Four Channels for BRL

Rate Differentials, the Dollar, the Harvest, and the Price of Risk

MARCOS MOLLICA

Weekly data, 2010–2026 · n = 869 · Newey-West HAC, 4 lags

USDBRL is usually talked about as a single story — "carry," or "risk-off," or "commodities." The data says it's at least four separate stories, moving largely independently of each other, and a regression that doesn't separate them will misattribute one channel's move to another.

$$\Delta USDBRL \approx \beta_1 \cdot \Delta(\text{BR–US 2y diff.}) + \beta_2 \cdot \Delta(DXY) + \beta_3 \cdot \Delta(\text{commodities}) + \beta_4 \cdot \Delta(\text{BR term premium}) + \varepsilon$$

This note builds that decomposition one channel at a time, using the same curve-based rates machinery as the companion term premium note as its foundation, and a single discipline applied without exception: no variable enters the model until it has cleared two bars — genuine orthogonality against everything already included, and a real, statistically meaningful contribution to explanatory power. One variable is kept in the final specification despite failing the second bar, and that exception is stated as plainly as the rule.

THE DISCIPLINE, STATED BEFORE ANY RESULT

Every asset in a macro dataset shares some exposure to a small number of global factors. Add enough right-hand-side variables to an FX regression and R^2 will always rise; the question is whether that rise reflects genuine, separable information or one factor wearing several names. The protocol used throughout this note, for every candidate variable without exception:

1. *Orthogonality first.* Report the correlation between the candidate and every regressor already in the model, before looking at whether it helps the regression. A candidate that is highly correlated with something already included is not a new channel, whatever its own coefficient looks like in isolation.

2. *Incremental contribution second.* Add the candidate to the existing model and report the R^2 gain and the candidate's own HAC-robust t-statistic. Both are shown — a variable that

raises R^2 by construction (it always will, with enough regressors) but carries an insignificant coefficient has not earned its place.

3. *Admit only if both clear the bar* — with one stated exception, addressed on its own terms at the end, not folded into this rule.

Why Changes, Not Levels: A Tested Choice, Not an Assumed One

Every regression in this note is run on weekly changes, not levels. That choice is usually asserted rather than tested — a reasonable-sounding appeal to the risk of spurious regression among trending series — and this note holds itself to a higher bar than that everywhere else, so it does here too.

Unit root tests (ADF and KPSS, run jointly since neither alone is fully reliable) confirm that USDBRL, DXY, Brent, iron ore, the 2-year differential, and the US term premium all behave as I(1), non-stationary series over this sample — trending, not mean-reverting, exactly the profile that makes a levels regression risky. That much simply confirms the working assumption. The more useful question is whether these series are cointegrated — whether some linear combination of them defines a stable long-run equilibrium that a levels regression with an error-correction term could exploit, giving a more efficient long-run estimate than a changes-only specification throws away.

They are not, on the evidence available. A Johansen test on the full system finds no robust cointegrating relationship: the trace statistic marginally rejects "no cointegration" at short lag lengths (1–2 lags) but fails to reject the same null at longer ones (3–4 lags) — a result that flips with an essentially arbitrary specification choice is not a finding worth building on. The max-eigenvalue statistic, generally the more reliable of the two Johansen tests in finite samples, never rejects "no cointegration" at any lag length tested. Pairwise Engle–Granger tests confirm the same conclusion more simply: USDBRL shows no cointegration with any single candidate driver on its own ($p > 0.3$ in every case, most above 0.7).

This makes the changes-only specification a tested econometric choice, not a cautious default. Forcing an error-correction framework onto series the data doesn't actually show a stable long-run relationship between would mean inventing an equilibrium the tests can't find — the same trap the US term premium section of this note deliberately avoids by differencing rather than regressing two persistent series in levels. Consistency matters here: a note that flags that risk explicitly in one place should not walk into a milder version of it in another.

One unit root result is worth a specific note rather than a passing mention. The Brazilian term premium's own classification is ambiguous — ADF rejects a unit root (suggesting stationarity) while KPSS also rejects stationarity (suggesting non-stationarity), tests that should agree and here do not. That is not test failure; it is close to the expected signature of a variable that is theoretically mean-reverting — a term premium has no structural reason to trend to infinity — but reverts slowly enough, across a sixteen-year sample spanning several genuinely distinct regimes, to look statistically like a near-unit-root process in a finite window. It is the same regime-dependence theme running through the rest of this research, showing up here as an econometric symptom rather than a headline result.

Channel 1: The Rate Differential

The foundation, carried over directly from the term premium work: the BR–US 2-year yield differential, built from the same DI curve construction. This is the starting point, not a candidate under test — its inclusion is the paper's premise, not a result.

Channel 2: The Global Dollar

Orthogonality: correlation between ΔDXY and Δ(2y differential) is 0.05 — essentially zero.

MODEL	R^2	DIFF. T-STAT
Rate differential alone	—	—
+ DXY	23.3%	8.72

DXY passes both bars cleanly: orthogonal to the regional rate story, and a large, significant independent contributor. This is Channel 1 + Channel 2, the baseline the rest of this note builds on.

Channel 3: Commodities (Terms of Trade)

The working hypothesis going in was that commodities would fail the orthogonality bar — most commodities are USD-denominated, so broad dollar strength and commodity weakness should move together close to mechanically, making a commodity factor largely redundant with DXY. Three components were tested, each USD-priced to avoid the circularity of using a BRL-denominated index to explain BRL: Brent crude (2nd-month generic), soybean (2nd-month generic, Brazil's largest single export), and iron ore (SGX 62% Fe CFR China).

	CORR. VS Δ(2Y DIFF)	CORR. VS ΔDXY
Brent	-0.078	-0.120
Soybean	-0.080	-0.194
Iron ore	-0.032	-0.092

The hypothesis was wrong. All three are modestly, not mechanically, correlated with DXY — the USD-denomination argument turned out to be much weaker in practice than in theory.

MODEL	R^2	INCREMENTAL
Rates + DXY	23.3%	—
+ Brent, Soybean, Iron ore	30.4%	+7.1pp

Soybean carries the strongest, most robust relationship ($t = -6.12$ in the joint model), Brent second ($t = -4.45$), iron ore weakest but still marginally significant ($t = -1.75$). A separate check confirmed commodities are not simply proxying for the rate differential itself: regressing the differential on each commodity gives R^2 under 1% in every case, mostly insignificant — this is a genuinely separate transmission mechanism, not the rates channel wearing a different name.

Channel 4: The Brazilian Term Premium

The companion note, *What the DI Curve Is Actually Pricing*, decomposes every Brazilian DI yield into an expectations component and a term premium — what the market is pricing Copom to do, versus what it charges to bet on that path; its appendix gives the complete estimation procedure. The rate differential used in Channel 1 captures the level of Brazilian yields relative to US yields; it does not distinguish which part of that level is expectations and which is risk compensation. This channel tests whether that distinction matters for BRL specifically, not just for the curve.

	CORR. VS Δ(2Y DIFF)	CORR. VS ΔDXY
BR term premium (36mo)	-0.001	+0.020
BR term premium (60mo)	+0.006	+0.025

Essentially zero against both. This is the cleanest orthogonality result of any candidate in this note — the term premium is not a repackaged version of the rate level or the dollar.

MODEL	R^2	INCREMENTAL	BR TP T-STAT
Rates + DXY + commodities	30.4%	—	—
+ BR term premium (60mo)	31.4%	+1.0pp	2.73

Smaller than the commodity channel in magnitude, but genuinely significant and genuinely orthogonal. The economic reading is direct: it is specifically the risk-compensation component of Brazilian rates — not the expectations component, and not the rate level as a whole — that carries independent information for the currency. A hiking cycle priced as high-conviction and credible should move BRL differently than the same level of rates priced with a wide, uncertain risk premium attached, even though both look identical in the raw differential. This is the direct FX analogue of the companion note's whole argument: separating expectations from premium is not just a curve-fitting exercise, it changes what the object explains.

The Deliberate Exception: The US Term Premium

Every reader of this note who has followed US rates markets in 2026 will have the same question: doesn't the disappearance of the US term premium as a stable, near-zero benchmark change all of this? The companion note, *What the DI Curve Is Actually Pricing*, addresses this directly for the curve itself; the same objection needs to be addressed here, for BRL.

Only the 10-year US term premium is used, not the 2-year. The 2-year point is dominated by near-term monetary policy expectations — precisely the component the BR–US 2-year differential in Channel 1 already captures on both sides of the pair. Including a US 2-year term premium alongside a BR–US 2-year differential risks testing two overlapping versions of the same near-term policy story rather than two genuinely separate ones. The 10-year is a cleaner test: it is where duration risk compensation is conventionally understood to concentrate, and it is the specific series the companion note's own US-control exercise treats as the primary reference. The series is sourced directly from the New York Fed's published monthly ACM decomposition of the Treasury curve — the identical methodology detailed in the companion note's appendix, applied by the Fed to the US curve, not re-estimated here.

	CORR. VS Δ(2Y DIFF)	CORR. VS ΔDXY

US term premium (10y)	+0.066	-0.025

MODEL	R^2	INCREMENTAL	US TP (10Y) T
Rates + DXY + commodities + BR TP	31.40%	—	—
+ US term premium (10y)	31.40%	+0.00pp	0.07

By the rule stated at the outset, this fails outright — the incremental R^2 does not round above zero, and the coefficient's t-statistic is 0.07. It is included in the final specification anyway, and stated here explicitly as the one deliberate exception to the discipline this note otherwise holds itself to.

The reason is not statistical — it's that omitting it, silently or otherwise, invites exactly the objection it's meant to preempt.

A reader who wants to know whether US term premium compression explains any of this deserves the regression that actually tests it, not an assurance that it wasn't necessary. The honest answer is that it was tested, on the specific series the object of the objection actually refers to, and found to explain literally nothing once the rest of the model is in place.

The companion DI note treats this same US term premium series as arguably stationary in levels, on theoretical grounds (a term premium is risk compensation, not a persistently trending quantity) rather than a settled statistical result, given how short the post-2024 trending window still is. That is a different choice than the one made here, and worth confronting directly rather than leaving as a quiet inconsistency between the two notes: re-estimating the full model above with the US term premium entered in levels rather than differenced changes essentially nothing — R^2 moves from 20.64% to 20.75% against a 20.67% differenced baseline (all on the underlying weekly panel rebuilt for this check), and the coefficient remains statistically indistinguishable from zero either way ($t = 0.52$ differenced, $t = -1.28$ in levels). The choice of specification for this one regressor does not change its conclusion in either note.

The Full Model

REGRESSOR	B	T (HAC)

Δ(BR–US 2y differential)	+0.0232	8.36
ΔDXY	+0.5015	6.64
Δ Brent	-0.0672	-4.79
Δ Soybean	-0.1162	-5.83
Δ Iron ore	-0.0262	-1.95
Δ BR term premium (60mo)	+0.0079	2.71
Δ US term premium (10y)	+0.0004	0.07

R^2 = 31.40%, up from 23.3% at the rates-plus-dollar baseline. Six of seven regressors are genuinely orthogonal and genuinely significant; one (US term premium) is included by deliberate exception, not by the rule; none were added to inflate the headline number — the sequence above shows exactly what each one contributed, in the order it was tested, including the one that contributed nothing at all.

Open Question: The Sign on the Rate Differential

Channel 1's own coefficient carries a sign that runs against the textbook carry-trade intuition. A wider BR–US 2-year differential should, on the standard story, make BRL more attractive and pull the currency stronger. Instead it comes back positive and strongly significant (β = +0.0232, t = 8.36): a widening differential coincides with BRL *depreciation*, not appreciation. This holds up on an independent rebuild of the underlying weekly panel from raw DI futures and Bloomberg data, so the result is not an artifact of the original construction.

A wider rate differential should make BRL more attractive. The data says the opposite — and the reason is still unresolved.

Two explanations were considered, both consistent with the discipline used throughout this note.

Reaction-function (reverse causality). Copom's decisions are not independent of BRL. If the central bank raises rates in response to currency weakness and the imported-inflation pressure that follows it, a given week can show both a wider differential and a weaker currency without any carry mechanism running backwards — the causal arrow runs

from BRL to the differential, not the reverse, and a contemporaneous weekly regression cannot separate the two directions on its own.

Fiscal risk deterioration. A channel closer to the literature on sovereign risk premia and debt-sustainability concerns under an unfavorable fiscal starting point (in the spirit of Sargent & Wallace's "unpleasant monetarist arithmetic"): with roughly half of Brazil's federal debt stock indexed to the Selic rate, a widening differential driven by rate hikes mechanically raises near-term debt service. If the fiscal backing for that debt is itself in question, higher rates can raise — rather than compensate for — the currency risk investors demand to hold BRL.

An interaction test was run to try to discriminate between the two: the differential interacted with the Brazilian term premium level (one-week-lagged, demeaned) as a candidate proxy for fiscal risk conditions. It came back null — incremental R^2 of +0.09pp, t = 0.64 on the interaction term — and a median split by term premium level showed no meaningful difference in the differential's coefficient across the two regimes. That result weakens one specific version of the fiscal-risk-deterioration explanation. It does not, on its own, support reverse causality: the standard identification device for this kind of puzzle in the literature — controlling for a plausibly exogenous inflation driver, in the spirit of the US "price puzzle" fix of adding commodity prices to a monetary VAR — does not transfer cleanly to Brazil. The natural Brazilian analogue to that exogenous driver is BRL itself, and BRL is precisely the variable under study here, endogenous to the policy decision by construction; using it as a control would reintroduce the problem it is meant to solve.

This remains an open question, not a resolved one. The next test under consideration is a narrow-window monetary policy shock series — the change in DI-implied pricing from just before to just after each Copom decision — which would isolate the surprise component of policy moves from BRL's pre-decision level by construction, and offer a cleaner read on which direction the causality actually runs.

Application: A Fair Value Read

A regression this well-specified invites an obvious application, and it is one of the most common chart types in a sell-side FX note: a model-implied "fair value" line plotted against the actual rate, with the gap between them read as a signal. Building it correctly from this specification takes more care than it looks, and the two ways it can be done wrong are worth showing, not just avoiding silently.

The specification above is estimated in weekly changes, not levels — deliberately, since USDBRL's level trended from roughly 1.7 to over 5 across this sample, and regressing a level on other levels risks the well-known spurious-regression problem when the underlying series are non-stationary. A fair value level has to be built by cumulating the model's fitted weekly changes forward from an actual starting point, which introduces a problem levels regressions don't have: compounding.

The naive version — cumulate every fitted weekly change, including the regression's own intercept, from the start of the sample in 2010 — fails badly. The intercept is small on a weekly basis (0.11%) but compounds to an implied 163% move over 863 weeks, unrelated to anything any regressor actually did; it is simply the average rate at which BRL depreciated over 16 years, restated as if the model had predicted it. The resulting "gap" grows to 80% and never oscillates, which is a symptom of the construction, not a finding about the currency. The opposite fix — drop the intercept entirely — is worse: with no intercept, the model's cumulated path barely moves at all over 16 years, implying a fair value near 2.0 in 2026 against an actual rate above 5, understating BRL's genuine trend depreciation entirely.

The fix used here is to re-anchor the model to the actual level every 26 weeks, and read only the cumulative deviation that opens up within each window before the next reset. This keeps the regression's own long-run drift assumptions from ever compounding far enough to dominate the picture, and it matches how these charts are actually used in practice — as a read on whether the currency has moved more or less than the specification would suggest over the past several months, not as a claim about its correct level across a full cycle.

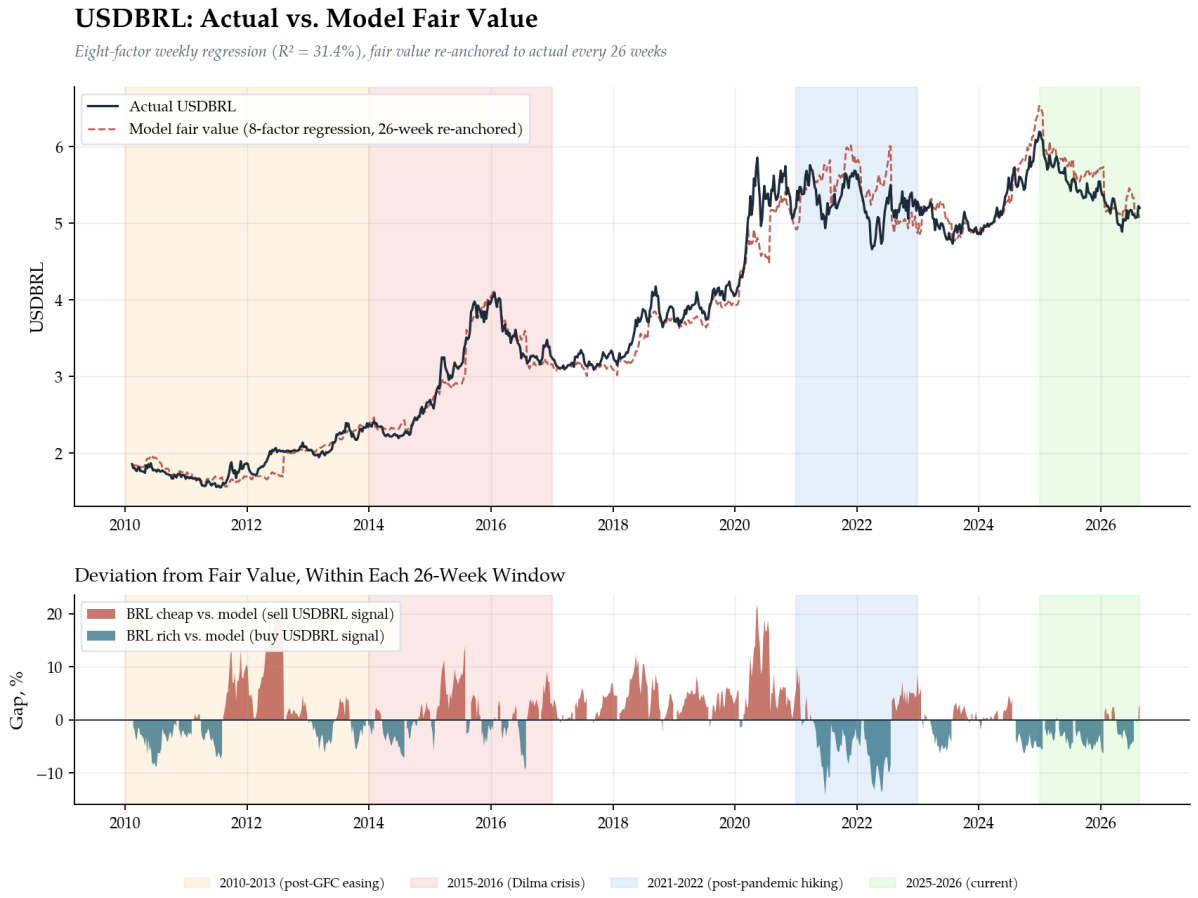

USDBRL: Actual vs. Model Fair Value — eight-factor weekly regression (R^2 = 31.4%), fair value re-anchored to actual every 26 weeks.

Reading it: the top panel tracks the two paths since 2010; the bottom panel shows the resulting gap, positive when BRL is cheaper than the model — the point at which a mean-reversion read would favor selling USDBRL — and negative when it is richer, favoring the opposite. The gap centers close to zero over the full sample (mean +0.7%, std 5.5%) and moves through sensible, recognizable episodes: BRL reads cheap against the model into 2012-13 and again through the 2020 COVID shock, and rich through stretches of 2021-22 and 2024. The most recent reading has BRL modestly cheap, at +2.9%.

With R^2 = 31.40%, roughly two-thirds of any given week's move is unexplained by this specification, so part of any reading is accumulated noise rather than signal, and this is in-sample fitted like everything else in this note — a level worth watching, not a level worth trading on its own.

Like the term premium estimates it builds on, this is an in-sample result over the full 2010–2026 window, not walk-forward validated. The relationships are economically coherent and orthogonality is unusually clean across the board, which is encouraging for a

first look — but a proper out-of-sample test is the natural next step before any of this informs live sizing at scale.

A Counterfactual: BRL Without the Dollar

The four-channel decomposition also makes a different kind of question answerable: how much of USDBRL's path is broad-dollar strength or weakness passing through mechanically, and how much would be there anyway? The same purge used for the US term premium control in the companion DI note — strip out a regressor's estimated contribution and look at what remains — applies here to Channel 2.

Each week's DXY-attributable move is the model's partial DXY coefficient times that week's actual ΔDXY. Subtracting it from the actual weekly change in USDBRL leaves a dollar-neutral weekly change, which is cumulated forward into a counterfactual path — re-anchored to the actual level every 26 weeks, for the same reason the fair-value chart above is: an unanchored cumulation lets the model's own drift assumptions compound over sixteen years into something no longer readable as a signal.

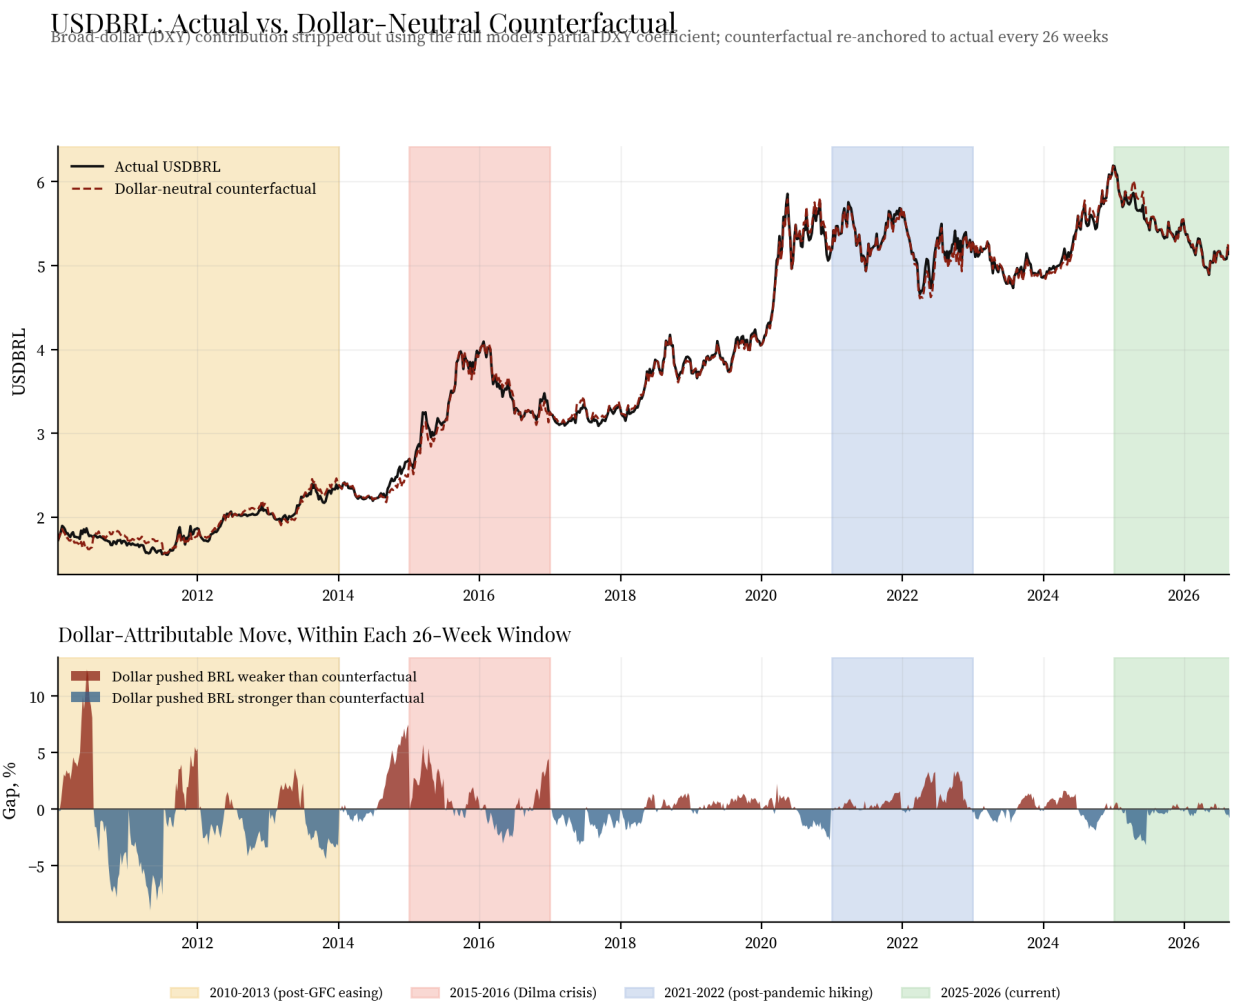

USDBRL: Actual vs. Dollar-Neutral Counterfactual — DXY's partial contribution removed, counterfactual re-anchored to actual every 26 weeks.

Reading it: the two paths track closely for most of the sample — DXY explains a real but bounded share of any given window's move, consistent with its 7.1pp standalone contribution to the model's R^2. The gap panel is where the channel shows up distinctly. 2010–2011 stands out as the sample's clearest episode: the dollar's post-GFC weakness was doing real work pushing USDBRL down, and stripped of it the counterfactual sits meaningfully higher — BRL was borrowing dollar-weakness tailwind, not purely a domestic story, in exactly the period the note's regime shading calls out as post-GFC easing. 2015–16 shows the reverse: dollar strength into the Dilma crisis was compounding a currency move that domestic fiscal and political stress was already driving, with the counterfactual running below actual through most of the episode. The most recent window (2025–26) shows a modest, not dramatic, dollar contribution — the current leg of BRL weakness reads as mostly non-dollar in origin, which is a different diagnosis than the popular "it's all the dollar" framing would suggest. Concretely: the latest reading has actual USDBRL

at 5.14 against a dollar-neutral counterfactual of 5.18 — the current dollar backdrop is a modest tailwind for BRL, and without it USDBRL would be trading closer to 5.18, roughly 0.8% weaker than the actual rate.

The same caveats as the fair-value chart apply without qualification: this is a partial-equilibrium purge from an in-sample linear model, not a structural counterfactual — it assumes the other channels' behavior is unchanged in a world without the dollar move, which is a simplification, not a general-equilibrium claim. It is a read on attribution within the model's own terms, not a claim about what BRL would have actually done.

This is one piece of a broader look at what actually explains BRL, alongside the term premium work — more to follow.

⁂