

BRIEFING · SEPTEMBER 2026

# Artificial Intelligence in *UK Justice, Courts and Probation*

What the evidence actually says, what the regulator expects, and what can safely be done now. England and Wales.

---

**Neil Catton**

Independent Technology Strategist · Technology, Ethics & Human Purpose

NEILCATTON.COM

---

Prepared as a market briefing. Every figure is traced to a department, an inspectorate, an auditor, Parliament, a statutory regulator, published legislation, or a government register.

## A note on the numbers

*This briefing is built from what can be checked, and it names what was left out.*

Almost every published statistic about artificial intelligence in a given sector is produced by a company selling artificial intelligence into that sector. None of that material appears here. Every figure in this briefing traces to the Ministry of Justice, HM Prison and Probation Service, HM Courts and Tribunals Service, the Parole Board, HM Inspectorate of Probation, the Forensic Science Regulator, the National Audit Office, the Committee of Public Accounts, the statute book, or the government's own algorithmic transparency register.

Three habits govern the rest of the document, and they are worth stating before any number appears.

### DATES ARE STATED, BECAUSE THEY MOVE

One central source is a cohort of people who entered the community between July and December 2018, followed for two years to November 2022, and published in February 2024. It is described that way every time it is used. A figure whose date is hidden is a figure that will be wrong eventually, and usually in the direction that flatters whoever is quoting it.

### AN ABSENCE IS SEARCHED BEFORE IT IS ASSERTED

Where this briefing reports that something is not there, a search was run — and a control search was run beside it, to prove the search itself worked. The claim made is the claim that was tested. Where the register is said to return nothing for a named tool, the same register returns seven records for "probation", which is how we know the search reaches the text of a record and not merely its title.

### NAMED AND EXCLUDED

A widely repeated figure for how many people are profiled each day by a Ministry of Justice system is excluded: it reaches this research through an advocacy organisation and could not be traced to a departmental source. Subgroup accuracy figures for one sexual-offending predictor are excluded: the smallest of them rests on three recorded reoffences in a group of 153 people, and the department's own report cautions against reading subgroups of that size for a rare offence type. Both would have made this briefing louder. Neither would have made it more true.

*Every vendor claim encountered in the research was excluded, without exception, and none is used anywhere in this document.*

Sources for each figure are given at the foot of the section in which it appears, and listed in full in section thirteen.

## The system this sits inside

*Large, under pressure at both ends, and joined in the middle by a decision that is partly scored.*

### 242,000

People on the probation caseload, 2025

### 87,332

Prison population at 24 November 2025

### 80,061

Open Crown Court cases at end March 2026

The Probation Service cost £1.45bn to run in 2023-24 and supervises a caseload that has stayed broadly stable since unification. The prison estate had a published useable operational capacity of 89,334 on the same date its population stood at 87,332 – headroom of roughly two per cent. The prison population is projected to reach between 98,000 and 103,600 by March 2030, with a central estimate of 100,600. Every scenario in that projection sits above today's capacity.

The courts are moving the same way. The Crown Court open caseload rose five per cent in a year to 80,061. The median time from receipt to completion is 168 days, and for a contested jury trial it is 410 days, up fourteen per cent. In the magistrates' courts, 386,056 cases were received in the first quarter of 2026 and 380,082 were disposed of: the caseload grows even when the system is working at volume.

#### WHY THIS MATTERS TO AN AI PROGRAMME

*The projection names three drivers, and one of them is the number of people recalled to custody.*

Recall is a probation decision, informed by a scored assessment. It is the point where the technology question and the capacity question turn out to be the same question.

Ministry of Justice and HMPPS, *Prison Population Projections 2025 to 2030* · Ministry of Justice, *Criminal court statistics quarterly, January to March 2026*, 25 June 2026 · National Audit Office, *Building an effective and resilient Probation Service*, HC 1354, 24 October 2025.

## The pressure point

---

*The quality of risk assessment has more than halved, and the plan assumes technology will help close the gap.*

In 2024, probation practitioners adequately assessed risk of harm in 28 per cent of cases. In 2018-19 the same measure stood at 60 per cent. That is the inspectorate's own finding, reported by the National Audit Office and then by the Committee of Public Accounts, and it is the single most important number in this briefing.

It does not stand alone. The number of people charged with a Serious Further Offence while on probation rose 55 per cent in 2023-24 against 2021-22. Probation officers have been working at an average of 118 per cent of capacity for several years, with highs of 126 per cent in some regions. At March 2025 the officer vacancy rate was 21 per cent, with 1,479 posts unfilled, against 14 per cent in September 2021. The proportion of staff with four years' experience or less rose from 28 to 38 per cent in four years.

*The Probation Service has too few staff, with too little experience, managing too many cases.*

HM Chief Inspector of Probation, Annual Report 2026

Against that record, HMPPS plans more than thirty digital initiatives under its improvement programme, including an AI transcription tool, and targets a 25 per cent reduction in practitioner workload by 2027. In the same period the auditor records that funding constraints "meant it scaled back and delayed planned digital improvements".

This is the contradiction a senior official has to hold. The efficiency case is real, the pressure is real, and the process the technology is being introduced into is one the inspectorate says is failing on quality in two cases out of three — where the output is a judgement about risk to the public.

Committee of Public Accounts, *Efficiency and resilience of the Probation Service*, 4 February 2026 · National Audit Office, HC 1354, 24 October 2025 · HM Inspectorate of Probation, *Annual Report 2026*, 20 July 2026.

## What the sector is facing

*A dozen tools are declared. The ones that have been scoring people for two decades are not among them.*

The Algorithmic Transparency Recording Standard is mandatory across central government, including arm's-length bodies that provide public or frontline services. A record is required where a tool has "a significant influence on a decision-making process with public effect", and the policy defines that term deliberately widely: "'Significant influence' includes where an algorithmic tool meaningfully assists, supplements, or fully automates a decision-making process." There is no justice or law-enforcement carve-out, and no deadline by which a record must be published.

The justice family is not a reluctant filer. Twelve records sit on the register across the Ministry of Justice, HMPPS, the Probation Service, HM Courts and Tribunals Service, the Office of the Public Guardian and Cafcass, with two more from the Crown Prosecution Service. They cover a conversational voice assistant, an internal chatbot, facial comparison for probation check-in, semantic search across probation contact logs, record linkage, judicial transcription, a prison violence estimator and the Effective Proposal Framework used at pre-sentence and pre-release stages. HM Courts and Tribunals Service is separately piloting AI-assisted listing.

A search of that same register returns nothing for OASys, nothing for recidivism, and nothing for parole. The actuarial instruments that have informed assessments about recall, categorisation and release for two decades — and which the department's own AI plan describes as tools that "routinely support human-led assessments of reoffending risk" — do not appear on it.

### THE COMPARATOR, AND IT IS THEIRS

HMPPS has already published a thorough record for a custodial risk tool. The Violence in Prisons Estimator produces an estimate of how many violent incidents a person is at risk of being involved in over a year, and it feeds the Prison Categorisation tool. Its record names the model type — an Explainable Boosting Machine — the training data, roughly 45,000 records from 2023, the cross-validation design, a published bias and ethics assessment, an eighteen-month review cycle, and a plain statement that "a human decision-maker is involved in every process where these estimates are used".

The template for what is missing already exists inside the organisation.

Government Digital Service, *ATRS mandatory scope and exemptions policy*, 17 December 2024, updated 9 September 2026 · Algorithmic Transparency Recording Standard register, searched 20 September 2026 · HMPPS: Violence in Prisons Estimator record, 28 April 2025 · Ministry of Justice, *AI Action Plan for Justice*, 31 July 2025.

# The challenges of using AI here

---

*Four, and none of them is about the technology working.*

|    |                       |                                                                                                                                                                                                                                                                                          |
|----|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 01 | <b>Transparency</b>   | The standard catches any tool that meaningfully assists a decision about a person. Twelve justice records exist and a search returns none for the reoffending instruments. This cannot be explained as an implementation backlog when two risk tools have already been filed.            |
| 02 | <b>Fairness</b>       | The department's own revalidation reports "large miscalibration" by ethnicity on the Risk of Serious Recidivism tool, underpredicting by 1.37 to 1.68 points across 81,258 people. Nothing was published beside it — no equality assessment, no response, no statement of what follows.  |
| 03 | <b>Accountability</b> | Since 5 February 2026 the question of whether human involvement in a decision was meaningful turns partly, by statute, on how far the decision was reached by profiling. That test now applies to recall, and it applies inside a service under the workload described in section three. |
| 04 | <b>Measurement</b>    | There is no route by which an AI-related error in a risk assessment, a listing decision or a pre-sentence proposal is recorded as one. Serious Further Offence reviews exist and do not capture it. The question cannot currently be answered at all.                                    |

Underprediction, in the revalidation's own terms, means the tool estimated a lower rate of reoffending than actually occurred. That direction matters and should not be blurred: it is a public protection failure rather than a liberty failure, and the two call for different responses.

*The cohort behind those figures entered the community in 2018 and was followed to 2022. The report is dated 1 February 2024. It is the most recent published revalidation of these instruments.*

Ministry of Justice, *Revalidation: Risk of recidivism tools*, 1 February 2024 · Data Protection Act 2018 s.50A, as substituted by the Data (Use and Access) Act 2025 s.80(3).

## The threat side

*It runs at the evidence, and the capability that would have to catch it is the least accredited part of forensic science.*

The threat model does not transfer between sectors, and it does not transfer here. In insurance it was fraud against the firm. In social housing it was better-equipped claimants. In legal services it was the firm's own unverified output reaching the court. In justice it is evidential, and the numbers come from the statutory regulator's own annual report.

### 2%

Indicative compliance, technical audio operations

### 19%

Specialist video multimedia

### 33%

Digital data capture

Those are indicative levels of compliance with the statutory Code of Practice, from the Forensic Science Regulator's 2025 survey, to which 80 of 136 onboarded organisations responded. If AI-manipulated audio or video reaches a court, these are the disciplines that would have to detect it.

CHECKED TWICE, ACROSS TWO PERIODS

*The Regulator's annual report contains no mention of artificial intelligence, machine learning, deepfakes, synthetic or manipulated media.*

Nor does the newsletter published after it, in December 2025. The report covers 25 July 2024 to 24 July 2025 and was published on 29 July 2026, so the absence is evidenced for that period rather than for the Regulator's current thinking.

There is a second direction, already established in the courts and not repeated at length here: judicial guidance records that litigants in person "rarely have the skills independently to verify legal information provided by AI chatbots", and the case law on fabricated citations is settled ground.

Forensic Science Regulator annual report for 2024 to 2025, published 29 July 2026 · Forensic Science Regulator newsletter FSR-NEW-0007, December 2025 · AI Guidance for Judicial Office Holders, 31 October 2025.

# The UK legal position

*There is a rulebook. It is not an AI rulebook, and the part that matters most is a subsection nobody quotes.*

Probation and prisons process personal data for law enforcement purposes under Part 3 of the Data Protection Act 2018, not the UK GDPR route that governs most sectors. That distinction decides which provisions apply, and it is routinely missed.

The Data (Use and Access) Act 2025, at section 80(3), substituted sections 49 and 50 of the Data Protection Act 2018 with new sections 50A to 50D. Full commencement was 5 February 2026.

*When considering whether there is meaningful human involvement in the taking of a decision, a person must consider, among other things, the extent to which the decision is reached by means of profiling.*

*Data Protection Act 2018, section 50A(2)*

|             |                                                                                                                                                                                                                                     |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Section 50A | Defines a decision as based solely on automated processing where there is "no meaningful human involvement", and defines a significant decision as one producing an adverse legal effect or a similarly significant adverse effect. |
| Section 50B | A significant decision based entirely or partly on sensitive processing may not be taken solely on automated processing unless the person has given explicit consent, or the decision is required or authorised by law.             |
| Section 50C | Requires safeguards enabling the person to be informed, to make representations, to obtain human intervention, and to contest the decision. No time limits are specified in the section.                                            |
| Section 50D | Lets the Secretary of State define, by regulations subject to affirmative resolution, what meaningful human involvement means and which decisions have significantly adverse effects.                                               |

A practitioner signing a recall decision that rests substantially on a scored assessment is inside that test now. Section 50D is the one dated thing still to come: the definition can be set for the sector, or the sector can write down its own first.

Data Protection Act 2018 ss.50A–50D, as substituted by the Data (Use and Access) Act 2025 s.80(3); commenced 5 February 2026 by SI 2026/82 regulation 2(j) · Information Commissioner's Office, law enforcement guidance, updated 19 June 2026.

## Beyond the UK

*The European Union classifies this exact class of tool as high-risk. That is a comparator, not an obligation.*

Annex III of the EU Artificial Intelligence Act lists the uses treated as high-risk. Point 6(d) covers AI systems used by law enforcement authorities "for assessing the risk of a natural person offending or re-offending... or to assess personality traits and characteristics or past criminal behaviour of natural persons or groups". Point 6(e) separately covers profiling in the detection, investigation or prosecution of offences. Point 8(a) covers systems used by or for a judicial authority "in researching and interpreting facts and the law and in applying the law to a concrete set of facts".

Read together, those three points describe the reoffending instruments, the profiling that feeds them, and the court-side tools now being piloted.

### STATE THIS PLAINLY

The EU Artificial Intelligence Act does not apply to UK bodies operating domestically. Nothing in this section describes a compliance gap, and it should not be presented as one. Its value is as a reference point: another jurisdiction, looking at the same class of tool, concluded that it warranted conformity obligations rather than a voluntary record with no deadline.

## The UK side of point 8(a)

Point 8(a) is worth holding against what the release decision actually looks like here. The Parole Board conducted 30,320 parole reviews in 2025/26 and concluded 6,646 oral hearings, directing release in 3,140 cases and refusing it in 12,528. It reduced its active caseload from 9,685 to 8,000 and cut the number of cases waiting ninety days for an oral hearing by 40 per cent.

### WHAT THE BOARD ITSELF USES

The only reference to artificial intelligence anywhere in the Board's annual report is an "AI-powered Policy and Guidance tool (POLI) to support members in accessing relevant information more efficiently". There is no mention of machine learning, algorithmic risk assessment, actuarial tools or automated decision-making in parole determinations.

That is not a criticism of the Board. It is a description of where the scoring sits: upstream of the body making the decision, produced by instruments the Board does not own, and — on the evidence in section four — not on the transparency register.

Regulation (EU) 2024/1689, Annex III, points 6(d), 6(e) and 8(a) · Parole Board for England and Wales, Annual Report and Accounts 2025/26, HC 439, 14 July 2026.

# Today

---

*Seven things that can be done now, none of which waits on new law, new funding or a supplier.*

- **Inventory every tool that meaningfully assists a decision about a person**, and record who currently answers for each one. This is usually a short piece of work and almost always a surprising one.

---

- **Test each against the transparency standard's own wording**, and publish a record where it applies. The Violence in Prisons Estimator record is the internal template, and it is a good one.

---

- **Publish an equality analysis for the actuarial instruments**, or publish what was done about the miscalibration finding. The Effective Proposal Framework already carries an Equality Analysis; the instruments with a published finding against them do not.

---

- **Write down what meaningful human involvement means** for recall, categorisation and pre-sentence proposals — before regulations under section 50D define it for you.

---

- **Record overrides and escalate them**, as the Effective Proposal Framework already does, with a documented rationale and a manager alerted. An override nobody sees is not a control.

---

- **Add a field that records AI involvement** to whatever captures a Serious Further Offence review, so that the question can be answered at all.

---

- **Set the risk thresholds for digital rollout** that the Committee of Public Accounts asked for in February 2026, and the monitoring that goes with them.

---

None of these requires a decision about whether to adopt AI. They are the conditions under which that decision can be defended afterwards, and six of the seven are documentation rather than technology.

Drawn from the ATRS scope policy, Data Protection Act 2018 ss.50A-50D, the HMPPS Violence in Prisons Estimator and Probation Service Effective Proposal Framework records, and Committee of Public Accounts recommendation 4, 4 February 2026.

# Tomorrow

Three horizons, and three levels defined by where accountability sits rather than by how capable the system is.

## Three levels

|            |                                                                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Assistive  | It prepares, summarises and surfaces; the person decides. Accountability is unchanged. Transcription and judicial note-taking sit here, and so does almost all defensible value today.                                                                                                                       |
| Augmentive | It recommends within bounds; the person confirms or overrides, and the override is recorded. The Effective Proposal Framework is the worked example, in its own words: humans at the beginning making the rules, humans at the end making the decision, with overrides documented and escalated to managers. |
| Adaptive   | It decides within a defined envelope. Oversight moves to the envelope and accountability to whoever sets and monitors it. Nothing in UK justice operates here today, and this level is a reasoned projection rather than an observation.                                                                     |

## Three horizons

|                      |                                                                                                                                                                                                                    |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0 to 12 months       | The inventory, the transparency records, and a written definition of meaningful human involvement for recall and pre-sentence proposals.                                                                           |
| 12 to 36 months      | Error recording that can answer whether AI was involved. Evaluation criteria that include the effect on the person, not only administrative effort. Regulations under section 50D are likely to land in this band. |
| 36 months and beyond | Whether the actuarial instruments are revalidated, replaced or retired, and on whose published evidence. Nothing currently names a successor to OASys.                                                             |

*The third band is a projection. It depends on funding decisions the auditor has already recorded as constraining the digital programme, and should be read as a direction of travel rather than a plan.*

The Probation Service: Effective Proposal Framework record, 10 February 2025, updated 2 September 2026 · National Audit Office, HC 1354, 24 October 2025.

# Benefit realisation

*A worked example from the sector's own record, and the measure that is missing from it.*

In March 2026 a minister set out, in a written answer, how HM Courts and Tribunals Service will judge its AI-assisted listing pilot. The criteria are "reductions in administrative effort associated with listing; improved timeliness and consistency of listing activity; the impact on listing accuracy and associated rework; and the effect on effective use of court capacity", together with feedback from judges and operational staff on usability and confidence.

It is a serious list, and every item on it is operational. None of it measures the effect on the person being listed.

|                                                                                                                                                    |                                                                                                                                                                             |
|----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>WHAT IS MEASURED</p> <p>Administrative effort. Timeliness. Consistency. Accuracy and rework. Court capacity. Staff and judicial confidence.</p> | <p>WHAT IS NOT</p> <p>Time on remand. Distance travelled to court. Witness and victim availability. Whether any error runs consistently in one direction, and for whom.</p> |
|----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Listing decides when a person's trial happens. A benefit case that measures only effort will report a success whatever happens to the people the decisions are about — and it will not survive the first committee session that asks the obvious question.

**The practical step is small.** Add one measure of the effect on the person to every AI business case in the portfolio. It costs little, it is usually available from data already held, and it converts a productivity claim into something defensible.

Hansard, written answer of 9 March 2026 from the Minister of State at the Ministry of Justice to a question of 27 February 2026.

# A way in

*Eight workstreams. The first three form a natural first step.*

|    |                                            |                                                                                                              |                   |
|----|--------------------------------------------|--------------------------------------------------------------------------------------------------------------|-------------------|
| 01 | <b>Readiness assessment</b>                | What is running, who answers for it, and which tools fall inside the transparency standard's own definition. | 4–6 weeks         |
| 02 | <b>Use case identification and triage</b>  | Sorted by where accountability sits, not by expected saving.                                                 | 3–4 weeks         |
| 03 | <b>Regulatory and legal position paper</b> | Part 3, sections 50A to 50D, and a written definition of meaningful human involvement for this estate.       | 3–4 weeks         |
| 04 | <b>Board and executive governance</b>      | Who approves, who reviews, and on what evidence.                                                             | 4 weeks           |
| 05 | <b>Threat and integrity response</b>       | Evidential integrity, and the forensic capability gap described in section six.                              | 4 weeks           |
| 06 | <b>Benefit realisation framework</b>       | Including the effect-on-the-person measure absent from current criteria.                                     | 3 weeks, parallel |
| 07 | <b>Build, buy or partner evaluation</b>    | Tested against the transparency and equality obligations, not only cost.                                     | 4–6 weeks         |
| 08 | <b>Twelve-month mobilisation roadmap</b>   | Sequenced against the section 50D timetable.                                                                 | 2 weeks           |

## THE FIRST STEP

*Workstreams one to three form a natural first engagement of six to eight weeks.*

Six to eight weeks also fits inside the window before regulations under section 50D define meaningful human involvement — which is the difference between describing a practice you already have and being told what it should have been.

nc@neilcatton.com · neilcatton.com

## Sources

*Every figure in this briefing, traceable to a named publication and a date.*

- Ministry of Justice, *AI Action Plan for Justice*, published 31 July 2025. Foreword by the Minister for Prisons, Probation and Reducing Reoffending.
- Ministry of Justice, *Revalidation: Risk of recidivism tools: an evaluation of the actuarial instruments developed to assess recidivism risk in England and Wales*, 1 February 2024.
- Ministry of Justice and HMPPS, *Prison Population Projections 2025 to 2030, England and Wales*.
- Ministry of Justice, *Criminal court statistics quarterly, January to March 2026*, 25 June 2026.
- National Audit Office, *Building an effective and resilient Probation Service*, HC 1354, 24 October 2025.
- Committee of Public Accounts, *Efficiency and resilience of the Probation Service*, Sixty-Fifth Report of Session 2024–26, 4 February 2026.
- Committee of Public Accounts, *Use of AI in Government*, HC 356, Session 2024–25, 26 March 2025.
- HM Inspectorate of Probation, *Annual Report 2026*, 20 July 2026.
- Parole Board for England and Wales, *Annual Report and Accounts 2025/26*, HC 439, 14 July 2026.
- Forensic Science Regulator, *Annual report for 2024 to 2025*, published 29 July 2026, covering 25 July 2024 to 24 July 2025; and newsletter FSR-NEW-0007, December 2025.
- Government Digital Service, *Algorithmic Transparency Recording Standard: mandatory scope and exemptions policy*, 17 December 2024, updated 9 September 2026; and the ATRS register, searched 20 September 2026.
- HMPPS: Violence in Prisons Estimator, algorithmic transparency record, 28 April 2025, updated 9 September 2026.
- The Probation Service: Effective Proposal Framework, algorithmic transparency record, 10 February 2025, updated 2 September 2026.
- Data Protection Act 2018, Part 3, sections 50A to 50D, as substituted by the Data (Use and Access) Act 2025, section 80(3); commenced 5 February 2026 by SI 2026/82 regulation 2(j).
- Information Commissioner's Office, guidance on the Data (Use and Access) Act 2025 for law enforcement agencies, updated 19 June 2026.
- Hansard, written answer of 9 March 2026, Minister of State at the Ministry of Justice.
- *AI Guidance for Judicial Office Holders*, 31 October 2025.
- Regulation (EU) 2024/1689 (the EU Artificial Intelligence Act), Annex III.

*Excluded on the record, and named here because the exclusions are part of the method: a widely repeated daily profiling figure with no departmental source; subgroup accuracy figures resting on three recorded reoffences; all vendor material.*

## About the author

---

Neil Catton is an independent technology strategist with more than thirty-nine years of technology leadership across twenty-plus sectors — from central government, policing and national health systems to banking, manufacturing and insurance. He has held CTO and Director-level roles at BT, Wipro, Fujitsu, Sopra Steria and Capita, and works across a fractional portfolio spanning public safety, public sector, insurance, health technology and early-stage startups.

He was Chief Technology Officer for Home Affairs and Criminal Justice at Fujitsu, and has worked on custody data and policing technology with UK forces and with the Police Service of Northern Ireland. He is the author of the technology ethics trilogy *The Next Evolution*, *The Cognitive Crucible* and *The Shadow System*, and writes at [neilcatton.substack.com](https://neilcatton.substack.com).

### How this briefing was produced

Only primary sources are used: a department, an inspectorate, an auditor, Parliament, a statutory regulator, published legislation, or a government register. Where this briefing reports an absence, that absence was searched for and a control search was run beside it to establish that the search itself worked. Figures that could not be traced to a named source are named as excluded rather than quietly dropped, because what a briefing leaves out is the better test of it.

In keeping with this report's own argument: this briefing was produced with AI assistance, used transparently and kept under direct human editorial control throughout. Every source was verified; every judgement was reviewed by Neil Catton before publication.

#### NOTICE

This briefing is not legal, regulatory, financial or professional advice. The position is stated as at 20 September 2026 and may since have moved; readers should check the current position before relying on anything here. Third-party sources are cited in good faith and should be verified before reliance. Named organisations appear only from regulator, inspectorate, auditor, parliamentary or court findings and are used illustratively; none has been engaged or approached in connection with this briefing. No advisory or client relationship is created by receipt of this document.

**Independence.** The author holds no vendor affiliations, reseller agreements, referral arrangements, commission interests or equity positions in anything referenced in this briefing, and has received no payment from any organisation named in it. That independence is what makes the exclusion discipline described above credible rather than merely stated.

© Neil Catton 2026. Internal use and circulation permitted with attribution.