

For People Development and People Regulation Personnel
within Financial Services


T-C NEWS

COMPETENCE • EXPERTISE • PROFESSIONALISM

OCTOBER 2026

Who needs training? We'll just use technology...

By Adrian Harvey from Elephants Don't Forget



IN THIS ISSUE

The complementary role of Artificial Intelligence in Insurance CPD

By Dr Nathan Matthews, General Manager
UKGI Learning Solutions

Designing training that survives contact with Monday

By Phil Ingle from Phil Ingle Associates

OCTOBER 2026

IN THIS ISSUE

How to contact us

Phone 01361 315 003

Or

Email office@t-cnews.com

Join our T&C Group on LinkedIn

Follow us on Twitter Editor_tcnews

Editor. Jeff Abbott© T-CNews. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording without the written permission of the copyright holder. At T-CNews we have tried to ensure the material in this publication is accurate. However, we cannot accept any responsibility for any errors, inaccuracies or omissions. The views and comments expressed in T-CNews by readers and contributors are not necessarily endorsed by T-CNews. The aim of T-CNews is to provide you with sufficient independent information to make informed choices. However, we cannot accept any responsibility for the quality or fitness of any goods or services advertised.

T-CNews is owned by

2be Development Consultancy Limited

6 Lundy Row, Amble NE65 0ZR

Telephone 01361 315 003

Email editor@t-cnews.com

Web site www.t-cnews.com

- 3 **Rockstead** – AI changes the decision. It must also change the audit
- 4 **Elephants Don't Forget** - Who needs training? We'll just use technology...
- 7 **Credit Services Association** – Breaking down language barriers: Improving support for customers in debt
- 8 **Aveni** – Does Agentic AI count as regulated advice?
- 10 **Compliance Matters** – Data subject access requests
- 12 **Phil Ingle Associates** – Designing training that survives contact with Monday
- 14 **Brand Financial Training** - AI can find the answer. But can it help you become an adviser?
- 16 **Age Wage** – Should DC pensions give retirement freedom or a steady income?
- 17 **Archer Training** – How to run a panel session that people actually remember
- 19 **UK Finance** - The complementary role of Artificial Intelligence in Insurance CPD
- 22 **Standards International** - Planning for 2027: 10 operational priorities every business should be thinking about

Welcome to the October edition of T-CNews. With half of our articles this quarter focusing upon different aspects of Artificial Intelligence it is clear which subject is dominating discussions at present. We still have non-AI based articles that serve as a useful reminder that life exists outside this world and is one we cannot overlook. Enjoy. **Jeff Abbott**

AI changes the decision. It must also change the audit



John Barbour
Chief Executive
Rockstead


For all the discussion about artificial intelligence in financial services, perhaps the more interesting question is not whether firms will use it, because increasingly that feels inevitable, but whether the governance and assurance surrounding its use is evolving at the same speed. The Mills Review is important in this respect because it reflects a regulatory conversation that is beginning to move beyond whether AI should be permitted and towards the much harder question of how organisations remain accountable for decisions when technology becomes an increasingly important part of making them. For lenders this matters enormously. Financial services has spent decades building compliance and audit frameworks around processes that were largely designed, operated and supervised by people. Policies were written, authorities established, samples reviewed and decisions tested against an understood set of rules. Even where technology was involved, there was usually a reasonably clear distinction between the system that processed information and the individual ultimately responsible for interpreting it. AI begins to blur that distinction as models might identify patterns, prioritise cases, interpret documentation, highlight anomalies or increasingly contribute to a recommendation. None of this necessarily removes human responsibility, but it does change what an auditor or compliance function needs to understand in order to determine whether the resulting decision was appropriate. The question is no longer simply whether

the correct process was followed but whether we understand how the process reached the answer in the first place and that changes our notions of assurance in a process. Traditional audit has often relied upon sampling historic decisions and checking them against policy. That remains important, but an AI enabled environment arguably requires something more continuous. Firms need to understand what information models are using, whether that information remains appropriate, how outcomes are changing over time and whether seemingly rational decisions are beginning to produce unintended consequences. Importantly, this is not simply a technology issue because a perfectly functioning model can still produce the wrong outcome if it is trained on inappropriate information or asked the wrong question. Equally, a model may consistently apply a policy that itself creates an undesirable customer outcome. Technology can make decisions more consistent without necessarily making them better and that is a fact that is likely to become increasingly important for compliance teams. Human decision making has always contained inconsistency and AI offers the possibility of reducing some of it, but consistency creates its own challenge because an error can now be repeated thousands of times rather than occurring occasionally. Experienced underwriters bring judgement, context and sometimes an entirely appropriate degree of discretion to a case. In this example, it's easy to see how quickly scale changes the nature of compliance risk. It also changes the importance of governance because firms need to be able to demonstrate who owns an AI enabled process, who is responsible for challenging it, what happens when its behaviour changes and where human intervention remains necessary. Somebody must still be accountable for the outcome. This is where, I suspect, the implications of the Mills Review may prove more profound than another layer of AI regulation. The direction of travel appears to recognise that financial institutions will increasingly use these technologies and that attempting to prescribe every possible application is unlikely to work. Responsibility

consequently moves back towards firms themselves and towards their ability to demonstrate effective governance, oversight and judgement. For boards, risk committees, compliance teams and internal auditors that means AI cannot simply sit within the technology function because if it influences customer outcomes, credit decisions, affordability assessments, servicing, collections or communications then it becomes part of the organisation's risk and control environment and needs to be treated accordingly. Audit will therefore need to evolve from asking whether the system did what it was supposed to do towards asking whether what it was supposed to do remains appropriate which requires a different type of assurance. It means looking not only at individual decisions but at patterns of decisions, understanding changes in data and behaviour and testing whether controls continue to work as models, markets and customers change. AI may ultimately make some forms of decision making easier to audit rather than harder, especially where human judgement can be remarkably difficult to reconstruct. We can record the outcome and perhaps the rationale but unconscious assumptions, inconsistencies and individual interpretation are much harder to identify retrospectively. Properly governed technology creates data, patterns and evidence that can be interrogated at scale. The opportunity is therefore not simply to control AI but to use the transparency it can create to improve assurance more broadly. Technology is exceptionally good at consistency and increasingly good at identifying patterns but wisdom, context and the willingness to challenge an apparently logical answer remain human responsibilities. The institutions that navigate this successfully will not be those that choose between people and technology. They will be those that understand what each is good at and construct their governance accordingly. For lenders, the next stage of AI adoption therefore needs to be accompanied by an equally serious conversation about assurance. The question for compliance and audit is no longer whether AI is coming into the decision making process but whether the controls around it are ready when it does.

Who needs training? We'll just use technology...

By Adrian Harvey from Elephants Don't Forget



Why deploying AI in the customer journey is not a substitute for a competent, confident workforce.

There's a seductive logic taking hold across financial services and beyond: why train people to a high standard when AI can simply supply the answer at the point of need? It sounds efficient. It sounds modern. It is also, I'd argue, a category error that firms will come to regret – commercially, legally, and morally.

1. The state of UK customer service

Let's start with the evidence, because the "AI will fix CX" narrative depends on customer service being broken enough to need fixing, and fixed by the thing being sold. The picture is more complicated than either story. The Institute of Customer Service's UK Customer Satisfaction Index has been tracking the nation's mood since 2008. After a sustained slide – the index dropped to 75.8 in July 2024, its lowest level since 2010 – there has been a modest recovery. By July 2026 the UKCSI stood at 78.3, up a point on the year but only 0.1 points higher than January 2026. That's not a turnaround. That's a plateau after a decade of underperformance, with satisfaction still trailing the highs of the early 2020s. What makes this sobering rather than merely disappointing is the backdrop against which it's happening: an unprecedented, multi-year wave of

investment in customer-facing technology. Chatbots, IVR, self-service portals, and now generative AI have been rolled into contact centres at scale, all promising to "transform" the customer journey. And yet the dial has barely moved. If technology alone were the answer, we would expect to see it in the numbers. We don't.

2. When it matters, people want people

The other inconvenient truth for the "AI will handle it" school of thought is that customers themselves don't buy it – not when the stakes are real. Survey after survey tells the same story with remarkable consistency.

A SurveyMonkey study found 79% of consumers strongly prefer speaking with a human rather than an AI agent, and 84% believe human representatives give more accurate support. Separate research puts the figure even higher for anything consequential: 85% of people would rather speak to a real person than an AI when contacting a business, rising to 89% for law firms – professions, tellingly, where competence and trust are inseparable. Crucially, this isn't blanket AI-scepticism; it's risk-weighted. Over two-thirds of consumers say they'd be uncomfortable using AI for medical advice (69%) or investment advice (68%), even while being perfectly happy to let a bot sort out a food order or a returns request. The pattern holds when researchers move from stated preference to scenario-testing too: when consumers are asked about calling their bank for support, 81% want some level of human involvement in the call.

“An employee who can retrieve an AI-generated answer but cannot explain, contextualise, or exercise judgment over it is not meeting that bar – whatever the screen in front of them says

This is the point that gets lost in the “just deploy the system” logic: customers aren’t indifferent to how a problem gets solved. They are actively, and increasingly, distrustful of firms that lean on AI for anything that matters – and financial services, almost by definition, is where things matter. Trust in a business falls for 57% of consumers when it relies mainly on AI for customer service, and 70% believe service quality would get worse if humans were removed in favour of AI. You cannot underwrite that risk away with a better prompt.

3. The legal reality of competence in UK Financial Services

Here the argument stops being about customer sentiment and becomes about law. The FCA’s Training and Competence sourcebook, the Senior Managers and Certification Regime, and – since 2023 – the Consumer Duty all rest on a shared premise: that the individual delivering advice or service to a customer must themselves be competent, not merely equipped with a tool that is.

The Consumer Duty in particular requires firms to ensure customers are given the support and understanding they need to make good decisions – and regulators have been explicit that outsourcing the substance of that support to a system does not discharge the firm’s obligation to have competent staff overseeing and standing behind it. Certification under SM&CR attaches personal accountability to individuals in customer-facing and advice roles. An employee who can retrieve an AI-generated answer but cannot explain, contextualise, or exercise judgment over it is not meeting that bar – whatever the screen in front of them says. If something goes wrong, “the AI told them to say that” is not a regulatory defence, and it will not read well in a Final Notice. Firms that treat competence as something the technology now owns are building a compliance liability, not removing one.

4. Cognitive outsourcing isn’t new – it’s just better branded

None of this should feel unfamiliar, because contact centres have been here before. Call scripting arrived decades ago on precisely this promise: strip judgment out of the interaction, hand the agent a decision tree, and consistency (and cost control) would follow. Voice prompting and IVR systems did the same thing one layer up – routing customers through menus so that, in theory, fewer decisions needed to be made by anyone at all, human or otherwise.

What’s notable is not that these things happened, but that they didn’t work. Two decades of scripting and prompting have coexisted with the flatlining satisfaction scores above. If stripping cognition out of the front line were the route to better CX, we’d have got there already – repeatedly, and at enormous expense. Generative AI is a more sophisticated version of the same bet: that the answer to inconsistent human performance is to need less from the human. The technology has changed; the wager hasn’t, and neither has its track record.

5. What this does to the people doing the job

There’s a human dimension to this that rarely makes it into the business case. Ask anyone who has worked a scripted contact centre floor what it does to their sense of purpose, and you’ll hear some version of the same thing: it is deskilling, demoralising, and it drives attrition. If the future role of an agent is to relay whatever the AI surfaces, without needing to understand why it’s right, you have not elevated the job – you have hollowed it out. You’ve turned a career into a mouthpiece.

Employees want more than a monitor to read from. They want mastery, some discretion, and the dignity of knowing their judgment matters to the outcome. Strip that away and you get exactly what contact centres have struggled with for years: high churn, low engagement, and a workforce that has quietly concluded the company doesn’t actually need them to be good at the job – only present for it. That’s a retention problem and a culture problem, long before it’s a regulatory one.



6. The law is starting to push back

If UK regulators haven't yet legislated a "right to a human," other jurisdictions are getting there. Spain's Ley 10/2025 – now in force – gives customers the right to request a human agent at any time, requiring bots and AI systems to provide a clear route to a person with full context transferred. The law goes further than sentiment: it requires 95% of calls to be answered within three minutes, on average, and applies not just to essential-service providers but to any company meeting certain size and turnover thresholds, including financial services. This matters beyond Spain's borders. It is a signal of direction. Regulators and legislators are increasingly treating "access to a competent human" not as a customer-service nicety but as a consumer right – which means firms betting their entire CX strategy on automation-first, human-optional models are building on ground that may not hold for long. It would be a brave UK strategy that assumes we will be the exception.

The allure of the system, and the discipline we've abandoned

Put all of this together and the appeal of "just deploy the technology" becomes obvious – and so does its hollowness. A system is measurable. You can point to a dashboard, a deflection rate, a cost-per-contact figure, and tell your board a story of progress. Improving people, by contrast, is slow, unglamorous, and resistant to a single quarterly metric. Too many firms have quietly concluded it isn't worth attempting – that competence is somebody else's problem now that the AI is in the loop. But the evidence above suggests this is precisely backwards. Customer satisfaction hasn't moved despite decades of exactly this kind of substitution. Customers themselves reject it for anything that matters.

The regulator expects genuine, individual competence, not a proxy for it. Employees experience it as an erosion of purpose. And the law, in at least one major economy, is starting to enshrine the opposite principle. The better bet – commercially and ethically – is not choosing between people and AI, but using AI to make people demonstrably, continuously better at their jobs: monitoring competence, closing gaps, and supporting wellbeing, rather than replacing judgment with retrieval. That's the harder path. It's also, I'd suggest, the only one that actually moves the dial. Firms like Elephants Don't Forget exist precisely because that gap between "deploying a system" and "improving your people" is where the real, durable advantage sits – and it's increasingly possible to close it without the trade-off firms have assumed they had to make.

Breaking down language barriers: Improving support for customers in debt

By Colleen Peel from Credit Services Association (CSA)

Effective communication is at the heart of good customer outcomes, yet for many people experiencing financial difficulty, language barriers can make accessing support significantly more challenging.

New research from the Credit Services Association (CSA) highlights the importance of recognising and responding to language needs across the debt recovery and advice landscape. With more than one million people in England and Wales reporting that they cannot speak English well or at all, language barriers remain a real and often overlooked challenge for both consumers and organisations.

Why language matters

The FCA identifies poor English language skills as a potential indicator of vulnerability, and under Consumer Duty, firms are expected to ensure customers can understand the information they receive. However, understanding everyday English is not the same as understanding complex financial terminology, which can create barriers to engagement, informed decision-making and access to support.

What the research found

The CSA engaged with debt advice organisations, language support providers and its member firms to better understand the scale of the issue and how it is being managed in practice.

The findings revealed that language barriers are regularly encountered across the debt advice sector:

- Half of UK debt advisers report working with non-English-speaking clients.
- PayPlan recorded hundreds of interactions requiring translation support.
- Citizens Advice Gateshead reported that around 10% of customers required interpreter services or other language assistance.

The research also found that support available across debt advice organisations varies considerably. While some charities can offer interpreters and translated materials, others rely on alternative approaches such as third-party referrals or authorised family members. This can make it difficult for firms to know which organisations are best equipped to support customers with specific language needs.

What this means for firms

The report emphasises that language needs are not one-size-fits-all. Some customers may only need help understanding written communications, while others may require ongoing interpretation support throughout their customer journey.

Rather than adopting a single standard approach, firms should focus on:

- Identifying potential language barriers early.

- Making it easy for customers to disclose language needs.
- Recording and escalating language-related issues consistently.
- Providing access to appropriate support, referrals or interpretation services.
- Regularly reviewing whether signposting arrangements are delivering positive outcomes.

A range of tools and services are already available to help organisations improve accessibility, including interpretation services, website translation technology and specialist translation providers. The most effective approach will depend on the nature of the business, customer demographics and the volume of cases requiring support.

“ The CSA engaged with debt advice organisations, language support providers and its member firms to better understand the scale of the issue and how it is being managed in practice.

Looking ahead and next steps

The CSA's research concludes that firms do not need a perfect or universal solution. Instead, they need a structured and proportionate framework that helps customers overcome language barriers and access the support they need.

By improving communication and ensuring signposting routes are appropriate, organisations can enhance customer understanding, support better engagement and help deliver fairer outcomes for consumers experiencing financial difficulty.

The next step as part of this ongoing research will be a signposting workshop - Supporting Customers with Language Barriers – Practical Guidance and Q&A - which is taking place on **Tuesday 29 September** from 11.00 – 12.00. The session will explore best practice for identifying and engaging with customers who have language barriers. Attendees will hear from **Citizens Advice** and **PayPlan** about their experience of adapting processes to better engage people with language needs, including how to signpost customers to debt advice effectively. The **Association of Translation Companies** and **LanguageLine** will also speak about the support and solutions available to firms, including the role that AI may play in helping to address language barriers. For further information about the report or the workshop, please contact CSA Policy Manager, Louise Allemange – louise.allemagne@csa-uk.com.

Does Agentic AI count as regulated advice?

By Kent Mackenzie, Chief commercial officer, from Aveni



Agentic AI does not automatically count as regulated financial advice in the UK. The FCA regulates activities rather than technologies.

Whether an AI interaction crosses the regulatory perimeter depends on what the system does, how personalised its output is, whether it recommends a specific course of action and whether it goes on to arrange or execute it.

The issue is getting harder to separate in practice. AI agents can already compare products, use detailed information about an individual, make personalised recommendations and act on a consumer's behalf.

The FCA's July 2026 Mills Review calls this emerging category "**advice-like support**": highly personalised financial support from AI models that sits outside the current regulatory perimeter but may look very similar to regulated advice. It recommends that the FCA examine how the advice and guidance boundary, arranging rules and financial promotion rules work in continuous AI-enabled journeys.

The debate now has two legitimate sides. More regulatory clarity could protect consumers and create a level playing field between regulated firms and general-purpose AI providers. Draw the boundary too widely or AI

too early, however, and firms may struggle to use provide cheaper and more accessible financial support.

Does agentic AI count as regulated financial advice?

The short answer is **sometimes, depending on what the agent is doing**

The FCA's regulatory framework is activity-based. Using AI does not create a separate category of financial advice, and using an AI interface does not remove an activity from regulation.

For investment advice, FCA guidance says a personal recommendation can include a recommendation about a particular investment that is presented as suitable for an individual or is based on that person's circumstances. Those circumstances can include income, objectives, needs and risk appetite. The medium used to deliver the advice does not determine whether it falls within the definition.

That creates a fairly clear position at the extremes.

An AI tool that explains an ISA is providing information. An AI system operated by an authorised firm that considers a customer's circumstances and recommends a particular investment as suitable may be carrying out an activity already covered by financial services regulation. The difficult area sits between those two examples. That is where agentic AI is developing fastest.

Why agentic AI makes the advice boundary harder to apply

Most digital financial journeys used to contain reasonably identifiable stages.

A consumer searched for information. They compared products. They spoke to an adviser or made their own choice. They then completed a transaction.

An AI agent can combine those stages in a single conversation.

A consumer could tell an agent:

"I want to get a better return on my savings, but I need access to the money next year."

The agent might then:

1. Analyse the consumer's finances
2. Identify the relevant products
3. Compare their features
4. Decide which ones appear the most suitable
5. Recommend one
6. Initiate or complete the transaction
7. Continue monitoring the consumer's position afterwards

The FCA describes this development as a move towards financial services that are more **continuous and delegated**, with AI systems progressing from providing recommendations to taking actions within agreed parameters. Its consumer research found that one in five UK adults is already open to AI making decisions within goals they have set.

The regulatory question therefore becomes more complicated than deciding whether one chatbot response constituted advice.

Firms may need to determine which regulated activities are taking place across an entire agent-led journey.

The FCA now has a term for the grey area: "advice-like support"

One of the most important concepts in the Mills Review is **advice-like support**.

The Review defines it as highly personalised support from frontier AI models that currently falls outside the FCA's perimeter, but which would be treated as regulated financial advice if it fell within that perimeter.

This matters because general-purpose AI can increasingly use an individual's own information to generate highly specific financial recommendations.

The Mills Review highlights three problems.

1. Personalisation can blur the line between guidance and advice

The FCA's current framework already distinguishes between different levels of financial support.

Under the Advice Guidance Boundary Review, for example, **targeted support** allows authorised firms to make suggestions to groups of consumers with common characteristics. Simplified advice can go further where a consumer needs a recommendation based on their individual circumstances.

AI changes what can happen between those categories.

A general-purpose model may know considerably more about one person than a traditional guidance service does. It can combine their income, holdings, goals, transaction history and preferences and produce a highly individualised answer. The Mills Review specifically warns that these hyper-personalised recommendations may deliver advice-like



Firms may need to determine which regulated activities are taking place across an entire agent-led journey.

2. Similar financial influence can carry different regulatory obligations

This creates a potential level-playing-field problem. A regulated wealth manager may face strict requirements when making personalised recommendations. A general-purpose AI platform may influence the same consumer's financial decision without carrying equivalent obligations if the activity falls outside the perimeter.

The Mills Review calls out the risk of **regulatory arbitrage**, where unregulated AI platforms exert comparable influence to regulated firms without equivalent consumer protections. HM Treasury's July 2026 Financial Services AI Adoption Plan reaches a similar conclusion. It recommends a review of financial guidance and advice-like outputs from general-purpose LLMs, including their impact on consumer outcomes, competition and the level playing field between regulated and unregulated providers.

3. Agents can move from recommending to acting

A recommendation is only one part of the issue. An agent might also submit an application, move money, arrange a transaction or route a customer towards a particular provider. The regulatory perimeter then extends into other areas. The Mills Review therefore recommends examining:

- Advice and guidance in conversational AI interfaces
- The "by way of business" test for general-purpose AI
- Financial promotions
- Arranging rules
- Continuous AI-enabled consumer journeys

This is why the debate around agentic finance is wider than whether a chatbot has technically given advice. The relevant question is increasingly, and still to be determined is: **what financial activity has the agent performed on the consumer's behalf?**
More to come on this in our next article.

Data subject access requests

By Ian Ashleigh from Compliance Matters



How often do you think about the data that you hold about your clients, where it is held, and the rights they have to see what is in your files.

What do you think you need to provide a client that requests sight of the data you hold about them?

The Data Protection Act 2018 and the UK General Data Protection Regulations (UK GDPR) give individuals the right to access the personal data that a firm holds about them.

These notes provide guidance for firms in the event that a client or their representative submits a request for the data that your firm holds.

The right of access.

The right of access, or subject access, gives individuals the right to obtain a copy of their personal information held by your firm, as well as other supplementary information. Some firms call these a Data Subject Access Request (DSAR) to differentiate them from a Suspicious Activity Report (SAR).

It is a fundamental right for individuals to see all data that you hold whether electronically or on paper. Electronic data includes the files you hold in your back-office systems, email exchanges (internal and external), WhatsApp messages, and, for firms that use Slack, those internal exchanges.

Any request must be in writing. If your firm receives a DSAR you have 30 calendar days to respond and must provide the information free of charge.

If the wording of the request does not clearly identify the information that the individual is seeking, you may write to the individual within three working days asking for further clarification to assist in locating the information or data required.

The reasons for holding client information will be set out in your firm's Terms of Business and in your Privacy Notice. You may refer to these when responding to a DSAR.

What is a person entitled to?

Individuals have the right to obtain the following from an organisation:

- confirmation that you are controlling and/or processing their personal information;
- a copy of their personal information; and
- other supplementary information.

In most cases, you can confirm whether you are processing a person's personal information in general terms. However, this will depend on the nature of the request. If the request is for a specific piece of information, you must confirm or deny whether you are processing this information unless an exemption applies.

What can be redacted?

Firms can redact third-party personal data, legally privileged material, and confidential references from a DSAR response.

Under UK GDPR, individuals have a right to see their own personal data, but not information that compromises others or falls under specific legal exemptions. For example.

- **Other People's Personal Data:** Names, email addresses, contact details, or opinions belonging to other identifiable individuals mentioned in the records. The rule of thumb is to redact information that identifies another person unless they give their permission.

- **Legally Privileged Material:** Communications and advice shared between an organisation and its internal or external legal counsel.

- **Confidential References:** References given (or received) by the organisation for employment, training, or education purposes where confidentiality is expected.

- **Management Intentions / Negotiations:** Records detailing corporate strategy or negotiations where disclosing personal data prematurely would prejudice those discussions. For example, those referring to a merger or acquisition.

- **Health and Safety:** Information that could harm someone's physical or mental health if released.

Crime and Taxation Exemption: Data that, if revealed, would seriously hinder the prevention or detection of crime or the apprehension of offenders.

When redaction is not required

Organisations should balance privacy rights carefully. You do **not** automatically redact third-party data if:

- The other person has given explicit consent to share their information.

It is entirely reasonable in the circumstances to disclose the information without consent. Consider your client's rights to see the data you hold against the third party's privacy.

What is the supplementary information?

At the beginning of this guidance, I listed that you need to provide your client with other supplementary information. Most of this information is within your firm's privacy notice.

- **Purpose:** Why you process clients' personal data.
- **Categories:** The specific types of personal data you hold.
- **Recipients:** Who you share the data with, including outside organisations or specific categories of recipients.
- **Retention periods:** How long you will store the data, or the specific criteria used to determine that timeframe. You may refer to the FCA's document retention rules to provide the relevant detail.
- **Source of data:** Where you obtained their information if you did not collect it directly from them.

- **Individual rights:** Information on the client's rights to:

- request rectification if data is incorrect or out of date;
- erasure if the client wishes your firm to erase data that your firm holds. This will be subject to the FCA's retention rules;
- restriction of processing of their personal data by your firm;
- objection to their data being processed by an automated decision-making process;
- the right to complain to the Information Commissioner's Office (ICO).

- **Automated decisions:** Details about any automated decision-making or profiling used, including the logic involved and the significance or likely consequences for your client. This may include underwriting processes.

Permitted exceptions for charging a fee

Your firm may charge a "reasonable fee" to cover administrative costs in these specific situations:

- **Manifestly unfounded or excessive requests:** in circumstances when the request is repetitive or clearly designed to cause disruption or harassment for your firm.
- **Extra copies:** in circumstances in which the data subject asks for additional copies of information after the initial report has already been supplied.

Rules for charging

- **Administrative costs only:** The fee must be based strictly on the actual administrative costs of providing the extra copies or dealing with the exceptional request. For example, staff time, photocopying, and delivery costs.
- **Burden of proof:** Your firm must be able to justify why a request is excessive or unfounded if challenged. The threshold for this is high.
- **Prompt notification:** Your firm must inform the client about the fee without undue delay. You are not obliged to supply the duplicate or additional report until the fee is paid.

Care when recording information

As you can see, a client is entitled to see everything your firm holds about them. This includes internal emails and internal messages sent over communication platforms such as Slack, Teams, and WhatsApp. Clients may submit a DSAR as part of a complaint when they have already expressed dissatisfaction with your firm. Be careful not to commit to a disclosable medium anything you would not want the client to see, or their Barrister to read out in Court should the client take it that far.

In Summary

Clients have the right to see the data you hold about them. Be aware that this is not restricted to factfinds, meeting notes, and correspondence with product providers. It will also include internal emails and other exchanges with colleagues. In these days of remote working, be aware when you correspond with a colleague your words may be seen by the client.

Designing training that survives contact with Monday

By Phil Ingle from Phil Ingle Associates



Most training is judged on the day itself. The real test is what people are doing differently three weeks later.

In 1885, a German psychologist called Hermann Ebbinghaus published the results of an experiment on the one participant he could guarantee would turn up: himself. He invented a pool of around 2,300 meaningless syllables (WID, ZOF and their friends), memorised lists of them, and then tested himself at intervals to see what had survived.

What he found became known as the forgetting curve. Loss is steepest straight after learning, then flattens out. On his figures, roughly a third of the original effort remained after a day. There are caveats, of course: one man, nonsense syllables, and meaningful material fades more slowly. But the shape has been replicated many times since, including a careful rerun in 2015 that produced a closely matching curve.

It describes, with uncomfortable accuracy, what happens to a great many training days in financial services. You know the scene. Delegates leave buzzing. On Monday they face 200 unread emails, a client call and an IT system that has decided to misbehave. By Friday the photo of the flipchart is in a folder called "Training",

where good intentions go to die.

We're measuring the wrong thing

Ask most firms how they evaluate training and you'll hear about attendance, completion rates and feedback forms. These tell you the room was comfortable and the biscuits were adequate. (If the training is online and delegates were WFH, neither may be the case.) They don't tell you whether anybody did anything differently.

I've reviewed various studies on training effectiveness, including suggestions that only around 10% works. Honestly, this figure is shaky, and varies by industry sector and research methodology. The World Economic Forum's Future of Jobs Report 2025 surveyed how employers plan to assess skills through 2030. The hierarchy:

- **81%** will rely on work experience — the most common method, and rising in importance.
- **48%** will use skills assessments.
- **43%** will continue to use a university degree as a requirement.
- **17%** will prioritize apprenticeships.
- **14%** will consider online certificates in hiring

But the pattern is similar:

That is a better picture, but look at its shape. It's Ebbinghaus with a lanyard. Whatever the true number is, and your organisation may well differ, learning fades unless something holds it in place: that is application. The evidence base is consistent:

- **Role-play and simulation:** 34% better skill application than lecture.
- **Spaced reinforcement:** 60% better knowledge retention.
- **Microlearning:** 80% improvement in long-term retention.
- **Peer-to-peer learning:** 55% increase in knowledge retention.
- **Regular coaching cadence:** 50% higher likelihood of exceeding quota.
- **Applying skills within 24 hours of learning:** 70% higher retention.

This matters beyond the training budget. The FCA's guidance in the Training and Competence sourcebook expects firms to assess training needs at regular intervals and to review the quality and effectiveness of the training they provide, and to review employees' competence regularly and frequently. The wider competent employees rule in SYSC applies well beyond retail activities. An attendance record shows that a session took place. It is not evidence that competence improved.

The event isn't the problem

Here is the twist: the workshop is usually the best-designed part of the process. The trainer (including this one) is skilled, content is sound and exercises are relevant. What's missing are the bits either side.

It's the gym membership problem. A superb personal training session in January, followed by no routine, no accountability and no next session, will produce a lovely afternoon and little else. People don't blame the personal trainer for the abandoned membership in March.

So if the event isn't the problem, what would a training programme that survives Monday look like? Here are four things you can do.

1. Brief the manager before, not after

Baldwin and Ford's model of transfer puts the work environment alongside the individual and the training design, and later research has kept pointing to manager support as one of the strongest influences on whether learning gets used. Yet managers are often the last to know what their team has been taught.

Give them one page: what the course covers, the two or three behaviours to look out for, and one question to ask afterwards. Ask each delegate to arrive with a real situation they want to handle better. It costs almost nothing.

2. Make it real

Generic case studies about a fictional company are a splendid way to learn about that fictional company. If you want a difficult client conversation to go better, practise the difficult client conversation: explaining a fall in a fund to an anxious client, or the compliance officer challenging a senior colleague's shortcut. Build scenarios from your own firm's situations, suitably anonymised, with awkward pauses left in.



An attendance record shows that a session took place. It is not evidence that competence improved

3. Space it out

Ebbinghaus didn't only identify the problem. His work also showed that review is the cure, and that each well-timed revisit slows the next round of forgetting. So build the revisiting in: a short prompt after two days, another after two weeks and another after six, a 30-day action commitment, or a practice partner from the same session.

A PDF summary is not a revisit, but the sort of thing people mean to read.

4. Ask on Monday

This is the cheapest and most powerful step. Within a week, the manager holds a five-minute conversation: *What have you tried? What happened? What will you try next?* No form and no system. A conversation that says, in effect, "this mattered enough for me to ask."

Measure what you actually want

If the aim is different behaviour, measure behaviour. Sample calls or meetings against two or three specific behaviours. Using AI meeting summaries and recordings makes this process way easier than the meetings I used to observe when I worked in banking. Ask managers what they've noticed. Have delegates rate themselves before and after, then compare those ratings with the manager's view. The gaps are often the most interesting bit.

Evidence like this is far more persuasive than a completion certificate when someone asks how you know your people are competent.

Designing the Monday

The forgetting curve is well over a century old. The cure (review, practice and support) is barely younger, and needs no new software licence. What it requires is deciding that the training event is the start of the process rather than the whole of it.

A simple test: before your next programme, write down what should be different by Monday week and who will notice. If you can't, you haven't designed a training programme. You've booked a day out.

References

- [The forgetting curve: Hermann Ebbinghaus \(1885\)](#)
- [Baldwin & Ford \(1988\), Transfer of Training: A Review and Directions for Future Research, Personnel Psychology](#)
- [Georgenson and the 10% figure \(Work-Learning Research\)](#)
- <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>
- [FCA Handbook, TC 2.1 Assessing and maintaining competence](#)
- [FCA: Training and competence](#)

AI can find the answer. But can it help you become an adviser?

By Jennifer Dugdale from Brand Financial Training



Artificial intelligence has quickly become part of everyday life. Now used in almost all areas of business; drafting emails, summarising documents, explaining technical concepts, recording meetings - AI is changing the way we work. It is also changing the way we learn.

It isn't surprising that candidates studying for professional qualifications are embracing these tools. Used well, they can explain difficult concepts in plain English, create revision plans, generate practice questions and even test understanding (though, usually unreliably). For busy professionals trying to fit study around work and family life, that's an incredibly attractive proposition. The question, however, isn't whether candidates should use AI.

The question is whether we're beginning to mistake *finding* knowledge for *owning* it.

Professional exams have always required candidates to demonstrate technical understanding. Increasingly, though, we're seeing a generation of learners who can access information almost instantly. Need to know the Annual Allowance? Ask AI. Forgotten how tapering works? Search for it. Unsure how Business Relief operates? A chatbot can produce an explanation in seconds.

On the face of it, this sounds like progress but it raises an interesting question.

Once you overcome the problems of hallucinated information, out of date rates and allowances, or answers relating to foreign legislation: If candidates never have to wrestle with the material themselves, are they actually learning it?

Psychologists refer to something called the "illusion of competence". It occurs when information feels familiar because we've recently seen or read it, leading us to believe we understand it more deeply than we really do. Reading an AI-generated explanation can create exactly this effect. The content is often clear, well-structured and convincing. It feels as though learning has taken place. Yet when candidates sit down in an exam without access to AI, many discover that recognising an answer and recalling it independently are two very different skills. Professional examinations are not designed to test whether candidates know where to find information. They assess whether candidates can retrieve relevant knowledge from memory, interpret the question, and then apply that knowledge appropriately under pressure.

Those are cognitive skills that still require practice. This is particularly true as candidates progress towards more advanced qualifications. At this stage, success depends less on remembering isolated facts and more on understanding how technical concepts interact. An adviser may need to consider taxation, pensions, trusts and investment planning simultaneously within a single client recommendation. AI can certainly explain each of those topics individually. Whether the learner has developed the judgement to bring them together is another matter entirely. That distinction becomes even more important once candidates enter practice because clients rarely ask textbook questions. They present complex, often incomplete situations requiring interpretation, prioritisation and professional judgement. The value of an adviser has never been simply knowing the legislation. It lies in understanding how that legislation applies to the individual sitting opposite them. AI cannot develop that judgement on someone's behalf. This is not an argument against using technology. Quite the opposite. Future advisers will almost certainly need to become highly proficient users of AI. Firms are already exploring how artificial intelligence can support research, summarise information, improve productivity and enhance client communications. Some are even using AI to write suitability reports and manage client communications. Some of these tools will undoubtedly become part of everyday professional life. However, the role of AI should be to support thinking, not replace it. There is an important difference between asking AI to explain a difficult concept and asking it to answer a practice question for you. The first supports learning. The second risks bypassing the very mental effort that creates learning in the first place. Educational research has consistently demonstrated that retrieval practice (forcing ourselves to recall information without prompts) is one of the most effective ways to strengthen long-term memory. The struggle involved in remembering isn't a sign that learning has failed, it's the process through which learning occurs. Candidates can now receive instant answers without ever experiencing the productive difficulty that helps knowledge become embedded. While this makes studying feel easier, easier does not necessarily mean more effective. There is another consideration too. Generative AI is impressive, but it is not infallible. Financial services is built on legislation, regulation and taxation that evolve continually. AI systems can present outdated information or generate entirely plausible answers that are simply wrong. Perhaps more concerning is that they often do so with remarkable confidence. Candidates without sufficient technical understanding may not recognise those inaccuracies. Ironically, the people most likely to benefit from AI are often those who already possess a solid technical foundation, because they have the knowledge needed to challenge, verify and interpret its outputs.

“ The question is whether we're beginning to mistake *finding* knowledge for *owning* it.

Rather than asking whether AI belongs in learning, we should be asking how candidates can use it responsibly. Can it help explain a concept? Absolutely. Can it generate revision questions? Certainly. Can it challenge understanding or identify weaker areas? Without doubt.

But should it become a substitute for building technical knowledge, practising retrieval or developing professional judgement? Probably not.

Professional qualifications have never simply been about passing an exam. They are about laying the foundations for a career in which clients place enormous trust in the advice they receive.

Technology will continue to change how we learn, just as it is changing how we work.

The challenge for candidates is ensuring that convenience doesn't come at the expense of understanding. Because ultimately, the best advisers won't be those who know how to ask AI the right question

Should DC pensions give retirement freedom or a steady income?



Henry Tapper

Chair of AgeWage,
Pension Playpen and
CEO of Pensions Mutual

This is not a new argument. We've been having it all century and in 2014 many thought it had been settled by George Osborne. But pension freedoms haven't been a success for the silent majority who do not have advisers, have no aspiration to be wealthy and just want enough to have a retirement with holidays and a little security if health becomes a problem. I remember in the 1980s when I started advising people to buy pensions (if they didn't get a company one) that three quarters of the pension was paid as a lifetime income and the rest was to take care of health problems. I suspect I am not the only one who considers my pot in these terms.

The trouble is that while the people who bought DC pensions in the 1980s were professionals who worried they could only be 17.5% of pensionable earnings (a lot), those who have been saving have not been saving anything like that.

Auto-enrolment had increased the numbers saving into pensions, but it had introduced a new level of contributions that slowly rose throughout the last decade. But it stopped increasing in 2018 and has stuck at 3+4+1 of a band of an employer's earnings. We thought that we would see a further increase, were promised an increase in minimum contributions to be made by employer and employee by the middle of this decade, but it hasn't happened.

When auto-enrolment was introduced, Steve Webb (the pensions minister at the time) told us that the DWP would monitor the amount of income that people could purchase through the following decades. The assumption in 2012 was that all auto-enrolled savings would convert to pensions (except for those who had enough income to be allowed to spend their pot as they chose).

Steve Webb was wrong, the 2014 budget abolished the link between pension saving and a lifetime income from retirement. For a time, people saw their pot as "wealth"

and freedom took over as the mantra of DC pension providers. The Government as recently as 2024 were keen to define "value for money" by the rate that a pot built up. With this focus on the size of the pot, the worries about adequacy were put to one side.

All this has changed since the return of a Labour Government; we have had a Pensions Commission tell the country that 15m of the country are set to retire on an inadequate income. The Pensions Dashboard has been unveiled to demonstrate what we have built up, not as a pot but as the income a pot could buy. In line with this we have had a Pension Scheme Act that requires that everyone who doesn't choose to exercise freedom be defaulted into "a guided retirement pathway". This sounds too much like the Liverpool Care Pathway for my liking.

For many advisers, the guided retirement pathway is an anathema. It requires no engagement or any "attention to the pension". It is designed to pay an income that lasts as long as the path taker and is the responsibility of trustees or the contract managers of workplace personal pensions. On the face of it, workplace pensions are back as a lifetime income. Except the longevity insurance has been pushed back into later life (Nest have opted for 85). The agreed guided retirement pathway is flex and fix with the flex a guided drawdown and the fix, a late life annuity.

What will be the view of a generation of workplace pension savers? Will they succumb to the default pathway or find an alternative that suits them better. If I was back advising, I would be asking people whether they want control of their savings or to allow it to pass to my workplace pension provider. There has long been an assumption that 80% of people have no interest in taking control of their retirement finances, while 20% will split between those who want to work with an adviser and those who want to take their own decisions. This Government is happy to see those who have a wish to DIY, take their money wherever it suits them. It is happy that 80% of the working population can follow pathways rolled out to greet them as they reach their normal retirement age. DC is now a choice for all employers. As it is, an employer can choose to move workplace pension from DC to CDC. The CDC workplace pension makes retirement freedom even harder to achieve and post budget 2014 few would have thought that such a pension would be being promoted. But there appears to be considerable demand for CDC as a collective pension when it comes on the market next year.

It's emergence as a non-guaranteed undefined pension is a mark of how far the mood of employers and the silent majority has changed. The 80% have moved away from freedom and towards a requirement for a steady lifetime income.

How to run a panel session that people actually remember

By Paul Archer from Archer Training



What a red sofa on a Friday night can teach us about running better conversations I've always thought Graham Norton makes it look far too easy.

He sits there on that famous red sofa, surrounded by Hollywood A-listers, a comedian who's had two coffees too many, and someone promoting a film nobody's quite sure about. Within minutes, everyone's relaxed, stories are flowing, the audience feels involved, and before you know it, you're laughing along.

No one in the audience asks a question, yet somehow you feel part of it. And then you remember the last panel discussion you sat through. Four bar stools on top of the stage. A row of polite experts. Sensible questions. Even more sensible answers. And everyone's looking at their phones after 10 minutes.

Over the years, I've come to see that running a tight panel is less about asking good questions and more about creating the right conditions. When it works, it feels effortless. So if the conference organiser ever asks you to run a panel and thinks it's easy and about getting AI to do you a list of questions, think again. Graham Norton makes it look really easy.

Recently a client asked me to chair a 30-minute panel session at their conference. They challenged me to bring my facilitation techniques to the fore because they didn't want the usual "fireside chat" format.

It was a fairly standard panel titled *What are the key challenges facing mortgage markets?* On the panel, I had two experienced Business Development Managers from the corporate world, senior people too, and a successful head of a boutique mortgage brokerage, so lots of experience.

So I did, and the look on the faces of the esteemed industry professionals on the panel was a hoot. Only joking. I briefed them beforehand, and they loved the energy. Here's how I approached it.

First, I involved the audience before the panel had even properly started.

"Raise your hand if you might have a question. Just to give me a feel for the room. That's great. My job's done because we already have an agenda."

This sets the audience up straight away as they are responsible for asking questions, not you. Warm them up some more.

"Stand up, or raise your hand, if you've used AI in your business this year. Stand up if you've had a client shocked by their monthly payments."

“Panels need direction, not a script. When everything is pre-planned, it shows

Leaving the audience until the final five minutes is usually too late. I prefer to involve them early, even if it's just a quick show of hands or a simple question to the room. They stop watching and start participating, even silently.

Then I introduced a Red Card. Give each panellist a red card and tell them that if somebody on the panel says something they disagree with, they can show the card and say why. Make sure they know this is tongue-in-cheek. If everyone politely agrees with each other, the audience learns very little. My job is to stir things a little but keep it respectful.

I also use what I call Audience Pick a Card. Print questions on individual cards and create a deck. Then walk the room, pounce on someone and ask them to pick a card. Any card. Then ask them to read the question out loud.

The question has to be aimed at one panel member, not all three or four. The audience member has to choose who answers. You can bounce it to the others afterwards if you want, but I rarely do. More often than not, it turns into an echo chamber, with everyone agreeing and adding very little new.

Speak to the room, not to me. I always remind panellists that this isn't an interview. The moment they turn their answers towards me, the energy drops. When they lift their eyes and talk to the audience, it feels more open, more inclusive, and more like Graham.

I also try to break the rhythm on purpose. Short answers. Quick follow-ups. Humour. A gentle put-down. Sometimes a challenge. You want an interesting discussion so keep them on their toes.

"That's what a corporate communications team would say, Bob; what's the real answer?"

Or my personal favourite: "I don't quite believe you. Have another go." Then turn to the audience and ask: "Who's got a question that might make the panel slightly uncomfortable?"

Panels need direction, not a script. When everything is pre-planned, it shows. Answers become rehearsed and slightly distant. That way you can engage in surprises and off-the-cuff remarks, real human interaction.

Sometimes I nominate one panel member and move their chair three feet ahead of the others. "For the next four minutes, Jane, you're on the hot seat, with no support from the others."

I also like the Empty Stool Rule. Have an empty bar stool alongside the panel. If someone in the audience strongly disagrees, give them the stool for two minutes. Some of the best insights in the room never made it onto the stage in the first place.

What struck me afterwards was that some of the most interesting observations didn't come from the people on the stage. They came from the audience.

It turns out that conference organisers spend months finding three experts and then sit 200 others facing the wrong way.

A well-run panel rarely earns you a round of applause. The panellists get the credit, while you quietly hold the whole thing together. It's a subtle craft, and when it's done well, nobody quite notices, apart from you.

Most panel sessions are remembered for who was on the stage. The best ones are remembered for what happened in the room. Which is probably why I keep coming back to Graham Norton. That slight grin, the naughty boy look, the well-timed nudge.

He makes it look effortless. It isn't.

About the author

Paul Archer facilitates workshops, chairs conferences and spends an unhealthy amount of time thinking about what makes people engage, contribute and remember. Over the years he's worked with organisations throughout the UK helping leaders, teams and businesses communicate more effectively.

If you'd like to continue the conversation, you'll find Paul on LinkedIn or at www.paularcher.uk

The complementary role of Artificial Intelligence in Insurance CPD

By Dr Nathan Matthews, General Manager UKGI Learning Solutions



Artificial Intelligence (AI) is seldom out of the headlines. Some suggest it promises a utopian future, whereas others predict far darker consequences. What is clear, however, is that AI is already having a transformative impact on our economy and society and we cannot ignore it. This includes the effect of AI on how we work and how we learn. This article explores how generative AI can complement workplace learning, with specific reference to how insurance professionals can use the technology effectively.

From the outset, it must be emphasised that education within the insurance industry should remain grounded in credible and authoritative sources. AI does not replace CII qualifications, textbooks or e-learning platforms. These provide structured, often expert-curated learning programmes aligned to industry needs and should remain the foundation for professional study. Workplace practice, mentoring, regulatory publications and the industry press also remain important in achieving professional competence.

Generative AI can, however, play a complementary role in developing knowledge and improving understanding. Applied carefully and purposefully, it can question,

challenge and explain, simulate situations and provide real-time feedback. Many of these uses also reflect established principles and approaches within learning theory.

AI has attracted legitimate criticism in education. It can generate unverified learning materials containing errors, invented information or dubious citations that learners then rely upon. These are often referred to as 'hallucinations'. Learners may also use AI to generate plausible answers to assignments without properly engaging with the questions. Irrespective of the veracity of the output, generating an answer is not the same as employing the knowledge and reasoning required to reach it independently. Used in this way, AI may reduce the cognitive effort involved in learning.

The educational value of AI therefore depends heavily on how it is used. A useful principle is to use it in ways that make the learner think.

Most of us have likely used a Large Language Model (LLM) such as ChatGPT, Gemini, Copilot or Claude. Useful outputs depend on effective prompts, entered through text or increasingly voice.

When prompting, think about what you want to achieve. It is useful to give the LLM a role, such as that of a supportive tutor. Give it clear instructions, relevant context and any source material it should use. You can also add constraints or parameters, such as telling it not to reveal an answer before you have attempted the question, or give it examples of how the output should be structured.

Suppose you wanted to assess your knowledge of a particular topic. You could ask the LLM to question you and provide feedback. Someone new to insurance could use the following prompt:

Assume the role of a tutor providing training to insurance professionals. I am new to the insurance industry and work as an apprentice in a general insurance brokerage in the UK. Ask me ten questions relating to the fundamental principles of insurance. Ask one question at a time, allow me to respond, provide feedback and then move on. Do not reveal the answer before I respond. Then assess my overall knowledge and identify any areas in which I need to improve.

This provides immediate feedback while requiring the learner to recall information from memory. Reading course resources or completing e-learning materials can create familiarity with a subject, but familiarity does not necessarily mean that knowledge can subsequently be recalled or applied.

This utilises retrieval practice, where learners actively recall information independently rather than having to refer to notes or learning materials. The feedback also provides useful formative assessment for the learner, helping them identify strengths, knowledge gaps, and misunderstandings while learning is taking place.

AI also allows the difficulty to increase at the learner's own pace. You could move from simple questions to scenarios requiring application. The AI can be prompted to provide 'scaffolding' by offering hints or explanations when needed, then gradually reducing that support as the learner becomes more confident and capable of applying the knowledge independently.

Another way AI can support understanding is through its ability to explain concepts in different ways. This is particularly useful in insurance, where legal terminology, policy language and industry jargon are frequently used.

Consider the following statement:

Where an insured under a non-consumer insurance contract has failed to make a fair presentation of the risk, the insurer will have a remedy only if it can show that, but for the qualifying breach, it would have declined the risk or accepted it on different terms. The remedy depends on the nature of the breach and, where it was not deliberate or reckless, what the insurer would have done had a fair presentation been made, in accordance with the Insurance Act 2015.

This may be clear to an experienced practitioner, but less so to someone early in their career. Repeatedly asking a human tutor or colleague for clarification can cause social discomfort but using AI allows this to be done multiple times in a virtual environment, where learners can ask for repeated clarification without fear of embarrassment or judgement.

“ The educational value of AI therefore depends heavily on how it is used. A useful principle is to use it in ways that make the learner think.

A learner could prompt:

Explain the following technical insurance information in plain English for someone who is new to the insurance industry. Explain technical terms clearly, but do not omit important qualifications or change the legal meaning.

The learner can ask for a different explanation if the first one does not help and continue asking questions until they understand the concept.

This is consistent with cognitive load theory, which recognises the limited capacity of working memory. Complex subjects can be easier to understand when broken into smaller components and introduced progressively.

The learner can then go further and explain the concept back to the AI and ask it to identify anything missed or misunderstood. This self-explanation requires the learner to organise and articulate their understanding before progressing with the original learning material.

AI can also support learning about a particular insurance product or service. Care is needed when uploading information to an LLM. Personal, confidential or commercially sensitive information should not be uploaded to a public LLM. Documents should be uploaded only where your organisation permits it, to AI systems authorised by your employer, and in accordance with its information security, data protection and AI policies.

In insurance, where permitted, you could provide approved policy documentation and ask the LLM to simulate a prospective customer interaction:

Using only the policy documentation I have provided, assume the role of a customer considering this insurance product. Ask me questions about the cover, significant exclusions and important conditions.

After each answer, compare my explanation with the documentation provided. Tell me if I have omitted an important point and identify the relevant section of the document.

This can provide coaching for client interactions in a controlled and low-risk setting. Relevant FCA rules or guidance could also be provided, or signposted, to the LLM and used to prompt review of aspects of the interaction relating to Consumer Duty outcomes, such as consumer understanding or consumer support. However, any regulatory interpretation produced by the AI should still be checked against the FCA Handbook, FCA guidance and the firm's own compliance procedures.

Using AI to simulate potential client conversations can allow learners to practise realistic situations in a safe environment. They can participate in a mock client conversation, receive feedback, consider what they could improve and repeat the exercise to develop their performance. Using different customers and circumstances also requires learners to adapt their approach. This reflects experiential learning, where skills are developed through practice in realistic contexts. AI can also encourage learners to reflect on their own thinking. After a scenario, it could ask why you chose a particular response, how confident you were or what

“ As such, AI works best as a tool to complement and enhance traditional learning resources and should not be used as a substitute

further information you would need before advising the customer. This helps learners identify the limits of their knowledge and examine their reasoning.

There is, however, an important qualification to these examples. An AI system can give incorrect content or feedback. It may tell a learner that a wrong answer is correct, oversimplify a technical principle or produce a confident but inaccurate regulatory interpretation. As such, AI works best as a tool to complement and enhance traditional learning resources and should not be used as a substitute. Where accuracy is important, the learner should return to the authoritative sources and the relevant learning materials to verify the position.

The educational potential of generative AI lies particularly in its ability to interact with individual learners. Used carefully, it can challenge understanding, simulate practical situations and provide immediate, non-judgemental feedback. It can also encourage learners to reflect on their reasoning and engage more actively with authoritative knowledge and established CPD methods. Professional qualifications, authoritative sources and workplace expertise should remain the foundations of insurance learning. Generative AI can complement them by helping learners make sense of complex concepts, retrieve knowledge, explain their reasoning, apply what they have learned to realistic situations and adapt their responses. Its educational value is greatest when it supports understanding while requiring the learner to think.

Dr Nathan Matthews is General Manager of UKGI Learning Solutions - the provider of the Development Zone e-learning platform and Searchlight Training Solutions for the UK general insurance industry.

Planning for 2027: 10 operational priorities every business should be thinking about

By Michelle Hoskin, CEO, Standards International & Business and Operations Management Network



Growth, consolidation, succession or simply running a better business – whatever 2027 looks like for you, your operations need to be ready for it.

There is something about a new year that makes businesses want to move.

New targets. New technology. New people. New projects. New ideas.

But before you start adding, **stop and look at what you already have.**

Because one of the biggest mistakes I see businesses make is planning for where they want to go without being honest about whether their current operation can actually get them there.

And going into 2027, that matters more than ever. Businesses are navigating rising expectations, tighter margins, rapid developments in AI and technology, changing workforce needs, increasing regulatory scrutiny and clients who expect things to happen faster and more seamlessly.

So before another year of *doing* begins, here are ten things worth getting right.

1. Know where your business really is

Before setting another target, establish your starting point.

Where are the bottlenecks? Where are people compensating for poor processes? What relies too heavily on one individual? Which systems no longer work for the business you've become?

A business can look successful from the outside while carrying significant operational weaknesses underneath.

You can't build a meaningful plan without an honest baseline.

2. Plan capacity, not just headcount

Resource planning isn't simply asking whether you need another person.

Look at **capacity, capability and demand.**

Where is work actually sitting? Which skills will you need next year? Are highly skilled people spending too much time on low-value activity? Could work be redesigned, automated, outsourced or removed altogether?

Recruitment may be the answer. But don't automatically assume it is.

3. Get serious about operational resilience

A business continuity plan sitting in a folder isn't enough. Ask what happens if a key person disappears tomorrow. A major supplier fails. Your systems go down. A cyber incident hits. Your biggest client leaves. Or demand suddenly increases.

Then ask the uncomfortable question:

Could the business actually keep operating?

Resilience needs to be tested, reviewed and owned – not simply documented.

4. Make AI part of the business plan – Not a side project

AI is no longer something to leave with the tech enthusiast in the business.

But neither should businesses rush to implement every shiny new tool.

Where could AI genuinely remove friction, improve productivity or create capacity? What data is being shared? Who is responsible for governance? How will outputs be checked? And do your people understand how to use it properly?

AI won't fix a bad process. It may simply help you do the wrong thing faster.

5. Challenge your technology stack

2027 should not be another year of paying for systems nobody fully uses.

Look at your CRM, workflow, communications, reporting, HR, finance and client-facing technology.

Are your systems integrated? Is information being entered multiple times? Are manual workarounds becoming normal? Are you paying for functionality you've never implemented?

Technology should make the operation **simpler**, not create another layer to manage.

6. Follow the money – and the waste

Cost control isn't about cutting everything.

Look for operational waste: duplication, unnecessary subscriptions, inefficient processes, rework, poor supplier arrangements and expensive people doing work that shouldn't require their expertise.

The question isn't simply:

"Where can we spend less?"

It's:

"Where are we spending without getting enough value?"

That's a much better conversation.

7. Put Your Client Proposition Under Pressure

Don't assume that because clients valued something three years ago, they value it today.

Review the complete client journey.

What do clients genuinely use? Where is there friction?

What takes too long? What could be simpler? Where are expectations changing?

And importantly, ask the people closest to the client.

They will often spot problems long before they appear on a management report.

8. Make Your Data Useful

Most businesses don't have a shortage of data. They have a shortage of **useful information**.

Decide which measures genuinely tell you whether the business is performing well.

“ Leadership teams can spend so much time working in the business that nobody has enough space to look properly at it

Capacity. Profitability. Service levels. Complaints. Client retention. Workflow. Productivity. Risk. People.

Then make sure the underlying information is accurate. A beautiful dashboard built on poor-quality data is still poor information.

9. Build a business that doesn't depend on heroes

If one person leaving would cause chaos, you don't have resilience – you have dependency.

Document critical knowledge. Develop people. Create clear ownership. Review succession. Make sure decisions and processes aren't sitting inside someone's head.

This becomes even more important if succession, acquisition or sale is anywhere on the horizon.

A valuable business should be able to operate without constantly relying on the same few people to save the day.

10. Create space to actually run the business

Perhaps the biggest one.

Leadership teams can spend so much time working in the business that nobody has enough space to look properly at it.

Don't make your annual planning session a two-hour meeting squeezed between everything else.

Create proper time to review your people, processes, technology, finances, risks, client proposition and priorities.

Decide what you're going to **start, stop, fix and protect**.

And then turn it into an operational plan with owners, measures and deadlines.

2027: Less doing. More deliberate.

Businesses don't usually struggle because nobody is working hard enough.

They struggle because effort is being spent in the wrong places, operational problems are tolerated for too long, and planning gets pushed aside by whatever feels urgent today.

So before 2027 gathers pace, stop.

Look properly at the business you've built.

Understand what's working. Call out what isn't. Decide what needs to change. And make sure your operation is genuinely capable of supporting whatever comes next.

Whether that's **growth, succession, acquisition, sale or simply running a stronger business**, good operations will sit underneath it.

Plan it. Test it. Strengthen it. Then move.



would like to thank



**The support from our sponsors help ensure we bring
the content of T-CNews to you for free.**