BTRM

The Certificate
of Bank Treasury
Risk Management

A Bootstrapped SOFR Curve with Rating-Based Corporate Spreads

BTRM Working Paper Series #28

Nicolò Lavaroni
Analyst, BA Labs

Tommaso Donda
Analyst, BA Labs

Jan Delegos, CFA
Head of Risk, Sky Frontier Foundation

August 2026

A Bootstrapped SOFR Curve with Rating-Based Corporate Spreads

A reproducible term-structure model built from freely available end-of-day data

Nicolò Lavaroni[1] Tommaso Donda[1] Jan Delegos, CFA[2]

[1] *Analyst, BA Labs* [2] *Head of Risk, Sky Frontier Foundation*

Working paper. Data as of 2026-07-24

Abstract

This note builds a proxy for a US dollar risk-free discount curve from freely available, end-of-day data, then adds an approximated credit-spread layer by rating. The proxy risk-free curve is bootstrapped from SOFR instruments (overnight fixings, futures, and overnight-indexed swaps) and reprices each of them exactly. We explain the mathematics of the construction, then discuss factors that are easy to get wrong when the only data you have is free and public: matching quoting conventions, keeping the Treasury basis internally consistent, and confirming the curve admits no arbitrage. An illustrative calibration and a full account of the limitations follow.

Keywords: *SOFR; yield-curve bootstrapping; discount factors; overnight-indexed swaps; credit spreads; option-adjusted spread; term structure.*

1. Introduction and purpose

The model has two layers. A risk-free curve comes from published SOFR-based, liquid and tradeable instruments [6], and a set of credit spreads by rating sits on top of each respective term. Together they give discount factors and forward rates out to ten years. Most of what follows concerns the risk-free curve and how to build it from data anyone can pull for free. The credit layer is there to show how a ballpark spread is added, not to study any particular borrower.

We kept the construction auditable. The risk-free part is pinned to traded instruments and reprices them exactly, with each assumption written down.

The term structure must be arbitrage-free: the return from investing at the derived zero-coupon rate to any maturity must equal the return from investing at the shorter derived zero-coupon rate and then rolling at the corresponding forward rate for the remaining period. For this reason, any derived discount rates, forward rates, zero rates must be consistent with the spot curve rather than assumed independently.

When constructing the curve, it is important to adjust for differences in day-count conventions between the instruments, the compounding logic and other factors that can produce minor, yet scalable errors.

2. The SOFR risk-free curve

Bootstrapping solves for the discount factors that reprice a set of market instruments exactly, one maturity at a time from the short end out [1, 2]. Instead of fitting a smooth curve through the quotes, it inverts them. The method is standard [3] and the implementation uses QuantLib [1, 4].

Three kinds of instruments each cover their own stretch of the curve. The overnight SOFR fixing [6], together with the realised fixings of the current month, anchors the front. One- and three-month SOFR futures reach about two and a half years. Overnight-indexed swaps carry the curve from three to ten years through their par rates. The free swap feed stops at ten years, so the curve stops there. We do not extrapolate past the last real quote.

3. Mathematical formulation

The calculated discount curve is a set of discount factors DF(0,T), that represent the today's value of one unit paid at T. From them we read the continuously-compounded zero rate z(T) and the instantaneous forward rate f(t):

$$DF(0,T) = \exp(-z(T)\cdot T), \qquad f(t) = -d \ln DF(0,t) / dt.$$ [1]

Bootstrapping picks the discount factors so that each calibrating instrument prices to its quote. An instrument maturing at T_n depends only on the factors up to T_n, so they are solved in order from the short end. Each instrument adds one unknown, $DF(0,T_n)$, found by a single root solve with the earlier factors held fixed. That ordering is why the fit is exact, rather than a least-squares approximation.

Assuming the arbitrage-free rule holds, each subsequent instrument's derived rate is calculated from the previously obtained rates. For example, a SOFR overnight-indexed swap, at the market midpoint, is worth zero at par (the net present value of both legs should be equal), which pins its fixed rate. When the swap starts, the fixed rate R is set so that what is paid on the fixed leg has the same present value as what is expected to be received on the floating leg. The floating side comes from compounding the actual daily projected SOFR rates that occur over the life of the swap:

$$R = (DF(0,T_0) - DF(0,T_n)) / \Sigma_i \tau_i DF(0,T_i) .$$

The denominator is the annuity, τ_i the accrual of period i, and the numerator the value of the daily-compounded floating leg.

If we'd look at the SOFR Term Rates, we can use them to bootstrap the data points up to 12 months. Subsequently, we can repeat the bootstrapping exercise to obtain implied discount and forward rates from the futures contract prices. It is worth noting that while the Term SOFR Rates use daily compounding, an exception lies in the market convention for 1M rates, which uses a simple arithmetic average instead.

Between nodes we interpolate log-linearly on discount factors, so ln DF(0,t) is linear in t and the forward f(t) is constant on each interval. The interpolation is a commonly used method of approximation, given that not each exact maturity (such as instruments maturing in exactly 2 years 4 months and 3 weeks from now) has a traded, liquid market. We then check if we can reproduce any spot rate from our calculated discount and forward rates. That gives a clean no-arbitrage test. The curve admits no arbitrage exactly when the discount factors strictly decrease, which the build checks after every step.

$$DF(0,T_{i+1}) < DF(0,T_i) \quad \Leftrightarrow \quad f(t) \geq 0.$$ [2]

---

[1] Verify the compounding and calculation methodology of the input instrument before using it in the model

[2] Some savvy readers would notice that this rule didn't hold in the negative rates environment of the 2010s, but this is a story for another time, as that age made a large chunk of assumptions used for rates modelling null and void (otherwise termed, "Alice in Wonderland finance", as Professor Moorad Choudhry noted).

If we want to produce forward rates, for each pair of maturities, all we need to do is to take their corresponding discount factors, use their ratio, invert and annualize it. Such a rate would inform us about the market-implied rate between T_i and T_{i+1}. For example, a geometric average of an annualized spot rate T(0,i) and an annualized forward rate T(i,i+1) would give us an implied, annualized spot rate for T(0,i+1).

4. The credit overlay

In this section we add the credit risk spread on top of the calculated rates, split by rating. The market data gives one option-adjusted spread (OAS) per rating, blended across maturities [7]. This number is simply a “ballpark” level of a credit spread for the entire broad basket of credit-exposed instruments. For this reason, any more specific sector/geography/product type credit spread will vary. That said, adding that ballpark single number into a term structure has three elements:

- The level is the rating's OAS, quoted over Treasuries;
- The slope comes from the investment-grade index's maturity buckets, which we reuse for every rating;
- The third piece is the Treasury-to-SOFR basis, which moves the spread from a pickup over Treasuries to a pickup over SOFR. Getting those pieces to combine consistently is noted next in Section 5.

Rating	OAS over Treasuries (bp)
AAA	43
AA	57
A	66
BBB	98
BB	166
B	294
CCC	991

Table 1. Blended index OAS by rating, as of July 2026.

5. Data selection, biases, and workarounds

Compounding is easy to get wrong here. OAS and the Treasury-SOFR basis are quoted semi-annually, but the discount curve holds rates continuously. A spread s quoted semi-annually equals $2 \cdot \ln(1 + s/2)$ continuously, and the gap grows with the spread: under a basis point at investment grade, about 24 bp at a ten-per-cent (CCC) level. Added as a continuous spread it would overstate the wider levels, so for this reason we add it in the semi-annual terms it is quoted in.

The basis needs the same care. The Treasury inputs are constant-maturity par yields, and in a sloping curve environment a par yield and a zero rate are not the same number. To give a simple example, a zero rate is a rate we calculate for a theoretical zero-coupon bond; i.e. the entire cash flow occurs at the maturity, so only changes of the discount rate at the maturity would impact the bond’s price. However, Treasury securities and almost all other commonly traded mid- and long-term bonds pay regular coupons (quarterly, semi-annually, etc). For this reason, a change in value of a discount factor of any corresponding coupon would impact the bond’s pricing.

In a more abstract way, a vanilla, coupon-paying Treasury can be replicated as a basket of cash-flow weighted, zero-coupon bonds of matching maturities. For these reasons, we bootstrap a Treasury zero curve from the par yields first (for a par bond priced at 100, $1 = (y/2)\cdot\Sigma_j\ DF(0,T_j) + DF(0,T_n)$), then take the basis zero against zero.

It is worth noting that interest rate futures don't reprice to the last decimal, and two effects result from this. These effects are easy to conflate, so we present them apart.

One is convexity: a margined future generally prices a touch above the equivalent forward. We handle that separately, by setting the convexity adjustment to zero, and return to it below.

The other is a small numerical residual. A piecewise-flat forward rate can't exactly reproduce a rate that settles on an average or a daily compounding basis, so up to about 0.15 bp of difference is left on the futures. Note that is not an arbitrage violation, as the discount factors strictly decrease. With this condition satisfied, the curve admits no arbitrage by construction, and the swaps reprice with expected precision. The residual shows up on a few monthly contracts and doesn't grow with maturity, which points to interpolation rather than a funding effect. We keep log-linear because a log-cubic scheme won't converge on this set.

It is also worth adding that for the same maturities, there exists a basis (pricing differential) between discount factors obtained from treasuries and from swap rates. This is due to market-specific factors (one being a spot asset, the other being a derivative, one being a collateral/investment item, the other being mostly a hedging instrument). This basis may create some small dislocations when trying to piece together a spread-over-treasuries on top of a SOFR-swap-derived curve.

We left the futures convexity adjustment at zero. Doing it properly needs a volatility input we cannot obtain easily, and the bias is small (under a basis point in the first year, a few by two and a half years) [5]. One thing we do not leave to chance is no-arbitrage: log-linear interpolation stays arbitrage-free only while the discount factors keep falling, so the build tests that on every run and stops if a forward turns negative.

6. Results and validation

As of writing this document, the SOFR curve is mildly humped: up through the first year, roughly flat from two to five years, then rising toward ten. Investment-grade all-in rates land within about a point of the risk-free curve, and the high-yield rungs spread out quickly (see Figure 1).

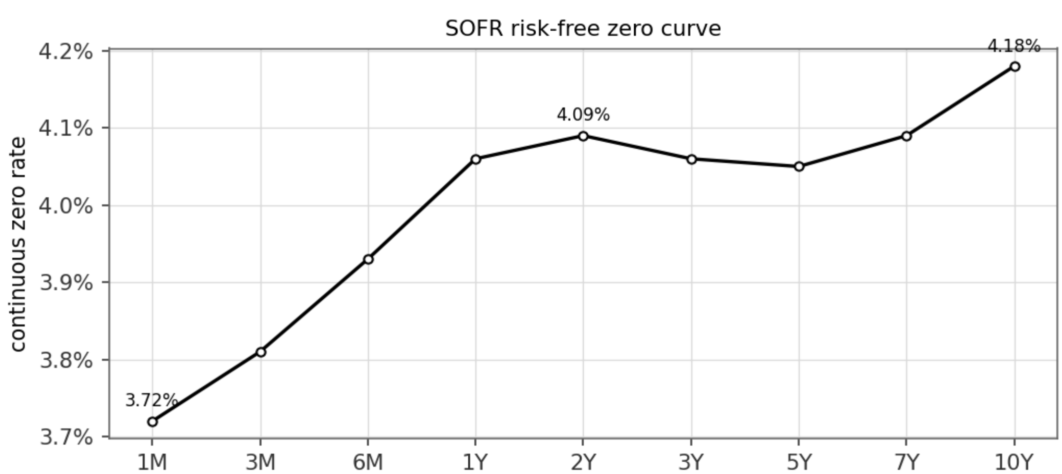

Figure 1. SOFR risk-free zero curve, continuous compounding, as at 24/7/2026.

Tenor	SOFR zero rate
1M	3.72%
3M	3.81%
6M	3.93%
1Y	4.06%
2Y	4.09%
3Y	4.06%
5Y	4.05%
7Y	4.09%
10Y	4.18%

Table 2. SOFR continuous-compounded zero rates as at 24/7/2026.

Tenor	AAA	AA	A	BBB	BB	B	CCC
1M	27	35	40	59	100	176	583
3M	30	39	44	63	103	179	587
6M	32	40	46	65	105	181	588
1Y	24	33	38	57	97	173	580
2Y	43	51	57	76	116	191	598
3Y	54	64	70	93	141	232	717
5Y	70	83	91	120	182	297	913
7Y	86	101	110	144	216	350	1068
10Y	96	113	123	161	241	390	1182

Table 3. Credit spread over SOFR by rating and tenor (basis points).

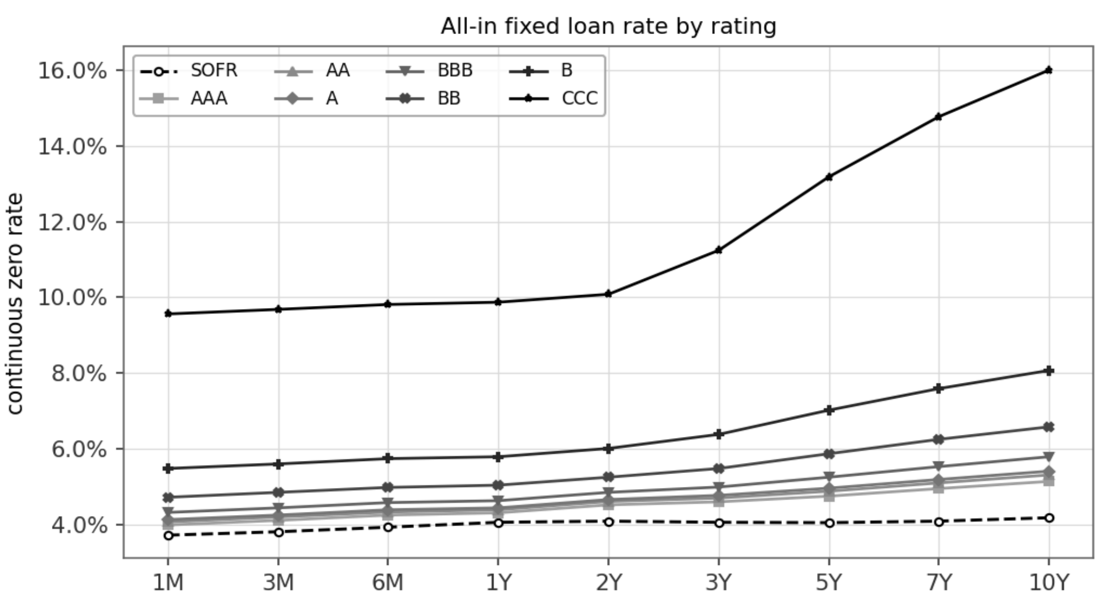

Figure 2. All-in fixed rate by rating; SOFR shown for reference.

Tenor	SOFR	AAA	AA	A	BBB	BB	B	CCC
1M	3.72	3.99	4.07	4.13	4.32	4.72	5.48	9.56
3M	3.81	4.11	4.20	4.25	4.44	4.85	5.60	9.68
6M	3.93	4.25	4.33	4.39	4.58	4.98	5.74	9.81
1Y	4.06	4.31	4.39	4.44	4.63	5.04	5.79	9.87
2Y	4.09	4.52	4.61	4.66	4.85	5.25	6.01	10.08
3Y	4.06	4.60	4.70	4.77	4.99	5.48	6.38	11.24
5Y	4.05	4.75	4.88	4.96	5.25	5.87	7.02	13.18
7Y	4.09	4.95	5.10	5.19	5.53	6.25	7.59	14.77
10Y	4.18	5.14	5.31	5.41	5.79	6.58	8.07	16.00

Table 4. All-in fixed rate by rating and tenor (per cent, continuous compounding).

The overnight rate and its history come from the New York Fed [6], the futures from the CME through a delayed public feed, the swap rates from Pensford, and the corporate spreads and Treasury yields from FRED [7]. Inputs are range-checked before each build, and the finished curve is checked to reprice its inputs and to admit no arbitrage.

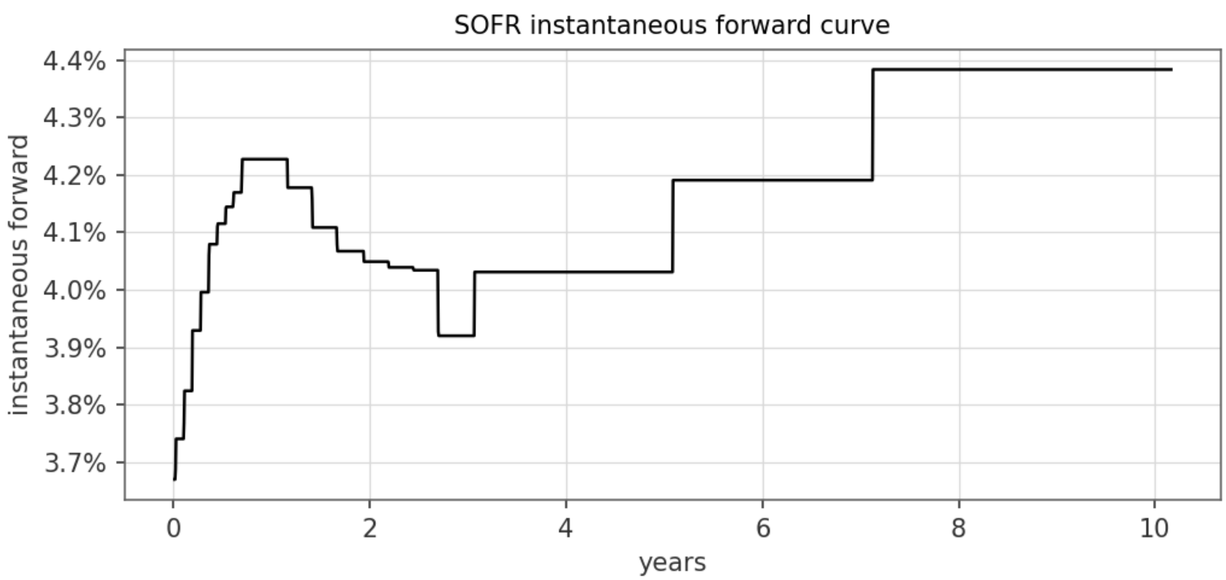

Figure 3. SOFR instantaneous forward curve shown for reference.

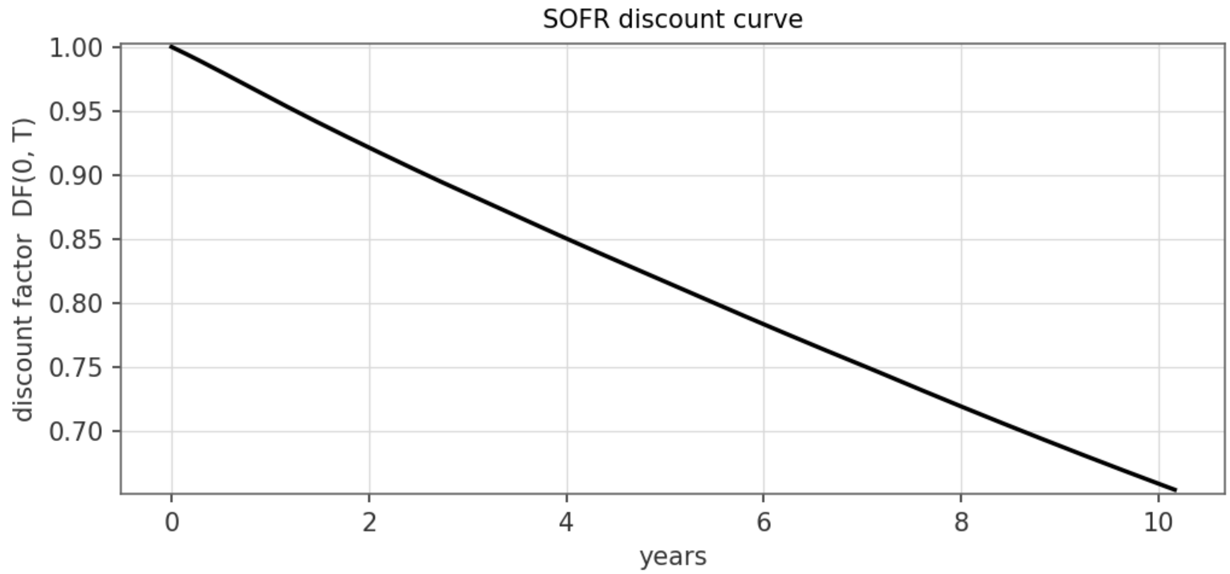

Figure 4. SOFR discount curve shown for reference.

7. Usage under known limitations

We note the following:

7.1 Modelling and data

- The curve stops at ten years (Section 5) and is not extrapolated past the last node.
- Log-linear interpolation gives piecewise-flat forwards and a sub-basis-point futures residual, which we checked is an interpolation effect (Section 5).
- The futures convexity adjustment is left at zero, a documented approximation.
- Turn-of-year and quarter-end spikes in the overnight rate are not modelled.
- The data is free, delayed, indicative, and pulled from several providers at slightly different times.
- The credit overlay reuses one investment-grade maturity shape for every rating; the high-yield long end is the weakest part, and the spreads are index averages, not any single borrower's.

7.2 Validation and reproducibility

- Repricing shows internal consistency. The instrument conventions match the SOFR-OIS standard but are not reconciled against a second source.
- Nothing has been cross-checked against a second pricing engine or an analytic benchmark.
- The data rests on single free providers; the fallback keeps a run alive through an outage, but the result is worse.

Three things would raise confidence without impacting the method: a test suite on a frozen input fixture with a golden-master reprice, a saved snapshot of the raw inputs so a run reproduces exactly, and a cross-check against a second pricing engine.

Given awareness of all the above, with care we suggest that our curve building approach may be used to determine relative value assessment, from trading or investment perspective, in credit-risky markets that are denominated in USD or pegged-USD.

References

[1] Ballabio, L. *Curve bootstrapping.* The QuantLib Guide. https://www.quantlibguide.com/Curve%20bootstrapping.html

[2] Keogh, K. *Bootstrapping a SOFR curve.* https://kevindkeogh.com/posts/bootstrapping-a-sofr-curve/

[3] Ametrano, F. M., and Bianchetti, M. (2013). *Everything You Always Wanted to Know About Multiple Interest Rate Curve Bootstrapping but Were Afraid to Ask.* Working paper.

[4] Ballabio, L. *Implementing QuantLib.* Leanpub.

[5] Hull, J. C. *Options, Futures, and Other Derivatives.* Pearson.

[6] Federal Reserve Bank of New York. *Secured Overnight Financing Rate (SOFR).* https://www.newyorkfed.org/markets/reference-rates/sofr

[7] Federal Reserve Bank of St. Louis. *FRED: ICE BofA US Corporate and High Yield Option-Adjusted Spreads.* https://fred.stlouisfed.org