

LLMs and Youth Labour Reallocation: Evidence from the United Kingdom

Claudio Costanzo*

October 2026

Abstract

After the release of ChatGPT in November 2022, young people in Britain began to enter different jobs. In the Labour Force Survey, the share of new hires aged 22–25 starting in elementary occupations rose from 0.19 in 2015–2022 to between 0.24 and 0.30 in 2023–2025, as much among recent graduates as among non-graduates, while their employment and pay changed little. The fall in demand is where language models are used. Within broad occupation groups, online job adverts fell after 2022 by 0.16 log points per standard deviation of observed exposure to the models. Hires of 22–25-year-olds fell by 28 percent in exposed occupations while they were unchanged elsewhere. The 2008–2011 downturn, a far deeper hiring slowdown, cut administrative entry jobs by as much but raised graduates’ elementary entry by a borderline three points, against eight after 2022. The evidence suggests a fall in the demand for young labour in the occupations whose tasks language models perform, with a rising minimum wage and migration shaping where the entrants went.

JEL codes: J23, J24, J31, O33.

Keywords: large language models, entry-level jobs, young workers, occupations, job adverts, minimum wage.

*KU Leuven. Financial support from the Research Foundation Flanders (FWO) is gratefully acknowledged. Contact: claudio.costanzo@kuleuven.be

1 Introduction

Large language models became a mass-market product with the release of ChatGPT on 30 November 2022. Within two years, more than a quarter of employed adults in the United States reported using generative artificial intelligence at work (Bick et al., 2024). In the United Kingdom, employer surveys recorded a rise in the share of businesses using text generation with language models from one in twenty in late 2023 to one in six by mid-2026.¹ These models draft, summarise, answer routine queries and write code, the tasks that have long been given to new entrants in professional and administrative careers. An early hypothesis, sharpened by Brynjolfsson et al. (2026), holds that language models reduce the demand for inexperienced cognitive labour before they reduce the demand for experienced labour, so that young workers show the effects of the technology first. The consequences, if it holds, reach beyond the young, because a first job shapes a career (Kahn, 2010; Oreopoulos et al., 2012; von Wachter, 2020).

This paper examines the hypothesis for the United Kingdom. It asks how the entry of young people into employment changed after November 2022, and whether the change was caused by language models or by the hiring slowdown, the minimum-wage increases and the migration flows that coincided with their diffusion. On the first question, the paper finds that young people began to enter different jobs. Among new hires aged 22–25, the share starting in elementary occupations rose above the range of the previous eight years, as much among recent graduates as among non-graduates, while their employment and pay moved little. On the second, it finds that the fall in demand is located where language models are used, in comparisons of occupations with their neighbours in the same broad group and net of the trend that preceded the release. In the 2008–2011 downturn, a far deeper hiring slowdown, graduates moved into low-paying jobs as much as after 2022 but into elementary jobs by much less. After 2022 the wage floor and migration shaped where the displaced entrants went.

Most of the evidence comes from the Quarterly Labour Force Survey (LFS), 81 quarterly files from 2006 to mid-2026 with the occupation, tenure, qualification and pay of every respondent. Entry is measured with new hires, employees who had been with their employer for less than three months, among people aged 22–25, the ages at which graduates enter the labour market and non-graduates are a few years into it. Online job adverts by detailed occupation measure the demand for labour at monthly frequency from 2017. An employer survey measures the adoption of language models by industry. Administrative payroll and

¹The share of businesses using text generation with large language models in the Business Insights and Conditions Survey rose from 0.051 in the wave of September 2023 to 0.171 in the wave of June 2026 (section 2.5).

vacancy data measure the hiring cycle. To these data the paper brings five measures of the exposure of occupations to language models. One is the ability-based measure that the Department for Education built from the abilities an occupation uses (Department for Education, 2023). Four are task-based measures from US sources, mapped here to British occupations: the task-feasibility rating of Eloundou et al. (2024), the observed-exposure index and the usage share of the Anthropic Economic Index (Handa et al., 2025; Anthropic, 2026), and the applicability score of Tomlinson et al. (2025).

Three facts describe what changed. First, the young started different jobs. Among new hires aged 22–25, the share in elementary occupations, which are cleaning, catering, bar work, storage and similar jobs, averaged 0.19 over 2015–2022 and rose to 0.30, 0.24 and 0.27 in 2023–2025 (0.25, 0.22 and 0.23 without survey weights). Over the same years the share in administrative jobs fell by a third and the share in professional jobs did not move. This shift in flows, which left the stock of young employees in elementary jobs barely changed, is as large among recent graduates as among non-graduates. Among graduates who were full-time students a year earlier, the share now working in an elementary job more than doubled in 2024–2025. Second, employment and pay moved less than the mix of jobs. The employment rate of 22–25-year-olds fell by two percentage points with survey weights and by a tenth of that without them. At these ages the graduate pay premium fell by five log points. The share of young employees paid within 5 percent of the rising minimum wage went from 0.12 to about one in six. Third, the fall in demand is where language models are used. Along the observed-exposure index, online job adverts fell after 2022 by 0.16 log points per standard deviation within broad occupation groups, by 0.14 without weights and by 0.08 net of the trend of 2017–2022. The task-feasibility rating gives a within-group fall of the same sign and larger size, whereas the ability-based measure, which sorts occupations largely by broad group, gives none once the trend is allowed for. Young hiring fell in the same occupations, where hires of 22–25-year-olds fell by 28 percent between 2015–2022 and 2023–2025 while hires in unexposed occupations were unchanged. The within-group gradient accounts for three fifths of that fall, its part net of trend for a third.

The third fact is the comparison that identifies an effect of the technology. It sets occupations beside their neighbours in the same broad group, before and after the release, with exposure allowed a linear trend of its own over the pre-period, on the assumption that, absent the technology, adverts and hiring in exposed occupations would have moved with those in less exposed occupations of the same group. If the assumption holds, the estimate net of the trend is the effect on demand. The body of the paper addresses the threats to this comparison, of which one qualifies the estimate. The end of the pandemic hiring boom in technology falls on the IT occupations that the observed-exposure index ranks highest.

Without the twelve IT unit groups the within-group gradient remains, at 0.10 log points per standard deviation, but the part of the fall that exceeds the pre-2023 trend does not, which places the fall beyond trend in the IT occupations, where language models are most used for the work and where the boom also ended.

The gradient is estimated within broad groups, but the occupations holding the most exposed third of employment and those holding the least exposed third also part after 2022 when set side by side. Hires of 22–25-year-olds per 100 young employees fell from 8.0 to 5.8 in the former and rose from 7.1 to 7.7 in the latter. The share of graduate entrants starting in the former fell by ten points and the share starting in the latter rose by ten. These contrasts are between broad groups as well as within them.

Three forces coincided with the diffusion and could produce parts of the pattern. The first is the hiring cycle, since the boom of 2022 ended as the models diffused. As a benchmark for what a hiring slowdown alone does to the same statistics, the paper uses the 2008–2011 downturn. With a far larger rise in youth unemployment, it cut administrative entry jobs by the same amount and moved graduates into low-paying jobs as much as the recent episode did, though mostly into sales and care jobs. Their elementary share changed by one point in 2009–2011, the three years from the trough in output, and rose by a borderline three points in 2011–2013, the three years from the 2011 peak in unemployment, against eight in 2023–2025. The second and third forces are the minimum wage, which rose by about a tenth in 2023 and again in 2024, and migration into elementary jobs, which rose after 2021. Each lowered the cost or raised the supply of elementary work, though neither accounts for the fall in demand within broad groups in the occupations the models perform. The evidence points to a fall in the demand for labour in the occupations whose tasks language models perform, a fall that began with their diffusion and is clearest in information technology, and as a reallocation of young entrants away from that work.

The paper makes three contributions, the first to the literature on the labour-market effects of language models. Experiments have measured productivity gains, largest for novices (Noy and Zhang, 2023; Brynjolfsson et al., 2025). Exposure indices have been built from the abilities an occupation uses (Felten et al., 2021, 2023; Department for Education, 2023), from ratings of its tasks (Eloundou et al., 2024; Gmyrek et al., 2023) and from the tasks people bring to the models (Handa et al., 2025; Tomlinson et al., 2025; Massenkoff and McCrory, 2026). Before the release, American vacancy data showed no clear relation between exposure and employment growth at the level of occupations (Acemoglu et al., 2022), while firm data show modest within-firm effects of task exposure over 2010–2023 (Hampole et al., 2025). After the release, Danish survey-to-register data show no effect on earnings or hours through 2024 and no loss of early-career jobs at workplaces that adopted chatbots (Humlum

and Vestergaard, 2026), whereas American payroll records show the employment of 22–25-year-olds in exposed occupations 19 percent below the path of their less exposed peers by mid-2026, through hiring, a gap that halves under a control for the education of the workforce (Brynjolfsson et al., 2026). The same pattern for juniors appears in résumé, Census and French hiring records (Hosseini Maasoum and Lichtinger, 2026; Tucker, 2026; Orr et al., 2026; Palladino et al., 2026) and in Swedish registers (Lodefalk et al., 2026). Estimates for Norway and Finland are imprecise or null (Hernæs and Kostøl, 2026; Facius and Iacono, 2026; Kauhanen and Rouvinen, 2026). These studies measure stocks of employment, counts of hires or job-finding by the exposure of the occupation. For Britain the evidence comes from vendor records of hires and postings (Klein Teeselink, 2025; Lambert and Schindler, 2026; Chandar and Klein Teeselink, 2026) and from online adverts (Henseke et al., 2026; Schlicht, 2026; Egan, 2026). The paper provides, to my knowledge, the first evidence from household-survey microdata on the occupations in which young people start work, and the first evidence of that kind for the United Kingdom. It decomposes the exposure gradient in young hiring into its parts between and within major occupation groups. Along the ability-based measure the gradient is a contrast between professional and manual groups, whereas the measures built from observed use, the observed-exposure index and the usage share, which I map to SOC 2010 and SOC 2020 unit groups, vary within groups and show the gradient inside them. Earlier British mappings cover abilities (Department for Education, 2023), tasks (Henseke et al., 2026; Dwan-O’Reilly, 2026) and usage at the level of sub-major groups (Schlicht, 2026). The unit-group mapping of observed use, its application to new hires by age and graduate status, and the within-group design are new.

The second contribution is to the literature on entering the labour market in a downturn (Kahn, 2010; Oreopoulos et al., 2012; Altonji et al., 2016; Schwandt and von Wachter, 2019; von Wachter, 2020; Rothstein, 2023). For Britain that literature finds that the graduates who left education in the 2008–2009 recession were more often in low-paying occupations for years afterwards (Clarke, 2019; Henahan, 2020), while the lasting employment losses fell on the less educated (Cribb et al., 2017; Clarke, 2019). The paper compares the two episodes with the same data and tests, on that literature’s definition of low-paying occupations and on the narrower elementary group, from both the trough in output and the peak in unemployment of the earlier episode. On the broad definition the two episodes moved graduates into low-paying entry jobs by similar amounts, mostly into sales, customer-service and care jobs in the downturn and into elementary jobs after 2022, when graduates’ employment rate, which had fallen by five points in both recession windows, did not move.

The third contribution is to the literature on the polarisation of work and the decline of routine cognitive jobs (Autor et al., 2003; Goos and Manning, 2007; Autor and Dorn, 2013;

Goos et al., 2014; Michaels et al., 2014; Cortes et al., 2017; Salvatori, 2018; Jaimovich and Siu, 2020; Dillender and Forsythe, 2022), to the work on the demand for cognitive tasks and the upgrading of hiring in downturns (Beaudry et al., 2014, 2016; Hershbein and Kahn, 2018; Modestino et al., 2020), and to the flows view of occupational decline, in which shrinking occupations stop recruiting the young (Autor and Dorn, 2009; Cortes et al., 2020). The administrative share of young hires fell in 2023–2025 at a pace that cannot be distinguished from its 2015–2022 trend, outside the recessions in which Jaimovich and Siu (2020) locate past routine losses. In the LFS since 2006, graduates’ elementary share of entry jobs was at most 0.139 in any year before 2023 and 0.215, 0.179 and 0.186 in 2023–2025. Against 2015–2022 the rise is 7.8 points, more than twice the borderline rise of 2.9 points in the downturn measured from 2011. Graduates moved into non-graduate work between 2001 and 2012, into clerical, sales and care jobs, while the share of all graduates in low-skill work was 3.0 percent in 2013 and 3.4 percent in 2017 (Office for National Statistics, 2013, 2017). After 2012 the share of young graduates in high-skilled work recovered (Department for Education, 2025a). Beaudry et al. (2016) describe graduates cascading down the occupational ladder when the demand for cognitive tasks fell after 2000. The within-group exposure gradient and the contrast between administrative and elementary entry test that mechanism against the technology. The appendix follows the workers above the entry layer, whom the seniority-biased studies treat as insulated or complemented (Brynjolfsson et al., 2026; Hosseini Maasoum and Lichtinger, 2026; Autor and Thompson, 2025; Ide and Talamàs, 2025). Their ranks grew along a trend that predates 2023, with no acceleration after 2022, which matches the American estimates for workers over 35 and gives the complementarity reading no support in British data.

The rest of the paper proceeds as follows. Section 2 describes the diffusion of language models, the British labour market after 2022, the data and the exposure measures. Section 3 sets out the empirical strategy with its identifying assumptions. The results are in section 4, which covers the change in the entry jobs, employment and pay of the young and its relation to the exposure of occupations to language models. Section 5 uses the 2008–2011 downturn as a benchmark for the hiring cycle, gathers the mechanisms, from the elementary detour to the wage floor, migration and the layers above entry, and states what the evidence does not settle. Section 6 concludes. The appendices give the construction of the data and the exposure measures, additional results and robustness checks, the analysis of the layers above entry, the benchmark under alternative windows, and the responses of students and graduates by subject.

2 Background and Data

This section describes what the models do and the three forces that acted on the entry labour market as they diffused, then the data and the measures of exposure.

2.1 Language models and entry-level work

The models released from November 2022 differ from the automation technologies of the previous three decades in the tasks they perform. Industrial robots and the computerisation of routine work replaced codifiable manual and clerical tasks (Autor et al., 2003; Autor, 2015; Acemoglu and Restrepo, 2020), whereas the aggregate consequences of the new models depend on which tasks they take over (Acemoglu, 2025). Language models read, write and reason over text and code, work that the task literature had classified as non-routine cognitive. The exposure indices built since 2023 agree on the broad ordering of occupations, in which administrative, professional and managerial occupations are exposed and trades, operatives and elementary occupations are not (Eloundou et al., 2024; Felten et al., 2023; Department for Education, 2023). Within those groups the indices disagree on the ordering, as section 2.4 shows, and the disagreement decides what a comparison along each measure identifies.

The hypothesis that young workers bear the first effects has two parts. The first is that the tasks language models perform are the tasks delegated to juniors, such as a trainee solicitor’s research memorandum, a junior analyst’s first draft, a graduate programmer’s boilerplate or a customer-service agent’s scripted reply. If a model performs these tasks at low cost, the firm has less reason to hire a junior to perform them. The tasks of experienced workers, which involve judgement, relationships and responsibility, are affected later or complemented (Brynjolfsson et al., 2026; Hosseini Maasoum and Lichtinger, 2026; Autor and Thompson, 2025; Ide and Talamàs, 2025). The second part is that entry jobs are where experience accumulates. If fewer juniors are hired into cognitive occupations, a cohort of graduates starts in jobs that do not build the skills the professional occupations reward (Deming, 2017; Deming and Noray, 2020), and the effect on their careers outlasts the shock. The first part is testable with data on hiring by occupation and age, whereas the second requires following cohorts over years and is beyond the horizon of this paper.

2.2 The British labour market after 2022

Three other forces, with which any attribution must reckon, acted on the entry labour market in the same period.

The hiring cycle turned. Vacancies peaked at 1.29 million in the three months to June 2022, the highest level recorded by the Vacancy Survey, and fell to an average of 749 thousand in 2025, a level between those of 2015 and 2019. Monthly inflows into payrolled employment in 2025 stood at 84 percent of their 2022 level, while payrolled employees aged 18–24 stood at 99 percent of the same base, against 105 for those aged 35–49.² By the standard of that recession the slowdown of 2023–2025 was mild in employment and unemployment, and sharp in vacancies because 2022 was a peak.

The wage floor rose. The National Living Wage, the minimum wage of adults, rose from 9.50 pounds an hour to 10.42 in April 2023 and to 11.44 in April 2024, when it was extended from workers aged 23 and over to those aged 21 and over, and to 12.21 in April 2025 (Low Pay Commission, 2024). These increases were among the largest since the introduction of the National Living Wage and brought the floor to the government’s target of two thirds of median hourly earnings (Low Pay Commission, 2024). They raised the pay of elementary jobs relative to professional jobs and, as section 4.4 shows, the share of young employees paid at the floor.

Migration into elementary jobs rose. After net migration to the United Kingdom reached record levels in 2022 and 2023 (Office for National Statistics, 2024b), the number of foreign-born employees in elementary occupations, which had fallen from 771 thousand in 2015 to 527 thousand in 2021, returned to 766 thousand by 2025. Young natives and recent migrants entered the same occupations at the same time. Section 4.5 documents the composition, which the paper treats as a second supply-side force on the jobs that took in young entrants.

2.3 The Labour Force Survey

The main source is the Quarterly Labour Force Survey of the Office for National Statistics, the household survey from which the official labour-market statistics of the United Kingdom are produced (Office for National Statistics, 2026c). For the main analysis I use the End User Licence person files for every quarter from the first quarter of 2015 to the second quarter of 2026, 46 files, while the 2008–2011 benchmark of section 5.1 uses the files from the first quarter of 2006 to the fourth quarter of 2014.³ Every respondent reports their labour-market status, their occupation and industry, the month they started their current job, their highest qualification, the age at which they left full-time education and, for a subset, their hourly

²Section 5.1 compares these movements with those of the 2008–2009 recession, when vacancies fell to the lowest annual average in the record of the Vacancy Survey, which begins in 2001.

³The April–June 2007 file is not distributed, which leaves 2007 with three quarters. Each quarterly file holds between 26,000 and 76,000 respondents aged 16–64. Because households stay in the survey for five consecutive quarters (Office for National Statistics, 2026c), a person appears in up to five files, all of which are used, with inference clustered by household within quarter, as explained in section 3.1.

pay.

The quantities that recur throughout the paper are defined as follows. A *new hire* is an employee, as opposed to a self-employed person, who has been with the current employer for less than three months at the interview. A *young* person is aged 22–25. In this range the median graduate has left full-time education and the median non-graduate has been in the labour market for several years, which places both groups at the entry stage. *Occupation* is the four-digit unit group of the Standard Occupational Classification 2010 (SOC 2010), which has 369 unit groups (Office for National Statistics, 2010). The code used is the unit group coded by the ONS for 2015–2020 and, from 2021, the ONS mapping of the new SOC 2020 code to SOC 2010.⁴ The first digit of the unit group is the *major group*, also called the broad group here: managers (1), professionals (2), associate professionals (3), administrative and secretarial (4), skilled trades (5), caring and leisure (6), sales and customer service (7), process and machine operatives (8) and elementary occupations (9). Groups 1–3 are called professional jobs below, and groups 5, 8 and 9 manual jobs. A *graduate* holds a degree or an equivalent qualification. *Years since education* is age minus the age at which the person left full-time education.

The April–June files include a retrospective module that asks every respondent what they were doing twelve months earlier. It yields a second measure of entry that does not depend on tenure, the share of people aged 22–25 who were full-time students a year earlier and now work as an employee in an elementary job. The two-quarter longitudinal files, which link the interviews of the same households in consecutive quarters, give quarter-to-quarter transitions between jobs and labour-market states (appendix A). An *entry job*, throughout, is the job of a new hire at the entry stage, aged 22–25 or, where stated, 0–3 years out of full-time education. A *recent graduate* is a graduate who left full-time education 0–3 years before the interview.

The sample of young new hires is small and shrank after 2020. Across 2015–2026 the files contain 89,246 interviews of employees aged 22–25 in occupations with an exposure score, of whom 6,343 were new hires. By year, the count of young new hires falls from 849 in 2015 to 452 in 2022 and 244 in 2023, before recovering to 460 in 2025 (appendix table A1). The fall reflects the collapse of LFS response rates during and after the pandemic. In 2024 the ONS reweighted the survey (Office for National Statistics, 2024a), and the new weights are more dispersed than the old. The Kish effective sample of young new hires is 0.67–0.78 of the interview count in 2023–2025 against 0.96 in the 2000s, while the largest person weights in 2023–2025 exceed 8,000 against 3,300–4,700 in 2012–2014. Six in ten young new hires

⁴Appendix A shows that the main results are the same in native SOC 2020 codes for 2021–2026, and documents the earlier change from SOC 2000 to SOC 2010 in 2011, which bears on the benchmark.

in 2023–2025 are proxy responses, reported by another member of the household, against five in ten in 2015–2022. For these reasons the paper reports the headline series on entry jobs, employment and hiring with and without weights and, for the main series, by type of response, and treats the unweighted estimates as the lower bound.

2.4 Exposure of occupations to language models

The paper uses five measures of the exposure of an occupation to language models, one ability-based and four task-based. The first is the language-model exposure score that the Department for Education constructed for each SOC 2010 unit group (Department for Education, 2023) by the method of Felten et al. (2021, 2023), in which each of 52 abilities in the O*NET database is rated for its relation to language-model capabilities and an occupation’s score is the prevalence-and-importance-weighted sum of the abilities it uses. This is the ability-based measure. Four task-based measures rate an occupation’s tasks. The feasibility rating of Eloundou et al. (2024) is the share of an occupation’s tasks for which a GPT-4-class model, with the software built around it, could halve the time needed, as rated by the model itself over O*NET task descriptions before the release. The observed-exposure index of the Anthropic Economic Index (Anthropic, 2026) combines the feasibility of an occupation’s tasks with the frequency with which those tasks appear in work-related conversations with Claude, and is reported for 756 detailed occupations of the US Standard Occupational Classification. From the same source, the usage share is the share of conversations from Great Britain in April–May 2026 by the O*NET occupation of the task, divided by the occupation’s share of British employment. The applicability score of Tomlinson et al. (2025) is built from the activities performed in Bing Copilot conversations. These names are used throughout.

The four US measures, defined on O*NET or US SOC occupations, are moved to SOC 2010 unit groups in two steps. The first step uses the O*NET-to-ISCO-08 crosswalk published with the European ESCO classification, the second the ONS proportional conversion from ISCO-08 to SOC 2010. At the first step each O*NET occupation is weighted by one over the number of ISCO codes it maps to, which gives less weight to occupations spread over many ISCO codes. At the second step each SOC 2010 group takes the mean of its ISCO codes’ scores weighted by the ONS proportions. Usage shares are allocated along the same route and summed. Appendix A.2 documents the crosswalks, the coverage of each step and the plausibility of the mapping for benchmark occupations. All five measures are standardised to mean zero and standard deviation one over employees aged 16–64 in 2015–2022, weighted, so that a coefficient is the effect of one standard deviation of exposure. Earlier

British mappings exist for the ability-based measure (Department for Education, 2023), for the task ratings of Eloundou et al. (2024), carried to SOC 2020 by Henseke et al. (2026), for an ILO task index (Dwan-O'Reilly, 2026), and for a composite that includes the observed-exposure index at the level of sub-major groups (Schlicht, 2026). The unit-group mapping of the observed-exposure index and of the usage share is new here.

The measures differ in a way that is central to the paper. Appendix table A2 decomposes the employment-weighted variance of each standardised measure into the part between SOC major groups and the part within them. For the ability-based measure 0.73 of the variance is between the nine major groups and 0.84 between the 25 sub-major groups, which means that it sorts occupations largely by broad group, placing administrative, professional and managerial jobs above trades, operatives and elementary jobs. For the observed-exposure index the between-group share is 0.39, for the applicability score 0.44, for the usage share 0.34, and for the feasibility rating 0.63. Within major groups the observed-exposure index, the applicability score and the usage share correlate with the ability-based measure at 0.51, 0.55 and 0.20. Among professionals the ability-based measure ranks solicitors, psychologists and teachers highest, whereas the observed-exposure index ranks programmers, web developers and IT analysts highest, and, outside the professions, call-centre staff and medical secretaries. A regression that compares occupations along the ability-based measure therefore compares broad groups, as section 4.7 shows for young hiring, whereas the task-based measures also order occupations within them. The mapping of the task-based measures to British occupations is imperfect, because many of the SOC 2010 unit groups that receive a US score inherit it from a single ISCO code whose O*NET occupations are spread over much of the exposure distribution. Differences between neighbouring unit groups are therefore measured with error of the order of the signal (appendix A.2), which attenuates the task-based estimates.

2.5 Online job adverts, employer adoption and other sources

The ONS publishes monthly counts of new online job adverts by SOC 2020 unit group from January 2017, derived from the Textkernel database of adverts scraped from job boards and employer websites (Office for National Statistics, 2026b). The paper uses the UK totals for 412 unit groups over the 115 months to August 2026 for which the ONS publishes the series, a run from which March 2026 is missing. Cells suppressed below the disclosure threshold, 813 of them and almost all in 2025–2026, are dropped. Because adverts measure the flow of demand for labour at the point where it is first expressed, in every occupation and at monthly frequency, which no survey can match, they have been used to study skill demands over the

cycle (Hershbein and Kahn, 2018; Forsythe et al., 2020). They do not measure hires, however, and the firms that advertise online are not all firms. The SOC 2010 exposure measures are moved to SOC 2020 through the ONS relationship table between the two classifications, weighting each SOC 2010 group by its share of the SOC 2020 group’s employment. Median hourly pay by SOC 2020 unit group for 2021–2025 comes from the Annual Survey of Hours and Earnings (Office for National Statistics, 2025) and is used here to describe the movement of pay across occupations.

Since September 2023 the Business Insights and Conditions Survey (BICS), a fortnightly survey of businesses, has asked which artificial-intelligence technologies a business uses (Office for National Statistics, 2026a). The adoption measure is the weighted share of businesses in each of fourteen industry sections that use text generation with large language models, over the twelve waves from September 2023 to June 2026, matched to the industry section of LFS employees. Across all businesses the share rose from 0.051 in the wave of September 2023 to 0.171 in the wave of June 2026. The HM Revenue and Customs Real Time Information system provides monthly counts of payrolled employees by age band and industry and of monthly inflows into payrolled employment (HM Revenue and Customs and Office for National Statistics, 2026). The Vacancy Survey provides three-monthly vacancies by industry (Office for National Statistics, 2026d). Appendix H uses two further sources, the Longitudinal Education Outcomes data on graduates by subject (Department for Education, 2025b) and UCAS applications by subject group (UCAS, 2026). Appendix A gives the details of each source.

3 Empirical Strategy

The evidence has two layers. One describes how the entry of young people changed after November 2022, in comparisons of the years after the release with the years before. The other asks whether the change is where language models are used, in comparisons across occupations by their exposure. This section sets out the estimators, the identifying assumptions and the threats to them, while sections 4 and 5 report the estimates.

3.1 Before-and-after comparisons

For a share y of young new hires, or of young people, the change after 2022 is summarised in two ways. The first is a year-level test in which each calendar year from 2015 to 2025 is

one observation and the annual share is regressed on an indicator for 2023–2025,

$$y_t = \alpha + \beta \mathbf{1}[t \geq 2023] + \varepsilon_t, \quad t = 2015, \dots, 2025, \quad (1)$$

with 9 degrees of freedom, and in a variant with a linear trend in t added. Beside it the tables report the range of the 2015–2022 values, which shows whether the post-2022 values fall outside it. The second is a person-level regression of the indicator for interview i in quarter q ,

$$y_{iq} = \alpha + \beta \mathbf{1}[2023Q1 \leq q \leq 2025Q4] + \mu_{m(q)} + \varepsilon_{iq}, \quad (2)$$

with quarter-of-year effects $\mu_{m(q)}$, person weights and standard errors clustered by household within quarter, because the household identifier is unique only within a quarterly file. It uses the quarterly variation and is the inference relied on when comparing magnitudes across samples or episodes. Annual shares in the figures come with 95 percent intervals from the same clustering. Both tests compare 2023–2025 with 2015–2022.⁵ Because the achieved sample of the LFS fell after 2020 and the 2024 reweighting gives young people a higher relative weight, the headline estimates on entry jobs, employment and hiring are reported with and without weights, and the main series by type of response.

A before-and-after comparison attributes to the period everything that changed in it. The 2008–2011 downturn provides a benchmark for what a hiring slowdown alone does to the same statistics. Equation (2) is repeated on 2009–2011, the three years from the trough in output, and on 2011–2013, the three years from the peak in unemployment, against 2006–2007, with the same definitions. A z statistic from the two standard errors compares the two episodes (section 5.1). The benchmark assumes that the earlier episode was a hiring slowdown without a new technology acting on the same jobs, in which case whatever a slowdown does to the entry jobs of the young appears in it. It covers the entry-job, employment and hiring statistics of the LFS only, because adverts begin in 2017 and the adoption survey in 2023, which leaves the exposure designs below without a counterpart in the earlier episode.

3.2 Comparisons across occupations by exposure

The second layer compares occupations by their exposure to language models, measured as in section 2.4, over the quarters from 2015Q1 to 2026Q2, with the post period running from 2023Q1. Adverts run from January 2017 to August 2026 with the post period from January 2023. Three designs are used.

⁵The first half of 2026 is reported separately, because it is a half-year and because the elementary share returned to its earlier range in it.

The first is descriptive: the SOC 2010 unit groups are cut into thirds of 2015–2022 employment by the observed-exposure index, after which the top and bottom thirds are followed side by side. Thirds give groups of equal employment with about two hundred young new-hire interviews a year in each. Appendix E shows the other cuts. For an outcome y_{ot} of unit group o in quarter t , the share aged 22–25 among its new hires or its young new-hire rate, the regression on the cells of the two groups is

$$y_{ot} = \alpha_o + \gamma_t + \sum_{k \neq 2022} \beta_k H_o \times \mathbf{1}[\text{year}(t) = k] + \varepsilon_{ot}, \quad (3)$$

where H_o is one for the high-exposure group, α_o and γ_t are unit-group and quarter effects, cells are weighted by the weighted number of hires or of young employees, and standard errors are clustered by SOC 3-digit group. Each β_k is reported re-centred on the mean of the 2015–2022 coefficients, with an interval from the full covariance matrix. The pre-period then averages zero and the years from 2023 read as departures from it. The difference-in-differences replaces the year indicators with a single indicator for the post period, in variants that add major-group by quarter effects or use interview counts in place of weights and, for adverts, add a high-exposure linear trend as a control. This design also applies to log monthly adverts by SOC 2020 unit group and to Poisson regressions of employee counts over the complete grid of unit groups and quarters. Graduate entrants’ destinations are indicators for starting in each group, estimated with equation (2). The figure draws one regression, of an indicator equal to one for a start in the high group, minus one for a start in the low group and zero otherwise, on year indicators. Because the two thirds are largely in different broad groups, this design describes the contrast between the most and least exposed occupations without by itself separating the technology from anything else that differs between professional and manual work.

The second design uses the continuous measure and identifies the effect. For young hiring,

$$y_{ot} = \alpha_o + \gamma_{g(o)t} + \beta z_o \times \mathbf{1}[t \geq 2023\text{Q1}] + \varepsilon_{ot}, \quad (4)$$

where z_o is the standardised exposure of unit group o and $\gamma_{g(o)t}$ are quarter effects in the main specification and major-group by quarter effects in the within-group specification, which compares each occupation with its neighbours. Sub-major-group by quarter effects are a further variant. The outcomes are the share aged 22–25 among the unit group’s new hires, weighted by the weighted number of hires, the young new-hire rate, weighted by young employment, and the count of young hires in a Poisson regression, with standard errors clustered by SOC 3-digit group. The coefficient β is the post-2022 change per standard

deviation of exposure. For adverts,

$$\ln A_{om} = \alpha_o + \gamma_{g(o)m} + \beta z_o \times \mathbf{1}[m \geq \text{Jan 2023}] + \varepsilon_{om}, \quad (5)$$

where m indexes months, cells are weighted by the occupation’s mean adverts in 2017–2022, and standard errors are clustered by SOC 2020 minor group. Event-study versions replace the post indicator with year indicators, drawn relative to 2022, to 2019 or to the pre-period mean, as each figure or table states. Pre-trends are handled in two ways. The trend-adjusted estimate first regresses log adverts on z_o interacted with a linear trend over 2017–2022 alone, subtracts the estimated trend from every month, and then estimates the post effect on the detrended series, which is the estimate called net of the pre-2023 trend. The alternative, which adds z_o interacted with a linear trend over the whole window as a control in one regression, is reported where the two differ. Placebo breaks move the post indicator to a date inside the pre-period and estimate on the pre-period alone. For the share of young among hires in the main specification, inference is also carried out by randomisation, in which the measure is permuted across unit groups within major groups, which preserves the contrast between broad groups and destroys the ordering within them, and the observed coefficient is compared with the permuted distribution.

The third design relates young hiring in industry section s to the adoption of language models by employers in that section,

$$y_{st} = \alpha_s + \gamma_t + \beta a_s \times \mathbf{1}[t \geq 2023\text{Q1}] + \delta' x_{st} + \varepsilon_{st}, \quad (6)$$

where a_s is standardised adoption, x_{st} are the section’s log vacancies and log payrolled employment, cells are weighted by employment, by hires or by young employment according to the outcome, and standard errors are clustered by section with a t distribution on 12 degrees of freedom. Randomisation inference over the permutations of adoption across sections is reported alongside. Its assumption is that no other shock of the period cut young hiring in sections in proportion to their adoption. Its controls remove each section’s hiring cycle. A further test asks whether the one section that dominates adoption drives the result.

3.3 Identifying assumptions and threats

The exposure gradient identifies an effect of language models under two assumptions. The first is that, absent the technology, young hiring and adverts in exposed occupations would have moved with those in less exposed occupations of the same broad group. Event studies show whether the groups diverged before 2023, the trend adjustment allows exposed occu-

pations a linear trend of their own, and placebo breaks show what a spurious break looks like. Where a pre-trend exists, as it does for adverts, the paper reports the post-2022 change net of it and treats only that part as evidence.

The second assumption is that no other shock of the period fell on occupations in proportion to their exposure within broad groups. Major-group by time effects remove shocks common to a broad group, such as the long decline of clerical work, but not the end of the pandemic hiring boom in technology, which falls on the IT occupations that the task-based measures rank highest within the professional groups, nor the hiring cycle of call centres, which the observed-exposure index ranks highest within sales and customer service. Every gradient is therefore re-estimated without the IT and customer-service unit groups, with a statement of which estimates survive. Two features of the measures affect both assumptions. Mapping the US measures to British occupations introduces error of the order of the signal in differences between neighbouring unit groups, which attenuates the task-based estimates (section 2.4). Usage data behind the observed-exposure index and the usage share were collected after the release, which means that usage could in principle respond to the outcome. The feasibility rating, which rates task descriptions without usage data, gives a within-group gradient in adverts of the same sign, which is why I read the two together. Pay regressions in section 4.4 apply the gradient of equation (4) to log hourly pay. Appendices F and H describe their own designs.

Two further points fix what the estimates mean. Under the reading that entrants were reallocated from exposed to unexposed occupations, the less exposed occupations are affected too, in which case a contrast between occupations measures the differential and overstates the fall in exposed occupations relative to the aggregate. The before-and-after statistics of sections 4.1 to 4.5 give the aggregate. The forces that coincided with the diffusion, the wage floor and migration, act mainly on the elementary jobs that took in the entrants, as section 5 shows in examining what each accounts for.

4 Empirical Results

This section reports the results in two parts. The first establishes, in five steps, that young people entered different jobs after 2022 while their employment and pay moved little. It covers the types of job that new hires aged 22–25 start, the entry jobs of graduates, the flows, stocks and extensive margin behind the shift, entry pay and the wage floor, and who entered the elementary jobs. The second part, from section 4.6, asks whether the fall in demand is located where language models are used, with the designs of section 3.2, and finds that it is. It describes the occupations holding the most exposed third of employment and




Figure 1: Elementary and administrative jobs among new hires and all employees aged 22–25, 2015–2025

Annual shares of employees aged 22–25 with less than three months with the employer (panels A and B) and of all employees aged 22–25 (panel C) whose SOC 2010 unit group is in major group 9 (elementary) or major group 4 (administrative and secretarial). Black: LFS person weights, with 95 percent intervals clustered by household within quarter; grey: unweighted. The shaded band is the range of the weighted 2015–2022 annual shares and the dashed line their mean; the vertical line marks 30 November 2022. Source: Quarterly Labour Force Survey, all waves.

those holding the least exposed third side by side, estimates the exposure gradient in online job adverts and in young hiring along the continuous measures, where the effect is identified, turns the gradient into hires, and ends with young hiring by industry and employer adoption.

4.1 Job types of new hires aged 22–25

Figure 1 plots, by calendar year, the share of new hires aged 22–25 whose job is in an elementary occupation, the share in an administrative and secretarial occupation, and the elementary share of all employees aged 22–25, with and without survey weights. Table 1 reports the series with their 2015–2022 mean and range, the three post-2022 values, the first half of 2026, the year-level test and the person-level estimate.

The elementary share of young new hires, 0.187 on average in 2015–2022 with annual values between 0.161 and 0.223, was 0.298 in 2023, 0.236 in 2024 and 0.267 in 2025, all three above the earlier range. The year-level test puts the change at 8.0 percentage points ($p = 0.001$), while the person-level regression gives 7.9 points with a 95 percent interval from 4.3 to 11.5. Without weights the pre-period mean is 0.178 and the three post-2022 values are 0.246, 0.216 and 0.226, a change of 5.1 points ($p = 0.004$). The change is present in both types of response, at 5.7 points among people who answered for themselves ($p = 0.097$ in the year-level test, interval from 0.3 to 11.7 in the person-level regression) and 9.3 points

Table 1: Job types of new hires aged 22–25, 2015–2026

Share of new hires aged 22–25	2015–2022			2023	2024	2025	2026 H1	Year-level test		Person-level	
	Mean	Range	Change					<i>p</i>	change [95% CI]		
<i>A. Weighted (LFS person weights)</i>											
Elementary (major group 9)	0.187	0.161–0.223	0.298	0.236	0.267	0.224	+0.080	0.001	+0.079	[+0.043, +0.115]	
Administrative and secretarial (4)	0.126	0.097–0.149	0.085	0.078	0.078	0.134	-0.046	0.002	-0.049	[-0.071, -0.027]	
Managerial, professional, associate (1–3)	0.355	0.302–0.409	0.304	0.401	0.350	–	-0.004	0.891	+0.004	[-0.034, +0.042]	
Trades, operatives, elementary (5, 8, 9)	0.284	0.248–0.325	0.382	0.342	0.354	–	+0.075	0.001	+0.073	[+0.033, +0.112]	
<i>B. Unweighted</i>											
Elementary (9)	0.178	0.159–0.220	0.246	0.216	0.226	–	+0.051	0.004	+0.048	[+0.020, +0.075]	
Administrative and secretarial (4)	0.126	0.100–0.152	0.070	0.088	0.091	–	-0.043	0.006	-0.044	[-0.064, -0.025]	
<i>C. Elementary share by type of response, weighted</i>											
Personal responses	0.175	0.126–0.207	0.236	0.155	0.305	0.168	+0.057	0.097	+0.060	[+0.003, +0.117]	
Proxy responses	0.191	0.164–0.237	0.331	0.283	0.237	0.259	+0.093	0.001	+0.091	[+0.045, +0.137]	
<i>D. All employees aged 22–25, weighted</i>											
Elementary (9)	0.136	0.122–0.150	0.139	0.139	0.149	–	+0.006	0.348	+0.006	[-0.002, +0.013]	

Note.—Each row is a share of employees aged 22–25 who had been with their employer for less than three months (panels A–C) or of all employees aged 22–25 (panel D), by calendar year. Mean and range refer to the eight annual values of 2015–2022; 2026 H1 is January–June 2026. The year-level change is the coefficient on a 2023–2025 indicator in a regression of the eleven annual shares on that indicator (9 degrees of freedom), with its *p*-value. The person-level change is the coefficient on the same indicator in a regression at the level of the individual interview with quarter-of-year effects and the row’s weights, with a 95 percent interval clustered by household within quarter. The sample contains 6,343 new hires aged 22–25 over 2015–2026, of whom 452 were interviewed in 2022, 244 in 2023, 328 in 2024 and 460 in 2025. Source: Quarterly Labour Force Survey.

among proxy responses. Among first-wave interviews only, which exclude any respondent conditioning on earlier answers, it is 6.0 points ($p = 0.013$; appendix table A3). The range to state is 5 to 8 percentage points, on a base of 19.

The administrative share moved the other way. It was 0.126 in 2015–2022, with a range of 0.097 to 0.149 and a downward drift within the period, and 0.085, 0.078 and 0.078 in 2023–2025, a fall of 4.6 points ($p = 0.002$) that is the same without weights. Allowing for a linear trend over 2015–2025 reduces the estimate to 2.6 points ($p = 0.122$), which makes part of the fall the continuation of the decline of clerical hiring that the earlier years already show. Young new hires in managerial, professional and associate professional jobs kept their share, at 0.355 in 2015–2022 and 0.304, 0.401 and 0.350 afterwards. The counterpart of the elementary rise is therefore the administrative fall together with smaller falls in sales, customer service and caring jobs. Manual jobs, trades, operatives and elementary occupations together, took a rising share of young entrants, from 0.284 to 0.382, 0.342 and 0.354.

The first half of 2026 is different. Among young new hires the elementary share was 0.224 on 199 interviews, inside the range of the first halves of 2015–2022, while the administrative share was 0.134, also inside its range. Three years of elevated values followed by a return are consistent with a level shift that ended and with a prolonged spike. Section 5.6 returns to the question.

Appendix tables A3 and A4 locate the change. It is present among women (3.8 points)

Table 2: Entry jobs of graduates and non-graduates, 2015–2026

	2015–2022				Weighted		Unweighted			
	Mean	Range	2023	2024	2025	2026	Change	<i>p</i>	Change	<i>p</i>
<i>A. New hires 0–3 years out of full-time education, all quarters (2026 = first half)</i>										
Graduates: elementary jobs	0.119	0.091–0.134	0.249	0.176	0.149	0.086	+0.072	0.004	+0.060	0.001
Graduates: groups 1–3	0.534	0.449–0.582	0.387	0.510	0.544	0.490	-0.053	0.175	-0.033	0.279
Graduates: above the exposure median	0.600	0.568–0.651	0.402	0.518	0.570	0.607	-0.103	0.009	-0.101	0.001
Non-graduates: elementary jobs	0.293	0.242–0.331	0.307	0.400	0.332	0.297	+0.053	0.051	+0.041	0.067
Non-graduates: above the exposure median	0.249	0.207–0.278	0.223	0.208	0.197	0.168	-0.040	0.035	-0.029	0.093
All employees, graduates: elementary jobs	0.072	0.065–0.082	0.081	0.095	0.096	0.092	+0.018	0.002	+0.014	0.007
All employees, non-graduates: elementary jobs	0.232	0.214–0.255	0.292	0.279	0.269	0.276	+0.048	<0.001	+0.044	<0.001
<i>B. Situation a year earlier, April–June quarters (2026 = April–June)</i>										
Graduates aged 22–25 who were students: employee in an elementary job	0.049	0.036–0.060	0.074	0.120	0.120	0.060	+0.056	<0.001	+0.035	0.032
same, among employees	0.093	0.062–0.132	0.137	0.189	0.212	0.109	+0.086	0.002	+0.049	0.080
Graduates 0–3 years out who were students: employee in an elementary job	0.060	0.043–0.083	0.132	0.120	0.098	0.050	+0.057	<0.001	+0.032	0.007
same, among employees	0.091	0.062–0.128	0.178	0.167	0.153	0.083	+0.075	0.001	+0.042	0.028
Non-graduates 0–3 years out who were students: employed	0.505	0.433–0.553	0.466	0.368	0.394	0.377	-0.096	0.018	-0.089	0.027
Employees aged 22–25 in elementary jobs: student a year earlier	0.090	0.063–0.116	0.130	0.218	0.184	0.097	+0.087	0.001	+0.043	0.045
Employees aged 22–25 in groups 1–3: same occupation a year earlier	0.637	0.595–0.673	0.670	0.684	0.715	0.686	+0.053	0.015	+0.052	0.013


Note.—Panel A: shares among new hires (and, in the last two rows, among all employees) who left full-time education 0–3 years before the interview, by whether they hold a degree, in all quarters; the 2026 column is January–June. Panel B: shares computed from the retrospective module of the April–June quarters among respondents aged 22–25, or 0–3 years out of education, who report having been a full-time student twelve months earlier; “among employees” divides the elementary workers by the employees among the former students, where the other rows divide by all former students; the last two rows condition on the current job and report the share who were students, or in the same occupation, a year earlier; the 2026 column is April–June 2026. The change is the year-level estimate for 2023–2025 against 2015–2022 (9 degrees of freedom), with its *p*-value, with and without person weights. Source: Quarterly Labour Force Survey.

and larger among men (12.3 points), among graduates (8.1 points) and non-graduates (9.2 points), and among people not in full-time education (8.7 points). It is absent at ages 26–30 and smaller and imprecise at ages 18–21, which places it at the ages at which graduates enter the labour market. Five unit groups account for most of the 8.0-point rise. Kitchen and catering assistants contribute 2.3 points, bar staff 1.6, cleaners 1.1, elementary storage occupations 1.1 and waiters 0.9. The administrative fall is spread across many unit groups, with receptionists and book-keeping clerks the largest. Figure A1 in the appendix plots the share of young people among new hires in each of the nine major groups. It falls in professional and administrative groups after 2022 and rises in the elementary group, with no movement in trades, operatives, caring or sales.

4.2 Graduates

Because the hypothesis of section 2 concerns graduates in particular, table 2 examines them with two measures, the first of which figure 2 plots.

The first measure compares graduates and non-graduates at the same career stage, 0–3 years after leaving full-time education. Among graduate new hires at that stage, the elementary share was 0.119 in 2015–2022 (range 0.091–0.134) and 0.249, 0.176 and 0.149 in 2023–2025, a rise of 7.2 points ($p = 0.004$; 6.0 points unweighted, $p = 0.001$). The



Shares by calendar year among people who left continuous full-time education 0-3 years before the interview, by highest qualification (degree or equivalent, or not): new hires (employees with under three months with the employer; A, B, D, E) and all employees (C, F), whose SOC 2010 unit group is in major group 9 (elementary) or above the median of DfE (2023) language-model exposure over weighted 2015-2022 employment. Black: LFS person weights, with 95 percent intervals clustered by household within quarter; grey: unweighted. Shaded band: range of the weighted 2015-2022 annual shares; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey (ONS), all waves.

Figure 2: Elementary and exposed jobs among early-career new hires and employees, graduates and non-graduates, 2015–2025

Annual shares among people who left continuous full-time education 0–3 years before the interview, by highest qualification: new hires (employees with less than three months with the employer; panels A, B, D, E) and all employees (panels C, F) whose SOC 2010 unit group is in major group 9 or above the median of the Department for Education language-model exposure score over weighted 2015–2022 employment. Black: LFS person weights, with 95 percent intervals clustered by household within quarter; grey: unweighted. Shaded band: range of the weighted 2015–2022 annual shares; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey, all waves.

share starting a job above the median of the ability-based measure fell from 0.600 to 0.402, 0.518 and 0.570, while the share in professional jobs fell by 5.3 points without reaching significance. Section 4.6 gives the same destinations by the observed-exposure index. Among non-graduates at the same stage the elementary share rose by 5.3 points ($p = 0.051$) from a base of 0.293. Both series were back inside their first-half ranges in 2026. The stock moved less than the flow and did not revert, as among all employed graduates 0–3 years out the elementary share rose from 0.072 to 0.081, 0.095 and 0.096 and was 0.092 in the first half of 2026, above its 2015–2022 range.

The second measure uses the retrospective module and does not depend on the tenure variable or on the hiring rate. Among people aged 22–25 with a degree who were full-time students a year earlier, the share who were now employees in an elementary job was 0.049

in 2015–2022 (range 0.036–0.060) and 0.074, 0.120, 0.120 and 0.060 in 2023–2026 (5.6 points for 2023–2025, $p < 0.001$; 3.5 points unweighted, $p = 0.032$). Among those in employee jobs, the elementary share went from 0.093 to 0.137, 0.189, 0.212 and 0.109. The cells are small, between 116 and 298 former students a year and 6 to 18 elementary workers among them in the post-2022 years. The 2023 value is just above the earlier range and the 2026 value is on its upper bound, which dates the peak of the graduate shift to 2024–2025 on this measure. Seen from inside the elementary jobs, the share of young elementary employees who had been students a year earlier doubled, from 0.090 to 0.130, 0.218 and 0.184, before returning to 0.097 in 2026. Among young employees in professional jobs the share who had been in the same occupation a year earlier rose from 0.637 to 0.670, 0.684 and 0.715, a pattern consistent with fewer entrants. Non-graduates who had been students a year earlier were less often employed (0.505 before, 0.466, 0.368, 0.394 and 0.377 after) and more often still students.

4.3 Flows, stocks and the extensive margin

Three further pieces of evidence show that the shift is a shift in entry. First, among all employees aged 22–25 the elementary share did not change, at 0.136 in 2015–2022 and 0.139, 0.139 and 0.149 afterwards (table 1, panel D). Second, the two-quarter longitudinal files, which follow each household across consecutive quarters, show no change in the quarterly rate at which young elementary employees left their job, 0.179 before and 0.186, 0.166 and 0.168 after, but a fall in the rate at which those employees moved into a professional job, from 0.034 to 0.023, 0.018 and 0.015 ($p = 0.049$; appendix table A6). Young people who were not employed in one quarter and were employees in the next started in an elementary job 0.201 of the time before and 0.344, 0.291 and 0.284 of the time after ($p < 0.001$). Third, the hiring of the young into elementary jobs rose in absolute terms while total hiring did not.

Hiring shows the same pattern. New hires per person aged 22–25 were 0.0516 a quarter in 2015–2022 and 0.0442, 0.0473 and 0.0486 in 2023–2025, a fall of half a percentage point that the year-level test does not distinguish from noise ($p = 0.211$) and the person-level test does (interval from -0.009 to -0.001 ; table 9). Its composition changed. Elementary hires per young person rose from 0.0095 to 0.0131, 0.0111 and 0.0129 ($p = 0.010$), while administrative hires per young person fell from 0.0065 to 0.0037, 0.0037 and 0.0038 ($p = 0.010$) and hires into professional jobs per young person did not change ($p = 0.311$). Elementary hiring per employee of all ages did not rise, which makes the increase in elementary hiring specific to the young. The rise in the elementary share of young entrants is therefore made of more



Weighted shares of all people aged 22-25 with (A-C) and without (D-F) a degree, by calendar year 2015-2025: ILO unemployed; neither employed nor enrolled on a course (ENROLL); employees with under three months with the employer. Bars: 95% intervals clustered by household within quarter. Shaded band: range of the 2015-2022 annual values; dashed grey line: their mean; vertical line: 30 Nov 2022. Source: Quarterly Labour Force Survey (ONS), all waves, person weights.

Figure 3: Unemployment, inactivity outside education and hiring of people aged 22–25, by degree, 2015–2025

Annual weighted shares of all people aged 22–25 with (panels A–C) and without (D–F) a degree: ILO unemployed; neither employed nor enrolled on a course; and employees with less than three months with the employer, as a share of the age group. Bars: 95 percent intervals clustered by household within quarter. Shaded band: range of the 2015–2022 annual values; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey, all waves, person weights.

elementary hires and fewer administrative hires in about equal measure, with professional hires unchanged.

Table 3, panel A, reports the activity of people aged 22–25. Figure 3 plots unemployment, inactivity and hiring by degree. The employment rate of the age group was 0.756 in 2015–2022 (range 0.745–0.770) and 0.737, 0.731 and 0.738 in 2023–2025, a fall of 2.1 points ($p = 0.002$) that continued into 2026. Without weights the fall for the age group is 0.2 points (table 9). The fall is concentrated among those without a degree, whose employment rate went from 0.721 to 0.691, 0.675 and 0.674, 4.1 points weighted and 2.0 unweighted, while the employment rate of young graduates did not move. The share of the age group neither employed nor enrolled rose from 0.152 to 0.169, 0.180 and 0.174. ILO unemployment rose by half a point and is within its earlier range.

Table 3: Activity and pay of people aged 22–25, 2015–2026

	2015–2022			2023	2024	2025	2026 H1	Weighted		Unweighted	
	Mean	Range	Change					<i>p</i>	Change	<i>p</i>	
<i>A. People aged 22–25</i>											
Employed	0.756	0.745–0.770	0.737	0.731	0.738	0.724	-0.021	0.002	–	–	
Employed, with a degree	0.817	0.792–0.837	0.804	0.814	0.826	0.828	-0.002	0.803	–	–	
Employed, without a degree	0.721	0.703–0.740	0.691	0.675	0.674	0.652	-0.041	<0.001	–	–	
ILO unemployed	0.063	0.055–0.073	0.064	0.069	0.071	0.081	+0.005	0.234	-0.001	0.878	
Neither employed nor enrolled	0.152	0.144–0.164	0.169	0.180	0.174	0.185	+0.022	0.001	+0.008	0.313	
New hires per person	0.052	0.038–0.058	0.044	0.047	0.049	0.039	-0.005	0.211	-0.005	0.173	
in elementary jobs	0.010	0.007–0.011	0.013	0.011	0.013	0.009	+0.003	0.010	+0.001	0.118	
in administrative jobs	0.007	0.004–0.008	0.004	0.004	0.004	0.005	-0.003	0.010	-0.003	0.015	
in groups 1–3	0.018	0.015–0.022	0.013	0.019	0.017	0.014	-0.002	0.311	-0.001	0.519	
Living with parents	0.487	0.443–0.539	0.524	0.537	0.535	0.536	+0.045	0.042	+0.074	0.067	
<i>B. Log real hourly pay of employees aged 22–25</i>											
Mean, all employees aged 22–25	2.218	2.153–2.280	2.235	2.285	2.276	2.344	+0.047	0.146	+0.051	0.126	
Degree premium (degree minus no degree)	0.208	0.173–0.236	0.139	0.153	0.180	–	-0.050	0.006	–	–	
Elementary minus groups 1–3	-0.417	-0.482–0.372	-0.350	-0.340	-0.371	–	+0.063	0.027	–	–	
New hires minus tenure of a year or more	-0.047	-0.110–0.003	-0.063	-0.101	-0.058	–	-0.027	0.242	–	–	
	Share paid within 5% of the minimum		2019	2020	2021	2022	2023	2024	2025	2026 H1	
	All employees aged 22–25		0.122	0.083	0.122	0.117	0.169	0.153	0.168	0.166	
	With a degree		0.079	0.042	0.088	0.063	0.136	0.107	0.144	0.122	


Note.—Panel A: annual shares of all people aged 22–25 (first five rows and the last row) and new hires per person aged 22–25 by the job type of the hire. Panel B: annual means of log real hourly pay (2015 prices) of employees aged 22–25 with derived hourly pay, income weights, and differences of means between groups. The change is the year-level estimate for 2023–2025 against 2015–2022, with its *p*-value, with and without weights. The lower table gives the share of employees aged 22–25 whose hourly pay lies within 5 percent of the National Minimum Wage or National Living Wage applicable to their age in the month of the interview. Source: Quarterly Labour Force Survey.

4.4 Entry pay and the wage floor

Panel B of table 3 and figure 4 describe prices. The premium of a degree in the hourly pay of employees aged 22–25 was 0.208 log points in 2015–2022 (range 0.173–0.236) and 0.139, 0.153 and 0.180 in 2023–2025, a fall of 5.0 log points ($p = 0.006$). Between an elementary job and a professional job at these ages the pay gap narrowed by 6.3 log points, from -0.417 to -0.350 , -0.340 and -0.371 ($p = 0.027$), whereas the gap between a new hire and an employee with a year or more of tenure did not change. In the Annual Survey of Hours and Earnings, which covers all ages, the median hourly pay of elementary occupations rose by 0.226 log points between 2022 and 2025 against 0.164 in professional occupations.

Part of the compression comes from the wage floor. The National Living Wage rose from 9.50 pounds in the year from April 2022 to 10.42 in April 2023, 11.44 in April 2024, when its coverage was extended from age 23 to age 21, and 12.21 in April 2025. The share of employees aged 22–25 paid within 5 percent of the minimum applicable to their age was 0.117 in 2022 and 0.169, 0.153 and 0.168 in 2023–2025, and among young graduates from 0.063 to 0.136, 0.107 and 0.144. A graduate who took a bar job in 2024 gave up less relative

Pay of employees aged 22-25 by job type, degree and occupational LLM exposure, UK, 2015-2025



Employees aged 22-25 with derived hourly pay (LFS waves 1 and 5), income weights; real pay at 2015 prices (CPI). A, B: differences of annual means of log real hourly pay between SOC 2010 major group 9 and groups 1-3, and between degree holders and others; bars: 95% intervals from household-clustered standard errors of the two means. C: coefficients of DfE (2023) language-model exposure (z) x year in a regression of log real pay on unit-group, quarter and age effects, relative to their 2015-2022 mean; 95% intervals clustered by SOC 3-digit group. Shaded band: 2015-2022 range; dashed grey line: its mean; vertical line: 30 Nov 2022. Source: Quarterly Labour Force Survey (ONS).

Figure 4: Pay of employees aged 22–25 by job type, degree and occupational exposure, 2015–2025

Employees aged 22–25 with derived hourly pay (LFS waves 1 and 5), income weights, real pay at 2015 prices. Panels A and B: differences between annual means of log real hourly pay of employees in SOC 2010 major group 9 and in groups 1–3, and between degree holders and others, with 95 percent intervals from household-clustered standard errors of the two means. Panel C: coefficients on the Department for Education language-model exposure score (standardised) interacted with calendar year, in a regression of log real pay on unit-group, quarter and age effects, relative to the 2015–2022 mean, with 95 percent intervals clustered by SOC 3-digit group. Shaded band: range of the 2015–2022 values; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey.

to a professional entry job than one who did so in 2019. Section 5.3 weighs the floor against the technology.

4.5 Who entered the elementary jobs

Two compositional facts frame the attribution. Among UK-born new hires aged 22–25 the elementary share went from 0.171 to 0.276, 0.239 and 0.280 ($p < 0.001$), whereas among the foreign-born it did not change (appendix table A5). At the same time the foreign-born flowed into the same jobs at other ages. Among elementary new hires aged 26–64 the UK-born share fell from 0.670 to 0.604, 0.576 and 0.575, while foreign-born employees in elementary jobs rose from 527 thousand in 2021 to 766 thousand in 2025. The share of young elementary new hires living with their parents was 0.532 in 2015–2022 and 0.629 and 0.600 in 2024 and 2025 ($p = 0.115$ weighted, $p = 0.046$ unweighted), while the share of all 22–25-year-olds living with parents rose from 0.487 to 0.524, 0.537 and 0.535 (table 3). Young natives, who were more often at home when they did so, and recent migrants entered elementary jobs together after 2022. Section 5.4 weighs migration against the technology.

4.6 High- and low-exposure occupations side by side

The first part of this section established that young people entered different jobs after 2022 while their employment and pay moved little. The rest asks whether the demand for labour fell where language models are used. It begins with a description, in which figure 5 puts the two ends of the exposure distribution next to each other. SOC 2010 unit groups are cut into thirds of 2015–2022 employment by the observed-exposure index, after which the top third, 83 unit groups with an employment-weighted mean standardised score of 1.23, is compared with the bottom third, 159 unit groups with a mean score of -0.94 . The largest unit groups in each are listed in the note to table 4. Panels A, B and D plot the coefficients of equation (3), the high-minus-low difference by year re-centred on its pre-period mean. Panel C plots its person-level counterpart and panels E and F the Poisson counterparts described in section 3.2. Table 4 gives the differences-in-differences. Appendix E gives the series in levels and the other cuts of the distribution.

The young were hired less in the high-exposure group after 2022 and more in the low-exposure group. The share aged 22–25 among new hires in high-exposure occupations averaged 0.196 in 2015–2022 and 0.154 in 2023–2025, with 0.129 in 2023. In low-exposure occupations it went from 0.149 to 0.169. New hires aged 22–25 per 100 employees of that age fell in high-exposure occupations from 8.0 to 5.8, with 4.8 in 2023 after 8.0 in 2022, and rose in low-exposure occupations from 7.1 to 7.7. The two differences-in-differences are -0.053 for the share (95 percent interval -0.084 to -0.022) and -2.2 hires per 100 for the rate (-3.7 to -0.8). On interview counts without person weights they are -0.035 ($p = 0.001$) and -1.6 ($p = 0.003$). The stock of young employees in high-exposure occupations did not fall relative to the low-exposure group ($+0.02$ log points, $p = 0.77$), the flow-and-stock pattern of section 4.3 seen by exposure. Total employment in high-exposure occupations, which had risen from 83 in 2015 to 100 in 2022, stood at 104 to 105 in 2023–2025.

Panel C shows where graduate entrants went. Among new hires aged 22–25 with a degree, the share starting in a high-exposure occupation was 0.448 in 2015–2022 and 0.351 in 2023–2025, a change of -9.9 points ($p < 0.001$), while the share starting in a low-exposure occupation rose from 0.228 to 0.323, $+9.5$ points ($p < 0.001$). Without person weights the changes are -7.2 points ($p = 0.001$) and $+7.6$ points ($p < 0.001$). In 2023 graduate entrants were as likely to start in the least exposed third as in the most exposed third (0.36 against 0.35), and in 2024 more likely (0.36 against 0.32), where in 2022 the shares had been 0.24 and 0.44. The shares come from 51 and 43 graduate hires in the two groups in 2023. The low-exposure third includes nurses and medical practitioners as well as elementary occupations, which makes the panel evidence of a move away from exposed work and not, by itself, of a move into elementary jobs.

High- minus low-exposure occupations (top third against bottom third), UK, 2015-2025



Quarterly Labour Force Survey (ONS), person weights; ONS online job adverts. Each panel: event-study coefficient on high × year with unit-group and quarter (adverts: month) effects, re-centred on the mean of the 2015–2022 (adverts 2017–2022) coefficients, shaded 95% interval clustered by SOC 3-digit group (graduates: person level, household clusters; employment: Poisson on weighted counts). Dashed line: 30 Nov 2022.

Figure 5: High- minus low-exposure occupations: young hiring, graduate destinations, adverts and employment, 2015–2025

High and low exposure are the top and bottom thirds of SOC 2010 unit groups by Anthropic observed exposure, cut at the exposure values where cumulative 2015–2022 employment of employees aged 16–64 reaches each third. Each panel plots the coefficient on high × year in an event-study regression of the outcome on year indicators interacted with the high-exposure indicator, re-centred on the mean of the 2015–2022 (adverts 2017–2022) coefficients so that the pre-period averages zero, with its 95 percent interval shaded. Panels A and B: SOC 2010 unit group by quarter cells of the two groups, unit-group and quarter effects, weighted by the cell’s weighted new hires (A) or weighted employees aged 22–25 (B), intervals clustered by SOC 3-digit group. Panel C: new hires aged 22–25 with a degree, person level, outcome 1 for a start in the high group, –1 for a start in the low group and 0 for a start in the middle third, regressed on year indicators alone, person weights, intervals clustered by household within quarter. Panel D: SOC 2020 unit group by month cells, unit-group and month effects, weighted by 2017–2022 mean adverts, intervals clustered by SOC 2020 minor group. Panels E and F: Poisson regressions of weighted employee counts over the complete grid of unit groups and quarters, with unit-group and quarter effects, intervals clustered by SOC 3-digit group. The series in levels are in appendix figure A7, and the first four outcomes under other cuts of the exposure distribution in appendix figure A8. Dashed line: 30 November 2022. Source: Quarterly Labour Force Survey; ONS online job advert estimates.

Table 4: High- and low-exposure occupations: young hiring, employment, graduate destinations and adverts, 2015–2026

Outcome	High exposure (top third)		Low exposure (bottom third)		High $\times$ post (2023Q1 onward)	
	2015–2022	2023–2025	2015–2022	2023–2025	Weighted	Unweighted
Share aged 22–25 among new hires	0.196	0.154	0.149	0.169	-0.053 [-0.084, -0.022], $p=0.001$	-0.035 [-0.054, -0.015], $p=0.001$
New hires aged 22–25 per 100 employees aged 22–25	7.97	5.78	7.14	7.66	-2.23 [-3.66, -0.79], $p=0.003$	-1.64 [-2.71, -0.56], $p=0.003$
Employees aged 22–25 (Poisson, log points)	97.2	95.8	103.6	99.3	+0.024 [-0.134, +0.181], $p=0.768$	+0.028 [-0.129, +0.185], $p=0.729$
Employees aged 16–64 (Poisson, log points)	91.5	104.3	101.8	100.5	+0.141 [+0.032, +0.250], $p=0.011$	+0.164 [+0.057, +0.272], $p=0.003$
Log online adverts (January 2017 to August 2026; post = January 2023 onward), month effects					-0.374 [-0.546, -0.202], $p < 0.001$	-0.312 [-0.407, -0.216], $p < 0.001$
major-group $\times$ month effects					-0.282 [-0.506, -0.058], $p=0.014$	-0.282 [-0.421, -0.143], $p < 0.001$
month effects, without the twelve IT unit groups (weighted)					-0.300 [-0.421, -0.178], $p < 0.001$	
<i>Adverts, trend checks (weighted)</i>					Month effects	Major-group $\times$ month
High $\times$ linear trend estimated on 2017–2022 (per year)					-0.047 [-0.073, -0.021], $p=0.001$	-0.043 [-0.071, -0.016], $p=0.003$
High $\times$ post with a high $\times$ linear trend over the whole window as a control					-0.125 [-0.303, +0.053], $p=0.166$	-0.137 [-0.306, +0.033], $p=0.112$
<i>Graduate new hires aged 22–25: share starting in the group, 2023–2025 against 2015–2022 (person level, quarter-of-year effects)</i>						
Destination	2015–2022	2023–2025			Weighted	Unweighted
Starting in a high-exposure occupation	0.448	0.351			-0.099 [-0.150, -0.047], $p < 0.001$	-0.072 [-0.115, -0.029], $p=0.001$
Starting in a low-exposure occupation	0.228	0.323			+0.095 [+0.046, +0.144], $p < 0.001$	+0.076 [+0.035, +0.116], $p < 0.001$

Note.—Groups as in figure 5, the top and bottom thirds of 2015–2022 employment by the observed-exposure index; the largest unit groups by employment in the top third are administrative occupations not elsewhere classified, sales account managers, secondary teachers, book-keepers, programmers, customer-service occupations, business and financial project management professionals and financial managers, and in the bottom third care workers, nurses, kitchen assistants, cleaners, nursing auxiliaries, lorry drivers, waiters and medical practitioners. Columns 1–4: means of the annual series over 2015–2022 and 2023–2025 (employment rows: index, 2022 = 100); in the lower panel, the pooled person-weighted share of graduate hires starting in the row’s group over 2015–2022 and 2023–2025. Columns 5–6: coefficient on high $\times$ post (2023Q1 onward) in regressions over 2015Q1–2026Q2 on SOC 2010 unit group by quarter cells of the two groups, with unit-group and quarter effects, weighted by the cell’s weighted hires (share) or weighted young employees (rate), and unweighted on interview counts; employment rows: Poisson regressions of the weighted count in thousands (column 5) or the interview count (column 6) over the complete grid of unit groups and quarters. Adverts: log monthly adverts by SOC 2020 unit group, January 2017 to August 2026, on high $\times$ post (January 2023 onward) with unit-group and month effects, weighted by 2017–2022 mean adverts (column 5) and unweighted (column 6); the second adverts row uses major-group by month effects, weighted and unweighted; the third gives the weighted estimate with month effects without the twelve IT unit groups of SOC 2020. The trend panel, weighted, gives the high $\times$ linear-trend coefficient estimated on 2017–2022 alone and the post effect with a high $\times$ linear trend over the whole window as a control, with month effects and with major-group by month effects. Lower panel: person-level regressions among new hires aged 22–25 with a degree of an indicator for starting in the group on an indicator for 2023Q1 onward, with quarter-of-year effects, with person weights and without. Brackets: 95 percent intervals clustered by SOC 3-digit group (cells) or by household within quarter (persons). Source: Quarterly Labour Force Survey; ONS online job advert estimates.

Online adverts fell in both groups after 2022 and further in the exposed one, to 49 and 50 percent of their 2022 level in 2024 and 2025 against 71 and 72 percent, a difference-in-differences of -0.37 log points, -0.30 without the twelve IT unit groups and -0.28 within broad groups ($p = 0.014$). Adverts in the exposed group had been falling relative to the other by about 0.05 log points a year over 2017–2022. With a high-exposure linear trend over the whole window as a control, the post-2022 contrast is -0.13 ($p = 0.17$) with month effects and -0.14 ($p = 0.11$) within broad groups. On the ability-based cut the hiring contrasts are similar, -0.039 for the share of young among hires ($p = 0.007$) and -2.1 per 100 for the young hire rate ($p = 0.001$). The shift of graduate entrants into low-exposure occupations is also similar ($+8.1$ points, $p = 0.001$), while the shift out of high-exposure occupations is smaller and not distinguishable from zero (-3.8 points, $p = 0.14$). The adverts contrast is -0.38 . By observed exposure, the two thirds are largely in different broad groups. Within the six major groups the two share, the hiring contrasts are -0.043 ($p = 0.19$) and -1.9 ($p = 0.23$). The comparison therefore describes the divergence, while the identification is in the gradient of the next subsection.

4.7 The exposure gradient in adverts and young hiring

Along the continuous measures, occupations are compared by their exposure before and after the release, within broad groups in the main design. Table 5 reports estimates of equation (5) for log adverts and table 6 estimates of equation (4) for young hiring, for the five exposure measures, with quarter (or month) effects and with major-group by quarter effects.

The ability-based measure. With the ability-based measure, the estimate of equation (4) with quarter effects is -0.0209 per standard deviation, with a 95 percent interval from -0.0335 to -0.0084 , over 10,576 cells and 47,517 hires (table 6). Between the top and bottom exposure quintiles the implied change is -0.057 , or 37 percent of the 2022 mean share of 0.152. The young new-hire rate of an occupation, hires aged 22–25 over employees aged 22–25, falls by 0.011 per standard deviation. A Poisson regression of the count of young hires gives -0.11 log points. The annual coefficients relative to their 2015–2022 mean are near zero before 2023 and -0.033 , -0.014 and -0.020 in 2023–2025.

Because 0.73 of the variance of the ability-based measure is between the nine major groups, equation (4) compares the hiring of young people in professional and administrative occupations with their hiring in elementary and manual occupations. Adding major-group by quarter effects, which remove any shock common to all occupations in a broad group, reduces the coefficient to -0.009 ($p = 0.47$). With sub-major-group by quarter effects it is zero, and without the elementary occupations -0.009 ($p = 0.15$). Permuting the measure

Table 5: Online job adverts by occupation and exposure to language models, January 2017 to August 2026

Measure	Unit-group and month effects			Plus major-group $\times$ month effects		
	Coef.	95% CI	p	Coef.	95% CI	p
<i>A. All unit groups</i>						
DfE language-model exposure	-0.136	[-0.194, -0.077]	<0.001	-0.083	[-0.164, -0.003]	0.043
Anthropic observed exposure	-0.170	[-0.239, -0.101]	<0.001	-0.155	[-0.231, -0.080]	<0.001
Microsoft AI applicability	-0.135	[-0.229, -0.041]	0.005	-0.095	[-0.209, +0.019]	0.100
Eloundou GPT-4 feasibility	-0.188	[-0.256, -0.121]	<0.001	-0.227	[-0.321, -0.133]	<0.001
Claude.ai GB usage per worker	-0.094	[-0.177, -0.012]	0.025	-0.026	[-0.092, +0.040]	0.431
<i>B. Major groups 1–3 only</i>						
DfE language-model exposure	-0.163	[-0.279, -0.048]	0.007	-0.155	[-0.254, -0.056]	0.003
Anthropic observed exposure	-0.182	[-0.286, -0.079]	0.001	-0.181	[-0.263, -0.099]	<0.001
Microsoft AI applicability	-0.157	[-0.342, +0.028]	0.095	-0.153	[-0.313, +0.006]	0.059
Eloundou GPT-4 feasibility	-0.274	[-0.402, -0.145]	<0.001	-0.265	[-0.378, -0.152]	<0.001
Claude.ai GB usage per worker	-0.089	[-0.236, +0.059]	0.233	-0.088	[-0.217, +0.040]	0.175

Note.—Coefficients on the standardised exposure measure interacted with an indicator for January 2023 onward, in regressions of log monthly online job adverts by SOC 2020 unit group (ONS Textkernel data, UK totals, 115 months; suppressed cells dropped) on unit-group and month effects (columns 1–3) or unit-group and major-group by month effects (columns 4–6). Cells are weighted by the unit group’s mean adverts in 2017–2022; 95 percent intervals are clustered by SOC 2020 minor group (104 clusters). Panel B restricts the sample to major groups 1–3. The US measures are moved to SOC 2010 through the O*NET–ISCO-08 and ISCO-08–SOC 2010 crosswalks and to SOC 2020 through the ONS relationship table. Source: ONS online job advert estimates; exposure measures as described in section 2.4.

across unit groups within major groups, which preserves the between-group contrast and destroys the within-group ordering, produces a distribution of placebo coefficients centred on -0.017 , against which the observed coefficient is not unusual (randomisation $p = 0.18$, table 6). Figure A1 shows the same thing without regressions. The youth share of hiring fell in the professional and administrative groups and rose in the elementary group, a contrast that the regression records and that restates section 4.1. Four other measures in the literature agree with it because they encode the same ordering of broad groups (appendix table A2). The patent-based measure of Webb (2020), which classifies manual and clerical occupations as the exposed ones, gives the opposite sign for the same reason.

Adverts. The task-based measures described in section 2.4 place 0.34 to 0.63 of their variance between major groups and retain, within groups, an ordering that reflects which tasks language models perform. Within the professions the observed-exposure index ranks programmers, web developers and IT analysts highest, and within sales and customer service call-centre staff and market research interviewers. Along every measure, adverts fell after 2022 in exposed occupations relative to others, by 0.09 to 0.19 log points per standard deviation with unit-group and month effects (table 5, column 1). The measures part company


Table 6: Young hiring by occupation and exposure to language models, 2015Q1–2026Q2

Measure	Share aged 22–25 among new hires			Young new-hire rate		Young hires (Poisson)	
	Main	Major $\times$ quarter	RI p	Main	Major $\times$ quarter	Main	Major $\times$ quarter
DfE language-model exposure	-0.0209 [-0.033, -0.008] $p=0.001$	-0.0091 [-0.034, +0.016] $p=0.473$	0.184	-0.0114 [-0.018, -0.005] $p=0.001$	-0.0084 [-0.020, +0.003] $p=0.155$	-0.1105 [-0.219, -0.002] $p=0.047$	-0.0244 [-0.193, +0.144] $p=0.776$
Anthropic observed exposure	-0.0238 [-0.035, -0.012] $p < 0.001$	-0.0167 [-0.034, +0.000] $p=0.054$	0.039	-0.0107 [-0.016, -0.006] $p < 0.001$	-0.0086 [-0.016, -0.001] $p=0.031$	-0.1523 [-0.262, -0.043] $p=0.006$	-0.1207 [-0.247, +0.006] $p=0.062$
Microsoft AI applicability	-0.0194 [-0.033, -0.006] $p=0.006$	-0.0146 [-0.030, +0.001] $p=0.071$	0.086	-0.0087 [-0.016, -0.002] $p=0.016$	-0.0080 [-0.016, +0.000] $p=0.054$	-0.1143 [-0.221, -0.007] $p=0.036$	-0.0762 [-0.199, +0.047] $p=0.225$
Eloundou GPT-4 feasibility	-0.0232 [-0.034, -0.013] $p < 0.001$	-0.0142 [-0.035, +0.007] $p=0.182$	0.087	-0.0110 [-0.016, -0.006] $p < 0.001$	-0.0087 [-0.018, +0.001] $p=0.064$	-0.1349 [-0.231, -0.039] $p=0.006$	-0.0697 [-0.228, +0.088] $p=0.387$
Claude.ai GB usage per worker	-0.0178 [-0.031, -0.004] $p=0.010$	-0.0034 [-0.020, +0.013] $p=0.687$	0.134	-0.0083 [-0.015, -0.001] $p=0.019$	-0.0012 [-0.009, +0.007] $p=0.768$	-0.0739 [-0.165, +0.017] $p=0.110$	+0.0127 [-0.087, +0.113] $p=0.804$

Note.—Coefficients on the standardised exposure measure interacted with an indicator for 2023Q1 onward. Columns 1–2: share aged 22–25 among the new hires of the SOC 2010 unit group in the quarter, cells weighted by the weighted number of hires (10,576 cells, 47,517 hires), with unit-group and quarter effects (main) or unit-group and major-group by quarter effects. Column 3: randomisation-inference p -value of the main estimate when the measure is permuted across unit groups within major groups (3,000 draws), two-sided around the mean of the permuted distribution. Columns 4–5: new hires aged 22–25 over employees aged 22–25 in the cell, weighted by young employment. Columns 6–7: Poisson regression of the weighted count of new hires aged 22–25. Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey; exposure measures as described in section 2.4.

when broad groups are held fixed. With major-group by month effects the ability-based measure gives -0.083 ($p = 0.043$), the observed-exposure index -0.155 with an interval from -0.231 to -0.080 , the feasibility rating -0.227 and the usage share -0.026 . Within the professional groups alone the observed-exposure coefficient is -0.181 . The observed-exposure result holds without weights (-0.138), holds when any one major group is left out (between -0.093 and -0.170), and holds within major groups 1, 2, 3 and 7 taken one at a time (appendix tables A8 and A9).

Timing and pre-trends. The timing requires care, because adverts in exposed occupations were already declining relative to their groups before 2023. Figure 6, panel C, shows the observed-exposure event study within major groups with the 2017–2022 trend extrapolated, in which the coefficients are below the trend in 2023–2025 and return towards it in 2026, whereas in panel D, without the IT unit groups, they do not differ from it. Appendix table A8 dates the decline. With 2019 as the reference year, the observed-exposure coefficients are -0.075 in 2020, -0.033 in 2021 and -0.046 in 2022, then -0.144 in 2023, -0.222 in 2024, -0.176 in 2025 and -0.156 in 2026. Exposed occupations lost 1.5 log points of adverts per standard deviation per year relative to their groups over 2017–2022. A placebo break at January 2020 estimated on 2017–2022 alone is -0.060 ($p = 0.001$). Net of the



ONS online job advert estimates (Textkernel), new adverts by SOC 2020 unit group. A: thirds of unit groups by Anthropic observed exposure, weighted by 2017–2022 adverts. C, D: coefficients on standardised observed exposure x year (2019 omitted) in regressions of log monthly adverts with unit-group and major-group x month effects, weighted by 2017–2022 mean adverts, with 95 percent intervals clustered by SOC 2020 minor group; dashed line: exposure x linear trend estimated on 2017–2022 alone, drawn through the mean of the 2017–2022 coefficients and extrapolated. 2026 covers January–August; March 2026 is missing from the source. Dotted line: 30 Nov 2022.

Figure 6: Online job adverts and observed exposure to language models, 2017–2026

New online job adverts by SOC 2020 unit group (ONS Textkernel data, UK totals), January 2017 to August 2026. Panel A: adverts by third of unit groups ranked by observed exposure (thirds weighted by 2017–2022 adverts), twelve-month moving average, December 2022 = 100. Panel B: adverts by SOC 2020 major group, same index. Panels C and D: coefficients on standardised observed exposure interacted with calendar year, 2019 omitted, in regressions of log monthly adverts on unit-group and major-group by month effects, weighted by 2017–2022 mean adverts, with 95 percent intervals clustered by SOC 2020 minor group; the dashed line is the exposure by linear-trend coefficient estimated on 2017–2022 alone, drawn through the mean of the 2017–2022 coefficients and extrapolated; panel D drops the twelve IT unit groups. 2026 covers January to August; March 2026 is missing from the source. Dotted line: 30 November 2022.

2017–2022 trend, the post-2022 effect along the observed-exposure index is -0.082 with an interval from -0.157 to -0.006 , whereas along the ability-based measure it is $+0.040$, with a placebo break at 2020 larger than its post-2022 effect. The post-2022 step therefore has two parts, a continuation of a decline that began in 2020, which the ability-based measure also captures, and a further fall of about 8 log points per standard deviation after 2022 that only the task-based measures capture and that, as the robustness paragraph below shows, is in the IT occupations.

Young hiring. The young-hiring estimates in table 6 are weaker and less robust. With major-group by quarter effects, the share of young people among an occupation’s new hires

falls by 0.017 per standard deviation of observed exposure ($p = 0.054$), against -0.009 ($p = 0.47$) for the ability-based measure. Permuting the observed-exposure index within major groups gives a randomisation p -value of 0.039 for the quarter-effects estimate, where the ability-based measure gave 0.18 (table 6). The young new-hire rate falls by 0.009 per standard deviation within groups ($p = 0.031$). Appendix table A11 shows what these estimates depend on. Dropping the five customer-service unit groups, which the observed-exposure index ranks highest within sales and customer service, moves the share estimate to -0.009 ($p = 0.30$) and the hire-rate estimate to -0.005 ($p = 0.17$), while dropping the IT and customer-service groups together moves it to -0.004 ($p = 0.68$). Since more than half of the cells contain one or two hires, the sample was also restricted to cells with at least three hires, which gives -0.015 ($p = 0.12$). With cells weighted by employment the estimate is -0.013 ($p = 0.10$). The within-group gradient in young hiring is therefore the right sign on the observed-exposure index and is not established at conventional levels once the call-centre occupations, whose hiring has reasons of its own, are set aside. The paper reports it as consistent with the adverts evidence, without treating it as independent evidence.

Robustness. Along the observed-exposure index, the within-group adverts gradient holds without weights, when any one major group is left out, and within major groups 1, 2, 3 and 7 taken one at a time (appendix tables A8 and A9). Without the twelve IT unit groups of SOC 2020 it is -0.10 log points per standard deviation ($p = 0.001$), but net of the 2017–2022 trend it is -0.01 ($p = 0.77$), against -0.08 with them. The feasibility rating behaves the same way, at -0.17 without the IT groups and -0.01 net of trend. Removing the customer-service groups changes nothing (appendix table A10). The fall in adverts beyond the pre-2023 trend is therefore in the IT occupations, where the end of the technology hiring boom also fell, while outside them exposed occupations continued the relative decline that began in 2020. The young-hiring gradient does not survive the removal of the customer-service groups, nor the restriction of the sample to cells with at least three hires, which gives -0.015 ($p = 0.12$) (appendix table A11). Teleworkability, measured by the ONS ability-to-homework score (Office for National Statistics, 2020), is correlated 0.78 with observed exposure across unit groups and 0.63 to 0.66 with each measure within major groups. With it as a control, the within-group adverts gradient along observed exposure is -0.134 against -0.155 and the young-hiring gradients grow, while the ability-based gradient does not survive, as in Henseke et al. (2026). Because the control accounts for the relative decline of exposed occupations in 2020–2022, the part of the gradient net of trend then depends on the window over which the trend is estimated, between -0.04 and -0.12 . Without the IT groups the within-group estimate holds under the main coding of teleworkability but not under every alternative. Its

part net of trend is imprecise (appendix tables A12, A14 and A13). The mapping of the task-based measures to British occupations is imperfect (section 2.4, appendix A.2), which attenuates their estimates.

4.8 What the gradient implies for young hiring

Table 7 turns the gradient into hires, to ask how many fewer young people were hired in exposed occupations and how much of the fall the gradient accounts for. Panel A sums the weighted young new hires of each unit group into quarterly averages over 2015–2022 and 2023–2025, by exposure and by major group. Young hiring fell from 170 to 152 thousand a quarter, by a tenth, with the fall in exposed occupations, those above the mean of the observed-exposure index. Hires there fell from 64 to 47 thousand, by 28 percent, while hires in unexposed occupations were unchanged in total at 106 thousand, a rise of 9 thousand in elementary hires offsetting a fall of 9 thousand in the other unexposed occupations. By major group, administrative hires of the young fell by 44 percent, sales and customer-service hires by 22 percent and professional hires by 11 percent, while elementary hires rose by 28 percent, from 32 to 41 thousand a quarter. In interview counts, which fell by 46 percent between the two periods as the achieved sample shrank, young hires in exposed occupations fell by 53 percent and in unexposed ones by 42 percent, while elementary hires fell by 32 percent and their share of young hires rose by 4.8 points, against 8.0 points weighted. Panel B applies each gradient to the 2015–2022 hires of each unit group as a proportional change of $\exp(\beta z_o) - 1$, as if young hires moved one for one with the outcome the gradient is estimated on, and sums the implied losses over exposed unit groups. Exposure z_o enters in full, which extends the elasticity estimated within broad groups to the differences in exposure between them. Net of the 2017–2022 trend, the within-group adverts gradient implies a fall of 6 thousand young hires a quarter in exposed occupations, a third of the observed fall there and two thirds of the rise in elementary hires, and between 0.5 and 11 thousand at the bounds of its interval. Its raw counterpart implies 11 thousand, three fifths of the fall, while in young hiring the Poisson gradient within groups implies 9 thousand, half of the fall and the whole of the elementary rise. The young-hire gradient has no detectable pre-2023 trend to net out, +0.0025 log points a year over 2015–2022 with a 95 percent interval from -0.0135 to $+0.0184$. Net of the point estimate, by the two-step method used for adverts, the gradient is -0.135 and implies 10 thousand, 55 percent of the fall, whereas with the trend estimated on 2017–2022 as for adverts it is -0.090 ($p = 0.17$) and implies 7 thousand. With the within-group component of exposure alone, each implied fall is half as large. The demand fall identified within broad groups therefore accounts for three fifths of the fall in

young hiring in exposed occupations, and its part net of trend for a third, a part that is in the IT unit groups and implies less than one thousand without them.

4.9 Employer adoption by industry

The third design asks whether the industries whose employers adopted language models hired fewer young people. In the Business Insights and Conditions Survey, the share of businesses using text generation with large language models averaged 0.273 over the twelve waves in information and communication, 0.183 in professional, scientific and technical services, 0.161 in education and 0.106 in administrative services, against 0.03–0.08 in construction, transport, manufacturing, hospitality and retail. A panel of industry section by quarter cells is formed from the LFS, 2015Q1–2026Q2, and equation (6) is estimated on it.

Table 8 reports the estimates. The young new-hire rate fell by 0.0046 per standard deviation of adoption after 2022 ($p = 0.047$; randomisation $p = 0.107$), and by 0.0042 with the cycle controls ($p = 0.052$). Young people’s share among new hires fell by 0.009 per standard deviation ($p = 0.21$) and their share among employees by 0.001, while at all ages the hire rate did not move. Nor does the elementary share of young hires vary with adoption (0.004, $p = 0.83$), because elementary entry rose across sections, in manufacturing, retail, administrative services, education and professional services. Year by year, the coefficients for the young hire rate are -0.005 in 2023, -0.004 in 2024 and $+0.005$ in 2025, none distinguishable from zero.

One sector drives the estimate. Without information and communication, the section with the highest adoption, whose young hires fell from 226 to 48 unweighted cases between the two periods, the coefficient is -0.0025 ($p = 0.41$). Without information and communication and education together it is $+0.0005$ (appendix table A15). A placebo break at 2020Q1 on 2015–2022 gives -0.0012 ($p = 0.73$). The forces that cut young hiring in information and communication after 2022, which plausibly include the end of the pandemic hiring boom in technology. The next section asks whether the forces that coincided with the diffusion could have produced the same changes.

5 Alternative Explanations and Mechanisms

Three forces acted on the entry labour market as language models diffused. The hiring boom of 2022 ended, the wage floor rose, and migration into elementary jobs rose. This section asks what each accounts for, and how the fall in demand documented in section 4 became the change in the jobs that young people start. Section 5.1 uses the 2008–2011 downturn as

Table 7: Young hiring by exposure and major group, and the reallocation each gradient implies, 2015–2025

A. Young new hires aged 22–25 per quarter, by group					
Group	2015–2022	2023–2025	Change	Percent	
Weighted, thousands					
All scored unit groups	170.2	152.3	-17.8	-10.5	
Exposed (observed exposure above the mean)	64.4	46.6	-17.7	-27.5	
Unexposed	105.8	105.7	-0.1	-0.1	
Elementary (major group 9)	31.7	40.6	+8.9	+28.0	
Unexposed other than elementary	74.6	65.4	-9.2	-12.4	
Administrative (major group 4)	21.8	12.2	-9.6	-44.0	
Professional (major groups 1–3)	60.1	53.7	-6.3	-10.5	
Sales and customer service (major group 7)	22.3	17.4	-4.8	-21.8	
Interview counts, unweighted					
All scored unit groups	159.8	86.0	-73.8	-46.2	
Exposed (observed exposure above the mean)	59.1	27.8	-31.3	-52.9	
Unexposed	100.6	58.2	-42.5	-42.2	
Elementary (major group 9)	28.8	19.6	-9.2	-31.9	
Unexposed other than elementary	72.3	38.7	-33.6	-46.5	
Administrative (major group 4)	20.7	7.3	-13.4	-64.6	
Professional (major groups 1–3)	55.1	31.7	-23.4	-42.5	
Sales and customer service (major group 7)	21.6	10.8	-10.8	-49.8	
B. Fall in exposed young hires implied by each gradient, thousands per quarter					
Gradient	Implied fall (share of their hires)	At the bounds of the 95% interval	Share of the observed fall	Share of the elementary rise	Within-group component only
Adverts within groups, net of the 2017–2022 trend $\beta = -0.082$ [-0.157, -0.006]	6.1 (9.5%)	0.5 to 11.1	35%	69%	3.1 (17%)
the same without the twelve IT groups $\beta = -0.009$ [-0.069, +0.051]	0.7 (1.1%)	-4.3 to 5.3	4%	8%	0.4 (2%)
Adverts within groups, raw post effect $\beta = -0.155$ [-0.231, -0.080]	11.0 (17.1%)	6.0 to 15.4	62%	124%	5.6 (31%)
Young hires (Poisson) within groups $\beta = -0.121$ [-0.247, +0.006]	8.8 (13.7%)	-0.5 to 16.3	50%	99%	4.4 (25%)
Young hires (Poisson) within groups, net of the 2015–2022 trend $\beta = -0.135$ [-0.262, -0.008]	9.7 (15.1%)	0.7 to 17.1	55%	110%	4.9 (28%)
Young hires (Poisson), quarter effects $\beta = -0.152$ [-0.262, -0.043]	10.8 (16.8%)	3.3 to 17.1	61%	122%	5.5 (31%)

Note.—Panel A: weighted new hires aged 22–25 per quarter in SOC 2010 unit groups with an observed-exposure score, averaged over 2015Q1–2022Q4 and 2023Q1–2025Q4; exposed unit groups have a standardised observed-exposure score above zero, the employment-weighted mean. The second block of panel A repeats the sums on interview counts, the unweighted number of new hires aged 22–25 per quarter; the achieved sample of the survey fell after 2022, so every count falls and the comparison is between groups. Panel B: each gradient β , in log points per standard deviation, is applied to the 2015–2022 young hires H_o of each unit group as $H_o(\exp(\beta z_o) - 1)$, as if young hires moved one for one with the outcome the gradient is estimated on; the implied fall sums these over exposed unit groups (in parentheses, as a share of their 2015–2022 hires), is recomputed at the bounds of the gradient’s 95 percent interval, and is compared with the observed fall in exposed young hires and the observed rise in elementary young hires. The exposure z_o enters in full, which extends the elasticity estimated within major groups to the differences in exposure between them; the last column uses only the within-major-group component of z_o , demeaned with 2015–2022 young-hire weights. The computation sets the change at zero exposure to zero. The gradients are the adverts estimates of tables 5 and A10 and the Poisson estimates of table 6; the Poisson row net of trend holds the 2015–2022 exposure $\times$ linear-trend coefficient of young hires, estimated by Poisson with the same effects and clustering, fixed as an offset while the post effect is estimated over 2015Q1–2026Q2, the two-step method of the adverts rows. Source: Quarterly Labour Force Survey; ONS online job advert estimates.

Table 8: Employer adoption of language models by industry and the hiring of young people, 2015Q1–2026Q2

Outcome	Mean adoption, September 2023 to June 2026		June 2026 adoption
	No controls	With vacancies and payroll	No controls
Young new-hire rate (hires aged 22–25 over employees aged 22–25)	-0.0046 [-0.009, -0.000] $p=0.047$, RI 0.107	-0.0042 [-0.008, +0.000] $p=0.052$, RI 0.160	-0.0046 [-0.010, +0.000] $p=0.061$, RI 0.123
Share aged 22–25 among new hires	-0.0089 [-0.024, +0.006] $p=0.214$, RI 0.164	-0.0095 [-0.024, +0.005] $p=0.165$, RI 0.152	-0.0103 [-0.023, +0.003] $p=0.110$, RI 0.123
Elementary share of new hires aged 22–25	+0.0039 [-0.034, +0.042] $p=0.827$, RI 0.872	+0.0067 [-0.032, +0.045] $p=0.710$, RI 0.750	+0.0118 [-0.027, +0.050] $p=0.517$, RI 0.613
Share aged 22–25 among employees	-0.0014 [-0.004, +0.001] $p=0.260$, RI 0.302	-0.0016 [-0.004, +0.001] $p=0.210$, RI 0.295	-0.0008 [-0.004, +0.002] $p=0.538$, RI 0.553
New-hire rate, all ages	-0.0007 [-0.002, +0.001] $p=0.306$, RI 0.326	-0.0002 [-0.001, +0.001] $p=0.701$, RI 0.718	-0.0002 [-0.002, +0.002] $p=0.866$, RI 0.849

Note.—Coefficients on standardised adoption, the weighted share of businesses in the SIC 2007 section using text generation with large language models (Business Insights and Conditions Survey, mean of the twelve waves from September 2023 to June 2026, or the June 2026 wave), interacted with an indicator for 2023Q1 onward, in regressions of the section by quarter outcome on section and quarter effects, with and without the section’s log vacancies and log payrolled employees as controls. Thirteen sections with an adoption value, 598 cells. Cells are weighted by employment, or by young employment and young hires for the young-hire outcomes. Brackets: 95 percent intervals clustered by section (t with 12 degrees of freedom); RI: randomisation-inference p -value from 5,000 permutations of adoption across sections. Source: Quarterly Labour Force Survey, BICS, Vacancy Survey, PAYE Real Time Information.


a benchmark for what a hiring slowdown alone does to the entry jobs of the young. Sections 5.2 to 5.5 gather the mechanisms: the elementary detour that took in the entrants, the wage floor that lowered its cost, the migration that raised the supply of labour to the same jobs, and the layers above entry, which were neither drawn in nor, on the evidence so far, displaced.

5.1 The hiring cycle: the 2008–2011 downturn as a benchmark

The changes of section 4.1 coincided with the end of the sharpest hiring boom in the history of the Vacancy Survey. To see what a hiring slowdown does to the entry jobs of the young when no new technology is diffusing, the paper uses the 2008–2009 recession and its aftermath. The LFS series were extended back to 2006 so that the entry-job, employment and hiring statistics of the paper can be computed through that episode. The three years from the trough in output, 2009–2011, and the three years from the peak in unemployment, 2011–2013, are each compared with 2006–2007, using the same definitions and the person-level test of equation (2).⁶

⁶Occupation is coded in SOC 2000 up to 2010 and in SOC 2010 from 2011, two classifications whose nine major groups correspond, although group 1 adds directors to its title in SOC 2010 and group 6, personal

Employer adoption of LLM text generation and the hiring of people aged 22-25, by SIC 2007 section, UK




Sections (SIC 2007 letters) covered by the ONS Business Insights and Conditions Survey: C manufacturing, E water, F construction, G retail, H transport, I accommodation and food, J information, L real estate, M professional, N administrative and support, P education, Q health, R arts, S other services. Adoption: weighted share of businesses using text generation with LLMs, mean of twelve BICS waves. Hiring: Quarterly Labour Force Survey, employees aged 16-64, new hires with under three months with the employer, person weights; points sized by the number of young hires.

Figure 7: Employer adoption of language models and the hiring of people aged 22–25, by industry section

Each point is an SIC 2007 section covered by the Business Insights and Conditions Survey: C manufacturing, E water supply, F construction, G retail, H transport, I accommodation and food, J information and communication, L real estate, M professional services, N administrative and support services, P education, Q health and social work, R arts and recreation, S other services. Horizontal axis: weighted share of businesses using text generation with large language models, mean of the twelve waves from September 2023 to June 2026. Panel A: change in the share aged 22–25 among the section’s new hires, 2023–2025 minus 2015–2022. Panel B: log change in new hires aged 22–25 per employee aged 22–25 over the same periods. Points are sized by the number of young hires. Source: Quarterly Labour Force Survey, employees aged 16–64, person weights; BICS.

The two episodes differ in size. Vacancies fell from an average of 662 thousand in 2007 to 450 thousand in 2009, a ratio of 0.68 and the lowest annual average in the survey’s record, whereas at 749 thousand the 2025 level is 0.60 of the 2022 average and 0.91 of the 2019 level. The share of 22–25-year-olds who were ILO unemployed rose by 3.2 percentage points in 2009–2011 and by 0.5 in 2023–2025. The employment rate of the age group fell by 4.6 points in the recession, with a 95 percent interval from 3.9 to 5.2, against 2.1 in 2023–2025. New hires per young person fell from about 0.06 a quarter to 0.045 in the recession and from 0.052 to 0.047 in 2023–2025. First interviews’ share of the sample rose from 0.233 in 2015–2022 to 0.295 in 2023–2025. On first interviews alone the employment rate of the age group fell by 4.1 points in 2009–2011 and 3.3 points in 2023–2025, which places the employment contrast

service in SOC 2000, becomes caring, leisure and other service. Appendix table A29 uses the January–March 2011 file, which carries both codes, to show that 0.965 of employees in the SOC 2000 elementary group are in the SOC 2010 elementary group and 0.950 of the administrative group stay administrative, while groups 1–3 are reshuffled among themselves. Every file from 2011 to 2014 also carries the SOC 2000 code mapped by the ONS. On SOC 2000 codes throughout, the estimates from 2011 for graduates’ elementary and low-paying shares are of the same size, while the fall in the administrative share is larger, 4.8 points in 2011–2013 against 4.0 (appendix table A28, panel B). The weights are those of each file. Yearly counts of young new hires are 795, 853, 754, 818 and 816 in 2009–2013, against the 2023–2025 counts in the note to table 1.



Weighted annual shares, Quarterly Labour Force Survey (ONS), all waves; new hires are employees with under three months with the employer; elementary is major group 9 of SOC 2000 (to 2010) and of SOC 2010 (from 2011, dotted line). Shaded: April 2008 to September 2009 (recession quarters); dashed line: 30 Nov 2022. Bars: 95% intervals clustered by household within quarter. Weights PWT14 to mid-2011, PWT18 to 2018, PWT24 from 2019.

Figure 8: Entry jobs and employment of the young through the 2008–2009 recession and after 2022, 2006–2025

Weighted annual shares, all waves. Panel A: elementary share of new hires aged 22–25; panel B: the same among new hires with a degree; panel C: employment rate of all people aged 22–25. Elementary is major group 9 of SOC 2000 up to 2010 and of SOC 2010 from 2011 (dotted line). Shaded area: April 2008 to September 2009; dashed line: 30 November 2022. Bars: 95 percent intervals clustered by household within quarter. Weights are the person weights of each file. The April–June 2007 file is not available. Source: Quarterly Labour Force Survey.

in households' later interviews. On first interviews the ILO unemployment rate rose by 3.6 and 1.4 points, which leaves the recession as the deeper slowdown in the unemployment and hiring of the young. Unemployment at all ages peaked in 2011, two years after the trough in output (Cribb et al., 2017). Among 22–25-year-olds the ILO unemployment rate kept rising, from between 0.081 and 0.088 in 2006–2008 to 0.128, 0.133 and 0.135 in 2011–2013, the highest three-year average of the episode, while their employment rate was lowest in 2011 and 2012.

Figure 8 plots the three central series across both episodes, whose estimates table 9 reports side by side. The elementary share of new hires aged 22–25 was 0.205 and 0.176 in 2006 and 2007 and 0.225, 0.212 and 0.225 in 2009–2011, a rise of 2.4 points with an interval from -0.1 to 4.9 , against 7.9 points in 2023–2025. Measured over 2011–2013 the rise is 1.7 points, with an interval from -0.8 to 4.1 . The difference between the two episodes is 5.5 points with an interval from 1.1 to 9.9 . Without weights the recession estimate is unchanged and the recent one falls to 4.8 points, which shrinks the difference to 2.4 points and removes its significance. Weights alone raise the recent estimate from 4.8 to 7.9 points while doing little to the earlier ones. The administrative share fell by the same amount in both episodes, 4.4 points in the recession, from 0.182 and 0.175 to 0.137, 0.142 and 0.125, and 4.9 points in 2023–2025, a fall that, at between 3.1 and 4.4 points with weights, holds in every recession window tried, including the windows from 2011 (appendix table A27). The share of young

Table 9: The two episodes: person-level estimates, 2009–2011 against 2006–2007 and 2023–2025 against 2015–2022

	2009–2011 vs 2006–2007	2023–2025 vs 2015–2022	Difference	z (p)
Elementary share of new hires aged 22–25	+0.024 [-0.001, +0.049]	+0.079 [+0.043, +0.115]	+0.055	+2.45 (0.014)
unweighted	+0.024 [-0.000, +0.048]	+0.048 [+0.020, +0.075]	+0.024	+1.27 (0.204)
graduates	+0.011 [-0.015, +0.038]	+0.078 [+0.037, +0.119]	+0.067	+2.67 (0.008)
graduates, unweighted	+0.016 [-0.010, +0.042]	+0.071 [+0.038, +0.105]	+0.055	+2.56 (0.010)
non-graduates	+0.040 [+0.003, +0.077]	+0.094 [+0.034, +0.154]	+0.054	+1.51 (0.131)
non-graduates, unweighted	+0.035 [-0.000, +0.071]	+0.039 [-0.008, +0.085]	+0.003	+0.11 (0.912)
Administrative share of new hires aged 22–25	-0.044 [-0.067, -0.022]	-0.049 [-0.071, -0.027]	-0.005	-0.31 (0.757)
Groups 1–3 share of new hires aged 22–25	+0.002 [-0.027, +0.030]	+0.004 [-0.034, +0.042]	+0.002	+0.08 (0.933)
Groups 5, 8, 9 share of new hires aged 22–25	+0.001 [-0.027, +0.030]	+0.073 [+0.033, +0.112]	+0.071	+2.89 (0.004)
Employment rate, ages 22–25	-0.046 [-0.052, -0.039]	-0.021 [-0.029, -0.012]	+0.025	+4.60 (<0.001)
unweighted	-0.045 [-0.052, -0.039]	-0.002 [-0.008, +0.005]	+0.044	+9.35 (<0.001)
with a degree	-0.049 [-0.061, -0.038]	-0.002 [-0.014, +0.009]	+0.047	+5.58 (<0.001)
with a degree, unweighted	-0.048 [-0.059, -0.037]	+0.012 [+0.003, +0.021]	+0.060	+8.23 (<0.001)
without a degree	-0.047 [-0.055, -0.039]	-0.041 [-0.053, -0.030]	+0.005	+0.75 (0.451)
without a degree, unweighted	-0.046 [-0.054, -0.039]	-0.020 [-0.029, -0.011]	+0.026	+4.33 (<0.001)
Share ILO unemployed, ages 22–25	+0.032 [+0.028, +0.036]	+0.005 [+0.001, +0.010]	-0.027	-8.79 (<0.001)
New hires per person, ages 22–25	-0.014 [-0.017, -0.011]	-0.005 [-0.009, -0.001]	+0.009	+3.57 (<0.001)
unweighted	-0.014 [-0.017, -0.011]	-0.005 [-0.008, -0.002]	+0.009	+3.90 (<0.001)

Note.—Each cell is the coefficient on an episode indicator in a regression of the individual-level outcome on that indicator and quarter-of-year effects, estimated separately on the years of the two periods compared, with person weights unless stated, and a 95 percent interval clustered by household within quarter. The difference is the 2023–2025 estimate minus the recession estimate, with a z statistic from the two standard errors (the samples do not overlap). New hires are employees with less than three months with the employer. Source: Quarterly Labour Force Survey, 2006–2025.

entrants in professional jobs moved in neither episode. Manual jobs, trades, operatives and elementary occupations together, did not gain share in the recession and gained 7.3 points in 2023–2025, a difference of 7.1 points ($z = 2.9$).

The contrast is largest for graduates. Among new hires aged 22–25 with a degree, the elementary share was 0.089 and 0.077 in 2006–2007 and 0.088, 0.083 and 0.120 in 2009–2011, a change of 1.1 points ($p = 0.40$), against 7.8 points in 2023–2025. The difference is 6.7 points ($z = 2.7$) and survives the removal of weights (5.5 points, $z = 2.6$). Among non-graduate new hires, the recession raised the elementary share by 4.0 points and 2023–2025 by 9.4 points, a difference that is not distinguishable from zero ($z = 1.5$) and that disappears without weights (3.5 against 3.9 points). The employment rate of young graduates fell by 4.9 points in the recession and did not move in 2023–2025, while that of young non-graduates fell by 4.7 and 4.1 points, so that in 2009–2011 the recession moved graduates out of employment with no significant change in their elementary share. The British studies of the 2008 cohorts, which count sales, customer-service, caring and leisure jobs with elementary ones as low-paying work, find that graduates traded down into them for years afterwards (Clarke, 2019; Henehan, 2020), while the lasting employment losses fell on the less educated (Clarke, 2019; Cribb et al., 2017). On that definition the share of graduate new hires aged 22–25 starting in low-paying jobs, 0.228 in 2006–2007, rose by 6.7 points in 2009–2011 (interval 2.7 to 10.8) and

by 8.9 points in 2011–2013 (4.9 to 13.0), against 6.9 points in 2023–2025 (1.9 to 11.9). The two episodes do not differ ($z = 0.04$ and -0.63 ; -0.69 and -1.25 without weights; appendix table A28). Of the rise of 2011–2013, 4.4 points were in sales and customer-service jobs, 1.7 in caring and leisure jobs ($p = 0.11$) and 2.9 in elementary jobs, whereas in 2023–2025 the sales and caring shares did not move and the rise was all in elementary jobs.

Two qualifications apply. A further difference between the episodes, taken up last, leads to sections 5.3 and 5.4. First, measured from the peak in unemployment, graduates’ elementary share rose by about three points, an estimate at the edge of significance. The share rose to 0.120 in 2011 and 0.132 in 2012 and fell back to 0.093 and 0.102 in 2013 and 2014. Against 2006–2007, the 2011–2013 window gives a rise of 2.9 points, with an interval from 0.1 to 5.6 ($p = 0.04$), which becomes 3.3 points without weights and 3.0 points with the first quarter of 2008 in the base period or with SOC 2000 codes throughout. The 2012–2014 window gives 2.4 points ($p = 0.08$). The rise comes from 2011 and 2012 and mainly from interviews in October–December, since without 2012 the estimate is 1.9 points ($p = 0.21$) and without those interviews 0.9 points ($p = 0.62$). A test on the annual shares gives $p = 0.13$, whereas three of the fourteen recession windows tried reach $p < 0.05$. The recent rise of 7.8 points is larger by 4.9 points ($z = 2.0$), a difference at the margin of significance that is 3.8 points ($z = 1.7$) without weights. A hiring slowdown may thus move graduates into elementary jobs, two years after the trough in output and by about three points, whereas the recent rise was more than twice as large, began in 2023, the year after vacancies peaked, and came with no fall in graduates’ employment rate. Second, the composition of young new hires changed between the episodes, as the share with a degree rose from 0.42–0.44 in 2006–2007 to 0.55–0.61 in 2022–2025. Since graduates’ elementary share is well below non-graduates’, the aggregate share is pulled down over time and understates the within-group shifts, which are the comparable ones.

Third, the two episodes differ in more than the presence of language models. The minimum wage rose by small steps in 2009–2011 and sharply in 2023–2025 (Low Pay Commission, 2024). Foreign-born workers in elementary jobs also grew in number after 2021. The benchmark on its own cannot separate these forces from the technology. Sections 5.3 and 5.4 examine what each accounts for. The benchmark shows that a hiring slowdown of the size of 2008–2011, acting on the labour market of its time, produced an administrative decline of the recent size, a larger fall in youth employment, as large a move of graduates into low-paying entry jobs, mostly into sales and care jobs, and a borderline rise of about three points in their elementary entry, two years after the trough in output.

5.2 The elementary detour

The reallocation took place at entry, while the stock of jobs barely moved. Between 2015–2022 and 2023–2025, young hires in exposed occupations fell by 18 thousand a quarter. Elementary hires of the young rose by 9 thousand, half as much, an increase offset by a fall of 9 thousand in the other unexposed occupations, so that total young hiring fell by a tenth, as much in number as hiring in exposed occupations did (table 7, section 4.3). The stock of young employees in elementary jobs barely changed. Quarterly exits from elementary jobs did not change, while the share moving from an elementary into a professional job each quarter fell from 0.034 to 0.019 (appendix table A6). Young people in elementary jobs left them at the earlier rate, and less often for a professional job. Graduates took the detour as much as non-graduates, since the share of graduates who were students a year earlier and now work in an elementary job more than doubled (section 4.2), and among graduate new hires aged 22–25 the share starting in the most exposed third of occupations fell from 0.45 to 0.35 while the share starting in the least exposed third rose from 0.23 to 0.32 (section 4.6). The fall in demand reached this cohort as a change in the jobs it started more than as a loss of employment.

5.3 The wage floor

Section 4.4 reported the facts, that the National Living Wage rose by about a tenth in April 2023 and again in April 2024, that the share of young graduates paid near the floor doubled, and that the pay gap between an elementary and a professional job at ages 22–25 narrowed. A higher floor, whose effects the minimum-wage literature leads one to expect in the distribution of pay and the mix of jobs more than in employment (Cengiz et al., 2019; Manning, 2021), lowers the cost of an elementary start for a graduate. The floor may therefore have drawn into elementary work some graduates who would otherwise have waited for a professional opening, which helps to explain where the entrants went.

The floor cannot account for the demand side. In principle it could affect hiring into the lowest-paid exposed jobs as well as the destination of entrants, because the share paid within 5 percent of it rose among young administrative employees as well as among elementary workers, from 0.07 in 2022 to 0.16, 0.21 and 0.19 in 2023–2025. It does not, however, reach the fall in adverts for exposed occupations within the professional groups, where the observed-exposure gradient is 0.18 log points per standard deviation (table 5, panel B), nor the relative decline of exposed occupations that began in 2020. Nor can the benchmark bound the floor’s role, because the floor rose by small steps in 2009–2011 while the technology was absent, which leaves the two unseparated there. The floor thus explains part of where

the entrants went. Another cause is needed for the fall in demand for the work they would have done.

5.4 Migration

Section 4.5 reported the composition, in which the number of foreign-born employees in elementary occupations rose after 2021, the UK-born share of elementary new hires aged 26–64 fell, and among new hires aged 22–25 the elementary share rose among the UK-born while that of the foreign-born, a small group, did not change (appendix table A5). Migration therefore raised the supply of labour to elementary jobs at the ages above entry, jobs that young natives and recent migrants entered together. Migration does not account for the shift among the young natives themselves, who are more than four fifths of young new hires and whose elementary share rose by as much as the aggregate. Nor does it account for the fall in adverts for exposed occupations within broad groups or in young hiring in the most exposed occupations relative to the least exposed. Like the floor, it affects where the entrants went, whereas the demand for the work they would have done fell for another reason.

5.5 The layers above entry

A hypothesis that language models thin the entry layer has a counterpart for the layers above it, where experienced workers could be drawn in as the complemented side of the technology or displaced in their turn. Appendix F, which applies the designs of section 4.7 to employees of all ages, finds that workers above the entry layer were not drawn in, since, once exposure is allowed a linear trend of its own, their employment in exposed occupations did not grow after 2022 and their supervisory shares and promotion rates did not change. It also finds that the evidence of displacement is weak, because the quarter-to-quarter separations of 35–49-year-olds in exposed occupations depend on a few hundred events and do not survive the removal of the IT and customer-service groups. The cross-sectional redundancy rate has a placebo break at the 2021 classification change. The change the data show is at the entry layer.

5.6 What the evidence does not settle

Seven limits bound the reading of the evidence. Adoption measured at the level of industries predicts a small fall in the hiring of the young, driven by one sector, and does not predict the elementary shift (section 4.9). The within-group gradient in young hiring depends on a few customer-service occupations and on small cells. The paper treats it as consistent with the

advertises evidence without treating it as independent (section 4.7). Within broad groups, the fall in adverts beyond the pre-2023 trend is concentrated in the IT occupations, where the technology hiring boom also ended. Without the twelve IT unit groups the gradient remains and the part beyond trend does not (appendix table A10). Adverts measure advertised vacancies, a step before hires. The task-based measures reach British occupations through a mapping whose error attenuates their estimates (section 2.4). Teleworkability is correlated with exposure across occupations and accounts for much of the exposure gradient in other British data (Lambert and Schindler, 2026; Henseke et al., 2026). With the ONS measure as a control, the within-group gradients along observed exposure hold and the ability-based gradient does not. The part beyond the pre-2023 trend then depends on the window of the trend, because the control accounts for the decline of 2020–2022. Where exposure and teleworkability coincide, as in the IT occupations, the two cannot be separated (section 4.7). The benchmark cannot by itself separate the wage floor and migration from the technology. Measured from the 2011 peak in unemployment, the downturn raised graduates’ elementary entry by a borderline three points against eight in 2023–2025, a difference itself at the margin of significance. Part of the recent rise may thus be what a hiring slowdown does (section 5.1).

6 Conclusion

Young people in Britain entered different jobs after November 2022. The share of new hires aged 22–25 starting in an elementary occupation rose from 0.19 in 2015–2022 to between 0.24 and 0.30 in 2023–2025, the share starting in an administrative occupation fell by a third, and the share starting in a professional occupation did not change. This shift is one of flows, as large among recent graduates as among non-graduates, coinciding with fewer moves from elementary into professional jobs. The employment and pay of the young moved less than the mix of their jobs. Workers above the entry layer were not drawn in.

Did language models do this? The comparison that can answer the question is the one between occupations and their neighbours in the same broad group, before and after the release, with exposure allowed a trend of its own. On it, the demand for labour fell in the occupations whose tasks language models perform, by 0.16 log points of adverts per standard deviation of observed exposure and by 0.08 net of the pre-2023 trend, a part concentrated in information technology. Hires of 22–25-year-olds in exposed occupations fell by 28 percent while hiring in unexposed occupations was unchanged, a fall of which the within-group gradient accounts for three fifths and its part net of trend for a third. The 2008–2011 downturn, which raised youth unemployment far more, produced the same decline in administrative

entry jobs and as large a move of graduates into low-paying jobs, mostly into sales and care jobs. Measured from the peak in unemployment, it raised graduates' elementary entry by a borderline three points, against eight after 2022, when their employment did not fall. The wage floor and migration shaped where the displaced entrants went and do not account for the fall in demand. The outcomes suggest as a fall in the demand for labour in the occupations whose tasks language models perform, a fall that began with their diffusion and is clearest in information technology, and as a reallocation of young entrants away from that work.

Two directions follow. The first is to follow the cohorts, because the paper observes entry jobs and stops short of careers. Whether a graduate who started in a bar in 2024 is in a professional job in 2028 is the question that decides whether the shift has lasting consequences, which panel elements of the LFS and the tax-linked graduate records can answer in a few years. The responses of students and graduates by subject, in appendix H, show one adjustment so far, a reversal of the rise in applications to computing, and no movement in graduate employment or pay by subject. The second direction is to price the detour. Both episodes are configurations of the same moments. One is a hiring shock that moved the young out of employment, cut administrative entry and moved graduates into sales and care jobs, with a smaller and later move into elementary jobs. The other is a shock that moved graduates into elementary jobs at once, with their employment unchanged, under a rising wage floor. A model of entry with search frictions, two kinds of entry jobs in which experience accumulates at different rates, and a wage floor could be asked to generate both, and to say what the elementary detour of the 2023–2025 cohorts will cost them. The evidence on the layers above entry, in appendix F, would fix where the shock fell, on entry, with experienced labour neither drawn in nor, on the evidence so far, displaced in any robust way.

References

- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Altonji, J. G., Kahn, L. B., and Speer, J. D. (2016). Cashier or consultant? Entry labor market conditions, field of study, and career success. *Journal of Labor Economics*, 34(S1):S361–S401.
- Anthropic (2026). Anthropic economic index: Occupational exposure and usage data. Labour-market exposure release of March 2026 and the Claude.ai usage-by-occupation dataset released on 26 June 2026 (April and May 2026 data), <https://www.anthropic.com/economic-index>.
- Autor, D. and Thompson, N. (2025). Expertise. *Journal of the European Economic Association*, 23(4):1203–1271.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.
- Autor, D. H. and Dorn, D. (2009). This job is “getting old”: Measuring changes in job opportunities using occupational age structure. *American Economic Review: Papers and Proceedings*, 99(2):45–51.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.
- Beaudry, P., Green, D. A., and Sand, B. M. (2014). The declining fortunes of the young since 2000. *American Economic Review: Papers and Proceedings*, 104(5):381–386.
- Beaudry, P., Green, D. A., and Sand, B. M. (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

- Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. Working Paper 32966, National Bureau of Economic Research.
- Brynjolfsson, E., Chandar, B., and Chen, R. (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab. Version of August 2026; first version August 2025.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942.
- Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.
- Chandar, B. and Klein Teeselink, B. (2026). How does AI change labor demand? Evidence from 41 countries. Working paper, Stanford Digital Economy Lab.
- Clarke, S. (2019). Growing pains: The impact of leaving education during a recession on earnings and employment. Technical report, Resolution Foundation, London.
- Cortes, G. M., Jaimovich, N., Nekarda, C. J., and Siu, H. E. (2020). The dynamics of disappearing routine jobs: A flows approach. *Labour Economics*, 65:101823.
- Cortes, G. M., Jaimovich, N., and Siu, H. E. (2017). Disappearing routine jobs: Who, how, and why? *Journal of Monetary Economics*, 91:69–87.
- Cribb, J., Hood, A., and Joyce, R. (2017). Entering the labour market in a weak economy: Scarring and insurance. Working Paper W17/27, Institute for Fiscal Studies.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4):1593–1640.
- Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. *Quarterly Journal of Economics*, 135(4):1965–2005.
- Department for Education (2023). The impact of AI on UK jobs and training. Technical report, Department for Education, London.
- Department for Education (2025a). Graduate labour market statistics: Calendar year 2024. Official statistics, 5 June 2025.
- Department for Education (2025b). Graduate outcomes (LEO): Tax year 2023 to 2024. Official statistics, Department for Education, London.

- Dillender, M. and Forsythe, E. (2022). Computerization of white collar jobs. Working Paper 29866, National Bureau of Economic Research.
- Dwan-O'Reilly, J. (2026). London's workforce exposure to generative artificial intelligence. Working Paper 103, GLA Economics.
- Egan, E. (2026). Generative AI: Degenerative for jobs? Bank Underground, Bank of England staff blog, 22 January 2026.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.
- Facius, D. and Iacono, R. (2026). Labor market consequences of generative AI: Early evidence from Norway. Working Paper 12752, CESifo.
- Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.
- Felten, E., Raj, M., and Seamans, R. (2023). How will language modelers like ChatGPT affect occupations and industries? Technical Report 2303.01157, arXiv.
- Forsythe, E., Kahn, L. B., Lange, F., and Wiczer, D. (2020). Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims. *Journal of Public Economics*, 189:104238.
- Gmyrek, P., Berg, J., and Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. ILO Working Paper 96, International Labour Organization.
- Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *Review of Economics and Statistics*, 89(1):118–133.
- Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.
- Hampole, M., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research.

- Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., and Ganguli, D. (2025). Which economic tasks are performed with AI? Evidence from millions of Claude conversations. Technical Report 2503.04761, arXiv.
- Henehan, K. (2020). Class of 2020: Education leavers in the current crisis. Technical report, Resolution Foundation, London.
- Henseke, G., Davies, R., Felstead, A., Gallie, D., Green, F., and Zhou, Y. (2026). Nine raters, one index: Carrying LLM disagreement into labour-market estimates. Technical Report 2507.22748, arXiv. Version 4, August 2026; earlier versions titled “How Exposed Are UK Jobs to Generative AI?”.
- Hernæs, Ø. and Kostøl, A. R. (2026). Has AI widened employment gaps? Tracking early-career employment by occupational exposure in Norway. Discussion Paper 179/26, RF-Berlin.
- Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.
- HM Revenue and Customs and Office for National Statistics (2026). Earnings and employment from Pay As You Earn real time information, UK. Release of September 2026.
- Hosseini Maasoum, S. M. and Lichtinger, G. (2026). Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. Working Paper 5425555, SSRN. First version August 2025, revised June 2026.
- Humlum, A. and Vestergaard, E. (2026). Still waters, rapid currents: Early labor market transformation under generative AI. Working Paper 33777, National Bureau of Economic Research. Revised March 2026; first version May 2025 as “Large Language Models, Small Labor Market Effects”.
- Ide, E. and Talamàs, E. (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133(12):3762–3800.
- Jaimovich, N. and Siu, H. E. (2020). Job polarization and jobless recoveries. *Review of Economics and Statistics*, 102(1):129–147.
- Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316.

- Kauhanen, A. and Rouvinen, P. (2026). AI has not impacted the youth labor market in Finland. Working Paper 135, ETLA.
- Klein Teeselink, B. (2025). Generative AI and labor market outcomes: Evidence from the United Kingdom. Working Paper 5516798, SSRN.
- Lambert, P. J. and Schindler, Y. (2026). The broken ladder: AI, remote work, and early-career hiring. Warwick Economics Research Paper 1615, University of Warwick. Also CAGE Working Paper 808 and RFBerlin Discussion Paper 228/26.
- Lodefalk, M., Löthman, L., Koch, M., and Engberg, E. (2026). Same storm, different boats: Generative AI and the age gradient in hiring. Working Paper 2/2026, Örebro University.
- Low Pay Commission (2024). National minimum wage: Low pay commission report 2024. Technical report, Low Pay Commission, London.
- Manning, A. (2021). The elusive employment effect of the minimum wage. *Journal of Economic Perspectives*, 35(1):3–26.
- Massenkoff, M. and McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Research report, Anthropic. 5 March 2026.
- Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, 96(1):60–77.
- Modestino, A. S., Shoag, D., and Ballance, J. (2020). Upskilling: Do employers demand greater skill when workers are plentiful? *Review of Economics and Statistics*, 102(4):793–805.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.
- Office for National Statistics (2010). Standard occupational classification 2010, volume 1: Structure and descriptions of unit groups. Palgrave Macmillan, Basingstoke.
- Office for National Statistics (2013). Graduates in the UK labour market: 2013. Article and reference tables, 19 November 2013.
- Office for National Statistics (2017). Graduates in the UK labour market: 2017. Article and time-series tables, 24 November 2017.

- Office for National Statistics (2020). Which jobs can be done from home? Article of 21 July 2020, with the ability-to-homework score by SOC 2010 unit group.
- Office for National Statistics (2024a). Impact of reweighting on Labour Force Survey key indicators: 2024. Methodology article, Office for National Statistics.
- Office for National Statistics (2024b). Long-term international migration, provisional: Year ending december 2023. Statistical bulletin, Office for National Statistics.
- Office for National Statistics (2025). Annual survey of hours and earnings: Occupation (4 digit SOC) table 14. Provisional 2025 release.
- Office for National Statistics (2026a). Business insights and conditions survey: Artificial intelligence adoption by industry. Wave 159 time series, waves to June 2026.
- Office for National Statistics (2026b). Online job advert estimates: Labour demand volumes by standard occupation classification (SOC 2020), UK. Experimental statistics based on Textkernel data, release of September 2026.
- Office for National Statistics (2026c). Quarterly labour force survey, 2006–2026: End user licence person files. UK Data Service, SN 5369 to SN 9607.
- Office for National Statistics (2026d). Vacancies by industry (VACS02). Release of September 2026.
- Oreopoulos, P., von Wachter, T., and Heisz, A. (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29.
- Orr, C., Tucker, L. C., and Warren, L. (2026). Graduating into disruption: Labor market outcomes for AI-exposed college majors. Working Paper CES-26-56, US Census Bureau, Center for Economic Studies.
- Palladino, M. G., Bergeaud, A., Bertheau, A., Bunel, S., and Degos, O. (2026). Generative AI: The first changes in the labour market are affecting young people. Blog post 463 and LIFT Working Paper 2026-02, Banque de France.
- Rothstein, J. (2023). The lost generation? Labor market outcomes for post great recession entrants. *Journal of Human Resources*, 58(5):1452–1479.
- Salvatori, A. (2018). The anatomy of job polarisation in the UK. *Journal for Labour Market Research*, 52(8).

- Schlicht, H. (2026). Canaries in the column? AI exposure and the UK’s hiring slowdown. Bank Underground, Bank of England staff blog, 6 August 2026.
- Schwandt, H. and von Wachter, T. (2019). Unlucky cohorts: Estimating the long-term effects of entering the labor market in a recession in large cross-sectional data sets. *Journal of Labor Economics*, 37(S1):S161–S198.
- Tomlinson, K., Jaffe, S., Wang, W., Counts, S., and Suri, S. (2025). Working with AI: Measuring the occupational implications of generative AI. Technical Report 2507.07935, arXiv.
- Tucker, L. C. (2026). You’re (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators. Working Paper CES-26-27, US Census Bureau, Center for Economic Studies.
- UCAS (2026). Undergraduate applications by subject group at the 30 june deadline, 2019–2026 cycles. UCAS data releases.
- von Wachter, T. (2020). The persistent effects of initial labor market conditions for young adults and their sources. *Journal of Economic Perspectives*, 34(4):168–194.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. Working paper, Stanford University.

A Data Appendix

A.1 Labour Force Survey

Files and sample. The paper uses the End User Licence person files of the Quarterly Labour Force Survey distributed by the UK Data Service: 46 quarterly files from January–March 2015 to April–June 2026, 12 files for 2012–2014 and 23 files for 2006–2011 (every quarter except April–June 2007, which is not in the catalogue). The sample keeps respondents aged 16–64 with a positive person weight. Each file is an independent cross-section of households in their first to fifth quarter in the survey. The household identifier is unique within a file, which is why all clustering is by household within quarter. Table A1 gives, by year, the number of interviews of new hires aged 22–25, the Kish effective sample implied by the person weights, the largest weight, the share of proxy responses and the share with a degree.

Table A1: Interviews of new hires aged 22–25, weights and composition, by year

Year	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026
Interviews, new hires aged 22–25	849	741	742	676	622	397	633	452	244	328	460	199
Kish effective sample	795	665	670	598	537	339	496	354	163	241	361	167
Largest person weight	3951	4769	5431	8008	5318	4598	4082	4736	11053	8450	8109	3795
Proxy-response share (weighted)	0.486	0.490	0.519	0.531	0.534	0.568	0.601	0.600	0.618	0.629	0.572	0.593
Share with a degree (weighted)	0.455	0.535	0.491	0.487	0.538	0.551	0.546	0.605	0.551	0.586	0.579	–

Note.—Interviews of employees aged 22–25 with less than three months with the employer, in occupations with an exposure score, by calendar year of the interview; 2026 covers January–June. The Kish effective sample is $(\sum w)^2 / \sum w^2$ over the person weights of those interviews. The proxy share is the weighted share of interviews answered by another household member. The last four rows are computed on the sample of the placebo scripts, whose yearly interview counts differ from the first row by at most two. Source: Quarterly Labour Force Survey.

Variables. Labour-market status is the ILO definition, with employees distinguished from the self-employed. A new hire is an employee whose reported length of time with the current employer is less than three months. Occupation is the four-digit SOC 2010 unit group of the main job, coded by the ONS up to 2020 and, from the first quarter of 2021, derived by the ONS from the SOC 2020 code. The results on entry jobs of section 4 are the same when native SOC 2020 codes are used for 2021–2026, as the independent re-derivation in the replication package documents. Industry is the SIC 2007 section of the main job. The highest qualification is the ONS derived variable, with degree or equivalent as the top category; respondents who do not know their qualification are treated as missing. Years since leaving full-time education is age minus the reported age at which continuous full-time education ended. Hourly pay is the ONS derived hourly pay of employees in their first and fifth

interview, deflated by the consumer price index to 2015 prices and weighted with the income weights. The retrospective module of the April–June files reports economic activity, and for those in work the occupation, twelve months before the interview. Country of birth distinguishes the four UK nations from all other countries, while relationship to the head of the family unit identifies young people living with their parents.

Weights and the 2021 and 2011 breaks. Person weights are those supplied with each file: PWT14 for 2006 to mid-2011, PWT18 for mid-2011 to 2018 and PWT24 from 2019. Those introduced in 2024 (PWT24, supplied for the files from 2019) give 22–25-year-olds a higher relative weight than the earlier weights and are the reason the weighted post-2022 changes exceed the unweighted ones. In 2011 the change from SOC 2000 to SOC 2010 barely affects the elementary group and moves the administrative share of young new hires by about a point (table A29, panel C, and table A28, panel B). Two-quarter transitions of table A6 use the two-quarter longitudinal files with their longitudinal weights, which follow each household across consecutive quarters.

A.2 Exposure measures and crosswalks

The Department for Education score is published for 365 SOC 2010 unit groups. Of the US measures, the Anthropic observed-exposure index covers 756 US SOC 2018 occupations, the Microsoft applicability score 785, and the Eloundou et al. rating 923 O*NET-SOC 2019 occupations. Anthropic’s usage data give the share of Claude.ai conversations from Great Britain by O*NET occupation for 444 occupations, averaged over the April and May 2026 releases. Two crosswalks link them to British occupations. The O*NET–ISCO-08 crosswalk published with the ESCO classification links 958 O*NET codes to ISCO-08 unit groups. The ONS proportional conversion links 420 ISCO-08 codes to 384 SOC 2010 unit groups with the proportions of each ISCO code’s job titles falling in each SOC group, normalised to one within the ISCO code. At the first step the score of an ISCO code is the mean over its O*NET occupations, each weighted by one over the number of ISCO codes it maps to. At the second step the score of a SOC 2010 unit group is the mean over its ISCO codes weighted by the normalised proportions. Usage shares are split by the same weights and summed, then divided by the unit group’s share of employees aged 16–64 in 2024–2025 and logged. After these steps 364 of the 365 DfE-scored unit groups have a score on each US measure and 339 a usage share. The median unit group draws on two ISCO codes. The mapping is coarse at the level of neighbouring unit groups. A programmer’s SOC 2010 group inherits, through ISCO 2514, the scores of US computer programmers and software developers together with those of numerical-control programmers and machinists,

and a nurse’s group inherits those of registered nurses and nursing assistants. Benchmark occupations land where they should, cleaners and chefs at the floor and call-centre staff and programmers near the top, while the aggregate statistics of table A2 are insensitive to the weighting at the ISCO step. Differences between adjacent unit groups within a major group, however, have measurement error of the order of the signal. All measures are standardised over employees aged 16–64 in 2015–2022 with person weights. The teleworkability control of appendix tables A12, A13 and A14 is minus the ability-to-homework score of Office for National Statistics (2020), the sum of five factors (exposure, interaction, location, physical activities, and tools and protective equipment) that the ONS built from O*NET work-activity and work-context ratings collected before the pandemic and mapped from American occupations through ISCO-08 to 365 SOC 2010 unit groups. Higher values mean work that can more readily be done from home. It is standardised and carried to SOC 2020 as the exposure measures are. It covers 364 of the 365 scored SOC 2010 unit groups and 406 of the 407 SOC 2020 unit groups with an observed-exposure score in the adverts panel. Because the mapping goes through ISCO-08, unit groups that share ISCO codes share a value, which leaves the 364 SOC 2010 unit groups with 190 distinct values.

Table A2: Exposure measures: variance between occupation groups, and the share of young new hires starting above each measure’s median

Measure	Unit groups	Variance between major groups	Between sub-major groups	Correlation with DfE within major groups
DfE language-model exposure	365	0.727	0.839	1.000
Anthropic observed exposure	364	0.386	0.611	0.510
Microsoft AI applicability	364	0.440	0.645	0.547
Eloundou GPT-4 feasibility	364	0.634	0.794	0.564
Claude.ai GB usage per worker	339	0.345	0.481	0.199
<i>Share of new hires starting a job above the measure’s median: 2015–2022 mean, 2023–2025 values, change (p)</i>				
	Ages 22–25		Ages 35–49	
DfE language-model exposure	0.451; 0.338, 0.404, 0.400	-0.070 (0.009)	0.501; 0.554, 0.538, 0.549	+0.046 (0.005)
ILO generative-AI task exposure (corr. 0.77)	0.428; 0.312, 0.396, 0.383	-0.065 (0.012)	0.474; 0.481, 0.483, 0.534	+0.025 (0.215)
Felten language-modelling AIOE (corr. 0.91)	0.436; 0.305, 0.384, 0.371	-0.082 (0.007)	0.491; 0.519, 0.525, 0.550	+0.040 (0.010)
DAIOE language-model index (corr. 0.88)	0.432; 0.285, 0.382, 0.372	-0.085 (0.008)	0.487; 0.513, 0.515, 0.534	+0.033 (0.029)
DfE AI exposure, all applications (corr. 0.98)	0.444; 0.345, 0.394, 0.383	-0.070 (0.006)	0.494; 0.532, 0.521, 0.545	+0.039 (0.010)
Webb patent-based exposure (corr. 0.22)	0.380; 0.353, 0.444, 0.393	+0.017 (0.407)	0.497; 0.539, 0.502, 0.540	+0.030 (0.037)

Note.—Upper panel: share of the employment-weighted variance (employees aged 16–64, 2015–2022) of each standardised measure that lies between the nine SOC 2010 major groups and between the 25 sub-major groups, and the employment-weighted correlation with the Department for Education score after removing major-group means. Lower panel: share of new hires aged 22–25 (and 35–49) whose unit group lies above the median of the measure over weighted 2015–2022 employment; the 2015–2022 mean and the 2023, 2024 and 2025 values, with the year-level change and its *p*-value. The correlation in brackets is the employment-weighted correlation of the measure with the Department for Education score across unit groups. The ILO, Felten–Raj–Seamans, DAIOE and Webb measures reach SOC 2010 through ISCO-88 and a bridge estimated on Understanding Society respondents whose job carries both codes. Source: Quarterly Labour Force Survey; measures as cited in the text.

A.3 Online job adverts


The ONS labour-demand tables derived from Textkernel data give monthly counts of new online job adverts by SOC 2020 unit group for the UK from January 2017. The panel has 412 unit groups over the 115 months to August 2026 for which the series is published. March 2026 is missing. The 813 cells suppressed for disclosure, which are concentrated in 2025–2026, are treated as missing. Each SOC 2020 unit group receives the employment-share-weighted mean of the SOC 2010 scores of the groups it maps to in the ONS relationship table, which gives a score to 408 of the 412 groups. The others map to SOC 2010 groups without a score or have less than half of their mapped employment scored. Log adverts are the log of the count, which is never zero. The exposure thirds of figure 6 are weighted by 2017–2022 adverts. Median hourly pay by SOC 2020 unit group is taken from the Annual Survey of Hours and Earnings, provisional 2025 release, table 14.6a.

A.4 Employer adoption, administrative series and education data

The Business Insights and Conditions Survey asks businesses which artificial-intelligence technologies they currently use. The adoption measure is the weighted share using “text generation using large language models” by industry section, from the twelve waves whose reference periods run from September 2023 to June 2026, and the share reporting that artificial intelligence reduced headcount from the four waves of September 2025 to June 2026. Water supply is suppressed in every wave, which leaves thirteen of the fourteen covered sections with a value. Because several sections are suppressed in some waves, the means use six to twelve waves. Real Time Information tables give seasonally adjusted monthly payrolled employees by age band and industry section and monthly inflows into payrolled employment. The Vacancy Survey gives seasonally adjusted three-monthly vacancies by section. For the industry panel, vacancies are those of the three months ending in the quarter’s last month, and payrolled employees the quarterly mean. The Longitudinal Education Outcomes tables of the Department for Education report, for UK-domiciled first-degree graduates of English providers, the share with sustained employment or further study and median earnings one, three and five years after graduation, by detailed subject and tax year. I map subjects to the 19 single-subject bands of the LFS and standardise the exposure of a subject as the mean language-model exposure of the jobs held by employed graduates of the subject aged 30–49 in 2015–2022. UCAS reports applications by subject group from UK 18-year-olds at the 30 June deadline of each cycle, and acceptances at the end of each cycle.

B Additional Results on Entry Jobs

Share aged 22-25 among new hires, by SOC 2010 major group, UK, 2015-2025



Weighted shares by calendar year 2015-2025; 95% intervals clustered by unit group; grey dashed line: mean of the 2015-2022 annual shares; vertical line: 30 Nov 2022. Panels ordered by the group's mean DfE (2023) language-model exposure (z, weighted by 2015-2022 employment). New hires: employees with under three months with the employer. Source: Quarterly Labour Force Survey (ONS), employees aged 16-64, all waves, person weights.

Figure A1: Share aged 22-25 among new hires, by SOC 2010 major group, 2015-2025

Weighted annual shares of new hires (employees aged 16-64 with less than three months with the employer) who are aged 22-25, in each major group; 95 percent intervals clustered by unit group; dashed line: mean of the 2015-2022 annual shares; vertical line: 30 November 2022. Panels are ordered by the group's mean Department for Education language-model exposure (standardised, in brackets). Source: Quarterly Labour Force Survey, all waves, person weights.

Table A3: The elementary share of new hires aged 22–25 by subgroup and age

Elementary share of new hires aged 22–25	Mean 2015–2022	Range	2023	2024	2025	Change	<i>p</i>	Outside
All	0.187	0.161–0.223	0.298	0.236	0.267	+0.080	0.001	3/3
Women	0.149	0.131–0.170	0.170	0.188	0.200	+0.038	0.004	2/3
Men	0.225	0.166–0.281	0.431	0.286	0.327	+0.123	0.003	3/3
With a degree	0.113	0.091–0.139	0.215	0.179	0.186	+0.081	<0.001	3/3
Without a degree	0.260	0.205–0.322	0.401	0.276	0.380	+0.092	0.014	2/3
Excluding full-time students	0.177	0.146–0.209	0.301	0.233	0.256	+0.087	0.001	3/3
First-wave interviews only	0.184	0.147–0.228	0.274	0.213	0.243	+0.060	0.013	2/3
Ages 18–21 (elementary)	–	–	–	–	–	+0.037	0.186	1/3
Ages 26–30 (elementary)	–	–	–	–	–	-0.003	0.832	0/3

Note.—Weighted annual shares of new hires aged 22–25 (or 18–21 and 26–30 in the last two rows) in elementary occupations, by subgroup: the 2015–2022 mean and range, the 2023–2025 values, the year-level change for 2023–2025 against 2015–2022 with its *p*-value, and the number of post-2022 years outside the 2015–2022 range in the direction of the change. “First-wave interviews only” uses respondents in their first quarter in the survey. Source: Quarterly Labour Force Survey.

Table A4: Unit groups contributing to the change in the job mix of new hires aged 22–25

SOC 2010 unit group		Change in share
Elementary occupations, total		+0.0800
9272	Kitchen and catering assistants	+0.0230
9274	Bar staff	+0.0158
9233	Cleaners and domestics	+0.0108
9260	Elementary storage occupations	+0.0108
9273	Waiters and waitresses	+0.0092
9241	Security guards and related occupations	+0.0079
Administrative and secretarial occupations, total		-0.0479
4216	Receptionists	-0.0082
4122	Book-keepers, payroll managers and wages clerks	-0.0066
4159	Other administrative occupations n.e.c.	-0.0054
4129	Financial administrative occupations n.e.c.	-0.0052


Note.—Change in the weighted share of new hires aged 22–25 in each SOC 2010 unit group, 2023–2025 pooled minus 2015–2022 pooled, for the unit groups with the largest contributions to the change in the elementary and administrative major groups. Source: Quarterly Labour Force Survey.

Table A5: Country of birth of new hires in elementary jobs, 2015–2026

	2015–2022		2023	2024	2025	2026 H1	Weighted		Unweighted	
	Mean	Range					Change	<i>p</i>	Change	<i>p</i>
UK-born new hires aged 22–25: elementary share	0.171	0.148–0.208	0.276	0.239	0.280	0.235	+0.094	<0.001	+0.060	0.001
Non-UK-born new hires aged 22–25: elementary share	0.264	0.183–0.317	0.404	0.219	0.207	0.167	+0.013	0.779	+0.001	0.980
New hires aged 22–25: UK-born share	0.841	0.805–0.889	0.847	0.817	0.807	0.845	-0.017	0.328	-0.004	0.769
New hires aged 26–64 in elementary jobs: UK-born share	0.670	0.636–0.710	0.604	0.576	0.575	0.607	-0.085	0.001	-0.061	0.018
Employees in elementary jobs: UK-born share	0.757	0.736–0.805	0.748	0.746	0.731	0.712	-0.015	0.341	-0.001	0.962


Note.—Weighted annual shares. UK-born: country of birth England, Northern Ireland, Scotland, Wales or UK not specified. The change is the year-level estimate for 2023–2025 against 2015–2022 with its *p*-value, with and without weights. Source: Quarterly Labour Force Survey.

Country of birth and household position of new hires in elementary jobs, UK, 2015–2025

**Figure A2:** Country of birth and household position of new hires in elementary jobs, 2015–2025

Weighted annual shares among employees with less than three months with the employer whose unit group is in major group 9. UK-born as in table A5; living with parents: child of the head of the family unit. Bars: 95 percent intervals clustered by household within quarter; shaded band: range of the 2015–2022 values; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey, all waves.

People aged 22-25 who were full-time students a year earlier: activity and job type in April-June, UK, 2015-2026



Weighted shares among respondents aged 22-25 of each April-June quarter who report having been a full-time student twelve months earlier (OYCIIRC), with (A-C) and without (D-F) a degree: employed (ILO); and, among employees (ILODEF1, STAT 1), in an elementary job (SOC 2010 major group 9) or in major groups 1-3. Bars: 95 percent intervals clustered by household. Shaded band: range of the 2015-2022 values; dashed grey line: their mean; vertical line: 30 Nov 2022. Source: Quarterly Labour Force Survey (ONS), April-June quarters, person weights.

Figure A3: People aged 22–25 who were full-time students a year earlier: activity and job type in April–June, 2015–2026


Weighted shares among respondents aged 22–25 of each April–June quarter who report having been a full-time student twelve months earlier, with (panels A–C) and without (D–F) a degree: employed; and, among employees, in an elementary job or in major groups 1–3. Bars: 95 percent intervals clustered by household; shaded band: range of the 2015–2022 values; dashed line: their mean; vertical line: 30 November 2022. Source: Quarterly Labour Force Survey, April–June quarters, person weights.

Table A6: Quarter-to-quarter transitions of people aged 22–25, 2015–2026

Status in the second quarter	Mean 2015–2022	Range	2023	2024	2025	2026	Change	<i>p</i>
Elementary employees aged 22–25: left the elementary job	0.179	0.154–0.202	0.186	0.166	0.168	0.165	–0.006	0.602
employee in groups 1–3	0.034	0.023–0.048	0.023	0.018	0.015	0.032	–0.015	0.049
not employed	0.064	0.045–0.105	0.073	0.068	0.080	0.075	+0.010	0.446
graduates: employee in a non-elementary job	0.114	0.049–0.192	0.116	0.077	0.088	0.106	–0.020	0.510
Employees aged 22–25 in groups 1–3: left groups 1–3	0.082	0.072–0.111	0.087	0.073	0.074	0.099	–0.004	0.594
Employees aged 22–25: not employed	0.035	0.029–0.044	0.036	0.034	0.033	0.024	–0.001	0.798
Not employed aged 22–25: employee	0.166	0.133–0.208	0.214	0.161	0.151	0.136	+0.009	0.593
elementary job among those becoming employees	0.201	0.175–0.246	0.344	0.291	0.284	0.324	+0.106	<0.001

Note.—Weighted shares of people aged 22–25 in the stated first-quarter state who are in the stated second-quarter state, by calendar year of the first quarter, from the two-quarter longitudinal files with longitudinal weights; the 2026 column is the first half. The change is the year-level estimate for 2023–2025 against 2015–2022 with its *p*-value. Source: Labour Force Survey two-quarter longitudinal datasets.

Quarter-to-quarter transitions of people aged 22-25, UK, 2015-2025




Weighted shares of people aged 22-25 in the stated quarter-1 state who are in the stated quarter-2 state, by calendar year of quarter 1 (two-quarter longitudinal files, longitudinal weights). Elementary: SOC major group 9 (SOC 2010 to 2020, SOC 2020 from 2021). Bars: 95 percent intervals from the weighted binomial variance. Shaded band: range of the 2015-2022 values; dashed grey line: their mean; vertical line: 30 Nov 2022. Source: Labour Force Survey two-quarter longitudinal datasets (ONS).

Figure A4: Quarter-to-quarter transitions of people aged 22–25, 2015–2025

Weighted shares of people aged 22–25 in the stated first-quarter state who are in the stated second-quarter state, by calendar year of the first quarter. Elementary: SOC major group 9 (SOC 2010 to 2020, SOC 2020 from 2021). Bars: 95 percent intervals from the weighted binomial variance; shaded band: range of the 2015–2022 values; dashed line: their mean; vertical line: 30 November 2022. Source: Labour Force Survey two-quarter longitudinal datasets, longitudinal weights.

Payrolled employees, vacancies and hiring inflows, UK, 2015-2026



A, C: HMRC PAYE Real Time Information, seasonally adjusted, release of September 2026 (A: monthly payrolled employees by age band; C: monthly inflows into payrolled employment, twelve-month trailing average). B: ONS Vacancy Survey VACS02, seasonally adjusted three-month rolling vacancies by SIC 2007 section, release of September 2026. Each series is divided by its 2022 average. Dashed line: 30 Nov 2022.

Figure A5: Payrolled employees, vacancies and hiring inflows, 2015–2026

Panels A and C: HMRC PAYE Real Time Information, seasonally adjusted, release of September 2026 (monthly payrolled employees by age band; monthly inflows into payrolled employment, twelve-month trailing average). Panel B: ONS Vacancy Survey, seasonally adjusted three-month rolling vacancies by SIC 2007 section, release of September 2026. Each series is divided by its 2022 average. Dashed line: 30 November 2022.


C The Ability-Based Exposure Regressions

Table A7: Regressions of young hiring on the Department for Education exposure score, 2015Q1–2026Q2

Outcome and specification	Coef. per SD	95% CI	<i>p</i> (RI <i>p</i>)
Share aged 22–25 among employees, main	-0.00243	[-0.005, +0.000]	0.069 (0.057)
Share aged 22–25 among new hires, main	-0.02092	[-0.033, -0.008]	0.001 (0.0006)
post from 2021Q1 (placebo timing)	-0.02511	[-0.043, -0.008]	0.005
dropping 2020–2022	-0.01976	[-0.032, -0.007]	0.003
first-wave interviews only	-0.02301	[-0.038, -0.008]	0.004
hires with a degree	-0.02673	[-0.058, +0.005]	0.096
hires without a degree	-0.01671	[-0.028, -0.006]	0.003
with exposure × linear trend	-0.02844	[-0.047, -0.010]	0.002
without IT occupations	-0.01943	[-0.032, -0.007]	0.003
placebo post from 2019Q1, 2015–2022 only	+0.00146	[-0.012, +0.014]	0.823
with major-group × quarter effects	-0.00913	[-0.034, +0.016]	0.473
with sub-major-group × quarter effects	+0.00220	[-0.029, +0.033]	0.889
excluding elementary occupations	-0.00894	[-0.021, +0.003]	0.149
exposure permuted within major groups (5,000 draws): permuted mean (sd)	-0.01678 (0.00319)		two-sided 0.193
top minus bottom exposure quintile, main (as a share of the 2022 mean of 0.1522)	-0.05650	-37.1%	
<i>Annual coefficients, share aged 22–25 among new hires, relative to the 2015–2022 mean</i>			
2015	+0.0035	[-0.005, +0.012]	
2016	-0.0022	[-0.012, +0.008]	
2017	+0.0010	[-0.012, +0.014]	
2018	-0.0062	[-0.018, +0.006]	
2019	+0.0007	[-0.011, +0.012]	
2020	-0.0068	[-0.019, +0.005]	
2021	-0.0008	[-0.012, +0.010]	
2022	+0.0108	[-0.002, +0.024]	
2023	-0.0328	[-0.054, -0.011]	
2024	-0.0138	[-0.031, +0.004]	
2025	-0.0202	[-0.035, -0.005]	
2026	-0.0091	[-0.031, +0.013]	

Note.—Coefficients on the standardised Department for Education language-model exposure score interacted with an indicator for 2023Q1 onward, in regressions on SOC 2010 unit group by quarter cells with unit-group and quarter effects. The first outcome is the share aged 22–25 among the unit group’s employees (16,739 cells, weighted by employment); the second is the share aged 22–25 among its new hires (10,575 cells, weighted by hires). Brackets: 95 percent intervals clustered by SOC 3-digit group (90 clusters); RI *p*: randomisation inference permuting the score across unit groups. The permutation within major groups reassigns scores among the unit groups of each major group (5,000 draws) and reports the mean and standard deviation of the permuted coefficients and the two-sided *p*-value of the observed coefficient around that mean. The annual coefficients come from the same design with exposure interacted with calendar year. Source: Quarterly Labour Force Survey.

Five exposure measures and the post-2022 change in young hiring and in job adverts, UK



Coefficients on the standardised measure x post (2023Q1 or January 2023 onward). A: share aged 22-25 among the unit group's new hires, LFS 2015Q1-2026Q2, cells weighted by the cell's hires, 95% intervals clustered by SOC 3-digit group. B: Poisson on weighted counts of new hires aged 22-25. C: log monthly online job adverts by SOC 2020 unit group (ONS, Textkernel), weighted by 2017-2022 adverts, clustered by minor group. Measures: DfE (2023) language-model exposure; Anthropic Economic Index observed exposure; Microsoft AI applicability score; Eloundou et al. GPT-4 beta; Claude.ai conversations from Great Britain per worker. US measures are moved to SOC 2010 through the O*NET-ESCO and ONS ISCO-08 conversions.

Figure A6: Five exposure measures and the post-2022 change in young hiring and in job adverts

Coefficients on the standardised measure interacted with the post indicator (2023Q1 or January 2023 onward), with 95 percent intervals. Panel A: share aged 22–25 among the unit group’s new hires, LFS 2015Q1–2026Q2, cells weighted by hires, clustered by SOC 3-digit group. Panel B: Poisson regression of weighted counts of new hires aged 22–25. Panel C: log monthly online job adverts by SOC 2020 unit group (ONS, Textkernel), weighted by 2017–2022 adverts, clustered by minor group. Measures: DfE (2023) language-model exposure; Anthropic Economic Index observed exposure; Microsoft AI applicability score; Eloundou et al. GPT-4 beta; Claude.ai conversations from Great Britain per worker. US measures are moved to SOC 2010 through the O*NET–ESCO and ONS ISCO-08 conversions.

D Robustness of the Adverts, Hiring and Industry Estimates

Table A8: Online job adverts and exposure: pre-trends, detrended estimates and sample variations

	DfE language-model	Anthropic observed	Eloundou GPT-4
Baseline (major-group $\times$ month effects, weighted)	-0.0835 [-0.164, -0.003]	-0.1555 [-0.231, -0.080]	-0.2270 [-0.321, -0.133]
Unweighted	-0.0638 [-0.129, +0.001]	-0.1381 [-0.197, -0.080]	-0.1705 [-0.243, -0.098]
Unit-group and month effects, unweighted	-0.0720 [-0.106, -0.038]	-0.1399 [-0.186, -0.094]	-0.1259 [-0.167, -0.085]
Measure $\times$ linear trend, 2017–2022 (per year)	-0.0260 [-0.043, -0.009]	-0.0155 [-0.026, -0.005]	-0.0237 [-0.037, -0.010]
Measure $\times$ linear trend, 2017–2019 (per year)	+0.0019 [-0.029, +0.033]	-0.0115 [-0.022, -0.001]	-0.0104 [-0.025, +0.004]
Post effect with measure $\times$ trend as a control	+0.0187 [-0.041, +0.078]	-0.0915 [-0.162, -0.021]	-0.1252 [-0.222, -0.028]
Post effect net of the 2017–2022 trend	+0.0403 [-0.041, +0.121]	-0.0816 [-0.157, -0.006]	-0.1138 [-0.208, -0.019]
Placebo break at January 2020, 2017–2022 only	-0.1163 [-0.172, -0.060]	-0.0600 [-0.095, -0.025]	-0.0937 [-0.143, -0.045]
Post effect, 2023	-0.0625 [-0.129, +0.004]	-0.1222 [-0.178, -0.067]	-0.1666 [-0.236, -0.097]
Post effect, 2024 onward	-0.0919 [-0.183, -0.001]	-0.1688 [-0.254, -0.084]	-0.2511 [-0.358, -0.144]
<i>Event study, measure $\times$ year, 2019 as the reference year</i>			
2017	-0.0098 [-0.074, +0.055]	+0.0187 [-0.004, +0.041]	+0.0182 [-0.012, +0.049]
2018	+0.0034 [-0.026, +0.033]	+0.0070 [-0.004, +0.018]	+0.0023 [-0.011, +0.016]
2020	-0.1591 [-0.232, -0.086]	-0.0751 [-0.115, -0.035]	-0.1166 [-0.179, -0.055]
2021	-0.1055 [-0.164, -0.047]	-0.0331 [-0.066, -0.001]	-0.0620 [-0.116, -0.008]
2022	-0.0909 [-0.154, -0.028]	-0.0461 [-0.083, -0.009]	-0.0821 [-0.134, -0.030]
2023	-0.1228 [-0.196, -0.049]	-0.1437 [-0.200, -0.088]	-0.2066 [-0.271, -0.142]
2024	-0.1812 [-0.270, -0.093]	-0.2224 [-0.312, -0.132]	-0.3379 [-0.437, -0.239]
2025	-0.1334 [-0.233, -0.034]	-0.1759 [-0.262, -0.090]	-0.2707 [-0.380, -0.162]
2026	-0.1314 [-0.244, -0.019]	-0.1561 [-0.239, -0.073]	-0.2409 [-0.353, -0.129]

Note.—Regressions of log monthly online job adverts by SOC 2020 unit group, January 2017 to August 2026, on the standardised measure interacted with the post indicator (January 2023 onward), with unit-group and major-group by month effects unless stated, weighted by 2017–2022 adverts unless stated, 95 percent intervals clustered by SOC 2020 minor group. The trend rows report the coefficient on the measure interacted with a linear trend in years, estimated on 2017–2022 (or 2017–2019) with the same effects. “Post effect net of the 2017–2022 trend” subtracts the estimated measure-specific trend from log adverts in every month before estimating the post effect. The placebo break sets the post indicator at January 2020 and uses 2017–2022 only. The event study interacts the measure with calendar year with 2019 as the reference. Source: ONS online job advert estimates.

Table A9: Online job adverts and observed exposure: leaving out one major group, and within single major groups

Major group	Leaving the group out		Within the group only	
	Coef. [95% CI]	<i>p</i>	Coef. [95% CI]	<i>p</i>
1 Managers	-0.158 [-0.237, -0.080]	<0.001	-0.095 [-0.145, -0.045]	0.002
2 Professional	-0.093 [-0.150, -0.037]	0.002	-0.213 [-0.332, -0.094]	0.001
3 Associate professional	-0.163 [-0.260, -0.067]	0.001	-0.134 [-0.220, -0.047]	0.004
4 Administrative and secretarial	-0.170 [-0.246, -0.094]	<0.001	+0.065 [-0.101, +0.231]	0.359
5 Skilled trades	-0.155 [-0.231, -0.079]	<0.001	-0.357 [-0.758, +0.044]	0.076
6 Caring, leisure and other service	-0.162 [-0.237, -0.087]	<0.001	+0.247 [-0.297, +0.791]	0.326
7 Sales and customer service	-0.158 [-0.239, -0.077]	<0.001	-0.129 [-0.223, -0.036]	0.018
8 Process, plant and machine operatives	-0.155 [-0.231, -0.080]	<0.001	-0.142 [-0.766, +0.481]	0.613
9 Elementary	-0.155 [-0.232, -0.079]	<0.001	-0.167 [-0.595, +0.261]	0.395

Note.—Coefficients on the Anthropic observed-exposure measure interacted with the post indicator (January 2023 onward) in regressions of log monthly adverts by SOC 2020 unit group, weighted by 2017–2022 adverts, with 95 percent intervals clustered by SOC 2020 minor group. Columns 1–2 leave out the stated major group and keep unit-group and major-group by month effects; columns 3–4 restrict the sample to the stated major group with unit-group and month effects. Source: ONS online job advert estimates.

Table A10: Online job adverts and exposure without the IT and customer-service unit groups

Measure	Sample	Month effects	Major-group $\times$ month	Net of the 2017–2022 trend
Ability-based (DfE)	All unit groups	-0.136 [-0.194, -0.077], $p < 0.001$	-0.083 [-0.164, -0.003], $p = 0.043$	+0.040 [-0.041, +0.121], $p = 0.325$
	Without the 12 IT groups	-0.109 [-0.148, -0.069], $p < 0.001$	-0.078 [-0.169, +0.013], $p = 0.092$	+0.041 [-0.050, +0.131], $p = 0.376$
	Without the 5 customer-service groups	-0.135 [-0.195, -0.075], $p < 0.001$	-0.079 [-0.160, +0.002], $p = 0.057$	+0.047 [-0.035, +0.128], $p = 0.257$
	Without both	-0.107 [-0.147, -0.067], $p < 0.001$	-0.073 [-0.164, +0.019], $p = 0.117$	+0.047 [-0.044, +0.138], $p = 0.309$
	Observed exposure	-0.170 [-0.239, -0.101], $p < 0.001$	-0.155 [-0.231, -0.080], $p < 0.001$	-0.082 [-0.157, -0.006], $p = 0.035$
Feasibility rating	All unit groups	-0.132 [-0.174, -0.091], $p < 0.001$	-0.104 [-0.164, -0.044], $p = 0.001$	-0.009 [-0.069, +0.051], $p = 0.766$
	Without the 12 IT groups	-0.176 [-0.246, -0.106], $p < 0.001$	-0.159 [-0.239, -0.079], $p < 0.001$	-0.082 [-0.161, -0.002], $p = 0.045$
	Without the 5 customer-service groups	-0.137 [-0.181, -0.092], $p < 0.001$	-0.104 [-0.170, -0.037], $p = 0.003$	-0.001 [-0.068, +0.066], $p = 0.973$
	Without both	-0.188 [-0.256, -0.121], $p < 0.001$	-0.227 [-0.321, -0.133], $p < 0.001$	-0.114 [-0.208, -0.019], $p = 0.019$
	Without both	-0.149 [-0.190, -0.108], $p < 0.001$	-0.172 [-0.259, -0.084], $p < 0.001$	-0.008 [-0.096, +0.080], $p = 0.860$
Applicability score	All unit groups	-0.189 [-0.258, -0.120], $p < 0.001$	-0.229 [-0.327, -0.131], $p < 0.001$	-0.112 [-0.210, -0.014], $p = 0.025$
	Without the 12 IT groups	-0.149 [-0.191, -0.107], $p < 0.001$	-0.171 [-0.266, -0.076], $p = 0.001$	+0.003 [-0.091, +0.098], $p = 0.944$
	Without the 5 customer-service groups	-0.135 [-0.229, -0.041], $p = 0.005$	-0.095 [-0.209, +0.019], $p = 0.100$	-0.007 [-0.121, +0.107], $p = 0.905$
	Without both	-0.083 [-0.137, -0.028], $p = 0.003$	-0.019 [-0.097, +0.059], $p = 0.628$	+0.075 [-0.003, +0.153], $p = 0.059$
	Without both	-0.140 [-0.241, -0.039], $p = 0.007$	-0.092 [-0.212, +0.028], $p = 0.132$	-0.001 [-0.120, +0.119], $p = 0.993$
Usage share	All unit groups	-0.082 [-0.142, -0.021], $p = 0.009$	-0.010 [-0.089, +0.070], $p = 0.813$	+0.089 [+0.010, +0.169], $p = 0.029$
	Without the 12 IT groups	-0.094 [-0.177, -0.012], $p = 0.025$	-0.026 [-0.092, +0.040], $p = 0.431$	+0.016 [-0.049, +0.082], $p = 0.624$
	Without the 5 customer-service groups	-0.049 [-0.093, -0.004], $p = 0.033$	+0.014 [-0.031, +0.059], $p = 0.530$	+0.054 [+0.009, +0.099], $p = 0.019$
	Without both	-0.098 [-0.181, -0.015], $p = 0.021$	-0.032 [-0.099, +0.035], $p = 0.349$	+0.013 [-0.054, +0.080], $p = 0.711$
	Without both	-0.052 [-0.097, -0.007], $p = 0.024$	+0.009 [-0.037, +0.056], $p = 0.689$	+0.051 [+0.005, +0.097], $p = 0.031$

Note.—Coefficients on the standardised exposure measure interacted with an indicator for January 2023 onward, in regressions of log monthly online job adverts by SOC 2020 unit group, January 2017 to August 2026, weighted by 2017–2022 mean adverts, with unit-group and month effects (column 1) or unit-group and major-group by month effects (column 2), and, in column 3, with major-group by month effects on log adverts minus the measure by linear-trend term estimated on 2017–2022 alone. The IT unit groups are SOC 2020 2131–2137, 2139, 2141 and 3131–3133 (twelve groups, 0.089 of 2022 adverts) and the customer-service groups 7211–7219 (five groups present, 0.020 of 2022 adverts). Brackets: 95 percent intervals clustered by SOC 2020 minor group. Source: ONS online job advert estimates.

Table A11: Young hiring and exposure within major groups: sample variations

Sample (major-group $\times$ quarter effects)	Share aged 22–25 among new hires		Young new-hire rate	
	DfE	Anthropic	DfE	Anthropic
Full sample	-0.0091 [-0.034, +0.016], $p=0.473$	-0.0167 [-0.034, +0.000], $p=0.054$	-0.0084 [-0.020, +0.003], $p=0.155$	-0.0086 [-0.016, -0.001], $p=0.031$
Without IT groups	-0.0081 [-0.033, +0.017], $p=0.524$	-0.0156 [-0.037, +0.006], $p=0.154$	-0.0078 [-0.019, +0.004], $p=0.183$	-0.0091 [-0.019, +0.001], $p=0.069$
Without customer-service groups	-0.0005 [-0.026, +0.025], $p=0.970$	-0.0088 [-0.025, +0.008], $p=0.298$	-0.0052 [-0.017, +0.007], $p=0.387$	-0.0051 [-0.012, +0.002], $p=0.169$
Without major group 7	-0.0045 [-0.031, +0.022], $p=0.734$	-0.0099 [-0.027, +0.007], $p=0.248$	-0.0068 [-0.019, +0.006], $p=0.279$	-0.0058 [-0.013, +0.002], $p=0.136$
Without IT and customer-service groups	+0.0007 [-0.025, +0.026], $p=0.959$	-0.0041 [-0.024, +0.016], $p=0.681$	-0.0045 [-0.016, +0.007], $p=0.445$	-0.0042 [-0.014, +0.005], $p=0.373$
Cells with three or more hires (columns 1–2) or young employees (columns 3–4)	-0.0144 [-0.046, +0.017], $p=0.364$	-0.0152 [-0.034, +0.004], $p=0.117$	-0.0071 [-0.020, +0.006], $p=0.294$	-0.0079 [-0.016, +0.000], $p=0.056$
Cells with five or more hires	-0.0143 [-0.060, +0.032], $p=0.536$	-0.0133 [-0.038, +0.012], $p=0.293$	–	–
Cells weighted by employment	+0.0004 [-0.039, +0.040], $p=0.986$	-0.0132 [-0.029, +0.003], $p=0.100$	–	–
Unweighted	-0.0044 [-0.033, +0.024], $p=0.762$	-0.0166 [-0.033, -0.000], $p=0.045$	-0.0058 [-0.018, +0.007], $p=0.353$	-0.0100 [-0.020, -0.000], $p=0.040$

Note.—Coefficients on the standardised measure interacted with the post indicator (2023Q1 onward) in regressions on unit group by quarter cells with unit-group and major-group by quarter effects, for the Department for Education score and the Anthropic observed-exposure measure. Columns 1–2: share aged 22–25 among the cell’s new hires, weighted by hires unless stated; columns 3–4: new hires aged 22–25 over employees aged 22–25, weighted by young employment unless stated. IT groups are SOC 2010 unit groups 2133–2139 and 3131–3132; customer-service groups are 7211, 7213, 7214, 7215 and 7219. Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey.

Table A12: Online job adverts and exposure with a teleworkability control

	(1) No control	(2) + Teleworkability × post	(3) + Teleworkability × year	(4) Both measures given 2017–2022 trends
<i>A. Log online job adverts, SOC 2020 unit groups, weighted by 2017–2022 mean adverts</i>				
Observed exposure: correlation with teleworkability 0.78, within major groups 0.65				
Within major groups	-0.155 [-0.231, -0.080]	-0.134 [-0.203, -0.065]	-0.134 [-0.203, -0.065]	–
Net of the 2017–2022 trend	-0.082 [-0.157, -0.006]	-0.060 [-0.130, +0.009]	-0.119 [-0.188, -0.049]	-0.124 [-0.194, -0.055]
Within major groups, without the 12 IT groups	-0.104 [-0.164, -0.044]	-0.091 [-0.154, -0.028]	-0.091 [-0.154, -0.028]	–
Net of the 2017–2022 trend, without the 12 IT groups	-0.009 [-0.069, +0.051]	+0.004 [-0.059, +0.067]	-0.049 [-0.112, +0.014]	-0.054 [-0.117, +0.009]
Feasibility rating: correlation with teleworkability 0.84, within major groups 0.66				
Within major groups	-0.227 [-0.321, -0.133]	-0.219 [-0.303, -0.134]	-0.219 [-0.303, -0.134]	–
Net of the 2017–2022 trend	-0.114 [-0.208, -0.019]	-0.105 [-0.190, -0.021]	-0.170 [-0.255, -0.086]	-0.177 [-0.262, -0.093]
Within major groups, without the 12 IT groups	-0.172 [-0.259, -0.084]	-0.171 [-0.248, -0.094]	-0.171 [-0.248, -0.094]	–
Net of the 2017–2022 trend, without the 12 IT groups	-0.008 [-0.096, +0.080]	-0.007 [-0.084, +0.070]	-0.064 [-0.141, +0.013]	-0.071 [-0.148, +0.006]
Ability-based (DfE): correlation with teleworkability 0.84, within major groups 0.63				
Within major groups	-0.083 [-0.164, -0.002]	+0.079 [-0.059, +0.217]	+0.079 [-0.059, +0.217]	–
Net of the 2017–2022 trend	+0.040 [-0.040, +0.121]	+0.203 [+0.065, +0.341]	+0.106 [-0.032, +0.244]	+0.097 [-0.041, +0.235]
Within major groups, without the 12 IT groups	-0.078 [-0.169, +0.013]	-0.011 [-0.107, +0.085]	-0.011 [-0.107, +0.085]	–
Net of the 2017–2022 trend, without the 12 IT groups	+0.041 [-0.050, +0.132]	+0.107 [+0.011, +0.203]	+0.011 [-0.085, +0.107]	0.000 [-0.096, +0.096]
<i>B. Log online job adverts, unweighted</i>				
Observed exposure: correlation with teleworkability 0.71, within major groups 0.55				
Within major groups	-0.138 [-0.196, -0.079]	-0.096 [-0.159, -0.034]	-0.096 [-0.159, -0.034]	–
Net of the 2017–2022 trend	-0.058 [-0.116, +0.001]	-0.016 [-0.079, +0.046]	-0.079 [-0.141, -0.016]	-0.082 [-0.144, -0.020]
Within major groups, without the 12 IT groups	-0.105 [-0.159, -0.052]	-0.061 [-0.118, -0.004]	-0.061 [-0.118, -0.004]	–
Net of the 2017–2022 trend, without the 12 IT groups	-0.026 [-0.080, +0.028]	+0.018 [-0.039, +0.076]	-0.047 [-0.104, +0.010]	-0.050 [-0.107, +0.007]
Feasibility rating: correlation with teleworkability 0.77, within major groups 0.55				
Within major groups	-0.170 [-0.243, -0.098]	-0.124 [-0.197, -0.050]	-0.124 [-0.197, -0.050]	–
Net of the 2017–2022 trend	-0.047 [-0.119, +0.025]	0.000 [-0.074, +0.073]	-0.063 [-0.137, +0.010]	-0.067 [-0.140, +0.007]
Within major groups, without the 12 IT groups	-0.133 [-0.201, -0.066]	-0.085 [-0.154, -0.017]	-0.085 [-0.154, -0.017]	–
Net of the 2017–2022 trend, without the 12 IT groups	-0.006 [-0.074, +0.061]	+0.042 [-0.027, +0.111]	-0.022 [-0.091, +0.047]	-0.025 [-0.094, +0.044]
Ability-based (DfE): correlation with teleworkability 0.79, within major groups 0.59				
Within major groups	-0.063 [-0.128, +0.002]	+0.046 [-0.015, +0.106]	+0.046 [-0.015, +0.106]	–
Net of the 2017–2022 trend	+0.004 [-0.061, +0.069]	+0.113 [+0.053, +0.174]	+0.015 [-0.046, +0.075]	+0.009 [-0.051, +0.070]
Within major groups, without the 12 IT groups	-0.058 [-0.121, +0.006]	+0.027 [-0.026, +0.081]	+0.027 [-0.026, +0.081]	–
Net of the 2017–2022 trend, without the 12 IT groups	+0.008 [-0.055, +0.071]	+0.093 [+0.039, +0.147]	-0.001 [-0.055, +0.052]	-0.006 [-0.060, +0.048]

Note.—Teleworkability is minus the ONS ability-to-homework score, standardised and carried to SOC 2020 as the exposure measures are (appendix A.2); its correlations with each measure are across the SOC 2020 unit groups of the adverts panel, weighted by 2017–2022 mean adverts in panel A and unweighted in panel B, overall and after removing major-group means. Coefficients on the standardised exposure measure interacted with an indicator for January 2023 onward in regressions of log monthly online job adverts, January 2017 to August 2026, on unit-group and major-group by month effects; the rows net of trend subtract the exposure by linear-trend term estimated on 2017–2022 before the post effect is estimated, as in appendix table A10. Column 2 adds teleworkability interacted with the post indicator; in the rows net of trend the exposure trend is the one estimated without teleworkability. Column 3 adds teleworkability interacted with calendar year, and in the rows net of trend the exposure trend is estimated with that path as a control. Because the panel is nearly balanced and its weights are fixed, exposure interacted with the post indicator has the same residual on either control, so columns 2 and 3 coincide in the rows within major groups and differ in the rows net of trend only through the exposure trend. Column 4 estimates linear trends in exposure and in teleworkability together on 2017–2022, removes both, and adds teleworkability interacted with the post indicator; removing the teleworkability trend changes the exposure coefficient by at most 0.00014, so column 4 too differs through the exposure trend alone. The estimates net of trend depend on the window on which the trend is estimated; with a 2017–2019 window, the observed-exposure estimate of column 3 is -0.042 (-0.111 to $+0.027$). Under other codings of teleworkability (its rank, its top third, quintile indicators, five factors, or a slope per major group) the within-group observed-exposure estimate without the IT groups has an interval that includes zero in several cases. All columns use the unit groups with both values, one fewer than appendix table A10: 406 for observed exposure and the feasibility rating and 407 for the ability-based measure, twelve fewer without the IT groups. Brackets: 95 percent intervals clustered by SOC 2020 minor group. Source: ONS online job advert estimates; Office for National Statistics (2020).

Table A13: Young hiring and observed exposure with a teleworkability control

	(1) No control	(2) + Teleworkability × post	(3) + Teleworkability × year	(4) Both measures given 2015–2022 trends
Share aged 22–25 among new hires				
Weighted by hires	-0.0167 [-0.0336, +0.0003]	-0.0238 [-0.0428, -0.0049]	-0.0236 [-0.0425, -0.0047]	–
Unweighted	-0.0166 [-0.0328, -0.0004]	-0.0254 [-0.0432, -0.0077]	-0.0252 [-0.0429, -0.0074]	–
Without the customer-service groups, weighted	-0.0088 [-0.0254, +0.0079]	-0.0156 [-0.0333, +0.0022]	–	–
Young new-hire rate				
Weighted by young employment	-0.0086 [-0.0164, -0.0008]	-0.0115 [-0.0209, -0.0021]	-0.0116 [-0.0210, -0.0022]	–
Unweighted	-0.0100 [-0.0196, -0.0004]	-0.0135 [-0.0248, -0.0022]	-0.0136 [-0.0249, -0.0023]	–
Young hires (Poisson)				
Weighted counts	-0.121 [-0.247, +0.006]	-0.162 [-0.308, -0.015]	-0.162 [-0.309, -0.016]	–
Interview counts	-0.091 [-0.230, +0.049]	-0.135 [-0.302, +0.032]	-0.134 [-0.302, +0.033]	–
Young hires (Poisson), net of the 2015–2022 trend				
Weighted counts	-0.135 [-0.262, -0.008]	-0.176 [-0.323, -0.030]	-0.163 [-0.310, -0.017]	-0.171 [-0.317, -0.024]
Interview counts	-0.079 [-0.218, +0.061]	-0.123 [-0.290, +0.045]	-0.100 [-0.267, +0.067]	-0.110 [-0.277, +0.058]

Note.—Coefficients on standardised observed exposure interacted with an indicator for 2023Q1 onward in the young-hiring regressions of table 6, on SOC 2010 unit group by quarter cells, 2015Q1–2026Q2, with unit-group and major-group by quarter effects, without and with the teleworkability controls of appendix table A12. The share rows weight cells by hires and the hire-rate rows by young employment unless stated; the Poisson rows use the weighted count of new hires aged 22–25 or the interview count. The rows net of trend hold the exposure trend estimated on 2015–2022 fixed as an offset, as in table 7; column 4 estimates the exposure and teleworkability trends together on 2015–2022 and holds both fixed. The customer-service groups are those of appendix table A11. All columns use the unit groups with a teleworkability value, 363 of the 364 for the share and the hire rate and all 330 for the Poisson rows. Brackets: 95 percent intervals clustered by SOC 2010 3-digit group. Source: Quarterly Labour Force Survey; Office for National Statistics (2020).

Table A14: Online job adverts and observed exposure by year, with a teleworkability path

	All unit groups		Without the 12 IT groups	
	No control	With teleworkability × year	No control	With teleworkability × year
<i>Observed exposure × year, 2019 = 0</i>				
2017	+0.019 [-0.004, +0.041]	+0.035 [+0.004, +0.065]	+0.020 [-0.011, +0.052]	+0.042 [+0.007, +0.077]
2018	+0.007 [-0.004, +0.018]	+0.003 [-0.013, +0.020]	+0.009 [-0.005, +0.024]	+0.007 [-0.013, +0.026]
2020	-0.075 [-0.115, -0.035]	-0.009 [-0.063, +0.045]	-0.100 [-0.145, -0.054]	-0.031 [-0.093, +0.030]
2021	-0.033 [-0.066, -0.001]	+0.032 [-0.017, +0.080]	-0.054 [-0.094, -0.014]	+0.012 [-0.045, +0.069]
2022	-0.046 [-0.083, -0.009]	+0.009 [-0.044, +0.061]	-0.055 [-0.106, -0.003]	-0.003 [-0.072, +0.066]
2023	-0.144 [-0.200, -0.088]	-0.088 [-0.150, -0.027]	-0.112 [-0.170, -0.054]	-0.062 [-0.136, +0.013]
2024	-0.222 [-0.312, -0.132]	-0.159 [-0.254, -0.064]	-0.162 [-0.236, -0.089]	-0.106 [-0.199, -0.013]
2025	-0.176 [-0.262, -0.090]	-0.130 [-0.206, -0.053]	-0.131 [-0.212, -0.051]	-0.093 [-0.173, -0.013]
2026	-0.156 [-0.239, -0.073]	-0.107 [-0.175, -0.038]	-0.125 [-0.210, -0.040]	-0.084 [-0.163, -0.006]
Exposure × linear trend, 2017–2022 (per year)	-0.0155	-0.0033	-0.0199	-0.0088
<i>Coefficient minus the 2017–2022 trend line</i>				
2023	-0.068	-0.089	-0.013	-0.035
2024	-0.131	-0.156	-0.043	-0.071
2025	-0.069	-0.123	+0.008	-0.049
2026	-0.034	-0.097	+0.034	-0.031
<i>Teleworkability × year, 2019 = 0</i>				
2017		-0.030 [-0.067, +0.007]		-0.039 [-0.076, -0.002]
2018		+0.007 [-0.013, +0.028]		+0.005 [-0.016, +0.026]
2020		-0.125 [-0.189, -0.061]		-0.123 [-0.190, -0.057]
2021		-0.122 [-0.177, -0.068]		-0.119 [-0.175, -0.062]
2022		-0.103 [-0.167, -0.039]		-0.093 [-0.158, -0.028]
2023		-0.104 [-0.173, -0.035]		-0.091 [-0.161, -0.020]
2024		-0.120 [-0.211, -0.029]		-0.102 [-0.194, -0.009]
2025		-0.087 [-0.203, +0.029]		-0.069 [-0.183, +0.045]
2026		-0.093 [-0.217, +0.032]		-0.074 [-0.196, +0.047]

Note.—Coefficients on standardised observed exposure interacted with calendar year (2019 omitted) in regressions of log monthly online job adverts by SOC 2020 unit group on unit-group and major-group by month effects, weighted by 2017–2022 mean adverts, without and with teleworkability interacted with calendar year (2019 omitted), on the 406 unit groups with both values (394 without the IT groups). The trend row is the exposure by linear-trend coefficient estimated on 2017–2022 with the same effects and, in the columns with the path, with teleworkability by year as a control. The gaps are each year’s coefficient minus that trend drawn through the mean of the 2017–2022 coefficients, as in figure 6. Brackets: 95 percent intervals clustered by SOC 2020 minor group. Source: ONS online job advert estimates; Office for National Statistics (2020).


Table A15: Employer adoption and the young new-hire rate: leave-one-out and placebo estimates

Young new-hire rate on adoption $\times$ post	Coef. [95% CI]	<i>p</i>
Baseline	-0.0046 [-0.009, -0.000]	0.047
Unweighted	-0.0046 [-0.009, +0.000]	0.055
Without J and P	+0.0005 [-0.006, +0.007]	0.861
Placebo break at 2020Q1, 2015–2022 only	-0.0012 [-0.009, +0.006]	0.728
<i>Leaving out one section</i>		
without C	-0.0053 [-0.010, -0.001]	0.024
without F	-0.0040 [-0.009, +0.001]	0.095
without G	-0.0045 [-0.009, +0.000]	0.057
without H	-0.0051 [-0.010, -0.001]	0.032
without I	-0.0044 [-0.009, +0.001]	0.078
without J	-0.0025 [-0.009, +0.004]	0.407
without L	-0.0046 [-0.009, +0.000]	0.052
without M	-0.0061 [-0.010, -0.002]	0.004
without N	-0.0049 [-0.009, -0.000]	0.037
without P	-0.0038 [-0.009, +0.002]	0.163
without Q	-0.0050 [-0.010, -0.000]	0.035
without R	-0.0048 [-0.009, -0.000]	0.046
without S	-0.0048 [-0.009, -0.000]	0.044

Note.—Coefficients on standardised adoption (mean of the twelve BICS waves) interacted with the post indicator in regressions of the young new-hire rate on section and quarter effects, weighted by young employment unless stated, with 95 percent intervals clustered by section (t with $G - 1$ degrees of freedom). Section letters as in figure 7. Source: Quarterly Labour Force Survey, BICS.

E High- and Low-Exposure Occupations: Levels and Alternative Cuts

Young workers in high- and low-exposure occupations, UK, 2015-2025




Quarterly Labour Force Survey (ONS), person weights, and ONS online job adverts. High and low exposure: top and bottom thirds of SOC 2010 unit groups by Anthropic observed exposure, cut where cumulative 2015-2022 employment reaches each third. New hires: employees with under three months of tenure. A-C, F: annual values. D: new hires aged 22-25 with a degree. E: annual adverts of the group's SOC 2020 unit groups. Dashed line: 30 Nov 2022.

Figure A7: Young workers in high- and low-exposure occupations, levels, 2015–2025

The series behind figure 5 in levels. High and low exposure are the top and bottom thirds of SOC 2010 unit groups by Anthropic observed exposure, cut at the exposure values where cumulative 2015–2022 employment of employees aged 16–64 reaches each third. Panel A: share aged 22–25 among new hires of all ages. Panel B: new hires aged 22–25 per 100 employees aged 22–25. Panel C: employees aged 22–25, 2022 = 100. Panel D: among new hires aged 22–25 with a degree, the share starting in a high-exposure and in a low-exposure occupation. Panel E: for each of the group's SOC 2020 unit groups the mean monthly adverts of the year, summed over the unit groups, 2022 = 100. Panel F: employees aged 16–64, 2022 = 100. Person weights; 2026 omitted. Dashed line: 30 November 2022. Source: Quarterly Labour Force Survey; ONS online job advert estimates.

High- minus low-exposure occupations under four cuts of the exposure distribution, UK, 2015-2025 (fifths: top v bottom fifth; thirds: top v bottom third; halves: above v below the median)



Quarterly Labour Force Survey (ONS), person weights; ONS online job adverts. Each panel: event-study coefficient on high $\times$ year with unit-group and quarter (adverts: month) effects, cells weighted by weighted new hires (first column), weighted employees aged 22–25 (second) or 2017–2022 mean adverts (fourth), re-centred on the mean of the 2015–2022 (adverts 2017–2022) coefficients, shaded 95% interval clustered by SOC 3-digit group (graduates: person level, household clusters). Cuts of SOC 2010 unit groups by Anthropic observed exposure at employment-weighted quantiles. Dashed line: 30 Nov 2022. Source: Quarterly Labour Force Survey; ONS online job advert estimates.

Figure A8: High- minus low-exposure occupations under four cuts of the exposure distribution, 2015–2025

Each panel shows the coefficient on high $\times$ year in an event-study regression with unit-group and quarter (adverts: month) effects, cells weighted by weighted new hires (first column), weighted employees aged 22–25 (second) or 2017–2022 mean adverts (fourth), re-centred on the mean of the 2015–2022 (adverts 2017–2022) coefficients, with its 95 percent interval clustered by SOC 3-digit group; the graduate column is a person-level regression of an indicator (1 for a start in the high group, –1 in the low group, 0 otherwise) on year indicators alone, person weights, clustered by household within quarter. Rows: the cut of SOC 2010 unit groups by Anthropic observed exposure at employment-weighted quantiles, top fifth against bottom fifth, top third against bottom third (the cut used in figure 5), above against below the median, and top fifth against all other unit groups. Columns: the share aged 22–25 among new hires; new hires aged 22–25 per 100 employees aged 22–25; among graduate new hires aged 22–25, the share starting in the high group minus the share starting in the low group; log online adverts. Dashed line: 30 November 2022. Source: Quarterly Labour Force Survey; ONS online job advert estimates.

F The Layers above Entry: Mid-Career and Senior Staff


This appendix reports the secondary analysis summarised in section 5.5. The claim that language models reshape the pyramid of employment has two halves, one that the entry layer of exposed occupations thins, which section 4 documents, the other concerning the layers above it. One reading is that they hold or grow, because experienced workers are complemented by the technology and are the ones who use it. Another is that they too are displaced once the technology does the work of a team. This appendix tests both readings with the same exposure measures and designs, applied to employees of all ages, looking at employment, hiring, pay and hours by age band, at the seniority structure inside occupations, at the exits of mid-career workers, and at the demand for senior and junior roles in adverts. A classification change falls in the middle of this material, because the LFS codes occupations to SOC 2010 until 2020 and to SOC 2020 from 2021. For current jobs it continues to derive a SOC 2010 code after 2020, but the last job of the non-employed and the two-quarter longitudinal files have only the classification current at the interview. Where the appendix uses those variables at the unit-group or minor-group level, SOC 2020 codes are mapped to SOC 2010 by the modal mapping of the ONS relationship table, which gives one classification for 2015–2026. Where it uses major groups, which correspond across the two classifications, no mapping is needed. Keeping the two classifications as separate units, which has no quarter in common between them, identifies the post-2022 change from 2021–2026 alone. The appendix reports what that alternative gives and why it differs.

Employment by age band

Figure A9 and table A16 describe employment in exposed occupations by age band, of which panel C of the figure gives the simplest view. With 2022 as 100, the weighted headcount in the top quintile of observed exposure stood in 2025 at 85 for ages 16–21, 100 for 22–25, 111 for 26–34, 109 for 35–49 and 107 for 50–64, against 107, 109, 99, 107 and 102 in the bottom quintile. Exposed occupations lost young workers and gained workers over 25 relative to unexposed ones. Poisson regressions of employee counts by unit group and quarter, one per age band, with unit-group and major-group by quarter effects, put the post-2022 change per standard deviation of observed exposure at -0.044 for ages 16–21 ($p = 0.055$), $+0.023$ for 22–25 ($p = 0.40$), $+0.066$ for 26–34, $+0.044$ for 35–49 and $+0.068$ for 50–64, the last three with $p < 0.005$ (table A16). For the three older bands the estimates are close to these on unweighted interview counts (0.061, 0.042 and 0.068) and among graduates alone (0.063,

0.040 and 0.071). For ages 16–21 interview counts give -0.048 ($p = 0.039$) and the few graduates -0.076 ($p = 0.36$).

Employment by age band in occupations exposed to language models, UK, 2015–2026



Quarterly Labour Force Survey (ONS), employees aged 16–64 with a SOC 2010 unit group, person weights. A, B: coefficients on the standardised exposure measure x calendar year (2022 omitted) in Poisson regressions of weighted employee counts by unit group and quarter, one regression per age band, with unit-group and major-group x quarter effects; bars: 95% intervals clustered by SOC 3-digit group; 2026 covers January–June. C: weighted headcount in the bottom and top quintiles of Anthropic observed exposure (quintiles weighted by 2015–2022 employment), mean of the four quarters of 2025 over the 2022 mean.

Figure A9: Employment by age band in occupations exposed to language models, 2015–2026

Employees aged 16–64 with a SOC 2010 unit group, person weights. Panels A and B: coefficients on the standardised exposure measure interacted with calendar year (2022 omitted) in Poisson regressions of weighted employee counts by unit group and quarter, one regression per age band, with unit-group and major-group by quarter effects; bars are 95 percent intervals clustered by SOC 3-digit group; 2026 covers January–June. Panel C: weighted headcount in the bottom and top quintiles of Anthropic observed exposure, mean of the four quarters of 2025 over the 2022 mean. Quintiles of unit groups are cut at the exposure values at which cumulative 2015–2022 employment reaches each fifth. Source: Quarterly Labour Force Survey.

Table A16: Employment in exposed occupations by age band, 2015Q1–2026Q2

Age band	Employment 2025, 2022 = 100		Poisson, Anthropic observed exposure			DfE	
	Bottom quintile	Top quintile	Main	Major × quarter	Net of trend	Major × quarter	
16-21	106.9	84.8	-0.1052 [-0.1699, -0.0405], $p=0.001$	-0.0444 [-0.0898, +0.0010], $p=0.055$	-0.1486 [-0.1950, -0.1023], $p < 0.001$	-0.0004 [-0.0796, +0.0789], $p=0.993$	
22-25	109.4	99.8	+0.0117 [-0.0617, +0.0851], $p=0.754$	+0.0232 [-0.0310, +0.0775], $p=0.401$	-0.0353 [-0.1083, +0.0377], $p=0.343$	+0.0804 [-0.0080, +0.1688], $p=0.075$	
26-34	98.7	111.0	+0.1086 [+0.0503, +0.1668], $p < 0.001$	+0.0658 [+0.0215, +0.1100], $p=0.004$	+0.0423 [-0.0091, +0.0937], $p=0.107$	+0.0510 [+0.0099, +0.0922], $p=0.015$	
35-49	107.4	109.2	+0.0673 [+0.0212, +0.1133], $p=0.004$	+0.0437 [+0.0149, +0.0726], $p=0.003$	-0.0048 [-0.0412, +0.0316], $p=0.796$	+0.0742 [+0.0326, +0.1157], $p < 0.001$	
50-64	101.6	107.2	+0.0749 [+0.0178, +0.1321], $p=0.010$	+0.0683 [+0.0247, +0.1119], $p=0.002$	+0.0164 [-0.0218, +0.0546], $p=0.400$	+0.0839 [+0.0373, +0.1304], $p < 0.001$	

Note.—Columns 1–2: weighted headcount in the bottom and top quintiles of Anthropic observed exposure, mean of the four quarters of 2025 with the 2022 mean set to 100; quintiles of unit groups are cut at the exposure values at which cumulative 2015–2022 employment reaches each fifth. Columns 3–6: coefficients on the standardised exposure measure interacted with an indicator for 2023Q1 onward in Poisson regressions of weighted employee counts (thousands) by SOC 2010 unit group and quarter, one regression per age band, over a complete grid of unit groups and quarters; “main” has unit-group and quarter effects, the other columns unit-group and major-group by quarter effects; “net of trend” adds the measure interacted with a linear trend in years as a control. Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey.

The event studies in panels A and B of the figure show where the gains of the older bands

come from, since their coefficients climb through the pre-period and flatten after 2022. For ages 35–49 the observed-exposure coefficient relative to 2022 is -0.074 in 2015, -0.021 in 2019 and -0.005 in 2021, then $+0.024$, $+0.016$, $+0.008$ and $+0.008$ in 2023–2026. For ages 50–64 it is -0.092 in 2015 and -0.036 in 2021, then $+0.013$, $+0.032$, $+0.002$ and -0.004 . Exposed occupations had been drawing in older workers relative to their broad groups for a decade, a drift that the post-2022 coefficients record. When the measure interacted with a linear trend is added as a control, the post-2022 effect for ages 35–49 is -0.005 ($p = 0.80$), for ages 50–64 $+0.016$ ($p = 0.40$) and for ages 26–34 $+0.042$ ($p = 0.11$) (table A16). The one band whose post-2022 path departs from its trend is the youngest. Ages 16–21 are at -0.153 , -0.122 and -0.113 in 2023–2025 relative to 2022, after values between -0.12 and 0.00 in 2015–2021 (table A17), and at -0.149 net of the trend. The ability-based measure gives the same shape with steeper pre-trends, net of which only the fall for ages 16–21 remains (-0.109 , $p = 0.041$).

Table A17: Event studies of employment by age band in exposed occupations

Year	16-21	22-25	26-34	35-49	50-64
<i>A. Anthropic observed exposure</i>					
2015	-0.122 [-0.187, -0.058]	-0.095 [-0.185, -0.006]	-0.039 [-0.094, +0.015]	-0.074 [-0.136, -0.011]	-0.092 [-0.149, -0.035]
2016	-0.089 [-0.164, -0.014]	-0.091 [-0.179, -0.003]	-0.038 [-0.089, +0.013]	-0.042 [-0.108, +0.024]	-0.073 [-0.140, -0.007]
2017	-0.123 [-0.203, -0.042]	-0.061 [-0.139, +0.017]	-0.047 [-0.104, +0.010]	-0.046 [-0.118, +0.026]	-0.067 [-0.137, +0.003]
2018	-0.078 [-0.176, +0.021]	-0.022 [-0.090, +0.045]	-0.040 [-0.107, +0.028]	-0.046 [-0.105, +0.013]	-0.072 [-0.126, -0.018]
2019	-0.071 [-0.141, -0.001]	-0.018 [-0.092, +0.055]	-0.033 [-0.097, +0.032]	-0.021 [-0.082, +0.040]	-0.064 [-0.114, -0.013]
2020	-0.080 [-0.168, +0.009]	-0.051 [-0.119, +0.018]	-0.015 [-0.074, +0.044]	-0.009 [-0.061, +0.043]	-0.052 [-0.095, -0.009]
2021	-0.001 [-0.088, +0.087]	-0.031 [-0.089, +0.027]	-0.013 [-0.050, +0.024]	-0.005 [-0.039, +0.029]	-0.036 [-0.062, -0.010]
2023	-0.153 [-0.254, -0.052]	-0.011 [-0.118, +0.096]	+0.070 [+0.017, +0.122]	+0.024 [-0.013, +0.060]	+0.013 [-0.024, +0.050]
2024	-0.122 [-0.210, -0.033]	-0.029 [-0.099, +0.041]	+0.024 [-0.031, +0.079]	+0.016 [-0.028, +0.060]	+0.032 [-0.008, +0.072]
2025	-0.113 [-0.179, -0.048]	-0.015 [-0.079, +0.049]	+0.024 [-0.021, +0.068]	+0.008 [-0.034, +0.050]	+0.002 [-0.035, +0.039]
2026	-0.034 [-0.116, +0.049]	-0.046 [-0.105, +0.012]	+0.036 [-0.012, +0.083]	+0.008 [-0.036, +0.051]	-0.004 [-0.047, +0.038]
<i>B. DfE language-model exposure</i>					
2015	-0.180 [-0.307, -0.054]	-0.160 [-0.300, -0.020]	-0.111 [-0.209, -0.012]	-0.113 [-0.202, -0.024]	-0.099 [-0.167, -0.032]
2016	-0.095 [-0.207, +0.018]	-0.162 [-0.294, -0.031]	-0.119 [-0.207, -0.031]	-0.076 [-0.157, +0.005]	-0.077 [-0.155, +0.000]
2017	-0.113 [-0.202, -0.023]	-0.175 [-0.319, -0.030]	-0.126 [-0.217, -0.034]	-0.052 [-0.136, +0.032]	-0.083 [-0.161, -0.004]
2018	-0.104 [-0.215, +0.007]	-0.110 [-0.220, -0.001]	-0.130 [-0.256, -0.003]	-0.055 [-0.126, +0.017]	-0.105 [-0.177, -0.032]
2019	-0.113 [-0.225, -0.001]	-0.118 [-0.243, +0.008]	-0.100 [-0.211, +0.010]	-0.019 [-0.089, +0.052]	-0.083 [-0.149, -0.018]
2020	-0.146 [-0.299, +0.006]	-0.164 [-0.298, -0.030]	-0.043 [-0.138, +0.051]	-0.008 [-0.084, +0.068]	-0.047 [-0.112, +0.018]
2021	-0.052 [-0.179, +0.075]	-0.093 [-0.188, +0.001]	-0.005 [-0.066, +0.056]	+0.026 [-0.021, +0.073]	-0.027 [-0.070, +0.016]
2023	-0.135 [-0.240, -0.031]	-0.071 [-0.181, +0.039]	+0.002 [-0.077, +0.082]	+0.038 [-0.008, +0.083]	+0.033 [-0.026, +0.092]
2024	-0.152 [-0.247, -0.057]	-0.042 [-0.142, +0.057]	-0.055 [-0.136, +0.026]	+0.040 [-0.034, +0.115]	+0.021 [-0.039, +0.081]
2025	-0.053 [-0.167, +0.061]	-0.020 [-0.122, +0.081]	-0.038 [-0.106, +0.030]	+0.040 [-0.026, +0.106]	+0.019 [-0.040, +0.077]
2026	-0.029 [-0.132, +0.075]	-0.035 [-0.148, +0.079]	-0.004 [-0.077, +0.069]	+0.029 [-0.042, +0.101]	-0.006 [-0.070, +0.059]

Note.—Coefficients on the standardised exposure measure interacted with calendar year (2022 omitted) in Poisson regressions of weighted employee counts by SOC 2010 unit group and quarter, one regression per age band, with unit-group and major-group by quarter effects; 95 percent intervals clustered by SOC 3-digit group; 2026 covers January–June. Source: Quarterly Labour Force Survey.

Two further measures qualify the asymmetry. The age composition of exposed occupa-

tions shifted upward after 2022 relative to their broad groups, by 0.6 points in the share aged 50–64 ($p = 0.002$) and 0.16 years in mean age ($p = 0.010$) per standard deviation of observed exposure, with no change in the ratio of employees aged 22–29 to those aged 35–49. When the cells are given equal weight the two estimates are -0.2 points ($p = 0.65$) and 0.07 years ($p = 0.45$). The shift therefore comes from the largest unit groups (table A18). Hiring of mid-career workers did not fall in a way that survives a change of design. The new-hire rate of employees aged 35–49 changed by -0.0012 per standard deviation within broad groups ($p = 0.046$), 4.6 percent of its pre-period mean of 0.0253, and by -0.0002 ($p = 0.75$) with quarter effects alone. For ages 22–25 the two designs give falls of 0.0086 ($p = 0.031$) and 0.0107 ($p < 0.001$) on a mean of 0.0735. Ages 26–34 and 50–64 did not move (table A19). Prices did not move between the bands, since within exposed occupations the hourly pay of 35–49-year-olds changed by -0.010 log points relative to that of 22–25-year-olds ($p = 0.29$). Usual weekly hours of employees aged 35–49 fell by 0.53 hours per standard deviation relative to those of employees aged 22–25 ($p = 0.002$), the opposite of what absorbing the work of absent juniors would produce. The difference combines a rise of 0.26 hours for 22–25-year-olds in exposed occupations ($p = 0.022$) with a fall of 0.27 for 35–49-year-olds ($p = 0.055$), both with quarter effects. Within broad groups the two are 0.10 ($p = 0.54$) and -0.22 ($p = 0.11$).

Table A18: Age composition of exposed occupations, weighted and unweighted

Outcome (unit group $\times$ quarter cells)	Mean 2015–2022	Main	Major $\times$ quarter	Major $\times$ quarter, unweighted
Share of employees under 30	0.258	-0.0049 [-0.0115, +0.0016], $p=0.139$	-0.0017 [-0.0071, +0.0037], $p=0.540$	-0.0042 [-0.0100, +0.0016], $p=0.156$
Share of employees aged 50–64	0.273	+0.0041 [+0.0001, +0.0082], $p=0.046$	+0.0060 [+0.0022, +0.0099], $p=0.002$	-0.0017 [-0.0092, +0.0058], $p=0.647$
Mean age of employees (years)	39.876	+0.186 [+0.031, +0.341], $p=0.019$	+0.165 [+0.040, +0.290], $p=0.010$	+0.073 [-0.117, +0.263], $p=0.447$
Log employees aged 22–29 over employees aged 35–49	-0.667	+0.008 [-0.036, +0.052], $p=0.714$	+0.020 [-0.026, +0.066], $p=0.395$	-0.003 [-0.040, +0.035], $p=0.894$

Note.—Unit group by quarter cells of employees aged 16–64 with a SOC 2010 unit group, person weights within cells. Coefficients on Anthropic observed exposure (standardised) interacted with an indicator for 2023Q1 onward, with unit-group and quarter effects (main) or unit-group and major-group by quarter effects, cells weighted by employment or, in the last column, unweighted. Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey.

Table A19: New-hire rates, pay and hours by age band in exposed occupations

Age band	New-hire rate		Log hourly pay		Usual weekly hours	
	Main	Major $\times$ quarter	Main	Major $\times$ quarter	Main	Major $\times$ quarter
16-21	-0.00008	-0.01194	-0.0099	-0.0158	+0.223	-0.157
	[-0.01182, +0.01166], $p=0.989$	[-0.02216, -0.00172], $p=0.023$	[-0.0299, +0.0100], $p=0.325$	[-0.0380, +0.0065], $p=0.162$	[+0.016, +0.430], $p=0.035$	[-0.529, +0.216], $p=0.405$
22-25	-0.01068	-0.00859	-0.0098	+0.0089	+0.257	+0.097
	[-0.01575, -0.00561], $p < 0.001$	[-0.01640, -0.00078], $p=0.031$	[-0.0289, +0.0094], $p=0.314$	[-0.0114, +0.0293], $p=0.385$	[+0.038, +0.475], $p=0.022$	[-0.214, +0.408], $p=0.537$
26-34	-0.00069	+0.00085	-0.0100	-0.0002	-0.237	-0.259
	[-0.00267, +0.00128], $p=0.488$	[-0.00077, +0.00248], $p=0.300$	[-0.0224, +0.0024], $p=0.113$	[-0.0154, +0.0151], $p=0.984$	[-0.565, +0.091], $p=0.155$	[-0.552, +0.034], $p=0.082$
35-49	-0.00020	-0.00115	-0.0194	-0.0065	-0.271	-0.222
	[-0.00145, +0.00104], $p=0.749$	[-0.00229, -0.00002], $p=0.046$	[-0.0297, -0.0092], $p < 0.001$	[-0.0146, +0.0015], $p=0.110$	[-0.549, +0.006], $p=0.055$	[-0.496, +0.051], $p=0.110$
50-64	-0.00020	+0.00008	-0.0116	+0.0027	-0.076	-0.103
	[-0.00166, +0.00126], $p=0.789$	[-0.00123, +0.00138], $p=0.908$	[-0.0242, +0.0010], $p=0.071$	[-0.0060, +0.0115], $p=0.536$	[-0.245, +0.093], $p=0.374$	[-0.288, +0.081], $p=0.269$
35-49 minus 22-25, pooled			-0.0097			-0.528
			[-0.0275, +0.0082], $p=0.288$			[-0.868, -0.188], $p=0.002$

Note.—Coefficients on Anthropic observed exposure (standardised) interacted with an indicator for 2023Q1 onward. New-hire rate: new hires aged in the band over employees in the band, by unit group and quarter, weighted by the band’s employment. Log hourly pay: employees with derived hourly pay, income weights, with unit-group, quarter and single-year-of-age effects. Usual weekly hours: person weights, same effects. “Main” uses quarter effects, “major $\times$ quarter” major-group by quarter effects. The last row is the difference between the 35–49 and 22–25 coefficients from a pooled model with band-specific occupation and quarter effects. Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey.


Seniority inside occupations

The LFS asks every employee whether they are a manager, a foreman or supervisor, or neither, whether they supervise other employees, and how long they have been with their employer. In the April–June quarters it also asks about managerial duties a year earlier. Table A20 uses these to describe the layers inside unit groups.

Table A20: Seniority structure inside occupations: supervision, tenure and promotion, 2015Q1–2026Q2

Outcome (unit group $\times$ quarter cells)	Mean 2015–2022	Anthropic observed exposure		DfE
		Main	Major $\times$ quarter	Major $\times$ quarter
Share with managerial or supervisory status	0.357	+0.0000 [-0.0056, +0.0057], $p=0.991$	+0.0041 [-0.0027, +0.0110], $p=0.235$	-0.0026 [-0.0100, +0.0048], $p=0.485$
Share responsible for supervising others	0.363	+0.0027 [-0.0026, +0.0080], $p=0.316$	+0.0064 [-0.0005, +0.0134], $p=0.069$	+0.0026 [-0.0057, +0.0110], $p=0.530$
Log non-supervisory over supervisory employees	0.751	-0.0110 [-0.0430, +0.0210], $p=0.495$	-0.0239 [-0.0554, +0.0077], $p=0.136$	-0.0034 [-0.0467, +0.0400], $p=0.877$
Share with under one year of tenure	0.171	-0.0067 [-0.0101, -0.0034], $p < 0.001$	-0.0060 [-0.0102, -0.0018], $p=0.005$	+0.0031 [-0.0036, +0.0098], $p=0.360$
Share with one to five years	0.355	+0.0007 [-0.0044, +0.0057], $p=0.793$	+0.0013 [-0.0063, +0.0089], $p=0.734$	-0.0030 [-0.0138, +0.0078], $p=0.580$
Share with five to ten years	0.172	+0.0032 [+0.0001, +0.0062], $p=0.044$	-0.0003 [-0.0037, +0.0032], $p=0.885$	-0.0014 [-0.0080, +0.0053], $p=0.688$
Share with ten or more years	0.302	+0.0029 [-0.0010, +0.0068], $p=0.143$	+0.0050 [-0.0019, +0.0118], $p=0.152$	+0.0013 [-0.0127, +0.0152], $p=0.859$
Mean tenure (years)	7.948	+0.0679 [+0.0163, +0.1195], $p=0.010$	+0.0894 [+0.0079, +0.1709], $p=0.032$	-0.0232 [-0.1896, +0.1432], $p=0.783$
Promotion: duties now, none a year ago (April–June, individual level)	0.068	-0.0014 [-0.0059, +0.0032], $p=0.559$	+0.0014 [-0.0050, +0.0079], $p=0.663$	-0.0014 [-0.0069, +0.0042], $p=0.623$
Loss of duties: none now, duties a year ago	0.073	+0.0012 [-0.0043, +0.0067], $p=0.668$	+0.0011 [-0.0049, +0.0072], $p=0.717$	+0.0039 [-0.0044, +0.0122], $p=0.356$
Share in NS-SEC class 1, employees in major groups 1–3	0.389	+0.0255 [+0.0013, +0.0497], $p=0.040$	+0.0259 [+0.0016, +0.0501], $p=0.037$	+0.0124 [-0.0371, +0.0619], $p=0.616$
with exposure $\times$ (2021 on) as a control			+0.0055 [-0.0001, +0.0111], $p=0.055$	-0.0002 [-0.0058, +0.0054], $p=0.938$
2021–2026 only			+0.0056 [+0.0013, +0.0098], $p=0.011$	-0.0016 [-0.0059, +0.0027], $p=0.448$
Share with under one year of tenure, pre-period 2015–2019 only		-0.0056 [-0.0091, -0.0021], $p=0.002$	-0.0037 [-0.0081, +0.0007], $p=0.097$	+0.0037 [-0.0030, +0.0104], $p=0.281$
Mean tenure (years), pre-period 2015–2019 only		+0.0304 [-0.0410, +0.1018], $p=0.399$	+0.0876 [-0.0206, +0.1958], $p=0.111$	-0.0174 [-0.2529, +0.2182], $p=0.884$

Note.—Rows 1–8: shares of employees in the SOC 2010 unit group by quarter (and the log ratio of non-supervisory to supervisory employees, and mean tenure from band midpoints of 0.125, 0.375, 0.75, 1.5, 3.5, 7.5, 15 and 25 years), regressed on the standardised exposure measure interacted with an indicator for 2023Q1 onward, with unit-group and quarter effects (main) or unit-group and major-group by quarter effects, cells weighted by employment. Rows 9–10: individual-level regressions in the April–June quarters among employees in paid work a year earlier without (row 9) or with (row 10) managerial or supervisory duties then, of an indicator for holding such duties now (or no longer), on exposure interacted with post, with unit-group, year and single-year-of-age effects (main) or unit-group, major-group by year and age effects. Rows 11–13: share of employees in major groups 1–3 classed as higher managerial and professional in NS-SEC, whose derivation changes with the occupational classification in 2021Q1; row 12 adds exposure interacted with an indicator for 2021 onward as a control and row 13 uses 2021–2026 only. Rows 14–15: the tenure outcomes with 2015–2019 as the pre-period (2020–2022 dropped; the mean column is left blank because it refers to 2015–2022). Brackets: 95 percent intervals clustered by SOC 3-digit group. Source: Quarterly Labour Force Survey.



Quarterly Labour Force Survey (ONS), employees aged 16-64 with a SOC 2010 unit group, person weights; annual weighted shares in the bottom and top quintiles of Anthropic observed exposure (quintiles of unit groups weighted by 2015-2022 employment). A: reported managerial status is manager, foreman or supervisor. B, C: time with the current employer under one year, and ten years or more. Dashed line: 30 Nov 2022. Source: Quarterly Labour Force Survey.

Figure A10: Seniority structure of employees in occupations exposed to language models, 2015–2025

Employees aged 16–64 with a SOC 2010 unit group, person weights; annual weighted shares in the bottom and top quintiles of Anthropic observed exposure (quintiles of unit groups weighted by 2015–2022 employment). Panel A: reported managerial status is manager, foreman or supervisor. Panels B and C: time with the current employer under one year, and ten years or more. Dashed line: 30 November 2022. Source: Quarterly Labour Force Survey.

The supervisory layer did not thicken, since among employees the share with managerial or supervisory status was 0.357 in 2015–2022 and changed by +0.004 per standard deviation of observed exposure within broad groups ($p = 0.23$), while the share who supervise others and the ratio of non-supervisory to supervisory staff did not move. Nor did the supervisory share in any age band. The tenure structure moved as the hiring results imply, by less than the headline. A fall of 0.6 points per standard deviation ($p = 0.005$) in the share of employees with less than a year of tenure, on a base of 0.171, came with a rise of 0.09 years in mean tenure ($p = 0.03$). The comparison year accounts for part of this, because in 2022 hiring peaked in exposed and unexposed occupations alike (figure A10). The coefficients on exposure by year, with 2022 as the reference, are between -0.0065 and -0.0107 in 2015–2020 and between -0.0110 and -0.0161 in 2023–2026. With 2015–2019 as the pre-period the fall in the short-tenure share is 0.37 points ($p = 0.10$) within broad groups and 0.56 points ($p = 0.002$) with quarter effects alone, and the rise in mean tenure 0.09 years ($p = 0.11$). Promotion did not quicken. Among employees who had been in paid work a year earlier without managerial or supervisory duties, the share holding such duties a year later was 0.068 before 2023 and did not change with exposure (-0.001 , $p = 0.56$), in any age band. The share losing such duties did not change either. Among employees in major groups 1–3 the share classed as higher managerial and professional in NS-SEC rose by 2.6 points per standard deviation ($p = 0.037$). The NS-SEC derivation changes with the occupational classification in 2021Q1, a change that shifts class shares by exposure. With exposure interacted with

an indicator for 2021 onward as a control the estimate is 0.55 points ($p = 0.055$). Over 2021–2026 alone it is 0.56 points ($p = 0.011$), which is why I do not count it as evidence of up-tiering.

Exits of mid-career workers

If the junior layer thins because the work is done by machines, the layers above might be expected to be safe. Two sources allow a check, both showing small rises in the exits of mid-career workers from exposed occupations, which survive some checks and fail others. The first source is the cross-section, in which people who are not in employment and report having been made redundant in the previous three months also report the occupation of the job they lost. Table A21 relates their number, per thousand employees of that minor occupation group and age band, to the group’s exposure, on a panel of SOC 2010 minor groups that runs through 2015–2026. The redundancy rate rose by 0.47 per thousand per standard deviation of observed exposure ($p < 0.001$; 0.66 with sub-major by quarter effects) on a base of 2.5 per thousand per quarter. By age band the rise is 0.48 for ages 35–49 ($p = 0.007$; 0.83 with sub-major by quarter effects), 0.57 for 26–34 ($p = 0.003$) and 0.27 for 50–64 ($p = 0.27$), against -0.11 ($p = 0.77$) for ages 22–25. The ability-based measure gives 0.44 for all ages with quarter effects ($p = 0.003$) and 0.50 with a wide interval within sub-major groups ($p = 0.31$). Panel B of the table shows what the estimate depends on. With 2015–2019 as the pre-period it is 0.36 for all ages ($p = 0.004$) and 0.38 for ages 35–49 ($p = 0.047$). A placebo break at the classification change of 2021Q1, estimated on 2015–2022, gives -0.37 ($p = 0.018$) and -0.51 ($p = 0.034$). Relative redundancy rates in exposed groups were thus falling into 2022, which shows that the mapping does not remove every trace of the break. Keeping the two classifications as separate units gives 0.73 and 0.86 ($p < 0.001$), in a design that identifies the post-2022 change from 2021–2026 alone, when 2022 was the trough of redundancies. Leaving out one sub-major group at a time moves the all-ages estimate between 0.41 and 0.55. The cells are sparse, with 759 redundant interviews in 2023–2026 across all ages and 269 for ages 35–49, and 59 and 81 percent of the cells empty.

Table A21: Redundancies by exposure of the last job and age, 2015Q1–2026Q2

	Rate 2015–2022	Anthropic observed exposure		DfE
Age band of the person made redundant	per 1,000	Minor group and quarter effects	Sub-major × quarter	Sub-major × quarter
<i>A. One classification (SOC 2010) throughout, 2015Q1–2026Q2</i>				
All ages 16–64	2.52	+0.471 [+0.217, +0.725], $p < 0.001$	+0.656 [+0.276, +1.037], $p = 0.001$	+0.495 [-0.462, +1.453], $p = 0.307$
22–25	2.27	-0.106 [-0.837, +0.625], $p = 0.773$	-0.226 [-1.083, +0.632], $p = 0.602$	-1.945 [-3.805, -0.086], $p = 0.041$
26–34	1.83	+0.567 [+0.193, +0.941], $p = 0.003$	+0.690 [-0.126, +1.506], $p = 0.096$	+0.888 [-0.301, +2.078], $p = 0.141$
35–49	2.26	+0.476 [+0.136, +0.816], $p = 0.007$	+0.834 [+0.255, +1.412], $p = 0.005$	+0.803 [-0.572, +2.177], $p = 0.249$
50–64	3.41	+0.273 [-0.217, +0.763], $p = 0.271$	+0.523 [-0.142, +1.187], $p = 0.121$	+0.743 [-0.414, +1.899], $p = 0.205$
<i>B. Robustness, minor-group and quarter effects for both measures</i>				
		Anthropic observed exposure		DfE
All ages, 2015–2019 pre-period		+0.357 [+0.116, +0.597], $p = 0.004$		+0.311 [+0.015, +0.607], $p = 0.040$
All ages, placebo break 2021Q1		-0.368 [-0.673, -0.064], $p = 0.018$		-0.078 [-0.396, +0.241], $p = 0.629$
All ages, minor groups by classification		+0.728 [+0.434, +1.021], $p < 0.001$		+0.494 [+0.145, +0.844], $p = 0.006$
All ages, leave one sub-major group out (range of 25)		+0.415 to +0.545		
35–49, 2015–2019 pre-period		+0.377 [+0.005, +0.750], $p = 0.047$		+0.222 [-0.303, +0.748], $p = 0.403$
35–49, placebo break 2021Q1		-0.515 [-0.989, -0.040], $p = 0.034$		-0.175 [-0.787, +0.438], $p = 0.573$
35–49, minor groups by classification		+0.861 [+0.386, +1.335], $p < 0.001$		+0.498 [-0.159, +1.155], $p = 0.137$
35–49, leave one sub-major group out (range of 25)		+0.347 to +0.557		

Note.—The outcome is the weighted number of people not in employment who report having been made redundant in the previous three months, per thousand employees of the minor occupation group of their last job and of the same age band, by quarter. Minor groups are SOC 2010 throughout: the last-job codes of 2021–2026 are SOC 2020 and are mapped to SOC 2010 minor groups by the modal mapping of the ONS relationship table, and the employee denominators use the SOC 2010 code that the LFS continues to derive for current jobs from 2021. The exposure of a minor group is the 2015–2022 employment-weighted mean of its unit groups’ standardised scores. Panel A: coefficients on exposure interacted with an indicator for 2023Q1 onward, with minor-group and quarter effects or minor-group and sub-major-group by quarter effects, cells weighted by employment, 95 percent intervals clustered by minor group. Panel B, minor-group and quarter effects: the pre-period restricted to 2015–2019; a placebo break at 2021Q1 estimated on 2015–2022; the two classifications kept as separate units (SOC 2010 minor groups to 2020, SOC 2020 minor groups from 2021, with adverts-weighted SOC 2020 exposure), which identifies the post-2022 change from 2021–2026 alone; and the range of estimates when one sub-major group at a time is left out. Source: Quarterly Labour Force Survey.

The second source is the two-quarter longitudinal files, which follow 594,261 quarter-1 employees with an occupation into the next quarter. Table A22 reports, by age band, the change after 2022 in the probability of each outcome per standard deviation of the exposure of the quarter-1 occupation, within broad groups. For employees aged 35–49 the probability of being out of employment a quarter later rose by 0.16 points ($p = 0.034$) on a base of 1.39 percent and the probability of having left through redundancy by 0.08 points ($p = 0.021$) on 0.55 percent, while the probability of having moved to a new employer fell by 0.22 points

($p = 0.014$) on 1.82 percent. The probability of being ILO unemployed (+0.04 points, $p = 0.39$) and of gaining managerial or supervisory status (−0.22 points, $p = 0.37$) did not change. Employees aged 22–25 show none of these changes. Employees aged 26–34 show a rise in the loss of supervisory status of 0.8 points ($p = 0.022$), and employees aged 50–64 a fall of 0.11 points in moves to a new employer with quarter effects alone ($p = 0.040$). Hours did not change in any band, whereas the ability-based measure shows none of it. Table A23 tests the 35–49 estimates, which against a 2015–2019 pre-period keep their size (0.14 points, $p = 0.076$; 0.09, $p = 0.038$; −0.24, $p = 0.016$). Without the IT and customer-service unit groups none is distinguishable from zero (0.08, $p = 0.25$; 0.06, $p = 0.25$; −0.27, $p = 0.053$). A placebo break at 2018Q1 estimated on 2015–2019 gives a redundancy estimate of −0.17 points ($p = 0.003$), twice the post-2022 estimate with the opposite sign. The pre-period thus contains swings of this size on its own. The estimates depend on 235 redundancy exits and 559 exits from employment among 35–49-year-olds in 2023–2026, 122 and 206 of them in the fifth of unit groups with the highest scores, unit groups counted equally. Keeping occupations by classification, with 2021–2022 alone as the pre-period, gives 0.25, 0.15 and −0.29 points, and a rise in unemployment of 0.17 points ($p = 0.006$) that the consistent panel does not show.

Table A22: Quarter-to-quarter transitions of employees by age band and exposure of the occupation, 2015Q1–2026Q1

Outcome in the following quarter	All	22–25	26–34	35–49	50–64
Not employed	+0.0006 [-0.0013, +0.0024] $p=0.545$; pre 0.0267	-0.0021 [-0.0084, +0.0043] $p=0.514$; pre 0.0349	-0.0000 [-0.0020, +0.0019] $p=0.963$; pre 0.0188	+0.0016 [+0.0001, +0.0031] $p=0.034$; pre 0.0139	+0.0010 [-0.0019, +0.0039] $p=0.494$; pre 0.0281
ILO unemployed	-0.0000 [-0.0008, +0.0008] $p=0.946$; pre 0.0095	+0.0002 [-0.0043, +0.0046] $p=0.938$; pre 0.0173	-0.0001 [-0.0014, +0.0012] $p=0.886$; pre 0.0078	+0.0004 [-0.0006, +0.0015] $p=0.393$; pre 0.0065	-0.0002 [-0.0015, +0.0011] $p=0.759$; pre 0.0071
Left through redundancy	+0.0008 [+0.0000, +0.0016] $p=0.048$; pre 0.0058	+0.0005 [-0.0013, +0.0022] $p=0.583$; pre 0.0058	+0.0014 [-0.0007, +0.0034] $p=0.197$; pre 0.0046	+0.0008 [+0.0001, +0.0015] $p=0.021$; pre 0.0055	+0.0008 [-0.0005, +0.0020] $p=0.230$; pre 0.0070
Employee with a new employer	-0.0009 [-0.0022, +0.0003] $p=0.138$; pre 0.0244	+0.0003 [-0.0045, +0.0051] $p=0.908$; pre 0.0439	+0.0006 [-0.0023, +0.0035] $p=0.684$; pre 0.0309	-0.0022 [-0.0040, -0.0005] $p=0.014$; pre 0.0182	-0.0009 [-0.0021, +0.0003] $p=0.123$; pre 0.0119
Employee in a different unit group	+0.0027 [-0.0029, +0.0084] $p=0.338$; pre 0.1513	+0.0005 [-0.0154, +0.0164] $p=0.946$; pre 0.1701	+0.0008 [-0.0075, +0.0091] $p=0.848$; pre 0.1587	+0.0032 [-0.0038, +0.0102] $p=0.367$; pre 0.1469	+0.0026 [-0.0044, +0.0095] $p=0.464$; pre 0.1397
Gained managerial or supervisory status	-0.0001 [-0.0036, +0.0034] $p=0.956$; pre 0.0608	+0.0005 [-0.0087, +0.0097] $p=0.916$; pre 0.0516	-0.0005 [-0.0072, +0.0061] $p=0.870$; pre 0.0714	-0.0022 [-0.0070, +0.0026] $p=0.369$; pre 0.0699	-0.0009 [-0.0064, +0.0047] $p=0.755$; pre 0.0562
Lost managerial or supervisory status	-0.0007 [-0.0047, +0.0032] $p=0.711$; pre 0.0808	-0.0225 [-0.0591, +0.0142] $p=0.226$; pre 0.1499	+0.0084 [+0.0012, +0.0155] $p=0.022$; pre 0.0960	-0.0027 [-0.0092, +0.0037] $p=0.403$; pre 0.0664	+0.0006 [-0.0051, +0.0063] $p=0.830$; pre 0.0740
Change in usual weekly hours	-0.02 [-0.05, +0.02] $p=0.396$; pre 0.0543	-0.07 [-0.22, +0.08] $p=0.361$; pre 0.3079	-0.04 [-0.10, +0.01] $p=0.085$; pre 0.0217	-0.01 [-0.07, +0.04] $p=0.689$; pre 0.0147	-0.02 [-0.10, +0.06] $p=0.663$; pre -0.1173

Note.—Quarter-1 employees aged 16–64 with an occupation in the two-quarter longitudinal LFS datasets, longitudinal weights. Occupations are SOC 2010 unit groups throughout; the SOC 2020 codes of 2021–2026 are mapped to SOC 2010 by the modal mapping of the ONS relationship table. Each cell is the coefficient on the Anthropic observed exposure of the quarter-1 occupation (standardised) interacted with an indicator for quarter 1 in 2023Q1 or later, in a regression of the quarter-2 outcome with unit-group, major-group by quarter and single-year-of-age effects, with its 95 percent interval clustered by SOC 3-digit group and the 2015–2022 mean of the outcome. Redundancy is reported in quarter 2 as having been made redundant in the last three months or as the reason for leaving the last job. Managerial or supervisory status is the reported managerial status of the main job. Source: Labour Force Survey two-quarter longitudinal datasets.

Table A23: Quarter-to-quarter transitions of employees aged 35–49 and of all ages: robustness

Outcome in the following quarter	Pre-period 2015–2019	Without IT and customer service	Occupations by classification	Placebo break 2018Q1, 2015–2019
<i>A. Ages 35–49</i>				
Not employed	+0.0014 [-0.0002, +0.0030], $p=0.076$	+0.0008 [-0.0005, +0.0021], $p=0.251$	+0.0025 [+0.0002, +0.0048], $p=0.035$	-0.0013 [-0.0030, +0.0003], $p=0.105$
ILO unemployed	-0.0000 [-0.0013, +0.0012], $p=0.985$	+0.0001 [-0.0011, +0.0014], $p=0.814$	+0.0017 [+0.0005, +0.0029], $p=0.006$	-0.0010 [-0.0022, +0.0001], $p=0.077$
Left through redundancy	+0.0009 [+0.0000, +0.0017], $p=0.038$	+0.0006 [-0.0005, +0.0017], $p=0.251$	+0.0015 [-0.0001, +0.0030], $p=0.064$	-0.0017 [-0.0029, -0.0006], $p=0.003$
Employee with a new employer	-0.0024 [-0.0043, -0.0004], $p=0.016$	-0.0027 [-0.0053, +0.0000], $p=0.053$	-0.0029 [-0.0054, -0.0004], $p=0.022$	-0.0014 [-0.0032, +0.0004], $p=0.134$
Employee in a different unit group	+0.0056 [-0.0026, +0.0138], $p=0.180$	+0.0026 [-0.0086, +0.0138], $p=0.646$	+0.0005 [-0.0063, +0.0073], $p=0.879$	+0.0009 [-0.0065, +0.0084], $p=0.803$
Gained managerial or supervisory status	-0.0003 [-0.0060, +0.0055], $p=0.929$	-0.0030 [-0.0106, +0.0046], $p=0.434$	-0.0067 [-0.0123, -0.0010], $p=0.021$	+0.0084 [+0.0032, +0.0135], $p=0.002$
Lost managerial or supervisory status	-0.0012 [-0.0077, +0.0052], $p=0.705$	-0.0024 [-0.0098, +0.0050], $p=0.521$	-0.0041 [-0.0128, +0.0045], $p=0.344$	+0.0024 [-0.0018, +0.0066], $p=0.260$
<i>B. All ages 16–64</i>				
Not employed	+0.0004 [-0.0016, +0.0023], $p=0.715$	-0.0001 [-0.0015, +0.0014], $p=0.941$	+0.0013 [-0.0007, +0.0033], $p=0.209$	-0.0007 [-0.0021, +0.0008], $p=0.367$
ILO unemployed	-0.0001 [-0.0010, +0.0009], $p=0.908$	-0.0003 [-0.0014, +0.0008], $p=0.559$	+0.0002 [-0.0008, +0.0013], $p=0.697$	-0.0007 [-0.0016, +0.0001], $p=0.094$
Left through redundancy	+0.0007 [-0.0002, +0.0016], $p=0.130$	+0.0002 [-0.0007, +0.0011], $p=0.664$	+0.0013 [+0.0001, +0.0025], $p=0.035$	-0.0011 [-0.0017, -0.0004], $p=0.002$
Employee with a new employer	-0.0009 [-0.0021, +0.0003], $p=0.133$	-0.0009 [-0.0027, +0.0008], $p=0.301$	-0.0013 [-0.0035, +0.0009], $p=0.234$	-0.0004 [-0.0022, +0.0013], $p=0.614$
Employee in a different unit group	+0.0043 [-0.0021, +0.0107], $p=0.187$	+0.0035 [-0.0035, +0.0105], $p=0.324$	+0.0009 [-0.0043, +0.0062], $p=0.725$	+0.0016 [-0.0032, +0.0064], $p=0.515$
Gained managerial or supervisory status	-0.0009 [-0.0046, +0.0029], $p=0.650$	-0.0026 [-0.0071, +0.0019], $p=0.249$	+0.0011 [-0.0024, +0.0046], $p=0.535$	+0.0019 [-0.0012, +0.0051], $p=0.222$
Lost managerial or supervisory status	-0.0004 [-0.0048, +0.0041], $p=0.874$	-0.0012 [-0.0063, +0.0039], $p=0.634$	+0.0002 [-0.0049, +0.0053], $p=0.941$	+0.0043 [+0.0006, +0.0080], $p=0.023$

Note.—Coefficients on the Anthropic observed exposure of the quarter-1 occupation (standardised) interacted with an indicator for quarter 1 in 2023Q1 or later, in regressions of the quarter-2 outcome with unit-group, major-group by quarter and single-year-of-age effects, longitudinal weights, 95 percent intervals clustered by SOC 3-digit group. Column 1 drops quarter 1 in 2020–2022, so that 2015–2019 forms the pre-period. Column 2 drops the IT unit groups (2133–2139, 3131, 3132) and the customer-service groups (7211, 7213, 7214, 7215, 7219). Column 3 keeps occupations by classification (SOC 2010 unit groups to 2020, SOC 2020 unit groups from 2021, with SOC 2020 exposure scores obtained from the SOC 2010 scores through the ONS relationship table), so that the post-2022 change is identified from 2021–2026 alone. Column 4 estimates a placebo break at 2018Q1 on 2015–2019 with the consistent classification. Source: Labour Force Survey two-quarter longitudinal datasets.

A change in the composition of the unemployed points the same way, with a caveat. Among the ILO unemployed aged 35–49 with a known last job, the share whose last job was in a managerial, professional or associate professional occupation was 0.282 in 2015, 0.357 in 2022 and 0.416, 0.375 and 0.443 in 2023–2025; the change of 7.5 points against 2015–2022 ($p = 0.008$) falls to 1.9 points ($p = 0.50$) once a linear trend is allowed. The unemployed were becoming former professionals before 2023 and continued to (table A24).

Table A24: Last job of the ILO unemployed by age band, 2015–2026

Age band	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026
26–34, groups 1–3	0.203	0.202	0.223	0.217	0.258	0.318	0.295	0.316	0.332	0.351	0.362	0.370
group 4	0.101	0.105	0.096	0.098	0.104	0.121	0.093	0.088	0.125	0.092	0.109	0.074
group 9	0.296	0.281	0.236	0.268	0.223	0.229	0.224	0.275	0.219	0.238	0.209	0.214
unemployed with a known last job	1344	1200	1060	925	814	692	744	518	348	541	816	478
35–49, groups 1–3	0.282	0.312	0.325	0.350	0.341	0.390	0.337	0.357	0.416	0.375	0.443	0.468
group 4	0.109	0.102	0.103	0.101	0.086	0.105	0.117	0.116	0.104	0.118	0.109	0.079
group 9	0.231	0.240	0.174	0.189	0.205	0.156	0.180	0.203	0.172	0.178	0.175	0.172
unemployed with a known last job	1769	1607	1373	1277	1112	1070	1184	738	534	713	1099	607
50–64, groups 1–3	0.333	0.321	0.350	0.341	0.371	0.359	0.347	0.361	0.399	0.403	0.446	0.477
group 4	0.109	0.096	0.115	0.132	0.125	0.123	0.127	0.079	0.089	0.080	0.117	0.095
group 9	0.170	0.191	0.172	0.180	0.175	0.155	0.150	0.193	0.172	0.194	0.172	0.135
unemployed with a known last job	1502	1432	1304	1276	1163	1283	1591	927	668	652	1109	590

Note.—Weighted shares of ILO unemployed people in the age band with a known last job whose last job was in SOC major groups 1–3, 4 or 9, by calendar year (2026 is January–June), and the number of unemployed with a known last job. The last-job code is SOC 2010 to 2020 and SOC 2020 from 2021; the major groups correspond. Source: Quarterly Labour Force Survey.

Adverts for senior and junior roles

Online adverts cannot see seniority within an occupation, but within a field they separate managers and directors from professionals, associate professionals and administrative staff. Six fields are built from SOC 2020 titles (IT, finance and accounting, human resources, marketing and public relations, legal, engineering; the unit groups are listed in table A25), and log monthly adverts are regressed on level interacted with post, with unit-group and field by month effects and the professional level as the reference. Table A26 reports the results. Adverts for managers and directors fell after 2022 by 0.28 log points less than adverts for professionals in the same field ($p = 0.003$), and by 0.17 log points less than all other levels together ($p < 0.001$). Two checks take most of this away. Adverts for managerial roles had been gaining on the other levels before 2022, by 0.029 log points a year over 2017–2022 ($p = 0.004$). Relative to 2022 the manager contrast is -0.13 in 2017, -0.09 in 2019 and -0.01 in 2021, then $+0.03$, $+0.10$, $+0.20$ and $+0.25$ in 2023–2026. With the trend as a control the post-2022 effect is -0.015 ($p = 0.44$), while a placebo break at January 2020 estimated on 2017–2022 gives $+0.13$ ($p < 0.001$). The contrast also contains the exposure gradient of section 4.7. Across major groups 1–4 the managerial unit groups are less exposed on the observed-exposure index than the professional ones (mean standardised scores of 0.05 and 0.15). Once exposure interacted with post enters the regression, the manager contrast falls from 0.17 to 0.13 in the six fields and from 0.12 ($p = 0.15$) to 0.03 ($p = 0.52$) across all unit groups in major groups 1–4. Within the latter it is 0.17 among unit groups above the median exposure ($p = 0.062$) and -0.05 below it. Adverts fell at every level. With 2022 as 100, the

mean unit group in 2025 posted 54 for managers and directors, 48 for professionals, 53 for associate professionals and 41 for administrative staff. In IT, where adverts for professionals fell by 0.855 log points between the two periods against 0.339 for IT directors, the contrast is +0.52, although its interval, on five clusters, should not be read closely.

Table A25: Adverts by level within fields: unit groups and field-level estimates

Field	Contrast	Coef.	95% CI	<i>p</i>
IT	level 1 minus professionals	+0.523	[+0.499, +0.546]	<0.001
IT	level 3 minus professionals	+0.560	[+0.409, +0.710]	<0.001
finance and accounting	level 1 minus professionals	+0.030	[-0.122, +0.183]	0.643
finance and accounting	level 3 minus professionals	-0.027	[-0.179, +0.126]	0.684
finance and accounting	level 4 minus professionals	-0.037	[-0.190, +0.116]	0.575
marketing and public relations	level 1 minus professionals	+0.079	[-0.006, +0.165]	0.060
marketing and public relations	level 3 minus professionals	-0.106	[-0.191, -0.021]	0.028
legal	level 3 minus professionals	-0.011	[-0.021, -0.001]	0.039
legal	level 4 minus professionals	-0.242	[-0.255, -0.229]	<0.001
engineering	level 3 minus professionals	+0.123	[+0.082, +0.165]	0.006

Note.—Fields are built from SOC 2020 unit-group titles: IT (1137, 2131–2137, 2139, 2141, 3131–3133, 3573), finance and accounting (1131, 2421–2423, 2433, 2440, 3532–3534, 4122–4124, 4129, 4132), human resources (1136, 3571, 4136), marketing and public relations (1132, 1133, 2432, 2493, 2494, 3554), legal (2411, 2412, 2419, 3520, 4212), engineering (2121–2127, 2129, 2481, 3113, 3114, 3119). Within each field, log adverts are regressed on the level of the unit group interacted with post with unit-group and month effects and professionals as the reference; human resources has no professional level. Intervals are clustered by SOC 2020 minor group, with three to seven clusters per field, and should not be read closely. Source: ONS online job advert estimates.

Table A26: Online job adverts for managerial and other levels within fields, January 2017 to August 2026

Contrast (level $\times$ post, log adverts)	Coef.	95% CI	<i>p</i>
<i>A. Six fields, unit-group and field $\times$ month effects</i>			
Managers and directors minus professionals	+0.279	[+0.111, +0.447]	0.003
Associate professionals minus professionals	+0.230	[-0.016, +0.477]	0.065
Administrative minus professionals	+0.057	[-0.104, +0.219]	0.464
Managers and directors minus all other levels	+0.166	[+0.114, +0.218]	<0.001
managers $\times$ linear trend, 2017–2022 (per year)	+0.029	[+0.010, +0.047]	0.004
post effect with the managers $\times$ trend control	-0.015	[-0.056, +0.026]	0.441
placebo break at January 2020, 2017–2022	+0.130	[+0.071, +0.189]	<0.001
with observed exposure $\times$ post as a control	+0.132	[+0.056, +0.209]	0.002
observed exposure $\times$ post in that model	-0.231	[-0.464, +0.002]	0.052
<i>B. Managers and directors minus all other levels, by year relative to 2022</i>			
2017	-0.127	[-0.226, -0.027]	0.015
2018	-0.077	[-0.150, -0.004]	0.039
2019	-0.086	[-0.166, -0.007]	0.035
2020	+0.110	[+0.018, +0.201]	0.021
2021	-0.010	[-0.047, +0.026]	0.561
2023	+0.032	[+0.004, +0.061]	0.027
2024	+0.101	[+0.066, +0.135]	<0.001
2025	+0.203	[+0.152, +0.254]	<0.001
2026	+0.253	[+0.192, +0.314]	<0.001
<i>C. All unit groups in major groups 1–4, unit-group and month effects</i>			
Managers and directors $\times$ post, all	+0.117	[-0.043, +0.278]	0.149
unit groups above the median observed exposure	+0.173	[-0.009, +0.355]	0.062
unit groups at or below the median	-0.048	[-0.202, +0.106]	0.531
all, with observed exposure $\times$ post as a control	+0.031	[-0.065, +0.126]	0.520

Note.—Log monthly online job adverts by SOC 2020 unit group (ONS Textkernel data, UK totals, suppressed cells dropped), weighted by 2017–2022 mean adverts. Panel A: coefficients on the level of the unit group (major group 1 managers and directors, 3 associate professionals, 4 administrative; professionals, major group 2, as the reference) interacted with an indicator for January 2023 onward, with unit-group and field by month effects, over 54 unit groups in six fields; the fourth row contrasts managers and directors with all other levels, and the indented rows give the managers by linear-trend coefficient over 2017–2022, the post effect with that trend as a control, a placebo break at January 2020 estimated on 2017–2022, and the post effect with Anthropoc observed exposure interacted with post as a control. Panel B: the managers minus others contrast interacted with calendar year, 2022 omitted. Panel C: managers and directors interacted with post against all unit groups in major groups 2–4, with unit-group and month effects, for all scored unit groups, by the median of Anthropoc observed exposure, and with exposure interacted with post as a control. Brackets: 95 percent intervals clustered by SOC 2020 minor group. Source: ONS online job advert estimates.

What the layers above entry show

Half of the pyramid claim holds. The junior layer of exposed occupations thinned, with fewer hires of 22–25-year-olds, a smaller short-tenure share and a fall in the employment of 16–21-year-olds beyond its trend. The layers above were not enlarged by it. Employment of workers aged 26–64 in exposed occupations grew relative to other occupations throughout 2015–2026 and did not accelerate after 2022, while their hiring did not fall in a way that survives a change of design, supervisory shares and promotion rates did not change, pay did not gain on juniors’ and hours fell relative to juniors’. The evidence that mid-career workers were displaced is weaker than the evidence that the young were not hired. In the two-quarter files, 35–49-year-olds in exposed occupations within broad groups were more often made redundant and more often out of employment a quarter later, and less often moving to a new employer, by an eighth to a seventh of their low baselines. In the cross-section, the redundancy rate of exposed last jobs rose by a fifth with minor-group and quarter effects. The cross-sectional estimate holds against a 2015–2019 pre-period and when sub-major groups are left out one at a time, while the two-quarter estimates keep their size against a 2015–2019 pre-period, one of them at $p = 0.076$. Against this, the two-quarter estimates are not distinguishable from zero without the IT and customer-service groups, they depend on a few hundred events, a placebo break in 2018 produces a redundancy estimate of the opposite sign and twice the size, and the ability-based measure shows none of them. The cross-sectional estimate has a placebo break at the 2021 classification change. Adverts for managerial titles held up relative to professional ones within fields on a pre-trend that the post-2022 years continue, a gap that is as much an exposure gradient as a seniority gradient. The stock of exposed occupations aged because entry slowed. The data suggest, without establishing it, that the experienced workers who remained were displaced as well.

G The 2008–2011 Benchmark: Windows, Definitions and Checks

Table A27: Recession estimates under alternative windows

Pre years	Post years	All	Graduates	Non-graduates	Administrative
2006–2007	2009–2010	+0.022 [-0.005, +0.049]	+0.002 [-0.027, +0.030]	+0.044 [+0.003, +0.084]	-0.040 [-0.065, -0.016]
2006–2007	2009–2011	+0.024 [-0.001, +0.049]	+0.011 [-0.015, +0.038]	+0.040 [+0.003, +0.077]	-0.044 [-0.067, -0.022]
2006–2007	2009–2012	+0.019 [-0.004, +0.042]	+0.021 [-0.004, +0.047]	+0.028 [-0.007, +0.062]	-0.040 [-0.061, -0.019]
2006–2007	2010–2012	+0.017 [-0.008, +0.041]	+0.026 [-0.001, +0.053]	+0.022 [-0.015, +0.059]	-0.040 [-0.063, -0.018]
2006–2007	2009–2013	+0.018 [-0.004, +0.040]	+0.018 [-0.005, +0.042]	+0.030 [-0.003, +0.063]	-0.040 [-0.060, -0.019]
2006–2007	2011–2013	+0.017 [-0.008, +0.041]	+0.029 [+0.001, +0.056]	+0.022 [-0.015, +0.059]	-0.040 [-0.062, -0.018]
2006–2007	2012–2014	+0.007 [-0.017, +0.031]	+0.024 [-0.003, +0.050]	+0.016 [-0.021, +0.053]	-0.039 [-0.061, -0.017]
2006–2008Q1	2009–2010	+0.022 [-0.005, +0.049]	+0.004 [-0.024, +0.032]	+0.042 [+0.002, +0.082]	-0.034 [-0.058, -0.010]
2006–2008Q1	2009–2011	+0.024 [+0.000, +0.048]	+0.014 [-0.012, +0.040]	+0.038 [+0.002, +0.074]	-0.039 [-0.060, -0.017]
2006–2008Q1	2009–2012	+0.020 [-0.003, +0.042]	+0.023 [-0.001, +0.048]	+0.027 [-0.006, +0.061]	-0.034 [-0.055, -0.014]
2006–2008Q1	2010–2012	+0.017 [-0.006, +0.041]	+0.028 [+0.001, +0.054]	+0.022 [-0.014, +0.058]	-0.035 [-0.057, -0.014]
2006–2008Q1	2009–2013	+0.018 [-0.003, +0.040]	+0.020 [-0.003, +0.043]	+0.029 [-0.003, +0.061]	-0.034 [-0.053, -0.014]
2006–2008Q1	2011–2013	+0.017 [-0.007, +0.041]	+0.030 [+0.004, +0.056]	+0.022 [-0.014, +0.059]	-0.034 [-0.056, -0.013]
2006–2008Q1	2012–2014	+0.007 [-0.017, +0.030]	+0.025 [-0.001, +0.051]	+0.016 [-0.021, +0.052]	-0.031 [-0.053, -0.010]

Note.—Person-level estimates of the recession episode indicator for the elementary share of new hires aged 22–25 (all, with a degree, without a degree) and for their administrative share, under alternative pre and post windows, with quarter-of-year effects, person weights and 95 percent intervals clustered by household within quarter. “2006–2008Q1” adds the first quarter of 2008 to the pre period. Occupation is coded in SOC 2000 major groups up to 2010 and SOC 2010 from 2011. The windows 2011–2013 and 2012–2014 start in and after 2011, the year in which unemployment at all ages peaked; for 22–25-year-olds, 2011–2013 has the highest unemployment rate of the three-year windows. Source: Quarterly Labour Force Survey, 2006–2014.

Table A28: The benchmark from the 2009 and the 2011 troughs: entry jobs on two definitions, employment and unemployment

	Downturn, against 2006–2007				Recent episode		
	Mean 2006–2007	2009–2011	2011–2013	2012–2014	Mean 2015–2022	2023–2025	Difference from 2011–2013 (<i>z</i>)
<i>A. New hires aged 22–25: share of entry jobs by type</i>							
Elementary (9)	0.192	+0.024	+0.017	+0.007	0.186	+0.079	+0.062 (+2.79)
		[-0.001, +0.049]	[-0.008, +0.041]	[-0.017, +0.031]		[+0.043, +0.115]	
Elementary, graduates	0.084	+0.011	+0.029	+0.024	0.113	+0.078	+0.049 (+1.97)
		[-0.015, +0.038]	[+0.001, +0.056]	[-0.003, +0.050]		[+0.037, +0.119]	
unweighted	0.082	+0.016	+0.033	+0.027	0.109	+0.071	+0.038 (+1.74)
		[-0.010, +0.042]	[+0.007, +0.060]	[+0.001, +0.053]		[+0.038, +0.105]	
Elementary, non-graduates	0.272	+0.040	+0.022	+0.016	0.259	+0.094	+0.072 (+1.99)
		[+0.003, +0.077]	[-0.015, +0.059]	[-0.021, +0.053]		[+0.034, +0.154]	
Low-paying (6, 7, 9), graduates	0.228	+0.067	+0.089	+0.083	0.305	+0.069	-0.021 (-0.63)
		[+0.027, +0.108]	[+0.049, +0.130]	[+0.043, +0.124]		[+0.019, +0.119]	
unweighted	0.230	+0.077	+0.094	+0.084	0.309	+0.057	-0.037 (-1.25)
		[+0.037, +0.118]	[+0.054, +0.134]	[+0.045, +0.123]		[+0.014, +0.099]	
sales and customer service (7)	0.091	+0.028	+0.044	+0.038	0.113	-0.003	-0.046 (-2.09)
		[-0.000, +0.056]	[+0.014, +0.073]	[+0.010, +0.067]		[-0.034, +0.029]	
caring, leisure and other service (6)	0.053	+0.028	+0.017	+0.022	0.079	-0.007	-0.024 (-1.38)
		[+0.005, +0.050]	[-0.004, +0.039]	[+0.000, +0.043]		[-0.033, +0.020]	
Administrative (4)	0.179	-0.044	-0.040	-0.039	0.128	-0.049	-0.009 (-0.58)
		[-0.067, -0.022]	[-0.062, -0.018]	[-0.061, -0.017]		[-0.071, -0.027]	
<i>B. The same on SOC 2000 major groups throughout 2006–2014</i>							
Elementary, graduates	0.084	+0.011	+0.030	+0.024	0.113	+0.078	+0.048 (+1.93)
		[-0.015, +0.038]	[+0.002, +0.057]	[-0.002, +0.051]		[+0.037, +0.119]	
Low-paying, graduates	0.228	+0.065	+0.088	+0.080	0.305	+0.069	-0.019 (-0.58)
		[+0.025, +0.106]	[+0.047, +0.129]	[+0.040, +0.120]		[+0.019, +0.119]	
Administrative	0.179	-0.047	-0.048	-0.043	0.128	-0.049	-0.001 (-0.09)
		[-0.070, -0.025]	[-0.070, -0.026]	[-0.065, -0.021]		[-0.071, -0.027]	
<i>C. People aged 22–25</i>							
Employment rate	0.744	-0.046	-0.049	-0.039	0.756	-0.021	+0.028 (+5.20)
		[-0.052, -0.039]	[-0.056, -0.042]	[-0.046, -0.032]		[-0.029, -0.012]	
unweighted	0.747	-0.045	-0.050	-0.039	0.760	-0.002	+0.048 (+10.20)
		[-0.052, -0.039]	[-0.056, -0.043]	[-0.046, -0.033]		[-0.008, +0.005]	
Employment rate, graduates	0.821	-0.049	-0.054	-0.043	0.817	-0.002	+0.051 (+6.09)
		[-0.061, -0.038]	[-0.065, -0.042]	[-0.054, -0.031]		[-0.014, +0.009]	
unweighted	0.826	-0.048	-0.052	-0.042	0.820	+0.012	+0.064 (+8.77)
		[-0.059, -0.037]	[-0.063, -0.041]	[-0.053, -0.031]		[+0.003, +0.021]	
ILO unemployment rate	0.083	+0.042	+0.049	+0.041	0.077	+0.008	-0.041 (-10.82)
		[+0.037, +0.046]	[+0.044, +0.054]	[+0.036, +0.046]		[+0.003, +0.014]	
ILO unemployment rate, graduates	0.068	+0.028	+0.030	+0.023	0.067	+0.009	-0.021 (-3.65)
		[+0.020, +0.036]	[+0.022, +0.038]	[+0.015, +0.031]		[+0.001, +0.017]	

Note.—Coefficients on a period indicator in regressions of the individual outcome on it and quarter-of-year effects, on the years of the two periods compared, with person weights unless stated; 95 percent intervals clustered by household within quarter in brackets. The recession windows are against 2006–2007 and the recent episode is 2023–2025 against 2015–2022; the means are those of the two base periods, with the row's weights. Low-paying occupations are major groups 6 (caring, leisure and other service; personal service in SOC 2000), 7 (sales and customer service) and 9 (elementary), as in Clarke (2019). Panel A codes occupation in SOC 2000 up to 2010 and SOC 2010 from 2011; panel B uses, for 2011–2014, the SOC 2000 code that the ONS maps from SOC 2010; 2015–2025 are on SOC 2010 in both. The difference is the 2023–2025 estimate minus the 2011–2013 one, with a *z* statistic from the two standard errors. Clustering by the household's quarter of entry, which groups its five interviews, widens the panel C intervals for 2011–2013 and 2023–2025 by factors of 1.2 to 1.6, with the differences still significant, and does not widen those of the graduate new-hire shares. Source: Quarterly Labour Force Survey.

Table A29: The two episodes by type of response, with trimmed weights, and the 2011 classification break

	2009–2011 vs 2006–2007	2023–2025 vs 2015–2022	Difference (z)
<i>A. Elementary share of new hires aged 22–25 by type of response</i>			
Personal responses	+0.023 [-0.011, +0.057]	+0.060 [+0.003, +0.117]	+0.037 (+1.09)
Proxy responses	+0.025 [-0.011, +0.061]	+0.091 [+0.045, +0.137]	+0.065 (+2.19)
Graduates, personal responses	+0.015 [-0.022, +0.052]	+0.051 [-0.006, +0.108]	+0.036 (+1.04)
Graduates, proxy responses	+0.012 [-0.026, +0.050]	+0.094 [+0.039, +0.150]	+0.082 (+2.41)
All, with response-type indicators	+0.024 [-0.000, +0.049]	+0.078 [+0.042, +0.115]	+0.054 (+2.40)
<i>B. Weights</i>			
2023–2025 estimate, weights trimmed at the 99th percentile	+0.074 [+0.040, +0.108], $p < 0.001$		
2023–2025 estimate, records above the 99th weight percentile dropped	+0.069 [+0.035, +0.103], $p < 0.001$		
2009–2011 estimate, weights trimmed at the 99th percentile	+0.024 [-0.001, +0.048], $p = 0.058$		
<i>C. The 2011Q1 classification break: employees with both codes in the January–March 2011 file</i>			
Major group 9	staying: 0.965 (of 4595)	coming from the same SOC 2000 group: 0.995 (of 4452)	
Major group 4	staying: 0.950 (of 4835)	coming from the same SOC 2000 group: 0.924 (of 4968)	
Major group 1	staying: 0.601 (of 6095)	coming from the same SOC 2000 group: 0.999 (of 3695)	
Major group 2	staying: 0.987 (of 5619)	coming from the same SOC 2000 group: 0.700 (of 7997)	
Major group 3	staying: 0.725 (of 5783)	coming from the same SOC 2000 group: 0.791 (of 5219)	
Young new hires, 2011Q1 (n=157)	elementary: 0.217 vs 0.217	administrative: 0.121 vs 0.108 (SOC 2010 vs mapped SOC 2000)	

Note.—Panel A: person-level episode estimates for the elementary share of new hires aged 22–25 among personal responses, among proxy responses, and with indicators for the type of response. Panel B: the 2023–2025 estimate with person weights trimmed at their 99th percentile over the years compared, or with the records above that percentile dropped, and the recession estimate with trimmed weights. Panel C: among employees aged 16–64 in the January–March 2011 file, which carries both the SOC 2010 code and the SOC 2000 code mapped by the ONS, the share of each SOC 2000 major group whose SOC 2010 code is in the same major group, and the share of each SOC 2010 major group that comes from the same SOC 2000 group; the last row reports the elementary and administrative shares of new hires aged 22–25 in that quarter under the two codings. Source: Quarterly Labour Force Survey.


H Supply Responses: Subject Choice and Graduate Outcomes

If language models reduce the return to a field of study, applicants and graduates should respond. This appendix examines two margins, the subject choices of school leavers and the outcomes of recent graduates by subject, and finds a response in one subject only.

UCAS records applications from UK-domiciled 18-year-olds by subject group at the 30 June deadline of each admissions cycle. Figure A11 plots the share of applications to six subject groups over the 2019–2026 cycles, while table A30 reports the shares. Computing is the one subject whose trajectory turned. Its share of applications rose from 0.048 in the 2021 cycle to 0.056 in 2022, 0.062 in 2023 and 0.066 in 2024, the latter two being the first cycles decided after the release of ChatGPT, and fell to 0.057 in 2025 and 0.050 in 2026, below its 2022 level. Law and engineering rose throughout the period, from 0.058 to 0.071 and from 0.062 to 0.084 between 2022 and 2026. Languages continued a long decline. A regression of log applications on the exposure of each subject’s graduate occupations interacted with the post-2023 cycles, across 20 subject groups with group and cycle effects, gives a coefficient of

0.003 ($p = 0.89$; randomisation $p = 0.89$). The same regression for acceptances gives -0.001 . School leavers did not move away from exposed subjects in general, although they moved away from computing after two cycles in which they had moved towards it.

Share of applications from UK 18-year-olds by subject group, cycles 2019-2026



Applications at the 30 June deadline to each subject group (CAH1) as a share of applications to the 20 groups (combined and general studies left out); UK-domiciled applicants aged 18. Exposure in brackets: LLM exposure of graduates' jobs, SD across subject groups (as in fig_ucas_subject_choice). Dashed line between the 2022 and 2023 cycles. Source: UCAS 30 June deadline 2026 release.

Figure A11: Share of applications from UK 18-year-olds by subject group, admissions cycles 2019–2026

Applications at the 30 June deadline to each subject group (CAH level 1) as a share of applications to the 20 groups, UK-domiciled applicants aged 18. The exposure in brackets is the mean language-model exposure of the occupations held by employed graduates of the subject in 2015–2022, standardised across subject groups. Dashed line: between the 2022 and 2023 cycles. Source: UCAS, 30 June deadline data, 2026 release.

The Longitudinal Education Outcomes data give population outcomes by subject for every cohort of first-degree graduates of English universities. I regress, across 33 subjects and the tax years 2017–18 to 2023–24, the share of graduates one year after graduation with no sustained employment or study, and their log median earnings, on subject exposure interacted with the 2023–24 tax year, with subject and year effects. The exposure coefficient is 0.09 percentage points for the share with no sustained destination ($p = 0.31$; randomisation $p = 0.40$) and 0.002 for log earnings ($p = 0.71$). A triple difference against graduates five years out gives -0.10 points ($p = 0.26$). The movements the data show are general. The share of graduates with no sustained destination one year out rose from 6.9 percent for the 2022–23 tax year to 7.3 percent for 2023–24. It rose in computing (8.7 to 10.1 percent), software engineering (6.2 to 8.3) and computer science (8.1 to 9.3), the subjects whose graduates enter programming, and by about a point in business, English, history, languages, economics and

law. In the LFS, the exposure of a graduate’s subject predicts a lower probability of a job above the exposure median after 2022 (coefficient -0.011 per standard deviation, $p = 0.07$) and a higher probability of wanting longer hours (0.012 , randomisation $p = 0.02$) and of part-time work (0.009 , randomisation $p = 0.03$), but two of nine job-quality outcomes pass at the 5 percent level. The adjustment to language models, in so far as the subject data see it, was in the type of job graduates started, while their employment and pay by subject did not move.

Table A30: Applications by subject group, exposure regressions for applications and acceptances, and graduate outcomes by subject exposure

<i>A. Share of applications from UK 18-year-olds at the 30 June deadline, by subject group</i>						
	2019	2021	2022	2024	2025	2026
Computing	0.0471	0.0484	0.0556	0.0658	0.0573	0.0501
Engineering and technology	0.0627	0.0633	0.0615	0.0685	0.0755	0.0837
Law	0.0561	0.0576	0.0583	0.0641	0.0677	0.0714
Languages	0.0480	0.0392	0.0374	0.0351	0.0325	0.0292
<i>B. Subject exposure $\times$ post, log applications and acceptances (20 subject groups)</i>						
	Coef.	95% CI	p	RI p		
Applications, 2024–2026	+0.0034	[-0.045, +0.052]	0.887	0.892		
Applications, pre = 2019 and 2022 only	+0.0012	[-0.050, +0.052]	0.963	0.965		
Acceptances, 2024–2025	-0.0008	[-0.044, +0.042]	0.968	0.968		
<i>C. LEO: subject exposure $\times$ 2023–24 tax year, first-degree graduates one year after graduation (33 subjects)</i>						
No sustained destination (pp)	+0.0901	[-0.089, +0.269]	0.312	0.402		
Sustained employment or further study (pp)	-0.1735	[-0.462, +0.115]	0.230	0.349		
Log median earnings (real)	+0.0017	[-0.008, +0.011]	0.708	0.743		
No sustained destination, one year out minus five years out (triple difference)	-0.1006	[-0.279, +0.078]	0.259			

Note.—Panel A: applications at the 30 June deadline from UK-domiciled 18-year-olds to the subject group as a share of applications to the 20 CAH level-1 groups. Panel B: coefficients on the standardised exposure of the subject group’s graduate occupations interacted with the post cycles, in regressions of log applications (or log acceptances) on group and cycle effects across 20 groups, with 95 percent intervals (t with 19 degrees of freedom) and randomisation-inference p -values from permuting exposure across groups. Panel C: coefficients on subject exposure interacted with the 2023–24 tax year in regressions of the outcome of first-degree UK-domiciled graduates one year after graduation on subject and tax-year effects across 33 subjects, tax years 2017–18 to 2023–24, weighted by the number of matched graduates, with intervals on 32 degrees of freedom; the triple difference subtracts the same coefficient for graduates five years after graduation. Source: UCAS; Department for Education, Longitudinal Education Outcomes.