

Limited Price Indexation

The largest option position in sterling, assembled by statute, held by people who did not buy it, and hedged with an instrument that cannot replicate it

Research note — 2 August 2026 · companion to the TEC10, US CMT, Formosa and Zero Floor notes

What LPI is, and who has it

Limited Price Indexation is the rule that a UK company pension must rise with inflation each year, but not by more than a fixed maximum, and never fall. That is the whole of it. If inflation is 3%, the pension rises 3%. If inflation is 9%, the pension rises by the cap — 5% or 2.5% depending on when the benefit was earned. If prices fall, the pension does not fall; it stays where it is.

Written as a payoff, it is an inflation collar. The pensioner is long inflation between 0% and the cap and has no exposure outside that range. Whoever owes the pension is the other side: **short a floor at 0%, long a cap at 5% or 2.5%.** Because the increase applies to the pension *as already increased*, the collar resets every year on a notional that ratchets upward — which is why this is a strip of annual options and not one long-dated option.

Why it exists. Before 1997, a company pension could be eroded to nothing by inflation with no legal remedy. The Pensions Act 1995 fixed that by requiring annual increases. But an *uncapped* promise would have been open-ended for employers, so Parliament capped it. The cap was cut from 5% to 2.5% in 2005 for newly earned benefits, for the same reason. **LPI is not a financial product. It is a compromise between protecting pensioners and not bankrupting sponsors, and the option structure is a by-product nobody designed.**

Who has it, and how much. Around **10.5 million people** hold private-sector defined-benefit pensions. The PPF records that **about 85% of schemes** provide capped inflation increases on benefits earned after April 1997, most commonly capped at 5%. The liability sits with roughly **£1.4 trillion** of UK DB schemes, hedged through **over £1 trillion** invested in liability-driven investment products, against a market of **£688.5bn** of index-linked gilts.

Who is on the other side of the option, and this is the problem. Nobody, in any deliberate sense. The pensioner did not buy the collar and mostly does not know its strike. The scheme did not choose to sell it. The sponsor inherited it. And the natural supply that would let a scheme lay it off — bonds or swaps whose payments are collared the same way — barely exists: **£1.16bn of collared sterling bonds in twenty-five years**, of which 86% was placed by one company directly into its own pension scheme.

So the rest of this note is about one question: what happens when several million people hold an option position that was created by statute, is held by institutions that did not choose it, and cannot be transferred to anyone who wants it.

Executive summary

A UK defined-benefit pension in payment is a strip of annual inflation collars on a ratcheting notional, and nobody chose it. Every other family in this series begins with someone selling an option for yield. LPI does not. It is assembled from three provisions in two Acts that never mention each other's numbers: Pensions Act 1995 s.51(2) imposes the duty to *increase*; s.51ZA points at an external order; and Pension Schemes Act 1993 Schedule 3 paragraph 2(3) makes that order "*the lesser of*" realised inflation and a statutory maximum of **5 per cent** for service between 6 April 1997 and 5 April 2005, **2.5 per cent** thereafter. The cap lives in the 1993 Act. The duty lives in the 1995 Act. **The floor is drafted nowhere at all** — it exists because the statute uses the verb "increase" and the schedule uses the verb "add to". This is the third category from the Pipeline survey, the interesting one: not an illiquid market risk, not a model-marked parameter, but a payoff created by legislation and amendable by it.

The Bank of England diagnosed the problem precisely, in its own words, and then forgot it. The Quarterly Bulletin of 2009 Q3 states that there are "*few natural providers of protection against some form of constrained RPI inflation (ie LPI) and therefore **no natural supply of RPI inflation options***". The Quarterly Bulletin of 2012 Q3 is more explicit still: LPI "*gives them an incentive to buy 0% RPI floors and sell 5% RPI caps*", and "*at long maturities, particularly at 20 to 30 years, these structural flows **dominate the UK market***". Across the entire 2022 LDI corpus — the December 2022 Financial Stability Report, both Cunliffe letters, a 95-page Bank staff working paper, seven speeches, four Pensions Regulator documents, a Commons select-committee report and all 77 of its written evidence submissions — **the words "LPI", "limited price indexation", "inflation collar" and "capped indexation" do not appear once**. The liability is described as "linked to inflation", full stop. The one participant who told the authorities in writing that the liability is floored at zero, Insight Investment's submission LDI0076, was not quoted in the resulting report.

The collar is standardised boilerplate that has been sitting switched off in UK bond programmes for twenty years. It is not a bespoke clause: the "*Limited Indexation Factor / Maximum Indexation Factor / Minimum Indexation Factor*" module is pre-drafted in the programme documentation of **National Grid, Vodafone, Severn Trent, Thames Water, Anglian Water, Heathrow, United Utilities, Northumbrian Water, Places for People** and a dozen more — Law Insider indexes **163 documents** containing it. **In every set of Final Terms I could read to the end, both fields say "Not Applicable."** The instrument was never unavailable. It is declined. I can evidence the election in **two issuers, five bonds, £1.16bn of original notional, across twenty-five years**: Dŵr Cymru's 2001 Sub-Class A5 and B4 ("*Minimum Indexation Factor: 1.0 / Maximum Indexation Factor: 1.05*", on RPI) and BT's three 2018 CPI-linked tranches ("*105 per cent.*" and "*100 per cent.*"). Against £1.4 trillion of DB liabilities.

Enumerating the universe first, then working down it, gives the market rather than a handful of deals. FIRDS yields **88 GBP inflation-linked bonds, £11.35bn**; reading the documents settles **34 of them, 68.7% by value**. The result: **8 collared bonds, £1.51bn**, against **26 proven uncapped, £6.29bn**, with £3.56bn still open. The largest single line, **Network Rail 1.125% 2047 at £3,960m**, is **uncapped** — its Pricing Supplement is a scanned image, which is why three passes could not read it. Three collars were found that no published account mentions: **Ottoway/Campus Living Villages £158m, Warrington Borough Council £150m at a 3% cap** — the only non-5% strike found anywhere — and

ULiving@Gloucestershire £39m, HM Treasury guaranteed. Expanded into a fixing calendar, the resolved terms give **1,519 observation dates**, of which **142 are collared and 69 of those fall in September** — the same month as the statutory LPI fixing.

And the £1.16bn is smaller than it looks, in two ways. BT's £1.0bn — 86% of the total — **never reached the market**: Part B of the Final Terms names the purchaser as **Britel Scotland II L.P.**, a vehicle "*established for the purpose of holding certain investments for the BT Pension Scheme*", non-syndicated, one dealer, as part of BT's own triennial funding agreement. It extinguished one scheme's liability rather than feeding a market. And it is **CPI-linked**, whereas among schemes with capped post-1997 indexation the PPF records **66% RPI against 19% CPI. RPI-linked collared supply comes to £160m of original notional**. Meanwhile the sovereign was asked to issue LPI gilts in 2004 and CPI-linked gilts in 2011 and refused both, in writing, on liquidity-fragmentation grounds.

The economic content is a scheme that is short a cap strip and long a floor strip, hedging both with a straight line. An index-linked gilt or a plain RPI swap pays uncapped RPI. The liability pays $\min(\max(\text{RPI}, 0\%), 5\%)$. On my illustrative model the LPI(0,5) liability is worth **4.8% less** than the uncapped equivalent — a spread of about **27bp a year** — and LPI(0,2.5) is worth **21.4% less**, about **137bp a year**. Those are my numbers, not market quotes; **no public LPI price series exists against which to check them**. The inflation delta of the liability is **0.598** for LPI(0,5) and **0.196** for LPI(0,2.5). A scheme running a 100% RPI hedge against an LPI liability is over-hedged by a factor of nearly two on the first and five on the second.

The specific unhedgeable residual is the serial correlation of inflation, and it can be isolated. Hold every single year's marginal distribution of RPI exactly fixed — so that every annual cap and every annual floor, taken in isolation, is worth exactly the same — and vary only the dependence between years. A strip of independent collars cannot move under this experiment; its value is a sum of marginal expectations, and it does not, to within 0.07% of Monte Carlo noise. **The actual liability moves by 2.82%.** That is the model-risk band. It is unhedgeable by construction: the parameter has no quote, no instrument and no market, and two desks with identical vanilla inflation surfaces will disagree by that much on the same liability. It arises because PA 1995 s.54(3) defines the "annual rate" as the rate *as previously increased* — the cap compounds on the ratcheted-up base, so the notional is a function of the whole path.

Both tails have already bound, inside seven years, and the market prices them the wrong way round. The 2022 LPI fixing was **12.6%** — the year to September, which is what the statute specifies, not the 14.2% October print almost everyone quotes — against a 5% cap: the hedge paid 12.6% and the liability rose 5.0%, a **7.6-point over-hedge in a single year** on the indexed portion, and 10.1 points against a category Y liability. And the floor is not a modelling assumption — the Occupational Pensions (Revaluation) Order 2009, S.I. 2009/3267, specifies **0.0% for both the higher and the lower revaluation percentage** in the table at article 2, printed in a statutory instrument, in the year RPI to September ran at -1.4% . Meanwhile BlackRock's published observation is that quoted LPI prices make the **floor more valuable than the cap** even though the forward curve sits nearer the cap, which is a skew the history does not support; NatWest's is that LPI spreads "*widened every year since 2013, corresponding directly with the perceived reduction in market liquidity*".

In 2022 the correct hedge ratio collapsed towards zero at exactly the moment schemes were being asked for collateral. On my model the LPI(0,5) delta falls from 0.692

at 2% inflation to **0.070 at 8%** and **0.001 at 12%**. At the September 2022 fixing of 12.6% the liability was almost entirely insensitive to further inflation, while the hedge was fully sensitive — and the Bank's index-linked leg tells you what happened when that hedge had to be sold. The Bank had **never bought index-linked gilts** as a matter of policy and had "*no prior operational capacity*"; it took a separate market notice on **11 October 2022**, after 30-year real yields rose **63 basis points in one day**, to extend purchases to linkers, and it bought them at **≥3 years** residual maturity against **>20 years** for conventionals — a judgement that the whole linker curve was dysfunctional, not just the long end. It bought **£7.2bn** of them in four auctions, 37% of a £19.3bn intervention, using only **35.9%** of the capacity it had made available.

February 2030 is the part with no hedge at all, and it is not a market risk. The RPI will be redefined to CPIH methods by statutory process; the government has stated twice, in identical words, that "*the government will not offer compensation to the holders of index-linked gilts*"; and the judicial review brought by the BT, Ford and Marks & Spencer scheme trustees was dismissed in full on 1 September 2022, with the Chancellor granted a declaration that the gilts' cessation clause **will not be triggered**. The court recorded as common ground that RPI runs "*about one percentage point higher than the CPIH*" and that the impact is "*around £90 billion – £100 billion*". Holgate J's answer to who should have considered the hedging consequences is at [174]: "*Hedging or matching indexed liabilities is likewise a matter for the financial institutions and any regulator involved. There is no legal requirement for the UKSA to have regard to such matters.*" On my model the reform costs an uncapped RPI holder **5.88%** and an LPI(0,5) liability **3.55%** — **the asset falls 2.3 points further than the liability it hedges**, and it does so precisely because the cap was already truncating the upside the reform removes. Two official estimates of the step differ by about 50bp, and **no instrument has a payoff that is a function of which one is right**.

Part 1 — What LPI is, and why it is assembled from statutes that never mention each other

1.1 The three-link chain

The LPI payoff is not created by one provision, and no single document states it. It is built from three, each of which is unremarkable on its own.

Link one, the duty. Pensions Act 1995 s.51(2): where the section applies, "*the annual rate of the pension ... **must be increased annually by at least the appropriate percentage***". Section 51(1)(b) gates it to pension "*attributable ... to pensionable service on or after 6 April 1997*".

Link two, the definition. Section 51ZA, inserted by Pensions Act 2004 s.278(7), says the appropriate percentage is "*the **higher revaluation percentage***" for a category X pension and "*the **lower revaluation percentage***" for a category Y pension, both taken from the annual order made under paragraph 2 of Schedule 3 to the Pension Schemes Act 1993.

Link three, the cap. PSA 1993 Schedule 3 paragraph 2(3): the higher revaluation percentage is "*the **lesser of***" the increase in "*the general level of prices in Great Britain*"

and "*the higher maximum rate*"; paragraph 2(6) fixes that rate at "**5 per cent**" for a twelve-month period, and the lower maximum rate at **2.5 per cent**.

So the cap is in the 1993 Act, the duty is in the 1995 Act, and neither cross-refers to the other's number. A practitioner reading s.51 alone would not find a 5% cap in it.

Two mechanical details that matter for hedging and are routinely missed. The reference period under paragraph 2(5)(b) is the twelve months **to 30 September**, and the increase applies from 1 January. The LPI fixing is a September-on-September observation, not a December one. And s.54(3) defines "*annual rate*" as the rate "**as previously increased under the rules of the scheme or under section 51**" — the cap compounds on the ratcheted-up base, not on the original pension. That single clause is what makes the liability path-dependent, and Part 4.1 is about what it costs.

1.2 The cap split, and the wrinkle in it

The statute does not tranche by accrual date directly. It tranches by a category X / category Y test. Section 51(4A) makes a pension category X if it "*became a pension in payment before the commencement day*" **or** becomes payable later and "*the whole of which is attributable to pensionable service before that day*". Section 51(4B) makes it category Y if it becomes payable on or after the commencement day and is wholly attributable to service on or after it. Section 51(4C) splits mixed cases: "*each of those parts of the pension is to be treated as if it were a separate pension*".

The commencement day is **6 April 2005**, appointed by S.I. 2005/275 art. 2(7). (S.I. 2005/695 art. 2(4) appointed 15 March 2005 "*for the purpose only of conferring power to make regulations*" — a distinction worth keeping, because the two dates are often conflated.)

The wrinkle: a pension already in payment before 6 April 2005 is category X in its entirety — 5% cap on the whole of it, including any part attributable to post-2005 service. The clean "tranche by accrual date" description is only correct for pensions coming into payment on or after that date.

Below the two statutory tranches sit two tighter ones. Post-1988 GMP carries min(price increase, **3%**) under PSA 1993 s.109(3), for the tax years 1988-89 to 1996-97 only. Pre-6 April 1988 GMP carries no scheme-side indexation obligation whatsoever. And **pre-1997 non-GMP accrual carries no statutory indexation requirement at all** — s.51(1)(b) is an express gate. That it is nonetheless indexed in most schemes is a matter of scheme rules: the PPF's *Purple Book 2025* records that "*more than three quarters of schemes provide indexation on scheme benefits accrued before 6 April 1997*", most commonly a fixed 3% or RPI capped at 5%.

Tranche	Statutory minimum	Payoff
Pre-6 Apr 1988 GMP	none	fixed
6 Apr 1988 – 5 Apr 1997 GMP	PSA 1993 s.109	min(π , 3%), floored at 0
Pre-6 Apr 1997, excess over GMP	none	fixed unless scheme rules say otherwise
6 Apr 1997 – 5 Apr 2005	PA 1995 s.51 + s.51ZA(1)(a)	LPI(0,5)
From 6 Apr 2005	PA 1995 s.51 + s.51ZA(1)(b)	LPI(0,2.5)

A single member's pension is therefore a portfolio of up to five differently-struck collars with different start dates, and the weights are a function of that member's employment history.

1.3 The floor is written nowhere

There is no drafted 0% floor in any of the relevant provisions, and I looked for one. The floor is structural, and it emerges from four separate pieces of ordinary drafting:

1. PA 1995 s.51(2) requires the pension to be "**increased**". An increase of a negative percentage is not an increase.
2. PSA 1993 Sch. 3 para. 1(1): the final salary method is to "**add to**" the amount otherwise payable "*the additional amount*".
3. PSA 1993 Sch. 3 para. 2(1): the Secretary of State shall specify percentages "**(so far as it is necessary to do so)**".
4. PSA 1993 s.109(2), for GMP: "**Where it appears to the Secretary of State that that level has increased ...**".

Note also that paragraphs 2(3) and (3A) define the percentage as "*the lesser of*" inflation and the maximum. There is no symmetric "greater of" clause anywhere. The asymmetry is deliberate and it is the entire floor.

The empirical proof is a printed statutory instrument. The Occupational Pensions (Revaluation) Order 2009, S.I. 2009/3267, made 10 December 2009 and in force 1 January 2010, specifies in the table at **article 2**, for the revaluation period 1 January to 31 December 2009, a higher revaluation percentage of **0.0%** and a lower revaluation percentage of **0.0%**. RPI in the year to September 2009 was **-1.4%**. The floor is not a modelling convention; it is printed in an SI, and it has already bound once.

One further mechanism makes the strip carry memory. PA 1995 s.53(1) lets trustees who make a discretionary increase in one tax year "*deduct the amount of the increase from any increase which ... they would be required to make ... in the next tax year*", and s.53(3) then rolls the notional base forward "*as if they had been increased*". So a scheme paying above the statutory minimum buys at most one year's relief and never a permanent base reduction. For any scheme paying discretionary increases, the payoff is not even a clean sequence of annual collars — it has a one-period lag in it.

1.4 Deferred revaluation is a different instrument, and the conflation is expensive

Pensions in payment and deferred revaluation are structurally different options and are constantly treated alike. They run on different statutes, they changed on different dates, and they have different Greeks.

Deferred pensions revalue under PSA 1993 s.84 and Schedule 3, not under PA 1995 s.51. The cap moved from 5% to 2.5% for pensionable service on or after **6 April 2009** — four years later than the in-payment change — by Pensions Act 2008 s.101 and Schedule 2 paragraph 2, commenced by S.I. 2009/82 art. 2(2)(f).

The instrument itself is different. Paragraph 2(7) takes the revaluation percentage "*for the revaluation period which is of the same length as the number of complete years in the pre-*

pension period", and paragraph 2(6)(b) sets the maximum for a multi-year period as *"the percentage that would be the inflation percentage had the general level of prices increased at the rate of 5 per cent compound per annum"*. **Deferred revaluation is therefore one cap on the cumulative compounded index over the whole pre-pension period — a single long-dated option on an average, closer to an Asian. A pension in payment is a strip of annual caps.**

The cap binds far less often on the average than on the individual years, so the deferred option is worth more. On my model, for the same 5% cap:

Horizon	In-payment strip	Deferred cumulative	Gap
10 years	1.3715	1.4017	+2.20%
20 years	1.8160	1.9063	+4.97%
30 years	2.3934	2.5713	+7.43%

Getting these the same way round is the single commonest modelling error in the area, and at thirty years it is worth seven and a half points.

1.5 RPI-linked LPI is a legislated steady state, not a legacy artefact

Statute has never named an index. PSA 1993 Sch. 3 para. 2(3)(a) says only *"the percentage increase in the general level of prices in Great Britain"*, and para. 2(4) adds that the Secretary of State *"may estimate the percentage increase ... in such manner as he thinks fit"*. The 2010 switch from RPI to CPI therefore required no primary legislation. Steve Webb announced it in a Written Ministerial Statement on 8 July 2010; the first CPI order was S.I. 2010/2861.

And that is exactly why it did not reach scheme rules. The DWP consulted on whether to give schemes a statutory override and refused, in terms, in June 2011: *"The Government is satisfied the consultation process has not identified any new or compelling evidence to warrant the introduction of a statutory override"* (¶34), and *"the Government does not propose to introduce a modification power"* (¶43).

Then Parliament removed the pressure to switch. Pensions Act 2011 s.19 inserted PA 1995 s.51(4ZB)-(4ZG), letting a scheme that has continuously increased by RPI since the beginning of 2011 keep doing so, capped at *"the default percentage"* of 5% or 2.5%, without comparing against the CPI-based statutory figure. The Explanatory Notes say it plainly: the amendments *"ensure they need not carry out an annual comparison of the RPI under scheme rules and CPI under the statutory requirements and pay the higher of the two"*.

Absent s.19 there would have been a one-way max(RPI-capped, CPI-capped) ratchet, and schemes would have been pushed towards CPI. Section 19 killed that. **The statutory floor is CPI-based, a large stock of scheme liabilities is contractually RPI-based, and the two never have to be compared.** RPI-linked LPI is not a legacy problem awaiting cleanup. It is a legislated equilibrium.

Nor can most schemes switch unilaterally. The leading authority is *Barnardo's v Buckinghamshire* [2018] UKSC 55, and the scheme **lost**. The rule defined RPI as the published index *"or any replacement adopted by the Trustees"*; Lord Hodge held at [19] that

this means "*the RPI or any index that **replaces** the RPI and is adopted by the trustees*", and at [23] that it is "*that **official body** and not the trustees who are to effect the replacement*". The interpretive posture is hostile to rescue: at [27]–[29], "*Only by relying on hindsight can weight be given to this consideration; and that is not legitimate.*"

The outcome turns entirely on the individual scheme's drafting. *Danks v QinetiQ* ("any other **suitable cost-of-living index selected by the Trustees**") permitted a switch; *Arcadia* ("any **similar index satisfactory for the purposes of** HMRC") permitted one jointly; *Thales* gave a qualified answer with a fiduciary brake and a one-shot limit; *Britvic* was won by the employer on appeal. There is no general power, and the Supreme Court expressly noted at [11] that whether many schemes share the *Barnardo's* wording "*is not clear*".

Part 2 — The supply side: the clause exists, the deals do not

2.1 What I was looking for, and the phrase that finds it

Uncapped RPI-linked paper is abundant. Collared paper is what a scheme actually needs, and the practitioner literature says almost none exists. I wanted a number on "almost none", built bottom-up from offering documents.

The search key turns out to be a specific piece of drafting. UK index-linked bonds that carry an LPI collar do not say "LPI". They define an **"Indexation Factor"** and then constrain it with a **"Minimum Indexation Factor"** and a **"Maximum Indexation Factor"**, specified in the pricing supplement. A minimum of 1.0 is a 0% floor; a maximum of 1.05 is a 5% cap. Searching for "LPI" finds commentary. Searching for **"Maximum Indexation Factor"** finds instruments.

2.2 The denominator, and why document search could never produce one

A full-text search across prospectuses can never tell you what it missed. That is the structural defect in the first two passes of this register, and it is fixed by inverting the problem: instead of searching documents to find instruments, take the regulatory instrument database and enumerate the universe first.

FIRDS — the reference database of every instrument admitted to trading on an EU or UK venue — gives that universe. `01_register/firds_universe.py` reproduces it end to end.

Two things make the screen non-obvious, and the second one is a trap.

FIRDS carries **no index-linked flag**. The ISO 10962 CFI code distinguishes fixed from variable interest, and every inflation-linked bond is variable — but so is every SONIA floater and every RMBS tranche. Necessary, nowhere near sufficient. The usable discriminator is `bnd_fltng_rt_ndx_name`, the floating reference index: a SONIA floater names its benchmark, while an inflation-linked bond reports `FORMULA` or names RPI or CPI directly.

And **names are unreliable in exactly the wrong direction**. Only about 3% of GBP variable-rate debt carries any inflation marker in its FIRDS name. Of the five collared bonds established from documents, the two Dŵr Cymru lines are named **"DWRCYFIN FRN 31 03 2027"** and **"DWRCYFIN FRN 31 03 2031"** — they are literally labelled FRN. A name search finds neither. The SPIRE tracker's name-based linker detection works well on sovereigns, whose naming is standardised; **it does not transfer to corporates**, and assuming it did is what made the earlier passes miss things.

The result: 88 GBP inflation-linked bonds, £11.35bn of nominal. Working down that list document by document then settles the terms:

	Bonds	Nominal
Collared	8	£1.51bn
Uncapped — proven, not assumed	26	£6.29bn
Unknown — terms not established	54	£3.56bn

	Bonds	Nominal
Universe	88	£11.35bn

Coverage: 34 of 88 bonds — 38.6% by count, 68.7% by value. The register can now say what most of the market is, rather than what a handful of deals are. `01_register/lpi_terms.csv` carries the cap, floor, index, lag, base index figure, fixing months, payment dates, amortisation and 2030 fallback for every resolved line.

The screen also immediately surfaced issuers the document sweep never reached. The largest single line is **Network Rail 1.125% 2047 RPI at £3,960m** — four times the entire BT deal, absent from every previous pass. Behind it: Thames Water 2053 and 2055, National Grid Gas and Electricity Transmission across a dozen lines, Cadent, Severn Trent, Yorkshire Water, Affinity Water, United Utilities, Anglian Water, Bazalgette, Scottish Power, Eastern Power Networks, Western Power Distribution, Places for People, THFC, Warrington Borough Council, and a tail of PFI and student-accommodation issuers — UPP, Ottoway, Arlington, ULiving, Aysgarth, HHT. **Forty-five of the 88 are RPI-linked and eight CPI-linked.**

Network Rail is the single largest position and it is uncapped. The Pricing Supplement of 18 September 2007 is a scanned image with no text layer, which is why three passes could not read it; rendering it in a browser and reading it visually settles it. Minimum Rate of Interest and Maximum Rate of Interest both **"Not Applicable"**, redemption at outstanding nominal subject to the full Index Ratio, no par floor. Base Index Figure **206.54000**, RPI with a three-month lag and daily interpolation — which reconciles exactly to $207.3 + (19/30) \times (206.1 - 207.3)$ off the June and July 2007 prints. **The programme has no Limited Indexation machinery at all;** it uses the gilt-style full Index Ratio.

And that turns out to be the pattern. Whether a bond *can* be collared is often structural rather than an election. United Utilities' Final Terms for index-linked notes run Base Index Figure, Index Figure applicable to, *t*, Reference Gilt, Minimum Rate of Interest, Maximum Rate of Interest — and that is the complete list. Anglian's Indexation Bond Provisions carry only Base Index Figure, "any other terms", Reference Gilt. **Those programmes cannot elect a collar without bespoke drafting.** Severn Trent's, by contrast, carries the switch as a live drafted option — so a missing Severn Trent Final Terms is a real gap, while a missing United Utilities one is not.

The £3.56bn still unknown is unknown, not uncapped, and it is now concentrated in two places: the nine 2005-06 National Grid and Cadent private placements, whose bond-level Final Terms are on no free public venue at all, and a handful of pre-archive deals. One of them, **UPP Bond 1**, matters more than its £75m suggests: the programme is a Limited-Index-Linked structure, with Limited Index Ratio and Limited Indexation Month in Condition 7 — only the cap and floor values sit in non-public tranche terms. **That one must not be defaulted to uncapped.**

Two data-quality findings about FIRDS itself. HHT PLC, carrying £798m across two lines, was dissolved in March 2021 — its FIRDS rows are stale and the bonds should be treated as redeemed. And the "INSTS" tag on the National Grid family is "Instruments", the programme's own word for its notes, not instalments; those bonds are bullets. A register built from FIRDS alone would have got both wrong.

2.3 The fixing calendar, and where the collars actually observe

A register of instruments becomes more useful as a register of dates. 06_calendar/fixings.csv expands the resolved terms into **1,519 indexation observations across 26 bonds, running from 2026 to 2062**. Of those, **142 observations on 7 bonds are collared** — the dates on which a cap or floor can actually bind on a sterling corporate bond.

They cluster, and they cluster on the same month as the liability.

Observation month	Collared observations
September	69
June	39
March	34

The statutory LPI fixing is the year to **30 September** (PSA 1993 Sch. 3 para. 2(5)(b)). So the largest concentration of collared corporate observations sits on the same print that sets the pension increase — Dŵr Cymru's Limited Indexation Months are "*March and September*", Ottoway's collar is measured September-over-September, ULiving's is "*March and September*". That alignment is convenient for a hedger and uncomfortable for a market: a single September RPI print simultaneously moves the liability and the only instruments that match it, with no diversification across observation dates of the kind the TEC10 calendar showed.

The June cluster is BT's £1bn, on CPI — a different index, a different month, and a bond that never left the sponsor's own scheme.

2.3 The collar is boilerplate, and it has been sitting unused for twenty years

This is the finding that reframes the whole supply question. The collar is not a bespoke clause negotiated deal by deal. It is an *optional module* inside the generic UK "Index Linked Instruments" condition, itself descended from the index-linked gilt conditions, and it is pre-drafted in essentially every UK regulated-utility EMTN programme and trust deed:

"Limited Indexation Factor" means, in respect of a Limited Indexation Month ..., the ratio of the Index Figure applicable to that month ... divided by the Index Figure applicable to the month ... **twelve months prior thereto**, provided that (a) if such ratio is greater than the **Maximum Indexation Factor** specified in the relevant Final Terms, it shall be deemed to be equal to such Maximum Indexation Factor and (b) if such ratio is less than the **Minimum Indexation Factor** specified in the relevant Final Terms, it shall be deemed to be equal to such Minimum Indexation Factor;

"Limited Index Linked Instruments" means Index Linked Instruments to which a Maximum Indexation Factor and/or a Minimum Indexation Factor ... applies.

Law Insider indexes **163 documents** containing "Limited Indexation Factor" and 156 containing "Limited Indexation Month". The clause sits in the programme documentation of **National Grid, Vodafone, Western Power Distribution, Affinity Water, Northumbrian Water, National Gas Transmission, Severn Trent, Thames Water, Anglian Water, Heathrow Funding, United Utilities, Eastern Power Networks, Arqiva, Bazalgette,**

Places for People, Motability and BT — and, in near-identical form, in an Italian bank's programme. It is not UK-specific and it is not rare.

In every set of Final Terms I could read to the end, both fields say "Not Applicable." Southern Water's February 2024 CPI-linked sustainability bond: Minimum and Maximum Indexation Factor both "*Not Applicable*", Limited Indexation Months "*Not Applicable*". Western Power Distribution's 2052 index-linked notes: "*Not Applicable*". Severn Trent's programme carries the option and issues plain fixed.

So the constraint on LPI supply is not legal technology. The documentation to issue collared inflation paper has been sitting, pre-drafted and lawyer-approved, in the programmes of most UK regulated utilities for the better part of twenty years. What is missing is an issuer willing to sell the cap and a price at which anyone will buy it. **The option is enumerated in nineteen-plus programmes. I can evidence the election in two.**

2.4 The two deals that did elect it

Dŵr Cymru, 2001 — and I got this wrong the first time. An earlier pass of this note concluded that Welsh Water "discloses no cap or floor terms". That was wrong, and it was wrong because I read the semi-annual Investor Report, which carries a principal-balance table and no indexation terms, instead of the prospectus, which carries the machinery, and the pricing supplement, which carries the numbers. Condition 7 of the 2001 programme prospectus defines Limited Indexed Bonds exactly as above. The Pricing Supplement dated 9 May 2001 for the **Sub-Class B4** bonds completes it:

(ix) Minimum Indexation Factor: 1.0 (x) Maximum Indexation Factor: 1.05 (xi) Limited Indexation Month(s): March and September (i) Index/Formula: UK Limited Price Index

That is LPI(0,5), on RPI, written into a corporate bond, and the pricing supplement calls the index by name.

BT, 2018 — and this is the one everybody misses. Three CPI-linked tranches issued 25 June 2018 under BT's EMTN programme, £330m 2033, £330m 2039 and £340m 2042, **£1.0bn in total**. BT did not use the Final Terms fields; it hard-coded the values into a bespoke appendix: "*Maximum Indexation Factor*" means 105 per cent.; "*Minimum Indexation Factor*" means 100 per cent. Again LPI(0,5), annually collared and compounded.

Issuer	ISIN	Size	Index	Collar	Maturity
Dŵr Cymru B4	XS0129065446	£75m	RPI, 8-month lag	1.00 / 1.05	2027
Dŵr Cymru A5	XS0129065362	£85m	RPI	1.00 / 1.05 (<i>reported, not read</i>)	2031
BT plc	XS1833077545	£330m	CPI	1.00 / 1.05	2033
BT plc	XS1833078196	£330m	CPI	1.00 / 1.05	2039
BT plc	XS1833078600	£340m	CPI	1.00 / 1.05 (<i>reported</i>)	2042

£1.16bn of original notional, five bonds, two issuers, twenty-five years. Against a defined-benefit sector of £1.4 trillion and over £1 trillion invested in LDI products. The Dŵr Cymru pair has accreted to £344.8m by 31 March 2026 — about 6% of that group's own debt, and the rest of its index-linked stock is **uncollared**.

Two details make the £1.16bn smaller than it looks.

First, **the BT bonds never reached the market.** Part B of every set of Final Terms names the purchaser: **Britel Scotland II L.P.**, "*an investment holding vehicle established for the purpose of holding certain investments for the BT Pension Scheme*", with BT as scheme sponsor. Method of distribution: non-syndicated, one dealer. It was part of BT's triennial pension funding agreement. **The largest LPI bond family in sterling is a sponsor-to-own-scheme placement that happens to be listed.** It extinguished one scheme's liability; it never fed a market.

Second, **the index is wrong for most of the demand.** BT's £1.0bn — 86% of the total — is **CPI-linked**. The PPF's *Purple Book* records that among schemes with capped inflation on post-1997 accrual, **66% are RPI-linked against 19% CPI-linked**. So of the collared sterling paper that exists, the large majority references the wrong index for the large majority of the liability. **RPI-linked collared supply comes to £160m of original notional. That is the number.**

2.5 Tesco Property Finance is not an LPI bond

The other name in the received account is a straight error, and the prospectus settles it in one sentence. All six Tesco Property Finance issuers carry **fixed** coupons. The TPF 5 prospectus is titled "*£450,500,000 Secured **5.6611 per cent.** Bonds due 13 October 2041*" and states:

"**The Bonds and the Partnership Loan will bear interest at a fixed rate, whereas the Rental Income will be linked to a Retail Prices Index ("RPI"). Accordingly, prior to or on the Closing Date, in order to hedge the exposure which would otherwise exist:** (a) the Partnership and the Issuer (in its capacity as "Partnership Swap Provider") will enter ..."

The RPI is in the leases, it is swapped out inside the structure, and the bondholder never sees it. **There is no coupon indexation to collar.** The likely source of the confusion is that Tesco plc's own *unsecured* RPI-linked bonds (XS0137945373, XS0248395088) and Tesco Bank's RPI-linked retail bond do exist and are genuinely index-linked. Tesco *Property Finance* is not.

2.6 The sovereign refused, in writing, twice

Two rows in the register have a size of zero and they carry as much of the argument as the rest put together.

In December 2004 the DMO consulted on ultra-long gilts and recorded, at paragraph 13, a proposal for "*index-linked bonds with limited price indexation properties (LPI bonds), i.e. where the indexation of the cash flows to the reference price index is **capped to the upside and/or to the downside***". Paragraph 17 refused it:

"HM Government is therefore not inclined to issue instruments that are likely to appeal to a limited group of investors or that may lead to a fragmentation of the gilts market, with a resulting loss of liquidity. **It is, therefore, not inclined to pursue the option of issuing LPI bonds.**"

The market's reply, in the March 2005 response, is worth reading twice: *"Some saw a mix of conventional and index-linked gilts as being important to manage limited price-indexed (LPI) liabilities."* That is delta-hedging, described to the sovereign issuer, in 2005.

In November 2011 the DMO refused CPI-linked gilts, partly on *"fragmentation in the index-linked gilt market, and associated illiquidity"*. Paragraph 26 of the response is decisive for Part 7:

"Respondents were universally of the view that the prospects for CPI/RPI hedging products emerging as an alternative liability management tool to CPI-linked gilts is unlikely in the absence of the issuance of CPI-linked gilts."

Both were "kept under review". The Lords Economic Affairs Committee repeated the call for CPI-linked issuance in January 2019. Nothing has happened in fifteen and twenty-two years respectively.

2.7 What the supply picture actually is

Not "there is no LPI market because the instrument does not exist". The instrument exists, is standardised, and is declined.

Corporate issuers have carried the clause for two decades and left it switched off, because electing it means selling a 5% inflation cap — short optionality they would have to price, hedge and hold. The sovereign was asked twice and refused twice, on liquidity-fragmentation grounds. The dealers who once intermediated it have left: RBS told *Risk* in 2007 that *"an active two-way market in LPI [0,5] has developed out to 50 years"*, and by 2017 NatWest was reporting that *"all but a very small number of banks"* had withdrawn. And the one £1bn transaction that did happen was a sponsor handing paper directly to its own scheme.

Every party who could supply the collar has independently decided not to. That is a much stronger statement than a missing market, and it is the reason the residual in Part 4 has nowhere to go.

Two honest limits on all of this. **No full-length base prospectus was read end to end** — the fetch tooling truncates at around 120,000 characters, so the nineteen programmes carrying the clause are enumerated rather than adjudicated, and any one of them may contain an elected series I did not find. And **privately placed collared paper leaves no public trace at all**. The BT deal is visible only because it happened to be listed; a bilateral placement that is not listed is invisible, and the size of that population is **NONE FOUND from any source**. Given that the one large deal I did find was itself bilateral, this is likely the largest single source of error in the register, and it runs in the direction of **understating** supply. # Part 3 — What the scheme is actually short

3.1 Net of the whole chain

Work out who is long what and the picture is unusually lopsided, because one side of it does not trade.

The **member** is long a collared inflation-linked annuity. They did not choose the collar and in most cases do not know the strike.

The **scheme** is short that annuity. Decomposed against a plain RPI leg, it is **short a strip of annual RPI caps struck at 5% or 2.5%** and **long a strip of annual RPI floors struck at 0%**, both applied to a notional that ratchets up with realised capped inflation.

The **sponsor** stands behind the scheme, so ultimately holds the residual after assets.

The **LDI manager** sells the scheme a leveraged position in index-linked gilts and RPI swaps — a **linear** inflation hedge — plus, occasionally, an LPI swap.

The **dealer**, on the rare occasions an LPI swap trades, is short the collar to the scheme and must warehouse it, because there is no natural other side.

And that is where the chain stops. In the Formosa note the dealer sells the Bermudan on into the street. Here there is nobody to sell it to. The Bank of England said so in 2009 and the sentence has not been superseded:

"There are some institutions/organisations that are natural suppliers of RPI inflation protection, such as the UK government and utility companies ... **But in contrast there are few natural providers of protection against some form of constrained RPI inflation (ie LPI) and therefore no natural supply of RPI inflation options.** Instead, the supply of RPI floors has largely relied upon the ability of the sellers of those options to absorb the risk."

The same box records what happened when the risk stopped being absorbable: "*Trading in inflation floors was also affected once it became clear that RPI inflation would turn negative in 2009 leading to **losses for some market makers.***"

3.2 The collar, priced

The illustrative model in `lpi_model.py` is a one-factor Gaussian process for year-on-year RPI: 250,000 paths, forty years, a 1.5% annual volatility taken from LGIM's published working number, mean reversion on a four-year half-life, and a drift running from a 4.0% spot to a 3.4% long-run level stepping down to 2.9% after February 2030. It is not a pricing model and none of what follows is a market quote.

Per 1.0 of starting pension, discounted at 4.3%:

	Cap	PV	vs uncapped RPI	Spread to RPI
Uncapped RPI	—	32.895	—	—
LPI(0,5) , category X	5.0%	31.335	−4.79%	27bp
LPI(0,3) , post-88 GMP	3.0%	27.201	−17.31%	107bp
LPI(0,2.5) , category Y	2.5%	25.867	−21.41%	137bp

The "spread" is the flat annual deduction from the index that reproduces the same present value — the quantity an LPI swap quotes. **There is no public LPI swap price series to check these against.** That absence is the subject of the note, and it is the reason the figures are labelled as mine throughout.


Note the shape: the step from a 5% cap to a 2.5% cap is worth far more than half the first cap, because 2.5% sits close to the middle of the forward distribution while 5% sits well into its tail. Category Y liabilities are a substantially different instrument from category X, not a scaled version of one.

3.3 LPI is not a strip of independent collars

This is the sentence the rest of the note turns on. The standard shortcut is to value each year's collar independently and compound the results. It is wrong, and the model can size the error:

	True payoff, ratcheting notional	Strip of independent collars	Error
LPI(0,5)	31.344	31.174	-0.54%
LPI(0,2.5)	25.869	25.841	-0.11%

Half a per cent of a liability is not a rounding error at this size, but the level is not the point. **The point is that the error exists at all**, because it tells you the liability is a function of the whole path and not of the marginal distributions — and therefore that it depends on a parameter no vanilla inflation option market can pin down. Part 4.1 is what that costs.



The mechanism is s.54(3). Writing $\pi(t)$ for the year-on-year RPI observed in the year to September t , and $P(t)$ for the pension in payment:

$$P(t) = P(t-1) \times (1 + \min(\max(\pi(t), 0\%), c)), \text{ where } c = 5\% \text{ or } 2.5\%$$

The cap applies to the *increment*, but the increment applies to a base that has already ratcheted.

The market's own drafting confirms this is the right construction, which is a useful check on the model. The bond documentation in Part 2.3 builds the payment multiplier the same way: the *Limited Indexation Factor* is a **twelve-month** ratio, individually collared, and the *Limited Index Ratio* is then "the product of the *Limited Indexation Factor* for that month and the *Limited Index Ratio* as previously calculated in respect of the month twelve months prior thereto" — a running product of annually-collared factors. Not a cap on cumulative inflation. Eight per cent followed by two per cent pays less than five per cent twice, in the

bond exactly as in the pension. The `lpi_ratchet()` function in the model is that recursion. Two paths with identical marginal distributions in every year, differing only in how the years are ordered and correlated, produce different liabilities.

Part 4 — The residuals

4.1 Path dependence, and the model-risk band

The unhedgeable parameter is the serial correlation of inflation, and it can be isolated cleanly.

The experiment is the one the TEC10 note used on inter-date correlation, adapted. Build the inflation paths so that every year's marginal distribution is *exactly* fixed regardless of the dependence — take a standard normal AR(1), $z(t) = \rho \cdot z(t-1) + \sqrt{1-\rho^2} \cdot e(t)$, which is standard normal in every period for any ρ , and map it through a fixed marginal. Then vary ρ alone.


Under this experiment **every annual cap and every annual floor, taken in isolation, is worth exactly the same in every row.** Nothing a vanilla inflation option market could price has changed.

ρ	LPI(0,5)	LPI(0,2.5)	Independent strip
0.00	31.178	25.846	0.56231
0.30	31.242	25.857	0.56241
0.60	31.390	25.881	0.56252
0.90	32.057	25.994	0.56267
Band	2.82%	0.57%	0.07%

The independent strip does not move, and cannot, because its value is a sum of marginal expectations — the 0.07% is Monte Carlo noise. The actual liability moves by 2.82%.

That is the model-risk band, and it is unhedgeable by construction. Serial correlation of realised inflation is not quoted anywhere, has no instrument, and is not extractable from an option surface even if one existed. Two desks with identical vanilla inflation surfaces, identical forward curves and identical volatilities will disagree by 2.82% of the liability on LPI(0,5), and there is no trade that closes the difference.

Figure 2 — the model-risk band, with every vanilla price held fixed



Note the direction: **higher persistence makes the liability more expensive**. Inflation that stays high for several consecutive years ratchets the notional through the cap repeatedly; inflation with the same marginals but no persistence averages out against the cap. Anyone whose model assumes independent annual inflation is systematically under-reserving, and the error is largest for the 5% cap because that is where the ratchet has room to run.

4.2 The skew nobody can justify

The market prices the floor above the cap, and the history does not support it. BlackRock states it directly:

"If we look at the current RPI forward curve, expected inflation is, on average, **much nearer to the 5% cap than the 0% floor** ... A combination of these factors would suggest that **the cap should be more expensive than the floor**. However ... **we can see that the converse is true and floors are more valuable than caps** ... market pricing has been driven by technical supply and demand factors rather than the fundamentals of where inflation has typically printed."

NatWest's version, from 2017: **"The LPI spread to RPI ... implies that the underlying RPI distribution exhibits a strong negative skew ... Whilst historic RPI analysis does exhibit some skew ... it is far less pronounced than LPI pricing might imply, suggesting that demand for the floor (or lack of supply) has driven prices upward."**

Running the same marginal-preserving experiment on skew instead of correlation, holding mean and variance fixed:

Marginal skew	LPI(0,5)	Cap leg	Floor leg
-0.60	33.884	0.0168	0.0570

Marginal skew	LPI(0,5)	Cap leg	Floor leg
0.00	31.333	0.0714	0.0300
+0.60	29.834	0.0930	0.0000

The negative-skew row is the one that reproduces the market's floors-above-caps ordering, and it prices the liability **8.1% higher** than the symmetric case. So the observed pricing anomaly is not a curiosity — it is worth eight points of liability, and it is driven by a supply-demand imbalance rather than by a view about inflation.

4.3 There is no executable inflation volatility market above a linear market that barely functions

The option market is model-marked because the linear market underneath it cannot support an executable fixing. ICE Benchmark Administration consulted in March 2026 on launching inflation swap benchmarks at all, proposing **Level 2** for EUR HICPxT and GBP RPI — dealer-to-client prices displayed electronically — with Level 1 executable order-book pricing only "*when sufficient volumes are available*", and USD CPI deferred. **It launched them on 5 June 2026**, GBP UK RPI 1100 running 1Y to 30Y off Tradeweb Level 2 data, with USD CPI still deferred.

That is a genuine improvement and it is the single most important development in this subject in a decade. But note what it is: a **Level 2** fixing, built from dealer-to-client prices displayed by a venue, not from an executable order book. A linear benchmark that starts at Level 2 does not yet give the option market above it something to mark against. **If the linear market cannot support an executable benchmark, the option market above it is model-marked by construction** — and it remains so today. The Bank made the same observation in 2009 — "*market activity in UK inflation caps and floors has typically been **low** compared to other inflation-linked products*" — and in 2012 warned about the consequences for its own published implied distributions: "*Lack of liquidity does not mean that the implied pdfs contain no information. But they may be **noisy or slow to react to news**.*"

The sell side's account of who is left is uniform. LGIM, 2018: "*the UK LPI swap market has become increasingly illiquid; **most banks have withdrawn from the market** or are pricing these contracts at levels not justified by historic inflation.*" NatWest, 2017: "***all but a very small number of banks** having withdrawn from the LPI swap market*", with spreads that "*widened every year since 2013, corresponding directly with the perceived reduction in market liquidity*". BlackRock, 2024: "*over recent years, we have seen **more LPI trades unwound than new trades being initiated***".

It was not always so. RBS told *Risk* in September 2007 that "*an active two-way market in LPI [0,5] has developed out to 50 years*", and an IFoA paper of June 2007 found that "*the cost of hedging LPI(0,5) cashflows ... is similar, or slightly cheaper, than the cost of a full RPI hedge*". A market that existed in 2007 does not exist now. Whatever else this is, it is not a market that was never built — it is one that was abandoned.

4.4 The profession has no agreed way to measure the exposure

There is no standardised method for computing the inflation sensitivity of an LPI liability, and the body that set out to write one disbanded without publishing. The Institute and Faculty of Actuaries constituted an LPI Risk Working Party whose stated first task was *"a comprehensive review of all the existing alternative methodologies used to calculate the IE01 of LPI-linked benefits"*, noting that *"the presence of these caps and floors complicates the calculation of the inflation sensitivity (IE01) of these benefits. As a result many pension Schemes adopt a 'delta-hedging' approach."* No output was published.

What is used instead is what BlackRock calls *"so-called 3D-2D modelling"*, run *"either using market implied volatilities or a flat volatility assumption"*, of which it says: **"Both have their limitations ... this is an illiquid and heavily skewed market."**

LGIM quantifies the resulting dispersion: **"The delta increases by around 10% switching from a market consistent to real-world methodology at 20 years"**, and **"overall we estimate that LPI(0,5) risk could be in the region of 1.5%-2.0% p.a. ... a scheme that has 100% of benefits linked to year-on-year LPI(0,5) could be exposed to LPI risk of a similar magnitude to longevity risk."**

That last comparison deserves to be better known. Longevity risk gets a dedicated market, dedicated transactions, dedicated regulatory attention and a line in every trustee report. LPI risk of comparable magnitude gets a flat volatility assumption.

4.5 The benchmark is amendable by statute

The fourth residual is that the index itself is not a market object. It is Part 7.

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

5.1 Delta-hedging an option book with an instrument that has no convexity

A scheme holding index-linked gilts against an LPI liability is running a delta hedge on an option book using a linear asset, and the hedge ratio is a function of the thing being hedged.


Inflation delta of the liability, relative to uncapped RPI, for a parallel 10bp bump to the whole forward inflation curve:

Delta vs RPI	
LPI(0,5)	0.598
LPI(0,3)	0.268
LPI(0,2.5)	0.196

A ratio of 1.00 would mean a plain RPI hedge is exactly right. **It never is.** A scheme running a nominally 100% inflation hedge against a category X liability is over-hedged by a factor of about 1.7; against a category Y liability, by about 5.

And the ratio is not stable, which is why it has to be rebalanced:

Prevailing RPI	LPI(0,5) delta	LPI(0,2.5) delta
0.5%	0.623	0.409
2.0%	0.692	0.343
3.4%	0.615	0.211
5.0%	0.400	0.085
8.0%	0.070	0.006
12.0%	0.001	0.000



The delta peaks in the middle and collapses towards zero at both ends — towards zero as inflation rises through the cap, and towards zero as it falls through the floor. The correct hedge is largest exactly where the liability is *least* sensitive to further inflation. And a scheme that has hedged with linear instruments has no convexity with which to track the change; it can only trade.

5.2 Both tails, inside seven years

Mind the month. PSA 1993 Schedule 3 paragraph 2(5)(b) runs the reference period to **30 September**. The 2022 LPI fixing was therefore **12.6%**, not the **14.2%** October print that is almost universally quoted in commentary on this episode. The October number is the peak of the series; it is not the number any pension was uprated by. Both are in the table because the difference — 1.6 points — is itself a reminder of how loosely this liability is described.

	RPI leg (the hedge pays)	LPI(0,5) (the liability owes)	Leak
Year to September 2022 — the fixing	12.60%	5.00%	–7.60%
Same year, category Y	12.60%	2.50%	–10.10%
Year to October 2022 — the print usually quoted	14.20%	5.00%	–9.20%
Year to September 2009	–1.40%	0.00%	+1.40%

A scheme fully hedged with index-linked gilts or plain RPI swaps received 12.6% against a liability that rose 5.0% — a **7.6-point over-hedge in a single year** on the indexed portion. That is a gain, and it was booked as one. The problem is that it was a gain on a position the scheme did not know it had, in the same episode in which it was being asked for collateral on the rest of the hedge.

In the deflationary year the leak ran the other way: the hedge lost value while the liability did not fall at all, because the statutory order specified 0.0%.

Both tails of a collar, realised inside seven years, on a liability of over £1 trillion, in a market that most dealers had by then left.

5.3 What that did in 2022, and what the official record says about it

The index-linked leg of the 2022 intervention is the best-documented moment in the whole subject, and the documentation never once names the option.

The sequence matters. The Bank's first market notice, 28 September 2022, covered **conventional** gilts of more than 20 years' residual maturity. It took until **11 October** — a separate, later market notice — to extend purchases to index-linked gilts. Sir Jon Cunliffe explained why on 18 October:

"As a policy matter, **the Bank has never previously purchased index-linked gilts in its monetary policy operations ... the Bank had no prior operational capacity to make such purchases** and preparing to do so required significant operational changes."

Andrew Hauser gave three reasons on 4 November, of which the third is the one that broke: dealers had been expected to buy linkers from LDI funds and sell conventionals into the Bank's operations, "**warehousing the inflation basis risk**". On 10 October that stopped working. Thirty-year real yields rose **63 basis points in a day**; the Bank's press notice on 11 October called index-linked dysfunction "*a material risk to UK financial stability*".

Two details in the 11 October notice are diagnostic. The Bank bought index-linked gilts at **three years and above** residual maturity, against **more than twenty years** for conventionals — a judgement that the entire linker curve was dysfunctional, not just the long end. And it set "*a minimum yield*" per gilt, refusing to allocate below the real yields quoted at the close on 10 October.

The auction record, from the Bank's own published spreadsheet:

Date	Offered (£mn)	Allocated (£mn)	% of the £5bn cap
11 Oct	2,413.8	1,947.0	38.9%
12 Oct	2,017.8	1,979.8	39.6%
13 Oct	3,176.8	3,128.5	62.6%
14 Oct	129.4	129.4	2.6%
Total	7,737.8	7,184.6	35.9%

£7.2bn of index-linked against £12.1bn of conventionals — 37% of a £19.3bn intervention, bought on four of thirteen operating days, using barely a third of the capacity made available.

The official diagnosis was leverage and collateral-operational slowness. Margin and collateral calls exceeded **£70bn** against roughly £37bn of gilt sales. Thirty-year yields moved **140bp in four days**, "*more than twice as large as the largest move since 2000*", against an industry stress-test standard that Cunliffe records as "*up to a 100 basis point instantaneous increase*" — calibrated, in the event, to about 70% of what happened. Resilience buffers went from

about **150bp** before the crisis to **300-400bp** rebuilt by 30 November 2022, and the Pensions Regulator set a **250bp** market-stress minimum in April 2023, cumulative with an operational buffer.

Not one of those documents mentions LPI. I searched the December 2022 FSR in full, both FPC records, both Cunliffe letters, Bank Staff Working Paper 1019, seven speeches, four TPR documents, the Work and Pensions Committee report HC 826 and all 77 written evidence submissions. **NONE FOUND** — no "LPI", no "limited price indexation", no "inflation collar", no "capped indexation". The liability is "*linked to inflation*", and that is the whole of it.

The single exception is industry, not official, and it was not picked up. Insight Investment, submission LDI0076 of 15 March 2023:

"This is because **many pension payments linked to inflation are floored at 0% inflation**, while index-linked gilts are exposed to deflation, meaning that **fewer index-linked gilts are required to hedge pension scheme liabilities as inflation expectations fall.**"

One submission of seventy-seven. It does not use the word LPI, and the Committee's report does not engage with it.

The Bank knew this in 2009 and 2012 and wrote it down twice. By 2022 it had gone from the official record entirely.

Part 6 — The hedge, step by step

Take a scheme with **£1bn** of liabilities, of which **60% is inflation-linked**: £400m LPI(0,5) on 1997–2005 service and £200m LPI(0,2.5) on post-2005 service. The remaining £400m is fixed. All figures below are my arithmetic on the model, not a valuation.

Day one

The naive hedge. The scheme sees £600m of "inflation-linked" liabilities and buys £600m of inflation exposure — index-linked gilts, RPI swaps, or an LDI mandate that combines both with leverage. This is what a 100% inflation hedge ratio means in practice and it is what most schemes report.

The correct delta hedge. Applying the model's deltas: $400 \times 0.598 = \text{£239m}$, plus $200 \times 0.196 = \text{£39m}$, for **£278m** of inflation exposure. Against £600m of nominal liability the scheme should be running about **46%**, not 100%.

The gap is £322m of unwanted long inflation — more than half the notional the scheme thinks it needs. That gap is not a hedging error in the usual sense. It is the value of the caps the scheme is short and the floors it is long, expressed as a delta.

The option position, sized. The 27bp LPI(0,5) spread on £400m and the 137bp LPI(0,2.5) spread on £200m is roughly **£1.1m and £2.7m a year** of option premium respectively, embedded, never paid, never received, never marked to a traded price. Roughly **£3.8m a year on a £1bn scheme** — and the model-risk band on it, from Part 4.1, is 2.82% of the LPI(0,5) liability alone, or about **£11m of valuation uncertainty** attributable to a single unquoted parameter.

One year on, three states of the world

State one: inflation prints 2.0%. Both collars are live and neither strike binds. The LPI(0,5) delta rises to 0.692 and the LPI(0,2.5) delta to 0.343, so the correct hedge rises to $400 \times 0.692 + 200 \times 0.343 = \text{£345m}$, from £278m. The scheme should have *bought* £67m of inflation. A scheme running a static 100% hedge is now over-hedged by £255m rather than £322m and has done nothing to cause the improvement.

State two: inflation prints 8%. The LPI(0,5) delta collapses to 0.070 and LPI(0,2.5) to 0.006. The correct hedge is $400 \times 0.070 + 200 \times 0.006 = \text{£29m}$. The scheme should be running a **95% smaller** inflation position than it was a year ago. It should have sold £249m of index-linked gilts — and it should have done so into the market of September and October 2022, in which the Bank of England found that dealers could not warehouse the basis, real yields moved 63bp in a day, and it had to build the operational capacity to buy linkers from scratch. **The rebalancing trade the model says is correct is the trade the market could not absorb.** Meanwhile the hedge is producing large mark-to-market gains against a liability that is capped, which makes the funding level look better at the precise moment the collateral position is worst.

State three: inflation prints –1%. The liability increase is **zero**, by statutory instrument. The hedge loses about 1% of its inflation-linked value. The LPI(0,5) delta falls to 0.623 and LPI(0,2.5) rises to 0.409 — the floor is now closer to the money, so the category Y liability

becomes *more* inflation-sensitive, not less. The scheme is short deflation protection it cannot buy, in a market whose 2009 experience was that "*trading in inflation floors was also affected once it became clear that RPI inflation would turn negative ... leading to losses for some market makers*". The one time the floor bound, the market for the floor stopped functioning.

What the walkthrough shows that a static description cannot: the hedge's own size is a derivative. It ranges from £29m to £345m — a factor of twelve — on a liability whose nominal size never changes. And its largest required move is in the state of the world where the market is least able to accommodate it.

Part 7 — February 2030: a benchmark amendable by statute

7.1 What was decided, and by whom

The RPI will be redefined to CPIH methods in February 2030, and the decision trail is unusually clean.

Sir David Norgrove wrote to the Chancellor on 4 March 2019 recommending that RPI publication cease and, in the interim, that its shortcomings be addressed "*by bringing the methods of the CPIH into it*". His own characterisation of the result is the most quotable line in the file: "***The effect, at least initially, would be to turn the RPI into CPIH by another name.***"

The Bank's statutory response the same day, under s.21(2) of the Statistics and Registration Service Act 2007, found the change "*would constitute a **fundamental change** ... and that it would be **materially detrimental** to the interests of the holders of relevant index-linked gilt-edged securities*", and added: "***Bank staff estimate that the proposals would be expected to reduce RPI inflation by around one percentage point on average.***"

The consultation launched at the Budget on 11 March 2020, was extended to 21 August because of Covid, and drew **831 written responses**, of which 229 addressed the gilt-impact questions including **125 from DB scheme trustees**. Respondents proposed a LIBOR-style adjustment spread — CPIH plus X — explicitly analogising to the SONIA transition. It was refused at ¶47: "***there is no statistical basis for adding an additional amount 'X' to the reformed RPI.***" **That refusal is why the change cannot be hedged: it removed the one quantity that would have been tradable.**

On 23 October 2020 the Chancellor wrote that he "*will be unable to offer his consent to the implementation of a proposal ... before the maturity of the final specific index-linked gilt in 2030*". Note the act: **he did not approve 2030; he withheld consent for anything earlier**. The date was then fixed by UKSA once the Bank confirmed, on 13 November 2020, that consent would no longer be needed.

The legal mechanism is narrower than most commentary suggests, and nothing in the statute changes in 2030. Section 21(3) requires the Chancellor's consent only where the Bank finds material detriment to holders of "***relevant index-linked gilt-edged securities***" — defined at s.21(4) as those issued before commencement with an early-redemption clause. There were three. Once the last of them, the 4½% IL 2030, is beyond its final RPI reference date of November 2029, **the protected class is empty**, so the Bank cannot find material detriment, so consent is never required. Deputy Governor Broadbent put it in writing:

"any changes to the coverage or basic calculation of the RPI from 2030 onwards - **even those judged to be fundamental** - should not have a materially detrimental impact ... either because the redemption values of the relevant ILGs will not be affected ... or because, after July 2030, **there will be no relevant ILGs outstanding.**"

February is the ONS's usual month for changes to consumer price measures, and the next one after December 2029 is February 2030. **The date is set by a one-publication-cycle gap, not by policy.**

And there will be no compensation. The published response says so twice, in identical words, at ¶12 and ¶83:

"the government will not offer compensation to the holders of index-linked gilts. **The contractual terms of all index-linked gilts state that the RPI should be used to determine the index ratio** ... There is no change to this flowing from the implementation of the Authority's reform."

Read the structure of that argument, because it is exactly what makes the change unhedgeable in contract: **the gilt references the name of the index, not its methodology.** RPI will still exist, will still be published monthly under the same statutory duty, and will still be called RPI.

7.2 The judgment

The trustees of the BT, Ford and Marks & Spencer schemes brought judicial review and lost on every ground. *R (BT Pension Scheme Trustees Ltd & ors) v (1) UK Statistics Authority (2) Chancellor of the Exchequer* [2022] EWHC 2265 (Admin), Holgate J, heard 21–22 June 2022, judgment 1 September 2022. Permission to appeal was refused and none was sought from the Court of Appeal; the judgment is final, as the DMO's own auction prospectuses now record.

The facts the court treated as **common ground** are the ones worth carrying:

"[2] It is common ground that for reasons to do with statistical methodology, **the RPI produces an estimate of inflation about one percentage point higher than the CPIH and that this will continue in the long term.**" "[3] The impact of the long-run reduction of 1% in the RPI from 2030 onwards ... is said to be **around £90 billion - £100 billion.**" "[4] There are **10.5 million** people with private sector 'defined benefit' pensions, of whom a majority have pension schemes linked to the RPI ... they will receive reduced payments amounting to **around 4 - 9% of their lifetime benefits**, with women experiencing on average a greater reduction."

And the holding that matters most to anyone running a hedge is at [174]:

"Whether going beyond that it is also a good index to be used in a commercial contract such as a lease, or **in an investment such as a gilt, are matters for the parties involved in those transactions. Hedging or matching indexed liabilities is likewise a matter for the financial institutions and any regulator involved. There is no legal requirement for the UKSA to have regard to such matters.**"

The court was pressed on the consequence that nobody is then responsible, and answered at [177]: *"this is a matter for **Parliament.**"*

The private-law claim failed too, and it is the one that closes off the last route. The claimants argued that importing CPIH methods means RPI *"ceases to be published"*, triggering the cessation clause in post-September-2005 linkers and requiring the Chancellor to select a replacement index. Counsel conceded at [220] that *"a 'fundamental change' to the RPI is insufficient to trigger the cessation clause"*. Holgate J held at [234] that *"the cessation clause **does not ask whether the RPI has been changed in any respect** ... The clause does not provide protection for any given level of return."* The disposal at [238] granted the

Chancellor a declaration "*that ... the cessation clause **will not be triggered** by the implementation of the RPI decision in or after 2030.*"

7.3 Nothing spans it

There is no instrument whose payoff is a function of the methodology change alone. This is a documented absence, not a failure to look.

- **RPI swaps do not span it.** They reference RPI as published, so the floating leg absorbs the change identically to a linker. LGIM told the Lords that RPI swap fallbacks "*in most cases ... would be towards a reference index-linked gilt*" — a fallback pointing at an unprotected instrument.
- **CPI-linked gilts do not exist.** Considered in 2011, rejected partly because CPI *might later* include owner-occupier housing. Nine years later the government resolved that uncertainty in the opposite direction, by moving RPI onto CPIH, **without revisiting the instrument decision**. There has been no DMO consultation of any kind on index-linked issuance between December 2012 and January 2026.
- **Bespoke CPI hedges were rejected by the buy side as unusable** — the consultation response records at ¶116 that "*more bespoke instruments (such as CPI swaps) were unsuitable for pension funds to invest in*".
- **No adjustment spread exists**, by explicit government decision.
- **Prospectus terms give nothing.** Post-2005 linkers have only the cessation clause, now judicially declared not to bite. Since September 2022 the DMO has added a standing disclosure paragraph — disclosure, not protection.
- **And the official estimates disagree with each other by about 50 basis points.** The Bank says ~1.0pp and the court recorded that as common ground; the OBR and the DMO work on a step from a 0.9pp RPI-CPI wedge to 0.4pp, i.e. ~0.5pp. Both are official, both are current, and there is nothing to trade against the difference. **That dispersion is the quantitative shape of the spanning problem.**


The DMO has never stated a view of its own. Its "key events in the development of the index-linked gilt market" page records events from 1980 to 2012 and then a single passive 2020 entry pointing at the consultation. **25 November 2020 is not listed at all.** The Bank, for its part, has said one thing, and it is an admission: on its yield-curve pages it notes that "*no methodological change has been implemented*" and that the February 2030 alignment "*is likely to impact the fitting of the curve around this point*". Anyone reading Bank real or implied-inflation forwards across 2030 is reading a smoothed curve through a known discontinuity the Bank has chosen not to correct.

And there is no FPC or Financial Stability Report treatment of the 2030 change at all. Against £688.5bn of affected stock, the absence is itself a finding: the Bank has treated 2030 as a statutory-compliance question and a curve-fitting nuisance, not a stability question.

7.4 What it is worth, and why the collar makes it worse

On my model, applying the OBR's ~0.5pp step:

	With reform	Without	Change
Uncapped RPI — the hedge	32.921	34.977	–5.88%
LPI(0,5) — the liability	31.348	32.504	–3.55%
LPI(0,2.5) — the liability	25.871	26.250	–1.44%



The asset falls further than the liability it hedges — by 2.3 points for LPI(0,5) and 4.4 for LPI(0,2.5) — and it does so precisely because the cap was already truncating the upside the reform removes. A capped liability cares less about a reduction in expected inflation than an uncapped asset does, because part of the inflation being removed was inflation the scheme was never going to have to pay.

So February 2030 does not simply transfer value from linker holders to the Exchequer. **It widens the LPI basis**, on a liability whose basis risk nobody in the official record has acknowledged exists. The direction is robust to the calibration; the magnitude scales with the assumed step, which is itself disputed by 50bp between two official sources.

Part 8 — What to watch

1. Whether the new IBA inflation swap benchmark ever reaches Level 1. IBA launched GBP UK RPI and EUR HICPxT swap benchmarks on **5 June 2026**, at **Level 2** — dealer-to-client prices displayed by a venue — with Level 1 promised "*when sufficient volumes are available*" and USD CPI still deferred. **The waterfall level, not the launch, is the thing to watch.** If GBP RPI migrates to Level 1, the linear market has an executable fixing for the first time and an option market above it becomes conceivable. If it sits at Level 2 indefinitely, the model-marking documented in Part 4.3 is permanent. IBA publishes quarterly waterfall-level usage free; this is a standing monitor, not a research project, and it is now the cheapest live indicator in the whole subject.

2. The February 2030 documentation sweep, and it should be run now. The judgment establishes that a contract referencing "RPI" picks up the new methodology without amendment. But the response document also confirms at ¶52 that "*only an all-items RPI index and growth rates will be published*" — **RPIX and every RPI sub-index are discontinued**, not redefined. A contract referencing RPIX faces a reference that ceases to exist, which is a materially worse problem than one that changes definition. That is a documentation audit across the existing book, not a research question, and euro HICP rebased in February 2026 for the same reason.

3. Whether anyone reconstitutes an LPI market before 2030. The market existed in 2007 with two-way pricing to fifty years. It does not now. The 2030 step gives dealers a reason to re-engage — a known, dated, one-directional change in the underlying is the kind of event that creates a book — or a reason to stay away, because the change is unhedgeable for them too. Watch for LPI quotes reappearing on broker screens; one screen settles it.

4. Whether the official sector ever names the option. The Bank wrote the problem down clearly in 2009 and 2012 and had forgotten it by 2022. The Work and Pensions Committee received the point in writing from Insight and did not use it. The IFoA convened a working party to standardise the measurement and disbanded it. If the next inflation shock is a downside one, the floor binds, the 2009 experience repeats and the market that would price it is thinner than it was then. **The absence documented in Part 5.3 is the single most repeatable finding in this note, and it is the cheapest to re-test: search any new FSR for "LPI".**

5. Whether any issuer switches the clause on. This is the cheapest thing in the note to monitor and the most consequential. The collar module is already in nineteen-plus UK programmes; electing it requires no new documentation, no negotiation and no legal work — only a number in the Final Terms. **So watch for a populated "Maximum Indexation Factor" field.** A single regulated utility electing it on a benchmark-size series would be the first genuine third-party LPI supply in sterling since 2001, and it would be visible the day the Final Terms are filed. Set a standing search on the exact phrase.

6. And the correction that is already sitting there for someone to publish. Tesco Property Finance is not an LPI bond — those bonds are fixed-rate and the RPI is swapped out inside the structure — while BT's £1bn, the largest LPI-collared sterling issue ever done, appears in none of the standard accounts. The received wisdom names two issuers, one of which is wrong and the other of which is a tenth the size of the one it omits.

Sources and reproducibility

Statute and case law. Pensions Act 1995 ss.51, 51ZA, 51A, 53, 54, 67A; Pension Schemes Act 1993 ss.84, 109 and Schedule 3; Pensions Act 2004 s.278 with S.I. 2005/275 and S.I. 2005/695; Pensions Act 2008 s.101 and Schedule 2 with S.I. 2009/82; Pensions Act 2011 s.19 and its Explanatory Notes; Statistics and Registration Service Act 2007 s.21. All read as live consolidated text on legislation.gov.uk. The Occupational Pensions (Revaluation) Orders **S.I. 2009/3267** (the 0.0% year), **S.I. 2010/2861** (the first CPI order) and S.I. 2025/1211. *Barnardo's v Buckinghamshire* [2018] UKSC 55; *Danks v QinetiQ* [2012] EWHC 570 (Ch); *Arcadia* [2014] EWHC 2683 (Ch); *Thales* [2017] EWHC 666 (Ch); *BA v APS Trustee* [2018] EWCA Civ 1533; *Britvic* [2020] EWHC 118 (Ch) and [2021] EWCA Civ 867. **R (BT Pension Scheme Trustees Ltd & ors) v UKSA & Chancellor [2022] EWHC 2265 (Admin)** — paragraphs 1–181 read directly, 182–238 recovered from the same primary source by browser extraction.

The 2030 trail. Norgrove to the Chancellor, 4 March 2019; Broadbent to Norgrove, 4 March 2019 and 13 November 2020; Norgrove to Bailey, 5 November 2020; the UKSA statement of 4 September 2019; the National Statistician's advice of 26 February 2019; the Johnson Review (January 2015); the Lords Economic Affairs Committee *Measuring Inflation*, HL Paper 246; the UKSA/HMT consultation of 11 March 2020 and the **response of 25 November 2020**; OBR *EFO* October 2024 Box 2.3 and March 2015 Box 3.3; DMO Debt Management Reports 2023–24 and 2026–27, the ultra-long consultation of 2 December 2004 and its response, and the CPI-linked gilts response of 29 November 2011.

The 2022 episode. Bank market notices of 28 September, 3 October, 10 October and **11 October 2022**; the Cunliffe letters of 5 and 18 October 2022; Hauser, *Thirteen days in October*, 4 November 2022; Breeden, 7 November 2022; Bailey at the G30, 15 October 2022; the Financial Stability Report and Financial Policy Summary and Record of December 2022; APF Quarterly Report 2022 Q4; Quarterly Bulletin 2023; **the Bank's index-linked and conventional gilt purchase results spreadsheets**, parsed to auction level; TPR statements of 12 October and 30 November 2022 and the guidance of 24 April 2023; Work and Pensions Committee HC 826 and its 77 written evidence submissions, including **Insight LDI0076**.

On LPI specifically. Bank of England *Quarterly Bulletin* **2009 Q3**, box "UK RPI inflation options"; *Quarterly Bulletin* **2012 Q3**, Smith, "Option-implied probability distributions for future inflation"; LGIM, September 2018; NatWest/*The Actuary*, 7 August 2017; BlackRock, 2024; the IFoA LPI Risk Working Party page; Fulcher et al., IFoA, June 2007; Bahaj, Czech, Ding & Reis, *The Market for Inflation Risk*; PPF *Purple Book* 2024 and 2025.

Reproducing the numbers:

```
python3 00_note/lpi_model.py      # Parts 1.4, 3.2, 3.3, 4.1, 4.2, 5.1, 5.2, 6, 7.4
python3 00_note/charts.py         # the four figures
python3 00_note/build_any.py LPI_note.md note.html LPI_research_note.pdf
```

`lpi_model.py` needs only `numpy` and no network. The full research trail is in `01_register/research_mechanism_and_ldi.md` and `05_rpi_2030/research_rpi2030.md`, each with its own document table and its own NOT FOUND list.

Caveats

The model is illustrative and it is not a pricing model. It is a one-factor Gaussian process for year-on-year RPI with a flat discount rate, a deterministic drift and a constant volatility. It has no inflation smile, no seasonality, no term structure of volatility beyond a stationary Ornstein-Uhlenbeck shape, no stochastic real rates, no correlation between inflation and discount rates, and no month-to-month indexation lag. It values a level pension rather than a real cashflow profile and takes no account of mortality. **Every number in Parts 3, 4, 5, 6 and 7.4 is arithmetic on the assumptions in PARAMS, and none of them is a market quote.**

There is no public LPI price series against which any of it can be checked. This is not a limitation I can design around; it is the subject of the note. The spreads in Part 3.2, the model-risk band in Part 4.1, the deltas in Part 5.1 and the reform impacts in Part 7.4 are mine. Where a published practitioner figure exists — LGIM's 1.5% volatility, its 10% delta difference between market-consistent and real-world methods, its 1.5–2.0% p.a. LPI risk estimate — I have used it and attributed it, but those are the only external anchors available and they come from three sell-side and consultancy documents, not from a market.

The register covers 12.5% of its own universe by value, and that is now measurable rather than merely suspected. FIRDS gives 88 GBP inflation-linked bonds worth £11.35bn; terms are established for six. The other £9.93bn is **unknown, and must not be read as uncapped** — on the one issuer where both are known the split runs roughly 1:4.4 collared to uncapped, but generalising from a single issuer is exactly the move this note otherwise declines to make. The FIRDS pull is also **ESMA-scoped**: post-Brexit UK-only listings sit in FCA FIRDS and are not included, so 88 is a floor.

The register is also biased towards understating supply, and the bias is large. It covers public, listed, sterling paper. **Privately placed collared bonds leave no public trace at all** — and the single largest deal in the register, BT's £1bn, was itself a bilateral sponsor-to-scheme placement that is visible only because it happened to be listed. The size of the unlisted population is **NONE FOUND from any source**. Separately, **no full-length base prospectus was read end to end**: the fetch tooling truncates at around 120,000 characters, so the nineteen programmes carrying the collar clause are enumerated rather than adjudicated, and any one of them may contain an elected series I did not find. Within the register itself: Dŵr Cymru **B4** is VERIFIED from its Pricing Supplement, but **A5 is REPORTED** — the issuer's capital-structure slide and two S&P reports label it "LPI" and I could not obtain its Pricing Supplement. BT's **2033 tranche is VERIFIED** and the 2039 and 2042 tranches rest on a truncated capture. Tesco Property Finance 5 and 1 are VERIFIED; **TPF 2, 3, 4 and 6 rest on stated fixed coupons and a third-party schedule**, not their own prospectuses.

Three of the load-bearing 2030 numbers are estimates that disagree. The Bank's ~1pp, the court's ~1pp recorded as common ground, and the OBR's ~0.5pp step are all official and all current. I have used the OBR's step in the model because it is the most recent and the most explicitly derived, and I have flagged the discrepancy wherever the number appears rather than picking a winner. The £90–100bn figure at [3] of the judgment is what the claimants asserted and the court recorded as common ground; it is not an independent official estimate.

The inflation prints are now verified against ONS series CZBH. Year to September 2022: **12.6%**. Year to October 2022: 14.2%. Year to September 2009: **-1.4%**. The statutory reference period ends 30 September, so the September figures are the ones that govern, and the note uses them.

"NONE FOUND" in Part 5.3 means what it says, and it is the claim most worth attacking. I searched a defined corpus — the December 2022 FSR in full, both FPC records, both Cunliffe letters, Bank Staff Working Paper 1019, seven speeches, four TPR documents, WPC HC 826 and all 77 written evidence submissions — for "LPI", "limited price indexation", "inflation collar", "capped indexation" and "caps and floors". A document outside that corpus could falsify it. Bank Staff Working Paper 1019 was not searched for every one of those terms individually. **Do not attribute standard-deviation language to the Bank about 2022; there is none in any of those documents.** And there is no Treasury Committee report titled "The 2022 gilt market crisis" — it does not exist; the parliamentary output is WPC HC 826.

The largest correction in this note is to the note itself, and it is worth stating plainly. An earlier draft asserted that Dŵr Cymru "*discloses no cap or floor terms anywhere*" and concluded from that, and from the Tesco finding, that the LPI supply story collapsed to nothing documentable. **That was wrong.** The Dŵr Cymru bonds carry an explicit LPI(0,5) collar — "*Minimum Indexation Factor: 1.0 / Maximum Indexation Factor: 1.05*" — and the Pricing Supplement even names the index "*UK Limited Price Index*". The error was one of source selection: I read the semi-annual Investor Report, which carries a principal-balance table and no terms, rather than the prospectus and pricing supplement, which carry both. **The right search key was never "LPI" but "Maximum Indexation Factor",** and using it changed the finding from "no register can be built" to a much better one: the clause is standardised, present in nineteen-plus programmes, and almost never elected. The earlier version of Part 2 is superseded, and this paragraph records why.

A verification pass was run on the load-bearing claims and produced seven further corrections. All are carried above rather than quietly fixed.

1. **The 2022 LPI fixing is 12.6%, not 14.2%.** This is the one that mattered. The statutory reference period runs to 30 September; 14.2% is the October print, which is the peak of the series and is what almost all commentary on the episode quotes. An earlier draft of this note, and the Pipeline survey that commissioned it, both used 14.2%. The over-hedge in Part 5.2 is 7.6 points, not 9.2.
2. **The 0.0%/0.0% in the 2009 order is in the table at article 2,** not in a Schedule. That Order has no Schedule; the Schedule reference belongs to the parent Act.
3. **The IBA inflation swap benchmarks launched on 5 June 2026,** at Level 2, for GBP UK RPI and EUR HICPxT, with USD CPI deferred. An earlier draft said only that the March 2026 consultation had happened.
4. **The Tesco Property Finance correction is now VERIFIED,** from the TPF 5 prospectus and the TPF 1 investor report, having started as an inference from quoted coupons.
5. The 2009 revaluation order is S.I. 2009/**3267**, not 3053 as commonly cited.
6. The first CPI revaluation order is S.I. 2010/**2861**, not 2937.
7. The Chancellor did not approve February 2030; he **withheld consent for anything earlier**, and UKSA fixed the date. Relatedly, the 250bp LDI resilience buffer is a **Pensions Regulator figure from April 2023**, not a Bank recommendation from 2022 — the Bank's

300–400bp is descriptive of what funds had rebuilt, and the FPC recommendation names no number at all.

Confirmed unchanged in the same pass: the DMO's £688.5bn and 25.2% at end-December 2025 and the \$2.9 reference; the Bank's £19.3bn / £12.1bn / £7.2bn split, verbatim; the Purple Book 2025 sentence and that it is the current edition; the 5% and 2.5% maxima at PSA 1993 Sch. 3 para. 2(6); and the judicial review's citation, date, judge and disposal.

Finally, the framing itself should be held loosely. The Pipeline survey ranked LPI as the biggest available dive on the strength of the residual, the footprint and the data. The residual and the footprint are as advertised. **The data is worse than the survey estimated**, and a reader who weighted that more heavily would call this a mechanism study with a sizing layer rather than a market note. I would not argue very hard against them.