RATES STRUCTURED PRODUCTS · UNITED STATES

Constant-maturity Treasury as a structured product underlying

The US CMT complex: index, issuance landscape, issuer risk and hedging

Research note — 29 July 2026 · companion to the TEC10 note

Executive summary

The same trade is running in dollars, and it is smaller, later and structurally different from the French one. SEC-registered notes referencing a Constant Maturity Treasury rate went from **zero filings in 2023–24 to 186 in 2025 and 205 in the first seven months of 2026**. Goldman Sachs reintroduced the format on 3 February 2025 and eight other dealers followed. I can document **37 individual notes with a CUSIP and a source filing, of which 28 print a size, totalling $626.2m**, out of an estimated ~224 distinct notes across 438 filings. That is a floor. It is also, in money, **several times smaller than the €3.33bn French TEC10 wave** — the American version of this trade is the junior one.

The product is not the French product. Not one of the 37 is an autocallable, a digital or principal-at-risk. Every single one repays par. The dominant format — 22 of 37 — is a **callable range accrual**: a headline coupon of around 9% that accrues for each day the 10-Year CMT fixes at or below a barrier, paid quarterly, with the issuer able to call quarterly from year one. Where the French investor sells a strangle on one annual fixing, the American investor sells **roughly 2,520 daily digitals and a Bermudan call**.

A 5.00% barrier appears on twenty of the thirty-seven, and the forward curve crosses it at year 3.9. On my model the expected accrual is **72.5% of days, worth 6.52% a year against a 9.00% headline** — before the issuer's call, which is worth another 5.4 points to the issuer. The realised record so far splits sharply by vintage: the 5.00% cohort has accrued every day of every quarter since issuance, with 29bp of headroom left; the tighter 2026 vintage struck at 4.30% accrued **nothing at all in Q3 2026**.

The dealer's problem is the daily strip and the basis, not the vol surface. Booked as receiver spreads 5bp wide, the daily digitals come to **$1.1bn of gross replicating notional against a $5m note — 225 times the trade**, nearly six times the equivalent French figure, because the French note observes twelve times and this one observes 2,520. Meanwhile the range accrual is close to **vol-neutral**: it is long volatility where the forward sits below the barrier and short where it sits above, and on a 10-year note struck at the crossing point the two halves very nearly cancel. What does not cancel is the **Treasury-versus-swap basis**, which on my model moves the liability by about a third of a point per 20bp — comparable to a large move in the whole vol surface, and unhedged by anything in the swap book.

The sign of the residual does not flip against France, but the identity of the risk does. In both currencies the note pays when the government-bond yield is low, so the dealer receives fixed and ends up short the govvie-versus-swap spread. The French dealer is short

French sovereign credit with a redenomination tail. The American dealer is short a **US bank capital constraint** — the negative swap spread is an intermediation spread, and SLR reform, the Treasury clearing mandate and buyback policy are all live inside every note's life.

The dealer-to-dealer CMT market exists, it was born in September 2025, and it is the wrong shape. A direct search of DTCC's public Part 43 files across 214 business days and 4.77 million interest-rate records returns **95 CMT records, 43 new trades and $8.85bn of notional — and nothing at all before the week of 4 September 2025**. All of it is uncleared and bilateral; **exactly one is an option** in ten months. What trades is sub-one-year rate locks in $100–750m clips with the signature of pension risk transfer, not multi-year digital volatility. The note wave leads its own hedging market by seven months.

And the index itself was re-specified mid-contract. On 6 December 2021 the Treasury changed the CMT fitting method from a quasi-cubic Hermite spline to monotone convex. Contracts written before that date reference an index computed a different way afterwards, with no consent mechanism and no fallback trigger. The measured impact is small — a mean of −0.1 to +0.5bp — but the governance point is not: this is the mirror image of the French problem, where the index has no benchmark administrator at all.

Part 1 — What CMT is, and how it differs from TEC10

1.1 The index

The Treasury publishes the **Daily Treasury Par Yield Curve** and the Federal Reserve republishes it as the H.15 Treasury constant-maturity series. The construction is the essential contrast with the French index:

	TEC10 (France)	**CMT (United States)**
Method	linear interpolation between **two real OATs** bracketing the maturity	read-off from a **fitted curve** through the whole on-the-run set
Inputs	MTS France executable bid/ask mids, outstanding > €3bn	on-the-run **bid-side** yields, from the Federal Reserve Bank of New York
Fit	none — two bonds and a weight	quasi-cubic Hermite spline to 5 Dec 2021, **monotone convex from 6 Dec 2021**
Snap	11:00 CET	**3:30 p.m. ET**
Administrator	Banque de France, exempt from BMR under Art. 2(2)(a)	US Treasury; no benchmark-regulation regime applies
Reference pool	rotates monthly, two bonds at a time	the on-the-run set, rotating at each auction
Fallback	re-fix at 11:30 → poll five SVTs → previous day	FRBCMT → H.15(519) → comparable rate → **five-dealer bid poll at 3:30 p.m.** → previous period

Three consequences follow for a structured note.

The two-bond artefact disappears, and a fitting artefact replaces it. §4.2 of the French note describes a monthly pool review that makes TEC10 jump by the local curvature of the

OAT curve when the reference pair changes. CMT has no such discontinuity — but it has a spline. The value of the 10-year point depends on the whole on-the-run set, including the 7-year and the 20-year, so a dislocation anywhere on the curve propagates into the fixing. That is a subtler exposure and a harder one to hedge, because it is a curve-shape exposure with no natural instrument.

Bid-side, not mid. TEC10 is computed off the arithmetic mean of bid and ask. CMT is computed off bid-side yields, i.e. systematically a fraction of a basis point above a mid-market read. Small, deterministic, bookable — and it is the sort of convention that is easy to get wrong in a model calibration.

The method changed mid-contract, and nobody was asked. The 6 December 2021 switch from a quasi-cubic Hermite spline to monotone convex was a methodology improvement, publicly announced, with a documented mean effect of −0.1 to +0.5bp depending on tenor and a standard deviation of 0.2–1.0bp. It is not a scandal. But it happened inside the life of contracts referencing the index, there was no consent mechanism, and it did not trigger any fallback provision, because none of the fallbacks in a US pricing supplement contemplate the index continuing to be published under a different definition. The French problem in §4.8 of the companion note is *no administrator*. The American problem is *an administrator who is not a benchmark administrator and owes the market nothing*.

1.2 The plumbing, and the derivatives market that did not exist until September 2025

The ISDA machinery is real. The live floating rate options are `USD-CMT` and `USD-CMT Average 1W` in the 2021 Definitions, with `USD-CMT-T7051` and `-T7052` surviving from 2006. Of roughly 660 FpML floating-rate codes, **exactly two live families reference a sovereign bond yield at constant maturity: `USD-CMT` and `EUR-CNO TEC10`.** The documentation bar the French note says is low in euro is, in dollars, lower still — the language has been sitting in the definitions for twenty years.

The question is whether anything trades against it. It is answerable, because CFTC Part 43 requires public dissemination of swap transaction data and DTCC publishes the files. Searching the `UPI Underlier Name` field across **214 business days and 4.77 million interest-rate records** returns:

CMT records found	**95**
Distinct new trades	**43**
Aggregate notional	**$8.85bn**
Earliest	week of **4 September 2025** — and zero on every sampled day back through June 2024
Cleared	**none**; all uncleared, bilateral, off-facility
Options	**one**, in ten months, against ~510 swap-rate option records in three
Typical shape	sub-one-year rate locks, $100–750m clips, clustered expiry dates

Three things follow. First, **a dealer-to-dealer CMT swap market now exists** — the French note's claim that no such market exists for TEC10 has no American equivalent. Second, **it is the wrong shape for this product**: short-dated rate locks in institutional size, with the signature of pension risk transfer, not the multi-year digital volatility a range accrual needs. Third, **the note wave leads the hedging market by seven months.** Goldman printed the first note on 3 February 2025; the first CMT swap prints in September. Dealers were manufacturing the exposure for two and a half quarters before there was any interdealer venue for it at all.

There is also a listed hedge the euro market lacks, and it is closer than it looks and emptier than it should be. CME's **Micro Treasury Yield futures** settle in cash on the BrokerTec on-the-run 10-year *yield*, with a fixed $10 DV01 — structurally almost exactly the right instrument for a CMT payoff. Open interest as at 21 July 2026: **2,332 lots, against 6,439,641 in the 10-Year Note future.** The instrument exists and nobody uses it.

Part 2 — The issuance landscape

2.1 The prior life, and the correction it forces

CMT-linked registered paper is not new. Full-text search returns **123 filings across 2004–2012**, then nothing at all from 2013 until 2025. Two corrections to the standard account are worth making, because both are repeated in market commentary:

The "2004–2008 CMS steepener wave" is really 2007–2011 and 2013–14 in registered paper. The filings do not support the earlier dating.

The 2006 CMT vintage was curve-shape product, not range accruals. Morgan Stanley printed a 15× leveraged note on 10-year CMT minus 3-month Libor; JPMorgan printed a best-of on 2-, 5- and 10-year CMT. These were sold into a curve that inverted in 2006 and stayed inverted, and the payoffs were destroyed by the inversion rather than by any barrier. The 2010–2014 wave that followed *was* callable range accruals, and FINRA named both "steepener notes… high teaser coupon… then variable rates" and "range accrual notes" as complex products requiring heightened supervision in **January 2012**.

That matters for Part 5: **the structure now being sold was explicitly identified by the regulator as a complex product fourteen years ago**, and 24% of the current register runs to 2045 or beyond.

2.2 The current wave

	2023–24	2025	2026 to 29 July
424B2 filings mentioning "10-Year CMT"	**0**	**186**	**205**
All CMT phrases, union	0	~200	~238
Estimated distinct notes	0	~110	~114
Documented in the register	—	15	22

The zero in 2023–24 is a true zero, not a search artefact — though one search artefact is worth recording, because it nearly cost a whole issuer: **Morgan Stanley writes "10CMT", never "10-Year CMT"**, and 39 filings covering 24 notes were invisible until the phrase list was widened.

Filings by issuer, 2025 / 2026 YTD: JPMorgan 91/67 · Goldman 48/68 · Citigroup 36/15 · Morgan Stanley 31/8 · Bank of America 15/23 · Barclays 2/20 · RBC 0/7 · Jefferies 0/6 · BMO 0/1. Goldman overtook JPMorgan in 2026; Morgan Stanley has nearly stopped.

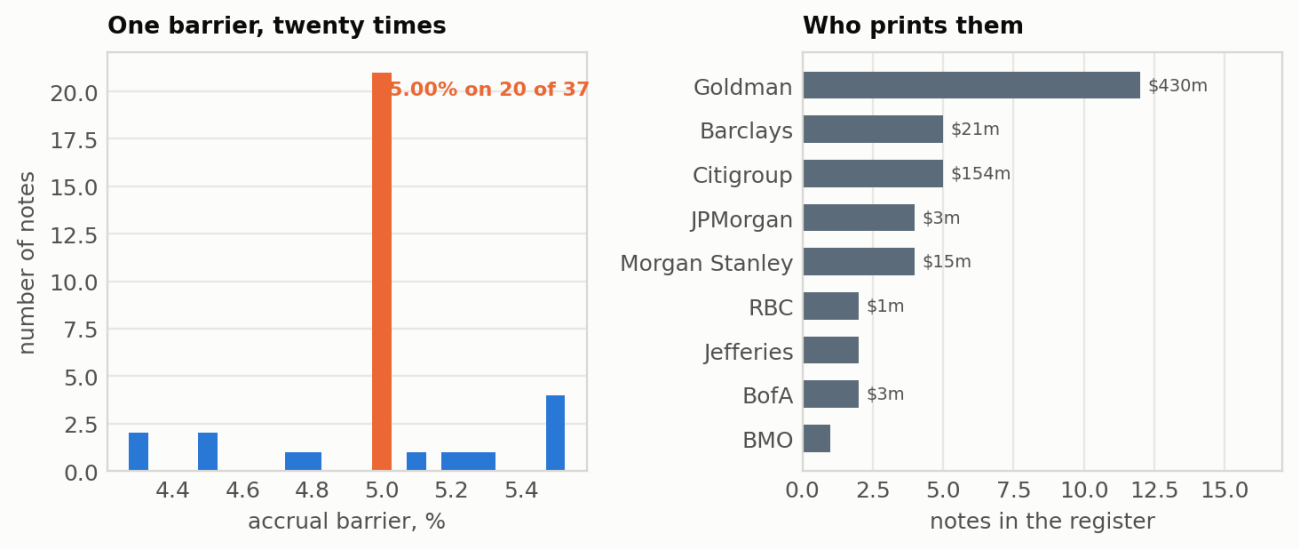

The register: barriers and issuers

The register

Thirty-seven notes with a CUSIP and a source filing. Structural mix:

Structure	Count	What the holder is short
Callable range accrual	**22**	~2,520 daily digitals + a Bermudan call
Fixed-to-floating range accrual	9	the same, after a fixed teaser period
Floored range accrual	3	the same, with a minimum coupon
Leveraged inverse floater	2	a levered short position in the rate, collared
Dual-tenor range accrual	1	a joint condition on 10Y and 30Y CMT

Not one autocallable. Not one digital-at-maturity. Not one principal-at-risk note. Every one of the 37 repays par. That is the single largest structural difference from the French market, where the autocall is the financing mechanism for the whole coupon, and it changes the risk profile completely: the American investor cannot lose principal, only opportunity, and the American dealer does not have a stochastic hedge horizon driven by an *investor-favourable* trigger — it has one driven by an *issuer-favourable* one.

Three deals that define the shape of it

Goldman 40058WV86 — $370,000,000, 5-Year CMT range accrual, cap 5.95%, priced 9 January 2026, **estimated value $993.50 per $1,000**. A 0.65% all-in cost, against 3–14% for genuinely retail structured paper. This is not a distributed offering; it is a single negotiated buyer. Together with Citigroup's **$150m** 5-Year note — whose underwriting fee is printed as "—", i.e. zero — two deals are **83% of the documented notional**.

Citigroup 17291W5S6 — $464,000, twenty years, paying `54.00% − 10 × (10Y CMT)`, capped at 20% and floored at zero. Leverage of ten. The coupon caps out at a 3.40% CMT and dies entirely at 5.40%. Quarterly discrete fixing, not a daily count — the only genuinely different payoff in the register, and it is a rounding error in size.

Barclays 06749FFE4 — $3,400,000, where the placement agents are **J.P. Morgan Securities LLC and JPMorgan Chase Bank, N.A.**, and the filing itself is named `..._424b2-8303jpm.htm`. Barclays manufactures; JPMorgan's private bank distributes. That is the distribution model for the long tail.

So there are two markets here, and they are not the same trade. A handful of very large, very cheap negotiated blocks for institutional buyers; and roughly two hundred $1–5m tranches placed into private-bank and adviser channels at estimated values as low as $860 per $1,000. The average of the two tells you nothing about either.

2.3 The barrier, and the vintage that is already paying nothing

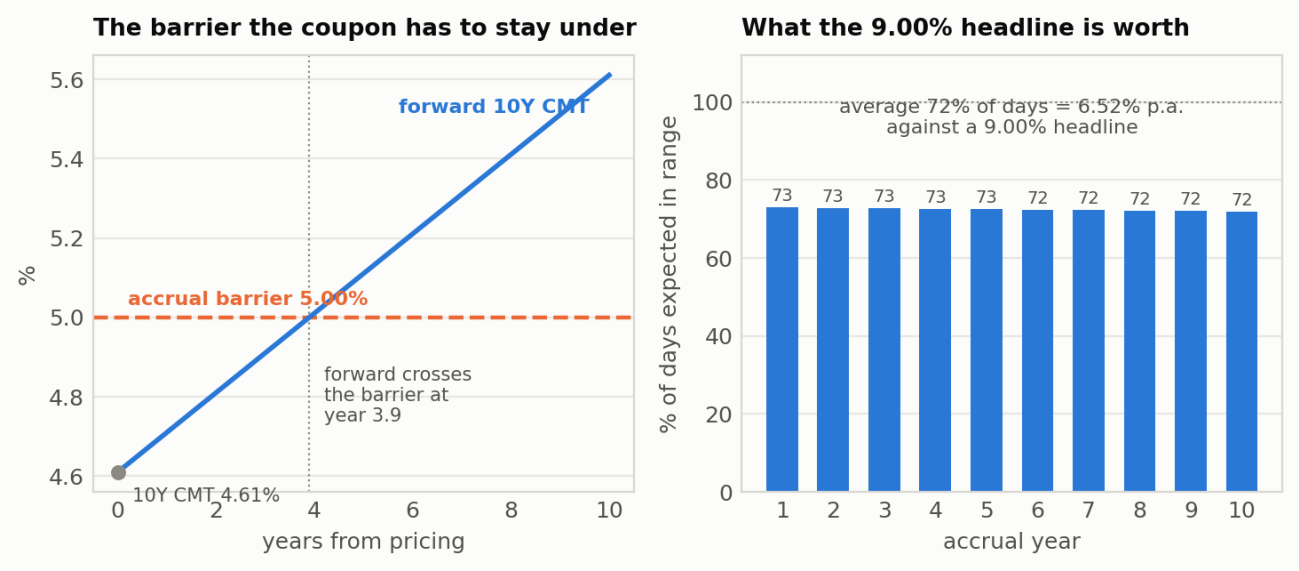

The barrier and what the headline coupon is worth

A **5.00% barrier appears on twenty of the thirty-seven notes**, with the rest scattered from 4.30% to 5.50%. That clustering is the American analogue of the round-number clustering the French note flags in §4.8, and it has the same consequence: a large aggregate digital notional struck at one number.

What it has produced so far divides sharply by vintage, and this is measurable from the Treasury's own published series rather than from a model:

Barrier cohort	Realised accrual	Annualised coupon on a ~9% headline
5.00% (the 2025 bulk)	**100% of days, every quarter since issuance**	full headline — with **29bp of headroom left**
4.50%	~15% of days in Q3 2026	**~1.4%**
4.30% (the 2026 tightening)	**0% of days in Q3 2026**	**nil**

The 2025 buyer has been paid in full and is 29 basis points from stopping. The 2026 buyer, sold a tighter barrier to support the same headline coupon in a higher-rate market, is already receiving nothing. Same product, same distributor, twelve months apart.

Part 3 — What the investor is actually selling

Take the modal note: **$5m, ten years, principal protected, 9.00% accruing for each day the 10-Year CMT fixes at or below 5.00%, paid quarterly, issuer-callable quarterly from year one.** Priced as at 28 July 2026 with the 10-Year CMT at 4.61% and the 10y10y forward at 5.61%.

3.1 The forward crosses the barrier at year 3.9

This is the same mechanism as §3.1 of the French note, in a different guise. The barrier is not struck below spot — 5.00% is 39bp *above* the 4.61% spot, and for the first four years the note looks comfortable. But the forward curve rises, and it **crosses 5.00% at year 3.9**. Everything after that is a bet against the forward.

On my model:

Expected accrual, year 1	73.0% of days
Expected accrual, year 10	71.9% of days
Average over ten years	**72.5% of days → 6.52% p.a.**
Liability, ignoring the call	96.27 per 100
Liability with the issuer's call	**90.85 per 100**
Value of the call, to the issuer	**5.41 points**

The flatness of the accrual profile is worth a moment. The French note's coupon probability collapses from 65% to 23% across its twelve observations, because it depends on a single annual fixing and the forward runs away from the barrier. Here the accrual barely moves, because a *daily count over a year* averages out the path: even in year 10, with the forward at 5.56% and the barrier at 5.00%, the rate spends 72% of its days below the barrier, because the distribution is wide and the count is linear in the probability rather than in the outcome. **The range accrual is a much gentler short position than the autocall.** It is also a much less transparent one — nothing in the payoff description tells the buyer that a 9.00% headline is a 6.52% expectation.

3.2 What the estimated-value disclosure reveals

US pricing supplements must print the issuer's own estimate of the note's value. Europe has no equivalent, and the numbers are informative:

Note	Tenor	Estimated value per $1,000	Implied all-in cost
Goldman 40058WV86 ($370m)	5y	**993.50**	0.65%
Goldman 40058QK65	6.5y	974.80	2.5%
Barclays 06749FX95	10y	957.00	4.3%
BofA 09711CBP8	10y	944.00	5.6%
Morgan Stanley 61766YD40	20y	903.00	9.7%
Morgan Stanley 61766YF97	20y	**860.00**	**14.0%**

Two readings. The benign one is that cost scales with tenor and with the size of the distribution chain, which is what you would expect. The less benign one is that **the same underlying, the same structure and the same week can carry a 0.65% cost or a 14.0% cost depending entirely on who is buying** — and that the difference, 13.35 points, is roughly twice the entire value of the issuer's call.

Part 4 — The issuer's risks

4.1 The daily digital strip

The coupon is not one option. It is a strip of daily digitals: on each of roughly 252 business days a year for ten years, the note pays $\frac{9.00\%}{252} \times 5$ = \$1{,}786\$ if the 10-Year CMT fixes at or below 5.00%, and nothing if it fixes at 5.01%.

Booked the way a desk books a digital — as a receiver spread of width w on the conservative side — the replicating notional per day is $\text{payout}/(w \cdot A)$, with A the 10-year annuity:

Strike spread	Notional per day	Over 2,520 days	Multiple of the note
2bp	\$1.1m	**\$2.8bn**	562×
5bp	**\$0.44m**	**\$1.1bn**	**225×**
10bp	\$0.22m	\$0.6bn	112×
25bp	\$0.09m	\$0.2bn	45×

Two hundred and twenty-five times the notional of the trade, at a 5bp width — against thirty-eight times for the French autocall, computed the same way. The ratio between the two is almost exactly the ratio of observation counts, 2,520 to 24, damped by the fact that each daily digital is 1/252nd the size.

The consolation is that the daily strip is far better behaved than a twice-yearly barrier. A single digital's gamma diverges as w^{-2} at its strike on its observation date; a strip of 2,520 of them, each 1/252nd the size, spread over every business day, has its gamma smeared across the calendar. **There is no cliff date.** The French desk faces 24 discrete barrier events over twelve years; the American desk faces a continuous, modest bleed. In risk-management terms that is a straightforwardly better book, and it is the strongest argument in the product's favour.

4.2 Why the range accrual is nearly vol-neutral

This is the result that most surprises people who have priced autocalls:

Normal vol	60bp	80bp	95bp	110bp	130bp
Liability per 100	91.76	91.17	**90.85**	90.60	90.33

The entire vol surface can move by 70 basis points of normal volatility and the liability moves by 1.4 points. The reason is structural. Each daily digital is long or short volatility depending on whether the forward on that day sits below or above the barrier: a digital that is in the money is hurt by volatility, one that is out of the money is helped. **The forward crosses the barrier at year 3.9, so roughly the first four years of the strip are long vol and the remaining six are short**, and on a note struck near the crossing point the halves very nearly cancel.

Bucketed by accrual year, on a \$5m note:

Accrual year	Delta, pts/bp	Vega, pts/bp of vol	Receive fixed, \$m	Straddle, \$m
1	−0.0418	−0.0157	2.7	−1.3

Accrual year	Delta, pts/bp	Vega, pts/bp of vol	Receive fixed, $m	Straddle, $m
2	−0.0140	−0.0045	1.0	−0.3
3	−0.0109	−0.0027	0.8	−0.1
4	−0.0088	−0.0015	0.7	−0.1
5	−0.0072	−0.0005	0.6	−0.0
6	−0.0066	−0.0001	0.5	−0.0
7	−0.0051	+0.0007	0.4	+0.0
8	−0.0041	+0.0013	0.4	+0.0
9	−0.0038	+0.0018	0.3	+0.1
10	−0.0023	+0.0021	0.2	+0.1
Total	**−0.105**	**−0.019**	**7.5**	**−1.6 net / 2.0 gross**

The vega column changes sign at year 7 — later than the year-3.9 forward crossing, because the near-dated buckets carry the issuer's call as well as the accrual. Net vega is **−$950 per basis point of normal vol on a $5m note**: essentially nothing. Gross is $2.0m of straddle notional against a $1.6m net, so even here the net understates the execution, but the whole vol position is an order of magnitude smaller than the French one.

The implication is uncomfortable for the standard justification. Dealers hedge CMT in swaps and swaptions because that is where the liquidity is, and accept the CMT-versus-swap basis as the price. But if the vol exposure is nearly nil and the basis exposure is not, the liquid instrument is hedging the risk the book barely has.

4.3 The basis: Treasury versus swap

Shock the CMT curve and leave the swap curve alone — a pure swap-spread move, which the swaption and swap hedge does not see at all:

Treasury cheapens by	Change in liability	On the $5m note	On the $626m register
+10bp	−0.17 pts	$8k	**$1.0m**
+20bp	−0.33 pts	$17k	**$2.1m**
+30bp	−0.50 pts	$25k	**$3.1m**
+50bp	−0.85 pts	$42k	**$5.3m**
−20bp	+0.33 pts	$17k	$2.1m

Compare that with the vol table above. A 30bp move in the swap spread — entirely ordinary over a year — costs half a point. A 35bp move in normal volatility — a large move — costs half a point. **The two are comparable, and the basis is the one nobody is hedging.**

The sign does not flip against France, but the identity of the risk does. In both currencies the note pays when the government-bond yield is low; the dealer therefore receives fixed and is structurally short the govvie-versus-swap spread, and to flatten it must buy the government bond asset swap. Measured today the French basis is the larger of the two: OAT–swap around **+76bp** against Treasury–swap **+41.8bp**. What differs is what the spread *is*. The French spread is sovereign credit, with a redenomination tail that would move TEC10 with no move at all in euro CMS. **The US spread is an intermediation spread** — Treasury yields have sat above swap rates since about 2015 because dealer balance sheet is

scarce, and the standard academic account of it (Klingler and Sundaresan) is pension duration demand against constrained intermediation, not credit. The American dealer hedging a CMT note in swaps is therefore not short the United States. It is short a US bank capital rule — and SLR reform, the Treasury clearing mandate phasing in over 31 December 2026 and 30 June 2027, and buyback policy are all live inside the life of every note in the register.

4.4 The issuer's call

Worth 5.41 points, and it is the mirror image of the French autocall in the one way that matters: **it is exercised in the state that is good for the investor.**

Year	Issuer calls if 10Y CMT is at or below
1	3.43%
3	3.63%
5	3.88%
7	4.17%
9	4.62%

The boundary rises through time because there is less coupon left to save. Two consequences.

For the investor: the note pays its full 9% when rates are low — and that is exactly when the issuer takes it away. The investor is short a call struck, in year one, about 118bp below spot. That is a real short position and it is not disclosed as one.

For the dealer: the hedge horizon is stochastic, as in France, but it shortens in a *rally* rather than in the state the investor wants. And unlike the French autocall, the trigger is discretionary — the issuer calls when it is economic to do so, which in practice means when its own funding has cheapened as well as when rates have fallen. Exercise therefore depends on the issuer's own credit spread, which is not a rates variable and which no swaption hedges.

4.5 Conduct: the structure nobody is looking at

FINRA identified range accrual notes as complex products in January 2012. What is supervised in 2026 is something else. **FINRA's May 2026 structured-note sweep is scoped to non-principal-protected worst-of notes** — and every CMT note in the register is principal-protected and single-underlying, so none of them is in scope. The 2011 SEC OCIE report that still serves as the template for structured-note examinations never mentions rate products at all. There has been **no 2024–26 enforcement on rate-linked notes**, and FINRA's 2026 Oversight Report has no structured-note section.

That is not evidence of a problem. It is evidence that the four questions this product actually raises are not being asked by anyone: that a 9.00% headline is a 6.52% expectation; that the 2026 vintage is accruing nothing; that the issuer's call is worth 5.41 points and is exercised precisely when the investor is winning; and that the same structure carries a 0.65% cost for one buyer and 14.0% for another.

Part 5 — The hedge, step by step

Same treatment as Part 6 of the French note, on the modal $5m ten-year note. All numbers from `cmt_hedge.py`, a one-factor normal model for the 10-Year CMT solved by backward induction on a rate grid — a grid rather than Monte Carlo because the issuer's call is Bermudan and a grid gives the exercise boundary exactly. Within each quarter the accrual is integrated analytically: conditional on the rate at the start of the quarter, each day's rate is normal, so the expected day count is a sum of normal CDFs and needs no simulation.

5.1 Day one

The desk takes in $5m of ten-year funding and books a liability of **90.85 per 100**, or $4.54m.

Delta. Liability DV01 **−$5,236 per basis point**: the liability falls as rates rise, so the book is short duration and the desk must receive fixed. Distributed across the accrual ladder that is **$7.5m of forward-starting receive-fixed against a $5m note**, concentrated at the front — $2.7m of it at the 1y × 10y node, because the first year is the one most likely to accrue in full and least likely to be called away.

Vega. Net **−$950 per basis point of normal vol**. Effectively flat, as §4.2 explains. Gross straddle notional $2.0m: sell $1.7m across years 1–6, buy $0.3m across years 7–10.

The digital strip. $1.1bn of gross receiver-spread notional at a 5bp width, spread across 2,520 daily observations. Nothing about that is executed; it is booked, reserved against, and managed as a residual.

The basis. To close the Treasury-versus-swap exposure the desk would buy roughly **$5.5m of 10-year Treasury asset swap for every $5m of note**. It does not.

5.2 One year on

	A. Sell-off 10Y CMT 5.60%	B. Unchanged 4.61%	C. Rally 10Y CMT 3.60%
Liability, from 90.85	**80.80**	91.86	**98.72**
Next year's expected accrual	**16.0%**	73.0%	**96.6%**
DV01, from −$5,236/bp	−$4,369	−$5,047	**−$1,966**
Vega, from −$950/bp	**+$3,873**	−$1,050	−$2,335
Receive-fixed ladder, from $7.5m	$6.6m → **pay $1.0m**	$7.7m → receive $0.1m	$3.3m → **pay $4.2m**

Case A, the sell-off. The liability drops ten points. The accrual collapses from 73% of days to 16% — the note has essentially stopped paying. And the vega swings from −$950 to +$3,873 per basis point, a sign change: with the rate now well above the barrier, every digital in the strip is out of the money and the whole strip is short volatility. The desk must **buy** volatility, having carried almost none. It is a small ticket in absolute terms, but it is the one direction in which this product does behave like the French one.

Case B, unchanged. Almost nothing happens. The liability accretes toward par, the ladder moves by $0.1m. This is the case that never occurs in the French walkthrough, where a 36bp drift generated the largest rebalancing of the three, and it is the clearest illustration of the

difference between a daily count and a barrier: **the range accrual has no cliff to sit next to.**

Case C, the rally. The liability goes to 98.72 and the accrual to 96.6% of days — the note is paying its full 9%. Which is why the issuer's call is now live: at 3.60% the year-one exercise boundary of 3.43% is 17bp away, and by year three the note is being called. The desk pays fixed on $4.2m of the ladder and the position starts winding down. **The unwind is a rally trade in both currencies — but in France it is triggered by the investor winning a coupon and being redeemed at par, and here it is triggered by the issuer deciding to stop paying.**

5.3 What the walkthrough shows

Gross-to-net is even more extreme than in France, and it matters less. $1.1bn of digital replication against a $5m note is 225×, six times the French ratio. But it is spread over 2,520 dates rather than concentrated on 24, so the risk it creates is a bleed rather than a series of cliffs.

The vol exposure is nearly nil and the basis exposure is not. That inverts the usual justification for hedging in swaptions. On this payoff, a desk that runs a tight vega book and a loose basis book is optimising the wrong variable.

The hedge horizon is stochastic in the opposite direction from France. The French note dies when the investor wins; the American note dies when the issuer decides the investor is winning too much. Both give the desk an unknown horizon; only the American one makes it discretionary and credit-dependent.

The market is too small to move anything. Aggregated across the $626m documented register, the whole book is roughly **$940m of forward-starting receive-fixed — about $820m of ten-year equivalent — and $656,000 per basis point of DV01**. The French wave, at €3.33bn, is an order of magnitude larger and is worth watching for its effect on OAT asset swaps. The American wave is not, yet. If it grows to the French scale the basis flow becomes interesting; at $626m it is noise in a Treasury market that trades that in seconds.

Caveats. The model is a one-factor normal process on a single rate with a linear forward curve, flat discounting, no smile, no swap-spread dynamics and no issuer-credit dynamics, so it cannot produce a price. Every market-wide figure is one archetype scaled linearly by notional, and the register contains a $370m block and a $464,000 tranche that have almost nothing in common. Signs and orders of magnitude carry; the decimals do not.

Part 6 — What to watch

1. **Does the CMT swap market grow into the right shape?** It exists, it is seven months younger than the note wave, it is entirely uncleared, and it has produced one option in ten months. If dealer-to-dealer CMT options start printing, this product becomes genuinely hedgeable and issuance should accelerate. If they do not, the basis stays warehoused.
2. **The 2026 vintage.** Notes struck at 4.30% and 4.50% are accruing nothing today. If rates stay where they are, a cohort of principal-protected notes will return exactly par after five to ten years, which is a real loss to inflation and a reputational event with no legal remedy — the notes will have done precisely what the documents said.
3. **The 29 basis points of headroom.** The 5.00% cohort has accrued 100% of days since issuance. The 10-Year CMT is at 4.61%. A 40bp sell-off flips $400m-plus of notes from full coupon to nothing, all on the same barrier, all at once.
4. **Whether anyone supervises it.** FINRA named this structure in 2012 and its 2026 sweep excludes it by construction. A single scoping change would alter distribution economics quickly.
5. **The swap-spread policy tail.** SLR reform, the Treasury clearing mandate and buybacks all move the Treasury-versus-swap spread, which is the dealer's largest unhedged exposure on this payoff. That is an unusual position: the residual risk on a retail rates product is a bank capital rule.

Sources and reproducibility

The deal register, the counts and the estimated values come from SEC EDGAR full-text search over forms 424B2, 424B5 and FWP; every row in `us_cmt_deals.csv` carries the Archives URL it was read from. The index methodology comes from the Treasury's yield-curve methodology pages and the H.15 release. The derivatives-market measurement comes from DTCC's public CFTC Part 43 files, 214 business days, 4.77m interest-rate records, searched on the UPI underlier field. Open interest is from the CFTC Commitments of Traders report of 21 July 2026. Bank structured-note liabilities are from 10-K disclosures. The full source-by-source record, with VERIFIED / REPORTED / INFERRED tags, is in `us_cmt_context.md` and `us_cmt_research.md`.

```
python3 cmt_hedge.py      # every number in Parts 3, 4 and 5
python3 charts_usd.py     # the figures
```

Caveats

- **All model output is illustrative.** One-factor normal, backward induction on a 641-point grid, quarterly call dates, accrual integrated analytically within each quarter. No smile, no swap-spread dynamics, no issuer credit. Order-of-magnitude, comparative statics, signs.
- **The register is a floor, and an unrepresentative sample.** 37 documented notes of an estimated ~224; $626.2m of documented size, of which two deals are 83%. JPMorgan's and BofA's distinct-note counts are estimates at two filings per note. Roughly 139 further 424B2 filings mention "CMT" without matching any tenor-specific phrase, so the true filing count is somewhere between 438 and 577.

- **Size context, honestly stated.** JPMorgan's 10-K shows $81.9bn of interest-rate structured notes, 49.6% of its $165bn book; Goldman and Morgan Stanley report $112.7bn and $132.5bn of fair-value-option borrowings; the three-firm total is around $410bn and grew 22% in 2025. Against that, published flow data put interest-rate underlyings at **0.27% of US retail structured-note volume**. The CMT wave is roughly 1% of annual US structured-note flow and has nonetheless become most of its own category.
- **Two facts I could not close:** whether any of the large negotiated blocks were bought by a single identifiable institution, and the true aggregate notional of the ~187 notes in the filing population that are not in the register.