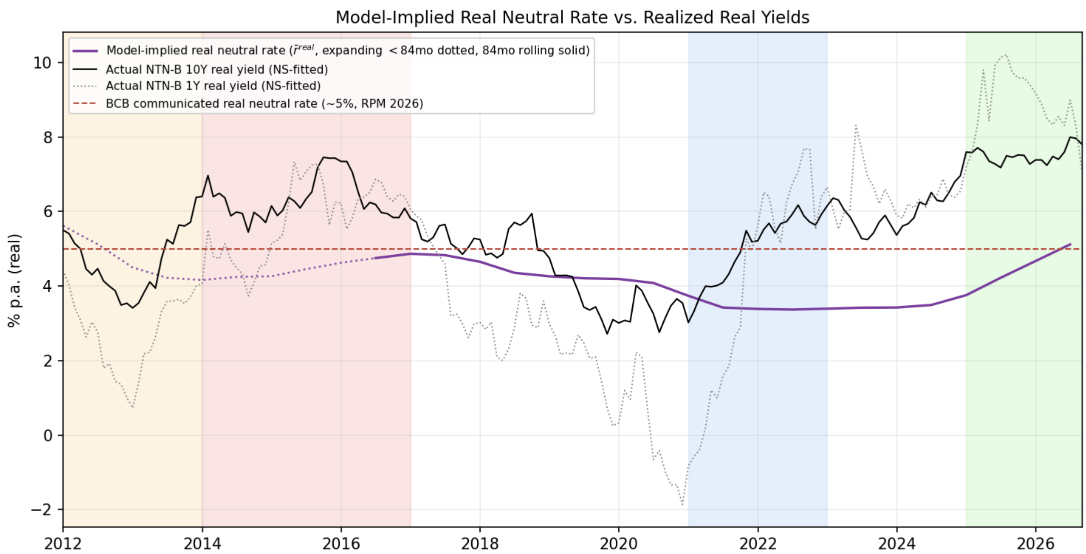

Figure 9: Model-implied real neutral rate (purple, expanding window from 24 months of history up to 84, dotted, then a fixed 84-month rolling window, solid, matching the same convention used elsewhere in this paper) against the NS-fitted NTN-B 1-year and 10-year real yields, and BCB's own communicated real neutral rate of approximately 5% (Relatório de Política Monetária, 2026).

Result. The full-sample (2010–2026) estimate is $\bar{r}^{\,\mathrm{real}} \approx 4.76\%$. The expanding-window portion of the series (2012–2017, before 84 months of history accumulates) starts near 5.6% and drifts gently down to roughly 4.8% – read with the usual caveat that an expanding window is progressively better-established but never as reliable as the full 84-month rolling estimate that follows it. The rolling 84-month window itself shows the estimate rising steadily from roughly 3.4% in 2021–22 to 5.1% by mid-2026, with the pace of the rise itself accelerating over the final two years – roughly a quarter-point per half-year through 2023, closer to half a point per half-year by 2025–26. This acceleration is mechanical, not a sign of growing estimation instability: the eigenvalue modulus of the rolling Φ stays flat (1.02–1.06) across the entire period, so it is not the near-singular-inversion problem described above resurfacing. A rolling window's average rises faster once the regime it is averaging over has itself shifted enough that older, lower-real-rate months rolling out of the window are being replaced by newer, meaningfully higher-real-rate months – Brazil's own real yields have risen substantially since 2024 as the hiking cycle progressed, and the rolling average is simply catching up to that shift, arriving at a level within a tenth of a point of BCB's own figure just as it finishes doing so. And the most recent reading sits within roughly 0.1 percentage points of BCB's own communicated real neutral rate of approximately 5%, stated in its 2026 Relatório de Política Monetária. That agreement is a genuine cross-check, not a built-in one: nothing in this estimation targets or references the BCB figure, which comes from an entirely different (semi-structural, survey-informed) methodology. The rolling series' own trajectory is informative independent of the level match: a neutral rate estimated to have risen roughly 1.7 percentage points since 2021–22 is consistent with the fiscal-risk-premium reading this research programme has argued for elsewhere – rising compensation for holding Brazilian duration mechanically raises what the curve says the real rate is anchored around, distinct from, but complementary to, the term-premium-based fiscal dominance diagnostic of the companion paper's Section 6.

This estimation used only the real (NTN-B) curve; a self-contained steady state doesn't need the joint machinery. A genuine joint reconciliation, however, does – and is now attempted directly, rather than deferred.

4.7 The joint reconciliation

The simplest way to reconcile a nominal and a real neutral rate is Fisher subtraction: $\bar{r} - \bar{r}^{\text{real}} \approx$ expected inflation. That is what a first pass at this section did, using Section 4.5's nominal $\bar{r}$ (10.05%) and this section's real $\bar{r}^{\text{real}}$ (roughly 5.1%) to imply an expected-inflation wedge near 4.9pp – a real number, but built from two curves estimated entirely independently of each other, with no guarantee their respective persistence, sample windows, or factor structures are mutually consistent. The genuine test of that number is whether it survives being estimated inside a single, jointly-fit model instead.

It was refit that way, using the same `pyacm` implementation of Abrahams, Adrian, Crump and Moench (2016) that the companion paper's Section 8.4 uses for the breakeven, jointly on Brazil's own nominal (DI) and real (NTN-B) curves rather than the US pair – three nominal factors, two orthogonalized real factors, and the liquidity factor already used in the companion paper's Section 8.4, six factors in total. One thing this joint refit ruled out rather than confirmed: the naive steady-state Fisher gap, redone with the joint model's own δ_0, δ_1, μ, produces an implausible nominal $\bar{r}$ under 1% – the six-factor state vector carries more near-unit-root persistence than either single-curve model on its own (three of six eigenvalue moduli exceed 0.90), and the same numerical fragility already flagged for the real curve's own steady state in isolation is worse, not better, in the higher-dimensional joint system. That number is not reported as a finding; a steady-state Fisher gap from this joint model isn't a reliable object, full stop, regardless of which state-averaging convention is used.

What *is* reliable, because it doesn't route through a steady state at all, is the model's own expected-inflation equation – a direct output of the joint estimation, not backed out by subtraction. Figure 10 plots it at three horizons.

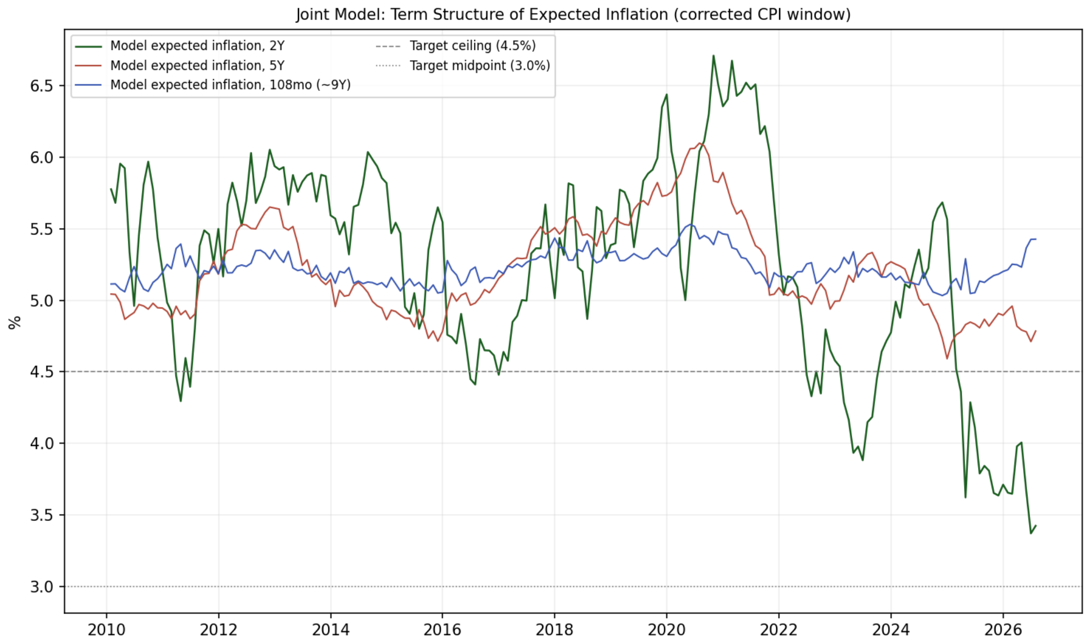

Figure 10: Joint-model expected inflation at 2Y, 5Y, and roughly 9Y (the longest point common to both curves), 2010–2026. The target ceiling (4.5%) and midpoint (3.0%) are Brazil's inflation-targeting band.

As of July 2026 – the most recent month with a genuine realized IPCA print, rather than one manufactured by forward-filling past it – the model's own expected inflation reads 3.4% at 2Y, 4.8% at 5Y, and 5.4% at the longest common point (roughly nine years) – close to target at the short end and still elevated further out, not uniformly above the ceiling throughout (full-sample means: 5.2%, 5.2%, 5.2% at the same three horizons). This is the more defensible reconciliation the two-curve Fisher approximation was standing in for:

expected inflation that has spent most of sixteen years running above target, spiked through the 2020–22 inflation surge (2Y expected inflation reached 6.7% in October 2020), and has since come back down, with the short end now reading closer to target than the medium and long end. It is also directly comparable to the companion paper's Section 8.4's breakeven table by construction, since the same joint model decomposes breakeven as expected inflation plus the inflation risk premium: the small positive IRP already reported there (roughly 0.5–0.8pp at the vertices reported) accounts for most, though not all, of the gap between this section's expected-inflation reading and the companion paper's Section 8.4's higher breakeven levels, with the residual most likely attributable to the real-curve construction differences already flagged in Section 4.6 and in this run's own documentation.

A fifth instance of the same estimation theme, caught before it reached the chart. An earlier version of this exercise reported a 2Y expected-inflation reading of 5.4% as of "August 2026" – a genuine outlier once checked (a 5.3-standard-deviation single-month move, five times the next-largest in sixteen years of history). The cause was neither the joint model nor the output-gap logic this time, but a data-alignment error one step upstream: the nominal and real curve panels feeding this model extend through August 2026 (market-based, essentially real-time), while realized IPCA – BCB releases inflation prints with roughly a ten-day lag – was only available through July. The original construction intersected three of the four required inputs and forward-filled the fourth (CPI) to match, silently manufacturing a zero-inflation August that had not happened yet rather than surfacing the mismatch. Restricting the joint estimation to the window where all four inputs are genuinely observed, rather than extending one past its own data, both removes the outlier and shifts the fitted parameters enough to move mid-sample values as well (the 2020 peak reported above, 6.7%, is the corrected figure; an earlier pass reported 7.7% from the contaminated fit) – a reminder that a data problem in a jointly-estimated model does not stay confined to the date it originates at.

4.8 Situating these estimates in the structural-versus-market-implied debate

The companion paper's Section 2 sets up two traditions for estimating a neutral rate: a structural one that filters it out of a macro model's IS and Phillips-curve relationships, and an asset-price one that reads it off the yield curve directly, of which this paper's Section 4 is an instance. The results above are not just an application of the latter tradition – two of them speak to the methodological debate between the two directly, on their own terms, rather than merely sitting on one side of it.

The first is Section 4.6's near-unit-root finding for the real curve's level factor, and the reason the textbook $(I - \Phi)^{-1}\mu$ steady state was abandoned there in favor of an average-state definition. That finding is, independently arrived at, the same phenomenon Bauer and Rudebusch (2020) report for the US Treasury curve: letting a long-run trend be directly spanned by observed yields produces implausibly volatile, numerically fragile estimates, which is precisely why their preferred specification treats the trend as an unspanned latent factor rather than a closed-form function of the state. This paper reaches the same conclusion from a different diagnostic (an eigenvalue check and a near-singular $\det(I - \Phi)$, rather than an observed blow-up in the trend series itself) on a different curve, a different country, and a different modeling choice (three-factor real ACM rather than a dedicated trend-spanning term structure model). That two independent routes into the same estimation problem land on the same fix is a small but genuine piece of evidence that the fragility is a structural feature of extracting a slow-moving trend from a persistent process on a finite sample – present in the asset-price tradition generally, not an artifact specific to either paper's implementation.

The second is Section 4.7's discovery that this same fragility gets *worse*, not better, once the nominal and real curves are estimated jointly rather than separately: the six-factor system's persistence sits closer to the unit circle than either three-factor single-curve system on its own. This is the reverse of what a reader might expect – more data and more structure imposed should, on priors, produce a more stable estimate, not a less stable one – and is worth flagging explicitly for anyone extending the joint nominal-real (Abrahams-Adrian-Crump-Moench) framework to a steady-state or trend-extraction application rather than the term-premium and breakeven decomposition it was originally built for.

Where this paper's results speak most directly to the reconciliation side of the debate is the agreement, unforced, between Section 4.6's curve-implied real neutral rate (about 5.1% as of mid-2026) and BCB's own communicated figure (approximately 5%, from its 2026 Relatório de Política Monetária) – a semi-structural, survey-informed estimate arrived at by an entirely different route, in the spirit of, though far narrower in scope than, Del Negro, Giannone, Giannoni and Tambalotti's (2017) attempt to make a structural model's r^* match what asset prices independently imply. One country, one moment in time, and a single point estimate is a data point, not a test – the natural way to actually test whether curve-implied and BCB-communicated readings track each other, rather than merely agreeing once, is the same cross-country extension of this paper's methodology already flagged as future work below: applying the identical construction to the other emerging markets in the mandate's universe and checking whether the same order of agreement holds there too.

5 Conclusion and Further Work

This paper states the choice of neutral-rate estimation method as a genuine three-way spectrum – structural, bridge, and reduced-form – rather than picking one silently, and tests where Brazil's own curves actually sit on it rather than assuming an answer. Whether Ang and Piazzesi's (2003) bridge is worth adopting is tested directly, maturity by maturity, using their own twelve-lag macro specification rather than a simplified approximation of it that this paper's own earlier pass initially found nothing significant with; done properly, inflation and the output gap significantly forecast excess returns on the nominal curve from 2 through 24 months and on the real curve from 24 through 60 months. The same no-arbitrage machinery, reused without modification from a companion paper, gives a model-consistent nominal neutral rate reading current Brazilian policy as clearly restrictive, and a real neutral rate landing within a tenth of a percentage point of BCB's own communicated figure – an agreement arrived at by an entirely different route, not built into the estimation. A specific numerical fragility in the real curve's own steady state is documented in full: near-unit-root persistence that makes the textbook closed-form steady-state formula numerically unreliable, a symptom already reported independently for the US Treasury curve (Bauer and Rudebusch, 2020), and one that gets worse rather than better once the nominal and real curves are estimated jointly.

Two extensions are natural next steps. First, and most directly tied to the paper's own headline result, the twelve-lag bridge specification tested in Section 4.3 used the output gap constructed via a joint Kalman filter matching Areosa's (2008) own methodology as closely as this paper's own data allowed; validating that construction more fully against BCB's own internal, unpublished output-gap series – rather than only its published Areosa-method line, which this paper's own construction already tracks reasonably closely – would tighten the evidentiary basis under the bridge test's own state variable. Second, the joint nominal-real reconciliation of Section 4.7 replaced an earlier two-curve Fisher approximation with a direct refit of the same joint (Abrahams-Adrian-Crump-Moench) model used in the companion paper's own breakeven decomposition, on Brazil's own nominal and real curves; what that joint refit could not do reliably was reproduce a steady-state nominal or real rate of its own, since the higher-dimensional, six-factor joint state carries more near-unit-root persistence than either single-curve model alone. Reconciling Section 4.7's expected-inflation reading precisely against the companion paper's own breakeven table, beyond the inflation-risk-premium accounting already offered there as a partial explanation, is the piece of this exercise still open.

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. Ang, A., & Piazzesi, M. (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, 50(4), 745–787.

4. Areosa, M. (2008). Combining Hodrick-Prescott filtering with a production function approach to estimate output gap. *Banco Central do Brasil Working Paper Series* No. 172.

5. Bansal, R., & Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4), 1481–1509.

6. Bauer, M. D., & Rudebusch, G. D. (2020). Interest rates under falling stars. *American Economic Review*, 110(5), 1316–1354.

7. 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.

8. Campbell, J. Y. (1993). Intertemporal asset pricing without consumption data. *American Economic Review*, 83(3), 487–512.

9. Cochrane, J. H., & Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1), 138–160.

10. Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, 2017(1), 235–316.

11. Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies*, 24(9), 2895–2934.

12. Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4), 937–969.

13. Holston, K., Laubach, T., & Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, 108, S59–S75.

14. Laubach, T., & Williams, J. C. (2003). Measuring the natural rate of interest. *Review of Economics and Statistics*, 85(4), 1063–1070.

15. Lucas, R. E., Jr. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

16. Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2), 145–161.

17. Mollica, M. A. (2026a). Brazilian Interest Rates: Term Premium, Fiscal Dominance, and the Real Curve. Working paper.

18. Orphanides, A. (2001). Monetary policy rules based on real-time data. *American Economic Review*, 91(4), 964–985.

19. Rudebusch, G. D., & Swanson, E. T. (2008). Examining the bond premium puzzle with a DSGE model. *Journal of Monetary Economics*, 55, S111–S126.

20. Rudebusch, G. D., & Swanson, E. T. (2012). The bond premium in a DSGE model with long-run real and nominal risks. *American Economic Journal: Macroeconomics*, 4(1), 105–143.

21. Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3), 401–421.

22. Wu, J. C., & Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2–3), 253–291.

Appendix: Estimation Detail Specific to This Paper

The complete, step-by-step derivation of the no-arbitrage curve machinery this paper reuses – state extraction, the Bauer-Rudebusch-Wu correction, the ACM risk-price recovery, and the resulting affine recursion – is in the companion paper's own appendix (Mollica, 2026a, Appendix A.1–A.4), which allows independent reproduction from the underlying DI and NTN-B data alone. What follows is specific to this paper's own additions: the output gap, the bridge state, and the validation items that apply to this paper's own results rather than to the shared machinery.

B.1 The output gap

Brazil's output gap is estimated via a joint Kalman filter matching Areosa's (2008) own BCB-preferred specification: natural employment, natural capacity utilization, and potential output estimated simultaneously, with the production-function identity $y_t = y_t^n + (1-\alpha)(e_t - e_t^n) + \alpha(c_t - c_t^n)$ imposed as an exact cross-equation constraint rather than a weighted combination of two independently-filtered series computed afterward. Employment and capacity utilization follow random walks (Areosa's lower-order "$r = 1$" filter, chosen specifically because it is less prone to end-of-sample distortion than a standard Hodrick-Prescott trend); potential output follows the usual local-linear trend. The capital share α is fixed at Brazil's own average labor share over the sample (Penn World Table, $1 - \alpha \approx 0.555$). The resulting series is validated against BCB's own published Areosa-method line (same order of magnitude, same broad decline toward zero over the shared window) and is subject to a documented real-time revision pattern: re-estimating on progressively truncated samples shows the same quarter's own reading revised down by roughly 0.2–0.3 percentage points as later data is added, consistent with the wider real-time output-gap literature (Orphanides, 2001).

B.2 The contemporaneous decomposition: a different question, not reported in the main text

The bridge state stacks year-on-year IPCA inflation, the output gap of B.1 (interpolated to monthly), and the three latent yield factors from the companion paper's own state, orthogonalized against the two macro factors. Before the genuine ex-ante forecast reported in Section 4.3, a different, contemporaneous test was run on this same state – Step 2 of the ACM procedure itself: excess returns regressed on the state's own realized shock v_{t+1} plus the macro factors, exactly as ACM's own risk-price recovery requires. This is a within-period co-movement question, not a forecasting one, and its baseline R^2 is mechanically high for that reason (bond returns nearly definitionally track the same-period factor shock). Conflating this with the genuine ex-ante question was an error an earlier pass at this exercise made; it is reported here, rather than in the main text, precisely to keep that distinction clean.

	R^2 latent-only	R^2 +macro	F-stat	p-value
Nominal short rate	0.9913	0.9930	20.45	$<$0.0001
Real short rate	0.9961	0.9962	3.09	0.048

Table 3: Short-rate equation, latent-only vs. latent+macro state. Same-period fit test, not a forecasting claim.

Maturity (mo)	R^2 latent-only	R^2 +macro	ΔR^2	F-stat	p-value
2	0.5747	0.6279	0.0532	11.65	<0.0001
9	0.9255	0.9302	0.0046	5.41	0.0053
18	0.9866	0.9873	0.0007	4.67	0.0107
24	0.9955	0.9956	0.0001	0.95	0.391
60	0.9969	0.9969	0.0000	0.47	0.629
108	0.9975	0.9976	0.0001	2.42	0.092

Table 4: Nominal curve, contemporaneous decomposition (ACM Step 2, includes v_{t+1}), selected maturities. Baseline R^2 is high throughout because the regression includes the same-period state shock by construction. Not comparable to the ex-ante R^2 in Section 4.3's tables.

Maturity (mo)	R^2 latent-only	R^2 +macro	ΔR^2	F-stat	p-value
24	0.9885	0.9885	0.0001	0.70	0.500
60	0.9987	0.9987	0.0000	0.48	0.621
120	0.9972	0.9972	0.0000	0.13	0.876

Table 5: Real curve, contemporaneous decomposition, selected maturities. No maturity shows a significant improvement in this specification – the reverse of Section 4.3's ex-ante finding for the real curve, since the two answer different questions.

The genuine ex-ante forecast reported in Section 4.3 drops v_{t+1} entirely and lags the macro variables to the start of the holding period, with twelve lags of each variable, matching Ang and Piazzesi's (2003) own preferred "Macro Lag Model" rather than a single-lag approximation of it.

B.3 A validation checklist specific to this paper

1. The rolling nominal and real $\bar{r}$ series (Section 4.5, 4.6) should sit within a plausible range of realized Selic and realized real yields respectively over the full sample – each is, by construction, a smoothed anchor for the policy rate, not an independent forecast. A window length shorter than the one used for the headline reading (84 months, expanding from 24) will track recent realized rates too closely to represent a genuine steady state; this was caught and corrected during this paper's own construction, and the check is recorded here so it is not missed again.

2. The bridge-versus-latent-only comparison (Figures 3–4) should show a small mean gap once evaluated correctly: fit once on the full common window, then evaluate at a genuinely different sub-window's own mean. Evaluating at the same window used for fitting returns identical values by construction – an OLS identity, not a finding – and is not a valid test.

3. The twelve-lag bridge forecast test (B.2) should show results stable across neighboring lag lengths (nine through thirteen, say). A result that appears only at one specific lag count and disappears at adjacent ones is more likely a specification artifact than a genuine relationship.