

# AI Adoption and Workforce Change in SMEs

**David Bharier**

British Chambers of Commerce

**Ben Etheridge**

University of Essex

**Paulo Morais**

University of Essex

No. 2026-1  
March 2026



INSTITUTE FOR SOCIAL  
& ECONOMIC RESEARCH

## **Non-technical Summary**

Artificial intelligence is spreading rapidly among UK small and medium enterprises. Using survey data from nearly 700 firms collected by the British Chambers of Commerce in early 2026, we find that over half now use AI in some form, up from around a third in 2025. Most rely on widely available tools such as ChatGPT or Copilot, but around one in ten firms have implemented bespoke AI systems tailored to their operations.

This distinction matters. Bespoke AI adoption, not AI use in general, is where workforce effects are concentrated. Roughly one in five bespoke adopters report staffing reductions directly attributable to AI, compared with just 3 per cent of firms using only generic tools. Bespoke adopters are also three times more likely to have restructured job roles, and restructuring is in turn linked to changing skills demands: not only growing requirements for AI literacy, but also for interpersonal and communication skills.

A notable finding is that firms investing in AI-related training are more likely to anticipate headcount reductions than those with no training plans. This suggests that, for some firms, training accompanies workforce reduction rather than preventing it.

These patterns have direct policy relevance. Headline AI adoption figures may overstate the extent of workforce disruption, since most use remains shallow. But as bespoke implementations spread, restructuring and skills shifts are likely to follow. Public support for retraining should not assume that employer-led training signals job preservation, and attention to geographic disparities in adoption, particularly lower rates in regions far from London and South East England, will be important for avoiding widening regional inequalities.

# AI Adoption and Workforce Change in SMEs\*

David Bharier<sup>†</sup>, Ben Etheridge<sup>‡</sup>, Paulo Morais<sup>§</sup>

## Abstract

This paper investigates Artificial Intelligence (AI) adoption and its labour market consequences among UK small and medium enterprises, using novel data from the British Chambers of Commerce Business Outlook Survey, collected in early 2026. AI adoption is increasingly widespread, with over half of responding firms currently using AI, up from around a third in 2025. Most users rely on generic tools such as ChatGPT or Copilot, but around one in ten firms have adopted bespoke AI implementations. We find that bespoke adoption in particular is associated with a coherent bundle of workforce adjustment. Approximately one-fifth of bespoke users report staffing reductions attributable to AI, and bespoke adopters are roughly three times more likely to have restructured job roles. Restructuring is in turn strongly associated with headcount reductions and shifts in skills requirements. Surprisingly, firms investing in AI-related training are significantly more likely to anticipate headcount reductions than those not investing in training. We also find that current AI users are substantially more optimistic about future productivity gains than non-users. Our findings provide a novel firm-level picture of how SMEs are reorganising work, adjusting workforces, and investing in skills in response to AI.

**JEL Classification:** J23, J24, L25, M51, O33

**Keywords:** artificial intelligence, AI adoption, small and medium enterprises, labour market, workforce adjustment, skills, job restructuring, productivity expectations

\*We thank funding from the ESRC Research Centre on Micro-Social Change (MiSoC), award number ES/S012486/1. We thank Emilia Del Bono and Paul Garcia for comments. All errors are our own.

<sup>†</sup>British Chambers of Commerce. E-mail: [d.bharier@britishchambers.org.uk](mailto:d.bharier@britishchambers.org.uk)

<sup>‡</sup>University of Essex. E-mail: [bsethe@essex.ac.uk](mailto:bsethe@essex.ac.uk)

<sup>§</sup>University of Essex. E-mail: [paulo.santosmorais@essex.ac.uk](mailto:paulo.santosmorais@essex.ac.uk)

# 1 Introduction

Artificial intelligence promises to be among the most transformative technologies of the coming decades. A large and growing literature has sought to forecast which occupations and tasks are most exposed to AI (e.g. [Webb, 2019](#); [Felten et al., 2021](#); [Eloundou et al., 2024](#)), and a number of recent studies have documented large productivity gains from AI adoption in experimental and quasi-experimental settings ([Brynjolfsson et al., 2025](#); [Noy & Zhang, 2023](#); [Peng et al., 2023](#)). Yet so far little is known about how firms are actually responding to these possibilities: how they are reorganising work, adjusting their workforces, and investing in new skills. This gap is especially acute for small and medium enterprises (SMEs), which employ the majority of workers in most advanced economies and encompass the young firms responsible for much of gross job creation and destruction ([Haltiwanger et al., 2013](#)). SMEs also experience AI adoption differently from larger firms: they are unlikely to employ dedicated AI integrators on which existing studies rely to identify firm-level adoption ([Hampole et al., 2025](#); [Hosseini Maasoum & Lichtinger, 2025](#)).

This paper provides new evidence on AI adoption and its labour market consequences among UK SMEs. It uses novel data from the British Chambers of Commerce (BCC) annual Business Outlook Survey, a long-running panel used to inform central government policy and with high participant adherence. We added a suite of AI-specific questions to the January 2026 wave of the survey, which received a little under 700 qualified responses, the vast majority of which are firms with 1–250 employees. A central feature of our instrument is the distinction between generic AI tools, specified as commercially available applications such as ChatGPT or Copilot, and bespoke AI implementations. While over half of firms in our sample report using AI in some form, only around one in ten have adopted bespoke solutions. It is these bespoke implementations that we interpret as analogous to the installations requiring AI integrators identified in larger-firm studies, and it is these implementations that are associated with meaningful employment effects.

Our contribution is twofold. First, our survey instrument distinguishes generic from bespoke AI adoption, a distinction that proves central to understanding workforce effects. Second, it covers the full pipeline from adoption through restructuring, skills shifts, and training within the same firms, providing a firm-level view that is difficult to obtain from, for example, worker-level data. We observe firms' AI adoption status, their productivity expectations, past and expected future staffing changes, whether job roles have been restructured, how skills requirements have shifted, and whether the firm plans to invest in AI-related training. Although our sample is modest in size, the findings are sufficiently stark that we are able to provide a number of robust, highly statistically significant results.

We begin by documenting adoption patterns. Over 54 per cent of firms currently use AI, with a further 17.5 per cent intending to adopt. Open-ended responses indicate that generic use is concentrated in content generation, writing, and general research, while bespoke implementations tend to involve administration, operations, data processing, and specialised design tasks. Adoption is substantially more prevalent in sectors with high ex-ante AI exposure, as measured

by [Felten et al. \(2021\)](#). There are also notable geographic disparities, with adoption lower in the regions furthest away from London and the South East: Yorkshire, the Humber and the North East, and the devolved nations of Wales, Scotland and Northern Ireland.

We next examine productivity expectations. Firms are overwhelmingly optimistic about AI's impact on productivity, consistent with recent UK survey evidence ([Yotzov et al., 2026](#); [Institute of Directors, 2025](#)). Importantly, we find that current AI users are substantially more optimistic than firms that merely intend to adopt. This experience-based pattern suggests that direct engagement with AI reinforces rather than dampens productivity optimism. This provides an important complement to studies documenting that the majority of firms report no measurable productivity gains from AI so far ([Yotzov et al., 2026](#)): Even in the absence of hard evidence of returns, the subjective experience of using AI appears to sustain or strengthen positive expectations, even compared to firms that plan to embrace AI in the future.

Our main focus is on the labour market consequences of adoption. We find no systematic effect of AI adoption on past staffing levels overall, but this aggregate masks a sharp distinction by type of implementation. Approximately one-fifth of bespoke AI users report that their workforce has decreased as a direct result of AI in the past twelve months, compared with only around 3 per cent of generic-only users. This difference is highly statistically significant and robust to a comprehensive set of controls. A key channel linking bespoke adoption to headcount reductions appears to be job restructuring: bespoke users are roughly three times more likely to have restructured or redefined job roles, and restructuring is in turn strongly associated with staffing reductions.

Perhaps our most novel results concern the interplay between skills, training, and workforce adjustment. A little under a third of AI-adopting firms report growing requirements for AI literacy, while over a sixth report increased demand for technical and analytical skills, and a similar share for interpersonal and communication skills. Importantly, these skills shifts are strongly associated with job restructuring, consistent with restructuring acting as a channel through which deeper AI integration alters skill demands. This pattern is particularly pronounced for AI literacy and interpersonal skills, and especially in AI-exposed sectors. The finding that interpersonal and communication skills are among the growing requirements speaks to complementarity between solely human capabilities and AI, consistent with evidence on the rising value of social skills ([Weinberger, 2014](#); [Deming, 2017](#); [Loaiza & Rigobon, 2024](#)).

A noteworthy manifestation of these patterns is the noticeable fraction of firms that are investing in AI-related training, but expecting to reduce headcount: a pattern we term 'replace and train'. Quantitatively, 14 per cent of firms with training plans expect decreases in headcount compared with 4 per cent of those without training plans. This suggests that, for a substantial body of firms, training accompanies workforce reduction rather than expansion, forming part of a coherent bundle in which bespoke AI adoption, job restructuring, skills shifts, and training investment are jointly associated. We note that all these results are cross-sectional and observa-

tional, and a strong causal interpretation should be taken with caution. Nevertheless they provide a novel picture that warrants future research.

The survey also provides information on the quantity of tasks performed by AI versus humans. We find that, unsurprisingly, the vast majority of AI-using firms (95 per cent) still report humans performing most tasks. However, a small but intriguing group, around 3.5 per cent of adopters, report that AI performs at least as many tasks as humans. While sample sizes preclude strong conclusions, we show robustly that these deeply integrated firms tend to be small, fast-growing enterprises that expect to expand their workforce. They therefore seem to be building their operations around AI from the outset, rather than removing workers from existing structures.

Our work contributes to a few main strands of the literature. We add to the growing body of firm-level surveys on AI adoption ([Bonney et al., 2024](#); [McElheran et al., 2024](#); [Yotzov et al., 2026](#)) by providing detailed evidence from SMEs that distinguishes generic from bespoke adoption. We contribute to the literature on AI's labour market effects ([Brynjolfsson et al., 2025](#); [Hosseini Maasoum & Lichtinger, 2025](#); [Hampole et al., 2025](#)) by documenting the specific channels, i.e. restructuring, skills shifts and training, through which deeper AI integration is associated with workforce adjustment. Finally we shift the focus from forecasting which jobs are exposed to AI toward documenting which skills face increased demand, and how firms are actively reorganising to accommodate these technologies, complementing nascent work in this area ([Loaiza & Rigobon, 2024](#); [Gathmann et al., 2024](#)).

The remainder of the paper is organised as follows. Section 2 describes the BCC Business Outlook Survey and the AI-specific questions we added to the January 2026 wave, along with the construction of our key variables and controls. Section 3 examines AI adoption patterns across the sample, including the determinants of adoption, the generic–bespoke distinction, and firms' productivity expectations. Section 4 turns to our main focus: the labour market consequences of AI adoption, covering headcount effects, job restructuring, skills shifts, training investment, and the replace-and-train pattern. Section 5 concludes. Appendix A provides further detail on the survey instrument and data processing, while Appendix B presents additional results referenced throughout the main text.

## 2 Data

We use data from the British Chambers of Commerce (BCC) annual Business Outlook Survey, administered in early January 2026 with responses received between January 13th and February 11th. BCC surveys are some of the longest-running private-sector business surveys in the UK, regularly cited by the Bank of England, HM Treasury and the Office for Budget Responsibility in

their economic assessments.<sup>1</sup> The Business Outlook Survey captures a broad cross-section of UK businesses through the BCC’s membership network.

Our sample comprises 672 qualified responses, of which 668 provide a response to the core AI adoption question. The sample is predominantly composed of small and medium enterprises: 561 firms (84%) employ between one and 250 workers. It additionally includes 68 sole traders and 39 larger enterprises with 250 or more employees. The focus on SMEs contrasts with the broader samples analysed in, for example, [Yotzov et al. \(2026\)](#), whose headline AI adoption rate of around 70% exceeds our finding of 54% discussed further below, consistent with a positive relationship between firm size and AI adoption ([Acemoglu et al., 2022](#); [McElheran et al., 2024](#)).

We report unweighted results throughout as cell sizes for key subgroups are too small for reliable weighted estimation. Our findings should therefore be interpreted as describing patterns within the responding sample rather than as population-representative estimates for UK SMEs. We note that BCC members are slightly under-represented in London and over-represented in the West Midlands and in manufacturing. In Appendix B we report main descriptive patterns for a strict SME subsample, excluding sole traders and larger firms, and applying sector and region weights to approximate population-representative proportions; patterns are not materially altered. For the regression estimates, our rich group of controls absorb the compositional imbalances that reweighting would address ([Solon et al., 2015](#)).

## AI-specific questions

In addition to the standard survey items, this wave included nine questions on AI implementation, listed in full in Appendix A. The central item (q2) asks whether the firm currently uses AI and, if so, whether it uses generic tools (such as ChatGPT or Copilot), bespoke organisational tools, or both. Firms reporting use of both are classified as bespoke users, on the logic that bespoke implementation represents a deeper level of organisational integration. This yields 298 generic-only users and 65 bespoke users among the 363 current AI adopters.<sup>2</sup> An open text field allows respondents to describe their AI use cases.

The remaining AI questions cover the extent of human vs AI involvement in the enterprises tasks (capturing especially ‘deep’ integration), past and expected future staffing changes attributable to AI, shifts in skills requirements, training investment plans, and beliefs about firm-level productivity effects of AI. Questions on staffing, task split, restructuring and on skills are routed only to current AI users; questions on training and productivity expectations are asked of all respondents.

---

<sup>1</sup>BCC survey data are used as an input to the Bank of England’s survey-indicator model for estimating underlying GDP growth (Bank of England, *Monetary Policy Report*, November 2025, Chart 1.7, available at <https://www.bankofengland.co.uk/monetary-policy-report/2026/february-2026>; see also the February 2026 Report). BCC forecasts are included in HM Treasury’s monthly *Forecasts for the UK Economy* comparison (available at <https://www.gov.uk/government/statistics/forecasts-for-the-uk-economy-february-2026>).

<sup>2</sup>We acknowledge that the distinction between generic and bespoke is not always sharp: firms that heavily customise commercial tools or build lightweight applications around foundation models may occupy an intermediate position. However, as we see when we present the results, our elicitation has meaningful empirical content.

## Background variables and controls

We draw on standard survey items covering firm size, revenue and profit margin (all reported in bands), international trade in goods and services (separately), sector, region and expectations about firm growth over the coming 12 months, which we interpret as expected revenue growth.

From these we construct a set of controls: binary indicators for positive and negative growth expectations; log firm size, log revenue per worker and operating margin (using band midpoints); and binary trade indicators. We impute missing values for revenue per worker and profit margin at the sample median, assign “no change” for missing growth expectations, and assign “no trade” for missing trade responses. In robustness exercises we drop all imputed observations, removing approximately one third of the sample. No key results are materially altered.

Because of limited coverage in some of the smaller regions, we aggregate the 12 government office regions to 8 coarser groups. Likewise we aggregate the BCC’s 19 sectors to 10 broader categories, with mappings reported in Appendix A. To interpret our stark results by sector, we additionally characterise sectors using two continuous measures. The first is the AI exposure index of Felten et al. (2021), a measure of sector-level predicted exposure to AI capabilities computed *before* the current wave of generative AI, converted to the BCC sectoral classification. The second is the *ex-post* leave-one-out sector-level adoption rate, computed from q2 responses in our sample. Details of the crosswalk and construction are in Appendix A.

## Sample description

Table 1 summarises the sample by AI adoption status. Just over half of firms (54.3%) currently use AI, with a further 17.5% intending to adopt and 22.3% reporting no plans. Adoption is strongly associated with firm outlook: 65% of firms expecting positive growth currently use AI, compared with 45% of those expecting contraction. International services trade is similarly associated with higher adoption, though goods trade is not. Among continuous characteristics, AI use is significantly associated with firm size and with sector-level predicted AI exposure, though not with revenue per worker or operating margin.

Appendix Table B.1 provides full distributional breakdowns by sector, region and detailed size and revenue bands. AI use is highest in IT & digital (78%) and creative & media (70%), and lowest in other services (31%) and manufacturing (45%). Geographically, adoption is highest in southern England and noticeably lower in the North East, Yorkshire, and the devolved nations. Appendix Table B.2 shows that the patterns reported in Table 1 are robust to restricting the sample to a strict sample of SMEs (with between 1 and 250 employees) and reweighting to match national sector and region distributions.

### 3 AI Adoption

We first assess the extent and nature of AI adoption across SMEs in the UK. Figure 1, which incorporates previous editions of the Business Outlook Survey, shows that AI adoption among BCC survey respondents has risen steadily, from 23% in 2023 to 54% in the January 2026 wave used in this paper. Figure 2 then expands on Table 1, discussed in the previous Section, by showing adoption status by firm size in full detail. It shows that sole traders and micro enterprises (1–9 employees) are far more likely to report no plans to implement AI than larger firms. While medium-size enterprises (up to 250 employees) are not noticeably more likely to report already using AI than micro enterprises, they are more likely to report intentions to adopt, indicating that they are taking more time to make plans. The overall pattern is consistent with literature that typically finds a positive correlation of AI use with firm size across the wider population of firms in the UK and elsewhere (Acemoglu et al., 2022; Bonney et al., 2024; McElheran et al., 2024; Yotzov et al., 2026).

To assess whether these patterns hold conditional on observable firm characteristics, we investigate wider determinants of AI adoption in a series of discrete choice regressions, presented in Table 2. Columns 1–4 model the probability of current AI use via binary logit. Across all four specifications, positive growth expectations are shown to be a robust predictor of current AI use. The same is true of international services trade, perhaps reflecting AI’s use in facilitating communication across languages. These results survive extensive controls for region, sector, firm size, revenue per worker and profit margin. Each is associated roughly with a doubling of the odds of current AI use.

Column 3 of Table 2 controls for sector via fixed effects, which capture substantial variation in adoption rates across industries. To move beyond these fixed effects towards a more interpretable measure, column 4 replaces sector dummies with the continuous and standardised measure of sector-level AI exposure converted from Felten et al. (2021). It shows that a one standard deviation increase in this measure is associated with a multiplicative increase in the odds ratio of 1.26, implying roughly a 6 percentage point increase in the probability of current AI adoption at the baseline rate of 54%, conditional on the other controls.

The relationship between firm size and AI adoption is made clearest in columns 5 and 6, which show a multinomial logit for current use and for intention to adopt against a baseline of no plans. These columns show a strong association with firm size for both outcomes conditional on controls. Similarly, positive expectations of firm output growth are associated with both current use and stated intentions to adopt.

A key focus of the debate on AI is uncertainty around effects on current and future productivity growth. Predictions range from at least a doubling of GDP over 10 years (e.g. Korinek & Suh, 2024) at the high end, to barely noticeable or even detrimental at the bottom end (e.g. Acemoglu, 2025). Our survey addresses this issue with a direct question on expected AI productivity impacts (q9 shown in Appendix A). In line with recent evidence from business surveys (Yotzov

et al., 2026; Institute of Directors, 2025), our data show that firms' expectations of future productivity growth are on balance positive. Figure 3 builds on this by showing a salient association of these expectations with AI adoption. While it is perhaps not surprising that those with no plans for AI adoption are less optimistic about the effects of AI, it is striking that current users are significantly more optimistic than intenders, suggesting that direct experience with AI reinforces rather than dampens productivity optimism.

We test this experience effect in greater detail with ordered logit regressions for the three main belief statements ('no change', 'some increase', 'significant increase'). Results are reported in Table 3. Columns 1 and 3–5 report results for current adopters, intenders, and those who are unsure against the omitted category of those with no plans, while column 2 groups the unsure in the omitted category. Across all specifications current AI adoption status has powerful predictive power for expectations of productivity effects. The very large coefficients reported here reflect the large swings in responses across categories.<sup>3</sup> We show a clean test of the difference between current adopters and intenders in Appendix Table B.3, which reports estimates with intenders as the base category.

Table 3 also shows a clear association of productivity expectations with firm-specific growth expectations. Additionally, beliefs are concentrated both in sectors more exposed to AI *ex ante* (column 3) and in those that have adopted AI more extensively *ex post* (column 5). Moreover, the association between adoption status and productivity optimism is amplified in more AI-exposed sectors (column 4).

We augment these findings with additional results presented in Appendix B. Appendix Table B.4 tests the robustness of several of our findings when removing firms with missing values of various covariates, and shows that these findings overwhelmingly hold up. Table B.5 complements Table B.1 by showing coefficients on regions from our main regression of AI adoption status. It shows that the devolved nations (Northern Ireland, Wales and Scotland) and Yorkshire and the North East clearly lag London and the South East, consistent with patterns of geographic diffusion discussed in Appel et al. (2025).<sup>4</sup>

Table B.6 shows transitions for the, admittedly small, set of firms we follow longitudinally. It shows high persistence in AI use: 96% of users in 2024 remained users in 2025. Net movement is strongly toward adoption, and reassuringly for the rest of the analysis, stated intention appears to be a meaningful leading indicator: 38% of intenders in 2024 had converted to users by 2025.

We close this section by introducing an important distinction between AI uses that we will explore more in Section 4: that between generic AI tools (e.g. ChatGPT, Copilot) and bespoke im-

---

<sup>3</sup>As a back-of-envelope illustration, among firms with no plans to adopt AI, only 2.3% expect AI to improve their firm's productivity compared to 97.7% expecting no change or worse, yielding raw odds of 0.024. For current users, the equivalent figures are 71.9% and 28.1%, with raw odds of 2.56, which is over 100 times larger. The conditional odds ratios in the table are somewhat smaller, reflecting the role of correlated firm characteristics, but remain large in magnitude.

<sup>4</sup>See also Bick et al. (2026) on the speed of the rollout of AI.

plementations suited to a specific organisation. Fully 10% of the sample report having adopted a bespoke AI solution. We interpret these as organisationally embedded implementations akin to those captured by [Hampole et al. \(2025\)](#) and [Hosseini Maasoum & Lichtinger \(2025\)](#), who identify installations associated with dedicated AI integrators. Given the SMEs in our sample, however, these may be managed by existing IT staff in conjunction with external vendors. Examples include computer-vision assisted quality control or specialised industrial maintenance systems. To assess these implementations more broadly, Table B.7 shows results of a simple text analysis of open-ended responses on AI use cases. Although the sample is small and we do not conduct formal analysis, generic use appears more associated with writing and general research, while bespoke implementations are more oriented toward administration and operations, graphic design, and other specialised uses.

## 4 Workforce Impacts

A salient debate on the economics of AI concerns its impact on jobs and the labour market. Using worker-level data, [Brynjolfsson et al. \(2025\)](#) and [Hosseini Maasoum & Lichtinger \(2025\)](#) provide evidence of recent negative labour market impacts on junior workers. Taking a longer historical view, [Deming et al. \(2025\)](#) provide evidence that the rate of occupational churn has also increased in recent years, following a long-term slowdown over the latter 20th century and early 21st century. In the long term, the effect of AI on the labour market likely depends on the extent to which it becomes fully agentic. In the short to medium term, however, its effect depends on the same forces shaping previous technological advances: the nature and extent of task automation by AI, the complementarity of these tasks with existing and new tasks, and the change in the nature of expertise required to supervise and use these new tools ([Autor, 2024](#); [Autor et al., 2024](#)). In this vein [Loaiza & Rigobon \(2024\)](#) and [Gathmann et al. \(2024\)](#) provide evidence effects of AI on the skill requirements and task composition of jobs.

### 4.1 Staffing and Organisational Restructuring

Our survey allows us to examine these dynamics from the firm's perspective and we begin by examining how AI use relates to staffing levels and workforce change. To this end we draw on questions asked to AI adopters covering past and expected headcount changes and internal job restructuring.

**Past headcount changes:** Figure 4 begins by showing responses to the question on AI-induced changes in staffing over the previous 12 months, split by generic vs bespoke users. While the vast majority of firms state that AI has had no effect, a noticeable finding is that firms with bespoke installations are more likely to report a negative effect, with around a fifth of these reporting a slight decrease in staffing.

We investigate this in more detail through ordered logit models for the five concrete responses,

with results in Table 4. The first column shows that bespoke installations are associated with staffing reductions even conditional on a full set of controls. The second and third columns replace sector fixed effects with the AI exposure index, which itself has little explanatory power. The fourth column then shows stronger evidence, however, that staffing cuts are prevalent in sectors with high *ex-post* adoption rates. Finally, we extrapolate to the broader sample by sensibly imputing that AI had “no bearing” on staffing levels among firms not currently using AI. Consistent with Figure 4 it shows that only bespoke status, rather than generic AI use, is significantly associated with staffing changes.

This distinction maps onto the task-based framework of [Hampole et al. \(2025\)](#), who decompose AI’s employment effect into mean task exposure, concentration of exposure, and firm-level productivity spillovers. Bespoke implementations presumably raise mean exposure across core tasks, tipping the balance toward worker substitution, compared to generic adoptions. The weaker productivity scale channel in most SMEs relative to the large firms studied by [Hampole et al. \(2025\)](#) may further explain why displacement effects are more visible in our sample.

We note that this pattern is also consistent with selection: firms already planning headcount reductions may choose bespoke implementations precisely as part of those plans. The cross-sectional, observational nature of our survey does not allow us to distinguish these accounts, though we return to this issue when discussing the broader pattern of co-occurring adjustments in the next subsection.

**Job restructuring:** Alongside staffing levels, the survey includes questions on whether firms have restructured job roles in their organization. We examine this by first including restructuring as an additional predictor of past changes to headcount (column 5 of Table 4). Restructuring is strongly associated with staffing reductions ( $OR = 0.236, p < 0.001$ ), consistent with restructuring serving as a channel from bespoke adoption to staffing reductions.

To provide more evidence on this point, we turn to determinants of restructuring itself. Accordingly, Table 5 presents ordered logit models for the extent of restructuring (none / few roles / most or all roles) among AI-using firms. Across all specifications, bespoke users show approximately three times the odds of restructuring relative to generic-only users ( $OR \approx 3.2, p < 0.01$ ). Two further predictors stand out: first, firms with a positive growth outlook are around three times more likely to restructure. Second, internationally trading services firms show particularly high restructuring rates ( $OR \approx 3.6$ ), indicating that the introduction of AI in fact goes deeper than purely facilitating communication with international clients discussed earlier. Finally, the interaction in Column 4 suggests the bespoke effect is somewhat attenuated in highly AI-exposed sectors, though this should be interpreted cautiously given the modest precision of sector-level exposure measures.

**Future headcount changes:** Turning to forward-looking beliefs, Figure 5 shows substantially more uncertainty about effects over the future 12 months than over the past ones. Whereas over

90% of firms with AI believed that it did not affect their staffing in the previous 12 months, this proportion reduces to 77% about beliefs for the future. This 13 percentage point difference is manifested in increased beliefs in all of the other categories: in job cuts, in job increases, and being unsure.

Beyond this increase in uncertainty, however, we find no reliable predictors of beliefs about future staffing changes in a particular direction, among the set of factors we have investigated so far. Appendix Table B.8 presents regression analysis mirroring that for past staffing changes; the most notable finding is some evidence of greater predicted future cuts among bespoke users in less AI-exposed sectors. We note that the forward-looking exercise is less clean than for past changes, since firms with intentions to adopt may hold systematic beliefs about future staffing but were not asked about this issue in the survey. We also note that we return to future staffing expectations below when examining training intentions, which reveal a more nuanced pattern.

## 4.2 Skills and Training

Having established that bespoke adoption is associated with restructuring, we now ask what consequences this restructuring has for skills requirements, and in turn for training intentions.

**Skills requirements:** As discussed previously the effect of AI adoption on skill requirements is a key pathway for its effect on labour market inequality more widely. Autor (2024) discusses in detail how AI's potential to democratise expert decision-making and hence attenuate inequality depends on scarcity of skills required for the tasks that AI replaces compared to the tasks that remain or the new tasks that are created. Eliciting information on skills requirements provides evidence on the areas in which these shifts are located.

Those firms currently using AI were asked (q7) to indicate any recognised shifts in skill needs across AI literacy, technical or analytical skills, and communication and interpersonal skills, with respondents permitted to select any that apply. Table 6 documents that, among AI-using firms, 30.6% report a greater need for AI literacy, while around 17–18% report shifts toward technical/analytical and interpersonal/communication skills respectively.

To provide additional evidence on how these skills shifts are related to AI, Table 6 also documents predictors of these skill shifts via logit regressions across three panels. We begin by noting that bespoke status, included in every specification, is not itself a significant predictor of any of the three skill shifts. The point estimates are positive throughout, and the lack of significance may reflect limited power given the small number of bespoke adopters. However, given the close relationship between bespoke adoption and restructuring established in Table 5, the results for restructuring in column 4 of each panel are revealing: it is strongly and significantly associated with both AI literacy shifts (OR = 5.163,  $p < 0.001$ ) and communication/interpersonal skills shifts (OR = 4.768,  $p < 0.001$ ), though not with technical or analytical skills. This pattern is consistent with restructuring serving as the active channel through which bespoke AI adoption

shapes requirements for certain skills.

Turning to sector-level patterns, Panel A shows strong evidence of a shift in AI literacy skills for firms in AI-exposed sectors (column 2), with weaker evidence of a shift in high *ex-post* adoption sectors (column 3). Panel C further shows that firms in high adoption sectors report a noticeably higher importance of interpersonal and communication skills. Overall, this increased need for interpersonal skills is consistent with evidence from [Loaiza & Rigobon \(2024\)](#), who emphasize the effect of AI on the growth in demand for such skills including empathy and opinion, as well as with a more established literature documenting the complementarity between recent technological change and social skills ([Weinberger, 2014](#); [Deming, 2017](#)). Despite evidence that AI technologies increase the need for high-level routine skills, specifically process monitoring ([Gathmann et al., 2024](#)), we find no significant differences across type of installation, restructuring level or sector for technical or analytical skills (Panel B). This may reflect differences in skill requirements for the current wave of generative AI technologies compared to those studied by [Gathmann et al.](#), or simply a lack of power with our small sample.

**Training intentions:** These shifts in skills requirements suggest the need for training, which is likely to form an important channel of labour market adjustment ([Hyman et al., 2025](#)). Across the sample of all firms, around 38% of firms state intentions to conduct some or significant AI-related training investment. As Table 7 Panel A shows, this again varies strongly across AI adoption status. The second column of this Table shows that there seem to be slightly stronger intentions among bespoke users, though the difference from generic users is not significant at conventional levels. Column 3 introduces restructuring alongside AI user status. Restructuring is highly significant (OR = 3.827,  $p < 0.001$ ), indicating it mediates at least part of the AI adoption effect on training. Although not shown here, it is also worth noting that training plans among those intending to adopt are very similar to those currently using AI.

**Implications for future headcount - replace and train:** We next examine the relationship between training plans and expectations of future staffing levels, building on the analysis discussed above. *A priori*, firms planning training investment might be expected to also plan workforce expansion. Equally, firms adopting AI may plan to save labour by reducing headcount while upskilling remaining employees — a ‘replace-and-train’ rather than ‘train-and-grow’ strategy. Figure 6 documents expected staffing changes by training intention status among AI users. It shows that firms with training intentions have more varied staffing expectations and higher reported uncertainty; Figure 6 also shows that the change in staffing expectations seems to be shifted, if anything, towards *reduced* employment, in line with the replace-and-train pattern. The Figure shows that a two-sided small-sample test of the ratios of expecting a decrease is significant at

conventional levels ( $p$ -value $<0.01$ ).<sup>5</sup>

Table 7 Panel B tests this more formally. Column 1 confirms that, conditional on a full set of controls, training intentions are significantly associated with expected headcount reductions (odds ratio of 3.531). Column 2 shows this holds when controlling for bespoke installation status, which itself strongly predicts expected reductions (odds ratio of 4.890). Column 3 adds job role restructuring, which is also highly significant and strong, and whose inclusion raises the pseudo- $R^2$  from 0.187 to 0.268. Training retains independent predictive power even after conditioning on restructuring ( $p = 0.051$ ), suggesting that the replace-and-train pattern is not entirely subsumed by organisational change. The picture that emerges is one in which AI adoption, and bespoke AI adoption in particular, is associated with a coherent bundle of firm-level adjustment. This often involves headcount reduction, role restructuring, shifts in skill requirement, and an accompanying need for training investment.

### 4.3 Early Evidence on Deep Integration

Given the patterns documented above, a natural question is how far AI integration extends within firms. With that in mind the BCC survey contains a question (q3) asking AI adopters the split between humans and AI in their organisation's work. We find that, for a small set of firms, a different set of workforce impacts seems to be at play.

We document first the extent of any majority-AI activities. Table 8 Panel A shows that, unsurprisingly, the vast majority of firms state that the bulk of work is human-performed. However, roughly 3.5% state either equal AI-human or majority AI activity, which we call 'deeper integration'. Panel B then provides a cross-tabulation of this deeper integration with bespoke status. Although this type of integration seems more common in bespoke installations, the difference in proportions is not here statistically significant.

Panel C, however, provides a further profile of these firms, yielding stronger evidence. This panel shows robust evidence that they are comparatively small (first row), expect to increase headcount in the next 12 months (row five), and expect to expand output (row six). There is weaker evidence that they are in more AI-exposed industries (second row,  $p$ -value on difference of 0.12). Furthermore, none of these firms reports having reduced headcount in the previous 12 months. This is a striking absence that is consistent with these being expanding firms building their operations around AI rather than using it to shed existing workers. The zero cell precludes a standard test of proportions, however. Nevertheless, this evidence contrasts with that from the previous subsections, showing that, within this narrow group of SMEs, this deep level of AI use is associated with expectations of staffing *increases* rather than reductions.

---

<sup>5</sup>It should be noted that training intentions are also generally associated with more varied patterns of future headcount. We do not emphasize the slightly higher fraction of firms expecting headcount growth when training because, unlike the fraction of firms expecting reductions, it is not significantly different from the proportion for firms not undertaking training. Overall, and although not shown in Table B.8, when included in an ordered logit regression of expected future headcount changes, training is not significant. Nevertheless, we emphasize the replace-and-train pattern we observe as noteworthy, even if it is only part of a heterogeneous pattern of adjustment to adoption of AI.

Overall, our takeaway is that these firms appear to be a small but non-negligible body of fast-growing firms building their operations around AI from an initially small scale, rather than shedding workers from existing structures. Given the nature of our survey it is difficult to tell whether the proportion of 3.5% is representative of the full population of SMEs, but this suggests an important avenue for future research.

## 5 Conclusions

This paper investigates AI adoption and its labour market consequences among UK SMEs, using novel data from the BCC Business Outlook Survey. Our central contribution is to distinguish between generic AI tools and bespoke organisational AI implementations, and to show that this distinction matters substantially for understanding how AI is reshaping work. AI diffusion is increasingly widespread, with over half of firms in our sample currently using some form of AI. It is strongly associated with firm size, positive growth expectations, international services activity, and sectors such as IT and professional services. Yet adoption currently remains limited in depth. Only around one in ten firms have adopted bespoke solutions, but it is these deeper implementations, not adoption of ChatGPT or Copilot, that are associated with meaningful employment effects. Bespoke adoption is strongly associated with job restructuring, which in turn predicts shifts in skills demands and training investment alongside staffing reductions. At the same time, these changes are associated with shifts in skills demands that extend beyond AI literacy: interpersonal and communication skills are among the growing requirements, consistent with complementarity between solely human capabilities and AI.

These findings carry implications for policy. The concentration of employment effects in bespoke adoption suggests policymakers should pay attention more to this metric than to broader metrics of use. If and when these deeper implementations diffuse throughout the economy, the restructuring and skills shifts we document are likely to become more widespread. The centrality of organisational redesign to our findings also carries implications for firms built around AI from inception, which may bypass the restructuring costs that incumbents face. Policymakers concerned with managing the labour market transition should also attend to the growing demand for AI literacy and interpersonal skills, which we presume cuts across sectors. The geographic disparities in adoption also warrant attention: if firms in the devolved nations and northern England continue to lag, the productivity and skills benefits of AI may accrue unevenly, reinforcing existing regional inequalities. The replace-and-train pattern we document suggests that public support for retraining should not assume that training and job preservation go hand in hand; some training investment will be directed at restructuring rather than expanding the workforce.

Several important caveats apply. Our data are cross-sectional, so the associations we document should not be taken as causal. Firms that adopt bespoke AI may differ systematically from those that do not in ways we cannot fully control for: many attributes, including optimism, innovation orientation, firm size, sector, and AI adoption itself, are bundled together and difficult

to disentangle. Furthermore, the sample, while sufficient to detect the large effects we report, is modest in size, and the number of bespoke adopters is small enough that some of our findings, particularly on deep integration with majority AI-led work, should be treated as indicative rather than definitive.

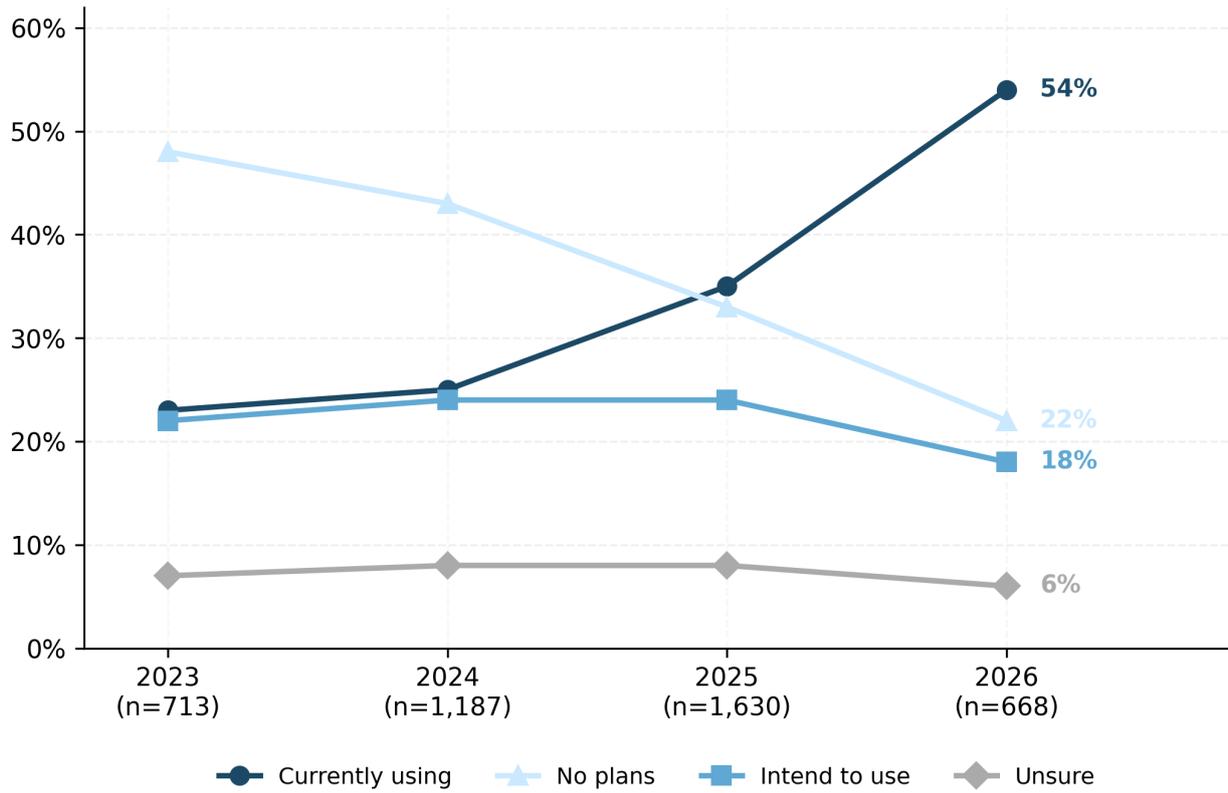
Our work shows the strength of business surveys in identifying patterns of labour market adjustment to AI. Resources and future research should be put into using these to track how adoption, restructuring, and workforce adjustment evolve over time, and to provide causal evidence. Linking larger firm surveys to administrative data such as Companies House and the Business Structure Database would permit direct measurement of employment changes rather than reliance on self-report, and would allow investigation of whether bespoke AI adoption is associated with changes in firm performance and survival. As the technology matures and bespoke adoption diffuses further, the patterns we document here in terms of the centrality of organisational redesign, the breadth of skills shifts, and the relationship between training and staffing, are likely to become increasingly important features of the labour market.

## References

- Acemoglu, D. (2025). The simple macroeconomics of ai. *Economic Policy* 40(121), 13–58.
- Acemoglu, D., G. Anderson, D. N. Beede, C. Buffington, E. E. Childress, E. M. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, et al. (2022). *Automation and the workforce: A firm-level view from the 2019 annual business survey*, Volume 30659. National Bureau of Economic Research.
- Appel, R., P. McCrory, A. Tamkin, M. McCain, T. Neylon, & M. Stern (2025). Anthropic economic index report: Uneven geographic and enterprise ai adoption. *arXiv preprint arXiv:2511.15080*.
- Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.
- Autor, D., C. Chin, A. Salomons, & B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics* 139(3), 1399–1465.
- Bick, A., A. Blandin, & D. J. Deming (2026). The rapid adoption of generative ai. *Management Science*.
- Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, & K. Savage (2024). Tracking firm use of ai in real time: A snapshot from the business trends and outlook survey. Technical report, National Bureau of Economic Research.
- Brynjolfsson, E., B. Chandar, & R. Chen (2025). Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence.
- Brynjolfsson, E., D. Li, & L. Raymond (2025). Generative ai at work. *The Quarterly Journal of Economics* 140(2), 889–942.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *The quarterly journal of economics* 132(4), 1593–1640.
- Deming, D. J., C. Ong, & L. H. Summers (2025). Technological disruption in the labor market. Technical report, National Bureau of Economic Research.
- Eloundou, T., S. Manning, P. Mishkin, & D. Rock (2024). Gpts are gpts: Labor market impact potential of llms. *Science* 384(6702), 1306–1308.
- Felten, E., M. Raj, & R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic management journal* 42(12), 2195–2217.
- Gathmann, C., F. Grimm, & E. Winkler (2024). Ai, task changes in jobs, and worker reallocation. Technical report, CESifo Working Paper.
- Haltiwanger, J., R. S. Jarmin, & J. Miranda (2013). Who creates jobs? small versus large versus young. *Review of Economics and Statistics* 95(2), 347–361.
- Hampole, M., D. Papanikolaou, L. D. Schmidt, & B. Seegmiller (2025). *Artificial intelligence and the labor market*. Number w33509. National Bureau of Economic Research.
- Hosseini Maasoum, S. M. & G. Lichtinger (2025). Generative ai as seniority-biased technological change: Evidence from us résumé and job posting data.
- Hyman, B. G., B. Lahey, K. Ni, & L. Pilossoph (2025). How retrainable are ai-exposed workers? Technical report, National Bureau of Economic Research.
- Institute of Directors (2025). Major blocker to ai adoption in british business. Technical report, Institute of Directors.

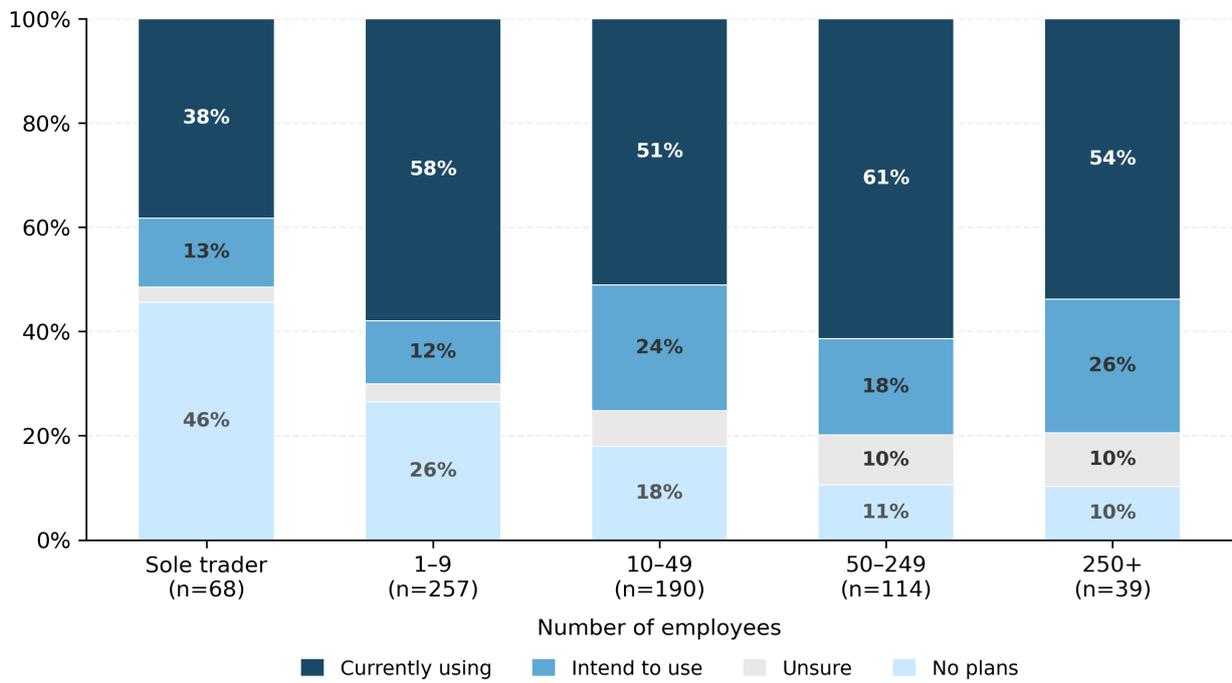
- Korinek, A. & D. Suh (2024). Scenarios for the transition to agi. Technical report, National Bureau of Economic Research.
- Loaiza, I. & R. Rigobon (2024). The epoch of ai: Human-machine complementarities at work.
- McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, & N. Zolas (2024). Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy* 33(2), 375–415.
- Noy, S. & W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654), 187–192.
- Peng, S., E. Kalliamvakou, P. Cihon, & M. Demirer (2023). The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.
- Solon, G., S. J. Haider, & J. M. Wooldridge (2015). What are we weighting for? *Journal of Human resources* 50(2), 301–316.
- Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.
- Weinberger, C. J. (2014). The increasing complementarity between cognitive and social skills. *Review of Economics and Statistics* 96(5), 849–861.
- Yotzov, I., J. M. Barrero, N. Bloom, P. Bunn, S. J. Davis, K. M. Foster, A. Jalca, B. H. Meyer, P. Mizzen, M. A. Navarrete, P. Smietanka, G. Thwaites, & B. Z. Wang (2026, February). Firm data on ai. Working Paper 34836, National Bureau of Economic Research. doi:10.3386/w34836.

**Figure 1:** AI adoption status over time (2023–2026)



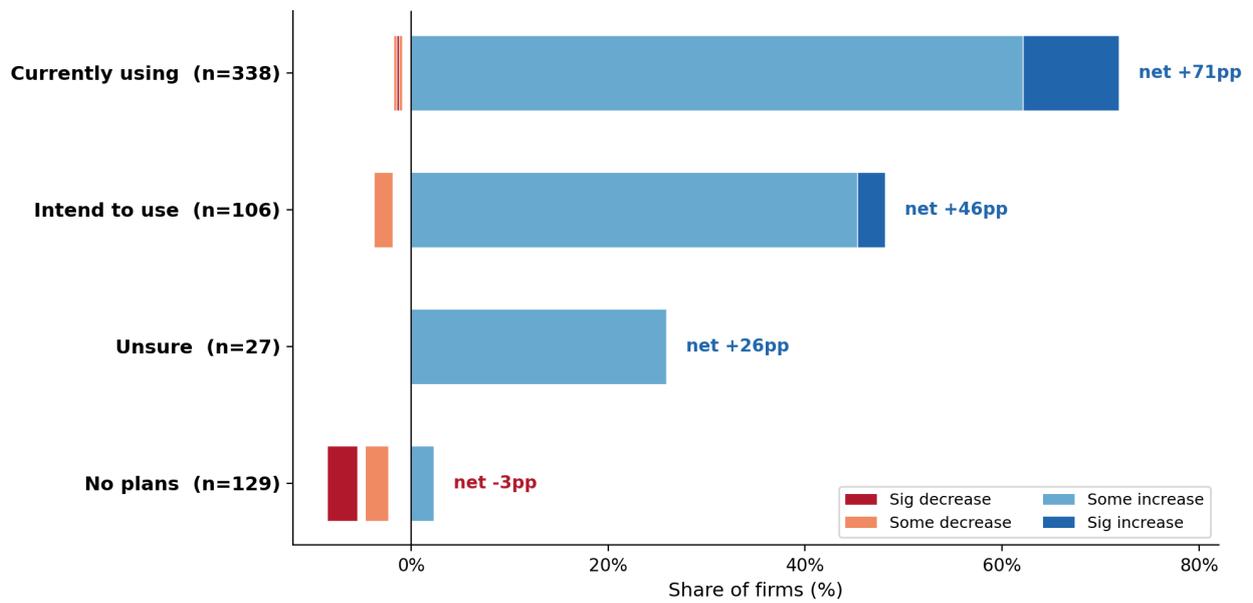
*Note:* Data for 2023–2025 are drawn from earlier waves of the BCC Business Outlook Survey. Sample composition may differ across waves; trends should be interpreted as indicative. The 2026 wave (N=668) is the basis for all other analysis in this paper.

**Figure 2: AI Adoption Across UK SMEs**



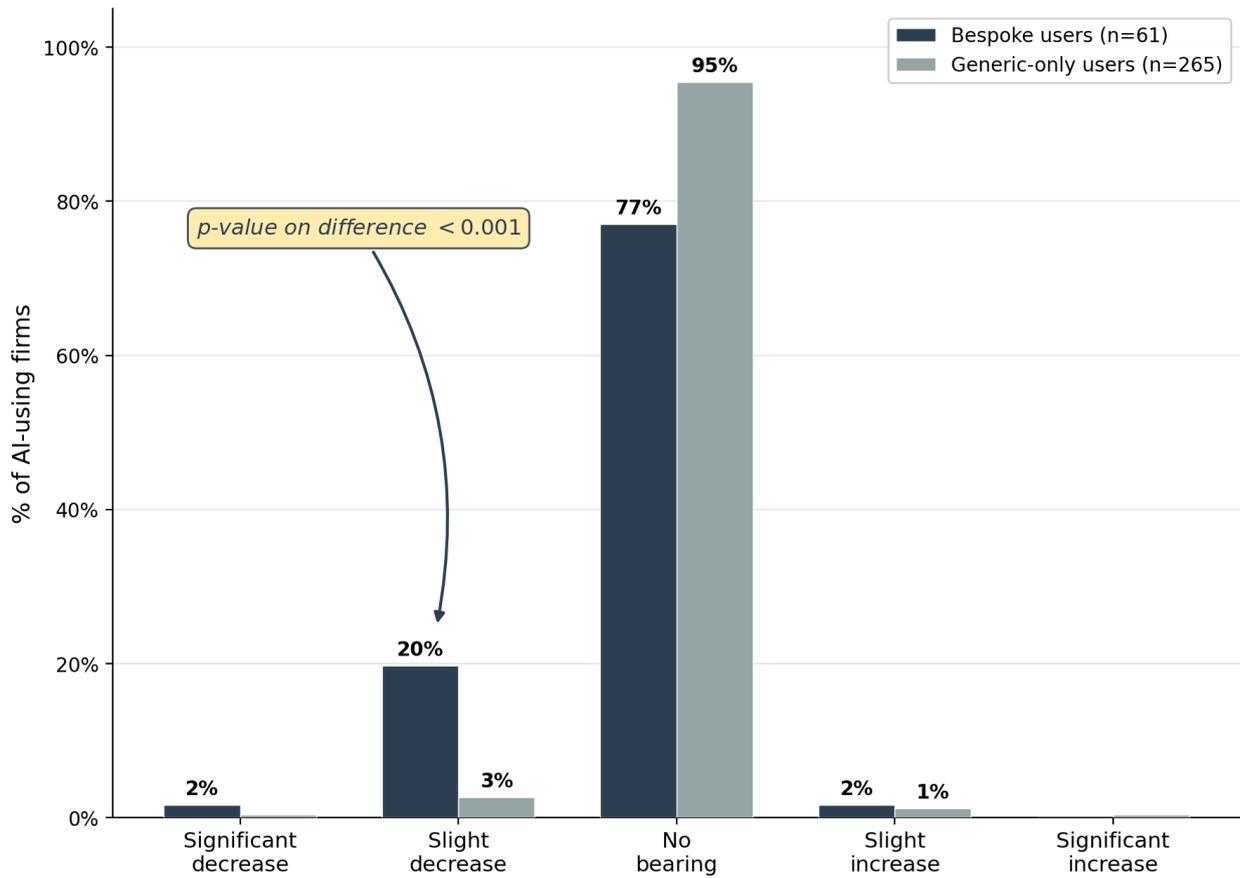
*Note:* From BCC Business Outlook Survey, January 2026. 4 firms with missing adoption status excluded.

**Figure 3: Experience of AI Breeds Optimism:**  
 Expected AI productivity impact over the next 12 months



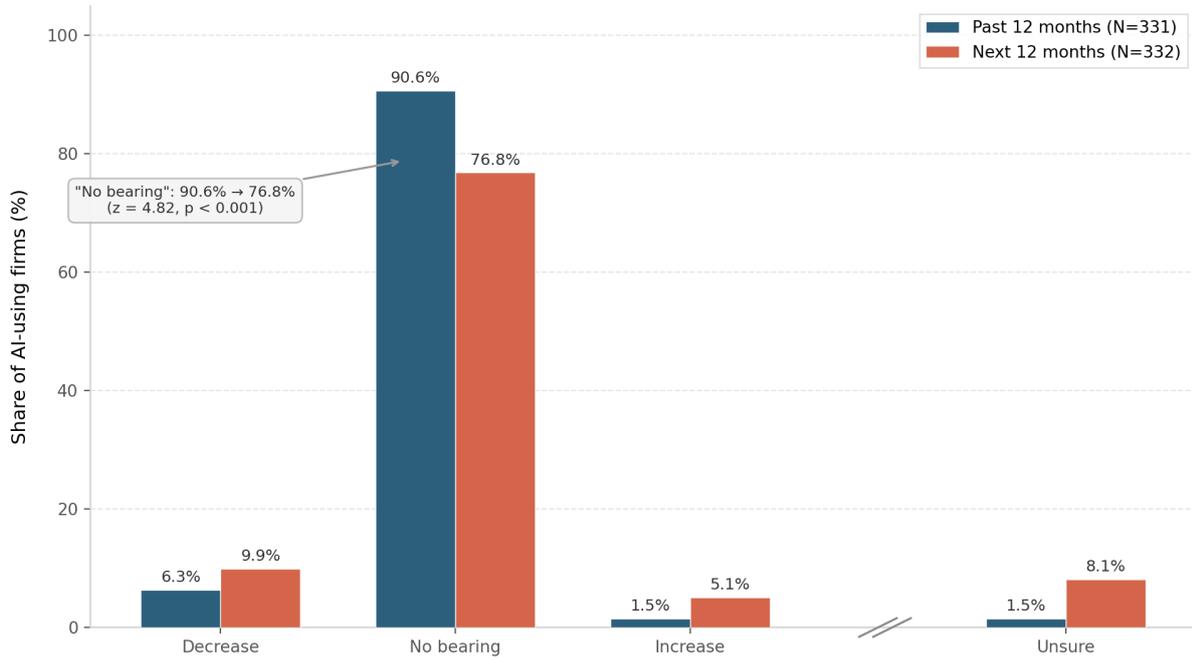
*Note:* 'No change' responses omitted (92% of firms with no plans to adopt AI expect no productivity change). Difference between current users and intenders: Mann-Whitney U test,  $p < 0.001$ . Source: BCC Business Outlook Survey, January 2026 (n=600, with responses of 'unsure' to productivity qn omitted).

**Figure 4:** Impact of AI on headcount in the past 12 months: bespoke vs generic-only AI users



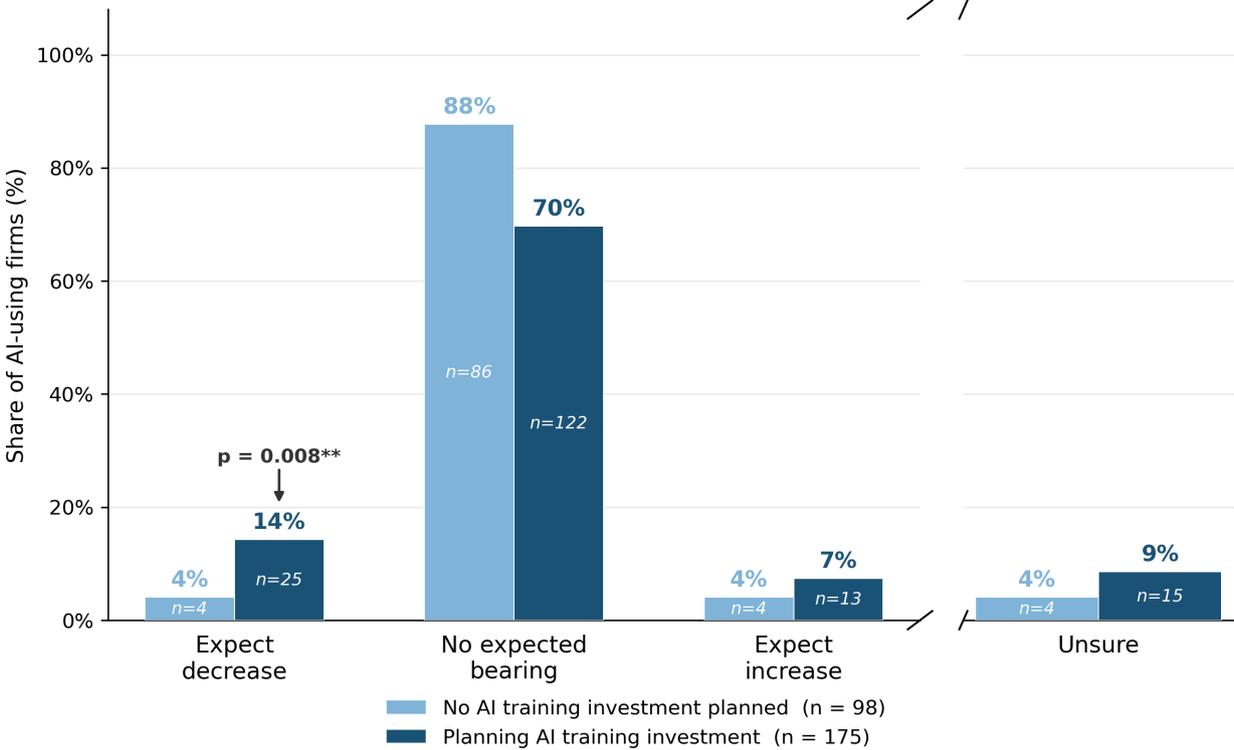
*Note:* Data from BCC Business Outlook Survey, January 2026. AI-using firms only (excludes “unsure” responses).  $p$ -value from Fisher’s exact test on the proportion reporting any headcount decrease (significant or slight) across the two groups.

**Figure 5:** Growing Uncertainty in Future Staffing Decisions: AI-related headcount changes, past experience vs. future expectations



*Note:* Data from BCC Business Outlook Survey, January 2026. Sample restricted to AI-using firms. “Decrease” and “Increase” combine slight and significant changes. Statistical significance assessed via two-proportion z-test on the “No bearing” category.

**Figure 6: “Replace-and-Train”:**  
 Firms investing in AI training more likely to anticipate headcount reductions



Notes: Among AI-using firms excluding q8 'Unsure' and q5 non-response (N = 273). \*\* Fisher's exact test, two-sided. The difference in expected headcount increases is not statistically significant (7.4% vs. 4.1%; Fisher's exact  $p = 0.311$ ).

**Table 1:** Sample Summary by AI Adoption Status

	Currently using	Intend to use	No plans	Unsure	<i>p</i> -value
Unconditional ( <i>N</i> = 668)	54.3%	17.5%	22.3%	5.8%	
<i>Conditional adoption splits</i>					
Positive growth outlook ( <i>N</i> = 301)	65.4%	18.6%	11.0%	5.0%	<0.001
Negative growth outlook ( <i>N</i> = 104)	45.2%	12.5%	38.5%	3.8%	<0.001
Trades goods internationally ( <i>N</i> = 248)	52.0%	17.3%	22.6%	8.1%	0.288
Trades services internationally ( <i>N</i> = 206)	66.0%	16.0%	13.6%	4.4%	<0.001
<i>Group means</i>					
Mean ln(firm size)	2.895	3.194	2.081	3.556	<0.001
Mean ln(revenue per worker)	4.180	4.243	4.197	4.291	0.935
Mean operating margin	11.642	10.974	12.772	11.513	0.368
Mean AI-exposure (Felten et al.)	-0.109	-0.313	-0.338	-0.561	0.003
<i>N</i>	363	117	149	39	

*Notes:* Conditional adoption splits show, for firms with each characteristic, what share fall into each adoption group (rows sum to 100%). *p*-values: Pearson  $\chi^2$  for binary rows; one-way ANOVA *F*-test for continuous rows. AI-exposure standardised (z-scores). Full distributional breakdowns in Table B.1. Weighted version shown in Table B.2.

**Table 2:** Determinants of AI Adoption

	Binary logit				Multinomial logit	
	(1) OR	(2) OR	(3) OR	(4) OR	(5) Intend RRR	(6) Using RRR
ln(Firm size)	1.051 (0.285)	1.071 (0.163)	1.138** (0.020)	1.090* (0.092)	1.552*** (<0.001)	1.436*** (<0.001)
Positive growth outlook	2.276*** (<0.001)	2.201*** (<0.001)	2.306*** (<0.001)	2.240*** (<0.001)	2.637*** (0.001)	3.757*** (<0.001)
Negative growth outlook	1.018 (0.939)	1.044 (0.859)	1.189 (0.499)	1.067 (0.792)	0.560 (0.138)	0.804 (0.450)
Trades goods internationally		0.708* (0.061)	1.070 (0.789)	0.860 (0.452)	0.574* (0.071)	0.678 (0.132)
Trades services internationally		2.016*** (<0.001)	1.811*** (0.003)	1.905*** (<0.001)	1.725* (0.080)	2.492*** (<0.001)
AI-exposure				1.257** (0.021)	1.060 (0.718)	1.298** (0.048)
Firm controls	No	Yes	Yes	Yes	Yes	Yes
Sector FE	No	No	Yes	No	No	No
Region FE	No	No	Yes	Yes	Yes	Yes
<i>N</i>	668	668	668	668	629	
Pseudo <i>R</i> <sup>2</sup>	0.0311	0.0500	0.1153	0.0748	0.1076	
Log-likelihood	-446.2	-437.5	-407.4	-426.0	-545.2	

Notes: Cols 1–4: binary logit (using AI = 1, else = 0); odds ratios reported. Cols 5–6: multinomial logit (base = No plans; Unsure dropped); relative risk ratios reported. Firm controls: ln(revenue per worker), operating margin mid-point (median-imputed where missing). AI-exposure is standardised (z-score). Robust standard errors throughout. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 3:** Productivity Expectations (Ordered Logit)

	(1) OR	(2) OR	(3) OR	(4) OR	(5) OR
Unsure (Base = No plans)	5.589*** (0.002)		6.271*** (<0.001)	10.576*** (<0.001)	5.734*** (0.002)
Intend to use	14.236*** (<0.001)		14.983*** (<0.001)	20.647*** (<0.001)	13.440*** (<0.001)
Currently using	34.048*** (<0.001)		38.096*** (<0.001)	49.467*** (<0.001)	32.695*** (<0.001)
Intender (Base = No plans + Unsure)		8.487*** (<0.001)			
User		20.250*** (<0.001)			
Positive growth outlook	2.050*** (<0.001)	2.113*** (<0.001)	1.966*** (<0.001)	1.938*** (0.001)	1.991*** (<0.001)
Negative growth outlook	0.696 (0.204)	0.667 (0.154)	0.693 (0.192)	0.651 (0.131)	0.694 (0.193)
AI-exposure			1.562*** (<0.001)	0.740 (0.271)	
Sector adoption rate					1.410*** (<0.001)
Unsure × AI-exposure				4.024* (0.057)	
Intend × AI-exposure				2.632*** (0.005)	
Using × AI-exposure				2.210*** (0.007)	
Firm controls	Yes	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	No	No	No
Region FE	Yes	Yes	Yes	Yes	Yes
<i>N</i>	602	602	602	602	602
Pseudo <i>R</i> <sup>2</sup>	0.2354	0.2275	0.2284	0.2363	0.2231
Log-likelihood	-454.1	-458.8	-458.2	-453.5	-461.4

Notes: DV: expected productivity impact (5-point ordered scale). Odds ratios reported. Base category for adoption status: No plans (cols 1, 3–5); No plans + Unsure (col 2). Firm controls (suppressed): ln(size), ln(revenue per worker), operating margin, trades goods, trades services. AI-exposure and sector adoption rate are standardised (z-scores). Joint test of interactions (col 4):  $\chi^2(3) = 18.58, p < 0.001$ . \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 4:** Determinants of Past Headcount Effects (Ordered Logit)

	AI users					Full sample
	(1) OR	(2) OR	(3) OR	(4) OR	(5) OR	(6) OR
Bespoke AI user	0.155*** (<0.001)	0.157*** (<0.001)	0.130*** (<0.001)	0.157*** (<0.001)	0.159*** (0.001)	
AI-exposure		0.727 (0.202)	0.541* (0.053)			
Bespoke × AI-exposure			2.207 (0.110)			
Sector adoption rate				0.445*** (0.003)		
Restructuring (ordered)					0.236*** (<0.001)	
AI user						0.614 (0.425)
Bespoke AI user						0.102*** (<0.001)
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Sector FE	Yes	No	No	No	Yes	Yes
Region FE	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	326	326	326	326	322	631
Pseudo $R^2$	0.2004	0.1697	0.1814	0.2054	0.2663	0.2299
Log-likelihood	-90.0	-93.4	-92.1	-89.4	-80.3	-100.3

Notes: DV: past AI headcount effect (5-point ordered scale). Unsure omitted (n=5). Odds ratios reported. Cols 1–5: AI users only. Col 5 adds restructuring (ordered) as mediator to col 1 specification. Col 6: full sample (non-users coded as “no bearing”); bespoke set to 0 for non-users. Firm controls (suppressed): ln(size), ln(revenue per worker), operating margin, growth outlook, trades goods, trades services. AI-exposure and sector adoption rate are standardised (z-scores). Robust standard errors. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 5:** Determinants of Job Role Restructuring (Ordered Logit)

	AI users			
	(1) OR	(2) OR	(3) OR	(4) OR
Bespoke AI user	3.226*** (0.002)	3.093*** (0.001)	3.102*** (0.001)	3.160*** (0.001)
AI-exposure		1.220 (0.272)		1.441* (0.073)
Sector adoption rate			1.235 (0.259)	
Bespoke × AI-exposure				0.512* (0.079)
Positive growth outlook	2.859** (0.011)	2.593** (0.016)	2.549** (0.017)	2.751** (0.012)
Negative growth outlook	2.437 (0.148)	2.006 (0.239)	2.001 (0.239)	1.955 (0.262)
Trades services internationally	3.292*** (<0.001)	3.591*** (<0.001)	3.602*** (<0.001)	3.675*** (<0.001)
Firm controls	Yes	Yes	Yes	Yes
Sector FE	Yes	No	No	No
Region FE	Yes	Yes	Yes	Yes
<i>N</i>	326	326	326	326
Pseudo <i>R</i> <sup>2</sup>	0.1571	0.1259	0.1262	0.1347
Log-likelihood	-152.1	-157.7	-157.7	-156.1

*Notes:* DV: restructuring level (0 = None, 1 = Few roles, 2 = Most/all roles). Unsure responses dropped. AI users only. Odds ratios reported. Firm controls (suppressed except growth outlook and trades services intl): ln(size), ln(revenue per worker), operating margin. AI-exposure and Sector adoption rate are standardised (z-scores). Robust standard errors. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 6:** Determinants of Skills Shifts (Binary Logit)

	(1) OR	(2) OR	(3) OR	(4) OR
<i>Panel A: AI literacy</i>				
Sample mean: 30.6%				
Bespoke AI user	1.347 (0.353)	1.304 (0.391)	1.304 (0.389)	1.000 (1.000)
AI-exposure		1.514*** (0.003)		
Sector adoption rate			1.307* (0.084)	
Any restructuring				5.163*** ( $<0.001$ )
N = 337				
<i>Panel B: Technical/analytical</i>				
Sample mean: 16.9%				
Bespoke AI user	1.428 (0.327)	1.477 (0.266)	1.452 (0.292)	1.285 (0.509)
AI-exposure		1.257 (0.170)		
Sector adoption rate			1.331 (0.143)	
Any restructuring				1.699 (0.178)
N = 337				
<i>Panel C: Communication/interpersonal</i>				
Sample mean: 18.4%				
Bespoke AI user	1.036 (0.924)	1.025 (0.946)	0.971 (0.937)	0.692 (0.352)
AI-exposure		1.102 (0.544)		
Sector adoption rate			1.490** (0.019)	
Any restructuring				4.768*** ( $<0.001$ )
N = 337				
Firm controls	Yes	Yes	Yes	Yes
Sector FE	Yes	No	No	Yes
Region FE	Yes	Yes	Yes	Yes

Notes: Sample: AI users only. DV = 1 if firm reports the given skill shift. Odds ratios reported. Any restructuring is a binary. Firm controls (suppressed): ln(size), ln(revenue per worker), operating margin, growth outlook, trades goods, trades services. AI-exposure and Sector adoption rate are standardised (z-scores). Robust standard errors. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 7:** Training Investment and the Replace-and-Train Pattern*Panel A: Determinants of training investment (ordered logit, DV: training level)*

	(1) OR	(2) OR	(3) OR
AI user	2.359*** ( $<0.001$ )		1.968*** ( $<0.001$ )
Generic AI only		2.241*** ( $<0.001$ )	
Bespoke AI user		3.099*** ( $<0.001$ )	
Any restructuring			3.827*** ( $<0.001$ )
Firm controls	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
<i>N</i>	656	656	626
Pseudo $R^2$	0.0652	0.0662	0.0788
Log-likelihood	-692.7	-692.0	-647.3

*Panel B: Replace-and-train pattern (binary logit, DV: expects any headcount decrease)*

	(1) OR	(2) OR	(3) OR
Any training investment	3.531** (0.012)	3.392** (0.016)	2.878* (0.051)
Bespoke AI user		4.890*** ( $<0.001$ )	4.146*** (0.006)
Any restructuring			5.728*** (0.006)
Firm controls	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
<i>N</i>	305	305	303
Pseudo $R^2$	0.1381	0.1870	0.2681
Log-likelihood	-90.1	-85.0	-74.8

Notes: Panel A: full sample; DV is training investment (4-point ordered scale). Col 3 adds restructuring (binary) as mediator. Panel B: AI users only; DV = 1 if firm expects any headcount decrease. Col 3 adds restructuring (binary) as mediator. Firm controls (suppressed):  $\ln(\text{size})$ ,  $\ln(\text{revenue per worker})$ , operating margin, growth outlook, trades goods, trades services. Robust standard errors. Wald test (Panel A col 2,  $H_0$ : generic = bespoke):  $z = -0.965$ ,  $p = 0.335$ . \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table 8:** The AI–Human Task Split*Panel A: Overall task-split distribution*

Category	<i>N</i>	%
Humans mostly	336	94.9%
Roughly equal	8	2.3%
AI mostly	4	1.1%
Unsure	6	1.7%
Total	354	100.0%

*Panel B: Task split by bespoke status*

	Human-led	Deeper <sup>†</sup>	% Deeper	<i>N</i>
Generic only	276	8	2.8%	284
Bespoke	60	4	6.2%	64

*Panel C: Profile of deeper-integration firms*

	Deeper ( <i>N</i> = 12)	Human-led ( <i>N</i> = 336)	<i>p</i> -value
Mean ln(size)	2.008	2.917	0.041
Mean AI-exposure	0.338	-0.134	0.126
% Bespoke	33.3%	17.9%	0.245
% past headcount decrease	0.0%	6.2%	N/A
% expect future HC increase	41.7%	3.6%	<0.001
% positive growth outlook	91.7%	53.6%	0.014

*Notes:* Based on survey question: “How would you characterise the current split between AI and human effort in your organisation?”. <sup>†</sup>Deeper integration = Roughly equal + AI mostly autonomous. Panel B: Fisher’s exact: OR = 2.30, *p* = 0.245. Panel C: *p*-values from Mann–Whitney *U* tests (continuous) and Fisher’s exact tests (binary). “% past headcount decrease”: N/A indicates test not computable (zero cells). *N* = 12 for deeper-integration firms limits statistical power.

## A Further Details on Data

### Appendix: AI Survey Items

Amid the standard battery of business outlook questions covering growth expectations, investment intentions, and trading conditions, firms were asked to respond to nine questions specifically relating to their use and experience of artificial intelligence.<sup>6</sup> These items, reproduced in full below, were designed to capture the breadth of the AI adoption pipeline, from current usage and task integration through to labour market effects, skills shifts, training investment, and productivity expectations. Questions 1, 2, 6, 8 and 9 were posed to all respondents. Questions 3, 4, 5 and 7 were routed only to firms reporting current AI (i.e. those selecting either generic or bespoke tools in Question 2). Although Question 6 (restructuring) was administered to the full sample, the majority of non-adopters left it blank; in the analysis we therefore restrict this item to current AI users.

Response options for each item are listed exactly as presented to respondents. Where questions permitted multiple responses, this is noted. Throughout the paper, we refer to these items as q1–q9.

**q1: Over the next 12 months, which phase do you expect your organisation to be in?**

*[All respondents; single response; N = 667]*

- Rapid growth trajectory
- Intend to grow
- Business as usual
- Intend to downsize
- Rapid downsizing trajectory
- Ceasing operations

**q2: Is your organisation currently using specific AI technology?**

*[All respondents; multiple responses permitted; N = 668]*

- Yes, we use generic tools e.g. ChatGPT/Copilot (*and please specify the main use*)
- Yes, we use bespoke tools built for our organisation (*and please specify the main use*)
- Not currently using but intend to use
- Not sure if we are
- No plans at all to use AI

**q3: What is the rough split between AI and humans in your organisation's work?**

---

<sup>6</sup>The full Business Outlook 2026 survey was administered by the British Chambers of Commerce (BCC) to its membership network in January 2026, yielding 672 qualified responses. The AI module was embedded within the broader survey instrument and appeared after questions on business growth outlook.

*[Current AI users only; single response; N = 354]*

- AI mostly handles tasks autonomously, with some human oversight
- Humans and AI handle tasks in roughly equal measure
- Humans mostly handle tasks, with AI providing support
- Unsure

**q4: Over the past 12 months, has your workforce size changed as a direct result of AI?**

*[Current AI users only; single response; N = 331]*

- Increased headcount significantly (more than 10%)
- Increased headcount slightly (1 to 10%)
- Had no bearing on headcount
- Decreased headcount slightly (1 to 10%)
- Decreased headcount significantly (more than 10%)
- Unsure

**q5: And over the next 12 months, will your workforce size change as a direct result of AI?**

*[Current AI users only; single response; N = 332]*

- Will increase headcount significantly (more than 10%)
- Will increase headcount slightly (1 to 10%)
- Will have no bearing on headcount
- Will decrease headcount slightly (1 to 10%)
- Will decrease headcount significantly (more than 10%)
- Unsure

**q6: Have any existing job roles in your organisation been restructured or redefined as a direct result of AI?**

*[All respondents; single response; N = 591]*

- Yes, most or all roles restructured or redefined
- Yes, a few roles restructured or redefined
- No, job roles have remained unchanged
- Unsure

**q7: Has AI changed the skills that your organisation needs?**

*[Current AI users only; multiple responses permitted; N = 332]*

- Greater emphasis on communication and interpersonal skills

- Greater emphasis on technical or analytical skills
- Greater emphasis on AI literacy
- Other skill set (*please specify*)
- No noticeable change
- Unsure or too soon to say

**q8: Over the next 12 months, will you invest in job-specific AI training?**

*[All respondents; single response; N = 658]*

- Significant investment
- Some investment
- None at all
- Unsure

**q9: Over the next twelve months, what impact do you think AI will have on your organisation's overall productivity?**

*[All respondents; single response; N = 655]*

- Significant increase
- Some increase
- No change
- Some decrease
- Significant decrease
- Unsure

Note that Questions 2 and 7 permitted multiple responses. Question 2 yielded a mean of 1.07 responses per respondent; Question 7 yielded 1.23 responses. For Question 2, firms selecting both generic and bespoke options are classified as “bespoke” users in the analysis, on the grounds that bespoke deployment represents a deeper level of organisational AI integration. Open-ended responses to the “please specify” prompts in Question 2 are coded and reported in Table B.7.

### **Region aggregation**

Because of limited coverage in some of the smaller regions, we aggregate the 12 government office regions to 8 coarser groups. London and the South East are combined into a single category, as are Yorkshire and the Humber with the North East. The three devolved nations (Scotland, Wales, and Northern Ireland) are combined into a single “UK outside England” category. The remaining English regions (East Midlands, West Midlands, East of England, North West, and South West) are retained as separate categories. London and the South East serves as the reference category throughout.

## Sector aggregation

The BCC survey classifies firms into 19 sectors. Several of these contain very few respondents, making it impractical to include 19 sector dummies in our regressions. We therefore aggregate to 10 coarser categories, merging small sectors with substantively similar neighbours while keeping analytically important sectors separate. Table A.1 reports the mapping. The smallest coarse groups (Transport & logistics and Hospitality & leisure) each contain 34 firms, which is adequate for sector fixed-effect estimation.

**Table A.1:** Sector aggregation: 19 BCC sectors to 10 coarse categories

Coarse sector	Fine BCC sectors included	<i>N</i>
Manufacturing	Manufacturing or production	137
	Pharmaceutical or scientific services	9
Finance & professional	Finance, accounting, or insurance	38
	Legal services	14
	Admin, support, or consulting	52
Construction & property	Construction, engineering, or trades	61
	Real estate, property, or development	29
Retail & wholesale	Retail or wholesale	64
Other services	Other services not listed above	49
	Utilities, waste, or energy supply	5
	Agriculture, forestry, fishing, or mining	7
Education & health	Education	28
	Health, social work, or third sector	25
	Public administration or defence	1
IT & digital	IT, data analysis, web or data services	45
Creative & media	Marketing, advertising, or communications	30
	Arts, entertainment, or recreation	10
Hospitality & leisure	Hospitality, catering, or tourism	34
Transport & logistics	Transport, logistics, or storage	34

## Sector-level AI exposure index from Felten et al. (2021)

To characterise sectors by their underlying exposure to AI capabilities, we draw on the AI Occupational Exposure (AIOE) index of Felten et al. (2021). The AIOE measures, for each occupation, the degree to which its constituent tasks overlap with advances in AI capabilities recent to their analysis. The original index is also provided at the level of 4-digit NAICS industries in the United States.

We construct a crosswalk from 4-digit NAICS industries to the 19 BCC sectors as follows. Most BCC sectors map straightforwardly to one or more 2-digit NAICS codes (e.g., Manufacturing maps to NAICS 31–33; Retail or wholesale to NAICS 42, 44–45). The main exception is the professional services sector (NAICS 54), whose 4-digit industries span several BCC categories: Legal services maps to NAICS 5411; Marketing, advertising, or communications to 5418–5419; Pharmaceutical or scientific services to 5417; and the remaining professional services codes

(5412–5416) are mapped to the BCC’s Admin, support, or consulting. For each BCC sector, we compute an unweighted average of the AIOE scores across its constituent 4-digit NAICS industries. We lack UK-specific employment weights at this level of detail. We therefore use unweighted averages.

In our regressions we standardise the index to have mean zero and unit standard deviation across sectors. When coarse sector fixed effects are included, the Felten index is omitted to avoid collinearity, since it varies only at the 19-sector level. Conversely, when the continuous Felten index is used, sector fixed effects are replaced by the index.

### **Leave-one-out sector adoption rates**

As a second sector-level characteristic, we compute the *ex-post* AI adoption rate within each sector, using question q2. For each firm  $i$  in sector  $s$ , we define the leave-one-out adoption rate as the share of *other* firms in sector  $s$  that report currently using AI (either generic or bespoke tools). This avoids the mechanical correlation that would arise if firm  $i$ ’s own adoption status entered both sides of the regression. The resulting variable captures the density of AI adoption in the firm’s immediate sectoral peer group and is standardised in the same way as the Felten index.

The adoption rate is computed at a 17-group sectoral classification: 16 of the original 19 BCC sectors are retained individually, while the three smallest (Agriculture, forestry, fishing, or mining ( $N = 7$ ), Public administration or defence ( $N = 1$ ), and Utilities, waste, or energy supply ( $N = 5$ )), are pooled into a single “Primary/public/utilities” group to stabilise the estimate. This intermediate classification is finer than the 10-category coarse sectors used for fixed effects, preserving more cross-sector variation in adoption rates, while avoiding the instability that would arise from computing a leave-one-out mean within cells as small as  $N = 1$ .

## B Additional Results

**Table B.1:** Full Sample Description by AI Adoption Status

	Currently using	Intend to use	No plans	Unsure	<i>p</i> -value
Unconditional ( <i>N</i> = 668)	54.3%	17.5%	22.3%	5.8%	
<i>Firm size</i>					
Sole trader ( <i>N</i> = 68)	38.2%	13.2%	45.6%	2.9%	<0.001
1-9 ( <i>N</i> = 257)	58.0%	12.1%	26.5%	3.5%	
10-49 ( <i>N</i> = 190)	51.1%	24.2%	17.9%	6.8%	
50-249 ( <i>N</i> = 114)	61.4%	18.4%	10.5%	9.6%	
250+ ( <i>N</i> = 39)	53.8%	25.6%	10.3%	10.3%	
<i>Sector</i>					
Construction & property ( <i>N</i> = 90)	50.0%	22.2%	23.3%	4.4%	<0.001
Creative & media ( <i>N</i> = 40)	70.0%	10.0%	17.5%	2.5%	
Education & health ( <i>N</i> = 54)	51.9%	24.1%	14.8%	9.3%	
Finance & professional ( <i>N</i> = 103)	68.0%	13.6%	16.5%	1.9%	
Hospitality & leisure ( <i>N</i> = 33)	69.7%	15.2%	12.1%	3.0%	
IT & digital ( <i>N</i> = 45)	77.8%	11.1%	11.1%	0.0%	
Manufacturing ( <i>N</i> = 146)	44.5%	23.3%	26.7%	5.5%	
Other services ( <i>N</i> = 59)	30.5%	18.6%	40.7%	10.2%	
Retail & wholesale ( <i>N</i> = 64)	54.7%	9.4%	25.0%	10.9%	
Transport & logistics ( <i>N</i> = 34)	47.1%	14.7%	23.5%	14.7%	
<i>Region</i>					
East Midlands ( <i>N</i> = 50)	56.0%	14.0%	26.0%	4.0%	0.050
East of England ( <i>N</i> = 86)	58.1%	17.4%	15.1%	9.3%	
London + SE ( <i>N</i> = 143)	60.1%	14.7%	19.6%	5.6%	
North West ( <i>N</i> = 97)	51.5%	16.5%	20.6%	11.3%	
South West ( <i>N</i> = 67)	62.7%	13.4%	22.4%	1.5%	
UK ex Eng ( <i>N</i> = 49)	34.7%	26.5%	32.7%	6.1%	
West Midlands ( <i>N</i> = 120)	55.8%	20.8%	20.8%	2.5%	
Yorkshire + NE ( <i>N</i> = 56)	41.1%	19.6%	33.9%	5.4%	
<i>Revenue band</i>					
Under £90k ( <i>N</i> = 109)	47.7%	12.8%	37.6%	1.8%	0.024
£90k–£249k ( <i>N</i> = 70)	62.9%	11.4%	24.3%	1.4%	
£250k–£499k ( <i>N</i> = 58)	51.7%	24.1%	17.2%	6.9%	
£500k–£749k ( <i>N</i> = 32)	71.9%	6.2%	21.9%	0.0%	
£750k–£999k ( <i>N</i> = 41)	65.9%	14.6%	17.1%	2.4%	
£1m–£2.49m ( <i>N</i> = 107)	49.5%	22.4%	21.5%	6.5%	
£2.5m–£4.99m ( <i>N</i> = 62)	51.6%	19.4%	22.6%	6.5%	
£5m–£9.99m ( <i>N</i> = 50)	56.0%	18.0%	20.0%	6.0%	
£10m–£19.99m ( <i>N</i> = 38)	57.9%	26.3%	13.2%	2.6%	
£20m–£49.99m ( <i>N</i> = 28)	57.1%	25.0%	10.7%	7.1%	
£50m+ ( <i>N</i> = 26)	76.9%	15.4%	3.8%	3.8%	
<i>N</i>	363	117	149	39	

Notes: Each row shows, for firms with that characteristic, what share fall into each adoption group (rows sum to 100%). *p*-values: Pearson  $\chi^2$  for each categorical block.

**Table B.2:** Sample Summary by AI Adoption Status (Strict SME-only, Weighted)

	Currently using	Intend to use	No plans	Unsure	<i>p</i> -value
Unconditional ( <i>N</i> = 561)	56.2%	15.5%	21.7%	6.6%	
<i>Conditional adoption splits (weighted %)</i>					
Positive growth outlook ( <i>N</i> = 260)	66.9%	18.7%	9.2%	5.2%	<0.001
Negative growth outlook ( <i>N</i> = 88)	41.2%	8.6%	45.5%	4.7%	0.001
Trades goods internationally ( <i>N</i> = 227)	58.1%	14.3%	20.8%	6.7%	0.277
Trades services internationally ( <i>N</i> = 184)	74.8%	11.3%	10.6%	3.4%	0.003
<i>Group means (weighted)</i>					
Mean ln(firm size)	2.903	3.178	2.483	3.527	<0.001
Mean ln(revenue per worker)	4.033	4.178	3.920	4.151	0.527
Mean operating margin	11.954	12.049	12.780	11.705	0.817
Mean AI-exposure (Felten et al.)	-0.046	-0.237	-0.258	-0.593	0.002
<i>N</i>	316	98	114	33	

*Notes:* SME-only sample (1–249 employees). Weighted by sector  $\times$  region raking to BPE population. Conditional adoption splits show, for firms with each characteristic, weighted share in each adoption group (rows sum to  $\approx$  100%). *p*-values: weighted logistic Wald test (binary rows); weighted WLS *F*-test (continuous rows). AI-exposure and sector adoption rate are standardised (z-scores).

**Table B.3:** Productivity Expectations (base = Intend to use)

	(1) OR	(2) OR	(3) OR	(4) OR
Unsure	0.417* (0.070)	0.444* (0.087)	0.535 (0.262)	0.457* (0.098)
No plans	0.084*** (<0.001)	0.079*** (<0.001)	0.058*** (<0.001)	0.090*** (<0.001)
Currently using	2.533*** (<0.001)	2.675*** (<0.001)	2.505*** (<0.001)	2.584*** (<0.001)
Positive growth outlook	2.082*** (<0.001)	1.995*** (<0.001)	1.960*** (<0.001)	2.027*** (<0.001)
Negative growth outlook	0.705 (0.220)	0.699 (0.202)	0.652 (0.132)	0.701 (0.205)
AI-exposure		1.563*** (<0.001)	1.980*** (0.002)	
Sector adoption rate				1.401*** (<0.001)
Unsure × AI-exposure			1.472 (0.582)	
No plans × AI-exposure			0.372*** (0.004)	
Using × AI-exposure			0.826 (0.440)	
Sector FE	Yes	No	No	No
Region FE	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
<i>N</i>	602	602	602	602
Pseudo <i>R</i> <sup>2</sup>	0.2304	0.2236	0.2318	0.2179
Log-likelihood	-457.0	-461.1	-456.2	-464.5

*Notes:* Base category = Intend to use. Cf. Table 3 where base = No plans. Firm controls: ln(size), ln(revenue per worker), operating margin midpoint, trades goods, trades services (suppressed). AI-exposure and Sector adoption rate are standardised (mean 0, SD 1). Robust standard errors. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table B.4:** Robustness: Non-Imputed Sample

Original table	Key coefficient	Full sample			Non-imputed		
		OR	$p$	$N$	OR	$p$	$N$
Table 2 col(4)	AI-exposure	1.257	0.021	668	1.329	0.016	501
Table 2 col(4)	Positive growth outlook	2.240	<0.001	668	2.466	<0.001	501
Table 2 col(4)	Trades services internationally	1.905	<0.001	668	1.817	0.006	501
Table 3 col(3)	AI-exposure	1.562	<0.001	602	1.579	<0.001	461
Table 5 col(1)	Bespoke AI user	0.155	<0.001	326	0.128	<0.001	254
Table 6 col(1)	Bespoke AI user	3.226	0.002	326	3.596	0.004	252
Table 7 col(2)	AI-exposure	1.514	0.003	337	1.438	0.018	260
Table 7 col(3)	Sector adoption rate	1.490	0.019	337	1.516	0.028	260
Table 8B col(1)	Any training investment	3.531	0.012	305	8.071	0.001	237
Table B8 col(1)	Bespoke AI user	0.454	0.061	305	0.558	0.237	237
Table B8 col(3)	Bespoke $\times$ AI-exposure	2.318	0.038	305	2.631	0.032	237

*Notes:* Non-imputed sample excludes observations where any of the following were imputed: ln(revenue per worker) (median), operating margin (median), growth outlook (zero-imputed), trades goods (zero-imputed), trades services (zero-imputed). Each row re-estimates the preferred specification from the indicated table.

**Table B.5:** Regional Variation in AI Adoption (Binary Logit)

Region	OR	95% CI	<i>p</i> -value
London + SE		[base]	
East Midlands	0.919	[0.464, 1.819]	0.808
East of England	0.884	[0.501, 1.561]	0.671
North West	0.698	[0.393, 1.240]	0.220
South West	1.361	[0.713, 2.598]	0.350
UK ex Eng	0.330***	[0.155, 0.706]	0.004
West Midlands	0.768	[0.443, 1.332]	0.348
Yorkshire + NE	0.440**	[0.217, 0.894]	0.023
Firm controls		Yes	
Sector FE		Yes	
<i>N</i>		668	
Pseudo <i>R</i> <sup>2</sup>		0.1153	
Log-likelihood		-407.4	

*Notes:* Binary logit (using AI = 1, else = 0). Region dummies with London + South East as base. Firm controls: ln(size), ln(revenue per worker), operating margin, growth outlook, trades goods, trades services. Robust standard errors. \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

**Table B.6:** Longitudinal Adoption Transitions (Panel Sub-sample)

<i>Panel A: 2024 → 2025 (N = 91)</i>				
From \ To	Unsure	No plans	Intend	Using
Unsure	0% (0)	20% (1)	20% (1)	60% (3)
No plans	3% (1)	54% (20)	22% (8)	22% (8)
Intend to use	4% (1)	4% (1)	54% (13)	38% (9)
Currently using	0% (0)	4% (1)	0% (0)	96% (24)

<i>Panel B: 2025 → 2026 (N = 170)</i>				
From \ To	Unsure	No plans	Intend	Using
Unsure	25% (2)	12% (1)	12% (1)	50% (4)
No plans	11% (4)	71% (27)	11% (4)	8% (3)
Intend to use	0% (0)	15% (6)	41% (16)	44% (17)
Currently using	1% (1)	1% (1)	9% (8)	88% (75)

<i>Panel C: Key transition statistics</i>	
AI use persistence (2024→2025)	96.0%
AI use persistence (2025→2026)	88.2%
Intender conversion rate (1-year, latest)	43.6%
Intender conversion (2-year cumulative)	51.2%
“No plans” conversion (2-year)	44.6%
Net movement (2024→2025)	Up: 30, Down: 4, Net: 26
Net movement (2025→2026)	Up: 30, Down: 20, Net: 10

*Notes:* Small panel sub-sample. Row percentages shown with counts in parentheses. Retrospective measurement; no controls. Transition matrices condition on non-missing adoption status in both waves.

**Table B.7:** Open-Ended AI Use Case Taxonomy

Use case category	Generic users		Bespoke users	
	<i>N</i>	%	<i>N</i>	%
Content & writing	51	23.6%	3	6.8%
Research & information	34	15.7%	2	4.5%
Admin & operations	24	11.1%	10	22.7%
Coding & technical	15	6.9%	2	4.5%
Data & analytics	6	2.8%	7	15.9%
Creative & design	16	7.4%	8	18.2%
Customer-facing & sales	29	13.4%	6	13.6%
Tool named only	68	31.5%	1	2.3%
Other/unclassified	34	15.7%	15	34.1%
Total respondents	216		44	

*Notes:* Categories are not mutually exclusive. Based on keyword matching of open-ended responses describing AI use cases. No formal tests — qualitative/descriptive.

**Table B.8:** Determinants of Expected Future Headcount Effects (Ordered Logit)

	AI users					Full sample
	(1) OR	(2) OR	(3) OR	(4) OR	(5) OR	(6) OR
Bespoke AI user	0.454* (0.061)	0.589 (0.196)	0.663 (0.306)	0.605 (0.218)	0.411** (0.041)	
Bespoke × AI-exposure			2.318** (0.038)			
AI-exposure		1.237 (0.233)	1.048 (0.811)			
Sector adoption rate				0.969 (0.865)		
Restructuring (ordered)					0.963 (0.916)	
AI user						0.503* (0.057)
Bespoke AI user						0.332** (0.021)
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Sector FE	Yes	No	No	No	Yes	Yes
Region FE	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	305	305	305	305	303	610
Pseudo <i>R</i> <sup>2</sup>	0.1033	0.0622	0.0736	0.0584	0.1214	0.1088
Log-likelihood	-170.0	-177.8	-175.6	-178.5	-164.3	-201.8

*Notes:* DV: expected AI headcount effect (5-point ordered scale). Odds ratios reported. Firm controls (suppressed): ln(size), ln(revenue per worker), operating margin, growth outlook, trades goods, trades services. Robust standard errors. Col 5 adds restructuring (ordered) as mediator to Col 1 specification. Col 6 includes non-AI users (coded as “no bearing”); bespoke set to 0 for non-users. AI-exposure and Sector adoption rate are standardised (mean 0, SD 1). \*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .