

Mexico: Term Premium, the Real Curve, and Neutral Rates

The Brazil Paper 1 / Paper 2 machinery applied to Mexico — results and charts only.

WHAT IS HERE

Everything in Paper 1 (term premium, factor diagnostic, US control, real curve, breakeven, basal premium) and the curve-based half of Paper 2 (predicted vs realized path, vintages, nominal and real neutral rate, expected inflation) has been run for Mexico, together with the output gap and the Ang–Piazzesi bridge tests. The real curve is built bond by bond from 32 individual Udibonos, anchored at Banxico's ex-ante real policy rate. Two real-curve models carried over from Brazil — a standalone Udibono ACM and a joint Mbono–Udibono RealACM — fail when checked month by month against Banxico's surveys of policy-rate and inflation expectations (Section 5). The headline real term premium, breakeven split and neutral rates therefore come from a new **survey-disciplined model** (Section 6): the policy path from the swap curve, a separate factor for the Mbono–swap spread, and dynamics estimated jointly with the surveys.

HEADLINES

- **Nominal term premium is high and rising:** 60m at 1.92pp (Aug), 2.04pp mid-September on ACM/BRW, about 1.3pp above its sample mean and the highest since 2014–15; the survey-disciplined model puts it at 1.77pp.
- **The rise is slope-driven, not level-driven** — a composition neither of the two Brazil calibration episodes shows; the Brazil diagnostic should not be carried over.
- **US link is short-run, not long-run** — the reverse of Brazil. Month-to-month US term premium moves pass through about one for one (R^2 0.2 on ACM/BRW, 0.33 on the survey-disciplined premium); no cointegration on either curve. On the real curve the short-run pass-through rises with maturity (β 0.18 at 5Y to 0.55 at 20Y).
- **The last 12 months are Mexico-specific:** the 60m premium rose 0.87pp (ACM/BRW) / 0.70pp (survey-disciplined) against a US-implied -0.09 / $+0.06$ pp. Since 2023 the real premium rose on Mexican grounds at 5–10Y; the 20Y followed US real rates.
- **Real term premium (survey-disciplined model — the headline):** 0.94 / 1.28 / 1.77 / 2.25 / 2.36pp at 2 / 3 / 5 / 10 / 20Y; about 0.4–0.5pp of it is the issuer spread Udibonos share with Mbonos (net: 0.55 / 0.82 / 1.22 / 1.77 / 2.10). Well above its sample mean at every maturity; the current rise shows no market dislocation.
- **Breakeven 4.2–4.8%** (2Y–20Y, Mbono minus Udibono) = expected inflation 3.95% $\rightarrow$ 3.55% + inflation risk premium $+0.22 \rightarrow +1.27$ pp. Expected inflation sits on Banxico's surveys out of sample (0.10pp error on the 1–4 and 5–8-year surveys).
- **Why a new model:** the standalone Udibono model essentially projects today's real rate forward (correlation 0.05 with a survey-anchored premium; -0.5 pp in 2023–24 against $+0.6$); the joint model's policy path is about 1pp above the survey today. The Mbono–swap spread is persistent and regime-like, but it is *not* what breaks the joint model.
- **The survey-disciplined model passes out of sample:** re-estimated each year 2016–2025, its expected policy path misses the next year's surveys by 0.24–0.31pp (12–36m), against 0.58–1.20 without the surveys and 0.45–0.49 for the joint model even in-sample.

- **Macro predictability splits by curve, differently from Brazil:** 12 lags of inflation and the output gap add nothing on the nominal curve ($p \geq 0.86$ at every maturity) but forecast real-curve excess returns at 2–7Y (p 0.004–0.047), under either gap. Brazil found both.
- **Output gap negative** (Banxico -1.2% in 2026Q2; our GDP-based baseline tracks Banxico’s series with correlation 0.99). The Areosa labor-and-capacity construction does not transfer to Mexico (+1.7%).
- **Neutral rates:** trailing-average $\bar{r}$ 8.0% nominal / 3.9% real; the survey-disciplined 5y5y is 5.74% nominal / 2.24% real, inside Banxico’s ranges (4.8–6.6% and 1.8–3.6%). With the target at 6.50%, the forward-looking measure reads policy as mildly restrictive (+0.8pp); the trailing one reads it as accommodative (–1.5pp).

A DEFINITIONAL ISSUE THAT ALSO APPLIES TO BRAZIL PAPER 2

With the short-rate equation fit on the same window it is evaluated on, $\bar{r} = \delta_0 + \delta_1 \cdot \text{mean}(X)$ equals the window mean of the 1-month rate exactly (OLS identity; verified to $1e-14$). The BRW mean-preserving intercept makes the closed-form $(I-\Phi)^{-1}\mu$ identical as well. Paper 2’s nominal 10.05% and real 5.1% are therefore the 84-month averages of Brazil’s 1-month nominal and real rates. Mexico is reported on the same definition for comparability, alongside the model 5y5y forward.

1 • DATA AND CURVE CONSTRUCTION

The nominal curve: TIE28 swaps spliced into F-TIE

TIE28 swap quotes stop on 31 Dec 2024; F-TIE OIS starts 5 Sep 2024. The four-month overlap gives a tenor-specific basis. TIE28 history is shifted down by it, and native F-TIE quotes are used from the switch date. The whole curve is on an overnight basis, anchored to the Banxico target. Zero rates are bootstrapped on 28-day coupons (policy rate as the period-1 anchor) and interpolated flat-forward to a 1–120 month grid. For the 20-year decomposition in Section 6 the same bootstrap adds the 20Y node (21H, quoted from Nov-2011) and extends the grid to 240 months; below 10Y the two grids differ by less than 2.5bp.

Code	Tenor (yrs)	Overlap days	Basis (bp)	SD (bp)
C	0.2	78	26.0	2.9
1A	1.0	80	28.0	2.5
2B	2.0	81	27.5	2.0
3C	3.0	80	25.9	2.1
4D	4.0	79	26.0	2.0
5E	5.0	79	25.5	1.3
7G	7.0	79	25.1	1.5
10J	10.0	80	24.6	1.6
16C	15.0	73	23.5	2.5
21H	19.9	72	24.6	2.1
32F	29.9	75	23.4	3.5

Table A. TIE28 - F-TIE swap basis over the Sep-Dec 2024 overlap. The TIE28 fixing sat 22.7bp above the target over the same window.




Figure A. Basis by tenor (left) and the spliced curve (right). No visible jump at the switch date.

Bond curves and the US store

Mbono zero curve: Nelson–Siegel with a per-date λ search on generic yields (median fit error 7bp), 1-month point anchored at the Banxico target. **Udibono real curve: a cross-sectional Nelson–Siegel fit across every Udibono outstanding on each day** — 35 bonds identified via SRCH (14 outstanding, 21 matured), 32 with Bloomberg history, those with at least one year to maturity; median 9 bonds per day, median fit error 8.7bp (Brazil’s NTN-B curve: 7bp). Daily quotes are too sparse before May 2009 (usually a single bond), so the real curve starts in mid-May 2009 and the models in June 2009. The real 1-month point is Banxico’s own ex-ante real policy rate throughout: target minus the survey mean of 12-month expected inflation (Banxico SIE, monthly since 1999; December is not surveyed from 2020, so November carries). Real GDP: the Banxico SIE level series supplied (base 2013) ends in 2023Q1; its quarterly growth matches the Bloomberg series used here (correlation 0.996), so the chained Bloomberg level is kept, being current. US series come from the shared us_core store: NY Fed ACM for the nominal control; the joint Treasury–TIPS model with the TIPS fitting-error liquidity proxy for the real control. The VIX version is kept only as a robustness check: its levels are implausible (US 5Y real TP $\approx -5\%$) and it explodes on the 20Y grid. The US pull used here is the 22 Sep 2026 refresh (us_core v2; own 5-factor ACM correlates 0.86–0.92 with the NY Fed series). **Surveys.** Banxico’s survey of private-sector specialists, monthly from Aug-2010: funding-rate expectations for the end of the current and next nine quarters and 28-day Cete expectations for the end of the current and next three years (mean, median, quartiles and cross-sectional standard deviation), plus the inflation expectations series above. Each month’s medians are turned into an expected average policy rate over the next 12, 24 and 36 months (linear between quarter-ends, Cete year-ends beyond the funding-rate horizon shifted by the funding–Cete gap, flat after the last point; median reach 30 months).

Factor stability across regimes

Regime	PC1 (level)	PC2 (slope)	PC3 (curvature)	Cumulative
2010–12 (post-GFC calm)	88.2%	7.9%	1.4%	97.5%
2016–17 (US election / peso)	93.9%	3.8%	1.7%	99.4%
2022–23 (hiking cycle)	87.9%	7.1%	2.5%	97.5%
2025–26 (current)	94.6%	3.6%	0.9%	99.1%

Table 1. Variance share by regime, daily changes of the F-TIIE zero curve, 13 tenors (3m–10y).

Level concentration is highest now (94.6%) and in 2016–17, but does not rise monotonically: the 2022–23 hiking cycle is as dispersed as the calm period. Loading shapes are stable except a front-loaded level factor in 2016–17.

Mexico F-TIIE swap curve: PCA loadings across regimes (daily changes)




Figure 1. PCA loadings across the four regimes.

The term premium

	1Y	2Y	3Y	5Y	7Y	10Y
Sep-26 (to 15th)	0.34	0.88	1.36	2.04	2.54	3.21
Aug-26	0.31	0.80	1.25	1.92	2.43	3.09
Sample mean	0.00	0.09	0.24	0.62	1.03	1.67

Table 2. ACM/BRW term premium, pp. 203 month-ends, Oct 2009–Aug 2026; September projected with the same parameters.

Model: 3 PCs (95.4%, 4.2%, 0.4%), BRW $\delta = 1$, largest Q eigenvalue 0.9987. Fit error 4–7bp at 3–10Y but 13–18bp at 3–12m (the policy-rate kink). 5Y zero 8.29% = expected path 6.26% (target 6.50%) + ~2pp term premium: the market is not pricing hikes. The survey-disciplined model of Section 6 gives a slightly lower 5Y premium (1.77pp) and a 6.26% path; this section keeps the ACM/BRW estimate for comparability with Brazil Paper 1.




Figure 2. F-TIIE zero yields (top) and term premium (bottom), 6/36/60 months, with episode shading.

Episodes (60m): taper 2013 peak 4.4pp (least reliable — see BRW sensitivity); 2016–17 roughly zero despite heavy hiking; 2022–23 deeply negative (trough –3.1pp), the credible-tightening pattern also seen in Brazil 2021–22; 2024 election/judicial reform still negative but turning.

Section 6 counterpart: factor composition

The current rise is carried by the slope factor (+2.5pp at 60m in September) with the level factor pulling the other way (−0.8pp). Brazil’s diagnostic reads level-driven as fiscal and curvature-driven as credible tightening; Mexico today is neither — consistent with Banxico easing to 6.50% while the long end does not follow. The factors are estimated separately for Mexico, so the Brazil two-case reading is not transferable.




Figure 3a. 60-month term premium by factor (time-varying part; constant 0.61pp).




Figure 3b. BRW vs OLS VAR. Today robust (1.92 vs 1.84); 2013-14 (BRW about 2pp higher) and 2022-24 (BRW about 2pp lower) depend on the correction.

A short-run link, not a long-run one

Mexico TP	EG p	KPSS resid p	α (t)	$\Delta \beta$ on US 2y; 10y (t)	ΔR^2
ACM/BRW 36m	0.169	0.10	-0.096 (-3.4)	-0.33 (-0.7); 0.93 (3.1)	0.22
ACM/BRW 60m	0.131	0.10	-0.101 (-3.7)	-0.64 (-1.0); 1.32 (3.1)	0.20
Survey-discipl. 36m	0.083	0.10	-0.106 (-2.7)	0.61 (3.1); 0.24 (2.1)	0.32
Survey-discipl. 60m	0.092	0.01	-0.106 (-2.8)	0.54 (2.3); 0.46 (3.2)	0.34

Table 3. Engle-Granger (MacKinnon p), KPSS on the levels residual, ECM speed of adjustment, and the differenced regression on NY Fed ACM 2y/10y (Newey-West, 6 lags). ACM/BRW: Oct-2009-Aug-2026; survey-disciplined (Section 6): Mar-2011-Aug-2026. US data: us_core pull of 22 Sep 2026.

On the ACM/BRW premium Engle-Granger does not reject a unit root in the residual while KPSS does not reject stationarity. Paper 1's standard requires both to agree, so cointegration is **inconclusive**, and the significant α is not valid inference without it. On the survey-disciplined premium the answer is the same (EG p 0.08–0.09) for a different reason: that series is itself stationary (ADF p 0.01–0.02, KPSS does not reject), because the survey removes the slow-moving component the ACM path produced, and a stationary Mexican premium cannot be tied in the long run to a persistent US one. The short-run link is strong on both: β on ΔUS 10y of 0.93–1.32 ($t \approx 3.1$) and R^2 0.20–0.22 on ACM/BRW; on the survey-disciplined premium both US maturities enter positively, the betas sum to about one (t 2.1–3.2) and R^2 rises to 0.32–0.34 — against ≈ 0.002 for Brazil. Consistent with a market heavily held by foreign investors, and the reverse of Brazil (long-run tether, no short-run link).

Window (60m TP)	Δ Mexico	US-implied	US share	Δ Mexico	US-implied	US share
Since Apr-2024 trough	+3.15	+0.82	$\approx 26\%$	+0.63	+0.54	$\approx 86\%$
Since Dec-2024	+0.76	+0.28	$\approx 37\%$	+0.13	+0.19	all
Last 12 months (since Aug-2025)	+0.87	-0.09	none	+0.70	+0.06	$\approx 8\%$

Table 4. US-implied change (differenced betas $\times$ change in NY Fed ACM 2y/10y) over stated windows. Columns 2–4: ACM/BRW premium; columns 5–7: survey-disciplined premium.

The two measures tell the same story about the last year and a different one about 2024. On both, the rise over the last 12 months is Mexican: +0.87pp (ACM/BRW) and +0.70pp (survey-disciplined) against a US-implied -0.09 and +0.06pp. The 3pp rebound from the April 2024 trough on ACM/BRW is mostly the hiking-cycle path overshoot of Section 5; on the survey-disciplined premium the trough and the rebound are small (+0.63pp) and largely US-driven.




Figure 4. ACM/BRW: domestic vs Mexico-specific term premium, 36m and 60m, and the cumulated US contribution.




Figure 4a. Survey-disciplined swap term premium, 36m and 60m: raw vs Mexico-specific.

The curve and the first real model (standalone ACM/BRW)




Figure 5. Udibono real zero curve by vertex (cross-sectional Nelson-Siegel on individual bonds).

Vertex	Real yield	Expected real path	Real term premium
2Y	3.49	2.41	0.92
3Y	3.87	2.45	1.27
5Y	4.31	2.48	1.78
10Y	4.64	2.33	2.35
20Y	4.64	1.53	3.06

Table 5. Standalone ACM/BRW on the Udibono curve – superseded by Table 13 (Section 6). Real zero yield, expected real path and real term premium, 31 Aug 2026 (% / pp). 207 month-ends from Jun 2009; PCs 94.8/4.7/0.5%; BRW δ scaled to 0.59 to keep Φ stationary (largest eigenvalue 0.9998; Q 0.998); fit SD 3–11bp.




Figure 6. Real term premium by vertex, standalone ACM/BRW on the bond-based Udibono curve from Jun 2009.

The two models agree on real yields but not on the real term premium (Table 6 below: –0.47pp at 2Y against 0.92). Section 5 tests both against Banxico’s surveys; neither survives, and Section 6 replaces them. For the record, the standalone premium rises with maturity (sample means 0.38pp at 2Y, 0.86 at 5Y, 1.44 at 10Y, 2.56 at 20Y), was high through 2013–15 (10Y peak 4.1pp, Sep-2015; 20Y 5.4pp, Nov-2013), turned deeply negative at the 2023–24 hiking peak (10Y trough –2.65pp, Feb-2024), and has climbed steadily since. Today it is 0.92 at 2Y, 1.78 at 5Y, 2.35 at 10Y and 3.06 at 20Y — below Brazil’s at 5–10Y (1.92–2.53pp) but slightly above it at 20Y (2.83pp). The expected real path is flat at about 2.3–2.5%, close

to Banxico's ex-ante real rate today (2.27%).

Robustness. The level factor is very close to a unit root, so BRW had to be scaled back to $\delta = 0.59$. The current 2–10Y readings barely depend on that choice or on the start date (δ from 0 to 0.6 and starts from Jun-2009 to Jan-2010: 2Y 0.81–0.96, 5Y 1.60–1.89, 10Y 2.14–2.41pp). The 20Y reading (2.35–3.07) and the 2013 levels (10Y 1.6–3.6) do depend on it. Compared with the previous run — fewer early bonds and realized inflation as the pre-2016 anchor — the current front-end premium is about 1.2–1.3pp higher at 2–5Y: the full survey anchor puts the expected real path near today's ex-ante real rate instead of above it.

Fiscal risk or technical dislocation?




Figure 7. 10Y real term premium and the cross-sectional fitting error.

With individual bonds this is now the same check as the NTN-B one, and it gives a sharper answer than Brazil's. The fitting error averages 9.7bp (90th percentile 16.4bp). In the current episode it is 4.9bp — lower than normal — so **today's rising real premium is a coherent repricing, not a dislocation**, as in Brazil. The 2013 taper tantrum is different: fitting error averaged 21.4bp (peak 26.1bp), above the 90th percentile, so part of that spike was technical dislocation — in line with the model sensitivity of the 2013 level. 2016–17 (7.8bp) and Covid (10.3bp) look normal; the one-month spike at end-2010 (56bp) is isolated. Flow data (foreign Udibono holdings) would still be the direct test.

The second real model: joint RealACM (Mbono + Udibono)

Vertex	Nominal TP (Mbono)	Real TP (joint)	Breakeven	IRP	Exp. inflation
2Y	0.30	-0.47	4.23%	+0.77	3.46%
3Y	0.52	-0.33	4.32%	+0.85	3.47%
5Y	1.08	0.24	4.42%	+0.85	3.58%
10Y	2.05	1.30	4.49%	+0.75	3.74%
20Y	2.83	2.11	4.57%	+0.72	3.85%

Table 6. Joint RealACM, Mbono (nominal, 1–20Y) and bond-based Udibono (real, 2–20Y), Apr 2011–Aug 2026, 31 Aug 2026 values. Liquidity factor = USDMXN 1M implied vol; fit 6.7bp nominal / 6.2bp real. Superseded by Table 13 (Section 6).

The nominal leg is Mbonos, not swaps: the Mbono–swap spread (about 0.5pp at 5–10Y today) would bias a swap-minus-Udibono breakeven. The raw fitting-error liquidity proxy used for Brazil still makes this model explode, even on the bond-based curve; with FX implied vol it is well behaved. Model expected inflation (3.46–3.85% today) sits within 0.3–0.4pp of Banxico’s survey on average (5Y vs next-1–4-years median: +0.29pp; 10Y vs next-5–8-years: +0.41pp) — the right level, a little high, and smoother in timing than the survey. In the previous run a smoothed-fitting-error version missed the survey by about 1pp, so the FX-vol split remains the one to use.




Figure 8. Breakeven (top) and inflation risk premium (bottom), 2/5/10/20Y – joint model.

The survey test, month by month

Banxico's surveys give two things neither model uses: an expected average policy rate over the next 12–36 months (from the funding-rate and Cete expectations) and expected inflation. Their difference is a survey real path; the Udibono real yield minus that path is a **survey-anchored real term premium** at 2–3Y, built without the swap curve, without Mbonos and without any VAR. It is not a perfect benchmark — forecasters are not investors, medians are sticky, and the 24–36-month inflation expectation is a blend of the 12-month and 1–4-year series — but it is independent of both models.




Figure 8a. Panel A: expected average policy rate over 24 months, survey vs swap model. Panel B: expected real short rate, survey vs the two models and the ex-ante real rate. Panel C: 2Y real term premium on each measure.

2Y real term premium, since 2016	Standalone	Joint	Survey-disciplined (Sec. 6)
Correlation with survey-anchored premium	0.05	0.64	0.90
Mean absolute gap (pp)	0.59	0.72	0.21
2023–24 average (survey-anchored: +0.57)	–0.53	+0.04	+0.18
31 Aug 2026 (survey-anchored: 1.00)	0.92	–0.47	0.94

Table 13a. Each model’s 2Y real term premium against the survey-anchored measure, monthly 2016–2026. At 3Y: correlations 0.01 / 0.57 / 0.85, gaps 0.72 / 0.68 / 0.26.

The standalone model fails on dynamics. Its expected real path has a 0.97 correlation with the current ex-ante real rate: with the level factor near a unit root it expects real rates to stay where they are, so its “term premium” is mostly the real yield minus today’s real rate — a slope, not a premium. It matched the survey in August only because the survey also expects Banxico to hold. In 2023–24, when analysts expected real rates to come down from about 7%, it reported a negative premium where the survey-anchored measure was +0.6pp. **The joint model** tracks the survey’s movements better (correlation 0.64) but misses its level, with a large swing in 2020, a deep negative in 2022 and a policy path about 1pp above the survey today.

Is it the Mbono–swap spread?

The joint model prices Mbonos, whose yields contain an issuer and liquidity spread over swaps, with the same factors that drive the policy rate. If that spread leaked into its expected path, it would explain the miss. The spread (Mbono zero minus F-TIE zero, monthly from Mar-2011) was modelled on its own:




Figure 8b. Mbono - F-TIE swap zero spread at 2, 5 and 10 years.

Tenor	31 Aug 2026	Sample mean	SD
1Y	0.26	-0.10	0.26
2Y	0.38	0.01	0.27
5Y	0.55	0.11	0.39
10Y	0.48	-0.03	0.44
20Y	0.26	-0.35	—

Table 13b. Mbono - swap zero spread (pp). 20Y from Nov-2011.

It moves in regimes rather than around a mean: negative in 2011–15, near zero in 2016–19, +0.5 to +1.0pp since 2022. A level factor explains 87% of it and a slope factor 10%. The level factor is highly persistent (autoregressive coefficient 0.96, half-life about 17 months; the ADF test cannot reject a unit root and KPSS rejects stationarity); the slope factor mean-reverts within about two months. Monthly changes are not explained by sovereign CDS, FX implied vol, USDMXN, the policy rate or Fed funds (R^2 0.04–0.06). Adding the swap curve factors lifts that to 0.23–0.32: the spread narrows when the swap curve steepens ($t \approx -3, -5$ at 10Y) and moves against swap-rate changes (correlation -0.49) — bonds lag swaps. Supply and foreign holdings, not yet in the data, are the likeliest drivers of the regimes.




Figure 8c. Gap between the joint and swap models' 36-month expected policy paths, against the 3Y spread.

But the spread does not explain the joint model's miss. The gap between the two models' expected paths moves only weakly with the spread month to month (correlation 0.26–0.34) and in the opposite

direction in levels (−0.22 to −0.40); today’s 2Y gap (0.92pp) is larger than the whole 2Y spread (0.38pp). The joint model’s error comes from its own dynamics. Over the full history its path is in fact closer to the survey (mean absolute gap 0.45pp at 24 months) than the swap ACM’s (0.68pp), which overshoots in hiking cycles by 1–2pp (2022–24). No curve-only model is reliable in every period; the surveys are the one consistent anchor.

6 • THE SURVEY-DISCIPLINED MODEL (HEADLINE)

Specification

One model covers the swap curve, Mbonos and Udibonos, with the surveys entering estimation as noisy observations of expected rates — the expectations side of Kim and Wright (2005):

```
state   $X_t = [N1\ N2\ N3 \mid S1 \mid R1\ R2]$ 
      swap PCs (1–120m) | Mbono-swap spread | real PCs orthogonal to N, S
P       $X_{t+1} = \mu + \Phi X_t + v_{t+1}, \quad \Phi[N, SUR] = 0$ 
short   $r_t = \delta_0 + \delta_N' N_t \quad (R^2\ 0.99)$ 
       $\tilde{r}_t = \rho_0 + \rho' [N_t, R_t] \quad (\text{real}; R^2\ 0.97; 1m = \text{target} - \text{survey 12m inflation})$ 
fit    min over  $\mu, \Phi: \sum \|\Sigma^{(-\frac{1}{2})} v_t\|^2 + \sum_h (\hat{s}_{h,t} - \tilde{E}_t r_{1..h})^2 / \sigma_{h,t}^2 + \sum_k (\hat{\pi}_{k,t} - \tilde{E}_t \pi_k)^2 / 0.25^2$ 
TP      $tp_t(n) = y_t(n) - \tilde{E}_t r_{0..n-1} \quad (\text{real: Udibono } y - \tilde{E}_t \tilde{r})$ 
       $BEI = y_{\text{Mbono}} - y_{\text{Udi}}, \quad \tilde{E}_t \pi = \tilde{E}_t (r - \tilde{r}), \quad IRP = BEI - \tilde{E}_t \pi$ 
```

The policy-path surveys enter at 12, 24 and 36 months, each month weighted by the cross-sectional standard deviation of the forecasts (floor 0.10pp; average 0.26–0.40pp); the inflation surveys (next 12 months, next 1–4 years, 5–8 years) enter with a 0.25pp error, as their dispersion is not in the file. Sample Mar-2011 to Aug-2026, 186 months; the swap PCs explain 96.3 / 3.3 / 0.4%, the spread factor 81%, the real factors 85 / 11% of the residual real curve. Largest eigenvalue of Φ 0.96. Term premia are yields minus model expectations — the expectations side of an affine model without a separately estimated set of risk prices; maturities to 20Y (the 20Y swap node from Nov-2011).

The out-of-sample test

Mean absolute error vs survey (pp)	Survey-disciplin ed	Same model, no survey	Joint (in-sample)	Swap ACM (in-sample)
Expected policy rate, next 12m	0.24	0.58	—	—
Expected policy rate, next 24m	0.30	0.91	0.45	0.68
Expected policy rate, next 36m	0.31	1.20	0.49	0.91
Inflation, next 12m	0.31	0.39	—	—
Inflation, next 1–4 years	0.10	0.51	—	—
Inflation, 5–8 years	0.10	0.87	—	—

Table 14. Expanding-window re-estimation at the end of each year 2016–2025; each fit is compared with the following 12 months of surveys (116 months). The two earlier models are estimated on the full sample, which favours them. Factor loadings are full-sample throughout; only the dynamics are re-estimated.




Figure 8d. Panel A: out-of-sample error on the 24-month policy path by year. Panel B: the error month by month at 12, 24 and 36 months.

The model beats both earlier models on months they were fitted to. Without the surveys the same model misses by 0.6–1.2pp, so the gain comes from the survey discipline, not the extra factors. The weakest year is 2024 (0.68pp at 24m): the easing came faster than a model estimated through 2023 expected. Expected inflation, which the joint model had 0.3–0.4pp too high, now sits on the 1–4-year and 5–8-year surveys out of sample (0.10pp).




Figure 8e. Expected average policy rate over the next 36 months: survey, survey-disciplined model (in-sample), swap ACM/BRW and joint RealACM.

The decomposition

31 Aug 2026 (% / pp)	2Y	3Y	5Y	10Y	20Y
Expected policy path	6.50	6.43	6.26	6.00	5.83
Swap term premium	0.78	1.18	1.77	2.66	3.37
Mbono – swap spread	0.38	0.45	0.55	0.48	0.26
Expected real path	2.55	2.59	2.55	2.39	2.28
Real term premium	0.94	1.28	1.77	2.25	2.36
Real TP net of issuer spread	0.55	0.82	1.22	1.77	2.10
Breakeven (Mbono – Udibono)	4.18	4.20	4.27	4.50	4.82
Expected inflation	3.95	3.84	3.72	3.61	3.55
Inflation risk premium	0.22	0.36	0.56	0.89	1.27

Table 13. Survey-disciplined decomposition – the headline measures. Real TP = Udibono zero - expected real path; net of issuer spread = real TP - Mbono-swap spread at the same maturity. Breakeven is bond on bond, so the issuer spread cancels in it.

The real premium rises with maturity and is well above its sample mean at every point (sample means 0.30 / 0.43 / 0.68 / 1.09 / 1.44pp). Unlike the standalone model, it was positive in 2023–24 (2Y +0.18), when the survey expected real rates to fall. Its two-to-three-year level matches the survey-anchored measure (1.00 and 1.34pp). About 0.4–0.5pp of the Udibono premium at 2–10Y is the issuer spread Mbonos also carry; net of it, the real term premium is 0.55–1.77pp. The expected real path is flat at about 2.3–2.6%, near today’s ex-ante real rate (2.27%) and inside Banxico’s neutral range. The 5Y swap premium, 1.77pp, is a little below the ACM/BRW estimate of Section 2 (1.92pp), whose bias-corrected dynamics expected more tightening.




Figure 8f. Panel A: 2Y real term premium on the four measures. Panel B: real term premium by maturity, survey-disciplined model.




Figure 8g. Panel A: 5Y breakeven, model expected inflation at 5Y and 10Y, and the 1-4-year and 5-8-year surveys. Panel B: inflation risk premium by maturity.

The inflation risk premium is modest: +0.22pp at 2Y rising to +1.27pp at 20Y. At 2-5Y it is 0.3-0.55pp below the joint model's because expected inflation there is higher, tied to the 12-month and 1-4-year surveys (4.23% and 3.80% in the latest rounds); at 10-20Y expected inflation is lower than the joint model's and the premium higher. It rose through 2021-23, peaked above 1pp at 5-20Y and has eased at the front end since.

Robustness

31 Aug 2026	RTP 2Y	5Y	10Y	20Y	Swap TP 2Y	5Y	10Y	20Y	IRP 2Y	5Y	10Y
Baseline (survey dispersion)	0.94	1.77	2.25	2.36	0.78	1.77	2.66	3.37	0.22	0.56	0.89
Fixed survey error 0.25pp	0.92	1.69	2.16	2.36	0.74	1.65	2.53	3.36	0.21	0.50	0.85
Fixed survey error 0.50pp	0.93	1.82	2.31	2.44	0.73	1.80	2.71	3.43	0.18	0.53	0.87
Fixed survey error 0.10pp	0.88	1.56	2.31	3.10	0.76	1.61	2.75	4.26	0.26	0.60	0.92
Three real factors	1.04	1.80	2.28	2.40	0.78	1.77	2.68	3.40	0.13	0.52	0.87
Two spread factors	0.95	1.78	2.26	2.36	0.78	1.78	2.67	3.37	0.21	0.54	0.89
Estimated to Aug-2025	0.80	1.70	2.21	2.33	0.74	1.80	2.68	3.37	0.32	0.64	0.95
No survey (restricted VAR)	0.98	1.77	2.01	1.95	0.70	1.49	2.04	2.51	0.11	0.26	0.50

Table 16. Real term premium (RTP), swap term premium and inflation risk premium under alternative specifications (pp).

At 2-10Y the readings barely move across survey weights, factor counts and an estimation ending a year earlier: real premium 0.80-1.04 at 2Y, 1.56-1.82 at 5Y, 2.16-2.31 at 10Y. The 20Y level is less secure. With a tight survey error (0.10pp) the model sits at its persistence ceiling and the 20Y real and swap premia rise to 3.1 and 4.3pp; the surveys reach only 8 years, so the 20Y expected path is an extrapolation of the dynamics. The no-survey variant is shown for reference only — it fails the out-of-sample test above.

WHAT CHANGED FROM THE EARLIER SPLIT

Real term premium: survey-disciplined model (Table 13) replaces the standalone Udibono model. At 2–5Y the two happen to agree today; over the history they do not.

Breakeven, expected inflation, IRP: survey-disciplined model replaces the joint RealACM; breakeven identical by construction; expected inflation now on the surveys (higher than the joint model's at 2–5Y, lower at 10–20Y), so the IRP is smaller at the front and larger at the long end.

Mbono–swap spread: modelled as its own factor and reported separately; it is not the source of the joint model's error.

US control: Tables 3, 4, 7 and 8 are re-run on these measures (Sections 3 and 7), with the US side rebuilt from the 22 Sep us_core pull.

No long-run link; a short-run one at the long end

Series	US liquidity proxy	EG p	KPSS resid p	α (t)	$\Delta \beta$ (t)	ΔR^2
Real 5Y	fit error	0.730	0.01	-0.101 (-3.1)	0.18 (2.4)	0.07
Real 5Y	VIX	0.200	0.08	-0.090 (-2.9)	-0.02 (-0.6)	0.00
Real 10Y	fit error	0.449	0.01	-0.114 (-3.3)	0.35 (3.9)	0.17
Real 10Y	VIX	0.197	0.03	-0.108 (-3.1)	0.09 (1.7)	0.02
Real 20Y	fit error	0.053	0.01	-0.146 (-3.0)	0.55 (5.2)	0.26
IRP 5Y	fit error	0.087	0.01	-0.155 (-4.4)	0.17 (2.4)	0.03
IRP 5Y	VIX	0.044	0.01	-0.118 (-3.7)	0.03 (1.2)	0.01
IRP 10Y	fit error	0.517	0.01	-0.113 (-3.7)	0.27 (3.6)	0.07
IRP 10Y	VIX	0.679	0.01	-0.073 (-3.6)	0.15 (4.2)	0.09
IRP 20Y	fit error	0.935	0.01	-0.042 (-2.6)	0.37 (2.9)	0.08

Table 7. Survey-disciplined real term premium and IRP (Section 6), Mar-2011-Aug-2026, against the US joint Treasury-TIPS model (us_core v2). Fitting-error liquidity is the baseline; VIX is the robustness check (its US levels are implausible – 5Y real TP about -5% – and it is not built at 20Y).

No cointegration anywhere. KPSS rejects stationarity of every levels residual, and Engle–Granger rejects only for the 5Y IRP against the VIX series (p 0.044) and nearly for the 20Y real premium (0.053); neither passes Paper 1’s both-tests standard. A basal real premium is therefore not defined, as before. **The short-run picture changes with the new measure.** On the superseded standalone premium the differenced betas were indistinguishable from zero; on the survey-disciplined premium, with the fitting-error US series, they are significant and rise with maturity: 0.18 at 5Y (t 2.4), 0.35 at 10Y (t 3.9), 0.55 at 20Y (t 5.2), with R^2 up to 0.26. The VIX version keeps little of this (10Y t 1.7), so the result rests on the baseline US measure. IRP: a partial short-run pass-through of 0.17–0.37 at all maturities, robust at 10Y under VIX (0.15, t 4.2) — unlike Brazil, where none survived.

What it means for the current rise. Since August 2023 the real premium rose 1.25pp at 5Y, 1.07pp at 10Y and 0.60pp at 20Y; the US-implied share is nil at 5Y, about 0.2pp at 10Y, and almost all of the 20Y move. The front and belly of the real curve repriced on Mexican grounds; the long end followed US real rates.




Figure 9. Panel A: Mexico 10Y real term premium against the US 10Y and 20Y. Panel B: raw vs Mexico-specific at 10Y and 20Y. Panel C: cumulated US-contributed component.

Putting it together: Table 8

	Observed	Expected path	TP total	Basal	Idiosyncratic	Yield if idio = 0
Nominal, 36m	7.61	6.43	1.18	0.68	0.50	7.11
Nominal, 60m	8.04	6.26	1.77	1.05	0.72	7.31
Real, 5Y	4.31	2.55	1.77	n/a	n/a	n/a
Real, 10Y	4.64	2.39	2.25	n/a	n/a	n/a
Real, 20Y	4.64	2.28	2.36	n/a	n/a	n/a

Table 8. Decomposition as of 31 Aug 2026 (% / pp), all rows on the survey-disciplined model (Table 13). Nominal basal is tentative: the level the premium would have at today's US premia on the levels fit, which is not a cointegrating relation (Table 3); real basal is not defined (Table 7).

Nominal 36m: 7.61% = 6.43% expected path + 0.68 basal + 0.50 Mexico-specific; 60m: 8.04% = 6.26% + 1.05 + 0.72. On this tentative reading, removing the Mexico-specific premium would put the 5-year zero at about 7.3%, 72bp lower. (On ACM/BRW the same split was 6.12 + 1.07 + 0.85, and 7.2%.) The more robust statement is Table 4's: the whole of the last year's rise is Mexican on either measure.




Figure 10. 36-month zero: actual, model-fitted, and the expected average short rate; term premium below.




Figure 11. The same at 60 months.




Figure 12. Vintage forecast paths (dashed) vs the realized Banxico target (solid).

2011-06: expected a gentle decline to about 3.7%; Banxico cut to 3% and then hiked through 2016–17 — the hiking is missed.

2016-11: expected a rise to about 7%; the target peaked at 8.25% in 2018–19.

2021-06: expected a plateau near 7%; the target reached 11.25% — the largest miss, as for Brazil’s 2021 vintage.

2024-06: expected a slow easing toward 9.7% by 2030; Banxico was at 6.5% within two years.

As in Brazil, the physical-measure VAR is surprised by moves that are larger and faster than expected, in both directions. For the current vintage, the model’s near-term path (6.56% over 12m, 6.48% over 24m) is close to the August survey; over the whole sample the survey shows the swap model overshooting in hiking cycles (Section 5), which is what the survey discipline of Section 6 corrects.

Nominal




Figure 13. Rolling $\bar{r}$ (84m, stepped 6m, expanding from 24m) and the swap ACM/BRW model 5y5y, against the Banxico target; policy gap below.

$\bar{r} = 8.01\%$ (window ending Mar-2026). With the target at 6.50%, the gap reads -1.5pp , “accommodative” — but only because the 7-year window contains the 11% years. The forward-looking measure is the 5y5y expected policy rate. On the swap ACM/BRW model it is 4.98% (Aug); on the survey-disciplined model it is **5.74%**, which reads policy as mildly restrictive ($+0.8\text{pp}$). **The choice of definition still flips the sign of the stance.**

Real




Figure 14. Rolling real $\bar{r}$ and the standalone model real 5y5y, against the Udibono 10Y and 1Y real yields and Banxico's published neutral real range.

Real $\bar{r}$ = **3.87%** (window ending May-2026), up from 1.3–2.4% in 2019–22 and a peak of 3.96% in mid-2025 — a 7-year average of the ex-ante real policy rate, the rate Banxico compares with its range. The survey-disciplined real 5y5y is **2.24%** (standalone model: 2.19%), close to today's ex-ante real rate (2.27%) and well below the 10Y Udibono (4.6%): the long real yield carries a large premium over where real short rates are expected to settle.




Figure 14a. Survey-disciplined 5y5y expected policy rate (nominal and real) against the target, the ex-ante real rate and Banxico's published ranges.

Banxico check (Paper 2's BCB counterpart). Banxico's published range for the long-run neutral real rate is 1.8–3.6% (midpoint 2.7%), unchanged since the August 2024 update (1.8–3.4% from 2019, 1.7–3.3% from 2016); the nominal range is 4.8–6.6%. The survey-disciplined 5y5y sits inside both (5.74% nominal, 2.24% real, just below the real midpoint); the trailing $\bar{r}$ sits above both. The closest comparison is Banxico's own affine curve model within its average (Kim–Wright type, nominal 10-year expectation minus the 3% target), at 2.8% with a 2.5–3.2% band; ours is about 0.3pp below its bottom. One caution: anchored by the 5–8-year inflation survey, the survey-disciplined 5y5y is very stable (5.5–5.9% nominal, 2.0–2.4% real over the whole sample), so it is informative as a level more than as a time series.

The output gap: validated against Banxico

Two constructions were built. The Areosa-style one follows Paper 2: potential output, natural employment and natural capacity utilization estimated jointly by a Kalman filter with the production-function identity imposed (employment = 100 – seasonally adjusted unemployment, MXUEUNSA; manufacturing capacity utilization MXMNTOTL from 2018; labor share 0.40). The second is a Hodrick–Prescott filter ($\lambda = 1600$) on log real GDP (chained from Bloomberg’s seasonally adjusted quarter-on-quarter series). Banxico’s own published gap is an HP filter with tail correction on GDP, and the comparison settles the choice:

Measure	Corr. with Banxico	Mean abs. diff (pp)	Mean diff 2024–26	2026Q2
HP on GDP (baseline)	0.993	0.35	+0.43	–0.04
Areosa-style	0.919	1.20	+2.09	1.71
Banxico (published)	—	—	—	–1.18

Table 9. Validation against Banxico’s GDP-based output gap (Informe Trimestral Abril-Junio 2026, Tabla 106), 2018Q1–2026Q2.




Figure 15. Baseline (HP on GDP), Banxico’s published gap, and the Areosa-style construction.

The HP gap reproduces Banxico’s series (correlation 0.99), including the turn negative from 2024Q4; it is less negative at the end (−0.04 vs −1.18 in 2026Q2), which fits Banxico’s tail correction and a different GDP vintage. The Areosa-style gap follows the shape but runs about 2pp above Banxico since 2024 and is still +1.7% today: Mexico’s unemployment rate is very low and trending down with informality, and capacity data cover only eight years, so both labor and capacity legs read “tight” while GDP says slack. The HP gap is the baseline; the Areosa-style construction is kept only as a documented divergence — the method was built for Brazil’s labor market and does not transfer.

Quarter	Banxico	HP baseline	Areosa-style
2024Q3	1.35	1.38	2.39
2024Q4	−0.15	0.03	2.46
2025Q1	0.01	0.33	1.66
2025Q2	−0.62	−0.15	1.77
2025Q3	−1.34	−0.70	1.47
2025Q4	−1.01	−0.21	1.67
2026Q1	−1.96	−1.00	1.04
2026Q2	−1.18	−0.04	1.71

Table 10. Last eight quarters, % of potential GDP.




Figure 16. Appendix variant: employment rate and capacity utilization against their jointly estimated natural levels, and the Areosa-style gap.

For the Areosa-style variant, refitting on truncated samples revises the reading down as data arrive (2024Q4: 2.84 → 2.46; 2025Q2: 2.04 → 1.77), and a labor share of 0.555 changes nothing (correlation 0.994).

Does the bridge change anything?

Maturity (m)	R ² lagged-only	R ² + macro (12 lags)	ΔR^2	F	p
2	0.060	0.128	0.067	0.56	0.954
3	0.075	0.128	0.053	0.44	0.991
6	0.103	0.166	0.063	0.54	0.960
9	0.070	0.131	0.060	0.50	0.975
12	0.053	0.111	0.058	0.47	0.984
18	0.081	0.151	0.070	0.59	0.933
24	0.077	0.142	0.065	0.55	0.959
36	0.078	0.139	0.062	0.52	0.971
60	0.077	0.139	0.062	0.52	0.970
108	0.079	0.159	0.080	0.69	0.857

Table 11. Nominal (F-TIIE) curve, genuine ex-ante forecast: one-month excess returns on lagged latent factors, then adding 12 lags each of INPC inflation and the baseline output gap (F-test, 24, 173 df).

Maturity (m)	R ² lagged-only	R ² + macro (12 lags)	ΔR^2	F	p
24	0.039	0.243	0.204	1.99	0.006
36	0.046	0.255	0.209	2.07	0.004
60	0.054	0.243	0.189	1.84	0.014
84	0.060	0.227	0.167	1.59	0.047
108	0.065	0.216	0.151	1.42	0.102
120	0.067	0.213	0.145	1.36	0.132
240	0.090	0.206	0.116	1.08	0.376

Table 12. Real (Udibono) curve, identical specification (24, 177 df).

Nominal curve: nothing at any maturity (p 0.86–0.99; no p-value below 0.60 at 9–13 lags). Real curve: **significant from 2 to 7 years** (p 0.004–0.047), not at 9–10 years (0.10–0.13) or 20; at 9–12 lags the 2–7 year segment stays significant throughout, and at 13 lags 2–5 years. The same test with the Areosa-style gap (previous run) gave the same shape, so the result does not depend on the gap choice. Brazil found predictability on both curves (nominal 2–24m, real 24–60m). In Mexico the nominal curve prices inflation and the gap without delay – consistent with the fast pass-through of US moves in Section 3 – while the real curve does not. Caveat shared with the Brazil test: the monthly gap is interpolated from smoothed quarterly values and carries information not available in real time.




Figure 17. Rolling nominal $\bar{r}$, bridge vs latent-only (short-rate equation fit once on the common window, averaged over rolling windows).




Figure 18. The same for the real curve. Mean absolute gap 0.017pp — with the ex-ante anchor the bridge adds little to the real short rate, as in Brazil.