How Fast Are We Going?

Technological Acceleration, Institutional Adaptation and Early-Career Training in the United Kingdom

Chas Burns

Copy Real

Independent Research White Paper | Illustrated Research Edition v0.8 | 7 November 2026

Abstract

How fast are we going, and can society keep pace? This paper asks whether artificial intelligence is entering economic life faster than the institutions that employ, train and regulate people can adapt; whether any such mismatch can be measured; and what would show that institutions are keeping pace. It separates three processes that public argument tends to merge: the development of technical capability, the diffusion and integration of AI in organisations, and the adaptation of institutions. Their speeds cannot be subtracted to give a single gap. The paper proposes instead to compare each institutional response with the length of the training cycle it is meant to protect.

The question cannot be tested in general. It can be tested in particular settings. The principal case is early-career training in the United Kingdom, because entry routes depend on tasks that AI can perform, and because employers' training decisions respond slowly and leave measurable traces. The proposed mechanism—that AI changes the net private cost of training a beginner—is not new; it rests on existing theory and on apprenticeship cost accounting. The contribution is a staged design. It establishes the wider labour-market baseline first, maps occupational signals, and only where a credible AI-related difference remains asks whether training economics, recruitment and progression point to temporary adaptation, persistent underinvestment or complementary expansion.

Current UK evidence is mixed, mostly observational and largely based on exposure rather than documented adoption. Nothing in it shows that institutions are failing to keep pace, and nothing shows that they are succeeding. The paper states what each stage of the design could and could not establish, including the conditions under which no causal conclusion would be available. The 2022–2032 period is a programme of observation, not a forecast.

Keywords: artificial intelligence; technological diffusion; institutional adaptation; United Kingdom; early-career hiring; net training cost; apprenticeship; wage floors; research design

1 Introduction

There is a peculiar difficulty in measuring how fast society is changing when the instruments of change fit inside an ordinary laptop. A factory requires a site, equipment, finance, employees and time before a new method of production affects the world outside its gates. An AI service can make a new capability available to millions through an update. Availability, however, is not the same as successful use. Nor does using a tool amount to reorganising an economy around it.

Public discussion regularly moves between three claims: that AI systems are becoming more capable; that businesses are adopting them; and that people must therefore be losing their jobs. Each proposition requires different evidence. None follows mechanically from the preceding one. A system may write competent code but fail an ambiguous production task. An employee may use AI daily

without their employer changing a single job description. A firm may redesign a workflow without dismissing anyone, yet recruit fewer beginners the following year.

The proposition investigated here is that artificial intelligence may be advancing, and entering working life, faster than the institutions responsible for accommodating it can adapt. The paper asks whether that mismatch exists, whether it can be measured, what consequences it might have, and under what conditions institutions adapt successfully. It does not assume that faster technology causes harm; adaptation may succeed. It does not assume that every technically exposed task will be automated, that every automated task eliminates paid work, or that the next decade has one inevitable outcome. Those assumptions would make the paper easier to write and less useful.

In my earlier paper *What Remains When the Tools Disappear*, I examined the implications of automated systems becoming part of ordinary human judgement and agency [26]. The present paper asks a different question: what happens when technical capability, organisational adoption and the institutions through which people develop expertise change on different timescales? The general question is too broad to test directly, so the paper tests it where the institutional response can be observed: the routes by which inexperienced people become experienced. Its most specific hypothesis is not mass unemployment. It is that some employers may make fewer places for beginners even while their output, and their demand for experienced judgement, increase.

The question has a practical urgency. An annual employment total tells us how many people are in work. It does not tell a fifteen-year-old which occupation will still have an accessible first step when they leave education. That does not invalidate the employment statistic. It suggests we have been looking at a more complicated landscape through one window.

Four research questions follow.

- Do technical capability, organisational integration and institutional adaptation proceed at distinguishable speeds, and can they be compared without constructing a misleading composite metric?
- What observations would show that institutions are keeping pace, and what would show a persistent mismatch?
- In UK early-career training, is there a credible AI-related change in entry routes once wider labour-market conditions are taken into account?
- If there is, does it arise through the economics of training, and does it persist?

The paper is a conceptual and evidence-synthesis paper with a proposed research design. It is not a completed causal study. Sections 2 to 5 set out the thesis, the evidence on capability and adoption, and what can be observed about institutional adaptation. Sections 6 to 8 introduce early-career training as the principal case and review the evidence and economics. Sections 9 and 10 set out the research design and how its results would bear on the wider thesis. Sections 11 to 13 cover policy, limits and conclusions.

2 The thesis and its testable form

2.1 Three processes

Technical capability refers to what a system can complete reliably under stated conditions: task complexity, required supervision, monetary cost, error rate and time to verified output. A benchmark success is not necessarily an economically useful result. The same model can be remarkably capable at one task and unreliable at another.

Diffusion and integration is the extent to which businesses, public services and professions alter actual working practices. This includes systems procurement, data access, accountability, employee

training, workflow redesign and hiring decisions. An office using AI occasionally to tidy correspondence is not equivalent to an organisation rebuilding customer service or software production around connected systems.

Institutional adaptation is the change in the rules, routines and resources through which people are employed, trained, certified and protected: employer training practices, apprenticeship and professional qualification routes, wage and tax rules, regulation and public services. It includes the opportunities through which people learn and earn. Human learning is not the slow element by nature; an apprentice may learn rapidly with an AI tutor. But the apprentice still needs training time, adequate feedback, a paying employer and access to meaningful responsibility. A retraining course that cannot pay the rent while it runs has a very real design flaw.

A mismatch between these processes is possible, but it is not guaranteed and need not run in one direction. AI reliability, infrastructure, regulation, weak demand or poor management may slow effective adoption. In other occupations, workers may develop new competencies more rapidly than employers offer positions using them. The object of study is which constraint binds, when, and for whom.

2.2 Why the speeds are not subtracted

Two separations are essential. First, the speed of a benchmark and the diffusion of a workplace practice are measured differently. Subtracting an exponential capability curve from an S-shaped adoption curve, even after standardisation, does not produce a meaningful ‘gap’ in years; any such number would reflect the choice of series, not a property of society. Second, the incentive to train junior staff may weaken even if technology and institutions adjust at identical speeds. The first-rung mechanism examined in section 8 does not logically depend on any gap between rates. This is a causal distinction, not a minor editorial one.

What can be done instead is to place observable events for each process on a common calendar and to ask a different question of each institutional response: did it arrive within the length of the pipeline it protects? An apprenticeship, a graduate training scheme and a professional qualification each renew expertise over a period of years. A response that arrives within one training cycle may be timely. A response that arrives after several cohorts have passed through a changed entry route is late, however fast or slow the underlying model was improving. This comparison needs no composite index, and it can be applied one occupation at a time. Figure 1 summarises the idea.

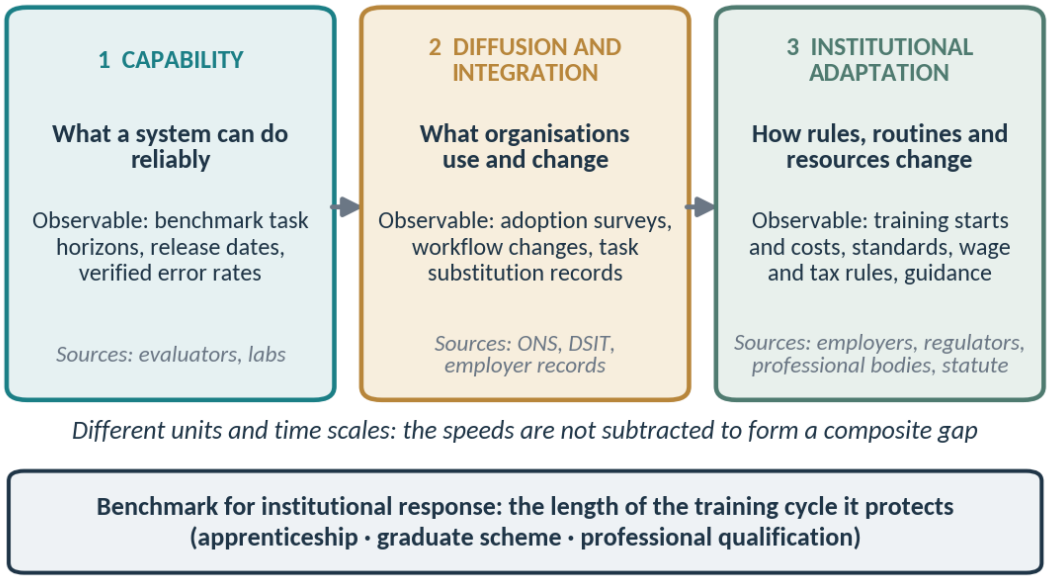

Figure 1. Three processes, their observables and the benchmark proposed for institutional response. Conceptual map, not measured data. No lag between the processes is estimated in this paper.

I have used *AI Velocity Gap* in earlier drafts as shorthand for the distance in time between technical availability, workplace integration and institutional adjustment. It is not a measured gap, a necessary condition for first-rung harm, or a claimed theoretical discovery. In this edition it is a subsidiary label for the timing comparison above and appears nowhere in the research design. Unless a timing measure predicts which occupations show persistent entry-route weakness better than diffusion and training-incentive theory already does, the label can be discarded without altering the research question.

2.3 What keeping pace, and falling behind, would look like

A thesis that no observation could contradict would not be worth testing. Table 1 states, for the early-career case, what each outcome would predict, what would count against it and which non-AI explanation must be excluded. 'Keeping pace' here means that entry routes, training and progression in occupations with documented AI integration hold up or recover within one training cycle, and that supporting institutions respond within the same period. A persistent mismatch means that they do not, despite adequate demand.

Table 1. What keeping pace, and falling behind, would look like in UK early-career training. Conceptual, not measured. Pathway labels follow section 9.4.

Outcome	What would be observed	What would count against it	Non-AI explanation to exclude
Institutions keep pace	Entry starts, training and progression in occupations with documented AI integration hold or recover within one training cycle; supporting standards and routes are revised in the same period.	A sustained shortfall in entry routes after integration, with no institutional response.	A demand recovery that would have lifted entry routes anyway.
A. Temporary adaptation	Entry weakens after documented workflow change, then recovers as supervised work is redesigned.	No recovery within the pre-specified window.	The general hiring cycle; reversal of labour-cost rises.
B. Persistent underinvestment	Paid starts stay weak relative to demand after integration; net training cost rises; more recruitment of trained staff from other firms; no institutional response within the cycle.	Net cost stable or falling; starts recover.	Weak demand, financing or policy costs that affect low-adoption firms equally.
C. Complementary expansion	Beginners' quality-adjusted value rises or supervision costs fall; entry routes hold or grow.	Entry routes contract where integration is deepest.	Service expansion unrelated to AI.
Baseline: wider weakness	Entry hiring falls comparably whatever the degree of AI integration.	A differential that survives controls for demand, cost and policy exposure.	Not applicable: this is the comparator.

The broad proposition would be weakened by any of the following: no credible AI-specific differential once wider demand and policy effects are accounted for; recovery or expansion of entry routes within the pre-specified window; or institutional responses that arrive within a training cycle. It would be supported, for this case, only by a sustained shortfall after integration, linked by independent evidence to the cost of training or to the absence of any institutional response. Even then it would be supported for one domain, not for society as a whole.

2.4 Earlier theory

The problem of technology outrunning social arrangements is not new. William F. Ogburn's account of *cultural lag* examined the tendency of aspects of social organisation to adjust at different rates [32]. Everett Rogers studied how innovations spread through social systems and why adoption is neither simultaneous nor inevitable [34]. Claudia Goldin and Lawrence Katz described a long historical race between educational attainment and technological demand for skills [31]. All three traditions warn against treating invention and adaptation as the same event.

Economics supplies more specific mechanisms. Brynjolfsson, Rock and Syverson's *productivity J-curve* explains how investment in new technology and the organisational assets needed to exploit it can delay measured returns [22]. Acemoglu and Restrepo distinguish the displacement of labour from tasks that automation takes over and the reinstatement of labour through new tasks that technology creates [33]. James Bessen's empirical work on learning-by-doing demonstrates that practical experience can constitute significant investment in human capital, even when it does not appear as a formal training course [35]. These accounts make an important correction to any argument that infers employment directly from a coding benchmark.

The training incentive problem has its own literature. Acemoglu and Pischke examine how labour-market imperfections affect employers' incentives to finance skills that workers may carry to other firms [23]. This mechanism can exist without AI. Garicano and Rayo start from the observation that juniors pay for training by doing menial tasks, and model what happens when AI performs a growing share of them [49]. Akin proposes a formal 'apprenticeship externality' in which firms that automate

junior tasks do not bear the social cost of closing the pathway to future experts [46], and Lovett develops a related argument about professional expertise as a shared pool whose regeneration can fail [50]. These are theoretical predecessors, not evidence that the predicted harm has occurred. What remains open is empirical: whether AI changes the net private cost of developing beginners in British workplaces, whether any change survives comparison with non-AI causes, and whether a profession as a whole is affected.

3 How rapidly is technical capability changing?

There is convincing evidence of rapid improvement in several defined technical tasks. The UK government's January 2026 assessment referred to a historical trend in which the length of certain tasks completed by frontier agents approximately doubled every seven months [5]. This remains useful historical context, but it is not the latest short-period estimate.

METR's *Time Horizon 1.1* update of January 2026 found a long-run median 50-per-cent-success doubling time of approximately **196.5 days**. On a fit restricted to the post-2023 period, its estimate was **130.8 days** with a reported interval of **107–161 days**; on a still shorter post-2024 fit it was **88.6 days**, without a comparable interval in the summary table [4]. Short-window fits are sensitive to the models and tasks included. The updated task collection and assessment methods also differ from earlier versions. These rates are estimates over studied periods, not physical constants or forecasts that all forms of intelligence will keep doubling on schedule.

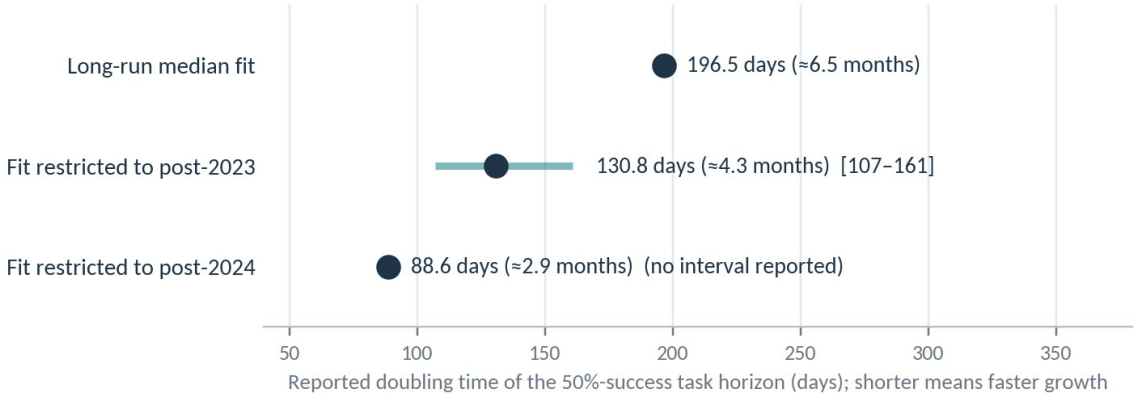

Figure 2. Reported doubling-time estimates for the length of tasks that frontier agents complete with 50 per cent success, by fit window. Shorter is faster. The estimates are reported values, not forecasts. Source: METR Time Horizon 1.1 [4].

The meaning of the measurement matters. METR estimates the human time normally required for a task at which an agent has a given probability of succeeding. A task labelled as taking a human several hours is not a claim that the system performs reliably for several unsupervised hours. A 50 per cent success rate may be impressive in research and unacceptable in a payroll department. The economically relevant measure is verified work after the costs of failed attempts, human checking, data preparation and liability have been included.

The possibility of AI accelerating AI research deserves separate treatment. Models increasingly assist coding, experiment design and parts of research engineering. If improvements in these systems help produce their successors, research cycles may shorten. Anthropic's account of its own development work describes AI assistance but does not establish that an independent system can design and deploy its successor without human direction [6]. RAND's October 2026 review similarly distinguishes assistance within research programmes from autonomous agenda-setting and open-ended self-improvement [7]. These claims also come from organisations with institutional interests and should be assessed accordingly.

Even if recursive feedback becomes important, it would not follow that an entire economy changes at the same rate. Hospitals, courts, councils and businesses must cope with reliability, privacy, procurement, cost and accountability. A software update can be distributed overnight. A trustworthy professional practice cannot be installed in the same fashion.

Capability is therefore the process with the best time series and the weakest direct bearing on institutions. A rising task horizon says what systems might do; it says nothing about what employers, regulators or training providers have decided to do about it.

4 The United Kingdom: rapid adoption, shallow integration

The Office for National Statistics (ONS) reported that **30 per cent of businesses** used at least one AI technology in **September 2026**, up from **9 per cent** when the question was introduced in September 2023. For businesses with **ten or more employees**, September adoption was **38 per cent**; for those with **250 or more**, **51 per cent**; and for those with **fewer than ten**, **29 per cent** [1]. This is rapid diffusion, but not evidence that a majority of British small businesses already uses AI. The ONS identifies the relevant Business Insights and Conditions Survey statistics as *official statistics in development*.

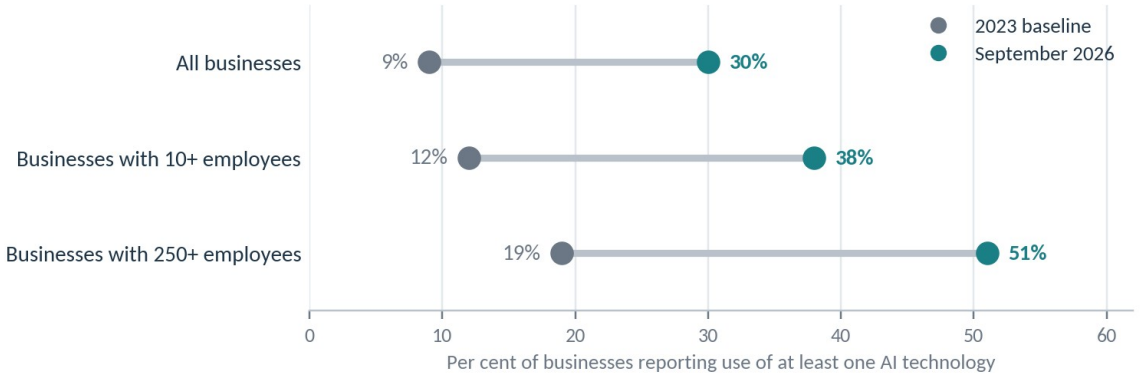

Figure 3. Percentage of UK businesses reporting use of at least one AI technology, 2023 baseline and September 2026. The groups overlap and must not be added together. The 2023 baseline for firms with ten or more employees is the late-2023 figure reported in [2]. Source: ONS Business Insights and Conditions Survey [1, 2].

The **June 2026** figure for businesses with ten or more employees was **35 per cent**, reported in an ONS analysis published on **20 July** [2]. It is therefore incorrect to combine the September headline with the June subgroup number as if they came from the same survey wave. The distinction is modest numerically and fundamental methodologically: a study of speed cannot afford to confuse the dates on its speedometer.

The Department for Science, Innovation and Technology's UK Business Data Survey reported **41 per cent AI use among businesses handling digitised data**, a different sample universe. Within its AI-using subgroup, **21 per cent** said AI was integrated into existing systems [3]. Such integration may consist of an AI function inside office, finance or customer-management software; it is not evidence that all those firms have redesigned their workforce. The ONS, meanwhile, found that the average number of AI technology types used per adopting business with ten or more employees had risen only from about **1.4 to 1.6** since late 2023 [2]. In the June sample, a relatively small minority described use as extensive.

Surveying individuals produces yet another population. The DSIT Public Engagement Survey, with fieldwork between November 2025 and March 2026, found that 59 per cent of adults had used generative AI in the previous three months, and that 54 per cent of online generative-AI users had used it for work [27]. These are not comparable denominators or time windows. We cannot subtract

the individual-use percentage from the business-use percentage and announce a measured institutional lag.

The ONS also asked businesses with ten or more employees using AI which roles were most affected. **Administrative and clerical work was named most often, by 38 per cent of these businesses** in September 2026, followed by creative and design roles at 24 per cent [1]. This is the proportion of surveyed AI-using businesses selecting a role category, not the proportion of administrative workers displaced. It identifies where workflow changes may be concentrated, not how many jobs have disappeared.

The credible conclusion is narrower and more interesting: AI access has diffused considerably, while deep organisational change is difficult to measure and evidently uneven. This leaves room for at least three outcomes. Firms may use AI to support existing workers; they may replace some labour through attrition and reduced recruitment; or lower costs may expand demand and create new positions. The existence of a chatbot on a desk does not select between those possibilities.

5 Institutional adaptation: what can be observed

Capability and adoption are measured, imperfectly, by evaluation groups and statistical offices. Institutional adaptation is not measured systematically at all, and this is the central empirical weakness of the broad thesis. The term needs disaggregating, because the layers adapt at different speeds and leave different traces.

- **Firm routines**: workflow redesign, supervision, hiring and training practice. Traces appear in employer records and, at aggregate level, in business surveys.
- **Training and qualification systems**: apprenticeship standards, professional qualification routes, university curricula and accreditation requirements. Revisions are dated and in principle public, held by regulators, awarding bodies and professional bodies. I have not checked how complete or consistent such records are.
- **Labour-market rules and public policy**: wage floors, employer contributions, levy arrangements and employment support. Statutory changes are dated and public.
- **Public services and regulators**: procurement, accountability and professional standards, traced through guidance and inquiries.

This review did not locate a systematic UK measurement of how quickly any of these layers has responded to generative AI. That is an open gap, not evidence in either direction. Dated institutional outputs can be listed—the government's assessment of AI capabilities and the labour market in January 2026 [5], or its entry-level hiring brief with LinkedIn in June 2026 [41]—but they measure attention, not adaptation of the routines that renew expertise.

The consequence for the thesis is plain. The proposition that society cannot keep pace is not, at present, testable as a whole. What is testable is whether the particular institutions that renew expertise adapt within the length of their own cycles. The next section chooses a case in which that question has observable answers.

6 Early-career training as the principal case

6.1 Why this case, and its limits

Early-career training suits the question for three reasons. Entry-level tasks such as summarising, drafting, checking and routine coding overlap with what generative systems do, so exposure is plausible. Entry routes depend on institutions with identifiable decision-makers: employers who finance training partly through trainees' productive work, qualification bodies, and public wage and

tax rules. And the outcomes are observable in principle: hiring flows, paid starts, progression and experience premia, with the cohort structure of the pipeline supplying the benchmark for response time described in section 2.2.

The case has equally clear limits. It is one domain, in one country with particular labour-market institutions. Technical exposure is not documented adoption. It says little about hospitals, courts or councils, whose adaptation problems are of a different kind. Evidence that entry routes hold up would not show that institutions elsewhere adapt; evidence that they weaken would not show that they fail. The case is a test of the thesis in one setting, not a proxy for the whole.

6.2 What precisely counts as job dilution?

Job dilution requires a definition precise enough to exclude changes that have little to do with AI. Treating every reduction in satisfaction, security or developmental value as dilution would encompass nearly any disappointing event in a labour market. The concept is useful only if its boundaries can be observed and tested.

For this paper, **job dilution** means a persistent decline in paid human labour input or paid entry opportunities **relative to a credible comparison path for similar output**, associated with independently observed AI-enabled workflow substitution. It may occur through fewer new hires, unfilled replacement vacancies, reduced contracted hours or a shrinking share of junior posts. A fall in job satisfaction, an increase in productivity or a change in task content alone does **not** establish dilution. Nor is every decline in a junior vacancy an AI effect.

Entry-route dilution is a narrower subtype: a reduction in the availability of paid, supervised opportunities through which beginners normally gain occupation-specific experience. Its primary measurements are the rate and level of junior hiring, paid apprenticeship or traineeship starts, access to supervised work and subsequent progression. Training quality must be assessed separately from headcount. A new route may be better than the old one; it should not be scored as a loss simply because it looks different.

Four outcomes should therefore be reported separately: **employment stocks**, **hiring flows**, **paid hours and earnings**, and **development opportunities**. The analysis must also record output, prices and demand. An employer may double its service volume with unchanged staffing and still create an economic benefit. Whether that change represents harmful entry-route dilution depends on the credible path for hiring, training and opportunity—not on a wish that office work remain inefficient.

These definitions impose an important exclusion. A firm reducing staffing during a broad downturn, without evidence of AI integration, is a case of falling employment. It is not a demonstrated case of AI-driven job dilution. Similarly, a senior workforce share can rise because seniors are growing faster, without the number of juniors falling. The difference matters.

7 The evidence on entry-level hiring

7.1 UK signals, contradictions and limits

British evidence no longer warrants a simple statement that AI has had little effect on hiring. It also cannot support a simple statement that AI is the principal reason for weak hiring. The most responsible reading places several kinds of evidence alongside one another.

Broad business surveys. In September 2026, only **6 per cent** of AI-using firms with at least ten employees told the ONS that AI had reduced their headcount, while fewer than **1 per cent** reported an increase [1]. A Bank of England survey found **nearly 90 per cent** of responding firms perceived no noticeable change in employee numbers associated with AI [10]. Such results are important checks against claims of economy-wide mass displacement. But they are self-reports about the current

workforce. They can miss a hiring decision never taken, a job not backfilled, or an employer unable to distinguish AI savings from general cost control.

Firm-level exposure. Bouke Klein Teeselink's UK study of job adverts and LinkedIn employment records spanning 2021–2025 found that firms with workforces more exposed to AI experienced an estimated **4.5 per cent reduction in total employment**, concentrated in junior positions, down **5.8 per cent**. Highly exposed firms were also **16.3 percentage points less likely to post a new vacancy** under the study's comparison [11]. These results are more directly relevant to the British first-rung hypothesis than a US graduate survey. They are nevertheless estimates based on differences in AI *exposure*, not an experiment randomly assigning AI adoption. Exposure may be correlated with business conditions or organisational changes that also affect hiring.

Occupational adverts. Haley Schlicht's August 2026 Bank of England staff analysis used British occupational composition and several exposure measures. Across **26 occupational groups**, the correlation between exposure and vacancy growth was **Spearman −0.70** in the reported comparison. Grouped into exposure thirds, online adverts fell **15 per cent** in the highest-exposure group, **10 per cent** in the middle and **6 per cent** in the lowest [12]. That graded pattern strengthens the case for investigation; a correlation across 26 occupations is not a causal estimate. Remote working, firm demand, sector composition and other changes may contribute. The article is staff analysis published on Bank Underground, not a formal position of the Bank's policy committees.

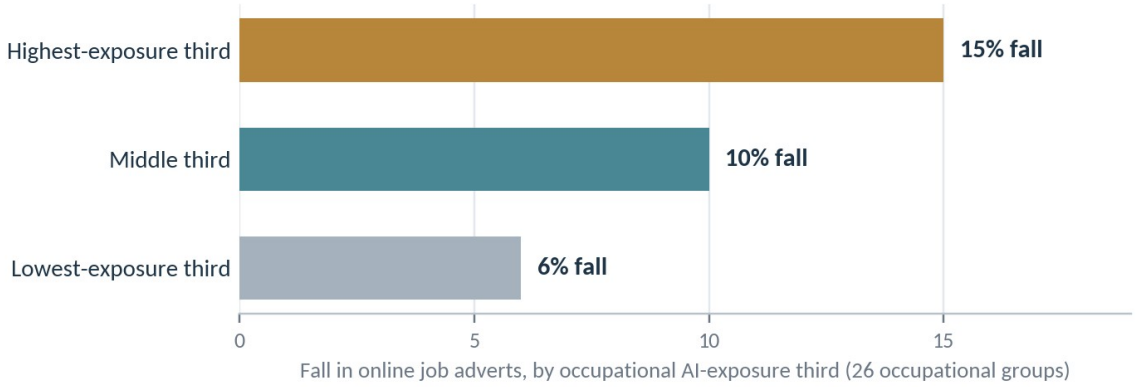

Figure 4. Declines in online advertisements by occupational AI-exposure third. Schlicht's exposure–vacancy correlation uses 26 occupational groups (Spearman −0.70). Association is not a causal estimate of jobs lost to AI. Source: Bank of England staff analysis [12].

Reported business restructuring. Morgan Stanley's first AlphaWise survey wave, released in early 2026, reported a net **8 per cent** reduction in positions attributed to AI among UK respondents [13]. Its sample covered selected sectors expected to experience a substantial AI effect and required firms to have used AI for at least a year. A subsequent **Wave 2** was reported as finding a **6 per cent** UK net loss, but it covered *different sectors* and the underlying primary report was not available for independent inspection in this revision [42]. The two percentages are not a trend measured in the same population and should not be pooled. They are reported changes in selected AI-using business samples, not independently verified UK-wide job-loss rates.

Bank intelligence. The Bank of England's July 2026 Monetary Policy Report states that contacts were gradually reducing demand for some highly automatable positions by slowing recruitment or not filling vacancies while increasing output [25]. This describes the mechanism at the heart of job dilution. It remains qualitative business intelligence and sits alongside the same report's finding that most surveyed firms saw no material employment impact.

The resulting picture is not uniform. Some datasets detect early selective contraction in precisely the occupations the hypothesis singles out. Broad employer surveys find modest change. Both could be true if early effects are concentrated, appear in recruitment before employment stocks, or differ

substantially across companies. They could also reflect different samples and measurement error. A proper study must actively try to distinguish these explanations.

7.2 The strongest competing explanation: wider UK hiring problems

The United Kingdom entered this transition with a weak youth labour market and rising employment costs. If these conditions are treated as a brief footnote, the AI argument becomes too convenient.

Indeed Hiring Lab reported in August 2026 that UK online job postings stood approximately **32 per cent below** their pre-pandemic baseline, unlike the United States and euro area, whose postings were at or above their respective baselines [15]. That comparison signals a UK-specific demand problem across many occupations. It cannot, by itself, explain which portion of the British decline—if any—AI caused. It does mean that a fall in graduate posts is not a clean natural experiment in technological substitution.

The Institute for Fiscal Studies (IFS) reported that the share of 16–24-year-olds in payrolled employment fell **4.3 percentage points**, from **54.9 to 50.6 per cent**, between December 2022 and December 2025 [16]. The decline was uneven: employment among **16–17-year-olds fell 7.3 percentage points**, compared with **4.8 points among 22–24-year-olds** and **1.1 points among 18–21-year-olds**. Many of the youngest workers were combining education with weekend or casual jobs. This broad youth-employment statistic is not a direct measure of graduate-entry hiring or paid professional training and must not be used as if the populations were interchangeable.

The IFS found that the recent deterioration in rates of young people outside education, employment or training was not disproportionately concentrated among recent graduates relative to recent non-graduate entrants [16]. That undercuts a simple national graduate-collapse narrative, while leaving open occupationally concentrated changes in legal, creative or technical entry routes. The IFS also found reasons to suspect a structural component rather than a purely cyclical fall. It did not establish AI as the cause, although it acknowledged that AI could be weighing on employment more broadly rather than only on young graduates. Neither possibility should be dropped from the competing explanations.

The United Kingdom also changed the price of employing people. Statutory minimum wages and employer National Insurance contributions can affect recruitment directly, especially in lower-paid occupations [17]. The IFS found no clear evidence that minimum-wage changes were a major driver of the youth employment deterioration, and judged the 2025 National Insurance reform unlikely to be its primary cause [16]. That does not settle a different question, whether a binding wage floor changes how employers respond when AI alters the productive contribution of beginners. Section 8.4 treats that as a separate hypothesis. National Insurance age and earnings thresholds must be modelled separately, not folded into a minimum-wage variable.

The government's June 2026 research brief using LinkedIn data found that **entry-level UK hiring was falling broadly in step with hiring at other seniority levels**, with no aggregate pattern of a distinctive entry-level collapse [41]. Beneath that headline, some information processing and professional occupations had markedly weaker entry hiring, but the same report explicitly declined to attribute the pattern causally to AI. LinkedIn members and postings are not a representative employment census. This is nevertheless genuine counter-evidence to a universal first-rung story and an argument for occupation-level rather than sweeping national estimates.

Other plausible explanations include high interest rates and financing costs, weaker consumer demand, changes in remote working and graduate supervision, post-pandemic hiring reversals, outsourcing, changes to immigration and labour supply, and the increasing demand for experienced staff. Some of these trends began before widespread generative AI use. A high exposure occupation

may also be a remote-work occupation, a sector hit by higher finance costs, or a type of role whose vacancy counts already peaked unusually high in 2021–2022.

These are not inconvenient details to be buried. They are competing hypotheses. If a firm's junior hiring fell at the same time as its AI use increased, the appropriate next question is whether comparable firms facing similar labour costs and demand—but without comparable AI integration—experienced the same decline. Until that comparison is made, the claim is suggestive, not causal.

7.3 International evidence: both contraction and expansion

The American evidence is important precisely because it is not the United Kingdom. In the August 2026 revision of *Canaries in the Coal Mine*, Brynjolfsson, Chandar and Chen found that employment among workers aged **22–25 in highly AI-exposed occupations** was approximately **19 per cent below the path implied by similarly aged workers in less-exposed occupations** [14]. This is a **relative employment gap**, not a 19 per cent fall in all young people's employment. The authors regard their findings as descriptive, note sensitivity to controls and pre-existing differences, and report concentration in their payroll sample. Firm-level controls place the visible divergence later than some simpler comparisons, around 2024. The figure should neither be applied mechanically to Britain nor presented as proof of a national youth-employment collapse.

The same study reports that the declines are concentrated in occupations where AI is used to automate work rather than to augment it [14]. If that pattern holds, it matters for section 8, because the distinction between substitution and complementarity is exactly what the training-cost mechanism turns on.

Other evidence runs in the opposite direction. Chandar and Klein Teeselink's September 2026 study of AI-adopting firms across **41 countries** found a declining *share* of junior employment associated principally with faster growth among senior workers, alongside suggestive overall employment growth [18]. This is particularly instructive because Klein Teeselink also authored the UK contraction study. Different samples, adoption definitions and research designs yield different estimates. The proper response is to investigate the differences, not to select the one that suits the argument.

Indeed Hiring Lab reported in July 2026 that the previously negative relationship between AI exposure and US job-posting change appeared to be reversing as highly exposed occupations took a stronger part in a more recent rebound [19]. Its later analysis found a marked shift towards senior advertisements in exposed US occupations, but the same evidence cannot establish whether jobs were upgraded or beginner places displaced. PwC's discussion of increasingly senior skill requirements in advertised entry-level jobs also draws substantially on American recruitment data [24]. None of these analyses is a direct count of missing junior posts in the United Kingdom. The underlying platforms and commercial reports have their own composition effects and interests; they identify changing patterns, not a complete labour-market census.

Westby, Modestino and Cheng find, in US vacancy data for software development, that the junior vacancies remaining after ChatGPT's release shifted towards problem solving, communication and attention to detail [52]. That suggests a raised bar for entrants rather than their disappearance, and it is compatible with both the substitution and the complementarity accounts.

At a broader scale, New York Federal Reserve economists found little distinct AI-driven decline in the US job postings they studied [28]. The Yale Budget Lab's October 2026 update also found no clear economy-wide labour disruption associated with AI in its employment measures [29]. The European Central Bank reported survey evidence in which heavier AI investment was often associated with increased staffing to operate the technology, while recognising possible substitution in particular settings [30]. None of these findings proves the absence of displacement among a narrow group of

British graduates. They do, however, challenge the claim that high technical exposure inevitably reduces employment everywhere.

The International Labour Organization's review of empirical evidence similarly describes uneven productivity effects, possible threats to young workers' entry opportunities, and limited support for declaring general mass displacement [8]. Its separate work on exposure indices warns against treating a technical exposure score as the realised probability of redundancy [9]. Exposure indicates tasks that *might* change. Actual hiring depends on adoption, demand, prices, institutions and employer choices.

The conclusion from these studies is not indecision. It is that **effects should be estimated at the level where the mechanism operates**: firms, occupations, new hires and training routes. An economy-wide null result can coexist with a genuine local loss; a local loss cannot establish a national catastrophe.

8 The economics of training and the missing first rung

8.1 The mechanism and its predecessors

Imagine a legal research department in which beginners traditionally learnt by summarising cases, verifying authorities and drafting preliminary memoranda. These activities do not constitute the whole profession, but they have formed part of its entry route. If a system takes over much of that work, the firm has choices. It can use the saved time to broaden apprentices' responsibilities and improve supervision. It can recruit fewer beginners. Or it can expand output and need more workers of several kinds. The software cannot tell us which managerial decision follows.

The missing first rung hypothesis is that, under some training arrangements, replacing commercially useful beginner tasks changes the economics of developing future experts. In conventional workplace training an employer pays wages, supervision and equipment costs but also receives value from the work the trainee completes. Garicano and Rayo make this the centre of a formal model: juniors pay for training by doing menial tasks, so AI that absorbs those tasks puts the bargain at risk, while AI that raises the value of a trained expert can strengthen it. Which effect dominates depends on how the AI-augmented value of a graduate compares with what AI produces on its own; careers survive above a critical threshold of that ratio and training compresses below it [49]. Bessen's finding that learning-by-doing is largely firm-specific, and comparable in scale to physical investment [35], matters here: an externality arises only to the extent that what beginners learn is portable between employers.

Germany's BIBB cost-benefit surveys and OECD apprenticeship research measure wages, supervision and productive work as separate components of net training cost, and treat retention benefits separately [43, 44]. Their findings show how such a calculation can be made. They do not establish that AI has raised training costs in Britain, and this review did not establish whether a routine UK equivalent exists.

8.2 Net private training cost

Let NC be the net private cost to an employer of one additional trainee over the training period:

$$NC = W + S + O - V$$

- **W**: wages and statutory on-costs paid to the trainee.
- **S**: employer-paid supervision, review and mentoring time.
- **O**: other direct costs, such as equipment, courses and administration.

- **V**: the quality-adjusted value of the trainee's productive contribution, measured at the cost of the firm's best alternative way of obtaining that work: a qualified employee, an outsourced provider or an AI system, including the cost of checking its output.

Expected retention benefits, R, are valued separately: avoided recruitment and induction costs, and any productivity premium of a firm-trained worker over the wage paid, discounted and weighted by the probability that the worker stays. A firm has a private incentive to train when expected R exceeds NC. Both are marginal concepts; surveys report averages, which can differ from the cost of the next trainee.

AI can act on each term. It can reduce V by removing tasks and by lowering the cost of the alternative against which they are valued; it can raise V by letting trainees perform higher-value work sooner. It can raise or lower S, since reviewing AI-assisted work takes time but may also reduce correction. It can leave W unchanged where pay cannot adjust (section 8.4). And it can raise R, by making trained experts more productive, or lower it, by making trained staff easier to recruit elsewhere. The effect on the decision to train is therefore ambiguous in theory, which is why the empirical question is a real one.

The valuation of V deserves emphasis because it is easily missed. A trainee who performs exactly the same tasks as before produces less value for the firm if an AI system can now produce the same output more cheaply, because V is measured against the best alternative [44]. The fall is a price effect, not only a quantity effect, and it can occur without any visible change in what trainees do.

This specification also clarifies what would make the mechanism AI-specific. The research needs independent records of which junior tasks changed, what output they produced before and after adoption, which costs shifted, and how that relates to paid training starts. A change in training budgets alone is inadequate: a firm could spend less while providing better instruction. The empirical contribution proposed here is narrower than the theory it follows [46, 49, 50]: to test the private-cost channel against complementary productivity gains, wage-floor interactions and profession-wide outcomes.

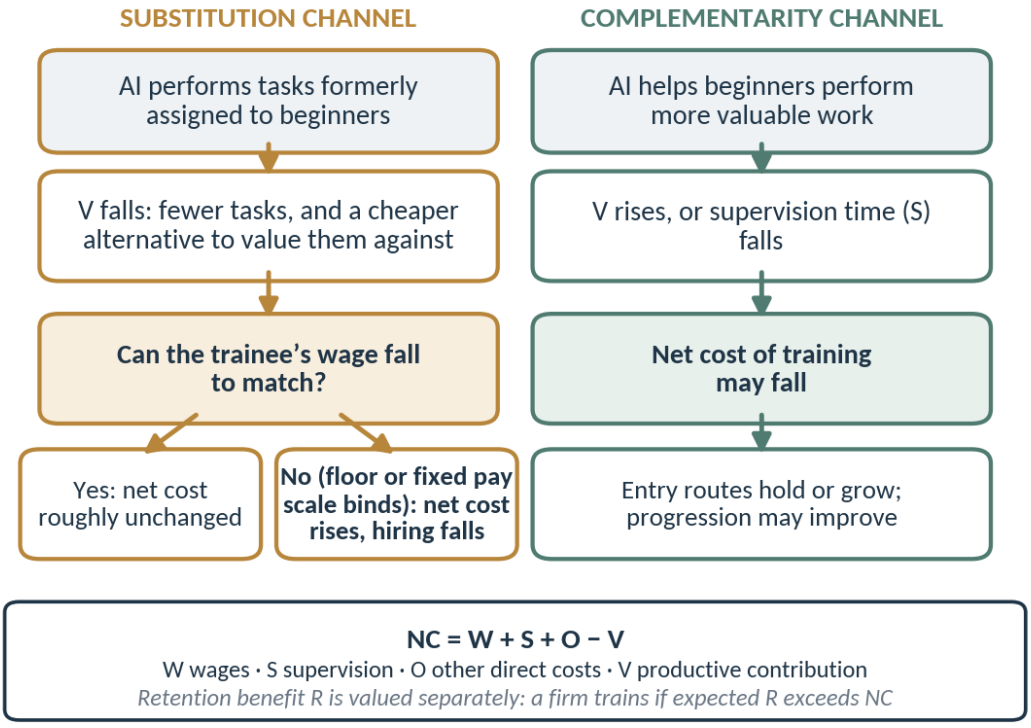

Figure 5. Two possible AI effects on training economics, including the role of wage adjustment. The arrows show hypotheses, not measured British changes. Productive contributions must be quality-adjusted; retention benefits are assessed separately.

8.3 What would discriminate substitution from complementarity

Table 2 lists observable quantities on which the two accounts make different predictions. No single row is decisive. A pattern across several, observed in the same occupations and periods, would be.

Table 2. Observable quantities that distinguish substitution from complementarity. Predictions are hypotheses; no row is decisive alone.

Observable	Substitution predicts	Complementarity predicts	Possible source
Share of trainee time on tasks AI now performs	Falls	Stable or rises, with higher-value tasks added	Time-use records; task logs
Value of trainee output against the best alternative (V)	Falls, even where tasks are unchanged	Rises	Costing exercises; quality-adjusted output
Supervision and review time (S)	Unchanged or rises	Falls	Supervisor records
Pattern across AI uses	Entry declines where AI automates work	Entry holds where AI augments work	Task-level exposure measures
Bite of statutory wage floors	Larger fall where pay is near a floor	No interaction	Pay data by age band
Paid starts and progression	Weak relative to demand	Stable or improving	Apprenticeship and employer records
Recruitment of already-trained staff	Rises as a share of senior hiring	No change	Vacancy and mobility data

8.4 AI and statutory wage floors

Three hypotheses must be kept apart. The first is a general cost effect: higher minimum wages or employer contributions reduce hiring of young people whatever the technology. The IFS found no clear evidence that this was a major driver of the recent deterioration [16]. The second is an adoption effect: higher wage floors make labour-saving technology more attractive. US evidence links higher minimum wages to greater industrial robot adoption in manufacturing [47], and earlier US work finds a related reallocation away from automatable jobs [51]. Neither concerns generative AI in UK entry routes.

The third is the interaction hypothesis. If AI lowers the value of a beginner's work, a competitive market would lower the trainee's wage and the net cost of training would be broadly unchanged. Where a statutory floor or a fixed apprentice pay scale prevents that adjustment, the cost rises and the employer responds by hiring fewer trainees, not by paying them less. OECD work already treats the relation between apprentice pay and productivity as central to the employer's cost-benefit balance [44]. The interaction is therefore a plausible implication of established accounting, not a new finding, and it is untested for the UK.

It should operate only under particular conditions. Entrants must be paid close to a binding floor or a fixed scale. That makes it more plausible in lower-paid clerical, customer-service and support roles, where the ONS found administrative and clerical work most often named as affected [1], than in the graduate professions that dominate public debate, where entrants earn well above any floor. Age-banded rates and their changes provide variation, but the bite of the floor is correlated with sector, region and employer size. A test would examine whether the AI-related change in entry hiring or paid starts differs by the share of entry posts paid near a floor, with National Insurance thresholds modelled separately. A null result must remain a live possibility, and the hypothesis should not be presented as an established explanation of UK hiring.

8.5 Counter-evidence

There is substantial counter-evidence against assuming that AI necessarily makes training more expensive. A field experiment involving 673 final-year IT apprentices in German vocational schools found that AI access increased success on tasks beyond their formal curriculum by 26–31 percentage points without measured deterioration in immediate comprehension [20]. Separately, Muehlemann's study of German establishments found that sustained AI adoption was associated with approximately 14 per cent more new apprenticeship contracts among firms already offering training, without a statistically significant change in whether firms provided apprenticeships at all, alongside a shift of continuing-training resources towards high-skilled employees [45]. Neither study establishes long-term UK professional outcomes. Together they make it essential to test complementarity as seriously as displacement.

The published *Quarterly Journal of Economics* study *Generative AI at Work* likewise found that AI assistance produced larger productivity gains for less experienced customer-support workers than for experienced ones [21]. Assistance can narrow immediate performance gaps. Whether that improvement also encourages companies to offer more entry jobs is a separate empirical question.

The danger, then, is not that the machine takes away boring tasks as such. There is no educational virtue in making a junior employee spend six hours doing something a reliable system can do in six minutes. The danger arises if employers remove the tasks **and** the paid opportunity to practise verification, judgement, responsibility and communication. Conversely, the opportunity arises if they replace repetitive labour with better supervised learning. It is the design of the apprenticeship, not nostalgia for the photocopier, that deserves protection.

The question becomes measurable by tracking training economics and developmental outcomes together: net cost per trainee, productive contribution, supervisory time, skill assessment, paid starts and progression. If firms that substitute AI for junior production also reduce entry hiring, is it because the value of supervised learning fell, because the net private cost rose, or because demand for the occupation changed? An employer claiming to have 'seniorised' its workforce must answer a straightforward question: where will its future seniors come from?

9 Research programme: from visible symptoms to an identifiable mechanism

The paper cannot establish an industry-wide training failure merely by observing that an AI-using firm hires fewer juniors. The empirical programme must therefore separate **what can be observed now**, **what requires employer cooperation**, and **what evidence would actually distinguish the proposed mechanisms**. The 2022–2032 horizon is a period for observing recruitment and adaptation, with earlier baselines where available; it is not a forecast of when disruption begins. Producing fewer future senior professionals is a longer-term claim and may require observations well beyond 2032.

9.1 Begin with occupations, not a promised employer database

Stage one is a feasible, primarily descriptive UK occupation-level study. It should map changes in paid entry routes by occupation and region using available ONS employment and vacancy series, apprenticeship starts and completions, age- and tenure-related employment measures where sample sizes permit, occupational classifications, published online job advertisements and available indicators of local demand. Official data should anchor the descriptive analysis; platform data can corroborate but cannot replace representative employment series. The LinkedIn–government entry-hiring brief [41] is especially useful as an aggregate counterpoint. The minimum publishable result of this first

stage is a transparent account of where signals occur and where they do not, not a causal claim that AI has hollowed out a training pipeline.

Every series requires a data dictionary: population, job or firm unit, occupation code, observation frequency, geographical coverage, entry-level definition, revision status and whether it measures an advert, a successful hire, paid employment or completed training. AI-task exposure is a pre-period occupation-level risk indicator, not a measure of actual adoption or automation. The ONS Business Insights and Conditions Survey measures business use in changing business samples and industry categories; it cannot be mapped mechanically to individual occupations or linked to payroll records by assumption [39]. Online adverts may be reposted and a vacancy is not a new employee. These distinctions determine what this first-stage study can conclude.

The first work package is an **access audit**. It will distinguish public tables from secure or restricted microdata, including possible ONS Research Accreditation and UK Data Service requirements [40]. If occupation-by-age cells are suppressed or too small, the study should aggregate transparently and report the loss of specificity. If no trustworthy junior-hiring count exists, the analysis must call its outcome an *entry-opportunity proxy*, not a directly observed flow of new professionals.

Age and tenure are imperfect substitutes for occupational seniority. A young employee can have considerable experience, and a fifty-year-old changing occupation may be a beginner. Where possible, classify entry roles using a fixed pre-period occupational job-title dictionary and validate a sample manually. Report separate results for graduate pathways, vocational apprenticeships and casual youth work; the IFS evidence makes clear that these are not a single labour-market population [16].

For each occupation, the descriptive comparison would place entry-opportunity indicators alongside pre-2022 trends, fixed pre-period AI-task exposure, general vacancy change, remote-working intensity, demand, pay and employer-cost exposure. It would test pre-trends, placebo windows and competing explanations. Because occupational observations are correlated within a small number of groups, conventional standard errors can give false precision. Cluster-aware uncertainty, feasible wild-cluster bootstrap sensitivity and appropriate small-sample or permutation checks should be considered, while recognising that even wild bootstrap procedures can fail with very few treated groups or uneven clusters [48]. Schlicht's informative association rests on 26 occupational groups [12]; no large count of individual job adverts creates 26,000 independent occupations. Stage one maps signals; only additional adoption and mechanism evidence could support a causal claim.

9.2 Define outcomes that reveal the mechanism

There is a difference between finding a weaker entry market and finding an economic underinvestment mechanism. The second requires observations about the costs and productive returns to training, employer choices and how professions replenish expertise. These are the proposed outcome families, with net private training cost per entrant the first mechanism target:

- **Training economics**: net cost per trainee (NC); wages, employer-paid supervision, equipment and other direct costs; quality-adjusted productive contribution valued against the best alternative; retention-adjusted expected returns assessed separately; and variance across occupations and training cohorts.
- **Paid entry and learning**: absolute new entrants, junior hires per 100 employees, paid starts, supervised hours, competence assessments, completion, progression to qualified roles and the ratio of newly qualified practitioners to departures.
- **Movement and recruitment**: external senior hires as a share of senior recruitment, inter-firm mobility and retention, vacancies for qualified people, pay premia linked to experience, and professional-body evidence about skill replenishment.

- **Potential modifiers**: share of entrants paid near a binding statutory wage floor, National Insurance liability, AI-verified task substitution or augmentation, employer training agreements, levy participation and professional accreditation obligations.

These outcomes are not independent proofs. A rise in net training cost after AI adoption may reflect more intensive or better-quality training, not necessarily harmful substitution. A fall in cost may conceal reduced supervised practice. Wage premia can rise when demand is healthy, and employee mobility may increase during expansion. The underinvestment explanation requires related changes in costs, training opportunities, employer retention expectations, professional recruitment and demand. An observable shortage by itself proves no economic externality.

Junior **shares and counts** must always be reported separately. An industry may increase its number of young entrants while expanding senior employment faster, reducing the junior share without eliminating a single junior job [18]. Paid work and training quality must also be distinguished: a higher number of apprentices completing trivial simulated exercises is not automatically a stronger supply of qualified professionals.

Stage one can measure only part of this package. Employer training budgets, poaching, verified integration dates and career progression generally require a **conditional stage-two employer and professional-body study**, not an imagined field inside public ONS tables.

9.3 Why the occupation or profession is the main unit

The first-rung problem is a potential **industry-wide training externality**. It cannot be demonstrated by a regression showing that one firm's junior payroll fell after installing software. That firm might simply have shifted its recruitment to more experienced staff while another employer absorbed the beginners. The national supply of trained people would be unchanged, although the burden of training had moved.

The primary unit for assessing a collective failure should therefore be an **occupation within a labour market**, with appropriately defined professional or sector boundaries. The evidence package is the total flow of new entrants, the amount and quality of paid training, qualification and progression, inter-employer transfers, experienced-worker vacancies and wage premia, all considered alongside occupational demand and output. Occupations and industries are not interchangeable: a bank employs cleaners as well as accountants, and the relevant training market may span several sectors. Schlicht's occupational results are stronger than the corresponding broad industry averages, illustrating that point [12].

The firm level remains valuable **for observing managerial mechanisms**—documented AI substitution, budgets, supervisory changes and decisions about training and replacement hiring. It is a supplementary explanation of *how* entry routes are altered, not the level at which to declare collective underproduction of skills. Where employers share a training system or qualification, the research should also study the professional body, college or apprenticeship consortium, not merely the firms easiest to recruit for interviews.

9.4 Test the baseline before classifying AI-related pathways

The proposed decision has two stages. First estimate wider occupational demand, wage and policy components, pre-existing trends and plausible spillovers. If a credible AI-related differential cannot be distinguished, record an aggregate baseline or an inconclusive result. Only where supported by actual integration or robust proxy evidence should the study examine the conditional patterns below. Figure 6 is conceptual, not observed employment data.

Path A — Temporary adaptation. Paid entry routes initially contract in occupations with independently documented workflow changes, then training starts, quality and progression recover as supervisors redesign entry work. Recovery must exceed what is explained by the general hiring cycle.

Path B — Persistent training underinvestment. The productive contribution that formerly financed part of junior learning falls relative to wages and supervision; the employer's net private cost or expected unrecouped cost of training rises; paid starts and replenishment remain weak relative to demand. Increased recruitment of already-trained workers from other firms adds mechanism evidence. The cost and recruitment sequence matters more than persistence alone.

Path C — Complementary expansion. The quality-adjusted productivity of beginners increases, or their training and supervision become less costly; supervised entry, completion or progression remain stable or improve. Higher demand and output may contribute, so expansion should not automatically be credited to AI.

Occupations need not belong to a single pathway. If one set of employers expands training while another reduces it, report both and measure their relative scale. If the study lacks convincing adoption information or cannot exclude macroeconomic explanations, the correct label is indeterminate. The non-AI baseline is a comparator, not a fourth causal mechanism.

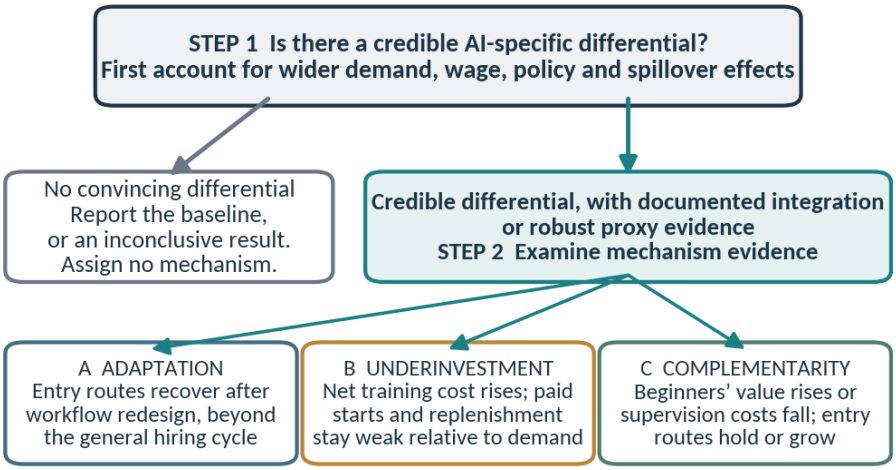

Mixed outcomes across occupations and employers are allowed; no pathway is assigned from vacancy trends alone.

Conceptual decision guide, not measured data.

Figure 6. Two-stage interpretation: estimate the wider labour-market baseline before distinguishing three conditional AI-related pathways. Conceptual decision guide, not measured data. Mixed and indeterminate outcomes are permitted.

For each outcome, the direction and period of comparison should be specified **before inspecting the relevant result**. A provisional *meaningful-effect planning rule* is a decline of two percentage points in the proportion of hires that are junior, accompanied by a separately material fall in absolute paid starts, sustained over four quarters. It is illustrative rather than discovered or preregistered. The occupation-level baseline, available number of hiring events and the smallest socially meaningful loss must set the final margin during a documented pilot.

An estimate whose confidence interval crosses zero is not evidence of no harm. Equivalence tests should ask whether data are precise enough to rule out pre-specified meaningful losses [38]. The report should distinguish meaningful contraction, meaningful expansion, sufficiently precise evidence of no practically material change, and an inconclusive result. No scenario should be awarded victory because a wide interval happens to include its preferred number.

9.5 Output and the meaning of substitution

A reduction in headcount is difficult to interpret without a corresponding output measure. A consultant's revenue depends on prices and client mix, not only work completed. A law firm's billable hours do not measure whether advice was good; a public agency can answer more enquiries while missing difficult ones. The study should seek task-appropriate output measures before analysing staffing changes: quality-checked cases, verified software changes, matters resolved, service volume adjusted for complexity, error rates, waiting times and real revenue where a defensible price adjustment exists. Supervision and correction time must be included.

When comparable output cannot be established, the outcome should be described as **entry-route contraction associated with documented AI integration**, not confidently as output-adjusted substitution. This preserves the difference between evidence and assumption.

9.6 Duration, cohort follow-up and response time

The **2022–2032 period** is sufficient for studying early adoption, recruitment changes and some training progression, with pre-2022 historical baselines added where reliable. It is unlikely to establish the eventual supply of senior professionals in every field. The lag between entering an occupation and practising independently varies considerably. Where the hypothesis predicts a shortage of future experts, a cohort extension into the late 2030s or beyond may be required.

Response time should be judged against the length of the training cycle, as in section 2.2. For each occupation the study should record the typical period from entry to independent practice, and ask whether the relevant institutional responses—redesigned supervised routes, revised standards or guidance, collective training arrangements—appeared within one cycle of documented integration. This keeps the comparison relative to the pipeline and avoids any composite measure of speed.

Early indicators include trainee entry, spending per entrant, completion, movement into qualified roles, experienced-staff vacancies, inter-firm recruitment, experience premia and measurable skills gaps. **They are primary outcomes for the mechanism even though some will be incomplete by 2032.** The study must distinguish indicators available within the observation period from the later causal claim about an actual shortage of future seniors. Each has alternative explanations, and none should be presented as a prophecy fulfilled before an appropriate cohort has had time to qualify.

9.7 Evidence weighting

The research must remain open to a result that contradicts the proposed training-cost channel. A temporary decline in paid entry followed by documented training redesign supports adaptation only if the recovery is distinct from general labour demand. A sustained increase in net private training cost accompanying lower productive junior output, reduced paid starts and greater dependence on recruiting experienced staff provides stronger grounds to investigate underinvestment. Falling net costs alongside strong progression can support complementarity. If an AI-specific differential cannot be distinguished from the wider baseline, decline to assign a mechanism. Even a convincing cost change is not proof of a future shortage until cohorts have had time to qualify.

Evidence should be weighted by fit to the question, study design and uncertainty: representative official series for employment levels and trends, independently documented employer and profession records for actual decisions, credible comparison designs where their assumptions survive, and commercial or platform samples for qualified corroboration. This is not a mechanical league table. A careful descriptive series may answer a question that a sophisticated causal design cannot. An observational UK study that shows fewer juniors cannot prove that firms collectively underinvest in training; that stronger claim needs direct evidence of decision-making incentives and net profession-wide outcomes. The aim is to distinguish what has been observed, what is inferred and what remains beyond available evidence.

10 Interpreting the results

UK studies already report pressure in particular exposed occupations and firms over **2021–2026** [11, 12]. The American study [14] concerns a different labour market and is not evidence of UK hiring dates. Without a measured adoption curve, cohort model and defined national threshold, it would be unsound to nominate 2028–2032 as the moment employment disruption suddenly becomes visible. Economic and social effects may unfold on several schedules, or may never become large in aggregate.

The more defensible question is **what changes become visible at different levels**. A new graduate may notice fewer genuine junior vacancies before an employed accountant experiences any change. A recruitment agency may notice a reduced flow of assignments while national employment remains steady. A professional body may not notice a missing supply of newly trained practitioners for several years. A household notices a change when it must support an adult child for longer than expected. These are different observations, with different time delays.

Interpretation follows the same sequence as the research design: first the non-AI employment baseline, then three conditional AI-related pathways if a credible differential exists. Different occupations, and employers within them, may display different outcomes. No numerical probability is assigned to a pathway.

Baseline or indeterminate finding — Where entry hiring falls comparably across occupations regardless of credible AI integration, demand and labour-cost explanations take priority. The June 2026 UK LinkedIn–government brief [41] and IFS evidence [16] argue against imposing a universal AI-specific junior collapse. Failure to distinguish an effect is not proof of zero AI influence; it may simply reveal inadequate measurement or statistical power.

A — Temporary adaptation lag. Following a credible AI-specific change, entry routes initially weaken, then recover as firms redesign supervised work and integrate new tools. Any recovery must be assessed against the wider labour market rather than attributed to technology merely because vacancies rise.

B — Persistent training underinvestment. Once AI is integrated, firms continue to offer too few paid entry positions relative to viable demand. The candidate AI-specific mechanism is a rise in the net cost of training when software replaces commercially useful junior tasks, alongside limited ability to capture the subsequent return on transferable skills. Increasing reliance on recruiting qualified people from competitors, weaker progression and eventual shortages are testable consequences, not assumptions.

C — Complementary expansion. AI increases beginners' productive contribution, reduces supervision costs or creates valuable new assignments. Paid entry starts and progression remain stable or increase, potentially alongside service expansion. This outcome has meaningful support in recent apprenticeship and productivity evidence [20, 45].

What the case can say about the wider thesis is limited but real. Under pathway A, institutions in this domain adapted within the cycle that matters, and the mismatch was temporary. Under pathway B, adaptation did not occur within that cycle, and the explanation is a private cost that firms do not recover; the thesis of a persistent mismatch then holds for this domain, with training incentives as its mechanism. Under pathway C, AI strengthened the institutions that renew expertise. Under the baseline finding, the thesis finds no support here, and weak early-career hiring is better explained by demand and policy. None of these results transfers automatically to other domains.

The practical policy question is not whether the United Kingdom is destined to cross one cliff in a particular year. It is whether evidence of selective hiring substitution can be detected early enough for education, employers and public services to respond while people still have real choices.

The question in the title remains worth asking: how quickly do technical capability, workplace integration and institutional adaptation change, and does the last keep up with the first two? The answer proposed here is to break the question into domain-specific tests against the length of the relevant training cycle. The title is a question, not a finding.

11 Policy and educational implications

Policy should protect the developmental function of paid entry work rather than preserve obsolete tasks. Where an AI system produces the initial research or draft, beginners can instead learn to validate it, manage exceptions, communicate with clients and exercise accountable judgement. Employers should measure the quality and cost of this redesigned training, including verified trainee output and the time supervisors actually spend. If productivity gains reduce net training cost, the case for new paid entry routes strengthens. If task substitution makes them privately more expensive while an occupation still needs future experts, collective financing or shared training institutions may be warranted. Both possibilities require measurement.

Employers who benefit collectively from a skilled profession may need sector-wide training commitments, portable credentials, apprenticeship partnerships or public support where the return to any single firm's investment is insufficient [23]. Yet subsidy should follow measured shortages or demonstrable training failures rather than vague predictions of AI collapse. Young people should not be instructed to retrain for occupations that employers have stopped offering at entry level.

Statistical and careers services should publish **hiring flows** alongside employment stocks. A dashboard distinguishing junior vacancies, actual hires, paid traineeships, progression, pay, working time and employer AI integration would be more useful than a single total of 'jobs at risk'. Such reporting must preserve occupational definitions and note where changes in data collection distort apparent trends. It should cover regions, disabled workers, gender, qualification levels and people who cannot afford unpaid training. An average can conceal very different life chances.

There is also an ethical issue of accountability. Firms remain answerable for the accuracy of their work and the consequences of staffing decisions. Using AI to verify a junior colleague's first draft is an organisational choice. Using AI as the explanation for every unfilled vacancy is another. Public debate should demand evidence of both.

Finally, the benefits count too. Lower service costs, shorter waiting times, new businesses, improved access and jobs made possible by AI are social outcomes of real value. An account that ignores them would be as selective as a company presentation that speaks only about efficiency while avoiding its recruitment figures. An honest policy design asks who gains, who pays, how long adjustment takes and whether people can exercise meaningful agency during it.

12 Limits and contribution

This is an independent conceptual and evidence-synthesis paper, not a completed causal study, a new statistical index or a forecast. The broad thesis cannot be tested as stated: institutional adaptation is not measured systematically, and the case examined here covers one domain. The most promising UK findings are observational and may contain unmeasured confounding. Technical occupational exposure is not documented firm adoption, and industry-level business surveys cannot be attached to occupational payroll records without further evidence. A business report of no job cuts can coexist with weaker entry recruitment. A decline in junior share can coexist with junior job growth. A foreign apprenticeship result does not establish British graduate outcomes. Cost accounting is proposed, not performed, and expected long-term returns are hard to value. These limits are conditions on the proposed contribution, not details to remove in pursuit of a cleaner story.

The initial UK occupation study can, at best, locate descriptive signals and counter-signals using consistent data definitions and uncertainty estimates robust to small numbers of clusters. It cannot identify the net private training-cost channel or an industry-wide externality without additional evidence. Employer records of validated output, supervisor time, wage costs, task redesign and retention may be unavailable or incomparable, and willingness to participate may produce selection bias. Wage floors may matter only in specific pay bands. The 3,200-per-group illustration is explicitly unclustered and is not an empirical sample-size claim. Inconclusive causal identification is an acceptable outcome.

The contribution is not a new theory. The training externality is old [23]; AI-specific versions are already proposed [46, 49, 50]; a measurable net-cost framework exists [43, 44]; German evidence demonstrates a credible complementary result [45]; and the interaction of wage floors with automation has US precedents [47, 51]. The paper adds a design: a comparison of institutional response with the length of the pipeline it protects, in place of a composite speed gap; a baseline-first inference that assigns no mechanism without a credible AI-specific differential; a net-cost specification that values trainees' work against the firm's best alternative and treats the wage-floor interaction as a separate hypothesis; and explicit conditions under which the answer would be inconclusive. Originality should be judged on executed measurement and on discrimination between mechanisms, not on a new label for a familiar concern.

13 Conclusion

How fast are we going? The honest answer is that we are travelling at several speeds at once, and we have not yet built instruments that measure every one of them equally well.

Technical capabilities in selected domains are improving quickly. UK businesses increasingly report using AI, though usage is not equivalent to wholesale redesign. Some UK studies suggest concentrated pressure on highly exposed hiring and junior positions, while broad surveys and international counter-evidence make claims of inevitable mass displacement unsupportable. The country also has a weak employment market, rising labour costs and longstanding difficulties facing young people. It would be a serious error to attribute all of them to the arrival of an AI assistant.

What deserves attention is the first paid opportunity to become good at something. A society can continue producing graduates while quietly diminishing the settings in which graduates acquire practical judgement. But it can also use AI to make apprentices more capable and open better routes into work. These futures depend on employers' incentives, educational design, public choices and the direction of demand—not simply on the length of time a model can complete a coding benchmark.

The proposed 2022–2032 observation programme will first identify wider demand and policy patterns and ask whether any AI-related change can credibly be distinguished. Where it can, the programme will test whether the private economics of training shifted towards temporary adaptation, persistent underinvestment or complementary expansion, and whether institutions responded within the length of the training cycle. If opportunities grow, record that. If junior work is automated but effective paid learning survives, record that too. If training costs rise and paid starts fall despite sufficient demand, investigate the incentives. If measurements are inadequate, accept uncertainty. A follow-up beyond 2032 may be necessary to learn whether a present shortage of entrants becomes a later shortage of experts.

Speed is impressive. Understanding who gets to move with it is rather more important.

Appendix A Technical research methods

A.1 Spillovers and comparison-group contamination

Conventional causal comparisons often rely on a version of the **Stable Unit Treatment Value Assumption (SUTVA)**: one unit's outcome is not changed by the treatment status of another. That assumption is particularly strained here. If an AI-adopting firm stops recruiting juniors, its unused training capacity may shift to other employers; graduates may compete harder for their vacancies; wage offers, client demand and the supply of experienced staff may change for both adopters and non-adopters. The supposed control group has itself been affected.

This is part of the mechanism, not merely a technical footnote. At firm level, estimates may describe redistribution between employers rather than the net loss of opportunities in an occupation. The analysis should map flows into other employers, out of the occupation and into training or inactivity; track hiring by both adopters and non-adopters; and report how conclusions change when the observation unit widens. Spillovers can cross occupation or regional boundaries, so aggregation reduces one form of interference without removing it.

Callaway–Sant'Anna and Sun–Abraham estimators may be appropriate in a carefully justified staggered-adoption panel [36, 37], but they require credible comparison groups and assumptions. **As adoption approaches universality, genuinely untreated or not-yet-treated controls may disappear.** These estimators cannot recover a causal effect through vocabulary alone. If treatment contaminates all plausible controls, the paper should report bounds, transparent associations and qualitative mechanisms rather than claim that a causal difference-in-differences estimate remains available.

A.2 Employer-panel feasibility and identification

A second-stage employer panel is conditional on documented access and informed participation. A pilot must determine whether willing employers and professional bodies can supply anonymised, comparable records of trainee wages, supervision time, assessed productive contribution, other direct training costs, paid starts, progression, recruiting sources and verified output. Prospective later benefits, including retained trained workers and avoided recruitment costs, must be collected and reported separately. Record how nonparticipating firms differ from participants, obtain ethical approval where required and protect personal data.

For each employer, independently document when and how AI changed task assignment, what output required correction, and whether verified junior production, supervisor time or training quality changed. A company calling itself 'AI-first' is not an integration measure. Fix event definitions from implementation records before examining subsequent recruitment. A firm's supposed integration-stabilisation date must not be chosen retrospectively when hiring happens to recover.

Employer outcomes should include net private training cost per entrant and its components, separately valued retention benefits, absolute junior and senior hires, hiring rates, paid starts, verified competence and recruitment from other firms. Zero-heavy junior-hire counts may require count models, aggregation or two-part methods, and training-cost estimates need sensitivity to task valuation and unmeasured overhead. The panel explains decisions; it does not by itself identify whether a profession produces too few trained people.

AI adoption is endogenous: expanding firms and firms under cost pressure may adopt sooner and alter recruitment independently. Where credible rollout variation exists, staggered-adoption designs can be considered subject to pre-trends and interference diagnostics [36, 37]. A possible instrument based on pre-period task exposure interacted with vendor rollout timing would require a defensible exclusion restriction, which is not currently established. Stronger causal language is not justified by

naming a modern estimator, especially when comparison firms share labour markets and adoption eventually becomes near-universal.

Sensitivity work should include placebo dates, alternative exposure indices, pre-2022 trends, negative-control occupations, regional demand, remote working, wage levels, occupation and firm size, age-specific minimum-wage and National Insurance exposure, and a separately pre-specified AI × binding-wage-floor interaction. For a firm or occupational panel, define the pre-period share of entry posts paid near a statutory wage threshold, validate the statutory age bands, and interact it with independently documented task substitution rather than occupational exposure alone. The parameter is not causally identified if AI integration is endogenous or costs and demand move together. A null wage-floor interaction must remain a live possibility; the IFS found policy effects difficult to separate even in a study devoted to them [16].

Figure A1. Conditional research sequence: occupational signal mapping, employer and training mechanism studies if access is obtained, and longer cohort follow-up. Conceptual study design, not observed data.

A.3 Statistical power and decision thresholds

A minimally credible preregistration must specify the primary outcome (net cost per trainee for the conditional mechanism study), expected baseline levels, a minimum effect of practical importance, number of independent occupational and employer clusters, pre/post periods, within-cluster correlation, treatment or adoption prevalence, likely spillovers and sample attrition. A power study must simulate the observable cluster structure. It must also specify what evidence would count against harmful underinvestment, including credible complementarity and sufficiently precise equivalence results.

As a deliberately simplified illustration, detecting a junior share of 10 per cent versus 8 per cent with 80 per cent power and a two-sided 5 per cent test requires roughly 3,200 independent hiring observations per group under equal-sample unclustered assumptions. This is not a sample size for the proposed occupation-level study. A small number of occupational clusters may make narrow causal contrasts infeasible regardless of millions of vacancies; wild-cluster bootstrap is a sensitivity method, not a cure when treated groups are very few or cluster sizes unequal [48]. Surveying employer-level net training costs will require a separate variance estimate and sample design. If access fails, the study ends honestly at descriptive signal mapping.

The final design should use an equivalence margin determined with sector experts and affected entrants, not selected after seeing the estimates. For small hiring counts, absolute opportunities and rates should be tested alongside shares, with uncertainty intervals and sensitivity analysis. The

primary outcome family should be kept small enough to avoid declaring success whenever one of many correlations reaches significance. When data are underpowered, **inconclusive** is an intellectually respectable result.

References

[1] Office for National Statistics (2026). *Business insights and impact on the UK economy: 8 October 2026* (September BICS reporting wave). https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/bulletins/businessinsightsandimpactontheukeconomy/8october2026

[2] Office for National Statistics (2026). *Artificial intelligence in UK businesses: 2023 to 2026*. Published 20 July; June survey data. https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/articles/artificialintelligenceinukbusinesses/2023to2026

[3] Department for Science, Innovation and Technology (2026). *UK Business Data Survey 2026*. https://www.gov.uk/government/statistics/uk-business-data-survey-2026/uk-business-data-survey-2026

[4] METR (2026). *Time Horizon 1.1*. 29 January. https://metr.org/blog/2026-01-29-time-horizon-1-1/

[5] UK Government (2026). *Assessment of AI capabilities and the impact on the UK labour market*. 28 January. https://www.gov.uk/government/publications/assessment-of-ai-capabilities-and-the-impact-on-the-uk-labour-market/assessment-of-ai-capabilities-and-the-impact-on-the-uk-labour-market

[6] Anthropic Institute (2026). *When AI builds itself*. https://www.anthropic.com/institute/recursive-self-improvement

[7] Sarma, G. P., Steratore, R., Jacob, M. and Bhatt, S. D. (2026). *Recursive Self-Improvement and Automated AI Research and Development*. RAND Corporation, PE-A5068-1, 12 October. https://doi.org/10.7249/PEA5068-1. https://www.rand.org/pubs/perspectives/PEA5068-1.html

[8] Merola, R., Ernst, E., Samaan, D., del Rio-Chanona, M. and Teutloff, O. (2026). *The impact of GenAI on jobs, productivity and work organization: a review of the empirical evidence*. International Labour Organization, 1 June. https://doi.org/10.54394/00034628. https://www.ilo.org/publications/impact-genai-jobs-productivity-and-work-organization-review-empirical

[9] Merola, R., Ernst, E., Samaan, D., del Rio-Chanona, M. and Teutloff, O. (2026). *Workers' exposure to AI: What indicators tell us – and what they don't*. International Labour Organization, 17 April. https://doi.org/10.54394/00033279. https://www.ilo.org/publications/workers%E2%80%99-exposure-ai-what-indicators-tell-us-%E2%80%93-and-what-they-don%E2%80%99t

[10] Bank of England (2026). *How will AI affect jobs and wages?* https://www.bankofengland.co.uk/explainers/how-will-ai-affect-jobs-and-wages

[11] Klein Teeselink, B. (2025). *Generative AI and Labor Market Outcomes: Evidence from the United Kingdom*. SSRN working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5516798

[12] Schlicht, H. (2026). *Canaries in the column? AI exposure and the UK's hiring slowdown*. Bank Underground, 6 August. https://bankunderground.co.uk/2026/08/06/canaries-in-the-column-ai-exposure-and-the-uks-hiring-slowdown/

[13] Morgan Stanley Research (2026). *Mapping AI's Rate of Change: Where the Rubber Meets the Road*. 11 February, especially Exhibit 69 (selected cross-country sector survey and UK employment results). https://www.morganstanley.com/content/dam/msdotcom/en/assets/pdfs/Research_AI-Rate-of-Change.pdf

[14] Brynjolfsson, E., Chandar, B. and Chen, R. (2026). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab, August revision. https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/

[15] Indeed Hiring Lab (2026). *Indeed's 2026 Mid-Year UK Jobs and Hiring Trends Report*. 3 August. https://hiringlab.indeed.com/uk/blog/2026/08/03/mid-year-uk-jobs-hiring-trends-report/

[16] Institute for Fiscal Studies (2026). *Why has the NEET rate risen? Understanding trends and drivers using administrative data*. May. https://ifs.org.uk/publications/why-has-neet-rate-risen-understanding-trends-and-drivers-using-administrative-data

[17] Institute for Fiscal Studies (2025). *Combined impact of minimum wage and tax increases may reduce opportunities for young people*. 1 April. https://ifs.org.uk/articles/combined-impact-minimum-wage-and-tax-increases-may-reduce-opportunities-young-people

[18] Chandar, B. and Klein Teeselink, B. (2026). *How Does AI Change Labor Demand? Evidence from 41 Countries*. Stanford Digital Economy Lab, September. https://digitaleconomy.stanford.edu/publication/how-does-ai-change-labor-demand/

[19] Gallacher, G. (2026). *AI and Job Postings: From Destruction to Creation?* Indeed Hiring Lab, 8 July. https://hiringlab.indeed.com/2026/07/08/ai-and-job-postings-from-destruction-to-creation/

[20] Brosch, H., Langer, C., Lergetporer, P., Mueller, S., Pfeifer, H. and Tschöpe, N. (2026). *Task Expansion with Generative AI: Experimental Evidence on Middle-Skilled, Early-Career Workers*. Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/publication/task-expansion-with-generative-aiexperimental-evidence-on-middle-skilled-early-career-workers/

[21] Brynjolfsson, E., Li, D. and Raymond, L. R. (2025). *Generative AI at Work*. *The Quarterly Journal of Economics*, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

[22] Brynjolfsson, E., Rock, D. and Syverson, C. (2018). *The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*. NBER Working Paper 25148. https://www.nber.org/papers/w25148

[23] Acemoglu, D. and Pischke, J.-S. (1998). *Beyond Becker: Training in Imperfect Labor Markets*. NBER Working Paper 6740. https://www.nber.org/papers/w6740

[24] PwC (2026). *2026 AI Jobs Barometer*. https://www.pwc.com/gx/en/services/ai/ai-jobs-barometer.html

[25] Bank of England (2026). *Monetary Policy Report: July 2026*. https://www.bankofengland.co.uk/monetary-policy-report/2026/july-2026

[26] Burns, C. (2026). *What Remains When the Tools Disappear: Ethics, Agency, and Human Life Inside Automated Systems*. Independent white paper, January.

[27] Department for Science, Innovation and Technology (2026). *DSIT Public Engagement Survey 2025/2026*. Published 16 July. https://www.gov.uk/government/statistics/dsit-public-engagement-survey-20252026/dsit-public-engagement-survey-20252026

[28] Audoly, R., Guerin, M. and Topa, G. (2026). *Do Job Postings Show Early Labor-Market Effects of AI?* Federal Reserve Bank of New York, Liberty Street Economics, 14 May. https://libertystreeteconomics.newyorkfed.org/2026/05/do-job-postings-show-early-labor-market-effects-of-ai/

[29] Yale Budget Lab (2026). *Tracking the Impact of AI on the Labor Market*. 14 October update. https://budgetlab.yale.edu/research/tracking-impact-ai-labor-market

[30] European Central Bank (2026). *Artificial intelligence: friend or foe for hiring in Europe today?* ECB Blog, 4 March. https://www.ecb.europa.eu/press/blog/date/2026/html/ecb.blog20260304~d9e34fc95f.de.html

[31] Goldin, C. and Katz, L. F. (2008). *The Race between Education and Technology*. Harvard University Press. https://lkatz.scholars.harvard.edu/publications/race-between-education-and-technology

[32] Ogburn, W. F. (1922). *Social Change with Respect to Culture and Original Nature*. B. W. Huebsch. https://books.google.com/books/about/Social_Change_with_Respect_to_Culture_an.html?id=tTlFAAAAIAAJ

[33] Acemoglu, D. and Restrepo, P. (2019). *Automation and New Tasks: How Technology Displaces and Reinstates Labor*. NBER Working Paper 25684. https://www.nber.org/papers/w25684

[34] Rogers, E. M. (2003). *Diffusion of Innovations*, fifth edition. Free Press. https://www.simonandschuster.net/books/Diffusion-of-Innovations-5th-Edition/Everett-M-Rogers/9780743222099

[35] Bessen, J. (1997). *Productivity Adjustments and Learning-by-Doing as Human Capital*. US Census Bureau, Center for Economic Studies Working Paper CES-97-17. https://www2.census.gov/ces/wp/1997/CES-WP-97-17.pdf

[36] Callaway, B. and Sant'Anna, P. H. C. (2021). *Difference-in-Differences with Multiple Time Periods*. *Journal of Econometrics*, 225, 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

[37] Sun, L. and Abraham, S. (2021). *Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects*. *Journal of Econometrics*, 225, 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

[38] Lakens, D., Scheel, A. M. and Isager, P. M. (2018). *Equivalence Testing for Psychological Research: A Tutorial*. *Advances in Methods and Practices in Psychological Science*. https://doi.org/10.1177/2515245918770963

[39] Office for National Statistics (2026). *Business Insights and Conditions Survey (BICS) QMI*. https://www.ons.gov.uk/economy/economicoutputandproductivity/output/methodologies/businessinsightsandconditionssurveybicsqmi

[40] UK Data Service (2026). *Apply to access ONS data*. https://ukdataservice.ac.uk/find-data/access-conditions/secure-application-requirements/apply-to-access-ons-data/

[41] Department for Science, Innovation and Technology and LinkedIn (2026). *A snapshot of entry-level hiring in the UK*. 8 June. https://www.gov.uk/government/publications/entry-level-hiring-in-the-uk-a-snapshot/a-snapshot-of-entry-level-hiring-in-the-uk

[42] Howes, G. (2026). *Britain Is Still AI Job Loss Capital: Morgan Stanley Research*. Pound Sterling Live, 13 May. *Secondary Wave 2 account; primary report not independently checked.* https://www.poundsterlinglive.com/economics/23287-britain-is-ai-job-loss-capital-confirms-new-morgan-stanley-research

[43] Wenzelmann, F., Schönfeld, G., Linckh, C. and Pfeifer, H. (2026). *Apprenticeship Training or External Recruitment of Skilled Workers: What Costs Should Firms Expect? Results of the 2022/2023 BIBB Cost-Benefit Survey*. BIBB Report 2. Bonn: Bundesinstitut für Berufsbildung. https://www.bibb.de/dienst/publikationen/en/21068

[44] OECD (2018). *Seven Questions about Apprenticeships: Answers from International Experience*, chapter 3, 'What is the right wage for apprentices?'. Paris: OECD Publishing. https://www.oecd.org/en/publications/seven-questions-about-apprenticeships_9789264306486-en

[45] Muehlemann, S. (2025). Artificial intelligence adoption and workplace training. *Journal of Economic Behavior & Organization*, 238, 107206. https://doi.org/10.1016/j.jebo.2025.107206

[46] Akin, M. S. (2026). *The Apprenticeship Externality: Generative AI, Entry-Level Work, and the Future Supply of Expertise*. SSRN working paper, posted 2 June; not peer reviewed. https://doi.org/10.2139/ssrn.6809082

[47] Brynjolfsson, E., Li, J. F., Miranda, J., Seamans, R. and Wang, A. J. (2026). *Minimum Wages and Rise of the Robots*. NBER Working Paper 34895. US industrial robot evidence, not an estimate for UK generative AI. https://doi.org/10.3386/w34895

[48] MacKinnon, J. G. and Webb, M. D. (2018). The wild bootstrap for few (treated) clusters. *The Econometrics Journal*, 21(2), 114–135. https://doi.org/10.1111/ectj.12107

[49] Garicano, L. and Rayo, L. (2025). *Training in the Age of AI: A Theory of Career Viability*. CEPR Discussion Paper 20634. https://cepr.org/publications/dp20634

[50] Lovett, N. (2026). The tragedy of the cognitive commons: how AI could disrupt the regeneration of professional expertise. *Human Resource Development Review*, online 26 July. https://doi.org/10.1177/15344843261470602

[51] Lordan, G. and Neumark, D. (2018). People versus machines: the impact of minimum wages on automatable jobs. *Labour Economics*, 52, 40–53. https://doi.org/10.1016/j.labeco.2018.03.006

[52] Westby, S., Modestino, A. S. and Cheng, P. (2026). *Generative AI and the Redefinition of Entry-Level Software Work*. IZA Discussion Paper 18723. https://www.iza.org/publications/dp/18723