BRIEFING · SEPTEMBER 2026

Artificial Intelligence in *UK Banking and Consumer Lending*

What the evidence actually says, what the regulators expect, and what can safely be done before the guidance lands.

Neil Catton

Independent Technology Strategist · Technology, Ethics & Human Purpose

NEILCATTON.COM

Prepared as a market briefing. Every figure is traced to a regulator, an ombudsman, a select committee, published legislation, or a firm speaking on its own published record.

ONE

A note on the numbers

Almost every published figure about AI in banking was produced by someone selling AI into banking. This briefing leaves those out, and says which ones.

An executive reading a briefing of this kind will check one figure. If it does not hold, the rest of the document dies with it. So the rule applied throughout is narrow: a number appears here only if it can be traced to a named regulator, an ombudsman, a select committee, published legislation, or a firm speaking on its own published record about its own operations. Supplier accuracy claims, conference statistics and adoption percentages of unknown provenance are absent.

Naming what was excluded is more useful than any statistic that survived, so here is the list.

EXCLUDED — AND WHY

"AI models give inaccurate financial advice 57 per cent of the time." In circulation in the trade press in September 2026 and attributed to a financial-technology company's own research. A vendor measuring the failure rate of the thing it sells against. Excluded.

Every consultancy adoption statistic encountered — proportions of banks deploying, scaling or piloting AI. None traceable to a regulator, a trade body's own data collection, or a firm's own disclosure. None used.

Agentic payments and agentic finance market forecasts. Uniformly supplier-produced, uniformly enormous, uniformly unsourced.

Bank headcount reductions attributed to AI. Widely reported for several large UK banks, and in every case traced back to press reporting of unnamed sources rather than to a disclosure. Where a bank has said it in its own annual report it would be usable; nothing was found at that standard.

ISO/IEC 42001 certification claims for named UK banks. Everything returned was certification-vendor marketing. Excluded, and the reason is set out in section ten.

The figures a reader is most likely to check are sourced as narrowly as possible. The adoption numbers come from the Bank of England and the Financial Conduct Authority's own survey of the sector, and its publication date is attached every time one of them appears. The consumer numbers come from the FCA's own research and its own Financial Lives survey, with the fieldwork dates stated. The complaints numbers come from the Financial Ombudsman Service, and the ombudsman's own explanation of why they moved is given alongside them rather than left for the reader to supply.

Firm-level figures are handled the same way. Seven large UK lenders were read directly. Five were read at annual-report standard and two could not be, and the document says which is which wherever a firm is named. One figure is kept with a caveat attached rather than excluded: Lloyds Banking Group's statement that generative AI "delivered around £50 million of value" in 2025 with "more than £100 million in additional value expected this year". A firm speaking on its own record is admissible under the rule above. The release does not say how the value is calculated, and that sentence appears every time the figure does.

One source is used with its commercial interest named. UK Finance's Annual Fraud Report 2026 contains the sector's most quotable passage on AI-enabled fraud — criminals "deploying deepfakes during account openings", synthetic identities, "selfie injections". That passage appears in a foreword authored by named executives of a fraud-detection vendor, in partnership with the trade body. UK Finance's own fraud figures are used here without qualification. The AI commentary is used only with the sponsorship stated in the same sentence, which is more useful to a reader than quoting it silently.

Position stated as at 15 September 2026. Several matters described here were unresolved at that date and are identified as such in the text.

TWO

The wider market

The sector has adopted AI faster than it has understood it, and it told the regulator so.

In the most recent joint survey by the Bank of England and the Financial Conduct Authority, 75 per cent of firms reported already using AI, with a further 10 per cent planning to within three years. In 2022 the same survey found 58 per cent using and 14 per cent planning. The median number of use cases per firm is expected to more than double, from nine to twenty-one, over three years. Foundation models already account for 17 per cent of deployments. The survey covered 118 respondents across six sectors of financial services.

Two findings inside that survey matter more than the adoption headline. The first is that 46 per cent of firms reported only **partial understanding** of the AI systems they were using, against 34 per cent claiming complete understanding. The second is that a third of all use cases are now third-party implementations, up from 17 per cent in 2022 — and understanding was lower where the implementation was bought in rather than built.

75%

Of firms already using AI, with a further 10 per cent planning to within three years. Up from 58 per cent in 2022.

46%

Reported only partial understanding of the AI systems they were using. 34 per cent claimed complete understanding.

1 in 3

Use cases are third-party implementations, up from 17 per cent. Understanding is lower where they are.

The same survey measured the supply side, and this is the strongest concentration evidence available anywhere in this series of briefings. The top three providers account for 73 per cent of cloud services, 44 per cent of models and 33 per cent of data services. A sector that has moved a third of its AI estate onto bought-in implementations has, at the same time, moved most of the infrastructure beneath it onto a handful of firms.

The date of this survey is itself a finding and is stated here rather than buried. It was published on **21 November 2024**. No later edition exists. It was still the authoritative adoption evidence cited by a select committee fourteen months later, and it remains the best measurement the sector has of its own AI estate. Every figure in this section is used with that date attached.

WHAT THIS MEANS FOR A BOARD

The most-quoted adoption figure in UK financial services is nearly two years old, and there is nothing newer at the same standard. Any internal paper that presents the 75 per cent as current is presenting a 2024 measurement of a market that has since doubled its median use cases. The honest position is that nobody — including the regulators — currently knows what the sector is running.

Bank of England and Financial Conduct Authority, *Artificial intelligence in UK financial services — 2024*, published 21 November 2024. 118 respondents across six sectors.

THREE

The pressure point

The rulebook is not being rewritten. The accountability has already moved.

The FCA's published position could not be plainer. "We do not plan to introduce extra regulations for AI. Instead, we'll rely on existing frameworks, which mitigate many of the risks associated with AI." The approach is described as "principles-based and focused on outcomes" because "with a fast-moving technology like AI, this is the best way to support UK growth and competitiveness." On the same page: "Our rules emphasise accountability for senior managers and are relevant to the safe use of AI."

Existing frameworks means two things: the Consumer Duty and the Senior Managers and Certification Regime. Neither was written with machine learning in mind, and neither is going to be amended for it.

The Treasury Committee arrived at the same place from the other direction. In January 2026 it concluded that "the Financial Conduct Authority, the Bank of England and HM Treasury are not doing enough to manage the risks presented by AI". It found that AI-driven decision-making in credit and insurance "lacks transparency", that AI decision-making and product tailoring "threaten financial exclusion for the most disadvantaged consumers", and that the lack of explainability of AI models conflicts directly with senior managers' requirement to demonstrate understanding of the risks in their area.

All three bodies rejected that characterisation in April 2026. The Bank of England was explicit: "We do not agree with the TSC's characterisation that the Bank is taking a 'wait and see' approach." The FCA told the Committee that firms understand "the complexity of AI does not diminish their accountability for its use", and accepted the recommendation to publish comprehensive guidance on applying existing rules to AI — including accountability under the Senior Managers regime — **by the end of 2026**.

> *A named executive already carries the accountability for systems the sector has told the regulator it only partly understands, under a rulebook that is not going to be rewritten, with the clarifying guidance arriving inside this financial year.*

That is the whole argument of this briefing and it does not require a single supplier statistic to make it. It is not a technology argument and it is not a compliance argument. It is a personal one. What a senior manager does in the next twelve weeks determines whether the guidance, when it lands, describes something they already do.

There is a second contrast worth holding alongside the first, because both parties to it are the reader's own government. Six months after a select committee said the regulators were not doing enough, HM Treasury published a *Financial Services AI Adoption Plan*, on 14 July 2026, opening with "Scaling AI is a strategic imperative for UK growth, competitiveness, and financial resilience." It commits the regulators to maintaining the principles-based, outcomes-focused approach, requires the FCA to review consumer impacts from general-purpose AI generating financial advice, and flags concentration risk on "a small number of global AI and cloud providers". It contains no numerical targets and no completion deadlines.

FCA, *AI and the FCA: our approach*, last updated 13 February 2026 · Treasury Committee, *Artificial Intelligence in Financial Services*, HC 684, 20 January 2026, paras 7, 12 and 21 · *Responses to the Committee's Fifteenth report*, HC 1791, 16 April 2026 · HM Treasury, *Financial Services AI Adoption Plan*, 14 July 2026.

FOUR

What the sector is facing

Two pressures. One is the customers the sector is not reaching. The other is an exercise already running, at a volume no one is doing by hand.

The gap between difficulty and support

The FCA's Financial Lives survey, with fieldwork between February and June 2024 across 17,950 adults, found 4.5 million adults — 8 per cent — in financial difficulty, having missed bills or credit commitments in three or more of the previous six months. A further 13.1 million, 24 per cent, had low financial resilience, unchanged from 2022. Some 26.4 million adults, 49 per cent, showed one or more characteristics of vulnerability.

Set against that, only **1.6 million adults, 3 per cent, received any support from a mortgage or credit lender over two years**. Where it was given, it worked: 43 per cent of credit holders who received support reported an improved financial situation as a result. Separately, 1.7 million adults used debt advice within twelve months.

The arithmetic is the argument. The support works, and it reaches about a third as many people as need it, over twice the measurement period. That gap is the strongest human case for AI in collections and forbearance in this sector, and it is made by the regulator rather than by a supplier. It is also the reason the same technology has to be got right, because the decision AI is being put inside is an affordability decline and an arrears call.

The exercise that is already running

In March 2026 the FCA published Policy Statement PS26/3, establishing a consumer redress scheme for motor finance commission. It covers two periods — 6 April 2007 to 31 March 2014, and 1 April 2014 to 1 November 2024 — with an estimated **12.1 million agreements eligible**, estimated 75 per cent uptake, £7.5bn of redress, £1.6bn of non-redress costs and £9.1bn of total cost to firms. Average redress is around £830 per agreement.

The policy statement contains no guidance on automation, technology or bulk processing. It describes a manual sequence: firms contact consumers, consumers respond, firms assess liability, payments are made. That sequence is specified for an exercise covering 12.1 million agreements. No firm is performing 12.1 million eligibility assessments by hand, and the rules governing the exercise are silent on the machinery that will actually do it.

> STATUS AT 15 SEPTEMBER 2026 — READ THIS BEFORE CITING THE SCHEME
>
> On 2 July 2026 the FCA announced that the scheme had been **partially suspended** following legal challenges brought by Consumer Voice (represented by Courmacs Legal), Volkswagen Financial Services, Mercedes-Benz Financial Services and Crédit Agricole Auto Finance. Firms are not required to calculate or pay redress, or to send communications about compensation owed, on the original timetable.
>
> What continues is the rest: identifying relevant complaints and agreements, gathering commission and disclosure data, notifying complainants who are *not* owed compensation, responding to non-scheme elements of complaints, and cooperating with brokers and the Financial Ombudsman. Notification of non-compensated complainants is due by **18 November 2026** for agreements from 1 April 2014, and 18 January 2027 for earlier ones. The Upper Tribunal will hear the challenges on 14–18 December 2026 or 16–26 February 2027. Payments, if the scheme survives, are expected in 2027.

The suspension makes the point sharper rather than softer, and the reason matters. The parts that have been suspended are the paying parts. The parts still running are identification, data gathering and notification across 12.1 million agreements — which is to say the mechanised parts. A firm is still required to work out who is in scope, and to write to the people who are owed nothing, by a date in November. The rulebook governing that work remains silent on how it should be processed.

FCA, *Financial Lives 2024 survey key findings*, published 16 May 2025, fieldwork 5 February to 16 June 2024, 17,950 adults · FCA, *Motor Finance Consumer Redress Scheme*, PS26/3, March 2026 · FCA statement, *Motor finance scheme partially suspended*, 2 July 2026.

FIVE

The record the technology is being fitted to

The sector's headline conduct measure improved dramatically in a single year. The ombudsman gives two reasons for it, and neither is that the harm fell.

The Financial Ombudsman Service published its annual complaints data for 2025/26 on 21 May 2026. New complaints fell to **214,600, down 30 per cent from 305,700**. The uphold rate was 30 per cent. Motor hire purchase complaints fell to 37,700 from 76,200. Current account complaints were 32,900, of which 18,900 concerned fraud and scams. Credit card complaints were 22,800, including 8,800 about irresponsible or unaffordable lending — down from 49,400 the previous year.

Cases brought by professional representatives fell to 18 per cent of the caseload, 38,600 cases, from 50 per cent and 152,800. Withdrawal and abandonment fell from 35 per cent to 18 per cent.

214,600

New complaints in 2025/26, down 30 per cent from 305,700. Uphold rate 30 per cent.

38,600

Cases from professional representatives — 18 per cent of the caseload, down from 152,800 and 50 per cent.

8,800

Credit card complaints about irresponsible or unaffordable lending, down from 49,400.

Alongside that fall sit two deliberate acts, and the ombudsman names both. Its *Plans and Budget 2026/27* of March 2026 attributes the change to the "introduction of a case fee for representatives… and the announcement of the FCA redress scheme for motor finance commission (MFC) cases". The fee was introduced, in the ombudsman's own words, to "encourage professional representatives to submit better-evidenced complaints, considering their merits more diligently before referring them to us." Professionally represented cases are expected to fall again, to 13,500 in 2026/27.

> WHY THIS BELONGS IN AN AI BRIEFING
>
> Complaint volume is the most common proxy for consumer detriment in this sector, and it is the series a firm is most likely to reach for when it wants to know whether an automated decision is working. In a single year that series moved 30 per cent, for reasons the ombudsman attributes to a fee change and a scheme announcement.
>
> A model trained on 2024/25 against 2025/26 complaint volumes would learn something false about how customers are being treated. So would a dashboard, a board pack, or a benefit case. The volume moved. Whether the harm moved is a separate question, and this data does not answer it.

Nothing here is a criticism of the ombudsman, which has published the explanation in its own documents and made no claim beyond it. The point is what happens downstream when a number this clean is picked up by something that does not read the footnote.

Financial Ombudsman Service, annual complaints data 2025/26, published 21 May 2026 · Financial Ombudsman Service, *Plans and Budget 2026/27*, March 2026.

SIX

What the regulator asked, and what the sector answered

Two things on the public record, about the same firms, that do not sit together.

Under the Senior Managers and Certification Regime a named individual is accountable for the area they run, and must be able to demonstrate that they understand its risks. That is the whole mechanism. It does not have a technology carve-out and it is not being given one.

The FCA told the Treasury Committee in April 2026 that firms understand “the complexity of AI does not diminish their accountability for its use”. The FCA’s own review of AI and the future of retail financial services, published in July 2026, put the question to firms directly and reported the answer: “No firm argued that the SMR accountability model should change… We agree that the SMR still applies.”

Set against that, the sector’s own answer to the regulators’ own survey: 46 per cent reported only partial understanding of the AI systems they were using, and understanding was lower still where the implementation was third-party — which a third of all use cases now are.

THE DUTY

A named individual, accountable for their area, required to be able to demonstrate that they understand the risks in it. Applied to AI without amendment, and confirmed twice in 2026 by the FCA.

THE ANSWER

46 per cent partial understanding. 34 per cent complete. Lower again where the system was bought rather than built, across a third of the estate.

The Treasury Committee named the same problem from the other side, finding that the lack of explainability of AI models conflicts directly with senior managers’ requirement to demonstrate understanding of risks. That is not a drafting quibble. It is a structural collision between a rulebook written for human decisions and a technology whose defining property is that the reasoning is not directly inspectable.

There is one uncomfortable turn in the evidence that a briefing of this kind should not leave out, because it cuts against the obvious remedy. The FCA’s own research note on explaining AI’s role in credit decisions, published in February 2025, found that more information “may not always be helpful for decision-making and could lead to worse outcomes for consumers”. Explanation is the answer to the senior manager’s problem. It is not automatically the answer to the customer’s. Those are two different obligations and they are being solved as though they were one.

Bank of England and FCA AI survey, 21 November 2024 · Treasury Committee HC 684, para 12 · FCA, *AI and the future of retail financial services* (the Mills Review), 6 July 2026 · FCA research note, *Credit where credit is due*, 24 February 2025.

SEVEN

Where AI actually is

This is the only sector in this series where the regulator has published how far up the decision ladder the technology has gone.

The joint survey measured decision authority directly. Some 55 per cent of use cases involve a degree of automation; 24 per cent are semi-autonomous, with human oversight retained over the critical decisions only; 2 per cent are fully autonomous. No other sector in this series has a published distribution of that kind, and it tells a board more than any adoption headline, because it says what the technology is allowed to decide rather than where it is installed.

55%

Of use cases involve some degree of automation. Foundation models account for 17 per cent of deployments.

24%

Are semi-autonomous, with human oversight retained over the critical decisions only.

2%

Are fully autonomous. Median use cases per firm expected to rise from nine to twenty-one in three years.

Testing live, with the regulator

The FCA's AI Live Testing second cohort is Barclays, Experian, Lloyds Banking Group's Scottish Widows, UBS, Aereve, Coadjute, GoCardless and Palindrome — covering investment guidance, consumer credit score information, agentic payments and financial crime, across agentic AI, small language models and neurosymbolic systems. Testing concludes at the end of 2026, with an evaluation report in early 2027. The Supercharged Sandbox ran 22 firms in January 2026, including Principality, Shawbrook and Virgin Money.

What seven lenders say in their own documents

Seven large UK lenders were read for this briefing. Five at annual-report standard — NatWest, HSBC, Santander UK, Nationwide and TSB. Two, Monzo and Starling, could not be, and the genre of each document is named below, because comparing an annual report against a marketing page is a false comparison. Section fifteen returns to what the comparison actually shows.

NatWest Group	More than 12,000 engineers with AI coding assistants; over 20,000 hours of AI learning in 2025. Names the Cora virtual assistant, and partnerships with AWS, OpenAI and Google. The only one to publish savings figures.
HSBC Holdings	"Over 100 GenAI solutions in use"; the HSBC Productivity Suite available to around 85 per cent of colleagues; 31,000+ engineers on an AI coding assistant. Adoption quantified, no monetary value. Group-level, not the UK ring-fenced bank.
Santander UK	"Enterprise solutions like ChatGPT" across the business; AI in financial crime with "a 50% alert reduction in some use cases". Names execution risk in its own risk section.
Nationwide	Names the most use cases of the seven and quantifies none: fraud detection, customer communications, quality assurance, app personalisation, chatbot. Microsoft Copilot rolled out.
TSB	Claims no deployment. Its only AI language sits in the emerging risks section: TSB "has started to investigate opportunities to safely utilise AI", with guardrails for its "safe exploration".
Monzo	Annual report not obtained. Its own engineering writing describes fraud detection by deep learning, credit risk models, and an LLM agent handling 150+ customer intents with "a 10 percentage point improvement in resolution for transaction queries".
Starling	Annual report not obtained. Its results release names Spending Intelligence, Scam Intelligence and the Starling Assistant, and claims the UK's first agentic AI financial assistant. Unquantified.

BoE and FCA AI survey 2024 · FCA AI Live Testing cohort two · the five annual reports named above, read directly.

The challenges of using AI here

Four risks, each one named by a regulator, a committee or the sector's own survey rather than inferred.

One — concentration, now designated, but only halfway

The survey's supply-side figures are the starting point: the top three providers hold 73 per cent of cloud, 44 per cent of models and 33 per cent of data services, against a third of use cases now being third-party implementations.

The regime built to answer this has now been switched on. On 10 July 2026 HM Treasury designated four entities as critical third parties, with oversight by the Bank of England, the PRA and the FCA beginning on 13 July: **Amazon Web Services EMEA SARL, Google Cloud EMEA Limited, Microsoft Ireland Operations Ltd and Oracle Corporation UK Limited**. Sarah Breeden, Deputy Governor for Financial Stability: "As critical third parties become increasingly embedded in the operations of financial institutions, they can introduce new forms of systemic risk."

All four are cloud providers. **No AI model provider was designated.** Four days after the designations, HM Treasury's own adoption plan flagged concentration risk on "a small number of global AI and cloud providers". The regime has been applied to the tier that was already the most concentrated, at 73 per cent, and not to the tier that sits closest to the decision. That is a live gap rather than a criticism of the designations, and it is the one a firm can map for itself today.

Two — explainability against the Senior Managers Regime

Set out in section six. The short form: a named individual must demonstrate understanding of risks in systems that 46 per cent of the sector says it only partly understands, and the most obvious remedy — more explanation to the customer — is one the FCA's own research has found can make consumer outcomes worse.

Three — exclusion in credit decisions

The Treasury Committee found that AI-driven decision-making in credit and insurance "lacks transparency", and that AI decision-making and product tailoring "threaten financial exclusion for the most disadvantaged consumers". The FCA's own review records around 900,000 adults in the UK as unbanked.

Underneath that sits a data problem the FCA has already measured. Its Credit Information Market Study found material differences in the data held on individuals by the three large credit reference agencies: only 30 per cent of matched individuals had consistent default information across all three. Data quality, in the FCA's words, "is likely important in preventing poor outcomes for consumers". A model is a function of its inputs, and the inputs disagree with each other.

Four — redress at volume, with no processing rulebook

Set out in section four. Identification and notification across 12.1 million agreements continue while payment is suspended, and the rules governing the exercise say nothing about automation or bulk processing.

Bank of England, *UK financial regulators to begin overseeing Critical Third Parties announced by HMT*, July 2026 · Treasury Committee HC 684, para 7 · FCA, *Credit Information Market Study* final report MS19/1.3, December 2023 · FCA Mills Review, 6 July 2026.

NINE

The threat side

The threat here is not fraud against the bank, though fraud is large and measured. It is advice-shaped output reaching customers from outside the regulatory perimeter, and customers believing it is protected.

The FCA surveyed 666 UK investors aged 18 to 40 with fieldwork on 24 July 2026. Four in five less experienced investors have used AI for investment help. Two in three use it occasionally or regularly. Some 56 per cent trust AI tools — ahead of television and radio at 47 per cent, the press at 46 per cent and social media influencers at 29 per cent.

44%

Mistakenly believe AI-generated financial information is regulated.

32%

Wrongly expect FSCS or Financial Ombudsman compensation if AI advice fails them. 38 per cent think investment decisions can be made on AI outputs alone.

Lucy Castledine, the FCA's Director of Consumer Investments: "Unlike regulated financial advice, AI-generated tips from general-purpose chatbots mean you are not covered if things go wrong."

> *A customer acts on unregulated output, loses money, and complains to the firm they hold an account with. The firm did nothing wrong and has no defence to offer, because the customer does not distinguish between the two sources.*

That is a conduct and complaints problem arriving from outside the perimeter, into a firm that cannot control the source and cannot see the volume. The FCA's own review names it as the first thing to fix. Its opening recommendation is to "secure and adapt the regulatory perimeter", and it asks whether AI systems "could provide services functionally equivalent to regulated activities such as advice or intermediation, while remaining outside the regulatory perimeter" — while also capturing "significant value from financial firms" from the same position.

The same review measured how far customers will go. Around 67 per cent use AI at all; 16 per cent for financial services; 24 per cent of those upload personal financial data to it. Twenty per cent said they would use autonomous AI acting within pre-set financial goals, rising to 28 per cent among existing AI users — the finding usually reported as "one-fifth of UK adults, around 11 million people". By comparison, 9 per cent use traditional financial advice.

Fraud, measured, and the passage that has to carry its sponsorship

UK Finance's Annual Fraud Report 2026, covering 2025, puts total losses at £1.28bn, up 4 per cent, across 4.06 million cases, up 11 per cent. Authorised push payment fraud accounted for £576.4m, up 19 per cent, across 248,070 cases. Unauthorised fraud was £703.4m, down 5 per cent. £1.68bn was blocked — 70 pence prevented for every £1 attempted. In the first year of the reimbursement regime, 88 per cent of in-scope APP losses were reimbursed. Some 66 per cent of APP fraud originated online and 17 per cent via telecoms.

The same report's commentary on AI and deepfakes — criminals using AI to scale and refine attacks, deploying deepfakes during account openings, synthetic identities and "selfie injections" — appears in a foreword authored by named executives of a fraud-detection vendor, published in partnership with the trade body. It is the sector's most quotable AI-fraud passage and it was written by a firm that sells the remedy. It is repeated here on that basis and on no other.

There is one AI-and-security finding in this sector that carries no commercial interest at all. On 15 May 2026 the Bank of England, the FCA and HM Treasury issued a joint statement on frontier AI models and cyber resilience: "The cyber capabilities of current frontier AI models are already exceeding what a skilled practitioner could achieve, and at a significantly higher speed, greater scale, and lower cost." The statement carries an expectation that boards understand the exposure. That sentence was written by the reader's own regulators and is the one to take to a board.

FCA research, fieldwork 24 July 2026, 666 UK investors aged 18–40 · FCA Mills Review, 6 July 2026, research by Yonder Consulting, 5,026 UK adults · UK Finance, *Annual Fraud Report 2026* · Bank of England, FCA and HM Treasury joint statement, 15 May 2026.

The UK legal position

There is no AI rulebook for financial services and there is not going to be one. There are obligations that already bite.

Consumer Duty	FCA	Requires good outcomes for retail customers. Applies to an AI-assisted decision exactly as it applies to a human one. No technology carve-out.
Senior Managers and Certification Regime	FCA and PRA	A named individual accountable for their area, required to demonstrate understanding of its risks. Confirmed twice in 2026 as applying unchanged to AI.
Data (Use and Access) Act 2025, s.80	In force 5 February 2026	Replaces Article 22 UK GDPR with Articles 22A–22D. Section eleven sets this out in full; it is the change that matters most to automated lending.
Critical Third Parties regime	HM Treasury, Bank of England, FCA	Four cloud providers designated 10 July 2026, oversight from 13 July. No AI model provider designated.
FCA guidance on applying existing rules to AI	Due by end of 2026	Accepted in April 2026, including accountability under the Senior Managers regime. Not published and not trailed as at 15 September 2026.
ICO statutory code on AI and automated decision-making	SI 2026/425, in force 12 May 2026	The Commissioner must prepare it. The regulations set no deadline for producing it.
FCA AI good and poor practice publication	Expected during 2026	The AI Input Zone was reopened 14 May to 19 June 2026 to gather evidence for it. Not yet published.

Bodies checked, and what was found

The **National Audit Office** has not examined AI in financial services. Its report *Financial services regulation: Adapting to change*, HC 313, 8 December 2023, contains one sentence on the subject: AI provides “opportunities for innovation but also risks to businesses and consumers which the FCA must plan for.” That is the whole of it, and it is not presented here as more.

The **Competition and Markets Authority** has published no position on AI in financial services. Its AI work — the April 2024 strategic update, the foundation models review, *Agentic AI and consumers* of 9 March 2026 — is horizontal and contains no financial services content. That gap sits awkwardly beside the FCA’s own review, which names competition and market power as one of four shifts AI is driving through retail financial services. Nothing was found from the Financial Services Consumer Panel either, and that is recorded as not found rather than confirmed absent.

ISO/IEC 42001, and why it is in this section

No UK financial regulator expects, requires or endorses ISO/IEC 42001, the management system standard for AI. The single mention across the FCA, PRA and Bank of England is in the FCA’s AI Sprint summary of April 2025, recording what attendees said: “It was noted that international standards for AI models are being developed and adopted, eg ISO 42001.” A record of a discussion, not a policy.

HM Treasury’s adoption plan proposes something else: that “the industry should explore developing an AI Third Party Assurance Framework”, benchmarked against assessments already in use, with SOC2 named as the comparator rather than ISO 42001. UKAS-accredited certification became available in the UK in late 2025 and firms may certify voluntarily. The gap between “a standard exists” and “a regulator expects it” is where a certification will be sold to this reader.

FCA approach page, 13 Feb 2026 · DUAA 2025 s.80 · SI 2026/425 · CTP designations, 10 Jul 2026 · NAO HC 313 · FCA AI Sprint summary, 23 Apr 2025.

ELEVEN

The change almost nobody has written about

On 5 February 2026 the rule on automated decisions changed, in the direction of permission. Eight days later the regulator confirmed it would not be writing rules.

Section 80 of the Data (Use and Access) Act 2025 removed Article 22 of the UK GDPR and replaced it with Articles 22A to 22D. It was commenced on **5 February 2026** by regulation 2 of the Data (Use and Access) Act 2025 (Commencement No. 6 and Transitional and Saving Provisions) Regulations 2026.

A "significant decision" is one that produces a legal effect for the data subject, or has a similarly significant effect. A creditworthiness assessment is squarely within that. So is an affordability decline, a pricing decision and, on any sensible reading, a decision about whether someone in arrears is offered forbearance.

Before the change, a solely automated significant decision was prohibited unless an exception applied. After it, such a decision may be taken where the controller has safeguards in place. The prohibition survives in one place only: where the decision is based, entirely or partly, on special category data under Article 9(1), and then only explicit consent or a specified legal condition unlocks it.

> THE FOUR SAFEGUARDS, WHICH ARE NOW THE WHOLE OF THE PROTECTION
>
> Provide the data subject with information about decisions taken about them. Enable the data subject to make representations. Enable the data subject to obtain human intervention. Enable the data subject to contest the decision.

The practical consequence for a lender is direct and easy to miss. A creditworthiness assessment does not ordinarily rest on special category data, so the provision that most constrained automated lending is the one that has moved. Solely automated credit decisions are now permitted with safeguards rather than requiring an exception to a prohibition. Permission arrived first; the explanation of how to do it properly has not arrived at all.

The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, SI 2026/425, were made on 16 April 2026, laid on 21 April and came into force on **12 May 2026**. Regulation 2(1) requires that "The Commissioner must prepare an appropriate code of practice giving guidance as to good practice in the processing of personal data… in relation to — (a) developing and using artificial intelligence, and (b) automated decision-making." Regulation 2(2) adds that the code must include guidance on processing children's personal data. **The regulations set no deadline for producing the code.**

So the sequence, on the public record, runs: the permission widened on 5 February 2026; the FCA's approach page confirming it will not write AI rules was last updated on 13 February 2026; the duty to write a statutory code came into force on 12 May 2026 with no date attached; and the FCA's own guidance on applying existing rules to AI is due by the end of 2026 and has not appeared. A firm automating a lending decision today is doing so in the widest permission the UK has granted and the thinnest guidance it has published, at the same moment.

This is a statement of the legal position as at 15 September 2026, not legal advice. The interaction between section 80 and the Consumer Duty, and the treatment of a decision that draws indirectly on special category data, are both matters on which a firm should take its own advice.

Data (Use and Access) Act 2025, s.80, inserting Articles 22A–22D UK GDPR; commenced 5 February 2026 by SI 2026/82, reg 2 · The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, SI 2026/425, made 16 April 2026, in force 12 May 2026.

TWELVE

Beyond the UK — and inside it

The same decision, about the same customer, is classified two different ways depending on which side of the Channel it is taken.

Annex III of the EU AI Act lists the use cases treated as high-risk. Point **5(b)** covers AI systems intended to be used to evaluate the creditworthiness of natural persons or to establish their credit score — expressly excluding systems used for the purpose of detecting financial fraud. Point **5(c)** covers risk assessment and pricing in relation to natural persons in the case of life and health insurance.

High-risk classification brings obligations: a risk management system, data and data governance requirements, technical documentation, record-keeping and logging, transparency to deployers, human oversight, accuracy and robustness requirements, and conformity assessment. For the Annex III categories these apply from **2 December 2027**.

EUROPEAN UNION

Creditworthiness evaluation and credit scoring named as high-risk in the legislation itself. Fraud detection expressly carved out. A specified obligations regime, with a date.

UNITED KINGDOM

"We do not plan to introduce extra regulations for AI." The same decision governed by the Consumer Duty, the Senior Managers Regime and section 80 safeguards. No AI-specific obligations and no date.

This is the sharpest UK–EU contrast available anywhere in this series, and it is not a matter of one regime being stricter than the other in the abstract. It is that the EU has specified what a firm must do and the UK has not. For a lender operating in both markets the practical consequence is that the EU obligations will end up setting the internal standard, because they are the only ones written down. That is a defensible outcome, and it should be arrived at deliberately, as a decision a board has taken, rather than by default because the specified regime is the easier one to build against.

Inside the UK there is a second boundary, and it affects who a firm is answering to. The Payment Systems Regulator's consolidation into the FCA is confirmed government policy, but it requires primary legislation that has not been enacted. As at 15 September 2026 the PSR still exists and has published a 2026-27 annual plan. A briefing that describes it as abolished would be wrong, and a firm planning on that basis would be planning against something that has not happened.

EU AI Act, Annex III, points 5(b) and 5(c); high-risk obligations for Annex III systems from 2 December 2027 · FCA, *AI and the FCA: our approach* · HM Treasury, *Financial Services AI Adoption Plan*, 14 July 2026 · Payment Systems Regulator Annual Plan 2026-27.

THIRTEEN

What can be done today

Seven, all defensible from the evidence in this briefing, none requiring the FCA's guidance to arrive first.

→ **Explainability evidence for each named senior manager.** For every AI system operating in a senior management function holder's area, a one-page record: what it does, what it decides, what it was trained on, what it was tested against, who reviewed it and when. This is what "demonstrate understanding" means when someone asks, and it is the single artefact most likely to be wanted when the guidance lands.

→ **A model inventory that includes the bought-in systems.** A third of use cases are third-party and understanding is lowest there. An inventory covering only what was built in-house documents the best-understood third of the estate and leaves the rest unexamined. Include the model provider, the deployment route and the contractual position on changes to the model.

→ **The four section 80 safeguards, evidenced rather than described.** For each solely automated significant decision, show where information is actually given, where representations are actually received, where human intervention is actually available and how contest actually works. A policy that says these exist is not the same as a record showing they operated.

→ **A processing control for the redress exercise.** Identification and notification continue while payment is suspended. Record how the in-scope agreements are being assessed, by what system, with what error rate, who reviews the exceptions, and what happens when the tool is uncertain. The notification deadline of 18 November 2026 is the one currently running.

→ **Complaint-source tagging.** Tag complaints that arise from AI output the firm did not produce. Without it, harm arriving from outside the perimeter is invisible until it is large enough to be a pattern, and the firm learns about its own exposure from the ombudsman.

→ **Concentration mapping against the designated four.** The critical third parties regime went live on 13 July 2026 for cloud. Map which AI systems depend on which designated provider, and separately which depend on a model provider that was not designated — because that second list is the one with no regime attached to it.

→ **Collections triage, tested on the decline path.** The strongest human case for AI in this sector is reaching more of the 4.5 million adults in difficulty than the 1.6 million currently reached. Test the model on who it stops offering support to, not only on who it reaches. The failure mode that matters here is silent.

None of these seven requires new technology, a supplier, or a budget cycle. Each of them produces a document that would be useful to a senior manager being asked to account for something, which is the position section three describes.

Each item traces to a source named elsewhere in this briefing. None depends on the FCA's end-of-2026 guidance having been published first.

FOURTEEN

Three levels, and tomorrow

This is the only sector in this series where the regulator has published how the estate is distributed across the three levels.

ASSISTIVE

The tool suggests. A person decides. Drafting customer correspondence, summarising a call, surfacing the relevant policy, triaging a complaint into a queue. Accountability is unchanged because the decision never leaves the person. The cost is verification time, and it is real. Most of the sector's defensible near-term value sits here.

AUGMENTIVE

The tool decides. A person can intervene. Fraud scoring with a review queue, affordability pre-assessment, collections prioritisation, redress eligibility with exception handling. The survey puts 24 per cent of use cases here. The question a board should ask is what share of those decisions a human actually looks at, measured rather than assumed.

ADAPTIVE

The system decides and changes itself. Two per cent of the estate today. The FCA is testing agentic payments with named firms now, and section 80 permits a solely automated significant decision with safeguards. Permission is not the same as readiness, and the figure to hold onto is the other one: 46 per cent partial understanding, at 2 per cent autonomy.

Near, medium and long term

Now to 12 months	FCA guidance on applying existing rules to AI, including Senior Managers accountability, due by the end of 2026. Motor finance notification of non-compensated complainants due 18 November 2026 and 18 January 2027. AI Live Testing concludes at the end of 2026 with an evaluation report in early 2027. The FCA's AI good and poor practice publication expected during 2026.
12 to 36 months	Upper Tribunal hearings on the motor finance scheme in December 2026 or February 2027, with payments, if the scheme survives, expected in 2027. EU AI Act high-risk obligations for creditworthiness systems apply from 2 December 2027. The Information Commissioner's statutory AI and automated decision-making code, on no published timetable. Critical third parties oversight beds in for the four designated cloud providers.
Beyond 36 months	Whether the regulatory perimeter is redrawn around general-purpose tools that shape financial decisions — the first of the Mills Review's seven recommendations. Whether AI model providers join cloud providers in the critical third parties regime. Whether the Payment Systems Regulator's consolidation into the FCA reaches the statute book. Whether a sector at 2 per cent autonomy and 46 per cent partial understanding closes the second number before it raises the first.

Dates here are taken from the FCA, HM Treasury, the Upper Tribunal, the EU AI Act and the commencement regulations. Nothing in this table is a forecast.

Bank of England and FCA AI survey 2024: 55 per cent of use cases involve some automation, 24 per cent semi-autonomous with human oversight of critical decisions, 2 per cent fully autonomous.

FIFTEEN

Benefit realisation

Seven of the UK's largest lenders were read at their own published standard. What they disclose about AI is not comparable, and that is the finding.

The sweep in section seven was set up to answer one question: what is AI delivering in UK lending, according to the firms themselves. It did not produce that answer, and the reason is more useful than the answer would have been.

Lloyds Banking Group is the only UK lender to have published a monetary value. Generative AI "delivered around £50 million of value" in 2025, with "more than £100 million in additional value expected this year"; over 50 solutions deployed; an internal knowledge tool used by 20,000 colleagues cutting search times by 66 per cent; 5,000 engineers on a coding assistant. The release does not explain how the value figure is calculated.

NatWest	Publishes hours, not money — and two figures for the same activity. "The equivalent of over 22,000 hours in 2025" saved in complaints handling at group level, and separately "another 17,500 hours through faster, high-quality complaint responses" within Retail Banking. Not reconciled in the document.
HSBC	Quantifies adoption and never outcome. Over 100 solutions, 85 per cent of colleagues, 31,000 engineers. No monetary value, no headcount effect.
Santander UK	One operational figure, hedged in its own words: "a 50% alert reduction in some use cases".
Nationwide	Names more use cases than any of the seven and quantifies none of them.
TSB	Claims no deployment. Its AI language is risk-register language about safe exploration.
Monzo and Starling	Publish operational claims through engineering writing and product releases rather than in reporting obtained for this briefing, so they are not comparable to the five above on genre either.

> *A board that benchmarks itself against a competitor's AI number is benchmarking against a disclosure decision, not a performance.*

There is no common unit across these firms, no common scope and no common method. Hours are not convertible to money without assumptions the bank is not making. Adoption is not outcome. One firm publishing a value and six not says nothing about which is further ahead, and that is easy to conflate.

Four rules follow, and they survive contact with a supplier.

→ **Do not benchmark against a competitor's AI figure.** Require the method alongside any number brought into the building, including from inside it.

→ **Measure before, not after.** A baseline taken after deployment is not a baseline — the most common failure in AI business cases, and the easiest to avoid.

→ **Count the verification cost.** Every assistive tool creates checking work. That time is a cost of the tool and belongs in the calculation, not outside it.

→ **Publish the method internally even though the market does not.** The sector has no shared measure. A firm that builds one has something its own board can rely on, whatever the market publishes.

Read directly: NatWest ARA 2025 · HSBC Strategic Report 2025 · Santander UK 2025 · Nationwide ARA 2026 · TSB ARA 2025.

SIXTEEN

A way in

Eight workstreams. The first three form a natural first step, and two real deadlines sit inside it.

01	**AI readiness and model inventory**	4–6 weeks	What is running, who owns it, what was bought and what was built, and where the understanding gap actually is.
02	**SMCR explainability evidence pack**	3–4 weeks	The one-page record per system, per senior management function holder, described in section thirteen.
03	**Section 80 safeguards review**	3–4 weeks	Where each of the four safeguards actually operates for each solely automated significant decision.
04	**Board and executive governance**	4 weeks	What reaches the board, in what form, and what the board is expected to understand under the May 2026 joint statement.
05	**Perimeter and complaint-source response**	4 weeks	Tagging, detection and a position on customers acting on advice the firm did not give.
06	**Benefit realisation framework**	3 weeks, parallel	Baseline before deployment, verification cost counted, method recorded.
07	**Concentration and third-party mapping**	4–6 weeks	Dependency on the four designated providers, and separately on the model providers that were not designated.
08	**Twelve-month mobilisation roadmap**	2 weeks	Sequenced against the dates in section fourteen rather than against a budget cycle.

Unusually for this series, there are two real deadlines and both sit inside the next twelve weeks. FCA guidance on applying existing rules to AI, including Senior Managers accountability, is due by the end of 2026. The motor finance notification window for non-compensated complainants closes on 18 November 2026. Six to eight weeks fits inside both. That is arithmetic rather than a sales line, and it is the reason this briefing has a close at all.

CONTACT

Neil Catton · nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

Independent technology strategist, working across a fractional and advisory portfolio.

SEVENTEEN

Sources

Every figure in this briefing comes from one of the following. Nothing else was used.

Regulators and central bank

- Bank of England and Financial Conduct Authority, *Artificial intelligence in UK financial services — 2024*, published 21 November 2024. 118 respondents across six sectors. No later edition exists.
- FCA, *AI and the FCA: our approach*, last updated 13 February 2026.
- FCA, *AI and the future of retail financial services* (the Mills Review), published 6 July 2026. Consumer research by Yonder Consulting, 5,026 UK adults and eight focus groups in London and Leeds.
- FCA, *Financial Lives 2024 survey key findings*, published 16 May 2025; fieldwork 5 February to 16 June 2024, 17,950 adults.
- FCA, *Motor Finance Consumer Redress Scheme*, Policy Statement PS26/3, March 2026.
- FCA statement, *Motor finance scheme partially suspended*, 2 July 2026.
- FCA, *Credit Information Market Study* final report, MS19/1.3, December 2023; and CP26/7, February 2026, consultation closed 1 May 2026.
- FCA research note, *Credit where credit is due: How can we explain AI's role in credit decisions for consumers?*, 24 February 2025.
- FCA, *Using Synthetic Data in Financial Services*, Synthetic Data Expert Group, 8 March 2024.
- FCA, AI Sprint summary, 23 April 2025; AI Lab and AI Live Testing pages, last updated 3 July 2026.
- FCA research on AI and retail investors, fieldwork 24 July 2026, 666 UK investors aged 18–40.
- Bank of England, FCA and HM Treasury, joint statement on frontier AI models and cyber resilience, 15 May 2026.
- Bank of England, *UK financial regulators to begin overseeing Critical Third Parties announced by HMT*, July 2026.
- Bank of England, climate-related financial disclosure 2026, 25 June 2026.

Government, Parliament and audit

- Treasury Committee, *Artificial Intelligence in Financial Services*, Fifteenth Report of Session 2024–26, HC 684, 20 January 2026.
- *AI in financial services: Responses to the Committee's Fifteenth report*, HC 1791, 16 April 2026.
- HM Treasury, *Financial Services AI Adoption Plan*, 14 July 2026.
- National Audit Office, *Financial services regulation: Adapting to change*, HC 313, 8 December 2023.
- Competition and Markets Authority, AI strategic update (April 2024) and *Agentic AI and consumers* (9 March 2026) — checked, and containing no financial services content.

Legislation

- Data (Use and Access) Act 2025, section 80, inserting Articles 22A–22D into the UK GDPR; commenced 5 February 2026 by the Data (Use and Access) Act 2025 (Commencement No. 6 and Transitional and Saving Provisions) Regulations 2026, SI 2026/82, regulation 2.
- The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, SI 2026/425; made 16 April 2026, laid 21 April 2026, in force 12 May 2026.
- Regulation (EU) 2024/1689 (the EU AI Act), Annex III, points 5(b) and 5(c).

Continued overleaf.

Sources

Ombudsman, trade body and firms' own records

- Financial Ombudsman Service, annual complaints data 2025/26, published 21 May 2026; and *Plans and Budget 2026/27*, March 2026.
- UK Finance, *Annual Fraud Report 2026*, covering 2025. The report's AI and deepfake commentary appears in a vendor-authored sponsored foreword and is used here only with that sponsorship stated.
- NatWest Group plc, Annual Report and Accounts 2025, approved 12 February 2026; and 2025 Climate Transition Plan Report.
- HSBC Holdings plc, Strategic Report 2025, 25 February 2026; and 2026 Interim Results, 4 August 2026.
- Santander UK Group Holdings plc, Annual Report 2025, 9 March 2026.
- Nationwide Building Society, Annual Report and Accounts 2026, 20 May 2026.
- TSB Banking Group plc, Annual Report and Accounts 2025, approved 5 February 2026.
- Lloyds Banking Group, own release on generative AI value; and Sustainability Report 2025.
- Monzo Bank, *Machine Learning at Monzo in 2025* (21 November 2025) and *Engineering the Future of Customer Operations: The Monzo Ops Agent* (4 June 2026) — the firm's own engineering publications, not its annual report.
- Starling Bank, results release of 21 May 2026 for the year to 31 March 2026, and Annual Report 2026 microsite.

Two gaps are recorded rather than filled. Monzo's Annual Report 2026 and Starling's Annual Report 2026 could not be obtained for this briefing, so those two firms are described from their own published material of a different kind, and this is stated wherever they appear. Nothing found from the Financial Services Consumer Panel is recorded as not found rather than as confirmed absent, because the Panel's publications index could not be read.

EIGHTEEN

About the author

Neil Catton is an independent technology strategist working across a fractional and advisory portfolio — fractional CTO and CIO roles, AI board briefings, AI readiness reviews, technology due diligence, data strategy and programme recovery. He is the author of *The Next Evolution*, *The Cognitive Crucible* and *The Shadow System*, and writes at writing.neilcatton.com on technology, ethics and human purpose.

More than thirty-nine years in technology across more than twenty sectors, including banking and financial services. That work includes designing and architecting two UK challenger banks from scratch — one built to serve military personnel, police officers, emergency responders and their families, the other an SME lending platform — in both cases designing the technology architecture end to end, selecting and onboarding the delivery ecosystem, working directly with investors and supporting the financial services licence process, and in one case advising the non-executive board. Earlier work covers Oracle financial platforms at a clearing bank and at a building society, and pan-European banking platform work with a banking software business. Current work includes a pre-launch business in the institutional investment and pension-adjacent long-term savings space, taking it through licensing and towards its first funded phase.

No organisation named in this briefing has been engaged, approached or consulted in preparing it.

NOTICE

1. This briefing is not legal, regulatory, financial or professional advice.

2. Position stated as at 15 September 2026. The FCA's end-of-2026 guidance, its AI good and poor practice publication, the motor finance legal challenges and the Information Commissioner's statutory AI code were all unresolved at that date and may since have moved.

3. Third-party sources are cited in good faith. Verify before relying on any of them.

4. Named organisations appear only where a regulator, ombudsman, select committee or court has placed them on the record, or where they published the material themselves. They are used illustratively; none has been engaged or approached.

5. No advisory or client relationship is created by receipt of this document.

6. Independence — no vendor affiliations, reseller agreements, referral arrangements, commission interests or equity positions in anything referenced, and no payment from any organisation named.

Produced with AI assistance, used transparently and under direct human editorial control — every source verified, every judgement reviewed by Neil Catton before publication.

Version history

v1.0	First publication. Eighteen sections and this version record.	15 Sep 2026