BRIEFING · SEPTEMBER 2026

Artificial Intelligence in *UK Grocery Retail*

What the evidence actually says, what the regulators have and have not said, and what can safely be done now.

Neil Catton

Independent Technology Strategist · Technology, Ethics & Human Purpose

NEILCATTON.COM

Prepared as a market briefing. Every figure is traced to a regulator, a trade body, published legislation, government statistics, named academic work, or an operator speaking on its own record.

ONE

A note on the numbers

Almost every published statistic about AI in retail is produced by a company selling AI into retail. This briefing leaves them out, and says which ones.

A senior executive picking up an AI programme in grocery will be handed a folder of adoption percentages, productivity uplifts and accuracy claims. Almost none of it can be traced. The discipline here is simple: if a number cannot be followed back to a regulator, a trade body, published legislation, government statistics, named academic research, or an operator speaking on its own record, it is not in this document.

What that leaves is narrower and considerably more solid. The Information Commissioner's Office and its published strategy. The Competition and Markets Authority and its annual plan and closed cases. The Groceries Code Adjudicator's own record of its meetings, its annual supplier survey and the Code itself. The Food Standards Agency's Science Council. The Home Office. Defra and the Office for National Statistics. The National Cyber Security Centre. The British Retail Consortium's survey work, and Usdaw's stated position. Which?'s freedom of information data on enforcement capacity. And, for what is actually running, the audited annual reports and results statements of Sainsbury's, Tesco, Marks & Spencer and the Co-op.

What was found and deliberately excluded

Naming the exclusions is more useful than any statistic that survived, because it tells the reader what the rest of their folder is made of.

- → Vendor and survey-house figures of the form "X per cent of retailers are now using AI". Every one located traced back to an organisation selling into the sector.
- → A technology supplier's assertion that its deep learning models are "up to 40 per cent more accurate than traditional retail forecasting systems". No baseline is defined, no method or period is given, and the company sells that technology to other grocers.
- → Sector research held behind a paywall, including material from the industry's own research charity, which could not be read in full and so could not be cited.
- → Consultancy productivity and margin claims that name no client and quote no named executive.
- → A widely circulated dollar figure for UK shoplifting losses, which is a currency conversion of a British Retail Consortium number. The British Retail Consortium is cited instead.
- → Claims that electronic shelf labels are enabling surge pricing in British supermarkets. No UK primary evidence was found, and the strongest available academic work points the other way.
- → The initial access route for the 2025 attacks on Marks & Spencer and the Co-op. It is widely reported and it is not confirmed in the National Cyber Security Centre's own record, so it is not stated here as fact.
- → One supplier accuracy figure, which appears in this document only as a claim being examined, never as a performance statistic.

One sentence, and one silence

Two things are flagged rather than excluded, because their weight is the finding. The Information Commissioner's Office has made exactly one public statement about facial recognition in shops, given to a trade magazine in February 2026 after a shopper was wrongly ejected from a supermarket. It is not among the five priorities in the ICO's published artificial intelligence and biometrics strategy, and no guidance, audit or enforcement sits behind it.

And across two publications in 2026 the Groceries Code Adjudicator makes no mention of technology, systems, data or artificial intelligence — while holding a binding duty over the accuracy of retailers' forecasts, which are now produced by machine learning. Both the sentence and the silence are quoted later.

Sources are named at the foot of every evidenced section and listed in full in section thirteen.

TWO

The wider market

Grocery is the one sector every household transacts with weekly, and the burden of getting it wrong is not evenly spread.

British households spent £303.3 billion on food and drink in 2025, measured at constant 2023 prices. Food and drink retailing generated £44.1 billion of gross value added in 2024, just over a quarter of the whole agri-food economy, and employed 1.1 million people in Great Britain in 2025. Retail sales as a whole ran at £9.6 billion a week in May 2026, with online at 28.8 per cent of the total, rising to 29.4 per cent by June — the highest share since April 2021.

£303.3bn
Consumer expenditure on food and drink, 2025, at constant 2023 prices.

1.1m
People employed in food and drink retail in Great Britain, 2025.

15.2%
Share of household spending on food for the poorest fifth, against 7.9% for the richest.

The last of those three is the one that should shape every design decision that follows. On average, food takes 10.9 per cent of household expenditure. For the poorest fifth of households it takes 15.2 per cent; for the richest fifth, 7.9 per cent. The households most exposed to a supermarket's pricing, promotion, personalisation and loss-prevention systems are the households spending nearly twice the share of their income in that shop, with the least capacity to absorb an error.

This matters more in grocery than in almost any other consumer market because the relationship is anonymous and repeated. A bank knows its customer. An insurer has a file. A supermarket serves millions of people it has never identified, dozens of times a year, and increasingly makes automated judgements about them on the basis of a face, a basket or a loyalty record. The volume is the point: even a very small error rate, applied to a very large number of interactions, produces a steady stream of people who have been treated badly and have nowhere obvious to go.

WHY THIS IS NOT THE USUAL MARKET SLIDE

Market size in this sector is not the argument. Every grocer knows what it turns over. The figure that changes an AI decision is the distributional one: who carries the consequence when a system is wrong, and how much room they have to absorb it.

Defra, Food Statistics in your pocket, published 29 July 2026, drawing on ONS Annual Business Survey, Consumer Trends, Labour Market Statistics and Family Spending in the UK. ONS, Retail sales Great Britain, June 2026, released 24 July 2026. House of Commons Library, Retail sales economic indicators.

THREE

The sector's pressure point

A record-low workforce, a record cost base, and a regulator that is looking somewhere else.

Retail employment in Great Britain has fallen to the lowest level ever recorded. There were 2.76 million retail jobs in March 2026. The four-quarter average for the first quarter of 2026 stood at 2.79 million — 66,000 fewer than a year earlier and 398,000 fewer than a decade ago, with full-time roles down 163,000 and part-time down 235,000 over that period. A quarter of a million retail jobs have gone in five years, 74,000 of them in the last twelve months.

Costs moved the other way. Employer National Insurance and the National Living Wage added roughly £5 billion to retail employment costs during 2025 and £6.5 billion over fourteen months. The cost of employing a full-time entry-level worker rose by 10 per cent, part-time by 13 per cent. In February 2026, 69 per cent of retail finance chiefs described themselves as pessimistic or very pessimistic about the year, against 56 per cent seven months earlier, and 84 per cent put labour and employment costs in their top three concerns, up from 21 per cent.

Their answer is visible in the same survey. Fifty-two per cent planned to cut hours or overtime, 32 per cent to freeze recruitment, 48 per cent to reduce head office headcount and 32 per cent to reduce store headcount. Sixty-eight per cent intended to drive higher productivity. And 61 per cent said they were investing in automation.

> THE CONTRADICTION
>
> *A sector is being pushed towards automated decision-making by arithmetic rather than pulled towards it by strategy — and every regulator holding a position on AI in this sector has aimed it somewhere other than the largest deployments.*

That last point deserves stating precisely, because the easy version of it is wrong. This is not an unregulated sector, and it is not a sector where nobody has thought about AI. Four bodies hold a position. The Information Commissioner's Office made proportionate police use of facial recognition one of five priorities in March 2026, while the facial recognition is in supermarkets. The Food Standards Agency's Science Council published a considered assessment of AI in the food system in June 2026 — advisory, and aimed at safety and authenticity rather than at commercial systems. The Groceries Code Adjudicator holds a binding duty over the accuracy of retailers' forecasts and, across two publications in 2026, has never mentioned the technology now producing them. The Competition and Markets Authority is looking at online checkouts.

The only thing a regulator has said about this in a shop

> *Retailers should carefully consider the risks of misidentification and have robust procedures in place to ensure the accuracy and integrity of the personal information they collect and process.*

That is the Information Commissioner's Office, speaking to a trade magazine in February 2026, after the first Sainsbury's misidentification. It is the regulator's entire published position on facial recognition in British shops. It asks for robust procedures. The operator's answer is a trained manager reviewing every match. Six months later that answer failed again.

So the pressure point is not absence of supervision. It is misdirection. A grocer buying automation to close a cost gap is making a governance decision whether or not it recognises it as one, and it will not be corrected from outside — not because no one is watching, but because everyone is watching something else.

BRC, Retail employment falls as costs rise, 18 June 2026. BRC CFO survey, 19 February 2026. BRC on Guaranteed Hours reforms, 13 August 2026.

FOUR

What the sector is facing

Four bodies hold a position on AI in or around this sector. Read all four and the pattern is unmistakable.

The Information Commissioner's Office published an update to its artificial intelligence and biometrics strategy on 17 March 2026. It names five priorities: giving organisations clarity on responsible AI and automated decision-making; ensuring high standards of automated decision-making in central government; setting expectations for automated decisions in recruitment; scrutinising foundation model developers; and supporting proportionate police use of facial recognition. Private-sector facial recognition does not appear. Draft automated decision-making and profiling guidance is in development, and the ICO is committed to producing a code of practice under the Data (Use and Access) Act.

The Competition and Markets Authority published its 2026–27 annual plan on 23 March 2026. Its consumer enforcement priorities are price transparency and misleading online choice architecture, in secondary ticketing, homeware and personal fitness. On algorithms it commits to deterring algorithmic collusion, using its in-house technology, data and AI capability. Its essential-spend work names infant formula, vets and road fuel. Grocery AI does not appear. Its review of loyalty pricing in groceries — the obvious place a competition regulator would have found an algorithmic story — closed on 27 November 2024 having analysed around 50,000 loyalty-priced products and found very little evidence of inflated regular prices. It wrote to supermarkets on compliance, raised concerns about alternating loyalty and "was/now" promotions, and asked retailers to improve access for shoppers without smartphones and for under-18s.

The Food Standards Agency's Science Council published a report on AI applications in food safety and authenticity on 25 June 2026. It is the most considered piece of regulatory thinking about AI anywhere in this sector, and its central conclusion is the same one this briefing keeps arriving at: *"AI should be treated as a decision-support tool, not a replacement for human accountability."* It found existing food safety law sufficiently robust and sought no new legislation. Its six recommendations include publishing guidance that AI is decision-support only with humans retaining accountability, promoting data standards for quality and provenance, and commissioning behavioural research on how workers actually interact with AI outputs. Its named risks are worth reading twice: embedded bias and drift without validation; outputs "difficult to explain or reproduce, undermining traceability"; weaker operators using AI to generate "convincing risk assessments" without substance; and "over-reliance on automation" leading to skill deterioration and false confidence. It is advisory, and it is aimed at safety and authenticity rather than at the commercial systems that make most of the decisions.

The Groceries Code Adjudicator's record of its June 2026 progress meetings with all thirteen designated retailers covers cost price increases, the seven golden rules, suppliers' reluctance to raise Code issues, delayed payments, de-listing and forecasting. Its July 2026 statement of priorities repeats the list: Mark White, Adjudicator — *"I will continue to focus proportionately on Code compliance in relation to delays in payment, forecasting, marketing costs, and de-listing practices."* Neither document contains any reference to technology, systems, data or artificial intelligence. The body holding a binding duty over the accuracy of retailers' forecasts has no stated position on the machine learning now producing them.

The people who would have to catch a problem in a store

Enforcement of practice inside a shop sits with local authority trading standards. Which? put a freedom of information request to all 187 services in Great Britain in July 2024. Between them they employ 2,284 full-time equivalent staff. The average service has 7.5 professionally qualified officers against a benchmark of 9.4 six years earlier; 63 per cent have six or fewer. Total funding was £117 million in 2023–24, a median of £2.04 per head of population.

> WHY THIS BELONGS IN AN AI BRIEFING
>
> A sector adopting automated decision-making faster than any consumer sector outside financial services is supervised, at store level, by a service that has lost roughly a fifth of its qualified officers per team in six years. Whatever governance a grocer builds, it will not be corrected from outside.

ICO, AI and biometrics strategy update, 17 March 2026. CMA, Annual Plan 2026 to 2027, 23 March 2026, and Loyalty pricing in the groceries sector, closed 27 November 2024. FSA Science Council, AI applications in food safety and authenticity, 25 June 2026. GCA, CCO progress meeting record, 12 June 2026, and Adjudicator's update, 16 July 2026. Which?, Trading standards resources and activities, 7 August 2025.

FIVE

The challenges of using AI here

Two stores in September 2025. Around two hundred by Christmas 2026. Twice in between, the wrong person was stopped.

The estate is wider than the argument that follows. On the operators' own record it runs from machine learning forecasting "embedded across all food products" at Sainsbury's, through automated warehousing at Daventry and mobile robots at Northampton, a decisioning engine issuing 500 million personalised offers a week and more than two hundred live AI use cases, to Tesco's agentic AI across 24 million Clubcard households and an assistant trialled with some 280,000 colleagues, and 11,000 Microsoft Copilot licences at Marks & Spencer covering every store manager. Most of that is unglamorous, defensible and delivering. Section seven returns to the part of it that carries a statutory duty. This section takes the one deployment where the failure has already happened in public, because it shows what every other part of the estate is being asked to rely on.

Now the specific system. Sainsbury's trialled facial recognition with the supplier Facewatch in two stores in September 2025. Its annual report, in the words of Tracey Clements, Chief Retail, Logistics and Supply Chain Officer, records "a 46 per cent reduction in logged incidents of theft, harm and aggression and anti-social behaviour and 92 per cent of offenders did not return". It reached six stores by February 2026, more than 55 by July, with up to 150 more announced before Christmas.

Two people, named, in public

In February 2026 Warren Rajah, a 42-year-old technology worker, was approached by three managers at the Elephant and Castle store and told to leave. Facewatch confirmed there were no incidents or alerts associated with him in its database and that it "did not play any part" in his being approached. Sainsbury's said this was "not an issue with the facial recognition technology in use but a case of the wrong person being approached in store", called it the first such incident and retrained staff. He was offered a £75 shopping voucher.

On 6 August 2026 Matt Arnold was stopped at a self-checkout in the East Dulwich store and prevented from completing his purchase. "The staff member said I'd been identified by the AI," he said, "and the cameras had flagged me." Sainsbury's response: "The incident was caused by human error, not the facial recognition technology… the Facewatch system has a 99.98 per cent accuracy rate, and every match is reviewed by a trained manager." Facewatch said its technology "was not at fault… A correct alert was sent to the retailer, but was subsequently subject to human error in the way it was handled." The system was suspended in that store; the wider rollout continued.

These are two different failures. In the first there was no alert about that person at all; in the second, a correct alert about someone, handled wrongly. Both were attributed to the human step. That is the finding, not the defence — the human step is precisely what the law now rests on.

The accuracy claim, examined

The 99.98 per cent figure, quoted by the retailer in its own defence, was examined on 4 May 2026 by Professor Alan Dix, Professor Emeritus at Swansea University and Director of its Computational Foundry. Shoplifting occurs in the order of one in ten thousand shopping visits, so a system that flagged nobody at all would score 99.99 per cent — better than the claim. On plausible assumptions he derives roughly one misidentified innocent person for every correct identification, and using official crime statistics closer to three. The point is not that the technology is bad, but that an accuracy percentage quoted without the base rate tells a board nothing about how often staff will approach the wrong person.

The same question sits under every other system in the estate, and section seven shows that for forecasting it is not a matter of good practice but of statutory duty.

J Sainsbury plc Annual Report 2026; Tesco plc, 9 April 2026; M&S, 25 March 2026; trade press, February and August 2026; Prof Alan Dix, 4 May 2026.

SIX

The threat side

In 2025 two of Britain's largest grocers were taken down. Neither attack defeated a control. Both persuaded a person.

It is tempting to assume the AI threat in retail is fraud against the till, because that is the shape the threat takes in insurance and banking. In grocery it is not. The 2025 attacks on Marks & Spencer and the Co-op were not clever technical defeats of a security product. They were failures of human verification, at scale, with consequences that show up in audited accounts.

The Co-op's own members' annual report for 2025 records an estimated £285 million direct revenue impact from "restricted systems, reduced promotional capability, operational inefficiencies", an estimated £86 million impact on profitability and a further £21 million of additional costs. Group revenue fell 2.3 per cent to £11.0 billion; a £131 million underlying operating profit became a £35 million loss, and a £161 million profit before tax became a £43 million loss. Shirine Khoury-Haq, Chief Executive: "In April, we identified malicious activity on our network. We then restricted certain systems to contain the threat. This kept stores, payments and online services operational, but impacted product availability."

Marks & Spencer reports £292.1 million of adjusting items for the year to 28 March 2026, including £131.3 million of incident-related costs, against £100.0 million of insurance proceeds. Technology investment in fashion planning, food warehouse management systems and e-commerce stopped and only restarted in the second half of the year. The National Cyber Security Centre puts the combined cost to Marks & Spencer and its insurers at over £300 million, and records that the data of all 6.5 million Co-op members was stolen by the DragonForce ransomware group — a statement the Co-op's own annual report does not make.

What the NCSC says the technology is for

> *Actors linked to China, Russia, Iran and the DPRK are using large language models to evade detection, support reconnaissance, process exfiltrated data, access systems through social engineering, and support vulnerability research and exploit development.*

The National Cyber Security Centre's annual review for 2025 records 1,727 incident tips, 429 incidents requiring its support, 204 of them nationally significant against 89 the year before, and 18 highly significant — a 50 per cent rise for the third consecutive year. It names AI-assisted vulnerability research and exploitation as the most significant near-term development.

The connection to the rest of this briefing is exact. The facial recognition failures and the ransomware losses are the same failure in two costumes: a human being, under time pressure, asked to verify an identity they cannot actually verify, and given a system that sounds authoritative. Generative AI makes the attacker's side of that exchange cheaper every month. It does nothing on its own for the defender's.

> STATED AS WHAT IT IS
>
> The specific initial access route for the 2025 retail attacks is widely reported and is not confirmed in the NCSC's own record. It is therefore not asserted here. The audited financial consequences, and the NCSC's own statements about AI and social engineering, are.

NCSC Annual Review 2025. Co-op Group Members' Annual Report 2025. M&S full year results, 52 weeks ended 28 March 2026.

SEVEN

The UK legal position

Section 80 did not tighten the rules on automated decisions. It loosened them — and moved the whole weight of the law onto the quality of the human review.

This is the change nobody in retail has framed, and it is the most consequential thing in this document. Before 5 February 2026, a decision made solely by automated processing that produced a legal effect, or a similarly significant effect, on a person was prohibited unless a narrow exception applied. Section 80 of the Data (Use and Access) Act 2025 substituted Article 22 of the UK GDPR. Since that date such decisions are permitted, provided four safeguards are met: the person is given information about the decision, can make representations about it, can obtain human intervention, and can contest it.

For a retailer this is not a relaxation. It is a transfer. The defence that a decision was not solely automated — the defence Sainsbury's runs when it says every match is reviewed by a trained manager — now rests entirely on whether that review is doing something real. In both published cases the company's own account is that the technology worked and the human handling did not. That is not a mitigation under section 80. It is the precise point at which the safeguard is supposed to hold.

Data (Use and Access) Act 2025, s.80	5 Feb 2026	Substitutes UK GDPR Article 22. Solely automated significant decisions permitted with four safeguards: inform, represent, intervene, contest.
Groceries Supply Code of Practice, para 10	In force	Full compensation for supplier costs from a retailer's forecasting error, unless prepared in good faith and with due care, after consultation, on a basis communicated to the supplier.
DMCC Act 2024, Part 4	6 Apr 2025	CMA enforces consumer law directly without going to court, with penalties up to 10 per cent of global turnover. First penalty issued April 2026: £4.2m and £760,000 of refunds for drip pricing.
Crime and Policing Act 2026	Apr 2026	Standalone offence of assaulting a retail worker; the £200 summary-only threshold for shop theft repealed, making all shop theft an either-way offence.
Employment Rights Act 2025	Pending	Guaranteed hours, reasonable notice of shifts, cancellation pay. Consultation closed 25 August 2026; outcome pending, no commencement date set.

The second duty, and the one an AI lead actually owns

Section 80 governs the customer. Paragraph 10 of the Groceries Supply Code of Practice governs the supplier, and has been in force far longer. A retailer "must fully compensate a Supplier for any cost incurred by that Supplier as a result of any forecasting error… attributable to that Retailer" — unless it prepared the forecast "in good faith and with due care, and following consultation with the Supplier", and has communicated "the basis on which it prepares any forecast".

Sainsbury's says machine learning forecasting is embedded across all food products. Three questions follow, and no UK grocer has answered any of them in public. Can a retailer demonstrate *due care* in a forecast produced by a model it did not build? Can it *communicate the basis* of a forecast when that basis is a set of learned weights? Has it *consulted the supplier* on a forecast the system regenerates nightly? The failure rate is already public: in the Adjudicator's April 2026 survey, costs from inaccurate retailer forecasting were the second most reported Code issue at 18 per cent, with a further 10 per cent reporting no compensation for forecasting errors. Both duties turn on the same question — did a person exercise real judgement, and can you show it.

A supplier harmed by a bad forecast has a regulator, a duty and a compensation route. A shopper wrongly identified as a thief has neither. They can complain to the ICO, or bring a civil claim as the 19-year-old wrongly flagged in a Home Bargains in Manchester did in 2024. There is no statutory grocery ombudsman; Retail ADR covers faulty goods, delivery and pricing, and only for participating members.

DUAA 2025 s.80, commenced by SI 2026/82. Groceries Supply Code of Practice, paragraph 10; GCA best practice statement on forecasting and promotions, updated 1 July 2026; GCA annual survey, 23 April 2026. CMA, AA and BSM decision, 15 April 2026. Home Office retail crime factsheet, 19 May 2026. ICO statement, February 2026. Big Brother Watch, 24 May 2024.

EIGHT

Beyond the UK

The EU has written down the duty this sector keeps failing — and then postponed it by more than a year.

It is commonly said that the EU has banned this. It has not, and a briefing that says so has not read the Article. Article 5(1)(h) of the EU AI Act prohibits the use of real-time remote biometric identification systems in publicly accessible spaces *for the purposes of law enforcement*. A supermarket is not a law enforcement body and is not caught by the prohibition.

What the Act does do is classify the technology. Annex III, point 1(a) makes remote biometric identification a high-risk AI system — the automated recognition of features to establish a person's identity by comparison against a database, as distinct from one-to-one verification, which is excluded. High-risk classification brings Article 26 deployer obligations, and one clause of it is the sentence this whole briefing has been circling: deployers "shall assign human oversight to natural persons who have the necessary competence, training and authority, as well as the necessary support". Deployers must also retain the system's logs for at least six months.

And then the Union deferred it. The Digital Omnibus agreement pushes the Annex III high-risk obligations to 2 December 2027, with Annex I product-embedded systems following in August 2028. Until those dates, a supermarket's facial recognition is governed on both sides of the Channel by data protection law alone.

> THE HONEST COMPARISON
>
> *The competence, training and authority test exists in writing, applies to nobody yet, and is the exact standard against which both UK incidents would be judged. The divergence between the two jurisdictions is not that one prohibits and the other permits. It is that one has written down what good oversight looks like and has not switched it on, and the other has not written it down at all.*

For a UK grocer with no European operation, none of this binds. It is still the most useful drafting available, because it is the only place a legislature has defined what the human in the loop must actually be. A board that adopts the Article 26 test voluntarily gets three things: a defensible answer to the section 80 question at home, a standard its store managers can be trained and measured against, and a position that does not have to be rebuilt if the UK's own guidance, when it arrives, lands in the same place.

Regulation (EU) 2024/1689, Article 5(1)(h), Article 26 and Annex III point 1(a). EU Digital Omnibus agreement on postponed high-risk deadlines, 2026.

NINE

Today

Eight things that are defensible now, need no new law, and would survive a regulator reading them back to you.

Four challenges specific to grocery

01	**The decision lands on a stranger**	Not a customer with a file and a reference number, but a member of the public being judged in front of other people.
02	**The safeguard is a duty manager**	Seconds, mid-shift, with a queue forming. It is the thinnest human oversight in any regulated consumer setting.
03	**Base rates make good systems look bad**	Rare events plus enormous volume produces a steady flow of wrong approaches from a system with an excellent error rate.
04	**There is no route back**	Retail ADR covers faulty goods, delivery and pricing, and only for participating members. Being wrongly identified as a criminal is not on its list.

None of this requires a new platform, a new supplier or a budget round. All of it is the work that looks unnecessary until the morning after an incident, when it is the only thing separating a company that was careful from one that was lucky.

→ **Write down every decision a system takes about an individual or a supplier**, and who signs it. Most grocers cannot produce this list, which means no one can say how many section 80 or paragraph 10 decisions the business is making.

→ **Test the human review, not the model.** Time it, observe it in a real store at peak, and record how often the reviewer actually disagrees. A review that never disagrees is not a review, and a regulator will reach that conclusion faster than a board will.

→ **Demand the base rate from every supplier, not the accuracy figure.** Not "how accurate is it" but "how many false approaches per thousand alerts, in a store of my footfall". A supplier who cannot answer that has not measured it.

→ **Build the due-care file for forecasting.** How the model was validated, what each supplier was told about the basis of their forecast, and when they were last consulted. Paragraph 10 asks for exactly these three things.

→ **Measure forecast error where it lands** — in supplier cost and waste, not only availability. The Adjudicator publishes the sector's score annually; a retailer that does not know its own is learning it from a survey.

→ **Publish a redress route of your own, before someone designs one for you.** Name it in store, staff it, log every use. The current published price of getting this wrong is a £75 voucher, set by the retailer.

→ **Give the person on the shop floor the authority to stop.** A manager who cannot overrule an alert without escalating is not oversight; they are a relay.

→ **Separate the loss-prevention, forecasting and personalisation estates** in governance. They carry different duties to different people, and one should never inherit another's risk tolerance because they share a data platform.

→ **Keep the logs.** Six months is the EU's floor for high-risk systems, and the minimum any UK operator will wish it had after an incident.

None of this is a pause on AI or a retreat from it. Markdown optimisation, vehicle routing, warehouse automation and colleague copilots all deliver real value and carry no duty to a named third party. The discipline above applies to the systems that judge a person or a supplier — and it is those which will define how this sector is regulated.

Derived from DUAA 2025 s.80, GSCOP paragraph 10, the GCA best practice statement updated 1 July 2026, the FSA Science Council's six recommendations, and EU AI Act Article 26.

TEN

Tomorrow

Three levels, defined by where accountability sits — not by how clever the system is.

The useful distinction for a board is not technical. It is about who answers for the outcome. Vendors sort systems by capability; regulators sort them by accountability, and so should the people signing them off.

ASSISTIVE

Prepares, summarises, surfaces

The person decides. Accountability is unchanged. Machine learning forecasting, markdown optimisation, range planning, colleague-facing search. Almost all defensible value in grocery today sits here, and most of it is unglamorous.

AUGMENTIVE

Recommends within bounds

The system gives reasoning and confidence; the person confirms or overrides, and the override is recorded. A facial recognition alert reviewed by a manager is meant to be here. Whether it is depends entirely on whether the review is demonstrably meaningful.

ADAPTIVE

Decides within an envelope

Oversight moves from the decision to the envelope, and accountability moves to whoever sets and monitors it. Nothing in UK grocery is credibly here yet. This level is a reasoned projection, not an observation.

Near, medium and long term

Zero to twelve months. The decision register, the timed review test, supplier base rates, a named redress route and log retention. None of this needs new technology and all of it is what a regulator would ask for first. It is also what makes any later deployment defensible rather than hopeful.

One to three years. The ICO's automated decision-making guidance lands. The Employment Rights Act consultation outcome sets the shape of scheduling, and with it whether rota systems become a section 80 question in their own right. EU high-risk obligations bite in December 2027, and the Deposit Return Scheme lands in October 2027, putting container-level data and reconciliation into every store estate. Grocers with a decision register will absorb these. Grocers without one will discover their estate during an investigation.

Three years and beyond. Pressure will build to let systems act without a human step in each case, because the human step is the expensive part and the workforce is shrinking. The board question then is not whether the model is good enough. It is who owns the envelope, how narrow it is, who monitors it, and what happens on the day it is wrong at scale rather than one shopper at a time.

THE HONEST CAVEAT

The first horizon is observation. The second is a reading of published timetables, and regulators miss timetables — the guaranteed hours consultation has already run past its own indicative dates, and the EU has just deferred its high-risk regime by more than a year. The third is a reasoned projection and nothing more. Any briefing that presents all three with equal confidence is selling something.

ICO, AI and biometrics strategy update, 17 March 2026. gov.uk consultation outcome pending. EU Digital Omnibus agreement.

ELEVEN

Benefit realisation

The two-store trial, and what it would take to make it carry a two-hundred-store decision.

Sainsbury's reported a 46 per cent fall in logged incidents of theft, harm, aggression and anti-social behaviour, and 92 per cent of identified offenders not returning. Those results are real, they are the company's own, and they are published in an audited annual report — which already puts them ahead of most evidence in this market. They came from two stores over one autumn.

To carry a decision a hundred times larger, that measurement needs three things it does not yet have in public. A control group of comparable stores without the technology, so the fall can be separated from everything else happening to retail crime in the same period — and something was happening, because violence and abuse across the sector fell from 2,000 incidents a day to 1,600 over roughly the same window. A count of alerts that led to an approach of the wrong person, which is the cost side of the ledger and appears nowhere. And a check on whether incidents fell or simply stopped being logged, because the presence of a camera changes reporting behaviour as well as offending behaviour.

Eight numbers, across two estates, reported to the board monthly

01	**False approaches per thousand alerts**	Counted by the store, not supplied by the vendor. The only number that measures harm done to innocent people, and the one no supplier volunteers.
02	**Override rate by the reviewing manager**	With a floor below which the review is presumed not to be happening. Set it in advance, in writing, and investigate stores that fall under it.
03	**Time from alert to approach**	A review that takes four seconds is not a review. The most direct evidence a regulator could be shown that human oversight is real.
04	**Complaints and their outcomes**	Including what was paid, to whom, and whether the person was removed from any watchlist. A count, not anecdotes.
05	**Forecast error as supplier cost**	Measured where it lands, in supplier cost and waste, not only as availability and service level. The Adjudicator publishes the sector's score annually.
06	**Planner override rate on forecasts**	And whether anyone reviews the overrides. The same test as line 02, applied to the system that touches every supplier rather than every shopper.
07	**Compensation paid under paragraph 10**	Reported to the board as a line of its own. If the answer is nil, the question is whether that is performance or an unexamined assumption of due care.
08	**Date each supplier was last consulted**	On the basis of their forecast. Paragraph 10 requires it and the Adjudicator asks for collaboration on it. Very few grocers could produce this list today.

The personalisation estate deserves the same discipline, where the harm is quieter and the volumes far larger. Five hundred million personalised offers a week is not a marketing statistic; it is five hundred million automated judgements about individual households, in a market where the poorest fifth spend nearly twice the share of their income. The question a board should ask of that engine is not what it lifted. It is which households systematically receive worse offers, and whether anyone has looked.

> THE NUMBER TO BE LEAST COMFORTABLE WITH
>
> The published price of wrongly ejecting a shopper is a £75 shopping voucher — set by the retailer, because no regulator has set one. The price of a bad forecast is set by statute and measured annually by the Adjudicator. The same company owns both systems, and only one of them has a number attached that anyone outside the business can check.

J Sainsbury plc Annual Report and Financial Statements 2026. BRC Crime Report 2026, 24 February 2026. Compensation reported in the trade press, February 2026.

A way in

Eight workstreams. The first three are a natural first step.

01	**Readiness assessment**	What decisions systems already take about individuals, where the data comes from, who signs each one, and what the review actually consists of in a store at peak.	4–6 weeks
02	**Use case identification and triage**	Sorted by the accountability ladder rather than by expected value, so the assistive work that is safe to accelerate is separated from the judgements that need governance first.	3–4 weeks
03	**Regulatory and legal position paper**	Section 80 exposure across the estate, paragraph 10 due-care exposure on forecasting, the special category data position on biometrics, and a defensible written answer to the ICO's "robust procedures" line.	3–4 weeks
04	**Board and executive governance**	Who owns automated decisions about individuals, what reaches the board monthly, and where the authority to stop a system actually sits.	4 weeks
05	**Threat and integrity response**	Identity verification for colleagues and third parties, tested against the failure mode that took down two grocers in 2025 rather than against a generic control catalogue.	4 weeks
06	**Benefit realisation framework**	The eight numbers in section eleven, instrumented and reported across both estates, with a control group designed in before the next rollout rather than reconstructed after it.	3 weeks, parallel
07	**Build, buy or partner evaluation**	Including the base-rate question put to every supplier in writing, and contractual clarity on who is accountable when an alert is correct and the handling is not.	4–6 weeks
08	**Twelve-month mobilisation roadmap**	Sequenced against the ICO's guidance timetable, the Employment Rights Act outcome, and December 2027 for anyone with European exposure.	2 weeks

The first three form a natural first step of six to eight weeks. That is not an arbitrary window. Sainsbury's moved from a two-store trial to around two hundred stores in fifteen months, and the sector's finance chiefs have already told the British Retail Consortium that 61 per cent of them are investing in automation this year. Six to eight weeks fits comfortably inside the pace this market is already moving at — which is the only sensible test of whether governance work is proportionate.

CONTACT

nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

THIRTEEN

Sources

Every figure traces to one of the following, read in full where the figure mattered. Incidents involving named individuals come from trade press carrying direct statements from the retailer and supplier, quoted rather than characterised.

Regulators and government

- Information Commissioner's Office — AI and biometrics strategy update, 17 March 2026; statement on retail facial recognition given to the trade press, February 2026.
- Competition and Markets Authority — Annual Plan 2026–27, 23 March 2026; Loyalty pricing in the groceries sector, closed 27 November 2024; AA and BSM drip pricing decision, 15 April 2026.
- Groceries Code Adjudicator — CCO progress meeting record, 12 June 2026; annual supplier survey, 23 April 2026; Adjudicator's update, 16 July 2026; best practice statement on forecasting and promotions, updated 1 July 2026.
- Food Standards Agency Science Council — Report of project: artificial intelligence applications in food safety and authenticity, 25 June 2026.
- Home Office — Crime and Policing Act 2026 retail crime factsheet, 19 May 2026. Defra — Food Statistics in your pocket, 29 July 2026; Deposit return scheme: producer and retailer responsibilities, obligations from 1 October 2027.
- Office for National Statistics — Retail sales, Great Britain, June 2026; Consumer Trends, Annual Business Survey, Labour Market Statistics and Family Spending, via Defra and the House of Commons Library. NCSC — Annual Review 2025.
- Department for Business and Trade — Make Work Pay: ending one-sided flexibility, consultation closed 25 August 2026, outcome pending.

Legislation

- Data (Use and Access) Act 2025, section 80, fully commenced 5 February 2026 by SI 2026/82, substituting UK GDPR Article 22.
- Groceries Supply Code of Practice, paragraph 10 (Groceries (Supply Chain Practices) Market Investigation Order 2009).
- Digital Markets, Competition and Consumers Act 2024, Part 4, in force 6 April 2025; Crime and Policing Act 2026; Employment Rights Act 2025; UK GDPR and Data Protection Act 2018.
- Regulation (EU) 2024/1689 (the EU AI Act), Article 5(1)(h), Article 26 and Annex III point 1(a), together with the Digital Omnibus agreement on postponed high-risk deadlines.

Trade bodies, consumer bodies and academic work

- British Retail Consortium — Crime Report 2026, 24 February 2026; Retail employment falls as costs rise, 18 June 2026; CFO survey, 19 February 2026; statement on Guaranteed Hours reforms, 13 August 2026.
- Which? — Trading standards resources and activities, 7 August 2025, from FOI requests to all 187 services in July 2024, benchmarked against the CTSI Workforce Survey 2018–19. CTSI approved ADR bodies register and Retail ADR scheme scope.
- Professor Alan Dix, Professor Emeritus and Director of the Computational Foundry, Swansea University — Facial recognition: what does accuracy mean?, 4 May 2026.
- Big Brother Watch — Landmark legal challenges launched against facial recognition, 24 May 2024. Usdaw — statement on AI, 9 September 2025.

Operators, on their own record

- J Sainsbury plc — Annual Report and Financial Statements 2026; preliminary results and presentation, 23 April 2026.
- Marks and Spencer Group plc — Full year results, 52 weeks ended 28 March 2026; Copilot rollout, 25 March 2026. Co-operative Group — Members' Annual Report 2025. Tesco plc — Sustainability Report 2025/26; AI assistant colleague trial, 9 April 2026; announcement with Adobe, 13 April 2026.

FOURTEEN

About the author

Neil Catton is an independent technology strategist, working across a fractional and advisory portfolio — fractional CTO and CIO, AI board briefings, AI readiness reviews, technology due diligence, data strategy, and interim and programme recovery. He has more than thirty-nine years inside the technology industry, across more than twenty sectors, and has held CTO and Director-level roles at BT, Wipro, Fujitsu, Sopra Steria and Capita.

In retail specifically, he led sector technology strategy at Sopra Steria covering core platform capability from front of house through to back office across major retail clients — ecommerce, store operations, supply chain, finance systems and digital customer experience — including a forward technology proposition for Tesco. At Wipro he acted as fractional CTO adviser to a large European food retail manufacturer and distributor, providing independent challenge to its last-mile transformation programme, which delivered a thirty per cent reduction in technical debt. Earlier work included mobile and loyalty systems for consumer-facing clients, and last-mile retail delivery models across multiple Middle East engagements.

He is the author of the technology ethics trilogy *The Next Evolution*, *The Cognitive Crucible* and *The Shadow System*, and the standalone *Working as Designed: Why systems that work still fail the people inside them*. He writes at writing.neilcatton.com on technology, ethics and human purpose.

CONTACT

nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

How this briefing was produced

Research proceeded outward from the regulator, then the legislation, then trade body and government statistics, then operators' own disclosures, and only then to named academic work. Where a figure mattered, the primary document was read in full rather than a summary of it. Publication dates were checked on every source, because search results routinely present old material as current. Figures that could not be traced to a named, checkable source were excluded, and the exclusions are listed in section one — naming what was thrown out is more useful to a reader than any statistic that survived.

Notice

1. This briefing is not legal, regulatory, financial or professional advice, and should not be relied upon as such.
2. The position is stated as at 2 September 2026. Several matters described here were live at that date and may since have moved — in particular the outcome of the Employment Rights Act consultation on guaranteed hours, which closed on 25 August 2026; the Information Commissioner's Office's draft guidance on automated decision-making and profiling; and the implementation timetable for the EU AI Act's high-risk obligations.
3. Third-party sources are cited in good faith and are listed in section thirteen. Readers should verify any figure against the primary source before relying on it.
4. Named organisations appear only where they are the subject of a regulator or government finding, published statistics, or their own public disclosures, and are used illustratively to describe publicly reported events. No organisation named in this document has been engaged, approached or consulted in its preparation. Individuals named are named because they have spoken publicly about their own experience.
5. No advisory or client relationship is created by receipt of this document.
6. Independence. The author holds no vendor affiliations, reseller agreements, referral arrangements, commission interests or equity positions in any product, supplier or organisation referenced in this briefing, and has received no payment from any organisation named in it. This is what makes the exclusion discipline described above credible rather than merely stated.