

What to look at next

Fourteen structured-note families, ranked by how much of the risk cannot be traded away — and what it would take to do each one

Survey note — 31 July 2026 · companion to the TEC10, US CMT, Formosa and Zero Floor notes

Executive summary

This is a shortlist, not a survey of the market. Four notes are done — TEC10, US CMT, Formosa, Zero Floor — and each landed on the same shape of finding: a payoff sold for yield, an embedded option nobody on the buy side priced, and a residual the dealer cannot lay off and therefore warehouses. This note ranks fourteen further families against that shape and says, for each, what a dive would produce and what it would cost. Scope is **rates and cross-currency**, per the existing book; equity autocalls and credit are noted at the end with the reason they are excluded.

The ranking is not by market size. It is by the product of four tests: how much of the risk is *structurally* unhedgeable rather than merely expensive; whether there is a dated, documented episode to anchor the note; whether a free deal-level dataset exists so the register can actually be built; and adjacency to what the desk already owns. Families that score high on residual risk but have no public footprint and no data — callable snowballs are the clean example — are honestly worse projects than families with a smaller residual and a full evidence trail.

Three candidates come out ahead, for different reasons.

The ICE Swap Rate fixing is the cheapest and it touches everything already on the book. IBA states in its own methodology that USD SOFR, EUR €STR and the GBP/EUR inflation swap rate settings are "expected to be calculated using input data at **Level 2 or Level 3** of the Waterfall". Level 1 — the level the methodology is built around — is *executable* prices and volumes from central limit order books. Level 2 is dealer-to-client prices displayed electronically by trading venues. Level 3 is movement interpolation off the previous day. **So the entire post-LIBOR CMS complex settles one rung or two below the executable prices the benchmark was designed to use**, and IBA publishes, free and quarterly, which level each setting actually ran at. That is a monitoring dashboard and a two-week note, not a project.

Callable range accruals and multi-callables are the generalisation of Formosa. The Formosa note established the vega channel through one wrapper; the payoff is the thing, and it has two dated episodes rather than one. Nomura's published street ladder shows dealers "longer vega in a sell-off and shorter vega in a rally" on Bermudan books — wrong-way, one-directional, no natural counterparty — and puts **>\$75bn** of FHLB 1y-NC-3m multi-callables into the market around SVB in March 2023, pushing USD 6m1y ± 100 risk reversals to all-time lows. The mortgage-convexity literature (Perli-Sack: hedging amplified rate moves by 16-30%) is the ready-made official-sector template for what the write-up should look like.

UK LPI is the largest genuinely unhedgeable residual in the rates complex. A capped-and-floored, path-dependent inflation liability delta-hedged with a linear inflation asset, in a market where the LPI swap has become illiquid and — per IBA's March 2026 consultation — the *linear* zero-coupon inflation swap market still cannot generate enough executable volume to support a Level 1 fixing in EUR or GBP, and USD is deferred entirely. **If the linear market cannot support an executable fixing, the option market on top of it is model-marked by construction.** Both tails have been realised inside seven years: RPI printed **14.2%** in October 2022 against a 5% cap, and the 0% floor bound in 2015–16. Against **£619bn** of index-linked gilts, 24.5% of the UK debt stock and roughly twice the next G7 country, held by a buyer the Bank of England has described as possibly price-inelastic. And RPI is statutorily redefined to CPIH methodology in **February 2030** with no compensation and no instrument that spans the change.

Two findings are negative and worth stating before anyone plans a dive. First, **there is no free swaption volatility surface.** CBOE's SRVIX, the one free swaption-vol index, stops on 11 February 2022; MOVE is proprietary; ICE Swap Rate is licence-only by explicit statement; CME settlement files are IP-blocked. DTCC's SDR gives real swaption strikes, premiums and expiries — 2,961 rows carried an option premium on 30 July 2026 — but that is trades on a ragged grid with capped notionals, not a surface. Any dive that needs to mark a surface needs a vendor. Second, **several of these families have no public size figure at all:** PRDC, quanto structures, rates TARNs, snowballs, dual-currency deposits, LPI swaps, EUR CMS-linked notes. Where that is true this note says so rather than repeating a number from a search snippet.

One data finding closes a gap in the Formosa note and is worth acting on today. That note recorded that "TPEx's complete professional-investor register, its factbook and its download endpoints all returned HTTP 403", and built its register bottom-up from 60 bonds. TPEx runs an **undocumented, unauthenticated OpenAPI**. A single GET on `tpex_international_bond_issue_org` returned, on 31 July 2026, **931 outstanding international bonds, every one with an ISIN**, each carrying issuer, issue and maturity date, currency, size, coupon, and the non-call period × call frequency. Of those, **856 are USD, totalling USD 216.7bn, of which 698 are callable for USD 164.7bn**; the modal structure is **NC5 called annually (324 of 698) and 75.5% of the callables have a tenor of 26 years or more**. That is the Formosa register, fifteen times larger, machine-readable, in one request — and it is a same-day snapshot with no date parameter, so the history is unrecoverable once missed. **Start a daily cron on it now.**

Part 1 — How the shortlist was built

Four tests, applied to each family. They are not weighted equally: a family that fails the data test is unworkable regardless of how interesting it is.

1. Residual risk. Not "is this complicated" but "what is left after a competent desk has hedged everything it can". The families that score highest are the ones where the residual is *structurally* untradeable — a correlation with no instrument, a fixing you cannot transact at, a basis the rate-setting body has explicitly told dealers not to lay off — rather than merely illiquid or expensive.

2. Footprint. A dated, documented episode where the hedging flow moved a market, ideally written up by an official-sector source. This is the test that separates a market note from a modelling paper. The mortgage-convexity literature is the benchmark: Perli and Sack at the Federal Reserve Board put the amplification of rate moves from MBS hedging at **16-30% over January to May 2003** — a measured episode, not a structural constant, and stated in their conclusions rather than as a headline result — the BIS wrote the episode up in September 2003, and the ECB cited both. Nothing in the structured-note complex has that standard of evidence, but some families come closer than others.

3. Data. Can the register be built from free, machine-readable, deal-level sources? This decides whether the note has a table of real ISINs or a paragraph of assertion. The four existing notes all live or die on this — TEC10 on a hand-built 325-row fixing calendar, Formosa on 60 bonds from offering documents.

4. Adjacency. How much of the existing four notes transfers. A dive that reuses the CMT convexity machinery, or the SPIRE repack plumbing, or the Taiwan regulatory chain, is worth materially more than its standalone score.

A note on what "high residual" means in practice. Three distinct things get called unhedgeable and they are not equally bad:

	Example	Can you price it?	Can you trade it?
Illiquid but real	5y→35y Bermudan vega	Yes, with a model	Yes, at a cost
Model-marked	Rate-FX correlation in a quanto; inflation vol	Only by convention	No
Not a market risk at all	Behavioural deposit floor; jurisdictional cancellation of a contractual floor	No	No

The third category is where the genuinely interesting work is, and it is the category the Zero Floor note found in EUR loans. It is also the category least served by a vol-surface dive.

Part 2 — The shortlist

Tiers are recommendations, not scores. Within a tier the order is the order I would do them.

#	Family	The residual you keep	Residual	Footprint	Data	Adjacency
A1	ICE Swap Rate fixing risk	Coupon settles on Level 2/3 indicative quotes; hedge settles on Level 1	High	Medium	High	Very high
A2	Callable range accruals & multi-callables	One-directional Bermudan vanna; street on the same side	High	High	High	Very high
A3	UK LPI and capped-and-floored inflation	Path-dependent collar delta-hedged with a linear asset	Very high	High	High	Medium
A4	Cross-currency basis in repacks	30 years of basis roll nobody prices at inception	High	High	Medium	Very high
B1	Leveraged CMS spread steepeners	Spread-option correlation with no quote; leverage multiplies hedge notional	High	Medium	High	High
B2	Term SOFR and the LIBOR-ISR fallback gap	Term/OIS basis dealers are <i>told</i> to warehouse; note/swap fallback asymmetry	High	Medium	Medium	High
B3	PRDC and long-dated FX vol	20–30y FX skew and FX-rates correlation	Very high	High	Low	Medium
B4	US CMT index construction	Bid-side indicative inputs, on-the-run bias, methodology change at Treasury's discretion	Medium	Medium	High	Very high
B5	The behavioural zero floor	A floor with no contract, no strike and no counterparty	High	High	Medium	Very high
C1	Inflation YoY caps and floors	No executable inflation vol market exists — by the administrator's own account	High	Medium	Medium	Medium

#	Family	The residual you keep	Residual	Footprint	Data	Adjacency
C2	TARNs, TARFs and accumulators	Path-integral knockout; market and credit risk perfectly correlated	High	High	Low	Low
C3	Callable snowballs and ratchets	Nested Monte Carlo; Greeks are computed, not observed	Very high	None	None	Medium
C4	Quanto CMS and quanto structures	Rate-FX correlation is a reserve policy, not a hedge	Very high	None	None	Medium
C5	Livret A / LDDS	max(€STR, inflation) with a hard floor and a discretionary override	High	Medium	High	Medium

Part 3 — Tier A

A1. The ICE Swap Rate fixing — the cheapest dive, and it is under every CMS position already on the book

What it is. ICE Swap Rate, formerly ISDAFIX, is the mid-price of the fixed leg, computed from a two-minute window split into 24 five-second blocks with a random snapshot in each. A synthetic order book is built at each snapshot, a volume-weighted bid and offer are computed by filling a hypothetical trade of Standard Market Size on each side, snapshots with insufficient volume or crossed quotes are dropped, VWAMPs outside the interquartile range are discarded, and if at least six snapshots survive they are quality-weighted by bid-offer and averaged.

The finding is in footnote 1 of the administrator's own methodology. IBA states that "USD SOFR ICE Swap Rate settings, USD SOFR ICE Swap Rate Swap Spread settings, EUR €STR ICE Swap, EUR and GBP Inflation Swap Rate settings are expected to be calculated using input data at **Level 2 or Level 3** of the Waterfall", and that it "expects to use Level 1 input data... when this is available in the future". **Level 1 is "executable prices and volumes" from central limit order books. Level 2 is "eligible dealer to client prices and volumes displayed electronically by trading venues". Level 3 is movement interpolation off the previous publication day.** The distinction is executable versus displayed-but-not-executable, and it is worth stating precisely: Level 2 is not a dealer poll and not an indicative quote in the ISDAFIX sense — it is real venue data, one rung below the executable prices the benchmark is built around.

Why that is a hedging problem rather than a documentation curiosity. The coupon on a CMS note, the exercise value of a cash-settled swaption, and the close-out on an early termination all settle against this number. The instrument you hedge with — a cleared SOFR or €STR swap — is priced off the executable market. The mismatch is not transitional; it is structural for as long as the post-LIBOR swap market does not generate enough central-limit-order-book volume to reach Level 1, which by IBA's own account it does not.

Three mechanics for a monitoring dashboard. Level 3 only applies where the rate was computed at Level 1 or 2 on the previous publication day *and* the tenors one year either side were themselves non-interpolated on both days — so it fails exactly when you need it. If Level 3 fails, the Insufficient Data Policy applies and there is a **No Publication**, which is a contractual event in the note, not a market event. And IBA publishes a free quarterly **ICE Swap Rate Transparency Report** giving, per setting, which waterfall level was used — in **10-percentage-point bands rather than counts**, which is coarser than the methodology's own wording implies but is still an ongoing quality metric for the fixing the book settles on, and as far as I can tell nobody tracks it.

Change risk sits with a regulator, not a market. ICE Swap Rate is a critical benchmark under the UK BMR and a significant benchmark under the EU BMR. IBA's own disclaimer states it "may be required, pursuant to applicable laws or an order of a regulatory or other competent authority, to make changes and/or continue to publish the affected benchmark" — and unrepresentativeness is explicitly not grounds to stop. The methodology of your CMS fixing is amendable by order.

The footprint is the manipulation record, and it is the reason the waterfall exists. CFTC orders for attempted manipulation and false reporting of USD ISDAFIX: Citibank **\$250m** (2016), RBS **\$85m** (2017), Deutsche Bank Securities **\$70m** (2018), Bank of America **\$30m** (2018) — and these are four of at least nine actions, the CFTC describing the Bank of America order as "the ninth CFTC enforcement action" on USD ISDAFIX. Others include Barclays \$115m, Goldman \$120m and ICAP Capital Markets \$50m. Conduct roughly 2007–2012. The incentive existed precisely because cash-settled swaptions and CMS notes settle on it.

A separate conduct precedent attaches to the day-one mark. The German BGH judgment of 22 March 2011 (XI ZR 33/10) on a CMS Spread Ladder Swap imposed a duty to disclose the product's initial negative market value as evidence of a conflict of interest. For a CMS-spread payoff, that day-one number is materially a function of the smile-extrapolation choice — which is the same unhedgeable object as the convexity adjustment. Conduct risk and model risk are the same risk here.

Data route. IBA methodology and benchmark statement (PDF, free). Quarterly waterfall-level usage from the ICE Swap Rate page. Cross-check against DTCC SDR trade counts by tenor to see whether venue volume plausibly supports Level 1. Nomura's 2023 volatility outlook, which is public, flags a related live basis: CMS delta accrues as ICE-fix-referenced received delta against paid LCH delta, matched only by convention, and a migration of the fixing methodology toward a blended LCH/Eurex basis "would give dealers an overnight headache". **I have not verified that claim against IBA or LCH — it is a lead, not a finding.**

What the dive produces. A short note — two weeks, not two months — plus a standing monitor. The deliverable is a table of every ICE Swap Rate setting the book touches, which waterfall level it actually runs at, how often, and what the note's contractual fallback does on a No Publication.

A2. Callable range accruals and multi-callables — the payoff behind the Formosa finding

The payoff. A fixed above-market coupon multiplied by the fraction of days the reference index fixes inside a range, callable by the issuer. Or, stripped of the range, simply a fixed or step-up coupon the issuer may retire at par on many dates. Either way the investor is short a strip of digitals, a Bermudan, or both.

Why this rather than Formosa again. The Formosa note established the channel through one wrapper and one buyer base. The vega supply is a property of the *payoff*, not the wrapper, and the second episode has nothing to do with Taiwan. In March–June 2023, US Federal Home Loan Banks issued more than **\$75bn** of very short-dated multi-callables — typically 1y-NC-3m-plus-1m at 5–5.25% — taken up largely by money market funds. Nomura's read is that this "seemed to double up on the skew position on the street", pushing the USD 6m1y ± 100 risk reversal to all-time lows, above 50bpv payer-minus-receiver. Two episodes, two continents, two buyer bases, one payoff.

The residual, stated precisely. The investor is short the option to cancel, and therefore short swaption volatility; the dealer is long the Bermudan and short the vanna. The consequence is directional and identical across the street: dealers "get longer vega in a sell-

off and shorter vega in a rally", so a rally in 30y swap rates mechanically creates vega demand for vanna hedging on dealer Bermudan books. That is wrong-way — the street is forced to buy vol into a rally, when vol is already bid — and there is no natural counterparty, because every structured-product seller holds the same position.

The range accrual adds a second, separate problem. Banks replicate the daily digitals with floorlet spreads of 5–10bp spacing rather than true digitals, which leaves "vega risks to very high strike options, risks that are very difficult to hedge" (Hagan). Hagan also names the modelling dead end: a model rich enough to fit both the floorlet smiles and the diagonal swaption vols "would be complex, calibration would be difficult, and the procedure would likely yield unstable hedges". The desk has already seen the US version of this payoff — the `04_us_cmt/` folder holds seven filings, four of them range accruals — so the machinery transfers directly.

Model risk on the Bermudan itself is genuinely unsettled, and the literature contradicts itself. Longstaff, Santa-Clara and Schwartz argue single-factor models significantly undervalue American-style swaptions and, more to the point, "can lead to severely biased estimates of hedge ratios and risk-management statistics". Andersen and Andreasen, in the same issue of the same journal, conclude Bermudan prices "depend only weakly on the number of factors" and attribute the LSS result to calibration. Twenty-five years on it is still an active topic. A note that resolves nothing but *documents* the dispersion — same trade, three calibrations, three hedge ratios — is worth more to a risk manager than another pricing paper.

Data route. SEC EDGAR full-text search over 424B2 and FWP is the register; the method is replicable and SLCG has already demonstrated it on the CMS subset. For the FHLB burst, the FHLBanks Office of Finance combined financial report gives consolidated obligations but **not a callable-only issuance series** — that is a gap, and the \$75bn figure is a dealer estimate. DTCC SDR gives the swaption tape.

What the dive produces. A register of US callable and range-accrual issuance by year and structure, the vega-per-billion arithmetic done the way the Formosa note did it, and the two episodes dated and quantified. The honest ceiling: the vega numbers will be your arithmetic, not a measurement, exactly as in the Formosa note — no public source quantifies this flow in any market.

A3. UK LPI and capped-and-floored inflation — the largest unhedgeable residual in the complex

The payoff. LPI(0,5) or LPI(0,2.5): the annual pension increase equals RPI or CPI floored at 0% and capped at 5% or 2.5%, applied to a compounding base. Statutorily — for **pensions in payment**, under Pensions Act 1995 s.51 — 5% for rights accrued 6 April 1997 to 5 April 2005 and 2.5% thereafter. (The equivalent cap split for revaluation *in deferment* changed on 6 April 2009, not 2005; conflating the two is the standard error and any register built for this dive should keep them apart.) Because the cap applies annually to a compounding base, **LPI is path-dependent — it is not a strip of independent collars**, and that is the whole difficulty.

Who is exposed and how much. UK defined-benefit schemes. Per the PPF Purple Book, 68% of schemes provide CPI- or RPI-linked increases for post-1997 benefits, for the vast

majority capped at 5%; only about 9% provide uncapped indexation. The dealer writing the LPI swap is short a strip of RPI caps at 5% (or 2.5%) and long a strip of floors at 0%, in a market with no traded vol surface at either strike, delta-hedged with a linear inflation asset.

Both tails have been realised, inside seven years. RPI peaked at **14.2% in October 2022** — the highest since 1980, though not an all-time high; RPI reached 26.9% in August 1975. At that print an LPI(0,5) liability increases by 5% while a linear RPI hedge delivers 14.2% — the scheme is massively over-hedged and whoever sold the cap eats the difference, with no cap market to buy back into. The 0% floor bound in the 2015–16 disinflation. The delta of LPI is not constant and needs rebalancing, and the hedging instrument itself has thinned: the UK LPI swap market has become increasingly illiquid, and using LPI swap prices to set the delta is described by practitioners as problematic because the quoted market implies unrealistic inflation behaviour.

The flow is one-way and the supply is shrinking. UK index-linked gilts stood at roughly **£619.0bn** in nominal uplifted terms at end-2024, **24.5%** of the government's debt portfolio — a share higher than any G7 peer and about twice the second-highest, which the DMO attributes to structural demand from the domestic pension sector. Government policy is to reduce linker issuance. The Bank of England has separately noted that scheme demand for linkers may have become relatively price-inelastic. Structurally short, into shrinking supply, against a buyer who does not care about price.

And the index itself is being redefined. RPI aligns with CPIH methodology from **February 2030**, and the Treasury's response is explicit that "the government will not offer compensation to the holders of index-linked gilts". Judicial review brought by five trustee claimants — the BT and Marks & Spencer schemes and three Ford vehicles — was dismissed by Holgate J on 1 September 2022 (*R (BT Pension Scheme Trustees Ltd & Ors) v UKSA and the Chancellor* [2022] EWHC 2265 (Admin)). Every RPI-linked structure crossing February 2030 has a step in its index definition that no derivative references. This is the same class of risk as the euro HICP rebasing of February 2026 — where AFT published the chaining rule and the sovereign bond is immune because it is defined as a ratio, but **a structured note that hard-codes an index level, a base year or an absolute threshold is not immune**. That is a documentation audit worth running now, independently of whether this dive happens.

Footprint. LDI, September–October 2022: the Bank bought **£19.3bn** of gilts between 28 September and 14 October, **£7.2bn of them index-linked**, and had to widen the operation to include linkers in a *separate, later* market notice on 11 October. The sequencing is the point — the inflation-linked leg needed its own intervention.

Data route. DMO index-linked gilt data and cash flows (CSV export endpoints, free), PPF Purple Book for the indexation-type denominator, ONS for the RPI series, BoE market notices and Quarterly Bulletin for the LDI episode, UKSA for the 2030 decision. **No public figure exists for the LPI swap market or the inflation cap/floor market** — say so rather than sourcing one from a consultancy deck.

What the dive produces. The cleanest "collar hedged with a linear instrument" case study available anywhere, with both tails realised and a statutory index change dated five and a half years out. Adjacency is medium — it does not reuse the CMT or repack machinery — which is the only reason it is not first.

A4. Cross-currency basis in repacks — the SPIRE risk nobody prices at inception

Not a payoff family; the funding leg every cross-currency structure rides on. The SPIRE levered-deals ledger already shows the shape: EUR notes on USD collateral (US Treasury, Oracle, AB InBev), EUR notes on JPY collateral (Government of Japan), JPY notes on USD collateral (Heineken) — all clustering at $1.02\text{--}1.08\times$ face once FX-converted, which the ledger correctly reads as the noise of the FX assumption rather than gearing. What the ledger does not price is that each of those compartments carries the basis to final maturity.

The mechanism is documented and directional. An issuer raises where spreads are tight and swaps to where it needs the currency; that flow is one-directional and moves the basis. BIS Working Paper 590 (Borio, Iqbal, McCauley, McGuire and Sushko, *The failure of covered interest parity*): the implied cross-currency position covering EUR-denominated reverse-Yankee debt of US non-financial corporates "has more than doubled since 2014, reaching approximately minus \$400 billion". The Australian case runs the other way — the RBA documents a persistent *positive* basis Kangaroo issuers receive as a premium, because Australian financials issue offshore and swap back. A surge in Samurai issuance pushed the 3-month USD/JPY basis to **−53.5bp** in October 2024.

The Taiwan leg is the one where the arbitrage is legally blocked, not merely expensive. TWD cross-currency basis has traded deeply negative almost continuously since the early 2000s, imposing on average an extra **100bp onshore and 200bp offshore** at the 3-month tenor on insurers hedging FX risk — and foreign arbitrage cannot compress it because an offshore investor's total holding of Taiwanese government bonds, corporate bonds, financial bonds, money market instruments and money market funds "must not exceed 30 percent of the net inward remittance". That is the concrete answer to why proxy hedging breaks: the arbitrage leg is illegal.

Why it is hard. The basis is a balance-sheet price, not a market risk. It does not mean-revert on a tradeable schedule, and the offsetting flow depends on someone else's issuance decision in another currency — or on a regulator changing an accounting rule in Taipei. There is no instrument that hedges that. A repack struck at a 30-year horizon carries 30 years of basis roll, and it is not on anyone's risk report at inception.

The call and the basis interact, which is the specific repack problem. Calling the note terminates the cross-currency swap, so the economic value of the call depends on where the basis is on the call date — a variable with no forward market beyond a few years. And early termination is a gap event, not a market move: rating criteria for CLN and repackaging vehicles are explicit that swap counterparty default "would typically lead to automatic termination at market value (even though underlying collateral may be performing)". The residual is the difference between the collateral's realisable price and the swap's mark, in a stress where both are moving and the collateral may be wrong-way to the swap.

Data route. The SPIRE tracker already has the compartments; this is an extension of an existing dataset, not a new one. BIS covered-interest-parity literature and the BIS Data Portal for the basis; RBA for the Kangaroo case; the CFR analysis for Taiwan. **No public aggregate for repack issuance exists** — the SPIRE tracker's ~1,900 notes is a census of one platform, not a market figure, and the note should say so.

What the dive produces. The most directly actionable of the four: a basis-roll exposure line for the existing SPIRE book, by currency pair and maturity, plus the call/basis interaction on the callable compartments. It is also the only Tier A candidate where the desk owns the primary data already.

Part 4 — Tier B

B1. Leveraged CMS spread steepeners

Payoff. After a teaser fixed coupon, $\text{Multiplier} \times (\text{CMS}_{\text{long}} - \text{CMS}_{\text{short}} - \text{spread})$, floored at 0% and usually capped, callable. A documented example: Barclays Callable CMS Steepener Notes due 2031 — 11% in year one, then $4.50 \times (\text{CMS}_{30y} - \text{CMS}_{2y} - 0.50\%)$, floor 0%, cap 11%, callable semi-annually from year one, 5.00% dealer commission.

The residual. CMS convexity on both legs, spread-option correlation and its smile, the cap and floor as spread options away from the money, and a Bermudan the dealer is long. The correlation is the unhedgeable piece: the market prices the two CMS marginals from their respective swaption smiles and couples them with a copula whose parameter is not quoted. A bivariate SABR treatment needs six correlations, several of which — the correlation between one CMS rate and the *volatility* of the other — are not observable at all.

Why leverage makes it worse mechanically. Hedge notional scales linearly with the multiplier, and multipliers went from 10–15× in 2005 to 50×+ by 2007 with client requests above 100×. With a 40–50× multiplier a very small fall in the spread drops the coupon to zero and a small rise pins it at the cap, so the dealer's delta to the spread is a narrow spike — and the issuer's rational call is triggered in exactly the regime where the hedge is largest.

Footprint. The 2005–07 flattening. USD 30s10s CMS was around 21bp in April 2007 against 60bp at end-2005, and \$2–3bn of highly leveraged CMS spread structures were sold in Asia in the nine months to May 2007. The curve facts are confirmed independently — BIS records 10y–2y UST turning negative in late 2005. **What does not exist is an official-sector document joining the product to a swaption-market move.** I looked; BIS and ECB attribute the 2003–05 US–euro swaption vol gap to *mortgage* hedging, not to structured notes. If this dive needs an official-sector hook it may have to be built from 424B2 filings plus swaption data rather than cited.

The ancestor is worth a paragraph in whatever gets written. Procter & Gamble's 5s/30s leveraged swap with Bankers Trust produced a \$152m pre-tax charge in April 1994 and the leading US case law on dealer duties in leveraged rate structures. 5s/30s is a leveraged CMS spread. This is the same trade thirty years apart, and the 1994 episode is the one with a full legal record.

Data. SEC EDGAR. SLCG counted **651 CMS-linked notes, at least \$4.95bn, 2015–2020**, from a full-text search of 424B2 filings — a replicable method with a stated denominator. My own calibration query on "constant maturity swap" in 424B2 for Jan-2025 to Jun-2026 returned **70 filings**, led by JPMorgan Chase Financial (39) and Goldman (12 + 12). Rate-linked notes are a thin tail of a mostly equity-linked market: plan for hundreds, not thousands.

B2. Term SOFR and the LIBOR-ISR fallback gap

Two related things, both live, both on any legacy structured book.

First, term SOFR is an unhedgeable basis by design. The ARRC "does not recommend that the dealer seek to hedge its own resulting SOFR Term Rate exposure with an additional SOFR Term Rate derivative"; dealers should "warehouse the risk... including through the use

of overnight SOFR derivatives", and interdealer trading of term SOFR, including term SOFR basis swaps, remains outside the recommendations. **The dealer holds the term/OIS basis because the rate-setting body wants it to.** That is the cleanest example in this entire survey of a residual risk that exists as a policy choice rather than a market failure.

Second, the ISR fallback gap, which is a repack problem specifically. The ARRC records "a small number of outstanding debt instruments tied to the USD LIBOR ISR with more problematic fallbacks, for example debt instruments which would fall back to the prior period fixing if the dealer poll(s) were to fail". Unpack that on a book: a CMS floater that freezes at its last fixing **stops being a floater**. Its DV01 goes from roughly zero to full duration overnight and its convexity position vanishes, while the hedging swap — if it adhered to ISDA Supplement 88 or the protocol — fell back to SOFR ISR + **26.161bp**. Note and hedge now carry different rates, different sensitivities and different fallback logic. And the legislative backstop does not reach it: neither federal nor state legislation covered transactions tied to the USD LIBOR ISR as at June 2022. **Whether the final Regulation ZZ rule of December 2022 changed that position is unverified and would be the first thing to check.**

Why it matters for an SPV. In a repack the note and the swap may adhere to different fallback regimes at different times, and the calculation agent is left exercising discretion — with the ARRC's guidance being that calculation agents "consider" its suggested formula. A recommendation, not a rule. This is a documentation dive with a real P&L consequence, and it maps directly onto SPIRE compartments.

The synthetic backstop is gone. All sterling LIBOR settings ceased permanently after 28 March 2024; the final synthetic USD LIBOR settings were published for the last time on 30 September 2024. Anything still relying on synthetic rates to avoid triggering a fallback has now triggered.

B3. PRDC and long-dated FX vol

The most extreme residual in the survey, and the worst data. A yen note paying a geared coupon on AUD/JPY or USD/JPY, floored at zero, Bermudan-callable. The dealer ends up long-dated FX vol and skew out to 20-30 years, FX-rates correlation in two currencies, rates vol in two currencies, and a call whose exercise boundary moves with spot — so the delta hedge has to be re-derived, not re-scaled, when spot moves.

The observable signature is that the dealer book makes the market it hedges in. Piterbarg's finding, via Risk: FX options "exhibit a significant volatility skew induced in longer-dated options by dealers trying to hedge their PRDC swap positions", and the structure is particularly sensitive to that skew. When the whole street is short the same 25-year AUD/JPY skew, unwinding is a self-inflicted move.

Footprint. The 2007-08 carry reversal. BIS measured AUD/JPY returning 12.5% annualised from January 2001 to September 2007 with a negatively skewed distribution. The delta-hedge unwind mechanics are visible in the option market: on roughly \$7bn of ¥90-strike puts, "the rapid unwind of the large delta hedges caused a sharp intra-day drop in USD/JPY of 2.9%".

Why it is Tier B despite the residual. No public size figure exists. PRDC notional is not reported by BIS, JSDA or the JFSA. There is no register to build. The nearest live dataset is

JSDA's quarterly complex structured bond (*shikumisai*) file, which is machine-readable and splits by underlying — Q1 2026: **¥696.5bn total across 3,358 issues, of which FX-linked ¥86.9bn across 354 issues**, and notably **every FX-linked issue in the eleven-quarter series was privately placed**, with 51% of FX-linked redemptions in Q1 2026 being early redemptions. That is the best live Japanese structured-note data there is, but it is quarterly aggregates by category, not deals. A PRDC dive would be a mechanism-and-episode note with no register.

B4. US CMT index construction — the extension of a note that is already written

The desk owns the CMT payoff note. This is the index. CMT is not a traded rate; it is a par yield read off a fitted curve. Since 6 December 2021 Treasury bootstraps instantaneous forwards so input instruments price without error, then applies monotone convex interpolation. The inputs are **indicative, bid-side price quotations — not transactions** — for the **most recently auctioned** securities, collected by the New York Fed at or near 3:30 PM.

Four separable residuals, worth splitting on a risk report. CMT-versus-CMS basis, which is the swap spread at the relevant tenor. On-the-run richness, inherited because every input is the most recently auctioned security and dropped discretely at each auction cycle. Bid-side bias, roughly half a bid-offer above a mid par yield. And a convexity adjustment computed on a *bond* par yield, where the annuity and the measure change differ from standard CMS.

The concrete failure mode: you cannot buy the fixing, and Treasury can change it. The 3:30 PM indicative snapshot can move for reasons that never print in a tradable market, and Treasury reserves the right to change inputs or methodology "in its sole discretion". That is not theoretical — before the 20-year bond returned in May 2020, Treasury used composite off-the-run proxies in the 20-year range, and at various times interpolated yields and rolled-down securities.

The 2021 methodology change is the dated episode, and the numbers are the interesting part. Average nominal CMT differences between the new and old methods over October 2020 to September 2021 ran **−0.1 to +0.5bp with standard deviations of 0.2 to 1.0bp**. For the **real** curve the averages were **−2.7 to −0.6bp with standard deviations up to 4.1bp**. A 4bp standard deviation on the real curve is a real P&L number on anything TIPS-linked. Two consequences the existing note may not carry: pre-change rates "remain official", so any long contract referencing CMT history references a **spliced index**; and any Treasury-determined rate averaging across December 2021 consists of rates produced by both methods.

Data. Treasury methodology page, the change information sheet, and the published MC-versus-HS comparison CSV. Daily CMT via the Treasury CSV endpoint. Small, cheap, and it upgrades an existing note rather than starting a new one.

B5. The behavioural zero floor

The Zero Floor note settled the contractual floor. The unhedgeable one is not in any contract. Retail deposit rates are truncated at zero because banks are unwilling to charge households negative rates; the ECB documents the truncation directly in the distribution of large euro-area banks' household deposit rates. There is no counterparty to buy that floor from and no strike to hedge at — and regulators nevertheless require it to be modelled.

The regulator's own diagnosis is the sharpest statement of the problem. In the ECB SSM 2017 sensitivity analysis of IRRBB, the finding was that "banks heavily rely on models of customer behaviour which were calibrated in a declining interest rate environment". The floor's delta is a model output, so when the regime turns the model changes and the hedge ratio jumps discontinuously with no market to trade into. The framework that forces the calculation is the EBA IRRBB and CSRBB guidelines (IRRBB from 30 June 2023) on top of the Basel IRRBB standard, which requires embedded caps and floors to be captured as automatic interest rate options.

A third loss channel is legal, and two of the three cannot be hedged at all. In Spain the floor the bank thought it owned was judicially cancelled retroactively; in Portugal it was legislated away prospectively in 2018; in the UK the PRA treated a zero or negative Bank Rate as an operational-readiness problem and gave firms a timeframe to build tactical solutions. **The residual on a zero floor is not principally a vol risk. It is jurisdictional legal risk on whether the floor survives, plus model risk on a floor you never contracted for.** That is the argument for a legal-and-model dive rather than another vol dive, and it is a natural sequel to the existing note.

Model-choice risk is the third piece and it is quiet. Once the underlying can go negative, Black vol ceases to exist at low strikes and the market quotes normal or shifted-lognormal. **The shift parameter is a convention, not a market price:** it changes the low-strike wing, and therefore the delta and vega bucketing of a 0% floor, without changing any quoted price.

Part 5 — Tier C, and why each is parked

C1. Inflation YoY caps and floors. The residual is real — inflation smile, zero-coupon-to-year-on-year convexity, seasonality decomposition, all model-marked. The evidence for "no liquid inflation vol market" is now primary rather than anecdotal: IBA is **only in March 2026 consulting on launching inflation swap benchmarks at all**, initially at Level 2 for EUR HICPxT and GBP RPI, with Level 1 added "when sufficient volumes are available" and a USD CPI benchmark deferred until volumes exist. Parked only because it folds naturally into A3 rather than standing alone.

C2. TARNs, TARFs and accumulators. The knockout is a path integral, so delta and vega step rather than glide, and with typical 2× gearing the dealer's notional doubles precisely when the client's ability to pay is impaired — market risk and credit risk correlated by construction. Two well-documented blow-ups: Taiwan TRF in 2014–16, where industry client losses were estimated in the hundreds of millions to billions and the FSC fined multiple banks; and Korea KIKO in 2008. **The Taiwan link is worth noting — the same regulator that supervises the Formosa buyer base is the one that policed the TRF market.** Parked because in *rates* the buyer base is thin and there is no register; the real market is FX corporate hedging, which is a different desk.

C3. Callable snowballs and ratchets. $\text{Coupon}(n) = \text{Coupon}(n-1) + K - L \times \text{Index}(n)$, floored at zero, callable. Piterbarg is explicit that these have no analytical underlying, must run in Monte Carlo in a forward-Libor model, and that computing the exercise value on a given date and path requires a **nested** simulation. Greeks are computed, not observed, and they move with the calibration. This is the highest residual in the survey and it has **no public market footprint and no public size at all** — the deliverable would be a modelling and P&L-attribution study, not a market note. Also a naming hazard: "snowball" now overwhelmingly means the Chinese equity autocallable, which will contaminate any search.

C4. Quanto CMS and quanto structures. The quanto drift adjustment is linear in rate-FX correlation, an object with no instrument, scaling with both vols and with time to expiry — so a 30-year quanto carries a correlation exposure an order of magnitude larger than a 2-year one, in the tenor where the historical estimate is least reliable. Quanto books are managed by reserve policy and concentration limits, not by hedging, and two dealers can mark the same trade materially differently without either being wrong. Parked because **there is no episode and no size**: quanto losses show up inside model-valuation-adjustment lines rather than as named events.

C5. Livret A and LDDS. The French oddity, and the one nobody writes about. The rate is set by formula under the *arrêté* of 27 January 2021 as the higher of (a) the arithmetic average of the semi-annual average €STR and French inflation excluding tobacco, and (b) a hard **0.5% floor**, computed by the Banque de France twice a year with a discretionary revision possible on 1 May or 1 November. It is simultaneously an averaged constant-maturity-style fixing, a **max() of a rate and an inflation index — a best-of, not a sum**, an absolute floor, and an explicit discretionary override, on a liability of very large scale sitting with Caisse des Dépôts and the banks. **I found no published analysis of how CDC hedges it.** That is either a gap in the literature or a gap in my search; it is the most genuinely novel item on this list and the least certain.

Part 6 — What has changed in the data since SPIRE was built

The SPIRE pipeline was built on ESMA FIRDS, OpenFIGI and the Euronext Dublin S3 bucket, with the documented pain point that the bucket refuses `ListBucket` so the index has to come from a JavaScript-rendered directory. Four things are now available that were not used, all free and unauthenticated, all verified in this session.

The ESMA prospectus register has an open Solr backend, and it is the European EDGAR. `registers.esma.europa.eu/solr/esma_registers_priii_documents/select` is queryable directly, no auth. Filtering out nested children it holds **45,474,087 parent documents** — final terms, base prospectuses, supplements — indexed by `instrument_isins`, `party_name` with LEI, `document_type`, `approval_filing_date`, `home_member_state_code` and `passporting`. Facets, wildcards and range queries all work. Home member state distribution is DE 43.2M, LV 1.40M, LU 400,904, NL 137,846, **IE 115,026. The limitation is decisive and should be stated up front: it is metadata only.** `document_rfss_id` does not resolve to a file; you get the national document ID and must go to the national register for the PDF — and every national register tried (AMF BDIF, CSSF, BaFin, Central Bank of Ireland) was dead or unreachable on the paths tested. So this replaces the *discovery* half of the Euronext problem, not the *retrieval* half.

SEC EDGAR full-text search has a JSON API and it is the best structured-note corpus in existence. `efits.sec.gov/LATEST/search-index` takes `q` with exact phrases, `forms`, and a custom date range, and returns 100 hits per page with the full `_source` **plus an issuer league table in the aggregations of the very first call.** Coverage starts 2001-01-01. Form types that matter: **424B2** is the priced pricing supplement, **FWP** the preliminary term sheet; 424B3 and 424B5 are minor and 424B4/424B7 are not used. The constraint is that `from + size` must be $\leq 10,000$, so bucket by month rather than paging. Deep links are constructible: `/Archives/edgar/data/{cik}/{adsh_no_dashes}/{filename}`.

TPEX publishes an OpenAPI, which closes a gap the Formosa note recorded as unclosable. Covered in the executive summary; the swagger is at `/openapi/swagger.json`, 225 paths, 21 bond-tagged. Beyond the issue register, TPEX publishes its **own daily theoretical prices** for USD fixed-rate callables and for USD zero-coupon callables including an implied rate — genuinely unusual, and directly relevant to the extension argument in the Formosa note. Send a User-Agent or Cloudflare returns 520. **Every endpoint is a same-day snapshot with no date parameter.**

DTCC's swap data repository is the only free source of actual swaption trades. The daily cumulative CFTC rates file carries strike, option premium, first exercise date, expiration and notional. On 30 July 2026 it held 21,423 rows, of which 2,961 carried an option premium. Caveats that matter: the legacy `Product name` / `Option Type` columns are entirely empty, so identification is UPI-only and you need ANNA DSB reference data to resolve UPI to underlying tenor; notionals are capped and rounded; history is a rolling two years.

And one that is worth ten minutes. `bis.org/api-search/search.json` is an undocumented full-text search over bis.org that **indexes PDF interiors**, and `bis.org/speeches/speeches.zip` is 129MB of every central bank speech with the complete text in a CSV column. Greppable offline, which is the cheapest official-sector search anyone will ever build.

What is definitively not available free, so budget for it or design around it. No swaption vol surface — SRVIX ends 11 February 2022, MOVE is proprietary, ICE Swap Rate is licence-only by explicit statement, CME settlement files are IP-blocked. SIFMA's downloads are now email-capture forms. ISDA SwapsInfo has no API. The CFTC Weekly Swaps Report was discontinued in March 2018. The ECB's CSDB micro data is ESCB-internal — only the derived SEC and SHS aggregates are public, and those stop at sector × country, so they cannot size a note family. And no rating agency criteria document could be fetched from any of S&P, Moody's, Fitch, Scope or DBRS; only KBRA server-renders, and the Fitch research sitemap is the one free rating-action feed that works from a script.

Part 7 — What is out of scope, and why

Equity autocallables are the other big structural dealer-vega story and they are deliberately excluded. The Korean ELS complex, the EuroStoxx and Nikkei autocall books, worst-of baskets, cliquets and the 2024 HSCEI losses are all genuinely in scope for the *question* — they are the largest single source of one-directional dealer vol positioning in any asset class. They are excluded here because the scope chosen for this pipeline is rates and cross-currency, because the hedging machinery does not transfer from the existing four notes, and because the data route is different (KSD via `data.go.kr` with a key, HKEX's daily quotation file, and the SEC corpus where the equity-linked notes vastly outnumber the rate-linked ones). If the scope changes, the natural entry point is Korean ELS: KSD exposes ISIN master **and autocall condition detail** through a free REST API, which is a better dataset than anything available on the rates side.

Credit is excluded for a different reason. CLNs, first-to-default baskets and bespoke tranches sit closer to the SPIRE wrapper than to the payoff question, and the repack-specific credit risk — early termination at market value on swap counterparty default, wrong-way collateral — is already folded into A4, where it belongs.

Sources

Everything in this note is sourced in `01_register/candidates.csv`, one row per family, and in `01_register/data_sources.md`, which carries the endpoints. The load-bearing primary documents are the ICE Swap Rate Calculation Waterfall Methodology and the March 2026 inflation swap benchmark consultation; the ARRC recommendations on USD LIBOR ICE Swap Rate contracts and on term SOFR scope of use; the US Treasury yield curve methodology page and its 2021 change information sheet; the PPF Purple Book and the DMO Debt Management Report; the UK Statistics Authority response on RPI reform; ECB working papers 2173 and 2283 and the SSM 2017 IRRBB sensitivity analysis; the BIS covered-interest-parity literature; and the CFTC ISDAFIX orders. Practitioner material — the Nomura 2023 volatility outlook, Hagan on range notes and convexity, Piterbarg on callable exotics, TARNs and multi-currency skew, and the Longstaff-Santa-Clara-Schwartz versus Andersen-Andreasen dispute — is cited as practitioner material, not as official-sector evidence.

The four endpoints quoted with figures were fetched in this session on 31 July 2026 and the numbers computed from what came back:

```
curl https://www.tpex.org.tw/openapi/v1/tpex_international_bond_issue_org
-> 931 records, all with ISIN; 856 USD / USD 216.7bn; 698 callable / USD 164.7bn
curl "https://registers.esma.europa.eu/solr/esma_registers_priii_documents/select?q=-type_s:child&rows=0&start=0"
-> numFound 45,474,087
curl "https://efts.sec.gov/LATEST/search-index?q=%22constant+maturity+swap%22&forms=424B2&startdt=2025-01-01"
-> 70 hits; JPMorgan Chase Financial 39, JPM 35, GS 12, GS Finance 12, Barclays 6
```

Caveats

- **Fourteen load-bearing claims were fact-checked against primary sources before publication; eleven verified clean and three needed correction.** The corrections are

in the text: the ICE Swap Rate Level 1/Level 2 distinction is *executable order-book prices versus dealer-to-client prices displayed by venues*, not "executable versus indicative", and the quarterly transparency report gives 10-percentage-point bands rather than counts; the reverse-Yankee -\$400bn figure is **BIS Working Paper 590, not the March 2019 Quarterly Review**; and the Perli-Sack 16-30% amplification is a January-to-May-2003 measurement stated in their conclusions, not a structural constant. Everything else in this note is one verification pass *behind* those fourteen.

- **This is a triage document, not research.** Every family here is summarised from a sourced sweep, not from primary documents read end to end. The four existing notes each took a week of document work; nothing in this note has had that. Treat every number as a starting point to verify, not as established.
- **The tiering is a judgement and the weights are not defensible arithmetic.** Residual, footprint, data and adjacency are scored High/Medium/Low by reading, not measured. A different reader who weighted adjacency lower would put LPI first.
- **Some claims are second-hand and are flagged in the text.** The Nomura ICE-fix-versus-LCH basis point, the FHLB \$75bn figure, and the Risk.net \$200m-a-year dealer P&L estimate are all practitioner statements I have not verified against a primary source. The Regulation ZZ question is explicitly open.
- **The absence of an official-sector footprint for CMS steepeners is a finding, not a gap in the search.** BIS QR 2005-07, IMF GFSR 2005, ECB Monthly Bulletin and ECB FSR were checked; the curve facts and the product facts exist separately and nothing joins them. Plan accordingly.
- **The TPEx figures are a single-day snapshot (31 July 2026) of outstanding bonds** and are not directly comparable to the TBAC's ~\$129bn USD callable figure from April 2024, which is Bloomberg-sourced and may use a different universe definition. Reconciling the two is a task, not a footnote.
- **No claim here about the Livret A hedging question should be relied on.** I did not find a published analysis; that may mean none exists, or that it is in French practitioner literature I did not reach.