RATES STRUCTURED PRODUCTS · FRANCE

Constant-maturity sovereign yields as a structured product underlying

The TEC10 complex: index, issuance landscape, issuer risk and hedging

Research note — 25 July 2026, revised 29 July 2026

Executive summary

The French State invented the instrument and then abandoned it; banks have now revived the index without reviving the instrument. The TEC10 — the Banque de France's constant-maturity 10-year OAT yield — was created in March 1996 to support the State's own OAT TEC 10 floating-rate bond. Only two lines were ever issued (€21.7bn combined at peak, both redeemed by 2009). For the State the product was formulaic: the coupon *was* the index, and the index *was* the State's own marginal cost of 10-year money. There was nothing to replicate.

Since mid-2024 the index has come back as a structured-note underlying, and this time the issuer has a genuine replication problem. I can evidence **thirty-one TEC10-linked notes** issued into the French savings market between October 2024 and August 2026, of which **twenty-five publish a size, totalling €3.33bn**. Read that as an order of magnitude rather than a placement total: three quarters of it is the registered nominal on Crédit Agricole Final Terms rather than evidence of what was placed, and €245m is a marketing cap — but against that, at least two further Crédit Agricole tranches are known to exist and are not enumerable from public sources, and six deals publish nothing at all. The archetype is a 8–12 year, 100% capital-guaranteed, autocallable note paying a conditional coupon if TEC10 fixes below a barrier. Two things about the issuer list matter more than the headline. First, **the market is three-quarters Crédit Agricole**: nine deals and €2.51bn against €815m for everyone else combined — €260m JPMorgan, €220m SG, €180m BNP Paribas, €125m Morgan Stanley, €30m CIC. Second, there are really **two markets, not one** — a wholesale/CGP channel printing €30m tranches in €1,000 denominations, and a retail branch-network channel printing €250–500m tranches in €100 denominations. Almost every published account of this wave describes the first and misses the second. The same phenomenon is running in parallel in the US: SEC 424B2 filings referencing "10-Year CMT" went from **zero in 2023–24 to 186 in 2025**.

The economic content is a retail investor selling term premium. Coupon barriers are struck below the *forward* TEC10 curve, not below spot. With the OAT 10y10y forward around 5.56% against a 3.4% spot (Barclays' September 2025 figures), a 3.80% barrier is comfortably in the money on day one and progressively out of the money thereafter. On an illustrative model the probability of earning the coupon falls from ~65% at the first observation to ~23% at the twelfth. This is the mechanism the Barclays authors describe: term premium re-emerging in sovereign curves, and being repackaged and sold on.

The issuer's problem is that the hedge does not exist. A CMS payoff can be statically replicated with swaptions (Hagan). A CMT payoff cannot, because the underlying is a *bond* yield: the replication chain runs through cash OATs, term repo out to twelve years, and a monthly-rotating pair of reference bonds. In practice dealers hedge the rate in swaps and swaptions and **warehouse the TEC10-versus-CMS10 basis**. A 20bp OAT asset-swap widening with the euro swap curve unchanged is worth roughly **€14m of unhedged P&L on a €500m book, and €90m across the €3.24bn in the fixing calendar**. Fully hedging it would mean buying on the order of **€5.3bn of 10-year OAT asset swap**. Aggregated across the issuance wave, that makes the dealer community either a structural buyer of French paper versus swaps or a structural warehouse of French spread risk — a feedback channel from retail notes into OAT valuations that deserves more attention than it gets.

Walked through as a trade, what the desk executes bears almost no relation to what its net risk report shows. Part 6 takes one €30m tranche through inception and then through three markets a year later. Day one generates a **€53m ladder of forward-starting receive-fixed swaps, a €7m gross straddle strip whose net is €1.9m, and — at a 5bp strike width — €1.15bn of receiver-spread notional per observation date**. A year later, a 106bp sell-off leaves the net DV01 barely changed while requiring €46m of gross swap tickets and €30m of straddle buying; a 36bp drift into the corridor between the barriers requires *more* gross execution than the sell-off does; a rally through the trigger kills the note and forces the whole ladder back into the market at once. **The hedge's own maturity is a derivative** — 5.2 years at inception, 8.6 after the sell-off, zero after the rally — which is not true of any other common rates product.

Three risks have no liquid hedge at all: the TEC–CMS basis, the inter-date correlation that governs the autocall (worth ~9 points of valuation range in my model with every marginal held fixed), and the benchmark itself — TEC10 sits *outside* the EU Benchmarks Regulation via the central-bank exemption, so a twelve-year retail note references an index with no BMR administrator authorisation, no Article 27 benchmark statement and no orderly cessation procedure. The Barclays article's proposals — repo-OIS basis swaps and CMT-anchored cash-settled swaptions — are a direct response to the first two.

Part 1 — What TEC10 is, and why the State's version needed no hedge

1.1 The index

`CNO-TEC n` is the annual actuarial yield of a *notional* French Treasury security with residual maturity exactly *n* years, obtained by linear interpolation between the two real OATs that bracket that maturity:

$$\text{CNO-TEC}\, n = y_1 + \frac{D_n - D_1}{D_2 - D_1}\,(y_2 - y_1)$$

Mechanics that matter for a structured product:

Eligible pool	OATs quoted on **MTS France** with executable bid/ask, each with an **outstanding above €3bn**. TEC1 also admits BTFs. Pool reset by the calculation agent on the **first business day of each month**; on-the-run preferred over off-the-run at equal maturity
Yield convention	annual actuarial, computed off the **mid price** (arithmetic mean of bid and ask), T+2 settlement, French market convention standardised by the CNO in June 1992
Fixing	**11:00 CET** every TARGET day; backup fixing 11:30
Liquidity test	per reference bond, bid/ask *yield* spread **< 10bp** → automatically valid; **10–30bp** → valid only if not more than twice the prior day's spread; **> 30bp** → rejected, and the nearest-maturity bond with acceptable liquidity is substituted
Fallback	(1) re-fix off MTS at 11:30 → (2) collect quotes from **at least five SVTs, drawn from the top ten of the annual ranking** → (3) total disruption: **the previous day's published fixing applies**
Publication	Banque de France ~11:30, rounded to two decimals; Bloomberg `BFRTEC10 Index`, Reuters `BDFCN0TEC`. AFT republishes on its own site around 15:00 — that lag is republication, not the fixing
Tenors published	TEC1, 2, 3, 5, 7, 10, 15, 20, 25, 30 — all published daily (there is no TEC4, 6, 8 or 9)

Governance. The CNO created the index in 1996 and subcontracted calculation to the Banque de France in December 2009; the BdF has been full **administrator, calculation agent and disseminator since March 2020**. Trademark title is ambiguous across the public sources — the CNO's 2013 deck asserts that TEC® is a community trade mark owned by the CNO, the BdF's technical note states that the BdF holds **exclusive trademark rights** requiring licensees to sign trademark agreements and pay royalties, and the AFT's page says flatly that the Banque de France owns the index and retains all intellectual property rights. Two of the three now point to the BdF; whether the CNO retains residual title is not resolvable from public sources.

BMR status — the important one. A direct query of the ESMA benchmarks register (93 administrator entities, 28,039 benchmark records as at 25 July 2026) returns **zero hits** for the Banque de France, the CNO, or any TEC benchmark; no central bank of any kind is registered, €STR included. The exemption is BMR Article 2(2)(a) — "This Regulation shall not apply to: (a) a central bank" — and ESMA Q&A 4.1 confirms that central-bank benchmarks need not appear in the Article 36 register while remaining usable by EU supervised entities. The practical consequence for a twelve-year retail note is set out in §4.8: no authorised administrator, no benchmark statement, no Article 21/23 cessation procedure. Users still owe Article 28(2) written fallback plans.

The move to the BdF in March 2020 is very likely a BMR response — the transitional regime for existing EU index providers expired 31 December 2019, and the CNO is a *loi 1901* association with no realistic path to AMF authorisation as an administrator. Neither body has published a rationale, so treat this as circumstantial.

1.2 The sovereign instrument, 1996–2009

	OAT TEC 10 #1	OAT TEC 10 #2
Announced / issued	27 Mar 1996; syndicated **9 April 1996** (BNP and Caisse des Dépôts), settled 25 Apr 1996	**March 1997**
Maturity	25 October 2006	25 January 2009
Outstanding at 31/03/2002	**€11.89bn**	**€9.85bn**

Combined peak outstanding €21.74bn. Line 2 was tapped for FRF 12.91bn (€1.97bn) between May and July 1998. Coupons were quarterly, paid 25 Jan/Apr/Jul/Oct, "predetermined revisable" — fixed in advance off the TEC10 five business days before the start of the coupon period:

$$c_{\text{quarterly}} = (1 + \text{TEC10} + m)^{1/4} - 1$$

with `m` a fixed *marge faciale*. The `^(1/4)` exponent preserves the actuarial character of the index. There was **no leverage** on TEC10 and no documented floor. The numerical value of `m` for either line is not in any public source I could reach.

Why this was easy for the State. The coupon was the index; the index was a synthetic 10-year OAT yield; the issuer *was* the OAT market. The State's liability and the reference were the same object, so the position was a perfect self-hedge with zero basis and zero optionality. The design goal — stated in the Senate's 2002 budget report — was that the bond "is auctioned at a price close to par, thereby reducing profits and losses." Neither the AFT nor the Senate ever published a reason for stopping; the outstandings were static by March 2002 and the AFT notes only that floating-rate OATs linked to long-term rate references existed "up until 2009." The plausible reading is consolidation onto plain-vanilla fixed-rate OATs plus the inflation-linked complex.

Note the asymmetry that sets up everything below. The State paid the index. A bank selling a *digital* on the index must manufacture the index, the optionality on it, and the joint behaviour of the index across twelve annual observations — with none of the State's natural offset.

Part 2 — The issuance landscape

2.1 The first wave, 1996–2012

TEC10 had a real, if modest, life as a floating-rate reference before it went dormant.

Issuer	Deal	Size	Coupon / structure
CADES	OAI 4.71% due 11 May 2008 (settled May 1998)	FRF 2.0bn	Fixed 4.71% coupon; **TEC10 indexes the amortisation schedule** (TEC10 <7.5% → 0% amortised, 8.5%+ → 100%, interpolated). Lead CCF, priced 20bp through the OAT curve
CADES	OAI 4.60% due 15 Jul 2010 (settled Jul 1998)	FRF 1.5bn (~€229m)	Same structure, 24bp through. Lead CCF
CADES	EUR TEC10 FRN due 25 Jul 2020 (priced 29 Sep 2004)	**€1.0bn**	`Max[0; (1 + TEC10 − 100bp)^(1/4) − 1]`, quarterly, floored at zero; issue price 98.55%, marketed as "TEC10 − 88bp". **Leads Barclays Capital, BNP Paribas, CDC IXIS.** The largest non-sovereign TEC10 bond found, and the CNO's own worked valuation example
IXIS CIB (ex-CDC IXIS → Natixis)	FR0000583213, maturity Jan 2009	FRF 990m	CNO TEC10 − 0.75%
CNCE / BFBP → BPCE	TSDI CNO TEC10 (Jul 1998) FR0000209793; plus FR0010176081 (€30m, 2005–2020)	—	TSDI perpetual-style; fixed-to-TEC10/CMS linked notes
LASER-Cofinoga	two subordinated bonds, 2004, 20y to Jan 2024	€50m + €30m	TEC10 + 1.15%
France Télécom	XS0720083244, issued **21 Dec 2011**, due 21 Dec 2021	**€272m**	**TEC10 + 0.50%**, quarterly, Act/360, no floor, no cap, no call. Non-syndicated, dealer Deutsche Bank. Fixing: Reuters `TREASTEC10` at 10:00 Paris
Areva	FR0011227339, issued **4 Apr 2012**, due 21 Mar 2022	**€200m**	**CNO-TEC10 + 2.125%**, quarterly, Act/360, no floor/cap. Issue price 99.017%. Syndicated, JLMs Barclays and Credit Suisse; calculation agent Société Générale

Two structural observations. First, **the first wave was all linear** — TEC10 plus or minus a spread, occasionally floored, once (CADES 1998) indexing an amortisation schedule rather than a coupon. There were no digitals. Second, **the capability never went away**: TEC10 remains an explicit permitted reference rate in the CAFFIL/SFIL *obligations foncières* base prospectuses (verbatim in the 2020, 2021, 2023 and 2026 editions), in Compagnie de Financement Foncier's, and in La Banque Postale's and Banque Palatine's EMTN programmes — although no executed covered-bond tranche was found.

A caveat on completeness: BALO full text starts in 2005 and the AMF filings dataset effectively in 2008, so the 1996–2004 vintage is largely unsearchable. A full-text search of the AMF corpus for "CNO-TEC" returns exactly two documents — France Télécom and Areva — which is strong negative evidence for the post-2008 period and equally a demonstration of the pre-2008 blind spot.

Then, roughly 2012 to 2024, nothing.

2.2 The current wave, October 2024 – August 2026

Thirty-one French-distributed notes with a confirmed ISIN, of which twenty-five publish a size. They split cleanly into two channels that have almost nothing in common except the underlying, so I have tabled them separately. Rows marked † I verified against the issuer or distributor's own document; the rest are solidly sourced at one remove.

2.2a The wholesale and CGP channel — €30m tranches, €1,000 denominations

Product	ISIN	Issuer / guarantor	Issued	Tenor	Size	Coupon	Coupon barrier	Autocall
ONYX 6 †	FR001400SDV1	**CIC** (direct, no SPV)	28 Oct 2024	12y	≤ €30m	6.30%	3.15%	≤ 2.35%, annual from yr 1
Afer Rendement Nov 2024 †	FR001400T3H1	JPMorgan Chase Financial Co / JPMorgan Chase & Co	06 Feb 2025	10y	**€200m**	5.50% (yr 1 unconditional)	3.50%	≤ 2.50%, yrs 2–9
MIF Structure Rendement 4	XS2755641334	JPMorgan Chase Financial Co	Q1 2025	12y	n/d	6.25%	3.50%	≤ 2.20%, yr 2+
(unbranded EMTN) †	XS2755637142	JPMorgan Chase Financial Co	30 May 2025	12y	€30m	7.50% (memory, in fine)	3.40%	3.00%, 11 dates
Rendement Taux France Oct 2025 †	FR0014010HP6	BNPP Issuance BV / BNP Paribas	26 Jun 2025	12y	€30m	7.50%	3.80%	≤ 3.30%, annual
Phoenix Taux France Semestriel Sep 2025 †	FR0014011946	BNPP Issuance BV	18 Jul 2025	12y	€30m	3.00%/sem	4.10%	≤ 3.20%, S2+
Phoenix Target France Fév 2026 †	FR0014014486	BNPP Issuance BV	25 Nov 2025	12y	€30m	3.00%/sem	4.00%	12% cumulative, S4+
OCCITAUX Mars 2026 †	FR0014014RE1	BNPP Issuance BV	09 Dec 2025	12y	€30m	3.00%/sem	4.00%	≤ 3.20%, S2+
CCF Objectif Taux TEC10 †	FRF0000002L7	Morgan Stanley Finance / Morgan Stanley	08 Jan 2026	10y	≤ **€125m**	8.25% (memory)	4.10%	≤ 3.25%, yrs 1–9

Product	ISIN	Issuer / guarantor	Issued	Tenor	Size	Coupon	Coupon barrier	Autocall
IMPULSION 13 †	FR0014015W56	BNPP Issuance BV	24 Feb 2026	12y	€30m	1.75%/qtr	4.10%	10.50% cumulative, Q6+
Phoenix Target France Mars 2026 †	FR00140157E1	BNPP Issuance BV	Mar 2026	12y	€30m	3.10%/sem (S1–S2 fixed)	4.00%	12.40% cumulative
Phoenix Target Sem. France Juil 2026	FR0014018F54	BNPP Issuance BV	18 May 2026	12y	n/d ‡	3.50%/sem	4.00%	10.50% cumulative, S3+
Phoenix Target France Juin 2026 †	FR0014017TM3	SG Issuer / Société Générale	Jun 2026	12y	€30m	3.50%/sem (S1 fixed)	4.15%	10.50% cumulative
EXCELTIS PATRIMOINE Juin 2026 †	FR0014015Y47	JP Morgan Structured Products BV / JPMorgan Chase Bank NA	30 Jun 2026	12y	€30m	1.50%/qtr	4.10%	9% cumulative, Q6+
ELAN Mai / Juin 2038 † ‖	FR0014018CV8	SG Issuer / Société Générale	06 Jul 2026	12y	€40m	8.75%/ elapsed yr, max 105%	— (payoff if TEC10 **>** 5.10% at maturity: +17.5%)	**issuer-callable**, yrs 2–11
Europe Patrimoine Rendement TEC 10 †	FR0014018DW4	SG Issuer / Société Générale	31 Jul 2026	12y	€30m	3.25%/sem	4.25%	≤ 3.50%, S2–S23
CCF Stratégie Taux TEC10 †	FR0014018P94	SG Issuer / Société Générale	19 Aug 2026	10y	≤ **€90m**	8.00%	4.28%	16% cumulative, yrs 2–9
Phoenix Target Tec 10 Août 2026	FR0014018T82	BNP Paribas Arbitrage Issuance BV	Aug 2026	—	n/d	terms not public	—	—
Phoenix Mémoire Target Tec 10 Nov 2026	FR0014019550	BNP Paribas Arbitrage Issuance BV	pending	—	n/d	terms not public	—	—

‖ The ELAN line is internally inconsistent across its own documents on name, ISIN and final observation date; §2.5 sets out the detail. ‡ FR0014018F54 publishes no size. Every other BNPP tranche in this table prints exactly €30m, so €30m is the obvious inference. It is **excluded from the €3.33bn issuance total**, which counts only published sizes; it *is* carried at €30m in the §2.5 fixing calendar, so that the deal's twenty-two observation dates appear on the schedule, and is tagged `inferred` in `fixings.csv`. The €3.24bn calendar total therefore contains €30m that the €3.33bn issuance total does not.

Two further deals belong to this channel and were missing from earlier accounts: **UNEP Phoenix Taux France Décembre 2024** (FR3CIBFS4889, Crédit Agricole CIB Financial Solutions, 27 Dec 2024, 10y, €30m) † and **Equitaux France 2025** (FR3CIBFS5118, same

issuer, 07 Feb 2025, 10y, €30m, coupon barrier 3.20%, autocall 2.55%) †. Both are €1,000-denomination CGP paper. Also resolved since the first draft: **FR0014018AI9** (SG Issuer, 27 May 2026, €30m, quarterly, 44 observations, 9% cumulative-coupon autocall) †, previously a LuxSE line with unknown terms — and the only sub-par issue price in the sample, at 99.80%.

What is still open in this channel. Two LuxSE listings, **FR1459ABD502** and **FR1459ABE294** (both Goldman Sachs Finance Corp International, "EUR-TEC10-CNO"), have no size and no terms: luxse.com returned HTTP 500 on every attempt across two sessions. And CACIB's *Athena TEC 10 Mars 2025* (10y, up to 5.97% p.a., quarterly autocall at 2.70%) is confirmed to exist and to have those terms, but no ISIN is published anywhere I could reach, so it is not counted.

2.2b The branch-network channel — €250–500m tranches, €100 denominations

This is the part of the market that has gone unreported, and it is the larger part. Crédit Agricole runs a second, entirely separate TEC10 franchise through **Crédit Agricole CIB Finance (Luxembourg)** and **LCL Emissions**, distributed through the LCL and Crédit Agricole branch networks rather than through independent advisers. Every one of these is verified against Final Terms.

Product	ISIN	Issuer	Issued	Maturity	Size	Coupon barrier	Autocall trigger	First obs.
LCL Vision TEC 10 (Janv 2025) †	FR1CIBLU0564	CACIB Finance Lux	13 Dec 2024	Apr 2033	**€500m**	3.40%	≤ 2.40%	Apr 2027
LCL Vision Tec10 (Sept 2025) †	FR1CIBLU1133	CACIB Finance Lux	18 Jul 2025	Nov 2033	**€400m**	3.60%	≤ 2.90%	Oct 2028
LCL TEC 10 Dynamique (Sept 2025) †	FR1CIBLU1141	CACIB Finance Lux	25 Jul 2025	Nov 2033	**€350m**	3.20% (terminal digital, no coupon)	≤ 3.20%	Oct 2027
LCL Amplitude TEC 10 (Janv 2026) †	FR0014014494	LCL Emissions	18 Dec 2025	Apr 2034	**€250m**	3.60%	≤ 3.60%	Apr 2028
CA Autocall TEC 10 (Janvier 2026) †	FR1CIBLU1406	CACIB Finance Lux	16 Dec 2025	Apr 2034	**€350m**	initial + 0.30% (4.04%)	≤ 4.04%	Mar 2029
CA Autocall TEC 10 (Mars 2026) †	FR1CIBLU1463	CACIB Finance Lux	18 Feb 2026	Jul 2034	**€300m**	initial + 0.30% (4.11%)	≤ 4.11%	Jul 2029
CA Autocall TEC 10 (Juin 2026) †	FR1CIBLU1570	CACIB Finance Lux	08 Jun 2026	Oct 2034	**€300m**	initial + 0.30%	≤ initial + 0.30%	Oct 2029

€2.45bn from seven deals, plus the €60m of CGP paper above: **€2.51bn from the Crédit Agricole group**, against **€815m from every other issuer combined** (€260m JPMorgan, €220m SG, €180m BNPP across six sized tranches, €125m Morgan Stanley, €30m CIC).

One qualification on all of these figures, and it cuts against the "floor" reading as well as for it. The €2.45bn is the *Montant Nominal Total* on the Final Terms, which is the amount the issuer registered, not proof that the full amount was placed with savers. A further €245m — ONYX 6 and the two CCF deals — are marketing caps rather than issued amounts (excluding them the total is €2.99bn). Against that, at least two 2025 Crédit Agricole tranches are known to exist and could not be enumerated, and six deals with an ISIN publish nothing at all. The figure is the best-evidenced number available and it is an order of magnitude larger than what has been reported; it is not an audited placement total. The second deal of the entire wave — CACIB's €500m on 13 December 2024 — was on its own larger than every non-Crédit-Agricole deal of 2024 and 2025 put together, which come to €380m.

Three differences from the €30m channel are analytically material, not cosmetic.

They start observing much later. The CACIB and LCL structures have **no observation at all for the first two to three years**. That is what buys the higher headline participation, and it means three quarters of the notional is untouchable by an early autocall until 2027–2029 (§2.5).

The 2026 vintage sets the barrier *relative* to the initial fixing — initial + 0.30% — rather than at an absolute round number. This is a better product from a fairness standpoint, since the trigger is not a bet on where the curve happened to sit in the marketing week, and it also removes those deals from the round-number clustering that §4.8 flags as a fixing-integrity surface.

Most of them pay no periodic coupon at all. Four of the seven pay nothing along the way and instead accrue a *gain per elapsed year* settled when the note redeems — a structure with no analogue in the €30m channel, and one that removes any interim cash flow the saver could use to judge how the product is doing. **LCL TEC 10 Dynamique** (FR1CIBLU1141, €350m) goes furthest: it has no coupon and no gain accrual either. Six early-redemption observations between October 2027 and October 2032 each test the same 3.20% level, and if it survives them the entire return — up to roughly €210m of upside on €350m — turns on **one TEC10 fixing in 2033, at a 3.20% strike**, against a spot of about 3.3–3.4% at issue. It is the purest expression in the dataset of the trade described in Part 3, and on any measure the largest single digital written on the index.

What the table says

Distribution is uniform in form and split in channel. Every deal in both tables is sold as a *unité de compte* inside assurance-vie, capitalisation or PER wrappers, several also via *compte-titres*. What differs is who sells it. The €30m tranches go through independent advisers and platforms — Abeille Vie / Abeille Épargne Retraite (AFER), KOREGE (CCF), Suravenir (Linxea), BoursoBank, Mutuelle d'Ivry–La Fraternelle; OCCITAUX is explicitly a public offering in France. The €250–500m tranches go through the LCL and Crédit Agricole branch networks, in €100 denominations, to savers who did not go looking for the product. That distinction matters for §4.8: a €100 minimum ticket sold across a branch network is a different conduct proposition from a €1,000 minimum sold by an adviser.

Issuer concentration is high and the SPV pattern near-universal. By count it is a tie — nine lines from the BNP Paribas vehicles, nine from Crédit Agricole's three. By notional it is not close: Crédit Agricole is three quarters of the market and BNP Paribas is 5%. CIC is the only issuer selling its own paper directly rather than through a Dutch or Luxembourg SPV with a parent guarantee.

Barriers have drifted with the OAT selloff, and the coupon with them. This is the cleanest trend in the data:

	Coupon barrier	Autocall barrier	Headline coupon
Oct 2024	3.15%	2.35%	6.30%
Nov 2024 – Jan 2025	3.50%	2.20–2.50%	5.50–6.25%
Mid 2025	3.40–3.80%	3.00–3.30%	7.50%
Late 2025 – early 2026	4.00–4.10%	3.20–3.25%	8.25%
Mid 2026	4.15–4.28%	3.50%	8.00%

For reference, TEC10 fixed at **3.97% on 23 July 2026**, against roughly 2.9–3.3% through 2024 and 3.2–3.6% through 2025. Historical span since March 1996 is about **−0.4% to +6.7%** (monthly average 6.65% in March 1996; official AFT fixings of −0.35% on 27 November 2020).

Design is diverging. Three distinct mechanisms are now in the market: the classic **dual-barrier** autocall (BNPP, JPM, CIC — separate coupon and autocall yield levels); the **cumulative-coupon** autocall (CCF/SG, IMPULSION, EXCELTIS — redemption triggers once accrued coupons reach 9–16%, so the autocall is a function of coupon history rather than a second yield barrier); and the **issuer-callable** note (SG's ELAN — the *issuer* owns the call, and the final payoff is inverted, paying more if TEC10 ends *above* 5.10%). Memory features are increasingly standard.

2.3 The parallel US wave

Worth a paragraph because it is much larger and independently confirms the driver. SEC full-text search on 424B2 filings:

Query	2023–24	2025	2026 to 1 July
`"10-Year CMT"`	**0**	**186**	**182**
`"Constant Maturity Treasury"`	4	105	83

The four 2023–24 hits are not structured notes — they are Morgan Stanley fixed-rate-reset subordinated notes off 5Y CMT, i.e. bank capital. **Retail CMT-linked issuance in the US was effectively zero before 2025.** The dominant format is a callable range accrual: RBC, Barclays, Citigroup, JPMorgan, Morgan Stanley, BofA, Goldman and Jefferies all print them, and a **5.00% upper barrier on 10Y CMT is near-universal** (BofA 5.20%, Barclays 4.75%, JPM 4.60% are the exceptions). Disclosed sizes are small — $3.0m, $3.7m, $4.25m — consistent with reverse-inquiry and private-bank distribution rather than wirehouse retail. Citigroup is the outlier, printing 20-year inverse floaters at `54.00% − 10 × (10Y CMT)` collared 0–20%.

So: **France sells digitals and autocalls to insurance-wrapper retail; the US sells range accruals to private banks.** Different payoffs, same underlying idea, same start date.

2.4 Why now

The Barclays authors' thesis is the right frame. Term premium returned to sovereign curves after 2022 as QE reversed; spot curves steepened; and because forward yields inherit and amplify spot steepness whenever term repo sits well below bond yields, the *forward* CMT curve steepened more. Their September 2025 numbers make the point: an OAT 10y spot of 3.42% against a 10y10y forward of **5.56%**, with gross carry of 2.14% decomposing into 0.55% slope, 1.09% financing and 0.50% convexity.

That gap is the raw material. A digital struck at 3.80% looks conservative against a 3.24% spot and looks aggressive against a 4.4% five-year-forward. The issuer sells the first framing and hedges the second.

Two things then made the CMT format specifically attractive rather than just CMS. First, the sovereign-versus-swap divergence the Barclays note describes: government curves steepening under term premium while swap curves flatten under pension and insurer duration demand, so a bond-yield underlying delivers materially more carry than a swap-rate underlying at the same tenor. Second, distribution: "the French 10-year government bond rate" is a phrase a retail saver in an assurance-vie contract will accept, and "the 10-year EUR constant maturity swap rate" is not. The CNO-TEC10 is also already a documented ISDA floating rate option — the JPMorgan final terms for XS2755637142 name the underlying verbatim as **"EUR-CNO TEC10 Floating Rate"** — which removes any documentation friction.

2.5 The fixing calendar

Because every one of these notes settles on a single daily print of a single index, the market has a *schedule*. It is worth writing down, because nothing about the risk is legible until you can see when the notional actually looks at the index.

I have reconstructed the observation schedule for **22 deals, €3.24bn of notional and 325 observation dates running from 21 October 2025 to 30 July 2038**, from the Final Terms and brochures in the source pack. 210 of those dates are printed verbatim in a document; 115 are derived from a printed rule (for example "chaque 12 mars, de 2027 à 2038") and are tagged as derived in the accompanying `fixings.csv`. Twenty-one of the 22 publish a size; the twenty-second, FR0014018F54, is carried at the €30m that every other BNPP tranche prints, which is an inference and is flagged as one. Nine further deals with a confirmed ISIN are excluded because they publish no schedule, no size, or neither.

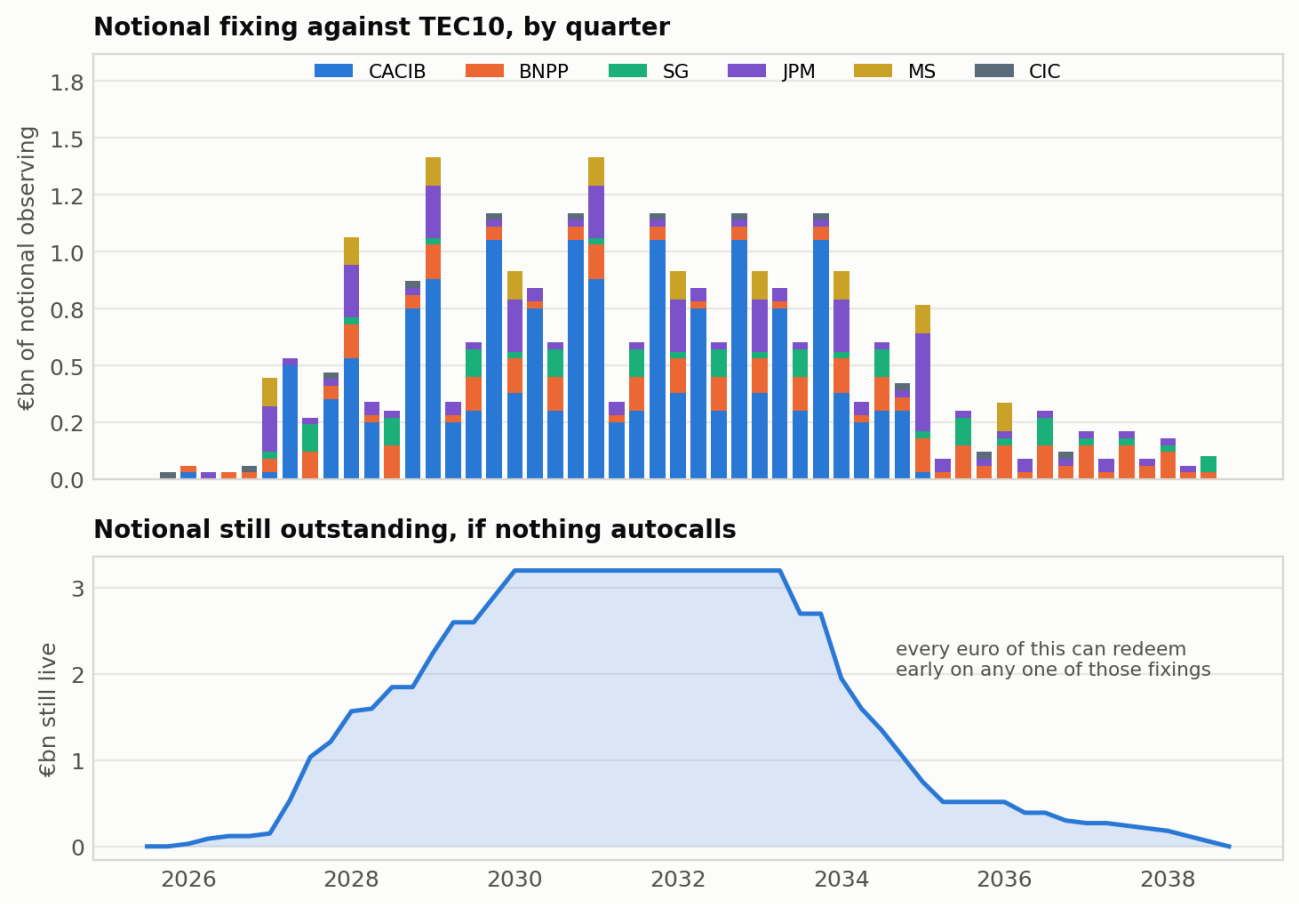

Notional observing against TEC10, by quarter, and the notional still outstanding

Four features of the calendar matter.

The book only really switches on in 2027. Almost nothing observes before then: the CACIB and LCL structures — three quarters of the notional — have **no observation at all for their first two to three years**. Deferred first observation is a design choice with a clear rationale (it lets the issuer sell a higher headline coupon by removing the cheap early autocall), and its consequence is that the crowded-unwind risk in Part 7 is not a 2026 or 2027 story. Counting each deal once per year, €1.57bn observes in 2027 and €2.25bn in 2028; from **2029 to 2033, €3.20bn — essentially the whole book — looks at the index every year**.

The dates cluster hard, and they cluster inside one banking group. Two calendar windows carry most of the observations. Summed across all years, **October alone accounts for €7.7bn of observation-notional and March for €5.9bn**, against €4.1bn for April, €2.9bn for July, €2.3bn for February, €1.3bn for January, €1.0bn for September and less than €0.4bn for each of May, June and December; there are no November observations at all. Take the two clusters as they actually sit — 23–31 October and 23 March–5 April — and they carry **€5.6bn and €8.0bn** respectively: 49% of all observation-notional in the calendar falls inside twenty-three days of the year. On a single date, 29 October in each of 2029, 2030 and 2031, **€780m of notional fixes against the same 11:00 print**.

Date	Deals	Notional fixing	Groups
29 Oct 2029, 2030, 2031	3	**€780m** each	CACIB, BNPP
30 Oct 2028; 29 Oct 2032; 31 Oct 2033	2	**€750m** each	CACIB

Date	Deals	Notional fixing	Groups
1 Apr 2033	2	**€750m**	CACIB
2 Apr 2027; 31 Mar 2028; 29 Mar 2029; 2 Apr 2030; 31 Mar 2031; 2 Apr 2032	1	**€500m** each	CACIB

The barriers do not cluster, so the book unwinds in layers rather than all at once. Coupon barriers run from 3.15% (ONYX 6) to 4.28% (CCF Stratégie), with SG's ELAN outside the range at 5.10% — and ELAN's 5.10% is not a coupon barrier at all but an inverted maturity payoff, so it does not belong in the comparison. Autocall triggers run from 2.35% to 4.11%. A single TEC10 level therefore does not call the whole market; it peels off a layer of it, and the layers are thick:

If TEC10 fixes at or below	Notional called	of which Crédit Agricole
4.04%	€650m	€650m
3.60%	€900m	€900m
3.20%	€1,495m	€1,250m
2.90%	€1,925m	€1,650m
2.40%	€2,655m	€2,180m

The two largest deals in the whole dataset — LCL Vision Jan 2025 at €500m and LCL Vision Sept 2025 at €400m — trigger only at **2.40% and 2.90%**. So a rally that stops at 3.20% still leaves **€1.23bn of the €2.48bn Crédit Agricole calendar book running**, and €1.74bn of the market as a whole. That is worth being precise about: the crowded-unwind risk is real but it is *graduated*. Against the 3.97% TEC10 fixed on 23 July 2026, the shallowest layer is 37bp away and the deepest — the €500m LCL Vision line — needs a **157bp rally** before it moves.

The 2026-vintage CACIB deals are a further exception. They set the barrier *relative* to the initial fixing (initial + 0.30%), which floats the trigger with the market at inception and, unlike everything else in the sample, does not embed a bet on where the curve happened to stand in a particular marketing week — nor does it contribute to the round-number clustering §4.8 flags as a fixing-integrity surface.

And the schedules themselves are not clean. In transcribing them I found: ONYX 6 printing `23/10/2024` and `30/10/2024` where 2034 is plainly meant; a Phoenix brochure printing `31 janvier 20230`; the Afer note giving two different final observation dates (08/02/2035 in the annual list, 06/02/2035 as *Constatation Finale*); the ELAN brochure printing an ISIN that does not match the ISIN on the Final Terms; and one BNPP brochure omitting its coupon payment rows altogether. Two are still unresolved rather than merely corrected. The Afer note's duplicate final observation is carried in the dataset as two rows. And the ELAN line (‖ in §2.2a) is inconsistent across its own documents on three counts — product name (Mai versus Juin 2038), ISIN (the brochure prints FR0014019M21 against the Final Terms' FR0014018CV8) and final observation (30 July 2038 against 29 June 2038). The calendar carries the brochure's 30 July 2038, which is why that is its stated end date; these may be two adjacent tranches rather than one document set, and the €40m size inherits the same doubt. None of this is fatal on its own, and a calculation agent will resolve it without drama. But they are typographic errors in the *only* documents a retail holder has, describing the *only*

dates on which the product does anything, in a note that runs for twelve years. Set against §4.8, they belong in the same file.

Part 3 — What the investor is actually selling

Take the BNPP archetype: 12 years, 100% capital guaranteed at maturity, 7.50% coupon if TEC10 ≤ 3.80% at the annual observation, autocall at par if TEC10 ≤ 3.30%. The barriers nest, so there are three states:

- **TEC10 > 3.80%** — no coupon, note survives. Investor gets par in twelve years, which at 2% inflation is a ~21% real loss.
- **3.30% < TEC10 ≤ 3.80%** — coupon paid, note survives. This 50bp corridor is the *only* state that is good for the investor.
- **TEC10 ≤ 3.30%** — coupon paid once, note dies at par.

The investor is short a strangle on the 10-year OAT yield, struck around 3.30%/3.80%, rolled annually for up to twelve years. **The autocall is not a sweetener; it is the short leg of that strangle.** It is what pays for a headline coupon that would otherwise be unfinanceable at par.

3.1 The barriers sit below the forward curve

This is the single most important picture. All numbers below come from a deliberately simple illustrative model — a normal (Bachelier) AR(1) for TEC10 at 90bp/yr vol, with the drift pinned to a forward curve rising from the 3.24% spot at launch toward the 5.56% 10y10y forward, discounted at 3.8%. It is an order-of-magnitude tool, not a pricing model, and the model file is included with this note.

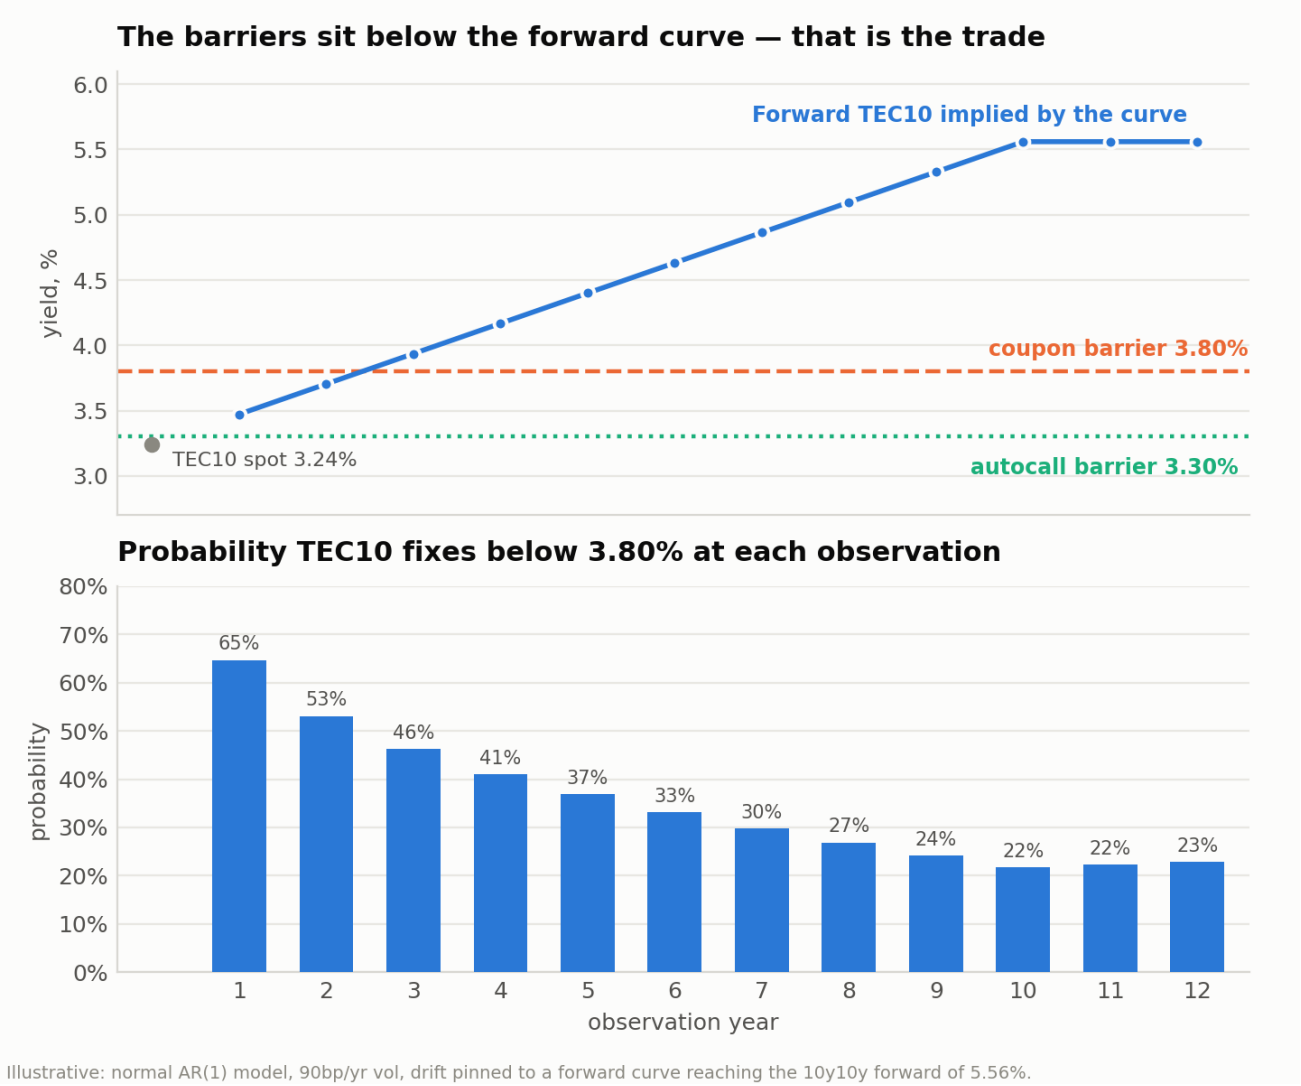

Illustrative: normal AR(1) model, 90bp/yr vol, drift pinned to a forward curve reaching the 10y10y forward of 5.56%.

Forward curve versus barriers, and coupon probability by observation

Expected life	**5.3 years**
P(autocall at year 1)	**42%**
P(runs the full 12 years)	**33%**
P(coupon) at observation 1 / 6 / 12	**65% / 33% / 23%**
Expected number of in-the-money observations (of 12)	**4.2**

Set the forward curve flat at spot instead — i.e. strip out the term premium — and the issuer's liability rises from 91.1 to 97.8 per 100. **Roughly seven points of the structure's value is the term premium embedded in the forward curve.** That is the trade, stated as a number.

The headline "7.50% per annum" therefore describes a coupon that on this model is earned about four times in twelve observations, on a note whose expected life is five years and which has a one-in-three chance of running to maturity paying nothing.

Part 4 — The issuer's risks

4.1 Why a CMT payoff is harder than a CMS payoff

The standard convexity result is well known. A payoff linear in a par rate observed at a future date is not worth the forward rate, because the measure under which the rate is a martingale (the annuity measure) is not the measure under which you discount (the T_p-forward measure). Writing $\alpha(S) = P(T, T_p)/A(T)$,

$$\mathbb{E}^{T_p}[S_T] - S_0 = \frac{\mathrm{Cov}^{A}\big(S_T, \alpha(S_T)\big)}{\mathbb{E}^{A}[\alpha(S_T)]} \;\approx\; \frac{\alpha'(S_0)}{\alpha(S_0)}\,\mathrm{Var}(S_T)$$

which is always positive. For a 10-year annual annuity at $S_0 = 3.35\%$ the multiplier α'/α is about **5.06**, giving:

Normal vol	1y	5y	10y	12y
60bp	1.8bp	9.1bp	18.2bp	21.9bp
90bp	**4.1bp**	**20.5bp**	**41.0bp**	**49.2bp**
120bp	7.3bp	36.4bp	72.9bp	87.4bp

This is first-order, not a refinement: the twelve-year adjustment at 90bp vol is roughly the same size as the entire distance from spot to the barrier. A 25% error in the assumed wing of the vol surface moves the effective forward at the far observations by 10–20bp.

Hagan's contribution is that this need not be modelled at all — it can be *replicated*. Applying Carr–Madan to $G(S) := \alpha(S)f(S)$ decomposes a CMS rate into a static portfolio of payer and receiver swaptions with weights $G''(K)$, so the desk books the strip against the vanilla book and the smile is handled consistently.

None of this survives intact when the underlying is a bond yield. Four things change:

1. **The annuity is not a traded numéraire.** For a swap rate, $A(t)$ is built from the same €STR discount bonds you value with. For a bond yield the natural annuity is the *bond's own* PV01 — a portfolio of OAT discount factors. The change of measure therefore drags in the OAT asset-swap spread and its dynamics, producing a $\mathrm{Cov}(y_T, \mathrm{ASW}_T)$ term with no CMS analogue and no hedging instrument.
2. **Repo replaces the discount curve in the forward.** $F = (P_{\text{spot}} - \mathrm{PV}(\text{coupons}))(1 + r_{\text{repo}}\tau)$. Specialness moves the forward yield directly: for a 10-year OAT of modified duration ~8.5, **20bp of specialness held for a year shifts the forward yield by ~2.4bp**. Over a twelve-year hedge with rolling, stochastic specialness this is a pricing input, not a nuisance.
3. **You cannot hold one bond.** Any fixed bond ages. Constant-maturity exposure *requires* a roll.
4. **Convention.** TEC10 is an annual actuarial yield on the French ACT/ACT convention; EUR CMS10 is 30/360 annual. Small — a few tenths of a basis point to about 1bp — but deterministic and bookable.

There is, notably, **no CMT analogue of the CMS convexity literature**. Hagan, Pelsser, Brigo–Mercurio and Andersen–Piterbarg's terminal-swap-rate framework all address swap

rates. That gap is itself a risk-management fact: every dealer's CMT model is proprietary and unbenchmarked, which is precisely what model validation and prudent valuation penalise.

4.2 The replication chain, and why nobody runs it

With bond A maturing before ten years and bond B after, $\mathrm{TEC10} = w\,y_A + (1-w)\,y_B$, and DV01-weighted positions $N_A = w \cdot \mathrm{DV01}_{\text{target}}/\mathrm{DV01}_A$, similarly for B, give a unit of yield exposure.

Direction: the note pays the investor when TEC10 is *low*, so the issuer's liability rises as yields fall. **The issuer is structurally short duration through the liability and must buy duration to hedge** — long OATs financed in repo, or receive fixed. One consolation: a long bond position *earns* specialness, so the repo leg carries positively, right up until the note autocalls and the position has to be unwound.

Three distinct roll mechanisms then leak value:

- **Continuous weight drift.** w moves daily as both bonds age; near $w \to 0$ or 1 the exposure concentrates in one leg and rebalancing is continuous.
- **Monthly pool review.** The reference *pair* can change discretely on the first business day of a month. When it does, the interpolated yield jumps by the local curvature of the OAT curve between the old and new pair plus the difference in idiosyncratic richness. A dealer holding the old pair takes an unhedged mark on the changeover date.
- **Benchmark rotation.** The AFT introduces a new 10-year line roughly twice a year. It arrives on-the-run and rich and cheapens as it seasons, so TEC10 embeds a systematically mean-reverting on-the-run premium of order 1–4bp that a swap rate simply does not have. Over twelve observations that is twelve independent draws on "where in the seasoning cycle was the reference bond on fixing day."

Add the transaction cost — 10-year OAT benchmark bid-offer of roughly 0.5–1.5 cents in normal conditions and several times that in stress (June 2024, the 2025 budget episodes) — paid through roughly 24 benchmark changes over twelve years, plus a balance-sheet charge (§4.9), and the cash chain becomes uneconomic. **In practice dealers hedge the rate in swaps, swaptions and OAT futures, and warehouse the basis.**

4.3 Delta and vega

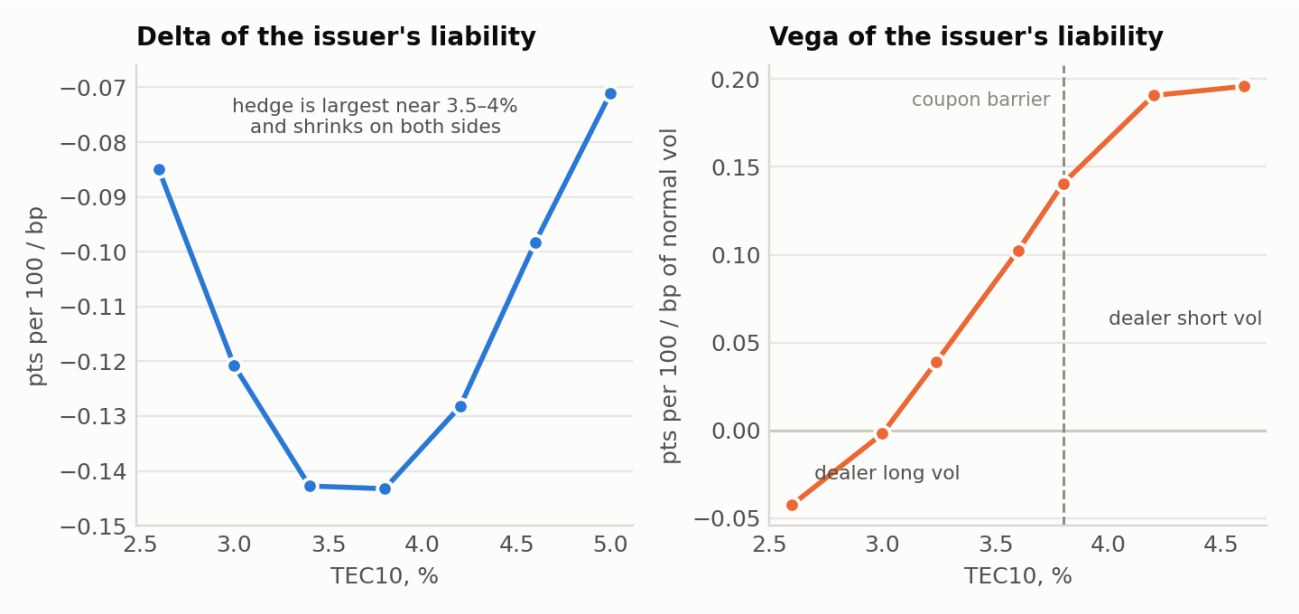

Delta and vega of the issuer's liability

Two features of the profile deserve attention.

Delta is hump-shaped, and the autocall is why. It peaks near 3.5–4.0% at about −0.14 points per 100 per basis point and falls to −0.085 at 2.8% and −0.071 at 5.2%. At low yields the note is expected to die almost immediately, so its value is pinned near par and its sensitivity collapses. **The required hedge therefore shrinks materially as rates rally** — and rates rallying is exactly when every dealer with the same book unwinds the same hedge. Expected life ranges from 2.5 years at a 2.4% spot to 11 years at 5.2%: the duration of the hedge is stochastic and systematically correlated across issuers, because the entire 2024–26 vintage was struck off the same TEC10 level with barriers clustered at round numbers.

Vega flips sign, and it does so near the autocall trigger rather than the coupon barrier. At low yields the dealer is long vol; at high yields, short. On this model the crossing sits around 3.0–3.2% — close to the 3.30% autocall level, not the 3.80% coupon level; §6.1 decomposes why. The swing is roughly 0.24 points per basis point of normal vol across a 200bp range in spot. Reversing a vega hedge is unavoidable, and it must be done exactly where gamma is largest and the market most likely to be gapping — in swaption nodes (10y10y, 12y10y, i.e. 22-year-final instruments) that carry 1–3 normal vol points of bid-offer.

Gamma has two regimes that are easy to conflate. Away from the barriers the issuer is long gamma — consistent with having bought the strangle. At each observation date there are two discontinuities, giving up to 24 barrier events over the life, where un-smoothed digital gamma diverges as $1/w^2$. The standard mitigation — booking each digital as a strike spread of width w on the conservative side — is a permanent, explicit P&L give-up.

4.4 The TEC10 versus CMS10 basis: the largest unhedged exposure

$$\text{TEC10} - \text{CMS10} = \underbrace{\text{OAT ASW}}_{\text{France credit, collateral, supply}} + \underbrace{\text{on-the-run premium}}_{\text{seasoning}} + \underbrace{\text{convention}}_{\text{ACT/ACT vs 30/360}} + \underbrace{\text{interpolation curvature}}_{\text{two-bond artefact}}$$

The dealer accepts this basis for one reason: liquidity. The EUR swaption grid supports hundreds of millions of vega at 0.25–0.5 normal vol points at liquid nodes. OTC OAT options are bilateral and thin; Eurex-listed options on the Euro-OAT future exist but trade a fraction of the Bund equivalent. There is no OAT vol surface with a credible skew.

The cost of that choice, quantified. Shock TEC10 by n basis points with the entire EUR swap curve unchanged — a pure France/ASW widening — and the swaption hedge does not move at all:

ASW widening	Change in liability	On a €500m book	On the €3.24bn fixing-calendar book
+10bp	−1.38 pts	€6.9m	**€45m**
+20bp	−2.79 pts	€14.0m	**€90m**
+30bp	−4.25 pts	€21.3m	**€137m**
+50bp	−7.19 pts	€36.0m	**€233m**

And note the sign: **a dealer who hedges a TEC10 note purely in swaps is long OAT-ASW widening — structurally short France credit.** To flatten it he must buy the OAT asset swap: buy OATs, pay fixed. Aggregated across the issuance wave, the dealer community becomes a natural bid for French paper versus swaps. Given that this wave began in the same quarter as the June 2024 dissolution and has run through two years of French fiscal and rating stress, that is a feedback channel between retail note distribution and OAT valuations that deserves more attention than it gets.

The exposure is also *amplified by the digital*. It is not a smooth duration sensitivity; near a barrier on an observation date, 20bp of ASW is the difference between a 7.50% coupon and zero. And there is no listed or OTC option on OAT asset-swap spread with any depth in which to cap it.

4.5 Inter-date correlation: the model risk no swaption pins down

Vanilla swaptions pin the *marginal* distribution of the 10-year rate at each expiry. They say almost nothing about the *joint* distribution across the twelve observation dates — and the autocall is entirely a statement about the joint distribution, since it is a first-passage-time question.

Holding every marginal fixed (each date's implied vol unchanged) and varying only the inter-date correlation:

Inter-date correlation	Liability PV	Expected life	P(runs 12y)
0.995	**83.5**	7.3y	57%
0.97	85.7	6.8y	50%
0.90	89.3	5.9y	36%

Inter-date correlation	Liability PV	Expected life	P(runs 12y)
0.80	**92.3**	5.1y	25%

An 8.8-point valuation range from a parameter that no traded vanilla instrument constrains. Mid-curve swaptions span part of it; most of it is unspanned. This is the single largest model-risk item in the book, and it is precisely the exposure a prudent-valuation additional valuation adjustment is designed to capture.

4.6 Term repo and financing

If any part of the delta sits in cash OATs, it must be financed for up to twelve years. Euro govvie repo is liquid to about a year, sporadic to two or three, and non-existent beyond. The desk prices a term-repo assumption at inception and then runs a rolling repo/€STR basis. On a €1bn OAT position, 1bp of repo-OIS basis for a year is €100k; a persistent 5bp adverse drift over twelve years is roughly €6m. The euro GC-versus-€STR basis has swung by tens of basis points since 2022 — the September and December 2022 collateral squeezes, then normalisation through QT.

This is exactly the exposure the Barclays authors single out: "*Every CMT-linked or yield-linked structure embeds a direct exposure to term repo and, crucially, to its evolution over time*." They also flag the reflexivity — when demand rises for forward bonds and CMT-linked notes, that demand itself pushes term repo up, and the term repo curve incorporates a transformation premium.

4.7 Sovereign credit, funding and balance sheet

France credit, separately from rates. Already quantified in §4.4. There is also a genuine tail: a French credit event or redenomination would move TEC10 with no move at all in CMS10, and would simultaneously make the OAT leg of any cash hedge unfinanceable.

Own funding. The note *is* twelve-year senior funding, priced off the issuer's own curve. Much of the apparent economics is funding value rather than structuring margin: on the illustrative model, moving the discount rate from 3.0% to 4.2% moves the liability from 94.4 to 89.6 — nearly five points from the funding assumption alone. Under IFRS 9, own-credit moves on fair-valued own liabilities go to OCI, so desk P&L and accounting diverge. If the note autocalls at year one, the issuer loses eleven years of pre-arranged term funding precisely when rates have rallied.

Balance sheet. French sovereign exposure is 0% risk-weighted under CRR Article 114(4), so a cash-OAT hedge is not RWA-constrained — it is **leverage-ratio constrained**, consuming leverage exposure with no offsetting risk weight, and repo netting under the LR framework is restrictive. This is the main reason the §4.2 cash chain is not run, and why desks push into OAT futures and total return swaps instead, trading balance sheet for cheapest-to-deliver optionality and counterparty risk.

4.8 Benchmark, fixing and conduct risk

These are specific to TEC10 and, in my reading, under-appreciated.

No BMR protection. Because the Banque de France is exempt under Article 2(2)(a), TEC10 carries none of BMR's administrator-authorisation, governance, cessation-consultation or orderly-wind-down machinery. For a twelve-year retail note, the fallback is whatever the note terms say — in practice calculation-agent discretion "acting in good faith", with the calculation agent being an affiliate of the issuer. That is a direct conflict of interest against the retail holder with no regulatory backstop. Compounding it: **there is no longer any cash instrument referencing the index** — the last OAT TEC 10 matured in 2009 — so the index survives as a statistical publication with no natural constituency defending its continued production.

Fixing integrity. A single-instant 11:00 CET fixing off MTS mid prices creates four distinct exposures. *Stale-fix risk*: the D-1 carry-over means a coupon can be determined off the previous day's level, an unhedgeable one-day gap that is asymmetric because disruption correlates with stress. *Thin-quote risk*: the 10bp bid-ask test and SVT polling fallback substitute a dealer poll for market prices exactly when market prices are unreliable — architecturally, a return to the pre-2012 ISDAfix model. *Gaming incentive*: with large aggregate digital notional struck at round numbers like 3.80% and 4.00%, there is a direct financial incentive to influence MTS mids in the seconds around 11:00 on observation dates, which is a live market-abuse surface. *No fixing-linked execution*: unlike the ICE Swap Rate, nothing settles against the TEC10 fix, so a dealer cannot execute "at the fix" and eats slippage at every observation.

Conduct. The AMF's Position-recommandation DOC-2010-05 on marketing complex financial instruments — and the ACPR's insurance-side twin, which is the binding text here since these are sold as *unités de compte* — applies a complexity grid whose best-known operational test is a presumption of excessive complexity where the payoff contains more than three mechanisms. Count them here: conditional digital coupon on one barrier; autocall on a second and different trigger; capital guarantee at maturity only; a non-standard underlying that is neither a familiar index nor a rate the saver can look up; and frequently a memory effect. That is four to five. How this wave sits against DOC-2010-05 is, at minimum, a live supervisory question, and the EIOPA value-for-money workstream on unit-linked products is the EU thread most likely to catch it.

Four specific disclosure gaps, none addressed by a PRIIPs KID: the 50bp winning corridor; the asymmetric horizon (roughly 40% chance of a one-year life against a one-in-three chance of twelve years at zero coupon); benchmark fallback resolved by issuer-affiliate discretion; and the fixing-integrity points above.

4.9 Capital

Under FRTB-SA the note decomposes into GIRR delta, vega and curvature on the EUR curve, **plus** CSR non-securitisation for the France sovereign component, **plus** the Residual Risk Add-On — 0.1% of gross notional for exotic underlyings and 0.01% for other residual risks, the latter explicitly capturing digital and barrier features. Under FRTB-IMA it is worse: TEC10 would almost certainly fail the Risk Factor Eligibility Test, since there are no observable transactions in TEC10-linked instruments to supply 24 real prices a year, making it a non-modellable risk factor with a punitive stressed-ES add-on. In practice this pushes the book to the standardised approach with the RRAO on top.

Part 5 — How it is actually hedged, and where the P&L leaks

5.1 The hedging stack

Instrument	What it covers	Depth	What it leaves
EUR swaptions, 1y–12y × 10y	bulk of vega and forward-rate delta	deep to ~10y expiry; 1–3 vol pts b/o at 10y10y+	the whole TEC–CMS basis; long-expiry skew
CMS caps/floors on EUR CMS10	pre-packaged Hagan replication — the right *shape*	far thinner than swaptions	CMS ≠ CMT: the entire ASW basis
Mid-curve / forward-starting swaptions	partially spans the inter-date correlation	thin in EUR	most of the 8.8-pt band in §4.5
Options on the Euro-OAT future (OOAT)	genuine OAT vol	small vs the Bund equivalent	not enough size; wrong tenor granularity
Options on Bund futures (OGBL)	very liquid vol	deep	leaves OAT–Bund spread vol unhedged
Euro-OAT future (FOAT)	OAT delta without cash balance sheet	liquid	CTD delivery optionality; 8.5–10.5y basket ≠ the TEC10 pair; quarterly implied-repo roll
Bond forwards / TRS on specific OATs	term OAT exposure, better LR netting than outright	bilateral, to ~1–5y	counterparty/CVA; no 12y tenor
Repo-OIS basis	§4.6	essentially nothing at term	everything

5.2 What Barclays is proposing, and why

The Risk.net article is explicitly a proposal to fill two of the gaps above.

Repo-OIS basis swaps. The authors decompose financing carry as `GrossCarry − TransformationPremium − BalanceSheetCost` and argue for trading the term repo-OIS basis directly — "a five-year basis swap (repo OAT vs €STR + spread), for example, reflects the expected average daily spread between the overnight government bond refinancing rate and the remuneration of overnight cash" — as the derivative analogue of equity index total return futures, isolating the economics of holding the asset from the cost of financing it without the same balance-sheet intensity. They describe the market as "still at an early stage." That matches what I could establish: euro infrastructure exists (Eurex GC Pooling, three-month €STR futures, the RepoFunds Rate France family) and bilateral term basis trades certainly occur, but I found no evidence of screen-quoted liquidity at five to twelve years. The USD analogue, SOFR versus Fed Funds, is genuinely liquid; euro is not there.

CMT-anchored volatility. The proposal is to adapt the architecture of euro cash-settled swaptions and anchor settlement to CMT fixings — "the mechanics remain identical, the same Greeks, the same hedging principles, with the underlying replaced by sovereign yield curves." Conceptually this is exactly the missing instrument: a cash-settled option settling on a *bond-yield* fixing rather than the ICE Swap Rate collapses the TEC–CMS basis into the option itself. In USD there is precedent — the Fed publishes CMT in H.15 and CMT swaps and caps/floors trade under long-standing ISDA floating rate options. In euro, the `EUR-CNO TEC10` floating rate option already exists and is used in live final terms, so the documentation

bar for a dealer-to-dealer CMT vol market is low. Whether the liquidity follows is the open question.

Their third leg — digital plays on the CMT-versus-CMS asset swap spread — is worth noting as a *product*, since it is the trade that monetises exactly the exposure §4.4 says dealers are involuntarily warehousing. Their worked example is a 10-year callable range accrual on the TEC−CMS spread, 100% capital protected, 5% annual coupon accruing while the spread stays below 1%.

5.3 Where the money actually leaks

The honest description is that the dealer is not replicating. He is **warehousing a twelve-year exotic against an approximate hedge**, funded by the embedded autocall and by twelve years of term funding, against leakages that plausibly consume several points over the life:

	Source	Rough scale
L1	**Convexity/CMT adjustment error** — 22 to 87bp at 12y depending on the vol wing; a 15bp forward error at an observation with 0.15–0.30% of notional per bp of digital delta	2–5 pts at that observation
L2	**Repo/financing carry** — no term repo curve past ~2y; 10bp adverse drift on €500m over 10y	~€5m, ~1% of notional
L3	**Uncertain hedge horizon** — ~40% chance of unwinding a 12-expiry swaption strip after one year, ~33% of carrying it to twelve; 1–2 vol pts b/o at 10y10y+ receiver nodes	0.5–1.5% of notional
L4	**Digital over-hedge reserve** — strike-spread of width w on the conservative side, ~½ × coupon × P(landing in the window), per observation, twice (two barriers), twelve times	0.3–1.0% of coupon PV per observation
L5	**OAT–swap basis run outright**	€14m per 20bp on €500m
L6	**Skew on a non-traded underlying** — the digital price is a derivative of the smile, and the TEC10 smile has never been observed	the largest valuation uncertainty
L7	**Fixing slippage and D-1 carry-over**	twelve unhedgeable instants

Part 6 works the same book through as a trade — what gets executed on day one, and what has to be re-executed a year later in three different markets — which is where these leaks become visible as tickets rather than as line items.

The corollary matters for the conduct discussion in §4.8: **the margin has to be large because the unhedgeable residual is large.** A structure whose hedging error is measured in points cannot be sold on a margin measured in tens of basis points. The complexity and the cost are causally linked — and the cost is borne by an assurance-vie saver whose KID discloses neither the 8.8-point model-risk band nor the value of the call they have written.

Part 6 — The hedge, step by step

Parts 4 and 5 describe the exposures and the instrument set. This part walks one deal through as a trade: what the desk books on day one, what it executes, and what it has to do a year later in three different markets. Every number in this part comes from `hedge_model.py`, which takes the same archetype and the same market parameters as the rest of the note and does one additional thing — it buckets the liability's risk **by observation date** and converts each bucket into the notional of an instrument a euro rates desk actually trades. It is a second script rather than an extension of the first, and the two are not identical: to make bump-and-revalue stable it smooths both digitals as 5bp strike spreads and builds the paths from an exact covariance matrix rather than by stepping the recursion. The base case therefore comes out at 91.44 and 5.24 years against Part 3's 91.13 and 5.33 — a third of a point and a month apart, which is the right order of agreement to expect and is worth stating rather than papering over. It is an order-of-magnitude tool, and the tickets are the point, not the decimals.

The deal is the BNPP archetype: **€30m, twelve years, capital guaranteed at maturity, 7.50% coupon if TEC10 fixes at or below 3.80% on the annual observation, autocall at par if TEC10 fixes at or below 3.30%.** Launch spot 3.24%, which is *below* the autocall barrier. It does not follow that the note is odds-on to die at the first observation, and this is the first place the term-premium point from §3.1 bites: the one-year *forward* is 3.47%, not 3.24%, so the model puts P(autocall at year 1) at **43%**, not better than even. Forty-three per cent on a single date is still far more probability mass than any other observation carries, and that concentration drives everything the desk does.

6.1 Day one

The desk takes in €30m of twelve-year funding and books a liability worth **91.44 per 100**, or €27.4m. The €2.6m gap is funding value, structuring margin and reserves in a proportion the desk knows and the buyer does not. Treat the size of that gap with care rather than as a fee estimate: §4.7 shows the liability moving from 94.4 to 89.6 as the discount rate goes from 3.0% to 4.2%, so on the funding assumption alone the gap ranges from about €1.7m to €3.1m, and it does not close entirely until the discount rate falls to roughly 1.7%.

Then it looks at the risk, and the risk is not one number. The liability depends on twelve different forward 10-year rates — the rate as it will be observed in one year, in two years, and so on — and the natural hedging instrument for each is a forward-starting swap and a swaption at that expiry. The desk's risk report therefore looks like this:

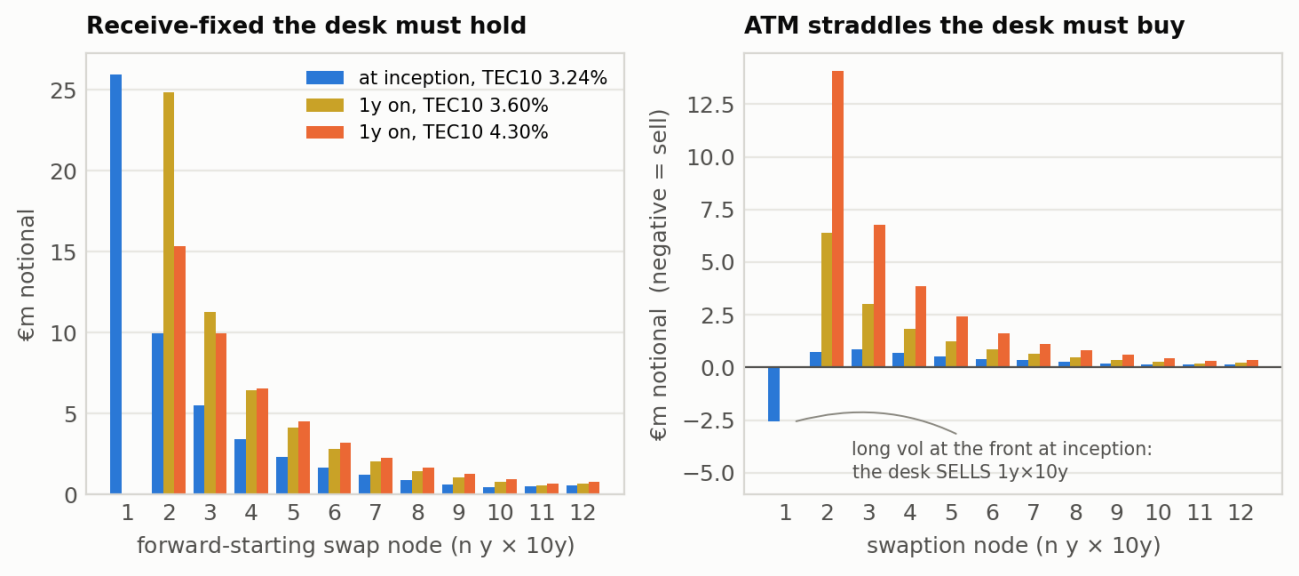

The hedge ladder at inception and one year on

Observation	Node	Liability delta	Liability vega	Receive fixed	Straddle
		pts/bp	pts/bp of vol	€m	€m
1	1y × 10y	−0.0702	−0.0055	**26.0**	**−2.6**
2	2y × 10y	−0.0260	+0.0022	9.9	+0.8
3	3y × 10y	−0.0139	+0.0030	5.5	+0.9
4	4y × 10y	−0.0084	+0.0027	3.4	+0.7
5	5y × 10y	−0.0054	+0.0022	2.3	+0.5
6	6y × 10y	−0.0038	+0.0018	1.7	+0.4
7	7y × 10y	−0.0027	+0.0016	1.2	+0.3
8	8y × 10y	−0.0019	+0.0012	0.9	+0.3
9	9y × 10y	−0.0013	+0.0010	0.6	+0.2
10	10y × 10y	−0.0009	+0.0007	0.5	+0.1
11	11y × 10y	−0.0009	+0.0007	0.5	+0.1
12	12y × 10y	−0.0010	+0.0008	0.5	+0.1
	Total	**−0.136**	**+0.013**	**53.0**	**1.9 net / 7.0 gross**

Read across the "receive fixed" column and the shape of the trade is obvious. **Half the delta sits at the first observation.** €26m of the €53m ladder is a one-year-forward 10-year swap, because that one date carries a 43% chance of ending the trade and a 66% chance of paying a coupon. The remaining eleven observations, spread over eleven years, carry between them slightly less risk than the first one does alone.

The vega column is the one that surprises people. Net vega is small — €3,754 per basis point of normal vol on a €30m note — but that nets a **long-vol** front against a **short-vol** back. At the first observation the liability *falls* when vol rises, so the desk is long volatility there and *sells* €2.6m of 1y × 10y straddles. Everywhere else it buys. Gross straddle notional is €7.0m against a €1.9m net: the net figure is useless as a description of what has to be executed.

The front-node sign is worth pulling apart, because the obvious explanation is the wrong one. It is tempting to say that more vol makes a near-certain autocall less certain and that this helps the issuer. The model says the opposite: **it is the coupon digital, not the autocall, that makes the front node long vol.** Strip the coupon out and price the autocall alone and node-1 vega is +0.0124 — short vol; add the coupon back and it swings by **−0.0179** to the full note's −0.0055. (The two features cannot simply be priced separately and added: removing the coupon changes the continuation value at the autocall barrier from 91.4 to 83.4, which nearly doubles the autocall's digital amplitude. The marginal contributions above are the meaningful decomposition; the standalone figures are not additive.) The economics are that the coupon digital is 66% in the money at the first observation while the autocall is 43%: a digital that is comfortably in the money is *hurt* by volatility, and the issuer, who is short it, is correspondingly helped.

And then there is the digital, which is where the notional stops being polite. The coupon pays 7.50% × €30m = €2.25m if TEC10 fixes at or below 3.80% and nothing if it fixes at 3.81%. Booked as a receiver spread — the standard treatment, and the one §4.3 describes as a permanent give-up — the replicating notional is:

Strike spread	Receiver spread	Notional per observation date
2bp	3.80% / 3.78%	€1,343m
5bp	**3.80% / 3.75%**	**€537m**
10bp	3.80% / 3.70%	€269m
25bp	3.80% / 3.55%	€107m

The autocall is a second digital and a larger one. At a 3.30% fixing the desk must pay par, while the eleven-year note it is extinguishing is worth 91.4 — an amplitude of **8.6 points, or €2.6m**, which at a 5bp spread is another **€617m** of receiver-spread notional. The two digitals point the same way, so they add:

> **€1,154m of gross receiver-spread notional, per observation date, against a €30m note** — thirty-eight times the notional of the trade, on one date, before any of the twelve are added together.

The multiple is a function of the strike width the desk chooses, and the table above is the sensitivity: at 2bp it is 96×, at 25bp 8×. That is the point rather than a caveat on it. The desk picks the width, the width sets both the gross notional it must manage and the size of the permanent give-up it books, and nothing in the market tells it what the right number is. It is also why the replication is not really run: the desk books the digitals conservatively at a wide spread, takes the give-up, and manages the residual.

Finally, what the desk does **not** do:

	Ticket	Direction	Size	Traded?
1	EUR forward-starting 10y swaps, 1y–12y	**receive fixed**	€53m across 12 nodes (€49m 10y-equivalent)	yes, deep
2	1y × 10y ATM straddle	**sell**	€2.6m	yes, deep
3	2–12y × 10y ATM straddles	**buy**	€4.4m	thins past 10y expiry
4	Receiver spreads at both barriers, 12 expiries	**buy**	≈€1.15bn per date	booked, not fully executed

	Ticket	Direction	Size	Traded?
5	10y OAT asset swap	**buy**	€49m	**no** — warehoused (§4.4)
6	Cash OAT / repo replication chain	long OATs, term repo to 12y	—	**no** — §4.2, §4.6

Ticket 5 is the one to hold onto. The swap and swaption hedge is a hedge against the *euro swap curve*. The note references an *OAT yield*. To close the gap the desk would have to buy roughly **€49m of 10-year OAT asset swap for every €30m of note** — buy the bond, pay fixed. It does not, because the swap market is where the liquidity is; and so, as §4.4 says, it ends up structurally short France.

6.2 One year on, case A — the sell-off. TEC10 fixes at 4.30%

No coupon. No autocall. Eleven observations left.

The mark is good: the liability goes from 91.44 at inception to **78.08** a year later. The 13.4-point difference is not a clean P&L — it mixes the rate move with a year of accretion toward par at the 3.8% discount rate, so the rate-driven part is larger than 13.4 points — but the direction and the order of magnitude are right, and on €30m it is a few million euros. It is worth being clear that this is the investor's loss, not value the desk has created, and that it is only realised to the extent the hedge tracked.

What changes in the risk is more interesting than the mark.

The hedge has to be rebuilt further out. Expected life goes from 5.2 years to **8.6 years**, and the probability of running the full twelve from 32% to 68%.

Two tickets follow, and they point in opposite directions. The 1y × 10y swap does *not* expire at the observation — it becomes a spot-starting 10-year swap — but the risk it was hedging is gone, so it has to be actively closed: **pay fixed on €26.0m**. Then, across the 2y–12y ladder, the desk **receives a further €20.2m**. Gross swap tickets **€46.2m**; the total ladder goes 53.0 → 47.2, so on a net basis almost nothing has happened, which is exactly why net risk numbers are the wrong way to think about the execution cost. The shape changes completely: where the front node was half the risk, the 2y × 10y node is now €15.3m out of €47.2m and the tail beyond five years has roughly doubled.

Vega moves an order of magnitude more than delta. Liability vega goes from €3,754 to **€34,686 per basis point of normal vol** — a factor of nine — while the DV01 moves by 15%. In instruments, the desk buys back the €2.6m front-node straddle and then **buys €27.9m of straddles across the strip, €13.3m of it at the 2y × 10y node alone**: gross vol tickets of €30.5m against a strip that was €7.0m gross the year before.

The reason is the same reason the mark moved. At 3.24% the note was a near-certain autocall — an almost-decided outcome, and almost-decided outcomes have little vega. At 4.30% every one of the eleven remaining observations is a genuine question about whether TEC10 comes back down through 3.80%, and the value of the structure is now dominated by volatility rather than by the near-certainty of early death. The desk has gone from owning a short-dated question to owning eleven open ones.

This is also where the position bites. The desk is **buying volatility after a sell-off**, which is normally when implied vol is bid, and it is buying it at 2y × 10y and beyond where the bid-

offer is 1–3 normal vol points (§5.1). It carried very little vega when vol was cheap and has to add a lot when it is not — a vanna effect that costs real money and does not appear in any first-order risk report.

Scale that to the whole market and it is a large number — of the order of **€5bn of gross swap tickets and €3bn of straddle buying** across the €3.24bn in the fixing calendar, from a dealer community that is largely the same five names. Take the scaling as an order of magnitude and no more. It is one archetype multiplied up, and three quarters of that market is eight- to nine-year Crédit Agricole paper with no observation for two to three years, a different payoff and a different expected life; the direction of the flow generalises, the size does not.

6.3 One year on, case B — the small move. TEC10 fixes at 3.60%

The coupon **is** paid — 3.60% is below the 3.80% barrier — and the note is not called, because 3.60% is above the 3.30% trigger. This is the corridor Part 3 identifies as the only good state for the investor, and it is also the case that generates the **largest delta rebalancing of the three**, on a move less than half the size of case A's. That combination is the non-obvious result. (Case A is still the more expensive scenario overall, because its extra €17m of straddle buying is executed at 1–3 vol points of bid-offer against a few tenths of a basis point on the swaps.)

The liability is **87.28**; expected life 6.3 years. And the delta ticket is the largest of the three cases. The same two trades as case A — close the €26.0m front node, then rebuild along the ladder — but the rebuild is bigger: **receive €29.0m across 2y–12y**, of which €14.9m at the 2y × 10y node alone. **Gross swap tickets €55.0m against case A's €46.2m** — about a fifth more execution, for a move less than half as large.

Vega rises too, but far less than in case A — €3,754 to €16,924 per basis point, and €13.6m of gross vol tickets against case A's €30.5m.

The mechanism is digital gamma. With TEC10 at 3.60%, the next observation sits between a coupon barrier and an autocall trigger that are 50bp apart — well inside one year's worth of vol. Both digitals are live, both are close, and the delta of a digital is largest exactly at its strike. A desk that sizes its hedging capacity and its bid-offer budget off large-move stress scenarios will under-provision for the small-move case, which is both more likely and more demanding.

There is a second-order point worth flagging. In both non-triggering cases the desk closes the same €26m front node and then rebuilds *further out along the curve* — the flow is the same shape whether rates rose 36bp or 106bp. What reverses it is not the direction of the move but whether a barrier was crossed, which brings us to case C.

6.4 One year on, case C — the rally. TEC10 fixes at 2.90%

The coupon is paid, the autocall triggers, and the note redeems at par. Twelve years of pre-arranged funding disappear after one (§4.7), and the entire hedge is unwound in a day:

- **pay fixed on the whole €53.0m ladder**, front node included;

- **close the €7.0m gross straddle strip** — buying back the front-node short, selling the eleven long nodes;
- tear up the €1.15bn of notional receiver-spread replication and release the digital reserves.

Note what has happened to the size of the ticket. In cases A and B the desk closed €26m and rebuilt €20–29m, so the *ladder* survived. Here the whole thing goes, at once, on the day the barrier is crossed. That is the digital discontinuity of §4.3 — not the smooth long-gamma region away from the barriers, but the point where un-smoothed digital gamma diverges — and it is why the position cannot be unwound gradually.

The direction of the delta leg is at least benign: paying fixed into a rally pushes swap rates back up, which leans against the move rather than amplifying it.

The volatility leg is the problem, and it is not diversified. Unwinding means every desk in the cohort selling the same long-dated straddle strip on the same morning. Each of them individually would be selling into a market that wants receiver optionality after a rally; collectively they are the supply, and the price they get is set by how much of it arrives at once. That is leak L3 in §5.3 crystallising, and it is where the crowded-door problem lives.

How crowded? §2.5 gives the arithmetic. On 29 October in each of 2029, 2030 and 2031, **€780m of notional observes against the same 11:00 fixing**; six further dates between 29 March and 2 April, one in each year from 2027 to 2032, carry €500m each. Scaling this deal's unwind across the whole €3.24bn in the calendar puts the theoretical maximum at roughly **€5.7bn-equivalent of received-fixed and €4.4m per basis point of DV01** — with the same caveat as §6.2 about multiplying one archetype up.

But it will not all go at once, and §2.5's trigger ladder is the reason. The autocall levels are spread from 2.35% to 4.11%, so a rally peels the book off in layers: **a fixing at 3.60% calls €900m, at 3.20% €1.5bn, at 2.90% €1.9bn, at 2.40% €2.7bn.** The two largest deals in the market — the €500m and €400m LCL Vision lines — do not trigger until 2.40% and 2.90%, which are 157bp and 107bp below the 3.97% TEC10 fixed on 23 July 2026. The realistic version of the crowded unwind is therefore not one morning but a sequence of them, each triggered by a further leg of a rally, each concentrated on the same two weeks of the calendar, and each hitting a dealer community of about five names.

6.5 What the walkthrough shows that a static description does not

Gross risk, not net risk, is what the desk actually trades. €30m of note generates a €53m swap ladder, a €7m gross straddle strip whose *net* is €1.9m, and — at the desk's chosen 5bp strike width — €1.15bn of receiver-spread notional per observation date. Every one of the three scenarios has a large gross ticket sitting behind a small net change. Structuring fees are quoted in tens of basis points against that.

The hedge's maturity is itself a derivative. 5.2 years at inception, 8.6 after a sell-off, zero after a rally. There is no other common rates product where the desk cannot tell you, within a factor of two, how long it will be carrying the position. Every bid-offer assumption in the pricing is conditional on an unknown horizon.

Vega moves an order of magnitude more than delta. Between inception and the case A sell-off the liability's DV01 moves by 15% and its vega by a factor of nine; in case B, DV01 barely moves at all while vega more than quadruples. A desk that thinks of this as a delta product will be right about the direction and wrong about the size of what it has to trade.

The largest rebalance comes from the smallest move. Case B, a 36bp drift into the corridor, needs about a fifth more gross swap execution than case A's 106bp sell-off — €55.0m against €46.2m. Digital risk does not scale with the size of the move; it scales with proximity to the strike.

And none of the four instruments hedges the thing the note is actually written on. Every ticket above is a euro swap-curve instrument. The underlying is a French government bond yield. The gap is §4.4, it is worth roughly €90m across the fixing-calendar book for a 20bp OAT asset-swap widening, and it is carried, not hedged.

Two caveats on all of it. The model is a normal AR(1) on a single rate with a forward curve, digitals smoothed at 5bp, discounting flat. It has no smile, no OAT–swap basis dynamics, no term-repo curve and no stochastic vol, so it cannot produce a price. What it can do — and what a fuller model would not change qualitatively — is show the *shape* of the risk: where on the curve it sits, how it redistributes when the market moves, and how large the gross tickets are relative to the trade. Second, every market-wide figure in this part is one €30m twelve-year annual dual-barrier note scaled linearly by notional. Three quarters of the market is eight- to nine-year Crédit Agricole paper with a deferred first observation and a different payoff, whose delta, vega and expected life bear no fixed relation to the archetype's. The signs and the orders of magnitude carry across; the specific billions do not.

Part 7 — What to watch

1. **Does the CMT vol market actually form?** The Barclays proposal is coherent and the ISDA plumbing already exists. If euro CMT-anchored cash-settled swaptions start trading dealer-to-dealer, the economics of this product improve materially and issuance should accelerate. If they do not, the basis stays warehoused and the wave is capped by dealers' risk appetite for French sovereign spread.
2. **The dealer bid for OAT asset swaps.** The aggregate book is now demonstrably large enough to matter: €3.33bn of *stated* size, and that is a floor. Hedged properly it implies roughly **€5.3bn of 10-year OAT asset swap to buy** (§6.1) — a persistent, non-economic bid for OATs versus swaps and for long-end receiver optionality. Worth testing against the 2024–26 behaviour of OAT ASW and 10y10y receiver skew.
3. **The first crowded unwind, and when it can happen.** §2.5 dates it. Almost nothing observes before 2027, because three quarters of the notional defers its first observation by two to three years; from 2029 to 2033 essentially the whole €3.2bn book observes every year, concentrated in two narrow calendar windows — October and late March. On 29 October in each of 2029, 2030 and 2031 there is **€780m fixing against a single 11:00 print**. But the triggers are laddered, not clustered: a fixing at 3.60% calls €900m, at 3.20% €1.5bn, at 2.90% €1.9bn, at 2.40% €2.7bn, so the unwind arrives in layers as a rally extends (§6.4). The equity autocall market has been through the concentrated version of this; the rates version has not been tested, and cannot be until 2029.
4. **The supervisory question.** DOC-2010-05 and the ACPR's insurance twin, plus the EIOPA value-for-money work on unit-linked products, are the live threads. A single AMF or ACPR communication on rate autocalls in assurance-vie would change distribution economics quickly.
5. **Benchmark governance.** An index with no BMR administrator, no cash instrument referencing it, a poll-based fallback and now **over €3bn of retail digital notional** — of which roughly €2.4bn is struck at absolute levels, many of them round numbers; only the three 2026 CACIB deals, €950m between them, use a barrier set relative to the initial fixing — is an obvious candidate for either a formal governance upgrade or an unpleasant incident.

Sources

Primary — index and sovereign instrument

- AFT — OAT TEC 10 (EN) and (FR); AFT — today's TEC 10 index
- Banque de France — *Note technique sur les indices CNO-TEC©*, rev. 07/05/2020
- CNO — *TEC10® Principle and Mechanism*, 6 June 2013
- Sénat — PLF 2002 report on State debt

Primary — deals

- AFER Rendement Novembre 2024 (FR001400T3H1, JPMorgan, €200m)
- CCF Stratégie Taux TEC10 (FR0014018P94, SG Issuer, ≤€90m) and CCF Objectif Taux TEC10 (Morgan Stanley)
- JPMorgan Final Terms XS2755637142 — underlying "EUR-CNO TEC10 Floating Rate"

- ONYX 6 (CIC, FR001400SDV1) · Rendement Taux France Octobre 2025 (BNPP) · OCCITAUX Mars 2026 · IMPULSION 13 · EXCELTIS Patrimoine Juin 2026 · ELAN (SG / BoursoBank)
- Historical: France Télécom Final Terms, 19 Dec 2011 · Areva Final Terms, 2 Apr 2012 · CADES €1bn TEC10 FRN 2020 · CADES 1998 OAI · BALO — Areva bondholder convocation, Sept 2016 · CAFFIL EMTN Base Prospectus 2026
- US: SEC EDGAR 424B2 filings, e.g. Barclays callable range accrual, Mar 2026 · JPMorgan dual-tenor 10Y/30Y CMT, Oct 2025 · Citigroup inverse CMT floater, Mar 2026

Analytical

- Samy Ben Aoun and Hamza Hoummady (Barclays), "The great term premium comeback", *Risk.net*, September 2025 (the attached reprint)
- P. Hagan, "Convexity Conundrums: Pricing CMS Swaps, Caps and Floors", *Wilmott*, March 2003
- A. Pelsser, "Mathematical Foundation of Convexity Correction", *Quantitative Finance* 3(1), 2003; D. Brigo & F. Mercurio, *Interest Rate Models*, 2nd ed., 2006; L. Andersen & V. Piterbarg, *Interest Rate Modeling*, Vol. III, 2010 (terminal swap rate framework)
- AMF Position-recommandation DOC-2010-05

Source pack

A companion pack accompanies this note. Alongside the note itself:

`fixings.csv`	the reconstructed fixing calendar — 325 observation dates across 22 deals, with size, barrier, autocall type and a printed/derived tag on every row
`fixings_notes.md`	every correction applied to a document typo, every deal excluded and why, and the totals
`tec_model.py`	the illustrative model behind Parts 3–5
`hedge_model.py`	the bucketed risk and hedge tickets behind Part 6
`TEC10_source_register.xlsx`	the claim-by-claim register

Plus **32 source documents** (every product brochure and set of final terms cited above, the CNO methodology deck, the 1996–2012 historical prospectuses and BALO notices, seven US SEC 424B2 filings), plus **`TEC10_source_register.xlsx`** — a claim-by-claim register mapping every substantive statement in this note to its evidence and confidence level, with a `Gaps` sheet naming everything that could not be retrieved or verified and how to close it. The two documents whose servers refuse direct download (the Banque de France technical note and the AFT page) are included as text extracts with their URLs.

Caveats

- **All model output in Parts 3–6 is illustrative.** It comes from a deliberately transparent normal AR(1) model (90bp/yr vol, drift pinned to a forward curve reaching the 5.56% 10y10y forward, 3.8% discounting, 400–500k paths), not from a calibrated production model. It is intended to establish signs, orders of magnitude and comparative statics, not prices. Part 6 adds a second script that buckets the same liability by observation date and

converts each bucket into an instrument notional; its hedge tickets inherit every limitation of the underlying model. In particular it has no smile, which matters for the *price* of the digital replication rather than its notional — the replicating notional is a function only of the payout, the strike width and the annuity — and it has no OAT–swap basis dynamics, which is the exposure Part 6 concludes is carried rather than hedged. Every market-wide figure in Part 6 is one archetype scaled linearly by notional across a market three quarters of which has a materially different payoff; treat those as orders of magnitude, not estimates. Both scripts accompany this note.

- **Six of thirty-one French deals have no published size**, and at least two further Crédit Agricole tranches from 2025 are known to exist but could not be enumerated, so the €3.33bn figure is a floor on the number of deals. It is not a placement total: €2.45bn of it is registered nominal on Final Terms rather than evidence of investor take-up, and €245m is a marketing cap (§2.2b). No trade-press or data-vendor aggregate exists in the public domain; structuredretailproducts.com is paywalled. For scale, the ACPR–AMF *Cartographie* (April 2025) puts total French structured-product *collecte* at €41.8bn in 2023, with rate-linked product at 13% of *encours commercialisés* — roughly €5.4bn — before this wave began.
- **The fixing calendar is reconstructed, not published.** Of its 325 observation dates, 210 are printed verbatim in a source document and 115 are derived from a printed rule; `fixings.csv` tags every row, and `fixings_notes.md` lists every correction applied to a document typo. Nine deals with a confirmed ISIN are absent from it entirely because they publish no schedule, no size, or neither.
- **Unverified items**, flagged deliberately: the ISINs of both OAT TEC 10 lines and the value of the *marge faciale* `m`; the ISIN of the CADES €1bn 2020 FRN; whether any CAFFIL/SFIL TEC10 tranche was ever drawn; the ISIN of the CACIB "Athena TEC 10 Mars 2025" deal and the terms of the two Goldman Sachs Finance Corp International lines, for which the Luxembourg exchange returned server errors on every attempt; whether the pre-2005 French TEC10 issuance list is complete (BALO full text begins in 2005, so it almost certainly is not).
- Where the BMR analysis, the trademark ownership question and the "why the State stopped issuing" question rest on inference rather than a sourced statement, the text says so.