

The BAM Curve, Priced and Hedged

Duration, Convexity, Key-Rate Risk, and an Out-of-Sample Hedge Test on Moroccan Sovereign Debt

Yirviel Somé

FMVA® · BIDA® · MSc in Finance candidate, ENCG Casablanca

Corresponding author: yirviell.some@gmail.com

Version 2.1 — 21 October 2026

Curve and hedges: 16 October 2026. Backtest: 1 April – 16 October 2026.

ABSTRACT

The methods here are standard. Nelson–Siegel–Svensson fitting, key-rate duration (Ho, 1992), and duration–convexity hedging are textbook material; the contribution is a reproducible application to Moroccan sovereign data, with the execution and out-of-sample evidence that a single-day exercise cannot supply. Using the 16 October 2026 BAM reference fixing, I fit an NSS curve to ten tenors. The in-sample RMSE is 3.85 bp, but six parameters on ten points leave four residual degrees of freedom; leave-one-out RMSE is 30.6 bp, so I also report a four-parameter Nelson–Siegel fit (4.23 bp in-sample, 8.87 bp leave-one-out) and check every hedge under it. The eight-line book is worth 81.4M MAD, with modified duration 6.99, curve DV01 58,386 MAD/bp, and dollar convexity 8.21B. A duration-only short of 76.9M in the 8.68-year line zeroes DV01 but leaves 3.40B of convexity and a –42.9k MAD/bp key-rate exposure at that tenor. A two-instrument duration–convexity hedge zeroes both but leaves +369k MAD to a +100 bp belly move at 12.76 years. A six-instrument key-rate hedge (3.5 to 28.5 years) cuts residual DV01 to zero and belly P&L to +86k, and is the only hedge that respects a 10M-per-line size cap (95% executable versus 13% for the duration short). In a 106-day out-of-sample backtest, the key-rate hedge cuts daily P&L volatility from 112k to 6k MAD; the duration-only hedge leaves 50k and still has the largest residual among the hedges. The lesson is that DV01 is a scalar summary of a curve, not a hedge: a parallel-neutral book is not a curve-neutral book, and a hedge that cannot be put on is not a hedge. Every figure in the paper is derived in Appendix A, where the formulas and the worked arithmetic are set out.

Keywords: Bank Al-Maghrib; Moroccan sovereign curve; Nelson–Siegel–Svensson; duration; convexity; key-rate risk; hedging; out-of-sample backtest

JEL classification: E43; G12; G15

NSS FIT RMSE 3.85 bp	PORTFOLIO VALUE 81.4M MAD	MODIFIED DURATION 6.99
DOLLAR CONVEXITY 8.21B	CURVE DV01 58,386 MAD/bp	KEY-RATE 6 RESIDUAL 0 DV01 / –0.25B

This is a risk-management exercise on public reference data, not a recommendation on any security. BAM reference yields are indicative secondary-market levels, not executable quotes. The author reports no conflict of interest. All errors are his own.

1. Introduction

Morocco's sovereign curve is where the country's rate risk is actually priced, and Bank Al-Maghrib (BAM) publishes a clean daily reference: the *Taux de référence des bons du Trésor* on the secondary market, with a CSV export per value date. The curve is a better laboratory than the equity index for a risk-management note: it has a fixed set of tenors, a published fixing, and a cash-flow structure that can be repriced exactly.

The methods are deliberately conventional. Nelson and Siegel (1987) and Svensson (1994) give the parametric curve; Ho (1992) gives key-rate duration; duration–convexity hedging is standard fixed-income practice. I do not claim a new estimator or a new hedge. The contribution is the Moroccan application, and four things that are usually missing from a one-day case study:

- **Fit honesty.** An NSS curve has six parameters and the BAM block has ten points. I report leave-one-out error, parameter standard errors and a four-parameter Nelson–Siegel robustness check, rather than presenting the in-sample RMSE as if it were out-of-sample accuracy.
- **A key-rate hedge, not just a recommendation.** The paper builds a three-line and a six-line key-rate hedge, shows what each leaves behind, and compares them with the duration and duration–convexity hedges.
- **An out-of-sample test.** A 106-day rolling backtest on daily BAM fixings, with rebalancing, transaction costs and financing, tests whether the hedges survive a moving curve.
- **Execution realism.** A 76.9M short in one line is 94% of the book. I show how each hedge degrades under a per-line size cap and what financing costs per month.

The note is still narrow. I do not forecast the curve, and I do not claim that one hedge is correct. I ask a smaller question: **when two books have the same DV01, what else do they own?** The answer is duration, convexity, and a set of key-rate exposures — and the difference between them is the difference between a hedge and a bet.

2. The BAM curve

The 16 October 2026 curve is upward sloping and steep at the front end: 10 observed points from 0.08 to 28.53 years, 2.13% to 4.38%. I fitted a Nelson–Siegel–Svensson (NSS) curve by grid-searching the two decay constants and solving for the four factor loadings linearly at each pair (Nelson and Siegel, 1987; Svensson, 1994).

Table 1. BAM reference yields and the NSS fit, 16 October 2026.

Maturity	Years	BAM yield	NSS fit	Residual
16 Nov 2026	0.08	2.130%	2.139%	+0.95 bp
15 Feb 2027	0.33	2.200%	2.176%	−2.37 bp
16 Aug 2027	0.83	2.290%	2.315%	+2.55 bp
14 Feb 2028	1.33	2.470%	2.461%	−0.90 bp
15 Apr 2030	3.50	2.890%	2.879%	−1.08 bp
20 Oct 2031	5.01	3.090%	3.091%	+0.08 bp
18 Jun 2035	8.68	3.480%	3.528%	+4.82 bp
18 Jul 2039	12.76	3.950%	3.865%	−8.50 bp
14 Aug 2045	18.84	4.090%	4.150%	+5.96 bp
19 Apr 2055	28.53	4.380%	4.365%	−1.51 bp

The in-sample RMSE is **3.85 bp**, with the largest miss at 12.76 years. That number is less impressive than it looks. Six parameters on ten points leave four residual degrees of freedom; the condition number of the normal equations is about 1.1×10^7 ; the residual standard error is 6.1 bp; and the parameter standard errors are large (b_2

1.03%, b_3 0.79%, τ_1 0.35, τ_2 0.52). The honest test is leave-one-out: the NSS LOOCV RMSE is **30.6 bp**, driven by the 28.5-year point (85.5 bp) and the three-month bill (26.3 bp). A four-parameter Nelson–Siegel curve has a slightly worse in-sample fit (4.23 bp) but a much better LOOCV (8.87 bp). The extra NSS curvature terms are not identified by ten points.

Table 2. The fit is thin. In-sample error understates out-of-sample error by a factor of eight.

Model	Params	In-sample RMSE	Max miss	LOOCV RMSE
NSS	6	3.85 bp	8.50 bp	30.58 bp
Nelson–Siegel	4	4.23 bp	9.24 bp	8.87 bp

I keep the NSS curve for continuity with the curve-fitting literature, but I do not treat it as a forecasting model. Every hedge below is re-solved under the four-parameter NS curve; the duration short moves from 76.94M to 77.03M, residual convexity from 3.39B to 3.40B, the six-line key-rate residual DV01 stays at zero, and the +100 bp belly P&L moves from +86.0k to +86.0k. The conclusions are not a product of the NSS overfit.




Figure 1. BAM reference yields and the NSS fit, with residuals in basis points. The fitted curve is smooth through the front end and undershoots the 12.76-year reference point by 8.5 bp. All pricing downstream uses the fitted curve, not the observed points.

3. The book

The portfolio is deliberately simple: eight of the reference lines (the two shortest bills are excluded), **10M MAD notional each**, annual coupons, ACT/365, coupon equal to the reference yield at the matching maturity. Prices are dirty — the present value of every remaining cash flow, with no accrued interest stripped out. Each bond is priced off the whole curve, not off a single yield, so the YTM in the table is the one rate that reproduces the curve price. The arithmetic behind every number in this paper is set out in Appendix A.

Table 3. Bond-level analytics. DV01 is curve DV01, in MAD per basis point.

Line	Years	Price	YTM	Mod. dur.	Convexity	Curve DV01
2.29% Aug 2027	0.83	100.34	2.343%	0.81	1.5	836
2.47% Feb 2028	1.33	101.62	2.490%	1.28	2.9	1,329
2.89% Apr 2030	3.50	101.37	2.910%	3.24	14.0	3,378
3.09% Oct 2031	5.01	102.91	3.120%	4.44	25.4	4,712
3.48% Jun 2035	8.68	100.71	3.535%	7.30	65.2	7,589
3.95% Jul 2039	12.76	102.07	3.836%	9.81	120.2	10,337
4.09% Aug 2045	18.84	100.78	4.080%	12.88	216.6	13,334
4.38% Apr 2055	28.53	104.27	4.246%	15.93	365.7	16,872
Portfolio				6.99	102.3	58,386

The book is worth **81.4M MAD**. Its modified duration is 6.99 and its weighted-average YTM convexity is 102.3, a dimensionless summary. The number used in the hedge is the **curve dollar convexity: 8.21B MAD** per unit of (decimal yield)², the second derivative of the repriced book. The two are not interchangeable: 102.3 is a per-bond average, 8.21B is the dollar exposure. When this note says “convexity” in a hedge context, it means the dollar number.

One methodological point matters for everything that follows: **DV01 here is curve DV01, not modified duration × market value**. I bump every fitted zero rate by 1 bp and reprice all cash flows. For the 8.68-year line the two measures differ by roughly 3%. On a curve this steep, the YTM duration is a useful summary, but it is not the risk.

4. The duration hedge

The obvious hedge is one instrument. I used the 8.68-year 3.48% line and solved for the notional that makes portfolio DV01 zero:

Short 76.9M MAD of the 8.68Y line → residual DV01 ≈ 0.

That looks like a clean hedge. It is not. The chart below splits the portfolio's DV01 by key tenor — a 1 bp triangular bump at each observed point, with the curve reprice reflected through the bond's cash flows (Ho, 1992).




Figure 2. Key-rate DV01 before and after the duration hedge, and after the six-line key-rate hedge. Duration zeroes the total by concentrating the exposure at 8.7Y; the key-rate hedge removes the shape.

The unhedged book has positive exposure at every tenor, rising to +12.8k at 18.8Y. After the duration hedge, the total is zero — but the 8.7Y point flips to **-42.9k MAD/bp** while the long end is untouched. The hedge did not remove rate risk; it concentrated it at one point on the curve, in the opposite direction. A 1 bp move at 8.7Y now costs more than a 1 bp move anywhere else gains.

Two other residuals matter:

- **Convexity.** The DV01-matched short leaves **+3.40B** of dollar convexity, 41% of the 8.21B book. The portfolio is more convex per unit of DV01 than the 8.7Y bullet, because it is a ladder across 0.8 to 28.5 years. At ± 100 bp the hedged book gains **+190k / +152k MAD** — positive in both directions. That is a long-volatility position nobody asked for.
- **Curve shape.** The single-instrument hedge has no opinion about slope; it simply shorts the 8.7Y point. A +100 bp steepener (short rates -100 bp, long rates +100 bp) costs the hedged book **-2.18M MAD**, while the unhedged book gains +501k. A +100 bp belly move at 12.76Y adds **+630k MAD**. The hedge can be worse than no hedge under the right curve move.

The execution problem is just as serious. A 76.9M short in one line is 94% of the book's market value and 7.7 times the book's 10M per-line holding. Section 7 shows that under a 10M per-line cap only 13% of this hedge can be placed, leaving 87% of the DV01 unhedged.

5. Duration plus convexity, and the shape that remains

To neutralize both DV01 and convexity, I used two instruments — the 5.01-year 3.09% line and the 18.84-year 4.09% line — and solved the 2x2 system:

Short 25.2M MAD of the 5.01Y line; short 34.9M MAD of the 18.84Y line.

Residual DV01 is 1 MAD/bp and residual dollar convexity is zero (to numerical precision). The stress table is the honest scorecard. I use a 12.76-year belly, not a 15-year hump, because 12.76Y is an observed key tenor and 15Y is not.

Table 4. Stress P&L by book, in thousands of MAD. The belly is a triangular bump at the 12.76-year reference tenor.

Scenario	Unhedged	Duration	Duration + convexity	Key-rate 6
Parallel -100 bp	+6,277k	+190k	+4k	-13k
Parallel +100 bp	-5,453k	+152k	-3k	-11k
Steeper +100 bp	+501k	-2,177k	+95k	+160k
Flattener -100 bp	-215k	+2,369k	+113k	-169k
Belly +100 bp (12.76Y)	-3,036k	+630k	+369k	+86k

At ± 100 bp parallel, the duration-plus-convexity book moves by **-2.9k / +3.8k MAD** — effectively flat. The duration-only book still gains in both directions, which is the convexity bet showing through. But a +100 bp belly at 12.76Y moves the two-instrument book by **+369k MAD**. The duration-only book is at +630k, and even the unhedged book loses 3.04M to that move. This is the trade: **a parallel-neutral book is not a curve-neutral book**. The two-instrument hedge buys parallel and convexity neutrality and sells shape neutrality.




Figure 3. Stress P&L for the unhedged, duration-hedged, duration-plus-convexity and six-line key-rate books. The unhedged book dominates the scale; the point is the residual shape risk in the other bars.

6. The key-rate hedge

The paper's earlier version said the next step was a key-rate hedge with three or more instruments, then stopped. This section builds it. With N instruments you can zero N key-rate exposures. The book has eight non-trivial knots (0.83Y to 28.53Y), so a three-instrument hedge is necessarily partial; the useful question is how much it leaves, not whether it is “the” hedge.

I solve a duration-neutral least-squares key-rate hedge: choose signed notionals that minimize the squared residual key-rate DV01 across all tenors, subject to the residual total DV01 being zero. The three-line version uses 12.76Y, 18.84Y and 28.53Y. The six-line version adds 3.50Y, 5.01Y and 8.68Y. Both are re-solved on the 16 October curve.

Table 5. The key-rate hedge comparison. Belly and steeper are P&L in thousands of MAD.

Hedge	Lines	Gross short	Residual DV01	Residual convexity	Belly +100	Steeper +100
Duration	8.68Y	76.9M	0	+3.40B	+630k	-2,177k
Duration + convexity	5.01Y, 18.84Y	60.1M	1	0	+369k	+95k
Key-rate 3	12.76Y, 18.84Y, 28.53Y	43.2M	0	-2.19B	+230k	+1,418k
Key-rate 6	3.50Y–28.53Y	63.0M	0	-0.25B	+86k	+160k

The three-line hedge is a real improvement on duration-only — it cuts gross short from 76.9M to 43.2M and belly P&L from +630k to +230k — but it leaves a negative convexity residual of -2.19B and a steeper exposure of +1.42M. It is not a finished hedge. The six-line hedge shorts roughly 10M in each of the six liquid lines. It zeroes DV01, cuts residual convexity to -0.25B, and reduces the belly move to +86k. Its residual key-rate exposure is the two short bills (0.83Y and 1.33Y), which are deliberately outside the hedge set: adding them would replicate the book, not hedge it.

The six-line hedge is also the most executable. At a 10M per-line cap it places 95% of its intended gross short and leaves 3.7% of DV01; the duration short places 13% and leaves 87% (Table 6). That is the practical argument for key rates: the hedge that looks best on a single DV01 number is the one that cannot be put on.

7. Size, liquidity and financing

BAM's reference table reports transaction volume as 0.00 for every line. It is a reference curve, not a depth measure, so I cannot prove that 77M is executable. I can show how each hedge degrades under a per-line cap. The caps below are illustrative, not observed: 10M is the book's own per-line holding, 20M is twice that, and 40M is four times.

Table 6. Size and financing. Gross/book uses the 81.4M MAD book; the cap column assumes each line can absorb 10M.

Hedge	Gross short	Lines	Gross / book	Cap 10M executed	Residual DV01 at cap	Financing / month (50 bp)
Duration	76.9M	1	1.95×	13%	50.8k (87%)	32k
Duration + convexity	60.1M	2	1.74×	33%	40.3k (69%)	25k
Key-rate 3	43.2M	3	1.53×	69%	17.8k (31%)	18k
Key-rate 6	63.0M	6	1.77×	95%	2.2k (3.7%)	26k

Financing is the other constraint. At a 50 bp net spread (the short rebate earns the policy rate, the repo spread is 50 bp), the six-line hedge costs about **26k MAD per month**, or 0.45 bp of book DV01. At a 2.75% all-in repo rate (the 2.25% policy rate plus 50 bp), the same hedge costs about **144k MAD per month**, or 2.47 bp of DV01. The duration short costs 32k / 176k per month at the same two assumptions. A hedge that saves less than 2–3 bp per month is not worth the financing; that is a real hurdle, not a rounding error.

Leverage is similar across the hedges — gross short is 1.5–1.9 times the book — but concentration is not. The duration hedge puts 7.7 line-equivalents into one line; the six-line key-rate hedge puts about one line-equivalent into each of six lines. If the book itself is executable, the key-rate hedge is the one that is executable at size.

8. Out-of-sample: does the hedge survive a moving curve?

A one-day hedge is a snapshot. To test the hedges as the curve moved, I downloaded 132 weekday BAM fixings from 1 April to 16 October 2026. Of these, 107 carry the full ten-point cross-section; I use those 107 fixings (106 daily intervals) and drop the sparse dates rather than fit a six-to-nine point curve whose decay constants are not identified. The book is a constant-maturity eight-line book; coupons are set from the first day's fitted curve, so no future information enters the position. At each rebalance the hedge is re-solved on that day's curve and held until the next rebalance. P&L is a full reprice, not a duration approximation, and includes coupon cash flows; without the coupon cash, the mechanical price drop at an ex-date would contaminate the hedge P&L. I charge 2 bp one-way on hedge turnover and a 50 bp annual net financing spread on the short notional.

The key-rate grid is fixed at the 16 October tenors for the whole backtest. The published set of reference maturities changes through time; re-basing the triangular bumps on each day's sparse set makes the hedge matrix ill-conditioned (on days with only six reference points it produced absurd notionals). That is a data caveat as well as a modelling choice.

Table 7. Rolling backtest, daily rebalance, net of transaction costs and financing. MAD, except annualized vol.

Book	Daily P&L std	Annualized vol	Worst day	95% VaR	Cumulative
Unhedged	112k	1.77M	−550k	−196k	+51k
Duration	50k	0.79M	−144k	−93k	−550k
Duration + convexity	8k	0.13M	−22k	−9k	+330k
Key-rate 6	6k	0.09M	−19k	−8k	+191k

The duration-only hedge cuts daily P&L standard deviation from 112k to 50k, but it still leaves the largest residual among the hedges; the reason is in the stress table, not the average: it is short 8.68Y, so a steepener or a belly move still moves the book. Duration-plus-convexity cuts the volatility to 8k. The six-line key-rate hedge cuts it to **6k**, a **95% reduction** relative to unhedged and an **89% reduction** relative to duration-only; its 95% VaR is −8k versus −196k unhedged. Weekly rebalancing leaves the risk profile essentially unchanged (6k standard deviation, +191k cumulative) and the transaction cost is 13.8k either way, so the operational recommendation is weekly, not daily.

This is not a profitability claim. The unhedged book ended +51k after carry; the duration hedge lost 550k after financing and costs; the duration-plus-convexity and key-rate books ended +330k and +191k. The point is the shape of the residual, not the sign of the cumulative P&L. A hedge that is flat on 16 October but has 50k daily volatility is still not a hedge a risk committee should accept.




Figure 4. Cumulative out-of-sample P&L, 1 April – 16 October 2026. Top: unhedged book. Bottom: the three hedges, net of 2 bp transaction costs and a 50 bp financing spread. The duration short leaves 50k of daily volatility; the key-rate hedge leaves 6k and stays closest to zero.

9. Limitations

- **The fit is thin.** NSS has six parameters and ten points. The in-sample 3.85 bp is not an out-of-sample number; LOOCV is 30.6 bp. The four-parameter NS robustness check gives the same hedges, but the curve level at the 28.5-year point is genuinely uncertain.
- **BAM reference yields are indicative, not executable.** They are a secondary-market reference, not a bid or an offer. The table reports zero transaction volume, so depth cannot be inferred from it.
- **The backtest is one sample.** It covers 6.4 months of 2026, not a full rate cycle. The book is constant-maturity and the coupons are set at the start; no taxes, coupon reinvestment, or intraday execution are modelled. Costs and financing are assumptions, not dealer quotes.
- **Short-sale constraints are not modelled.** The hedge assumes the lines can be borrowed. In practice, Moroccan sovereign repo and securities lending are thin; a desk may not be able to short the 28.5-year line at all.
- **Key-rate DV01 uses triangular bumps between reference tenors.** That is a convention, not a unique decomposition; a different bump shape changes the attribution, not the total DV01. The capacity caps are illustrative, not observed market depth.

The point of the project is not that one hedge is “correct.” It is that **DV01 is a scalar summary of a curve**, and two books with the same DV01 can have completely different risk. Duration gets you to zero. Key rates tell you what you still own. A size cap tells you whether you can own it at all.

10. Reproducibility

The engine, data and outputs live in the companion [mad-curve-lab](#) repository. The raw BAM CSV is committed under `data/bam_reference_2026-10-16.csv`; the daily backtest files are under `data/history/`; `outputs/` contains the analytics, the stress table, the backtest and the widget payload.

```
cd mad-curve-lab
python3 bam_curve.py selftest
python3 bam_curve.py analyze
python3 bam_curve.py analyze --model ns --out outputs_ns
python3 bam_curve.py fetch-history --start 01/04/2026 --end 16/10/2026
python3 bam_curve.py backtest
```

Data: Bank Al-Maghrib, *Taux de référence des bons du Trésor*, secondary market, value date 16 October 2026; daily history 1 April – 16 October 2026. Policy rate: Bank Al-Maghrib, 2.25%, accessed October 2026. Source page: [bkam.ma](#) → Marchés → Principaux indicateurs → Marché obligataire → Marché des bons du Trésor → Marché secondaire.

11. Conclusion

A duration hedge is a statement about the level of the curve. A duration-plus-convexity hedge is a statement about its curvature. Neither is a statement about its shape. On 16 October 2026, an 81.4M MAD Moroccan sovereign book could be made parallel-neutral with one short, and parallel-and-convexity-neutral with two. In both cases the residual risk moved, rather than disappeared: into the 8.68-year point in the first case, and into the belly in the second.

The key-rate hedge is the missing step. A three-line version is possible but leaves too much shape risk; a six-line version, sized at roughly 10M per line, zeroes DV01, cuts residual convexity to $-0.25B$, reduces the 12.76-year belly move to +86k, and survives a 10M-per-line capacity cap. In the out-of-sample backtest it cuts daily P&L volatility from 112k to 6k. That is what a hedge is supposed to do.

The practical implication for a Moroccan fixed-income desk is straightforward. Report DV01, but do not stop there. Report the key-rate profile, the dollar convexity residual, the P&L under a steepener, a flattener and a 12.76-year belly move, the financing cost, and the size at which each line can actually be shorted. A book that is flat to ± 100 bp parallel and +86k to a +100 bp belly is not flat. It is a set of curve views — and a hedge you cannot put on is not a hedge.

Appendix A. The calculations, step by step

Every number in this paper comes from four operations: fit the curve, price the cash flows, bump and reprice, solve for the hedge. This appendix writes out the formulas and the actual inputs, so the results can be checked by hand without running the code.

A.1 Fitting the curve

The Nelson–Siegel–Svensson curve, with t in years and y in decimal, is

$$y(t) = b_0 + b_1 f_1(t) + b_2 [f_1(t) - e^{-t/\tau_1}] + b_3 [f_2(t) - e^{-t/\tau_2}]$$

$$f_i(t) = (1 - e^{-t/\tau_i}) / (t/\tau_i)$$

At a fixed pair (τ_1, τ_2) the model is linear in the four loadings, so the design matrix and the least-squares solution are

$$X(\tau_1, \tau_2) = [1, f_1, f_1 - e^{-t/\tau_1}, f_2 - e^{-t/\tau_2}] \cdot b = \arg \min \|Xb - y\|^2$$

I grid-search τ_1 over $[0.30, 6.0]$ and τ_2 over $[\tau_1 + 0.35, 30.0]$, solving the linear system at every pair, then refine the six parameters jointly. The optimum is

$$b_0 = 0.047865 \cdot b_1 = -0.026479 \cdot b_2 = -0.016872 \cdot b_3 = -0.040431 \cdot \tau_1 = 0.5832 \cdot \tau_2 = 2.3495$$

Worked example at the 8.68-year reference point, where $t/\tau_1 = 14.8775$, $f_1 = 0.0672157$, $e^{-t/\tau_1} = 3.4 \times 10^{-7}$, $t/\tau_2 = 3.6929$, $f_2 = 0.2640450$ and $e^{-t/\tau_2} = 0.0248993$:

$$y(8.6767) = 0.047865 - 0.026479(0.0672157) - 0.016872(0.0672153) - 0.040431(0.2391457) = 0.035282 = 3.5282\%$$

which is the +4.82 bp residual of Table 1: 3.5282% against the 3.4800% reference. The in-sample error and its honest version, with $n = 10$ observations and $k = 6$ parameters, are

$$\text{RMSE} = \sqrt{(1/n) \sum (y_i - y_i^{\text{fit}})^2} = 3.85 \text{ bp} \cdot s = \sqrt{(1/(n-k)) \sum \varepsilon_i^2} = 6.09 \text{ bp}$$

$$\text{LOOCV} = \sqrt{(1/n) \sum (y_i^{\text{fit}(-i)} - y_i)^2} = 30.58 \text{ bp}$$

A.2 Pricing a bond

Coupons are annual, coupon dates are stepped back from maturity, and each t_i is actual days over 365. The dirty price is the present value of every remaining cash flow at the fitted zero curve, continuously compounded:

$$P = \sum CF_i \cdot e^{-z(t_i) t_i}$$

Applied to the 3.48% June 2035 line, that is Table A1. The yield to maturity is the single annual-compounded rate that reproduces the same price, and the three risk summaries follow from its discount factors $PV_i = CF_i/(1+y)^{t_i}$:

$$P = \sum CF_i / (1+y)^{t_i} \rightarrow y = 3.5354\%$$

$$D_{\text{mac}} = (\sum t_i PV_i) / P \cdot D_{\text{mod}} = D_{\text{mac}} / (1+y) \cdot C = (\sum t_i(t_i+1) PV_i) / (P (1+y)^2)$$

which for this line give 7.5530, 7.2951 and 65.17.

Table A1. The 3.48% June 2035 line priced off the fitted NSS curve, per 100 MAD of notional.

t (years)	Cash flow	z(t)	Discount factor	PV
0.671233	3.48	2.2667%	0.98490030	3.427453
1.673973	3.48	2.5486%	0.95823442	3.334656
2.673973	3.48	2.7492%	0.92912404	3.233352
3.673973	3.48	2.9049%	0.89877255	3.127728
4.673973	3.48	3.0450%	0.86734257	3.018352
5.676712	3.48	3.1782%	0.83492344	2.905534
6.676712	3.48	3.3039%	0.80204447	2.791115
7.676712	3.48	3.4210%	0.76903592	2.676245
8.676712	103.48	3.5282%	0.73628773	76.191054
Price				100.705489

A.3 Curve risk: DV01 and dollar convexity

The risk numbers come from bumping the whole curve, not from the yield to maturity:

$$\text{DV01}_i = -[P_i(z+h) - P_i(z-h)] / 2 \times n_i / 100$$

$$\Gamma_i = [P_i(z+h) - 2P_i(z) + P_i(z-h)] / h^2 \times n_i / 100 \cdot h = 1 \text{ bp} = 10^{-4}$$

P_i is the price of line i per 100 of face and n_i its notional in MAD; $P(z \pm h)$ reprices every cash flow against the curve shifted up or down by 1 bp, and Γ is the dollar convexity in MAD per unit of (decimal yield)², written Γ from here on, with Γ_i the same measure for line i . For the 8.68-year line at 10M MAD notional this gives DV01 = 7,588.90 MAD/bp and $\Gamma = 625.4$ MAD. The YTM shortcut $MV \times D_{\text{mod}} \times 10^{-4}$ gives 7,346.53, 3.2% lower, because a single YTM cannot reprice a cash-flow ladder on a steep curve.

A.4 Portfolio totals

The book-level numbers are sums, and one market-value weighted average:

$$MV = \sum n_i P_i / 100 \cdot DV01 = \sum DV01_i \cdot \Gamma = \sum \Gamma_i$$

$$D_{\text{mod}}^{\text{pf}} = \sum (MV_i / MV) D_{\text{mod},i} \cdot C^{\text{pf}} = \sum (MV_i / MV) C_i$$

Table A2. Portfolio aggregates and the arithmetic behind each one.

Quantity	Calculation	Value
Market value	sum of eight 10M lines × dirty price	81,406,654.81 MAD
Curve DV01	835.68 + 1,328.54 + 3,378.15 + 4,712.23 + 7,588.90 + 10,336.56 + 13,334.04 + 16,872.33	58,386.42 MAD/bp
Dollar convexity (Γ)	sum of the eight Γ_i	8,206,407,183 MAD
Modified duration	market-value weighted	6.9925
Convexity (dimensionless)	market-value weighted	102.276
Curve duration	DV01 ÷ ($MV \times 10^{-4}$)	7.1722

A.5 The duration hedge notional

One instrument, one unknown, one equation. The 8.68-year line carries 758.89 MAD/bp of curve DV01 per 1M MAD notional, so

$$w_{8.68} = -DV01_{\text{pf}} / DV01_{8.68} = -58,386.42 / 758.89 = -76.94 \text{M MAD}$$

The short is not convexity-neutral, because the bullet is not as convex as the ladder it hedges. The 8.68-year line carries 62.5401 MAD per unit of (decimal yield)² per MAD of notional — 62.54M per 1M notional — so the residual is

$$\Gamma^{\text{resid}} = 8,206,407,183 + (-76,936,653)(62.5401) = 3,395 \text{M MAD}$$

which is 41.4% of the book's convexity. The same arithmetic in key-rate space: the 8.68-year line has 687.18 MAD/bp of key-rate DV01 at its own knot per 1M notional, and the unhedged book has 9,941.32 MAD/bp there, so the hedged exposure is

$$9,941.32 - 76.9367(687.18) = -42,927.82 \text{ MAD/bp}$$

the -42.9k bar in Figure 2.

A.6 The duration-plus-convexity system

Two instruments, two equations — total DV01 and dollar convexity, both zeroed. Per 1M MAD notional, with w in millions of MAD:

$$[471.2234 \ 1,333.4035 \ ; \ 23,043,578 \ 218,582,639] \cdot [w_5 \ w_{18.8}]^T = [-58,386.42 \ , \ -8,206,407,183]^T$$

The determinant is 72,274,863,998, so Cramer's rule gives

$$w_5 = [(-58,386.42)(218,582,639) - (-8,206,407,183)(1,333.4035)] / 72,274,863,998 = -25.18 \text{M MAD}$$

$$w_{18.8} = [(471.2234)(-8,206,407,183) - (23,043,578)(-58,386.42)] / 72,274,863,998 = -34.89 \text{M MAD}$$

Substituting back, residual DV01 is 1 MAD/bp and residual dollar convexity is zero to numerical precision.

A.7 The key-rate hedge

Key-rate DV01 is a decomposition rather than a formula: at each of the ten reference tenors the knot is bumped by 1 bp on a triangular basis (one at the knot, zero at its neighbours, interpolated between) and the book is repriced. The portfolio's ten key-rate exposures, in MAD/bp, are

$$G = [0.13, 44.46, 940.41, 1,821.45, 4,341.08, 6,380.42, 9,941.32, 11,921.11, 12,781.00, 10,177.38]$$

and each hedge instrument has its own ten-vector H_i . The hedge is a constrained least-squares problem: minimise the squared residual key-rate DV01 across all ten tenors, subject to the residual total DV01 being zero.

$$\min_w \| G + \sum_i w_i H_i \|^2 \text{ subject to } 1^T (G + \sum_i w_i H_i) = 0$$

With H the 6×10 matrix of instrument key-rate vectors (scaled per 1M MAD notional to keep the normal equations well conditioned, and w in millions of MAD), $A = HH^T$, $c = HG$ and $a = H1$ the instrument total DV01s, the solution is

$$w_0 = -A^{-1}c \cdot w = w_0 - \lambda A^{-1}a \cdot \lambda = (a^T w_0 + 1^T G) / (a^T A^{-1} a)$$

The first term is the unconstrained least-squares fit; the second shifts it along the total-DV01 direction until the residual sums to zero. Table A3 gives the solution.

Table A3. The six-line key-rate hedge. DV01 and Γ are per 1M MAD notional.

Instrument	Notional (MAD)	DV01 per 1M	Γ per 1M
3.50Y (2.89% Apr 2030)	-11,013,742	337.82	11.63M
5.01Y (3.09% Oct 2031)	-10,683,350	471.22	23.04M
8.68Y (3.48% Jun 2035)	-10,384,281	758.89	62.54M
12.76Y (3.95% Jul 2039)	-10,293,986	1,033.66	120.85M
18.84Y (4.09% Aug 2045)	-10,252,077	1,333.40	218.58M
28.53Y (4.38% Apr 2055)	-10,337,391	1,687.23	381.55M
Gross short	62,964,827		

The residual key-rate vector is [0.09, 6.15, 831.43, 1,297.03, -372.79, -365.25, -354.80, -351.20, -347.27, -343.38]: the total is zero, and what is left is the two short bills plus a roughly -350 MAD/bp tail at every long knot. The three-line hedge is the same system with three columns instead of six, and gives -14.81M, -13.76M and -14.66M in the 12.76Y, 18.84Y and 28.53Y lines, 43.23M gross.

A.8 Stress and backtest arithmetic

Each stress scenario is a shift applied to the fitted zero curve, and every cash flow is repriced:

$$P\&L = \sum_i w_i [P_i(z + \Delta z) - P_i(z)]$$

The three shift families are: parallel, $\Delta z(t) = s$ at every tenor; slope, $\Delta z(t) = s(2t/30 - 1)$, so -100 bp at 0Y and +100 bp at 30Y; and belly, a triangular bump that is zero at 0Y and 30Y and peaks at the 12.76-year tenor. A +100 bp belly therefore gives +630,292 MAD to the duration hedge, +369,203 to the duration-plus-convexity hedge and +85,989 to the six-line key-rate hedge.

The daily backtest is the same reprice between two fixing dates, plus carry, minus financing and costs:

$$P\&L_t = \sum_i w_i [P_i(z_{t+1}) - P_i(z_t)] + \sum_i w_i C_i(t, t+1) - \text{financing}_t - \text{cost}_t$$

$$\text{financing}_t = (\sum_{\text{short legs}} |w_i|) \cdot 50 \text{ bp} / 365 \cdot \text{cost}_t = (\sum_i |w_i^{\text{new}} - w_i^{\text{old}}|) \cdot 2 \text{ bp}$$

The hedge is re-solved on each day's curve and held to the next rebalance. Over the 106 daily intervals the daily P&L standard deviation is multiplied by $\sqrt{252}$ for the annualised figure, and the 95% VaR is the 5th percentile of the daily P&L series. Financing for the six-line hedge is $62.96\text{M} \times 50 \text{ bp} / 12 = 26,235 \text{ MAD}$ per month at a 50 bp net spread, or 144,294 MAD per month at a 2.75% all-in repo rate — 0.45 and 2.47 bp of book DV01.

References

- Bank Al-Maghrib (2026). *Taux de référence des bons du Trésor — marché secondaire*, value date 16 October 2026, and daily fixings 1 April – 16 October 2026. Rabat: Bank Al-Maghrib. <https://www.bkam.ma>
- Bank Al-Maghrib (2026). *Taux directeur* (policy rate), 2.25%, accessed October 2026. <https://www.bkam.ma>
- Ho, T. S. Y. (1992). Key rate durations: Measures of interest rate risks. *Journal of Fixed Income*, 2(2), 29–44.
- Nelson, C. R., and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of Business*, 60(4), 473–489.
- Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992–1994. NBER Working Paper No. 4871.

Disclaimer. Personal risk-management exercise on public reference data, not investment advice. BAM yields are indicative, not executable. The author reports no conflict of interest.