

SÃO PAULO • SEPTEMBER 2026

Is the Auction Trade Worth the Trouble?

MARCOS MOLLIKA, PHD • MACRO BRASIL

Brazilian linker auctions leave a concession in the cash bond that takes two weeks to unwind. Nominal auctions leave nothing at all across days, and what happens inside the session is smaller than the cost of trading DI to capture it — except at the very largest offerings.

Every desk in São Paulo knows the trade. Buy DI ahead of a nominal auction — DI trades in rate, so buying is paying, and the position gains if yields rise — bid for the bond, then unwind both legs once the paper is in. It is a cumbersome operation with several moving parts and a balance-sheet cost, and the question that hangs over it is whether the concession being harvested is large enough to pay for the harvesting. This note measures both sides.

The answer differs sharply between the two instruments, and the difference is the most interesting thing here. The inflation-linked bond carries a real, well-identified concession that persists for days. The nominal complex does not — and the intraday move that does exist sits below the round-trip execution cost of the hedging instrument at all but the largest offerings.

THE DATA

Three sources, all covering the same market from different angles. The first is the Tesouro Nacional auction record: 19,813 individual auction lines running from January 2000 to September 2026, each carrying the instrument, the auction type, the round, the maturity, the quantity offered, the quantity actually accepted, and the accepted rate. Restricting to market sale auctions and excluding the Banco Central leg leaves the primary-market supply history at line-item resolution. Two facts from it shape everything that follows. The auction calendar is fixed and split by instrument: NTN-B auctions fall on Tuesday in 517 of 539 dates since 2011, while LTN and NTN-F auctions fall on Thursday in 741 and 568 dates respectively, and in 131 of 141 dates since 2024. And the quantity is disclosed on the morning of the auction itself, not with the quarterly calendar, which means the identity of the maturities being sold is anticipated while the size is genuine news on the day.

The second is a daily panel of bond-level NTN-B real yields by maturity date, 34 bonds running from November 2005, which supports the cross-maturity comparison at the heart of the linker section. The third is an intraday panel of DI1 futures pulled at five-minute resolution: 195,920 bars across 14 contracts spanning 0.3 to 8.6 years of maturity, covering 136 trading sessions from 27 February to 11 September 2026, of which 30 are nominal auction sessions and 28 are linker sessions. Bloomberg will not serve intraday bars beyond roughly 140 trading days, so this window is the maximum available at the time of writing and cannot be extended backwards — only forwards, by repeating the pull.

PANEL	OBSERVATIONS	CROSS-SECTION	PERIOD
Auction record	19,813 lines	6 instruments	2000–2026
NTN-B daily yields	7,396 bond-days	25 maturities, 13.6 live/date	2011–2026
DI intraday, 5-min	195,920 bars	14 contracts	Feb-Sep 2026
— nominal auction sessions	30	13.6 contracts/date	Thursdays
— linker auction sessions	28	13.6 contracts/date	Tuesdays

DI contract 0DZ26 excluded throughout: roughly 900 bars per range against 5,000 for the rest.

IDENTIFICATION

On any auction date the Treasury offers a subset of the maturities that are live and quoted. The rest trade in the same market, face identical aggregate shocks, and are untouched. That supports a within-date cross-sectional comparison in which date fixed effects absorb everything common — Copom, global duration, risk appetite, the currency — leaving only the difference between offered and non-offered points on the same curve on the same day.

SPECIFICATION


$$\Delta y(m,t) = \beta \cdot \text{Supply}(m,t) + f(\tau) + \delta_t + \varepsilon(m,t)$$

m = maturity or DI contract · t = auction date · δ_t = date fixed effect
 $f(\tau)$ = polynomial in maturity · SE clustered by date

For DI, $\text{Supply}(m,t)$ is the duration-matched offered size: each auctioned LTN or NTN-F is mapped onto the curve at its Macaulay duration and weighted by a Gaussian kernel in duration distance.

PART ONE: THE LINKER LEAVES A MARK

Yields at the NTN-B maturities about to be auctioned begin drifting up roughly a week ahead, relative to the rest of the curve on the same day. The run-up peaks at about 0.7 basis points per standard deviation of offer size on the auction day itself. The window eight to five days before shows nothing, so this is not a pre-existing trend.



Cumulative yield change at auctioned NTN-B maturities against the rest of the curve on the same day, within-date demeaned. The run-up peaks on the auction day and is given back over the

following fortnight.

Conditioning on which bonds are offered separates the two channels the disclosure regime creates. Among bonds that *are* offered, the size of the offering explains none of the pre-auction drift. The run-up is about anticipated identity — the calendar is public. The quantity has its own effect and it lands on the day it is disclosed: among offered bonds, a one standard deviation larger offering raises that maturity's yield by 0.44 basis points on the announcement day relative to the rest of the curve. Prior drift does not forecast the size the Treasury chooses, so the disclosed quantity is close to genuine news.

WINDOW	COEFFICIENT	T	
t-8 → t-5 <i>placebo</i>	-0.25	-1.84	—
t-5 → t-3	+0.41	2.97	***
t-3 → t-1	+0.48	4.93	***
t-1 → t <i>all bonds</i>	+0.07	0.89	—
t-1 → t <i>offered only</i>	+0.44	2.34	**
t → t+5	-1.00	-4.96	***
t+5 → t+10	+0.45	2.73	***
t-3 → t+10 <i>net</i>	-0.05	-0.17	—
t-3 → t+20 <i>net</i>	+0.08	0.21	—

NTN-B, basis points per standard deviation of offer size. Date fixed effects, cubic maturity control, SE clustered by date. Survives date × maturity-bucket fixed effects (t=2.88) and is stable across 2011-15, 2016-20 and 2021-26. A permutation test reassigning offered status at random within each date over 500 draws puts the observed effect at 0.495bp against a null centred on -0.008 with sd 0.147. *** p<0.01, ** p<0.05.

And then it is gone. Five days after the auction the cumulative effect has fallen by a full basis point, more than undoing the run-up. The path is not monotone — the following week gives back 0.45 basis points of that overshoot — but the destination is unambiguous: measured from three days before to ten days after, the net is -0.05 basis points, and at twenty days it is +0.08. This is dealer inventory pricing around a predictable event — the cost of warehousing a known quantity of paper for a few days — not a change in what investors require to hold duration.

PART TWO: THE NOMINALS LEAVE NOTHING ACROSS DAYS


Running the same design on the DI curve, with each auctioned LTN and NTN-F mapped onto its duration point, produces no usable result at daily frequency. The placebo window fires in most specifications and the estimates move with the choice of duration bucket. The reason is the calendar established above: a Thursday nominal auction has Tuesday's linker auction inside its pre-event window. At daily resolution the two events cannot be separated, and no amount of specification fixes a contaminated control period.

That is what sends the question intraday, where the windows shrink to minutes and the two auctions no longer overlap.

PART THREE: INSIDE THE SESSION

The auction is plainly visible in DI volatility, and the event window can be located from the data rather than assumed. On Thursdays the mean absolute five-minute change runs 97 per cent above the non-auction baseline at 11:30 and 110 per cent

above at 11:45. Tuesdays show no comparable DI signature, which is what one would expect: the linker is a real-rate instrument and DI is nominal.



Mean absolute five-minute change in DI rates by time of day, across the seven most liquid contracts. The nominal auction result window is visible as a clean excess between 11:25 and 12:00 on Thursdays.

Splitting the session at that window and running the cross-sectional test gives point estimates in exactly the shape the trade implies — a run-up into the auction, a move on the release, and a give-back afterwards — and none of them is statistically distinguishable from zero with 30 auction sessions.


SESSION SEGMENT	COEFFICIENT	S.E.	95% UPPER BOUND
Open → 11:25 <i>run-up</i>	+0.175	0.247	0.658
11:25 → 12:00 <i>release</i>	+0.118	0.169	0.448
12:00 → close <i>unwind</i>	-0.192	0.355	0.503
Open → close <i>net</i>	+0.101	0.469	1.021

DI, basis points per standard deviation of duration-matched offer size, 407 contract-days across 30 nominal auction sessions. Date fixed effects, quadratic maturity control, SE clustered by date. The same regression on non-auction sessions returns nothing.

The instinct is to read that table as a null and stop. That would be the wrong conclusion to draw from an underpowered sample. The useful question is not whether the effect differs from zero but whether it differs from the cost of capturing it, and for that comparison an imprecise estimate is still informative.

THE COST SIDE

The five-minute panel prices its own execution. The Roll estimator, which backs an effective round-trip spread out of the negative autocovariance in consecutive price changes, gives a median of 0.42 basis points across DI contracts — and it is strikingly flat, running from 0.29 at the front to 0.49 in the belly, and sitting at 0.44 in the four-to-seven-year band where NTN-F duration actually lives.




Roll implied effective round-trip spread by DI contract maturity, estimated from five-minute rate changes over 136 sessions. The shaded band marks where NTN-F duration sits at the tenth to ninetieth percentiles.

The flatness matters more than the level. It means there is no point on the curve where execution is cheap enough to change the arithmetic — the trade cannot be rescued by maturity selection.

The expected concession is smaller than the cost of harvesting it, and the harvesting cost measured here is only a floor.

THE ARITHMETIC

A typical-sized offering produces 0.13 basis points of run-up. An offering two standard deviations above the mean produces 0.35. One DI round trip costs 0.42, and that is before the bond leg — bidding, receiving the paper, and unwinding cash bonds whose bid-offer is wider than DI's. The measured cost is therefore a floor on the hurdle, not the hurdle.



Expected DI run-up as a function of offer size, against the measured round-trip execution cost. The lines cross at 2.4 standard deviations, which is the 96th percentile of observed offerings.

On the central estimate the two lines cross at 2.4 standard deviations above the mean offering — the 96th percentile of everything the Treasury put up in this sample. So the trade does clear its costs, but only for the very largest auctions, and only about one offering in twenty-five is that large. For everything else the expected concession does not pay for a single leg of the execution.

Putting the two legs together: the hedge earns an expected 0.13 basis points at a typical offering against 0.42 of round-trip cost, and the bond is acquired at fair value rather than at a discount, with the only positive expectation coming from an overnight hold whose dispersion is seven times its mean. Neither leg pays for itself on central estimates.

One caveat has to sit alongside that. At the upper edge of the confidence band the breakeven falls to 0.6 standard deviations, which would make the trade profitable at a large share of auctions. The data are consistent with the trade being marginal and also consistent with its being worthwhile for anything above an average offering. What can be said firmly is that the central estimate does not clear the cost, and that establishing the case for the trade would require the effect to sit near the top of its confidence interval.

THE BOND LEG

Everything above prices the hedge. But the hedge is not where the money is supposed to be: the operation exists to acquire paper at the auction, and the concession that matters is the cut rate against where the bond trades in the secondary market. That is directly measurable for linkers, since the auction record carries the accepted rate and the daily panel carries the bond's own yield.

CUT RATE MEASURED AGAINST	MEAN (BP)	MEDIAN	S.D.	SHARE > 0	T
Previous session close	+0.29	+0.22	6.37	53%	2.06
Same-day close	+0.09	+0.31	4.80	54%	0.89
Next-day close	+1.27	+1.57	9.05	62%	6.40

NTN-B, 2,087 auction lines 2011–2026. Positive means the auction cleared at a higher yield than the reference close – the buyer acquired paper cheap.

The first row is the run-up already documented, seen from the other side. The second is the one that matters for the auction itself, and it is essentially zero: measured against the same session's close, the cut clears within a tenth of a basis point of the secondary market. Participants do not buy cheap at the auction. They buy at market.

The third row is where a gain appears, and it is worth being precise about what produces it. Against the following session's close the cut is 1.27 basis points cheap, positive in 62 per cent of auctions. That is not a concession captured at the auction – it is the post-auction reversal, earned by holding the bond overnight and carrying one session of outright market risk. The standard deviation of that outcome is 9.05 basis points. An expected 1.27 against a dispersion of 9.05 is roughly one seventh of a unit of risk per trade, before paying the cash bond's bid-offer to get out.

CUT VS SAME-DAY CLOSE	N	MEAN (BP)	MEDIAN	T
Under 5 years	479	+0.29	+0.45	1.17
5 to 10 years	674	+0.39	+0.66	2.09
10 to 20 years	594	-0.22	0.00	-1.37
Over 20 years	340	-0.23	-0.08	-0.84
<i>2011–2015</i>	<i>688</i>	<i>-0.40</i>		<i>-2.14</i>
<i>2016–2020</i>	<i>552</i>	<i>+0.05</i>		<i>0.19</i>
<i>2021–2026</i>	<i>847</i>	<i>+0.52</i>		<i>3.76</i>


The auction cleared expensive to secondary in the first half of the sample and cheap in the most recent period, with the gain concentrated in the five-to-ten-year sector.

Two things stand out. The sector matters: the belly clears cheap and the long end does not, which is the opposite of where a duration-absorption story would put the concession. And the period matters more than anything else. Through 2011–2015 the auction cleared *expensive* to secondary by 0.40 basis points, a statistically clean negative. Since 2021 it has cleared 0.52 cheap. Whatever is being measured is a feature of the current placement environment rather than a structural property of auctions.

Participants do not buy cheap at the auction. They buy at market, and what looks like a concession is the following day's reversal, paid for with a session of outright risk.

A FOOTNOTE ON PLACEMENT

One descriptive fact deserves recording because it cuts against the intuition that concessions are demanded because paper must be absorbed. Across all NTN-B sale lines from 2011 to 2026, the median fill ratio — quantity accepted over quantity offered — is 0.45. Three quarters of offering lines are only partially filled and 18.1 per cent sell nothing at all. The Treasury routinely does not place what it puts up.



Distribution of the fill ratio on NTN-B offering lines, 2011–2026, $n=3,587$. Values are clipped at 1.0.

Despite that, the fill ratio does not predict yield moves in any window tested. Failing to place paper is not priced beyond whatever the announced quantity already did — which fits a world where the concession compensates dealers for short-horizon inventory risk rather than reflecting the market's appetite for duration.

WHAT THIS DOES NOT SETTLE

The intraday sample is 30 nominal auction sessions, which is what a 140-trading-day history limit allows. The standard errors above are a direct consequence, and the honest position is that the run-up cannot be distinguished from either zero or from a figure that would justify the trade. Repeating the pull every eight to ten weeks over a year would roughly halve them, and that is the only thing that resolves it — the history cannot be recovered retrospectively.

Date fixed effects absorb the curve-wide move by construction, so the volatility spike at 11:30 and the absence of a duration-localised effect are not in conflict. The auction moves the whole curve; what cannot be detected here is a component concentrated at the auctioned durations. If that survives a larger sample it is a

finding in its own right, and it would sit alongside the linker result rather than against it.

Finally, the design has power against auction-sized shocks only. Offerings are small relative to the outstanding stock at any maturity, so nothing here rules out a level effect from large, persistent shifts in supply.

REFERENCES

Lou, D., Yan, H. and Zhang, J. (2013). Anticipated and Repeated Shocks in Liquid Markets. *Review of Financial Studies* 26(8).

Greenwood, R. and Vayanos, D. (2014). Bond Supply and Excess Bond Returns. *Review of Financial Studies* 27(3).

Roll, R. (1984). A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *Journal of Finance* 39(4).