Who Will Become the Experts?

Part 3: Who Will Train Tomorrow's Experts?

The employer's dilemma: productivity today or expertise tomorrow?

By Dr Peter Fields

THE 2-MINUTE SUMMARY

AI may reduce the number of junior employees a business needs today without removing its need for experienced professionals tomorrow. But preserving traditional trainee roles simply because they have always existed makes little sense. Some junior work develops valuable expertise; other activities consume time without contributing much to learning.

Research suggests that AI can accelerate genuine learning, but producing better work with AI does not necessarily mean the person using it is becoming more capable. The challenge is to identify which activities develop professional judgement and whether AI can support that development more effectively than traditional methods.

The practical answer is to redesign training around the capabilities employers will actually need. This means combining AI with structured practice, appropriate supervision and feedback, while measuring genuine competence rather than simply productivity. Employers should also monitor their future skills requirements before shortages emerge. There is no established universal model, and neither regulation nor compulsory training automatically provides the answer. The most defensible approach is to experiment, measure what works and adapt accordingly.

Imagine running a professional-services business with 200 employees. You have invested in AI, and the results are encouraging. Documents are produced more quickly, routine analysis takes a fraction of the previous time, and experienced professionals can handle work that once required several junior colleagues.

Your finance director presents the next year's budget. Revenue is holding up, productivity is improving, and there is an obvious opportunity to reduce recruitment costs. Do you really need another twenty graduates when AI can undertake much of the work you would otherwise give them? The commercial argument is persuasive: employing people to perform unnecessary work is hardly a sensible business strategy.

But your head of department raises another question. Five years from now, who will be managing the clients, making difficult judgements, supervising complex work and leading the business? You could recruit experienced people from competitors. Perhaps that would be

cheaper than training them yourself. Of course, your competitors might reach the same conclusion.

This is the dilemma at the heart of the third article in our series. **AI may reduce the number of junior employees a business needs today without reducing its need for experienced professionals tomorrow.** The answer is not necessarily to preserve the existing career ladder. That ladder may be expensive, inefficient or poorly designed. But neither can businesses safely assume that the experience it once provided will somehow appear elsewhere. We need to examine the problem from the employer's perspective, recognising that commercial realities matter just as much as educational theory.

The problem with the traditional professional pyramid

Many professional-services businesses have historically followed a familiar structure: a small number of partners or senior specialists supervise experienced professionals, who in turn oversee more numerous junior staff. Senior people manage clients, exercise specialist judgement and supervise delivery. Junior employees undertake research, preparation, analysis and routine tasks. The firm earns revenue from work carried out at different levels of seniority.

This arrangement also serves an educational purpose. Juniors gradually encounter increasingly difficult problems and, with effective supervision, develop the knowledge and judgement required to progress. The productive and educational functions of employment have therefore been closely connected. But there is no reason to assume this represents the best possible system of professional development. Some junior work is repetitive without being particularly educational, and time recorded against client matters does not measure learning.

AI potentially changes the calculation. If experienced professionals can produce the same output using fewer juniors, businesses may no longer require such a broad base of staff. The need to develop future partners, managers and specialists has not necessarily vanished; its cost may simply have become more visible. **The question is whether developing expertise remains a commercially valuable by-product of junior employment, or becomes an investment that employers must make more deliberately.**

What are businesses actually doing?

The evidence is more varied than popular accounts of mass replacement might suggest. An Office for National Statistics analysis published in July 2026 found that reported AI use among UK businesses employing ten or more people had increased from approximately 12% in late 2023 to around 35% in June 2026. The ONS found no widespread change in overall headcount attributable to AI at that stage; among businesses reporting a lack of AI expertise, training

existing staff was a common response. These findings concern AI broadly, not exclusively generative AI or graduate recruitment. [1]

Employers nevertheless expect significant change. The World Economic Forum's *Future of Jobs Report 2025*, based on a survey of more than 1,000 substantial employers covering 55 economies, reported that 85% planned to prioritise upskilling and 63% regarded skills gaps as a major barrier to business transformation. More revealingly, when asked about their response to AI specifically, 77% planned to upskill or reskill existing staff, while 41% anticipated reducing staff where AI could replicate people's work. Those intentions can coexist: a firm may need fewer people for some tasks while investing in the people it retains. [2]

These are expectations reported in a survey conducted in 2024, not observed job losses or proven training outcomes. The survey principally reflects larger employers and does not establish what is happening to trainee recruitment in individual professions. A separate UK employer survey, published in January 2026 but based on fieldwork involving 801 employers in 2024, found that only 11% had provided AI-related training in the preceding year. The populations and dates differ, so the percentages cannot be read as a simple time series. [3]

There is a further distinction behind all these figures. **Training employees to use AI is not necessarily the same as using AI to develop competent professionals.** An accountant may learn to generate an analysis without improving their accounting judgement. A trainee lawyer may produce sophisticated research notes without being able to assess the authorities. Employers could invest heavily in AI literacy and still overlook the capabilities on which their future senior workforce depends.

Why might individual businesses underinvest in training?

The economic problem predates generative AI. Imagine a business paying to train a graduate for three years. Once the employee becomes more capable, a competitor may recruit them and benefit from an investment it did not make. Why should the original employer bear the cost?

There is no simple answer, and it would be wrong to suggest that firms inevitably refuse to finance transferable skills. Daron Acemoglu and Jörn-Steffen Pischke examined this question in their 1999 paper, *Beyond Becker: Training in Imperfect Labour Markets*. They explain why firms can rationally finance skills useful to other employers: recruitment frictions, imperfect information and wage structures can allow an employer to recover some of the value of higher productivity even if the skills are transferable. Their paper draws on economic models and existing evidence; it is not a forecast of AI-related recruitment. [4]

There can nevertheless be a difference between the training investment that suits an individual firm and that which benefits an industry. Historically, junior staff have supplied productive work while acquiring professional skills. If AI reduces the need for their immediate

contribution, the incentive to recruit them may weaken. A firm might decide to train fewer people and hire experienced professionals when required. That strategy could work for one business, but become much less reliable if its competitors make the same decision.

A shortage is only a possible outcome, not a prediction. AI may also reduce demand for some senior roles, shorten effective training pathways or create new occupations. The relevant decision is not whether to preserve historical headcount, but whether the organisation has a credible means of acquiring the capabilities it expects to need.

Should companies preserve traditional junior work?

One response would be to retain graduate recruitment and insist that trainees perform work which AI can already undertake more efficiently. I would resist that. Repetition is not automatically educational. A junior accountant transferring figures between spreadsheets may learn less than a colleague examining why the figures do not reconcile. A trainee lawyer need not manually read every document in a disclosure exercise to learn how unusual evidence is identified and evaluated.

Some effort develops capability; other effort merely consumes time. The employer's task is to establish which is which. A time-recording system can show how long a junior spends on research or drafting and what the client is charged, but it says very little about what the junior has learnt. Without that distinction, decisions about training risk being governed almost entirely by short-term productivity.

There are occasions when exposure to routine work does matter. Reviewing a representative selection of documents, for instance, may help a trainee appreciate how disputes arise and how apparently innocuous facts become significant. The right answer may therefore be partial automation accompanied by purposeful human review, rather than either complete manual processing or complete delegation. Employers need to find out what the task teaches before deciding how much of it to eliminate.

The alternative: redesign the work around learning

The studies discussed in Part 2 point in different directions. Brynjolfsson, Li and Raymond found productivity gains among customer-support staff and evidence consistent with employees learning from AI suggestions, including during periods when the assistant was unavailable. By contrast, an experiment involving patent lawyers found that AI improved drafting quality without a statistically significant average gain in junior lawyers' subsequent unaided judgement. Both studies have limits, but together they show why better output cannot be treated as proof of better learning. [5] [6]

A third experiment, involving school mathematics, found that unrestricted AI assistance improved practice performance but was associated with worse unaided test results. A specially designed tutor incorporating teacher-developed hints largely avoided that effect. This does not demonstrate what will happen in professional practice, but it reinforces the importance of how assistance is designed. [7]

The response should not be to preserve every traditional task. It is to **design AI-assisted work so that employees develop the capabilities they will eventually need**, using guided practice, feedback and increasing independence. Those are established educational principles. What AI changes is the risk that a workflow optimised for producing answers will quietly remove opportunities to develop the reasoning behind them.

What could an AI-supported professional apprenticeship look like?

Consider a trainee lawyer asked to analyse a contractual dispute. Under one approach, the trainee supplies the documents to AI, requests a complete opinion, checks the result and submits a polished memorandum. The partner sees the finished work but may have little idea how much understanding the trainee acquired.

A different process could preserve much of the productivity benefit while making learning more deliberate. The trainee first identifies the central issues and prepares an initial view, then uses AI to test arguments, locate potentially relevant authorities and expose weaknesses. The trainee must explain which suggestions were accepted or rejected and why. The supervisor reviews the reasoning as well as the final document. As competence improves, the amount of structured support can be reduced.

This is a proposed training design, not a validated method for trainee solicitors. It follows the general distinction between assistance that supports understanding and assistance that substitutes for it. The aim is not to manufacture difficulty. It is to ensure that the particular intellectual activity the trainee needs to practise is not routinely removed by automation.

It also requires a supervisor willing to explain why a proposed argument fails, rather than simply returning an improved draft. Trainees may learn as much from the reasons for rejecting a superficially convincing answer as from being shown the correct one. The employer must therefore decide whether the time saved in drafting should partly be redirected towards better feedback and exposure to more demanding problems. Otherwise the new system may deliver faster work without developing the judgement needed to supervise it.

Training should follow the capability, not the technology

Professional firms often describe AI training in terms of tools: prompting, document analysis, research platforms and automated workflows. Those skills are useful, but they are not the

same as professional competence. A lawyer must identify the real issue, assess evidence and consequences, communicate advice and recognise when further expertise is required. An engineer must understand the behaviour of the system being designed and spot implausible results. AI can assist with these tasks without assuming responsibility for them.

The better starting point is to identify the capabilities required at each stage of development. A trainee may need to identify a problem independently before seeking AI help; a more experienced professional may use AI to challenge a provisional analysis; a senior specialist may use it to test strategy. The appropriate assistance can therefore vary with the purpose of the task and the employee's competence.

The UK Skills for AI research programme, led by Dr Nisreen Ameen through a British Academy fellowship in partnership with Skills England, sets out the PRIMES principles for AI upskilling: **Practical, Reachable, Integrated, Modular, Expandable and Sustainable**. The programme drew on workshops, organisational case studies and a survey. Its recommendations concern training that employees can access, apply and maintain in ordinary work; they do not establish, experimentally, the best route to senior professional expertise. The survey sample was not nationally representative and included a high proportion of organisations already using AI. [8]

The programme's KPMG case study provides an illustration. It describes leadership support, an internal AI assistant called Spark, role-specific learning and apprenticeship arrangements as part of a wider approach to AI capability. The reported outcomes were supplied through the case-study process, not independently audited measures of improved judgement. The example shows how a large professional-services organisation is attempting to integrate AI training into work; it does not prove that its employees will become better senior advisers. [9]

How should employers measure success?

Productivity can be measured through completion time, output, error rates, revenue and profitability. None necessarily establishes that a trainee is becoming more capable. An employee may submit better work because AI performs more of the analysis, while their own judgement remains unchanged; another may initially become slower because the training deliberately requires independent reasoning.

One possibility is to assess employees periodically on unfamiliar problems both with and without AI. Assisted assessments examine how well they use the tools available in practice. Carefully chosen unaided tasks test the knowledge and judgement needed to supervise those tools, challenge errors and take responsibility. There is little point in requiring professionals to memorise information that can reliably be retrieved. More useful assessments might ask them

to detect a defective AI analysis, explain why a conclusion is unsafe or recognise when specialist advice is required.

Formal completion of a training programme is also different from mature workplace judgement. In the solicitors' profession, for example, the SRA distinguishes confirmation of qualifying work experience under the SQE route from the competence assessment made through the SQE examinations. Passing the qualifying threshold does not mean that a newly admitted solicitor possesses the judgement of an experienced practitioner. [10]

These are proposals for evaluation, not a universally validated test. They require suitable materials and supervision, and their value should be judged by whether they help identify learning and improve development decisions. **The objective is to measure capability, not merely participation in training or the appearance of competent output.**

The tests should reflect the occupation. In legal practice, an employee might be asked to identify a missing authority or explain why a seemingly persuasive analysis overlooks a factual distinction. In engineering, the task might involve detecting an implausible assumption in a simulation. Repeated assessments over time would offer more information than a single examination, although they would need to be proportionate to the cost and significance of the work. The intention is to detect gaps early enough for training to make a difference.

The role of the supervisor may become more important

AI may reduce the need for humans to check every intermediate step in routine work, but developing expertise is a different problem. A trainee can now submit an impressive document without necessarily understanding how its arguments were constructed. A supervisor may therefore spend less time correcting a draft and more time probing the employee's reasoning. That is a change in supervision, not necessarily its disappearance.

It has a cost. Senior professionals have demanding workloads, and time spent developing juniors may be less immediately profitable than chargeable work. Employers that want structured learning must recognise it in workloads and performance measures. Otherwise supervisors are being asked to prioritise long-term development while being rewarded for short-term output.

Practical arrangements might include explicitly allocating time for coaching, recognising successful supervision when evaluating senior employees, and involving suitably experienced staff rather than expecting every decision to reach a partner. The degree of review should also vary with the trainee's experience and the significance of the work. Requiring partner-level examination of every routine step would defeat many of the efficiency gains AI can offer; abandoning oversight of higher-risk judgements would be the opposite error.

In regulated legal practice, supervision is also an obligation. Paragraphs 3.5 and 3.6 of the SRA Code of Conduct for Solicitors require effective supervision of client work and competence among those managed. The SRA's guidance, updated in June 2026, discusses appropriate scrutiny and professional judgement when AI assists legal services. [10]

The legal distinction between delegated tasks and retained responsibility was clarified in *Chartered Institute of Legal Executives (CILEX) v Mazur* [2026] EWCA Civ 369. The Court of Appeal held that unauthorised individuals may perform tasks within the conduct of litigation for and on behalf of an authorised individual, where that individual retains responsibility and puts appropriate arrangements for direction, management, supervision and control in place. The intensity of supervision depends on the circumstances. **The judgment concerns delegation of litigation tasks, not the use of AI**, but it illustrates why the responsibility of the authorised professional cannot simply be transferred to a junior colleague. [11]

The SRA's warning notice *Misuse of AI*, published in August 2026, is more direct. It warns that those supervising junior or unauthorised colleagues may face regulatory consequences where fabricated authorities are submitted to court without adequate review or supervision. It also identifies risks of inaccurate work and breaches of confidentiality. A firm cannot discharge responsibility merely by relying on how convincing an AI-generated document looks. [12]

What about the commercial cost?

The proposal must survive contact with a profit-and-loss account. Suppose AI halves the time required for routine work. A business might reduce staffing costs, increase output, pass savings to clients or invest in training. The balance will depend on demand and pricing. Where work is billed by the hour, greater efficiency may reduce revenue from particular tasks; fixed fees may change the incentives but transfer some pricing and delivery risk to the firm.

A large partnership may spread the cost of supervision across many employees and clients. A smaller consultancy may train fewer juniors, but more intensively. An engineering business may need practical exposure to equipment and failures in a way a software business does not. There is no reason to prescribe identical arrangements for each. **Investment in development should reflect the capabilities the business expects to need, not simply the volume of junior work left after automation.**

The business model affects what is possible. A firm paid by the hour may find that reducing time spent on routine work changes the economics of employing and supervising juniors at the same moment. A firm selling a fixed-price service may have more flexibility to allocate some of the efficiency gain to training, but must manage the risk that the work proves more difficult than expected. Neither pricing model guarantees that investment in development will be maintained.

Who should bear the cost of training?

If businesses collectively reduce investment in junior training, should universities, professional bodies or governments intervene? There is a plausible economic concern: the benefits of professional training can extend beyond the employer paying for it. The OECD's *Employment Outlook 2023* recognises the risk of underinvestment in AI-related training where benefits spill over to others. That discussion concerns AI skills rather than directly proving underinvestment in professional apprenticeships. [13]

The remedy is not obvious. Acemoglu and Pischke explain that workplace training often includes mentoring, advice and practice whose quantity and quality are difficult to specify or verify. They also caution that subsidies may simply reward training an employer would have provided anyway, or have little effect where genuine training cannot be observed. Regulation and certification may help, but can also change employers' incentives in counterproductive ways. **The test of any intervention should be whether it increases real competence, not merely recorded training expenditure or course attendance.** [4]

Professional bodies might develop competency standards and encourage employers to demonstrate how relevant capabilities are acquired. That is more directly related to the problem than requiring a predetermined number of graduates or years of service. It would also make it easier to compare new pathways with traditional ones without assuming that shorter training is necessarily inferior.

I would want evidence of a developing problem before advocating major regulatory intervention. But waiting for an actual shortage of experienced professionals may be too late. Training decisions can take years to affect the supply of senior staff, and the consequences may take further years to correct. Earlier indicators could include trainee intake, progression into senior roles, assessment of professional competence and the time needed to recruit experienced employees. All must be considered against changing demand for expertise, including any reduction in senior work attributable to AI. If those indicators showed an emerging shortage of capability, the case for coordinated action would strengthen without waiting for a crisis.

Could AI make the traditional apprenticeship better?

The traditional apprenticeship can be slow and inconsistent. Two trainees working in the same organisation may receive very different opportunities: one works with an experienced professional who explains decisions and gives feedback; the other spends months completing routine assignments with little understanding of their significance.

AI could help reduce that inconsistency by providing explanations, generating practice problems or challenging a trainee's initial analysis before the work reaches a supervisor. It

might expose beginners to situations that would otherwise arise only after years of practice. Yet AI feedback can itself be wrong, and none of these possibilities removes the need for human oversight.

The opportunity is worth testing, rather than assuming. There is a difference between deliberately designing an AI-supported apprenticeship and simply allowing everyone to use the available chatbot. The first treats learning as an outcome to be assessed; the second risks assuming that it will follow automatically from greater productivity.

There may even be an opportunity to improve on traditional training. An AI tutor can be available when a human supervisor is occupied, explain material in different ways and offer low-stakes opportunities to practise. But a plausible explanation is not necessarily a correct one. Employers would need to select tools carefully, evaluate their educational effects and ensure that difficult questions ultimately reach an appropriately qualified person.

A practical approach for employers

I would start with the capabilities the business expects to need over the next five to ten years, not with a complicated new training system. It should then identify which existing tasks help develop those capabilities, which can be safely automated and which might be replaced by more effective practice. The objective is to retain useful learning, not historical workflows.

A modest pilot could compare ordinary AI-assisted work with a structured approach involving initial independent analysis, AI-generated challenges and supervised feedback. Where practicable, the firm should assess starting competence, allocate comparable tasks and judge progress using new problems after an interval, rather than only testing performance immediately after training. It may already have a natural comparison group in employees using AI in the ordinary way. The exercise still consumes supervisory and assessment time, but need not begin as a large research programme.

Confidentiality, fairness and employee consent where relevant must be considered, particularly if assessment results may affect employment decisions. The point is to improve training, not to punish employees for using technology their employer encourages. If the results show better capability development, the method can be refined and expanded; if not, the business can change course without having committed to an elaborate scheme.

What should employers avoid?

There are two unhelpful extremes. One is to preserve junior work largely unchanged because it has always been part of professional training, regardless of its educational value. The other is to automate as much junior work as possible, reduce recruitment and assume that expertise

will be available when needed. The first may waste resources; the second may create risks that become apparent only after several years.

Neither approach is warranted by the present evidence. **Automation and professional development should be treated as related but distinct decisions.** Employers should decide what capabilities are needed, determine how they can be developed efficiently and test whether that development is taking place. That is established management discipline applied to a different relationship between work and learning.

Who will train tomorrow's experts?

Part 1 established that AI is changing the economics of junior employment, although the evidence does not support attributing every reduction in recruitment to the technology. Part 2 demonstrated that improved AI-assisted performance does not necessarily translate into improved independent competence. Sometimes it does; sometimes it does not.

The employer's dilemma follows directly. Businesses can no longer safely assume that the work required to deliver today's services will automatically provide the experience necessary to develop tomorrow's professionals. That assumption may become increasingly unreliable as AI undertakes more of the routine intellectual work traditionally allocated to junior employees.

But this is not necessarily a reason for pessimism. **AI could help organisations develop professional expertise more efficiently than the traditional apprenticeship, provided they design and evaluate the process rather than simply automate the work.**

The evidence does not yet establish which arrangements will work best across different occupations. It does, however, justify treating professional development as a deliberate organisational objective rather than an incidental benefit of employing junior staff.

Some companies may need to preserve aspects of traditional training. Others may develop more effective approaches. Many will probably combine both.

There is no obvious reason why the same solution should apply to every profession or business. The critical question for employers is therefore not simply how many junior employees AI allows them to dispense with. It is whether they are developing the people who will eventually be capable of supervising, challenging and taking responsibility for the work AI helps them produce.

A business that overlooks that distinction may enjoy several years of improved productivity before discovering that its supply of experienced professionals is weaker than expected. A business that gets it right may gain something more valuable than immediate cost savings: a faster and more effective way of developing expertise.

In Part 4, we will examine the other side of the relationship. What should trainees, students and educational institutions do to ensure that AI becomes a means of developing genuine professional capability rather than merely producing increasingly sophisticated work?

Alex Asks

If I can do much of the work your trainees once performed, why should you continue employing them?

Perhaps because you will still need experts. But shouldn't you first establish whether there is a better way for me to help them become experts, rather than simply making them repeat everything their predecessors did?

References

[1] Office for National Statistics (2026). *Artificial intelligence in UK businesses: 2023 to 2026*. Published July 2026. ONS original source.

[2] World Economic Forum (2025). *The Future of Jobs Report 2025*. January 2025, particularly pp. 5-6, 48, 62-63 and Appendix: Report Methodology. WEF original source.

[3] Department for Science, Innovation and Technology (2026). *AI Skills for Life and Work: Employer Survey Findings*. Published 28 January 2026; Ipsos fieldwork March-June 2024. GOV.UK source.

[4] Acemoglu, D. and Pischke, J.-S. (1999). *Beyond Becker: Training in Imperfect Labour Markets*. *The Economic Journal*, 109(453), F112-F142, particularly pp. F119 and F126-F128. DOI: 10.1111/1468-0297.00405. Original paper.

[5] Brynjolfsson, E., Li, D. and Raymond, L. (2025). *Generative AI at Work*. *The Quarterly Journal of Economics*, 140(2), 889-942. DOI: 10.1093/qje/qjae044. Journal record.

[6] Autor, D., Rodchenko, T., Martin, J., Iscenko, Z., Strand, S., Pearl, D. and Ferere, M. (2026). *Does AI Assistance Enhance or Erode Expertise? Evidence from a Three-Month Field Experiment in Patent Drafting*. NBER Working Paper 35720, revised October 2026. DOI: 10.3386/w35720. NBER paper.

[7] Bastani, H. and colleagues (2025). *Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics*. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. DOI: 10.1073/pnas.2422633122. Full article.

[8] Ameen, N. and research team (2026). *What Works for AI Upskilling in the UK: Employer Guide* and supporting *Research Evidence, Analysis and Methodology*. British Academy Policy-Led Innovation Fellowship in partnership with Skills England; employer guide published via DWP / Skills England. Official employer guide.

[9] Ameen, N. and research team (2026). *What Works for AI Upskilling in the UK: Supporting Case Studies*, section 2.1.2 (KPMG) and its methodological limitations. British Academy / Skills England research programme. Publication landing page.

[10] Solicitors Regulation Authority (2026). *Code of Conduct for Solicitors, RELs, RFLs and RSLs*, paragraphs 3.5-3.6; and *Effective supervision*, guidance updated 12 June 2026. Code and guidance.

[11] *Chartered Institute of Legal Executives (CILEX) and others v Mazur and others* [2026] EWCA Civ 369, 31 March 2026, particularly [187(v)-(vii)] and [198]. Approved judgment.

[12] Solicitors Regulation Authority (2026). *Misuse of AI*, warning notice published 17 August 2026, particularly the discussion of supervisory responsibility and fabricated citations. SRA warning notice.

[13] OECD (2023). *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*, Chapter 5, *Skill needs and policies in the age of artificial intelligence*. DOI: 10.1787/08785bba-en. Chapter 5.