

# AI and the end of deflection

**The Drawing Board**  
*The drawing board for a prosperous future.*

|          |    |             |                   |
|----------|----|-------------|-------------------|
| Paper 05 | AI | 4,100 words | Rev. A · Aug 2026 |
|----------|----|-------------|-------------------|

# AI and the end of deflection

*The deflection model of complaint handling is finished. The customer on the other end now has the same AI the company does, and the arithmetic has flipped.*

---

By **Darach Ó Braonáin**

AUG 2026 · 4,100 WORDS

---

## THE ASK

That boards price their complaint-handling exposure in pounds before signing off any customer service strategy built on deflection, automated or otherwise.

Companies bought AI to cut the cost of customer service. Their customers got the same AI for free, and are using it to raise the cost of bad customer service. Most large institutions handle complaints by wearing people down until they give up. That approach now creates legal risk faster than it saves money, and the gap widens as more customers arrive with AI in hand. The organisations that come out of this well will do the opposite of the current trend: give the people who answer real power to fix things, sort problems out first time, and treat every complaint as free information about what is going wrong.

---

## The diagnosis

### Deflection was a design, not a failure

When I moved from Ireland to the UK in 1999, the sharpest difference I noticed was customer service. What I was used to was an empowered person on the other end of the phone whose job was to solve my problem. What I kept finding, as I set up accounts with utilities, banks, and telecoms, was a different machine. The first line seemed to exist to make me go away. The people answering had no authority to fix anything, and outcomes went to whoever persisted longest. I learned the techniques for navigating it, escalation paths, complaint procedures, the letters that trigger obligations. It worked, but it consumed so much time that I could only afford to challenge the things that bothered me most. Everything else I absorbed.

That machine made sense. Complaining was expensive for the complainant, in hours, in knowledge of rights most people never acquired, and in the stamina to sustain a process built to outlast them. Wearing people down was the cost control. The institution did not need to win disputes; it needed most disputes never to be properly started. None of this needs a conspiracy

theory, because the industry writes the design into its own numbers. Deflection rate, contact avoidance, and average handle time are standard performance measures, and every one of them rewards making the customer go away.

The same design shows up in data handling. A Subject Access Request (SAR) answered with more than 100 megabytes of disorganised material, the relevant items buried in irrelevance, is the same trick played with data. The obligation is technically met. The information is practically inaccessible. The complainant is expected to give up.

## AI erased the asymmetry

Every part of that design assumed the complainant's time and expertise were scarce. Neither is true any more.

A customer with a good AI model can now work out where they stand in minutes. They can find the levers: which regulator covers this, what deadline the law sets, which term of the contract was broken. Their first letter can cite the law that makes the complaint enforceable. The 100-megabyte SAR dump no longer works as a wall, because AI reads the lot in hours, cross-referencing and building timelines, where the tactic used to cost the customer weeks of unpaid evenings. The model will also point out that burying information sits badly with the transparency rules the SAR regime is built on, which turns the tactic itself into a fresh ground of complaint.

The effect on throughput matters as much as the effect on entry. Ordinary life generates more legitimate grievances than anyone has hours to pursue: the mis-sold policy, the withheld refund, the service paid for and never delivered. Most of them are simply absorbed. That was my own triage for twenty years: pick the issue that mattered most, pursue it properly, write off the rest. AI has removed the need to choose. Problems I would once have absorbed can now be seen through to resolution, properly and in parallel, because the model does the time-consuming work of understanding each one, drafting the correspondence, and tracking the deadlines. None of this is a professional practice, and none of it uses specialist legal tooling. It is what any equipped customer can now do in the ordinary course of life, and multiplied across every capable, motivated citizen it adds up to demand institutions have not yet planned for.

The homework now happens before any formal process begins. SARs, Freedom of Information (FOI) requests, Companies House filings, and published governance documents let a complainant read the institution's own record first. Institutions have lost the information advantage their complaint handling was built on, and they can no longer assume the person on the phone has not read the documents.

What this means for business is simple. An unhappy customer who would once have vented at a call centre and given up now arrives as a SAR, followed by a breach of contract claim, with the Information Commissioner's Office (ICO) copied in. The complaint has changed legal category, and the choice of category is the point. The same facts can arrive costing very different amounts, depending on how they are framed. A grievance about repeated service failures is feedback an institution can note and file; presented instead as a regulatory breach with an automatic compensation entitlement attached, it triggers legal duties, deadlines, and record-keeping, and

the cost of handling it multiplies. A complaint about how a manager behaves is one conversation; recast as a protected disclosure, where the facts meet the statutory tests, it carries emergency remedies. Knowing which framings exist, and what each obliges the institution to do, used to be professional knowledge. AI hands it to every complainant, and the model knows the most expensive framing available.

## The effort equation has flipped for institutions too

Deflection made sense while going formal cost the customer more than it cost the institution. That has now reversed, and every step up the ladder costs the institution more than it costs the complainant. Giving a customer information voluntarily is cheaper than answering a formal SAR on a legal deadline, and settling a dispute directly is cheaper than losing it in the small claims court. Briefing lawyers for a tribunal costs more again. Escalating now costs the complainant minutes of prompting; meeting the escalation costs the institution professional time, compliance risk, and management attention.

The arithmetic does not need many escalations to flip. Deflecting a contact saves the cost of a phone call, a few pounds at most. A single complaint that reaches the Financial Ombudsman costs a firm a £680 case fee before any staff time is counted,<sup>[1]</sup> a lost small claim costs the judgment plus the preparation around it, and a tribunal costs multiples more. For illustration, if deflection saves five pounds a call and one call in a hundred now turns into a formal dispute costing a thousand or more, the model is under water. AI just changed how often that happens.

A recent consumer case shows the mechanics at scale. In Feb 2026 a customer booked a £25.50 UK train ticket through the Uber app, whose train bookings are run by Omio. The charge went through in US dollars via a US entity, producing a 44 pence currency conversion fee that should never have existed. First-line support denied there had been a dollar charge at all, with the card statement showing it in black and white, described a journey the customer had never taken, and ignored every request to escalate. That went on for the best part of a month. Only then was the complaint taken out of the channel and reframed. Within days an email reached Uber's UK general manager, a complaint went to the Rail Delivery Group asking whether the retailer's licence permits charging UK domestic rail in a foreign currency, and a data protection rectification request started its own legal clock. A senior Omio support lead replied on a Sunday morning with a written admission on every point, a full refund of the charge for a journey already taken, and news of an internal review of payment routing and support handling. Uber's data protection team added a £5 goodwill credit, and the card issuer had already refunded the 44 pence itself, a refund being cheaper for the bank than processing a chargeback. The companies spent hours of staff time, executive attention, an internal review, and a cash refund many times the disputed fee to avoid making good a 44 pence error on day one. All of it was AI-assisted and took minutes to prepare; a few years ago the fee would have joined the pile of things not worth an evening of anyone's life.

The 44 pence was never the real number. Once the refund was refused, the cost of the refusal was set by the complainant, not the company. The handling manufactured the motivation. Nobody pursues 44 pence for the money; people pursue it after a month of being told their own bank statement is wrong. Deflection does not merely fail to close a grievance. It creates the determination to seek proper accountability, in customers who now have the tools to act on it.

Some organisations do the right thing because it is the right thing. Others do it only to avoid the consequences of doing the wrong thing, and for them the equipped complainant has become the incentive mechanism. A customer who can open several fronts at once can price a refusal at hundreds or thousands of times the sum in dispute, until doing the right thing first time is simply the cheaper choice. Nothing in that requires inventing pressure. Every front in the Uber case was an obligation the companies already carried; the complainant's new power is to invoke existing duties all at once, quickly, and without professional help. Deflection used to be the institution's cost control. It now hands control of the institution's costs to the customer it deflects.

When escalating is cheap for the customer and expensive for the institution, "solved first time" stops being a service ideal and becomes the cheapest option the institution has. Every pound spent fixing a problem at first contact now saves several later.

This is why automating bad service misses the point. AI can run a deflection front line for almost nothing, and boards are being shown exactly that saving. The saving is on the cost of talking to customers, which was never where the risk sat. The risk sits in what a mishandled complaint legally obliges the institution to do next, and that is set by the complainant's capability rather than the institution's headcount. Cheaper deflection is still deflection, aimed at customers who no longer have to put up with it.

## The friction was the injustice

Institutions will be tempted to complain that the system is being flooded and that the friction which once kept demand manageable is gone. That reading has the history backwards, because much of what the friction held back was justified.

Special educational needs and disabilities (SEND) tribunals are the clearest evidence. Of the appeals decided at a hearing in academic year 2024/25, 99% were found in favour of the family, the same proportion as the year before.<sup>[2]</sup> There is an honest caveat here. Only families with unusual persistence ever reached a hearing, and nobody can count how many who never appealed would also have won. That is the point. When 99% of those with the stamina to appeal win, stamina, not merit, was deciding outcomes. The rights existed on paper, the effort involved put them beyond most families' reach, and the budgets only balanced because of that.

The gate is now open, because the expertise and hours a tribunal appeal demands are exactly what AI supplies. It drafts the grounds, assembles the evidence bundle, and tracks the deadlines, compressing months of evenings into days. Registered appeals rose 18% last year, to 25,000, and the direction of travel is toward rights being exercised at something like the rate at which they exist.<sup>[2:1]</sup>

The pattern is not confined to education. The ICO has reprimanded a series of public bodies for subject access backlogs running into many thousands of requests, among them the Home Office with around 21,000 overdue.<sup>[3]</sup> Legal duties go unmet at scale because most requesters never push back.

The collapse of that friction is a good thing. People who never had the resources to hold institutions to account now do. If an institution only stayed afloat because people could not enforce their rights, its difficulties now are a correction. The question for leaders is how to run an institution that remains viable when its obligations are actually enforced.

## The demand is already here

The Economist's report of 6 Aug 2026 on Britain's employment tribunals shows the scale.<sup>[4]</sup> Applications for interim relief, an obscure emergency remedy that historically attracted around 20 applications a year across Great Britain, are now arriving at roughly 20 per month in most of the tribunal system's regional offices, according to the tribunal presidents' guidance of 22 Jun 2026.<sup>[5]</sup> The same guidance records that most of these applications fail, that other hearings are being postponed to make way for them, and that the surge often indicates the use of artificial intelligence. Employment tribunal claims rose 39% in the year to Mar 2026, to 50,000, and the backlog of pending single claims rose 55% in a single year, to 64,000.<sup>[6]</sup> Some cases filed today may not be heard until 2030. The same pattern is appearing wherever citizens meet bureaucracies, from Dutch municipal tax appeals to the Canadian privacy regulator to parking tribunals.

Demand is still being legislated upward. The Employment Rights Act 2025 cuts the qualifying period for unfair dismissal claims from two years to six months from Jan 2027 and abolishes the compensation cap for most claims at the same time, with the government's own impact assessment forecasting roughly 17% more claims.<sup>[7]</sup>

Most of the article's examples come from one kind of claimant: the user with no judgement of their own case, for whom AI produces volume and confidence in equal measure. The article expects well-constructed claims to arrive later, once specialist legal AI matures. That timing is wrong. What limits the quality of a claim is the judgement of the person directing the AI, not the capability of the model, and general-purpose models are already enough for a careful complainant to produce forensically structured, procedurally exact cases. My own experience is one data point, but nothing about it requires unusual skill or special tools, and everyone has the same models today. The precision wave is not coming. It is here, and claims of that kind do not strike out.

## The slop wave raises costs for everyone

The flood of poor-quality claims carries its own cost, and institutions should not comfort themselves that incoherent complaints can be dismissed. Volumes will rise steeply, and every AI-generated complaint costs money to read, log, and answer regardless of merit. One employment judge has already received a 282-page claim citing 67 grievances, prepared with Grok.

Ignoring the flood is not an option, because genuine and highly motivated complainants are hidden inside it, and they are precisely the ones who escalate and win. An institution that sorts complaints by how easy they are to dismiss will get its most dangerous cases wrong.



The tempting answer is to fight fire with fire and let AI answer the complaints AI generates. That is the riskiest option on the table. A model answering complaints at scale will make mistakes, and a mistake in a complaint response is not a service failure but a liability event. A wrong statement of the customer's rights, a false factual claim, or an unlawful refusal now lands in front of a complainant whose own AI will read it forensically and convert the error into the next ground of complaint. Machine mistakes at scale, read by the other side's machine, are a losing exchange.

## Autonomy favours the complainant

The current pattern, a person chatting with a model and sending the letters themselves, is the least of what is coming. The step from AI drafting the correspondence to AI managing the complaint is configuration, not research. An AI agent can already send the emails, read the replies, track the deadlines, and escalate on schedule, with the complainant doing little more than tuning it. Customer effort then falls to near zero, and the correspondence lands on the institution's side at machine speed, every message another piece of work someone must handle.

Who can safely automate is decided by the asymmetry of consequences. A complainant whose AI gets something wrong loses that complaint and stands roughly where they would have stood had they never raised it. An institution whose AI gets something wrong in response can breach a legal obligation and end up worse off than if it had never replied. The downside is bounded on one side and open-ended on the other. The bound is not perfect, since courts can award costs against hopeless claims and serial abuse attracts restraint orders, but across the complaint, ombudsman, and regulator stages where most disputes live, the complainant risks little beyond the position they started in. Customers can therefore hand their complaints to autonomous AI long before institutions can safely do the same with their answers, and an institution should only automate its responses if it is certain mistakes will not be made, which nobody deploying current AI at scale can honestly claim.

The result is the worst combination for the deflection model, complaint volume growing at machine speed on exactly the side that cannot automate its way out.

---

## The proposal

The two waves, precision and slop, converge on the same prescription from opposite directions. The defence against both is to resolve problems quickly and well, before customers reach for AI at all.

There is an affirmative case as well as a defensive one. In every walk of life things go wrong, and customers judge institutions on the response rather than the failure. The experience of getting a problem fixed is the brand. Companies spend heavily on marketing to build that brand, then hand its defining moments to the function they have cost-cut hardest. Apple built part of its premium on support that is empowered to solve problems on the spot; Timpson gives front-line colleagues authority to spend up to £500 resolving a complaint without asking permission. AI has

widened the gap between the two possible outcomes of a service failure. A bad response now compounds into a legal dispute, while a good one is becoming rarer, and therefore more valuable, as competitors automate their front lines.

Little in the six moves below is novel on its own; empowered front lines, first-contact resolution, and the attack on failure demand are customer service orthodoxy decades old. What is new is the price of ignoring them. The moves are defensive and offensive at once.

**Solve it before the onslaught.** Fix problems fast and fairly at first contact, before grievances harden into legal claims. Doing so also shrinks the pool of unhappy customers that feeds the slop. Deflection feeds both waves; resolution starves them.

**Empower the first line.** Give the people who answer real authority to resolve, refund, and commit. The empowered front line I met in Ireland in 1999 was never a theorised strategy. Under the new cost arithmetic it is simply the cheaper design, because a first line that can fix problems now costs less than a first line that deflects them. The deflected problem comes back with a legal deadline attached.

**Put AI behind the human, not instead of the human.** The person representing the institution should know at least as much as the complainant now does, with rights, records, history, and remedy options in front of them in real time. That is what AI is for in customer service. Replacing humans with autonomous AI in complaint handling raises the error rate at exactly the moment errors became expensive, and letting AI answer AI is the one set-up to rule out explicitly.

**Treat complaints and SARs as free audit.** A complainant who has read the documents is doing work the institution should want done. Complaints that cite the law are a map of where the organisation is falling short of it, delivered free, before a regulator arrives. Burying information in SAR responses should stop, because it no longer works and because it invites scrutiny under the transparency rules the SAR regime is built on.

**Fix the metrics, and put a price on the exposure at board level.** Deflection rate, contact avoidance, and average handle time now measure how fast the institution is manufacturing legal risk. The measures that matter are first-contact resolution and jeopardy created versus jeopardy retired. The exposure can be put in pounds with existing mechanisms. The Financial Ombudsman Service charges firms a case fee of £680 for each referred complaint, payable in most cases whether the firm wins or loses.<sup>[1:1]</sup> The Financial Conduct Authority's Consumer Duty requires firms to show evidence of good customer outcomes. Ofcom and Ofgem run automatic compensation schemes. The ICO acts on systematic SAR failure, and small claims exposure grows with every customer who can now draft a particulars of claim in an evening. These are UK examples; every developed jurisdiction has its own equivalents, and the argument travels. A board that has done this pricing will not sign off a customer service strategy built on wearing customers down.

**Reduce failure demand at source.** Much of what arrives at customer service is not really demand for service. It is failure demand, in John Seddon's term, contact created by something the organisation got wrong upstream. The billing error, the broken handoff, the process that only



works when nothing unusual happens all manufacture the contacts the front line then has to absorb; in many organisations, customer service exists to cope with dysfunction the business has chosen not to fix. The cost of leaving the dysfunction alone was never zero. It was just low, and hidden in the contact centre budget, and that is the cost AI-equipped customers have repriced. Reaching for a cheaper sticking plaster, automating the absorption of failure, answers the wrong question. The lasting answer is the harder work of digital business transformation, reorganising processes, systems, and data so the failure scenarios stop being created. AI belongs in that work too, accelerating the diagnosis, automating away internal friction, and making transformation cheaper than it has ever been; that, rather than the contact centre, is where the superior return on AI sits. "Solved first time" is the operational goal; the strategic goal is having less to solve, until the customer is happy first time and the contact never happens. The complaints map described above is the transformation backlog, already in priority order, delivered free.

---

## Objections

**Most customers still will not bother.** Escalation does not need to be common to sink the model. The arithmetic flips at low rates, and capable complainants are multiplying what they pursue even if no new complainants appear. Demand at the tribunals is already up 39% in a year.<sup>[6:1]</sup>

**We will use AI to spot the likely escalators and route only those to an empowered team.** This is the strongest version of the objection, and a reasonable transition step. It has two failure modes. Getting the prediction wrong is the expensive event, because the capable complainant looks like every other caller until the SAR arrives, and by then the liability exists. The tiered model also keeps generating the grievances it then has to catch, because the deflected majority remain the pool from which disputes and slop grow. Segmentation reduces the cost of the disease without treating it.

**We will use AI to triage and defend too.** This is an arms race the institution can only win by fixing the underlying practice. AI defending a bad position is a bad position argued for longer at higher cost, with new errors added along the way, and the asymmetry of consequences makes the exchange unequal, because the complainant's mistake costs one complaint while the institution's mistake can create a new breach.

**You celebrate accountability for institutions, then advise them how to face less of it.** Resolution is the accountability. The complainant escalates to get the problem fixed, and an institution that fixes it at first contact has delivered what the escalation exists to force. The advice is to meet obligations early, never to evade them; the difference between resolving a grievance and blocking access to remedy is exactly the difference that matters.

**This is scaremongering.** The tribunal statistics are not a forecast, and the interim relief surge shows how fast an obscure remedy becomes a mainstream tactic once the knowledge barrier falls.<sup>[5:1]</sup>

---

## SOURCES

1. Financial Ombudsman Service case fee, 2026/27 scheme year. ↩ ↩
2. Ministry of Justice, Tribunal Statistics Quarterly, Dec 2025; SEND appeal outcomes and registrations, academic year 2024/25. ↩ ↩
3. Information Commissioner's Office reprimands for subject access backlogs, including the Home Office. ↩
4. The Economist, "The tragedy of the commons, AI edition", 6 Aug 2026. ↩
5. Employment tribunal presidents' guidance on interim relief applications, 22 Jun 2026. ↩ ↩
6. Ministry of Justice, Tribunal Statistics Quarterly, Jun 2026. ↩ ↩
7. Employment Rights Act 2025, government impact assessment. ↩