Artificial Intelligence and the Labor Market*

Menaka Hampole
Yale School of Management

Dimitris Papanikolaou
Kellogg School of Management and NBER

Lawrence D. W. Schmidt
MIT Sloan School of Management

Bryan Seegmiller
Kellogg School of Management

July 21, 2025

Abstract

We utilize recent advances in natural language processing to develop novel measures of workers' task-level exposure to artificial intelligence (AI) and machine learning technologies from 2010 to 2023, capturing variation across firms and over time. We show that tasks exposed to AI subsequently experience lower labor demand. Employing a model that distinguishes between direct and indirect productivity effects of labor-saving technologies, we identify two variables that summarize the impact of AI on within-firm labor demand: an occupations mean task exposure to AI, and the degree to which this mean exposure is concentrated in a small number of tasks. Higher mean exposure reduces labor demand, whereas more concentrated exposures plays an offsetting role as it allows workers to reallocate their effort to non-displaced tasks. Leveraging exogenous variation in AI adoption linked to firms' pre-existing hiring practices, we find empirical support for these predictions. Overall, we observe relatively modest net employment effects due to countervailing forces: reduced demand in AI-exposed occupations is offset by productivity-driven employment increases across all occupations at AI-adopting firms.

*We are grateful to Huben Liu, Weizhe Sun, and Tim Zhang for their excellent research assistance. We also thank Daron Acemoglu, Philippe Aghion, David Autor, Antonin Bergeaud, Tarek Hassan, Wei Jiang (discussant), Yueran Ma, and participants at various conferences and seminars for their helpful comments. Dimitris Papanikolaou and Bryan Seegmiller thank the The Financial Institutions and Markets Research Center and the Asset Management Practicum for their generous financial support.

Recent advances in artificial intelligence have re-ignited the perennial concern that technology will automate away most tasks performed by workers, leading to large declines in labor demand, depressed wages, and diminished job opportunities for workers. In contrast to prior waves of technological change, which have largely exposed middle- and low-skilled occupations (Autor, Katz, and Kearney, 2006; Autor and Dorn, 2013; Kogan, Papanikolaou, Schmidt, and Seegmiller, 2023), AI exposure appears to be concentrated in white-collar jobs (Webb, 2020; Eloundou, Manning, Mishkin, and Rock, 2023). However, despite the fact that firm investments in artificial intelligence have been underway for well over a decade, measuring the impact of AI improvements on labor demand has been elusive.[1] Part of the challenge is that advances in AI related to an occupation's tasks may actually increase demand for that occupation, for instance, if it increases firm productivity.[2] Our goal is to shed light on the distinct channels through which AI affects overall labor demand by using theory to guide measurement.

The first challenge is measuring the intensity and direction of AI adoption by firms, and then identifying which worker tasks are potentially affected by these AI applications. We do so by leveraging recent advances in large language models (LLMs) and natural language processing (NLP) techniques applied to a rich corpus of resume and job posting data from Revelio Labs. To measure AI adoption by firms, we first identify the key employees in each firm that are responsible for developing AI applications based on their resume. Using LLMs, we extract detailed information on how these workers apply AI to their firms. Even though our measure of AI adoption only relies on resume data, and therefore misses instances where firms outsource their AI needs to other firms, it aligns well with survey evidence on AI adoption by firms (BTOS).[3] Consistent with the findings of Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas (2023a), firms with high AI utilization tend to be larger, more productive, and pay higher wages.

The next step is identifying which worker tasks are exposed to these AI applications. Using modern NLP methods, we estimate the semantic similarity between the AI applications we have identified in the first and the individual tasks performed by specific occupations from ONET. Our first result is that skills related to tasks that are highly exposed to AI application developed by a particular firm in a specific year are subsequently significantly less likely to be mentioned in job postings by the same firm. Notably, the granularity of our measure allows us to saturate the

[1] For example, while Acemoglu, Autor, Hazell, and Restrepo (2022) find that firms with AI-exposed workforces have reduced job postings for non-AI positions, any aggregate impacts of AI-labor substitution on employment and wage growth in more exposed occupations and industries have been too small to detect.

[2] See for instance Acemoglu and Restrepo (2018, 2021). In addition, some technologies may complement rather than substitute for labor (Autor, Chin, Salomons, and Seegmiller, 2024; Kogan et al., 2023), which further complicates the link between a job's exposure to AI and labor demand.

[3] The fraction of firms that develop their AI applications internally is not small. Using data from the Annual Business Survey, Chequer, Herkenhoff, Papanikolaou, Schmidt, and Seegmiller (2025) report that among intensive AI users, 49% of firms are doing R&D in-house, and that percentage rises to 77% when weighting by firm employment.

specification with a rich set of fixed effects: our most conservative specification compares two tasks that differ in their AI exposure that are performed by the same occupation in the same firm at the same point in time. We find that, within a treated occupation–firm pair, a one standard deviation increase in *task*-level AI exposure reduces by 4.8% the relative demand for skills related to that task, as a share of the total skill requirements for that job. Given that AI-exposed tasks experience a relative reduction in labor demand, our operating assumption in the remainder of the paper is that AI applications that are similar to specific tasks performed by a specific occupation are a *substitute* for these tasks.

Not all worker tasks are similarly exposed to AI advances. Tasks performed by higher-paid occupations (those at the 90th percentile of the pay distribution) are, on average, more exposed to AI. However, the fact that AI can substitute for labor in specific tasks does not necessarily imply that occupations whose tasks are exposed to AI will experience reduction in labor demand. The reason is the existence of two countervailing forces. First, when AI substitutes labor in certain tasks, workers may shift their effort toward tasks that AI cannot perform, thereby enhancing their overall productivity. Second, if AI adoption improves firm productivity, affected firms may consequently increase employment, even within occupations highly exposed to AI. These dynamics illustrate why the direct substitution effect of AI on tasks does not necessarily translate into diminished aggregate labor demand for exposed occupations.

We next introduce a model to translate task-level AI exposures into changes in overall labor demand within specific occupations and firms that captures both direct and indirect effects. In our model, each worker's output in a given occupation is an aggregate of multiple tasks, each potentially affected by technological advancements. These advancements are represented as improvements in task-specific intangible capital which serves as a substitute for labor in each task. The scope of these technologies can vary, influencing either most tasks within an occupation—such as a customer service chatbot——or only targeting specific tasks like automated expense report filing. Crucially, workers optimally allocate their time across tasks. Improvements in automation technology for one task directly influence its price but also indirectly impact other tasks performed by the worker. These indirect effects depend critically on the flexibility of task reallocation, the elasticity of substitution between labor and capital, and the degree of complementarity among tasks both within occupations and across occupations within the firm.

The central contribution of our model is identifying two primary variables that summarize the labor market effects of AI technologies. First, since AI substitutes for human labor in exposed tasks, the *mean exposure* of an occupation's tasks to AI is generally negatively correlated with labor demand for that occupation. Consequently, even modest advancements in technologies applicable across all tasks—such as a basic customer service chatbot—can decrease the demand for workers in the relevant occupation. Second, the extent to which the occupation's mean task-level exposure

$m(\varepsilon)$ is *concentrated* in a small number of tasks positively influences labor demand. For instance, an automated expense reporting system allows impacted workers to redistribute effort to unaffected tasks, potentially enhancing their overall productivity. Together, these two metrics comprehensively characterize labor demand shifts within firms. In addition, there is a third channel affecting labor demand across firms that hinges on productivity gains realized by firms from AI adoption.

In sum, the model implies that whether AI actually displaces an AI-exposed occupation is highly ambiguous, as it depends on the relative strength of direct substitution effects, indirect effects that operate through reallocation of effort in potentially complementary tasks, and the aggregate effects that operate through changes in firm productivity. Thus, the main part of the paper focuses on estimating these channels in the data. Using our estimates of task-level exposures to AI developed in the first part of the paper, we proceed to develop direct empirical analogues of the two key variables implied by the model: the mean and concentration of AI exposure within a job (an occupation–firm pair) at a particular point in time. Thus, the granularity of our data allows us to measure the AI exposure of different occupations employed in specific firms at a particular point in time.

A key challenge in our empirical analysis is that AI adoption is not necessarily random across firms and can be targeted to specific occupations. The first endogeneity challenge will lead to biased OLS estimates of AI adoption on productivity if adoption costs are correlated with the growth potential of the firm. The second endogeneity concern is that AI-based automation can be directed to specific occupations, for instance, if these workers are scarce. Selective implementation of AI that targets the tasks of specific occupations will bias the OLS estimates of AI exposure on occupation labor demand. Thus, to understand the full impact of AI on productivity and labor demand, we need to separate the selective adoption from the causal impact of AI.

To address the endogeneity of firm adoption of AI, we build on Babina, Fedyk, He, and Hodson (2024) and instrument for a firm's employment of AI-focused employees with the growth in AI graduates in universities from which the treated firms tend to hire workers. For example, if State Street tends to hire relatively more graduates from Boston University, while BNY Mellon tends to hire relatively more graduates from Wharton, the fact that Boston University produces more graduates with AI-related majors than Wharton is likely to lead to a lower relative cost of adopting AI in State Street than in BNY Mellon.[4] In addition, we construct a shift-share instrument for our firm–occupation measures of mean and concentrated AI exposure: we instrument for these measures using the mean and concentration in task exposure to AI across all applications in the same period and occupation (the shift) times the predicted intensity of AI adoption at the firm level using our instrument based on pre-existing tendencies to employ workers who graduated from

[4]Thus, the key variation in our IV strategy comes from labor market exposure to AI through persistent university hiring networks. Firms that historically hired from universities whose graduates later entered AI-related occupations are more likely to adopt AI themselves as they face a lower adoption cost due to their increased exposure to supply-side increases of AI-related labor via university pipelines.

particular universities.

We find that adoption of AI by firms leads to higher growth in revenue, measured productivity, profits, and employment. These patterns are consistent with Acemoglu et al. (2023a); Babina et al. (2024). Notably, our IV estimates are approximately 30 percent larger than our OLS estimates; a likely explanation is the selective adoption of AI by larger and more profitable firms, which tend to grow more slowly on average than smaller firms.

We next turn our attention to the demand for workers in affected firms and occupations. The granularity of our AI exposure measures allows us to saturate our specifications with a rich set of controls, including the interaction of firm and occupation with calendar-year fixed effects, with the estimates largely comparable across specifications. Thus, our specifications compare employment growth in differentially exposed occupations within the same firm, but also in the same occupation across different firms that adopt AI with different intensities or focus. Consistent with the predictions of our model, the mean exposure of an occupation's tasks to the AI applications adopted by their employer is negatively related to subsequent employment growth in that occupation in that firm. By contrast, greater concentration in the occupation's exposure to these AI applications is associated with greater employment growth. The IV results are comparable to our OLS results qualitatively, but are somewhat larger in magnitude. For instance, in our preferred IV specification, a standard deviation increase in our measure of mean task-level exposure to AI leads to approximately a 14.5 percent decline in the within-firm employment share of affected occupations over the next 5 years, while a standard deviation increase in the concentration measure leads to a 7.5 percent increase in the within-firm employment share.

In brief, our empirical results reveal a strong effect of AI on firm productivity and labor demand, but also evidence for within firm labor reallocation. These within-firm reallocation results are consistent with both the presence of strong AI–task substitution, but also productivity spillovers across tasks within an occupation, which considerably dampen this direct substitution effect. The result is that the overall effect of AI exposure on the relative demand for affected occupations within the firm is muted. To understand how these different forces impact overall labor demand for affected occupations, in the last part of our analysis we use our empirical estimates to quantify the net impact of AI on firm labor demand, both overall, but also across different occupations, job types, and worker earnings levels.

Overall, we find that the aggregate impact of AI on firm labor demand is muted, due to the presence of counter-veiling forces. Even though the direct substitution effect (the mean task exposure) is quantitatively stronger, there are significant labor-augmenting effects arising due to task reallocation (the concentration of AI task exposures) and the impact of increased firm productivity on firm labor demand. In addition, though these effects vary somewhat across the worker pay distribution, the net effects are more homogeneous across jobs than the underlying forces would suggest. In particular, the labor-substitution part of AI is stronger across higher-paid occupations, but then so is the reallocative effect due to concentrated task exposure. Thus, focusing on within-firm job reallocation, we find that employment of highly-exposed occupations (those at the 90th percentile of the pay distribution) declines by about 3.1% relative to the employment of the least exposed workers (those at the bottom of the pay distribution). After taking into account the firm productivity effect, this effect is mildly reversed, since the jobs that are more exposed to AI are more likely to exist in firms that adopt AI and realize productivity gains. Thus, the high-income jobs that are most exposed to AI actually experience a *slight increase* in their share of aggregate employment compared to low-income, less exposed jobs.

The fact that these aggregate estimates are muted does not necessarily imply there are no winners and losers. The most adversely impacted occupations fall under the business, financial, and engineering categories. In terms of share of overall employment, our estimates imply that these jobs experienced a decline of almost 2% over a five-year period. As before, the existence of task complementarities and a strong firm productivity effect act as offsetting factors to the significant substitution effect at the task-level. At the same time, however, our estimates also imply declines in the overall employment share of less exposed occupations, such as 'Food Preparation and Serving'. Even though these occupations face very little task exposure to AI, the fact that their employers do not utilize AI implies that they grow slower than the firms that do use AI, leading to lower labor demand overall.

In sum, this exercise illustrates why it may be difficult to detect the impact of AI technologies on labor reallocation overall, despite the strong evidence of AI–labor substitution at the task level. That said, however, we do find that our measures of AI exposure still have meaningful explanatory power for the realized reallocation of workers across jobs and tasks within jobs. Comparing the realized employment share growth at the occupation level with the total AI–predicted changes, we find that about 14% of the variation in observed aggregate occupational employment share changes can be attributed to the net effects of AI exposure. About half of this is driven by the reallocation away from AI–exposed occupations within the firm, with the remainder coming from the average firm–level AI use by occupations' employers.

Our work contributes broadly to existing work in economics studying the impact of technological change on the labor market. Closest to our paper is work that explores the impact of machine learning and AI on firm productivity and the labor market. Acemoglu et al. (2022) find some evidence that AI substitutes for labor at the establishment level, but they find essentially no impact of employment and wage growth for exposed occupations. Babina et al. (2024) find that AI adoption has an impact on firm growth, but find no impact that it leads to productivity gains or job automation. Acemoglu et al. (2023a) argue that the correlation between AI adoption and firm growth is driven by selection of which firms adopt advanced technologies. Gathmann, Grimm, and

Winkler (2024) argue that, in contrast to robots, AI has reduced the demand for abstract tasks and increased the demand for certain routine tasks. Aghion, Bunel, Jaravel, Timo Mikaelsen, and Sogaard (ming) is close to our work, as they find that what matters for labor demand is not only the exposure to AI, but also the type of the AI technology.[5]

Our measure of AI ends in 2023, and therefore it largely excludes the recent rise of generative AI (GenAI). Despite the relatively short period since GenAI first became broadly available in November 2022, some early work has studied its effect on firms and workers. Eloundou et al. (2023) construct an occupation-level exposure measure to generative AI (GenAI). They document that most occupations have significant exposure to GenAI, and unlike past instances of automation, high-wage occupations are significantly more exposed than low- or middle-wage jobs. Eisfeldt, Schubert, Taska, and Zhang (2023) build on Eloundou et al. (2023) and construct a firm-level measure of workforce exposure to GenAI. Using data on job postings and average wages at the occupation level, they argue that GenAI has helped firms improve profitability by reducing labor costs. In contrast to Eloundou et al. (2023); Eisfeldt et al. (2023), Auer, Köpfer, and Sveda (2024) argue that GenAI can be a complement for high-wage workers, and argue that low-wage jobs are more likely to be displaced in the future. Humlum and Vestergaard (2024) find that the adoption of GenAI tools by workers is widespread but uneven, with increased adoption among younger, higher-achieving and male workers. Using a version of Hulten's theorem, Acemoglu (2024) argues that generative AI will likely have small effects on productivity.

Compared to the existing body of work, our work has some key advantages on the measurement side. Specifically, our data and methodology allow us to construct direct measures of the exposure of a particular occupation in a specific firm to Artificial Intelligence. Our measures are grounded in theory, and, critically, they differentiate between the direct substitution effect of AI technologies on labor demand from the reallocative effect. Similar to Babina et al. (2024); Babina, Fedyk, He, and Hodson (2023), we use information in online resumes to determine whether a firm is actually adopting AI; in contrast to their work, however, our exposure measures are based on how each particular AI application adopted by a specific firm is related to the particular tasks a given occupation performs. The detailed nature of our measure allows us to separate the impact of AI from occupation-specific trends in employment, unlike existing measures that vary only across occupations (Webb, 2020; Felten, Raj, and Seamans, 2018; Brynjolfsson, Mitchell, and Rock, 2018; Eloundou et al., 2023; Eisfeldt et al., 2023). Further, our focus on actual changes in employment, rather than job postings (Acemoglu et al., 2022; Eisfeldt et al., 2023), helps address the concern that firms may hire multiple employees from one job posting. The increased granularity of our data helps us understand why the aggregate employment effects of AI have been muted, despite clear evidence for substitution at the

[5]Also related is Jiang, Park, Xiao, and Zhang (2025), who find that AI use leads to longer work hours, and Aum and Shin (2025), who document evidence that digital technology investments, including AI, have reduced labor demand for skilled workers in Korea.

micro level.

More broadly, our work connects to the literature studying the broad effects of labor-substituting technologies. An important stream has emphasized that automation can substitute for routine tasks, and therefore an occupation's share of routine tasks is indicative of its exposure to labor-saving technologies over the last several decades (Autor et al., 2006; Acemoglu and Autor, 2011; Goos, Manning, and Salomons, 2014). Another stream has emphasized constructing direct measures of labor-saving technologies and exploring their impact on the labor market (Webb, 2020; Jiang, Tang, Xiao, and Yao, 2021; Acemoglu and Restrepo, 2021; Humlum, 2019; Dauth, Findeisen, Suedekum, and Woessner, 2021; Koch, Manuylov, and Smolka, 2021; Bonfiglioli, Crinò, Fadinger, and Gancia, 2020; de Souza and Li, 2023; Benmelech and Zator, 2022; Kogan et al., 2023; Autor et al., 2024; Mann and Püttmann, 2023; Jiang et al., 2025; Dechezleprêtre, Hémous, Olsen, and Zanella, 2021; Aghion, Antonin, Bunel, and Jaravel, 2020).

1 Identifying Tasks Exposed to AI

We begin by estimating which worker tasks are likely to be exposed to advances in AI. First, we leverage resume data of AI developers to identify which applications are developed by specific firms. This approach delivers a granular measure of firm adoption of AI, which we subsequently validate based on survey data. Subsequently, we identify which worker tasks are more likely to be affected by these AI applications using natural language processing. Last, we show that our task-level exposure measures likely capture the substitutive effect of AI: skills that are related to tasks exposed to AI are less likely to be mentioned in subsequent firm job postings.

1.1 Data

Our primary dataset comes from Revelio labs, a leading workforce database provider that has collected the near-universe of Linkedin Profiles and job postings, which we link to publicly traded firms in Compustat. The resume data includes comprehensive details on individuals' educational and employment histories, including universities attended, fields of study, employers, job titles, employment dates, and self-reported descriptions of jobs. The job postings data includes company name and other identifiers, job taxonomies (e.g., SOC, NAICS), posting dates, removal dates, seniority, salary (as estimated by Revelio if missing), job location (state, city, MSA, and ZIP code), full posting text, and, if the firm is public, relevant stock market identifiers (e.g., CUSIP).

We focus our main analysis on job positions that were active between 2014 and 2023, which in addition to being the most recent full decade for which we have data, is also a period marked by both improved resume coverage and the rising prevalence of artificial intelligence in the workforce. However, when constructing our instrument, we also incorporate job positions dating back to 2005.

For the job postings data, we include records from 2010 to 2023, leveraging earlier years to construct our measures of task reallocation, as detailed in Section 1.4. To reduce the size of the raw job postings dataset, which exceeds a billion records before filtering, we randomly sample up to 10 postings per year for each occupation-firm pair. We further restrict our analysis to publicly traded companies, allowing us to merge labor market data with firm-level financial and performance metrics.

For tagging AI applications, we limit our analysis to resume job positions with a valid job description and a U.S. location, ensuring that position descriptions are in English. To count employment in non-AI positions within a firm, we additionally require a valid occupation identifier. In both cases, positions must be linked to a gvkey in Compustat. A position is classified as 'active' in a given year if the recorded start and end dates indicate that the worker held the position for at least six months within that year. After applying these restrictions, our dataset includes approximately 58 million LinkedIn profiles and 14 million job postings that are linked to Compustat firms, and meet the sampling criteria.

Additionally, we supplement our main data with O*NET and Compustat. O*NET's (Occupational Information Network) is a comprehensive database developed by the U.S. Department of Labor that provides detailed descriptions of occupations, including the specific tasks, skills, and knowledge required for each job. We use Compustat in order to identify publicly traded firms and other key financial information used in the analysis.

The Revelio resume data are not drawn from a random sample of the U.S. workforce; rather, they are constructed from publicly available resumes that individuals self-report and choose to post online. As such, the data tend to overrepresent workers in professional occupations, younger cohorts, and those with higher levels of digital engagement. Nevertheless, multiple studies suggest that these data provide broad and reliable coverage of the white-collar labor market. Tambe (2025) shows that the distribution of education, occupation, and age in Revelio closely tracks the CPS and ACS, particularly among managerial and professional roles. More broadly, Hershbein and Kahn (2018) validate the use of online job postings data (from Lightcast, formerly Burning Glass) by demonstrating that occupation- and industry-level trends in postings closely mirror those in the BLS's Occupational Employment Statistics. Taken together, these findings support the view that online labor market data—while subject to known selection biases—are a credible source for analyzing structural changes in the labor market. In particular, because Revelio overrepresents workers in high-skill, digitally intensive sectors such as professional services, tech, and finance, it is especially well-suited to studying the labor market impact of AI, which is disproportionately concentrated in these segments.

The Revelio data have very high rates of coverage precisely the population of employees who are involved with the development and deployment of AI technologies within their respective firms. Online resumes provide a useful forum for these employees to showcase their expertise and potentially attract new job opportunities by discussing projects with which they have been involved, hence employees have strong incentives to describe their activities in at least some detail.[6] Since deploying AI technologies frequently requires some degree of in-house investment, these detailed descriptions from online resumes provide a unique opportunity to study both the level of adoption and the direction of AI investments. It also allows us to capture adoption activities which are undetectable using patent data, another valuable source of textual information on new technologies often used in the literature. Capturing these additional activities is important given that many AI applications involve combining fairly standardized prediction algorithms with a firm's proprietary data and are therefore are unlikely to be patentable.

Much of our analysis relies on employment counts derived from resume data. When aggregated to the firm-year level, Revelio-based employment is highly correlated with firm-level employment reported in Compustat. Appendix Figure A.1 illustrates this relationship by plotting the log of resume-based employment against the log of Compustat employment, both in levels and in five-year changes—the horizon we use to measure firm growth. The correlation between the two data sources is 0.76 in levels and 0.56 in five-year changes, indicating that the resume data provide a reliable proxy for firm-level employment dynamics.

1.2 Extracting AI applications

We next describe our methodology for measuring firm adoption of AI.

Methodology

To identify which firms are using AI and for what purposes, we first examine job descriptions where workers explicitly mention AI usage. We provide a brief overview of this process here and detail it further in Appendix A.2. In the initial step, we search for AI-related keywords within job position descriptions in resumes from the Revelio database to identify roles where AI may be utilized. We then employ a large language model (Llama 3.1 70B) to refine and extract specific phrases within job postings that describe how AI is being used. This process involves filtering job positions through multiple queries to the language model to systematically tag and clean AI-related information.

The first query identifies and cleans relevant phrases. Since our focus is on *how* AI is being used rather than the specific tools being mentioned, the second query removes references to specific AI tools from the extracted text. It also filters out vague AI applications that do not specify a concrete use case (e.g., a phrase like "AI tools are being used to deploy computer vision models' lacks details on the intended purpose of the models and is therefore discarded"). A final, third query eliminates

[6] For example, as of mid-2023, LinkedIn's website reported that there were "117 job applications submitted every second", "8 people hired every minute", and "45% of hirers on LinkedIn explicitly use[d] skills data to fill their roles." Recruiting blogs frequently talk about the importance of having complete LinkedIn profiles for improving the efficiency of job search by generating more engagement from recruiters.

any remaining non-specific AI applications that were not successfully filtered in the second step. This process yields over 1 million distinct AI use cases derived from a little over 500,000 distinct job positions. We classify an AI application as being in use at firm f during a given year t if the corresponding job position on the resume was active for at least six months within that year.

Figure 1 illustrates the process of identifying AI applications adopted by specific firms with a specific example from the resume of a worker in an AI-implementing role at JP Morgan. The worker describes their position as follows:

> Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments.

From this description, our process initially screens this resume by identifying the terms "AI/ML" as AI-related. Subsequently, our LLM queries identify "Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses" and "development and deployment of quantitative risk models that serve regulatory and credit risk assessments" as phrases which encode distinct ways the worker uses AI. Our LLM queries then clean and process the AI applications. The first AI-related phrase becomes the following concrete AI application: "Forecast risk and fraud in various lending businesses, including auto, card, and home lending"; the second phrase becomes "Assess credit risk and provide regulatory compliance across different lines of business."

One potential shortcoming of our approach is that it only identifies the adoption of AI for firms that develop these applications internally. However, that is not necessarily a small sample: for instance, using data from the Annual Business Survey, Chequer et al. (2025) report that among intensive AI users, 49% of firms are doing R&D in-house, and that percentage rises to 77% when weighting by firm employment. Relatedly, Bonney, Breaux, Buffington, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, and Savage (2024) report that among the set of firms that report using AI in the BTOS, in the previous six months, about half of firms (on an employment-weighted basis) report having trained current staff to use AI, and around 10% report having hired staff trained in AI. Nevertheless, we next proceed to validate our AI adoption measure versus alternative data sources.

Validation

We validate our measurement of AI adoption at the firm level in several ways. First, we compare the frequency at which a firm uses AI technologies based off our resume data with survey-reported AI utilization rates in the Census Business Trends and Outlook Survey (BTOS). Bonney et al. (2024) show that this survey can be used to track firm AI use in real time. The BTOS surveys firms regarding their technology usage, including whether or not they are currently using AI technologies in production. They also report the survey reported rates in their data by sector (2-digit NAICS) by firm size bins. We compare with their data by checking whether we can observe an AI worker resume at given firm in the Revelio data. We compute the average probability at the same sector × firm size level and compare with the BTOS data.[7] The resulting scatterplot is in panel A of Figure 2, where we compare the share of firms in the BTOS who report using AI in that sector × size bin (y-axis) with the same rates in Revelio as of the end of 2023 (x-axis). We report the graph in a log scale to allow for better visual separation of the group of cells with low utilization rates. The two series have a strikingly high correlation (0.9 in levels). Thus the probability of a worker report using AI at their job in a given sector/size group is closely related to the rate at which employers self-report using AI in surveys.

Next, we compare the tendency of workers to report starting a new AI position at specific firm with the rate at which firms state demands for new AI positions at the same time. Specifically, in Appendix Figure A.2, we compare the AI-related job postings at firm f in year t with the *new* AI resumes that begin work at the firm in that year. We tag job postings as AI-related using the same keywords we used to find AI resumes. The figure plots a residualized binscatter of $\log(1 + \text{Newly-added AI Resumes}_{f,t})$ (y-axis) against $\log(1 + \text{AI-Related Job Postings}_{f,t})$ (x-axis). We control for the logs of total resume-based employment and total job postings in that year, as well as year fixed effects. The figure shows a strong positive relationship. The partial correlation between the two measures is 0.67. Thus, when workers report on their resumes that they began a new AI-related position at a given firm, this also coincides with periods in which their employer was stating its demand to hire employees in new AI-related job positions.

Next, we examine whether firms with workers who report AI use also patent in AI technologies. We measure AI patenting based off the USPTO AI patent database from Pairolero, Giczy, Torres, et al. (2025) (available here). In Appendix Figure A.3, we find a highly-significant partial correlation of 0.36 between AI patenting status and the extent of AI use at the firm (controls include the log of resume-based employment and a non-AI patenting indicator). Thus while not all AI-using firms patent in AI each year (as would be expected), the increased use of AI as represented in resumes is correlated with the tendency to do so.

Finally, Babina et al. (2024) also construct resume-based measures of AI employment at the firm using data from Cognism instead of Revelio, although they identify AI workers using related but slightly different criteria. Appendix Figure A.4 shows that our firm-level AI worker counts are

[7] The BTOS aggregated statistics can be found here. We focus on question 7. Since many "yes" response rates are suppressed, we use as our measure of adoption the average of one minus the "no" rate across the 2023-24 survey waves for a given sector × size bin. Since the survey includes private firms, here we rely on Revelio firm identifiers rather than restricting to Compustat firms, as in most of our analysis. To allow for additional comparability with the BTOS's sampling frame of single-unit employers, we also restrict the top size bin to firms that have fewer than 1,000 employees in the Revelio data.

highly-correlated with theirs (partial correlation of 0.75 after controlling for the log resume-based employment from both data sources).

Which firms adopt AI?

We next examine what types of firms tend to use AI and its impact on subsequent firm outcomes in the 2014-2023 period that is the focus of our analysis. Panel B of Figure 2 reports the distribution of resume-implied AI utilization rates by sector and firm size. We find that information and professional services report the highest rates, followed by finance and insurance; on the other end of the spectrum, accommodation and food services and other services have the lowest rates. Within each of the major NAICS sectors, we observe a monotonic relationship betwen firm size and AI utilization rates. In Table 1 we examine the relationship between $\log(1 + N_{f,t})$—where $N_{f,t}$ is the number of AI applications in place at firm f in year t—and log sales per worker; log sales; log profits (defined as sales minus cost of goods sold); log TFP (which we estimate for Compustat firms following the procedure outlined in İmrohoroğlu and Şelale Tüzel (2014)); and finally, the log average salary for each position at the firm based off Revelio–imputed salaries to see which types of firm tend to use AI technologies. We control for 3-digit NAICS industry by year fixed effects and the log of total Revelio employee resumes in each specification because of the potentially mechanical relationship between the number of resumes we can observe and the number of distinct AI applications we find at the firm level. We cluster standard errors by firm.

In Table 1, we find a strongly positive relationship between the extent of firms' AI utilization and productivity measures like log sales per worker or TFP, as well as sales, profits, and average pay. We scale the main independent variable to unit standard deviation. We find that a one standard deviation increase in AI utilization is associated with roughly 12% increases in both productivity measures, while sales and profits increase by about 31% and 42%, respectively; the marginal effect on log average salary to a standard deviation increase in AI use is about 11%. This table illustrates that AI–using firms are quite different from other firms, being larger, more productive, and higher paying. This echoes survey evidence on the cross-sectional relationship between advanced technology adoption and firm characteristics (Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas, 2023b; McElheran, Li, Brynjolfsson, Kroff, Dinlersoz, Foster, and Zolas, 2023). These correlations are consistent with the presence of economies of scale from AI adoption and/or larger firms having more data available for predictive algorithms. We emphasize that these relationships are only correlational, and could reflect both causal relationships as well as selection into which types of firms are able to take advantage of AI.

Heterogeneity in AI adoption across firms and sectors

A unique advantage of our resume-based approach to extracting AI applications is that it provides a highly granular picture of the *direction* of AI investments. Just as patents provide very detailed information about new products or processes, employees' position descriptions on online resumes offer a detailed view into the heterogeneous ways in which different firms are applying AI.

To illustrate this, we consider the case of two large, publicly-traded firms from our sample: JPMorgan Chase (JPMC) and Walmart. Given that both firms have a large number of AI applications that are tagged by our process, we use a clustering algorithm in order to summarize them. In particular, we use a k-means algorithm to partition the set of all applications at each firm into 5 categories, then use a LLM in order to supply labels for each cluster. The left panel of Table 2 provides several of these clusters. Intuitively, two of the key AI applications at JMPC are "Predictive Modeling and Financial Forecasting" (useful for both underwriting and asset management applications) as well as improving customer service via "Customer Engagement and Personalization" applications. By contrast, two examples for Walmart involve leveraging AI's capability to improve the efficiency of pricing and operations via "Forecasting, Pricing, and Supply Chain Optimization" and "Process Automation and Operational Efficiency". Both firms also have a cluster of applications related to fraud detection: "Fraud Detection, AML & Risk Mitigation" at JPMC and "Fraud, Security, and Anomaly Detection" at Walmart. Given that both firms process extremely high volumes of transactions, it is intuitive that both firms are making sizable AI investments to improve their ability to detect fraudulent ones.

We next conduct a similar analysis at a larger scale, again using a clustering algorithm to partition all AI applications from post 2010 into 20 categories. Once again, we provide examples from each cluster to a LLM, then ask it to supply labels for each cluster.[8] Figure 3 provides a heatmap which captures the fraction of AI applications (in each row) that appear in each cluster by major NAICS sector (in each column). Darker colors indicate a higher fraction of applications originating from a given cluster, where we normalize frequencies so that the total fraction of AI application sums to 1 in each column. Clusters are ordered vertically in descending order by their average prevalence across industries.

Examining Figure 3, we see some clear patterns regarding the breakdown of these applications across industries. The cluster which has the largest share of AI applications across most sectors is "Business Intelligence Insights", a sensible pattern given that AI-related predicted analytics are likely to be applicable for operations in most industries. However, some applications are significantly more concentrated in individual sectors, such as "Financial Risk Modeling" in Finance, "Personalized Recommendation Engines" in Retail, "Demand and Sales Forecasting" in Utilities, and "Healthcare

[8] In addition, Appendix A.3 includes longer descriptions as well as several representative examples of the applications which appear in each cluster.

Diagnostics and Genomics" in Health. Within Education, which notably includes researchers at higher education institutions, "Healthcare Diagnostics and Genomics", "Scientific and Industrial Modeling", "Image and Video Recognition", and "Autonomous Navigation and Robotics" emerge as the most prevalent AI applications. Again, this analysis points to the wealth of information available about AI use cases in online resumes, information which is likely to be considerably more detailed than what can be gleaned even from extant surveys of firms.[9]

1.3 Building a task-level exposure measure

After identifying AI applications in the previous section, we next measure the extent to which non-AI workers' tasks are exposed to these applications. Our starting point is that a specific worker's tasks that are semantically similar to a specific AI application is likely to be affected by that application. In the first part of this section we describe the methodology through which we estimate this semantic similarity and provide specific examples. In the second half, we show that skills related to tasks that are exposed to AI applications are less likely to be mentioned in subsequent job postings by the firm adopting the application, which strongly suggests that our exposure measure is identifying substitution between AI and labor in performing that task.

Estimating the similarity between AI applications and tasks

We define an occupation using the detailed 6-digit SOC codes and obtain the text of worker tasks from the O*NET database.[10] To measure the semantic similarity between AI applications and worker tasks, we use text embeddings. Text embeddings encode the semantic meaning of a document as a geometric vector representation. When generated from the same embeddings model, documents with similar textual content will exhibit high cosine similarity. We utilize the GTE-Large embeddings model, developed by Alibaba DAMO Academy, which encodes text into a 1096-dimensional vector and has demonstrated strong performance across a range of document similarity tasks compared to

[9]Consistent with our measure being comprehensive, many applications identified by our unsupervised learning procedure coincide with specific technologies about which the Census asks about in its AI supplement to the BTOS survey, which are (in descending order of the employment-weighted fraction of firms using them from Bonney et al. (2024)): 'data analytics using AI', 'robotics process automation', 'virtual agents or chat bots', 'text analytics using AI', 'speech/voice recognition using AI', 'machine/computer vision', 'marketing automation using AI', 'image/pattern recognition', 'natural language processing', 'biometrics', 'recommendation systems based on AI', and 'decision making systems based on AI'. Remaining categories involve the use of broader tools like 'machine learning', 'deep learning', and 'neural networks'. Only 'augmented reality', the least common choice, lacks a direct parallel on our list of clusters.

[10]While the O*NET database includes over 800 such codes, in our dataset, Revelio assigns resumes to 335 distinct 6-digit occupation codes, resulting in a slightly more aggregated classification in practice. We use task data from the O*NET 26.1 release (June 2022), published by the U.S. Department of Labor. For each 6-digit SOC code, we extract the list of Detailed Work Activities (DWAs)—short natural language descriptions of core tasks performed in the occupation. These DWAs form the textual foundation for our occupation-level measures of AI exposure. When an occupation in the Revelio data does not have a direct match in ONET, we assign it to the closest corresponding SOC code based on standardized crosswalks provided by ONET.

other models of similar scale.[11] Using the GTE embeddings, we represent each of the approximately 1 million cleaned AI applications as a 1096-dimensional vector, and we do the same thing for each of the roughly 20,000 job tasks in the O*NET database by representing them by a 1096-dimensional vector, as well. The end result is a matrix of 1 million cleaned AI applications times 20,000 job tasks of application–task cosine similarities $\rho(i, j)$.

Next, we compute a measure of task exposure for task j, performed by occupation o in firm f at year t,

$$\text{Exposure Probability}_{j,f,t} = \frac{1}{N_{f,t}} \sum_{i=1}^{N_{f,t}} I_{j,i}^{95}. \tag{1}$$

We impose that the vast majority of AI application/occupation task pairs should be unrelated. We only consider a task exposed to a particular AI application if the cosine similarity of the two text embeddings is above the unconditional 95th percentile in the overall task–AI application distribution of similarity scores. That is, the variable $I_{j,i}^{95}$ takes a value of 1 if the cosine similarity $\rho(i, j)$ of the GTE text embedding for the task j and AI application i are above the 95th percentile in the distribution across all potential (i, j) pairs.

Examples

The right panels of Table 2 report examples of O*NET tasks (middle column) and their associated occupations (right column) which have the highest average exposure probability to AI applications within each of the clusters we identified for JPMC and Walmart, respectively. Beginning with JPMC, the applications related to fraud detection, predictive modeling, and financial forecasting are most highly related to tasks performed by Other Financial Specialists as well as Accountants and Auditors. By contrast, the most exposed task in the "Customer Engagement and Satisfaction" cluster is to "monitor customer preferences to determine focus of sales efforts", a task performed by Sales Managers. Walmart's locus of AI exposure is meaningfully different from JPMC.

Within Walmart's fraud detection category, the most exposed tasks are to "analyze retail data to identify current or emerging trends in theft or fraud" (performed by Other Managers) or to "monitor machines that automatically measure, sort, or inspect products", which differ notably and intuitively from JPMC. In the first cluster related to pricing and supply chain optimization, Purchasing Managers as well as Wholesale and Retail Buyers emerge as the most highly exposed occupations. In the second cluster (process automation and operational efficiency), tasks related to Sales Engineers and Other Engineering Technologists are most highly exposed.

[11] Earlier word-specific embeddings models, such as GloVe and word2vec (Pennington, Socher, and Manning, 2014; Mikolov, Sutskever, Chen, Corrado, and Dean, 2013), have been widely used in economic research to measure document similarity (Seegmiller, Papanikolaou, and Schmidt, 2023; Kogan et al., 2023; Autor et al., 2024), but these models assign a fixed vector to each word regardless of context. Recent advancements in embeddings technology allow for models that capture the full contextual meaning of text.

These examples help to illustrate the rich heterogeneity in exposure which is detectable with our measure. Indeed, two workers employed in the same occupation at Walmart and JPMC may not be equally exposed to AI adoption because the firms direct their AI investments in different ways. These exposures are heterogeneous even when comparing two workers in the same occupation and industry. We exploit this highly granular information throughout our analysis.

1.4 AI Exposure and Demand for Tasks

Armed with a measure of the exposure of specific tasks to AI, we can then explore what our measure of exposure really captures: are tasks that are semantically similar to an AI application *complemented* or *substituted* by Artificial Intelligence? To answer this question requires data on how the demand for specific tasks has fluctuated over time. To do so, we construct a proxy for the the intensity of task utilization using the composition of skill demands in firms' online job postings. We obtain texts of online job postings for occupations from Revelio, and we tag each job posting with the list of skills the job posting requires using the Open Skills API provided by LightCast.[12] After tagging the skills in each job posting using the LightCast API, the average job posting in our dataset lists 17 distinct skills.

The Revelio job posting texts are linked to the firm that posted the job and the relevant occupation. Additionally, LightCast provides a textual description of each skill, which allows us to textually link the skills with associated occupation job tasks in O*NET. LightCast identifies approximately 30,000 distinct skills, and we impose sparsity in the textual linkages by connecting a skill as applicable to a given task if the cosine similarity of the task and skill description's GTE embeddings are in the top percentile of the distribution across all pairs, implying that the typical task has around 300 associated skills that could be applicable when performing the task. We then examine how the share of total skills demanded across job postings listed by firm f for occupation o are linked to task j,

$$Share_{j,o,f,\tau} = \frac{\text{\# of skills linked to task } j \text{ in occupation } o \text{ job postings at firm } f \text{ over time period } \tau}{\text{Total \# of skills in occupation } o \text{ job postings at firm } f \text{ over time period } \tau}$$

For a given year t, we take τ to be the 5-year window from $t-4$ through t inclusive. We then look at the Davis, Haltiwanger, and Schuh (1996) changes in these shares between consecutive 5 year periods. The Davis et al. (1996) (DHS) change is a second-order approximation to the log change, so coefficients can be interpreted in units of percentage changes, but it also accommodates cases

[12] A growing literature uses the LightCast (formerly Burning Glass) lists of job posting skills. A few examples include Deming and Kahn (2018), Deming and Noray (2020), Acemoglu et al. (2022), and Braxton and Taska (2023).

where one of the shares is equal to zero,

$$\Delta_{DHS}Share_{j,o,f,\tau+1} = 2 \times \frac{Share_{j,o,f,\tau+1} - Share_{j,o,f,\tau}}{Share_{j,o,f,\tau+1} + Share_{j,o,\tau}} \tag{2}$$

Here τ denotes the 5-year period that includes the years $t-4$ through t, inclusive. We use $\Delta_{DHS}Share_{j,o,f,\tau+1}$ to approximate the growth in the importance of task j to occupation o at firm f. Note that this measure now varies at the *task*–by–firm level, rather than the occupation–by–firm level as before. In a parallel manner, we construct a task–level exposure measure; since our unit of analysis is now at the task–level within an occupation, we only need to include the task–level exposure (proxy $\epsilon(j)$ in the model). To be consistent with our definition of occupation–level average exposure in (24), we allow task–level exposure to be a function of how likely a task is to be exposed to a given AI application at the firm, multiplied by how intensively the firm utilizes AI:

$$\text{Task-Level AI Exposure}_{j,f,t} = \text{Exposure Probability}_{j,f,t} \times \log(1+N)_{f,t}, \tag{3}$$

where Exposure Probability$_{j,f,t}$ is defined in (1). Similarly, our task–level instrument now becomes the likelihood of task j being exposed to an AI application from any firm, and interacted with the predicted number of AI employees:

$$\text{Task-Level AI Exposure}^{IV}_{j,f,t} = \text{Exposure Probability}_{j,t} \times \log(1 + \text{Predicted AI Employees})_{f,t}. \tag{4}$$

With these measures in hand, we estimate the following specification:

$$\Delta_{DHS}Share_{j,o,f,\tau+1} = \beta\, \text{Task-Level AI Exposure}_{j,f,t} + \alpha_{o,f,t} + \delta X_{j,o,f,t} + \epsilon_{j,o,f,t}. \tag{5}$$

The controls X include the task importance weight $\omega_{o,j}$ derived from O*NET task importance scores as explained in Section 2.2, which we take to be a rough notion of how much effort someone in this occupation would typically to allocate to this task on average. In specifications without the full complement of occupation interacted with firm interacted with calendar year fixed effects, we also control for the occupational average AI exposure to identify within-occupation task reallocation. We cluster standard errors at the occupation–firm level, and weight each observation by the number of workers in the occupation–firm cell multiplied by the task importance weight. We scale our task-level exposure measure (4) to unit standard deviation, and we multiply the dependent variable to 100 so that coefficient estimates can be interpreted as the percentage response to a standard deviation increase in task–level AI exposure.

Table 3 reports our estimates of equation (5). We see that our point estimates imply that an increase in the task–level AI exposure reduces the intensity at which firms demand skills related to that task. These estimates are consistently negative across specifications, and the magnitudes are largely comparable. The magnitudes are also economically significant substantial: focusing on column (3) which reports the estimates from our preferred specification which focuses on within occupation–firm variation, we see that a one standard deviation increase in the task–level AI exposure reduces the intensity at which firms demand skills related to that task by around 4.8 percent over the next five years. Given that the dependent variable corresponds to a share, our estimates imply that labor effort reallocates away from skills that are relatively more exposed to AI within an occupation towards skills that are less exposed. We conclude that our task-level exposure is primarily identifying the substitutive effect of AI on labor tasks.

Appendix Table A.8 lists the five most exposed tasks for selected occupations and reports the subsequent changes in demand given by equation (2).[13] Examining the table, we see that in the majority of case we observe large declines in the demand for highly exposed tasks, consistent with AI being a substitute for human labor in performing them. Paralleling an example mentioned in the introduction, preparing purchase orders and expense reports is a highly exposed task for Bookkeeping, Accounting, and Auditing Clerks. In our sample period, we observe a 34 percent average decline over five years in skills related to this task. Many of the other tasks, which relate to various forms of data analysis also experience double-digit declines. Relatedly, we observe very sharp declines in mentions of skills which relate to investigating claims and reviewing documents for Claims Adjusters, Appraisers, Examiners, and Investigators. As a final example, consistent with the large number of business intelligence-related AI applications, we see that large declines in mentions of skills related to tasks involving planning and tracking operational decisions for Industrial Production Managers.

2 Measuring the Exposure of Jobs to AI

So far, we have constructed a measure of task-level exposure to AI. The fact that tasks that are semantically similar to AI applications are likely to experience a reduction in labor demand suggests that our task-level exposure measure is primarily measuring substitution between AI and labor. However, workers can perform several tasks, and these tasks may be complements. Hence, the fact that a worker's certain tasks can now be performed by AI does not necessarily imply that that worker's occupation will experience a decline in labor demand. To map how these task-level exposures translate into shifts in occupational demand, we develop a model based on Acemoglu and Autor (2011); Acemoglu and Restrepo (2018); Caunedo, Jaume, and Keller (2023); Kogan et al.

[13] Specifically, we isolate the five most AI-exposed tasks for several white collar occupations, then report a measure of the dependent variable from equation (5) which is residualized with respect to the controls and the occupation-firm fixed effects then averaged (weighted by employment in each firm-occupation-year cell) across firm-years. As a result of this residualization, this variable averages to zero across all tasks within each firm-occ-year cell. We perform an analogous construction for residualized task-level exposure in order to identify the 5 most exposed tasks for each of these occupations.

(2023). Using the model as a guide, we proceed to develop measures of AI exposure at the job level.

2.1 A Model of AI Exposure

The model features workers of different occupations who each perform different tasks and has the following key features: labor and capital are substitutes in production, and capital is specific to each task. Technological improvements are reflected in declines in the (quality-adjusted) price of capital. Within a job, workers allocate their time among different tasks. Workers optimally choose across jobs taking occupation- and firm-level wages as given along with idiosyncratic preference shocks. Relative to prior work, our model describes the productivity effects of concentration in exposure to technological change within an occupation, and allows for endogenous re-organization of effort to emphasize different tasks in response to labor-substituting technological shocks to exposed tasks.

The model has a simple prediction: improvements in technology that substitute for labor in particular tasks can either increase or decrease labor demand for a specific occupation depending on how these technology improvements affect the capital that is specific in her tasks. In particular, if a technology uniformly improves the capital that is specific to most of the worker's tasks, then labor demand for that occupation decreases, since capital substitutes for associated labor tasks. By contrast, if the technology improves capital in a disparate fashion—some tasks are greatly affected while others are not—then labor demand for an occupation is likely to increase. The reason is that tasks are complements and workers can endogenously allocate their time among tasks. Thus, if a certain improvement is significantly better than labor in performing a specific task, the productivity of the other tasks will increase and workers will re-allocate more time to them.

Setup

There is a continuum of firms that produce intermediate goods Y_f. Aggregate output $\bar{Y}$ as a CES composite of the output Y_f of different firms,

$$\bar{Y} = \left(\int_{\mathcal{F}} Y_f^{\frac{\theta-1}{\theta}} df \right)^{\frac{\theta}{\theta-1}}. \tag{6}$$

Here, θ captures the elasticity of substitution across firms. Each firm produces a differentiated good by combining the output of many occupations,

$$Y_f = \left(\int_{O} Y(o,f)^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}. \tag{7}$$

Firms make profits because of imperfect competition, reflecting both monopolistic competition in product markets and monopsonistic power in labor markets. We denote the firm's markup over marginal cost by $\Theta = \frac{\theta}{\theta-1}1$. Due to the presence of monopsony power, the firm's marginal cost will exceed its average cost and the firm will mark down the wage it pays below the marginal cost of labor.

Workers in occupation o employed in firm f produce output $Y(o, f)$ by combining the output of J individual tasks. The total output of occupation o in firm f is given by

$$Y(o, f) = \left(\sum_j y(j)^{\frac{\psi - 1}{\psi}} \right)^{\frac{\psi}{\psi - 1}} \tag{8}$$

Here, ψ denotes the elasticity of substitution across tasks within a given job, which determines the elasticity of labor demand for each task. To simplify the notation, we will suppress the firm subscript and occupation subscripts unless needed.

Each task j in job (o, f) is produced by a labor input $l(j)$ and a capital input $k(j)$,

$$y(j) = \left(\gamma_j \, l(j)^{\frac{\nu - 1}{\nu}} + (1 - \gamma_j) \, k(j)^{\frac{\nu - 1}{\nu}} \right)^{\frac{\nu}{\nu - 1}}. \tag{9}$$

In the context of our application, we should think of $k(j)$ as intangible capital (e.g. software algorithms) that can substitute for labor in a specific task. Here, ν gives the elasticity of substitution between capital $k(j)$ and labor $l(j)$, while ψ denotes the elasticity of substitution across tasks within an occupation. In what follows, we will be assuming that $\nu > \psi$, which will imply that improvements in the technology that is specific to task j are likely to be labor-saving.

We model the impact of technological innovation as a reduction in $q(j)$, the quality-adjusted price of intangible capital $k(j)$ that is specific to task j,

$$\Delta \log q(j) = -\varepsilon(j). \tag{10}$$

A specific technology is potentially applicable to several tasks within a job. A given technological improvement that is applicable to job (o, f) can therefore be represented as a firm- and occupation-specific vector $\varepsilon \equiv [\varepsilon_1 \dots \varepsilon_J]$ of weakly positive random variables. If $\varepsilon(j) > 0$, that implies that the firm is adopting an improved (or cheaper) labor-saving technology that is specific to task j.

Workers in job (o, f) optimally choose the amount of time they allocate in each task. The effective supply of labor by worker i in task j is given by

$$\ell(j) = h(j)^{1-\beta}. \tag{11}$$

Here, the parameter $\beta \in (0, 1)$ captures the degree of decreasing returns to effort at the task level. If $\beta \to 0$ then there are no decreasing returns at the task level, whereas if $\beta \to 1$ then the quantity of effort is essentially fixed in each task. More generally, a smaller value of β implies that there is more scope of reallocating effort across tasks.

The total number of hours a worker can supply across all J tasks is equal to one. Given the above, the optimal time a worker in job (o, f) will allocate to task j is equal to

$$h(j) = \frac{w(j)^{\frac{1}{\beta}}}{\sum_{j \in J} w(j)^{\frac{1}{\beta}}}, \tag{12}$$

where $w(j)$ is the firm- and occupation-specific wage in task j.

There is a continuum of measure one of ex-ante identical workers who chose a firm–occupation pair based on the total earnings of that occupation and an idiosyncratic taste shock. As in Eaton and Kortum (2002), each worker draws a set of job-specific taste shocks that are independent and identically distributed according to a Fréchet distribution with scale parameter 1 and a shape parameter $1 + \zeta$. Using the properties of the Fréchet distribution, the measure of workers that choose job (o, f) is given by

$$N(o, f) = \bar{\zeta}\, W(o, f)^{\zeta}, \tag{13}$$

where $\bar{\zeta}$ is a constant defined in the Appendix and

$$W(o, f) \equiv \sum_{j \in J_o} h(j)^{1-\beta} w(j) \tag{14}$$

refers to the worker's total earnings in job (o, f), which are a function of her allocation of time and the (job-specific) task prices $w(j)$.

Model Implications

To derive the model's implication for measurement, we can examine the impact of a given technology ε on equilibrium earnings and employment. Denote by

$$\eta_o \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j)} \quad \text{and} \quad \eta_c \equiv \frac{\partial \log w(j)}{\partial \varepsilon(j')} \tag{15}$$

the own- and cross-elasticity of task-specific prices $w(j)$ to improvements $\varepsilon(j)$ and $\varepsilon(j')$ to along the symmetric equilibrium with $\gamma_j = \gamma$, $w(j) = w$, and $q(j) = q$ for all $j \in J$. We solve for these elasticities by a log-linear approximation around the symmetric equilibrium—see Appendix A.1 for details.

The first elasticity η_o captures the impact of the technology improvement that is specific to task j on the price of task j. As long as the elasticity of substitution between capital and labor ν is sufficiently high, the elasticity η_o is negative: improvements in the technology that substitutes for labor in task j leads to lower labor demand for that task and therefore a lower task price $w(j)$. In the limit where the technology becomes a perfect substitute for labor ($\nu \to \infty$ then the elasticity η_o converges to minus one as wages fall one to one with technology improvements. More broadly, a sufficient condition for the elasticity to be negative is that the number of tasks is sufficiently large and $\nu > \psi$.

The second elasticity η_c captures cross-task spillovers. In general, improvements in technology that are specific to task j can have spillovers on tasks that are not directly affected. The strength, and sign, of these spillovers is in general ambiguous as they depend on the elasticity of substitution across tasks ψ and the rate of decreasing returns to effort at the task level β. For most parameter values, the cross-elasticity η_c is negative—unless tasks are sufficient complements, that is ψ is sufficiently low. In addition, as long as the elasticity of substitution between capital and labor ν is sufficiently high, it is increasing as the ease of reallocation increases (β falls).

Given the above, we can express the resulting change Δ_ε of equilibrium employment in job (o, f) in response to technology ε using a second order approximation around the symmetric equilibrium,

$$\Delta_\varepsilon \log N(o, f) \approx \zeta \Big(\eta_o + \eta_c\,(J-1)\Big) m(\varepsilon) + \frac{\zeta}{2\,\beta}\,(\eta_o + \eta_c)^2\; C(\varepsilon) \tag{16}$$

where $m(\varepsilon)$ denotes the mean improvement of the technology across all tasks,

$$m(\varepsilon) \equiv \frac{1}{J} \sum_{j \in J} \varepsilon(j), \tag{17}$$

and $C(\varepsilon)$ denotes the concentration in exposure across tasks,

$$C(\varepsilon) \equiv \frac{1}{J} \sum_{j \in J} \Big(\varepsilon(j) - m(\varepsilon)\Big)^2. \tag{18}$$

The growth in employment is directly related to the growth in wages through the labor supply—equation (13).

Examining equation (16), we see that it has two terms that are related both to the mean level of technology improvement and its concentration, respectively. The first term combines two effects: the direct effect of the average improvement of the technology on worker earnings, which is a function of the two elasticities η_o and η_c. The latter elasticity is multiplied by the number of tasks $J-1$ because η_c is decreasing in J, while the product converges to a positive constant as $J \to \infty$.

The second term in equation (16) is increasing in the dispersion of the degree of technological improvements $\varepsilon(j)$ across tasks. This term captures the extent to which the occupation's average task-level exposure $m(\varepsilon)$ is *concentrated* in a small number of tasks. Labor demand is increasing in this term since technologies that greatly improve the (intangible) capital used only in some tasks but not others allow workers to reallocate their effort to unaffected tasks. Note that this term is over and beyond the standard Jensen's inequality effect that arises because we are approximating the growth in log wages: crucially it is a function of β, the coefficient capturing the decreasing returns to scale at the task level. Holding the elasticities η_c and η_o constant, a decrease in β (easier reallocation across tasks) implies a larger impact of the concentration term on worker earnings. In general, as long as the elasticity of substitution ν is sufficiently high, both elasticities decrease with β implying that the overall term rises as β falls.

The above discussion holds productivity improvements constant. We can also examine the impact of a change in firm productivity on labor demand for occupation o. In our model, we can express the productivity of firm f as

$$Z_f \equiv \left(\int_O X(o,f)^{\chi-1} \right)^{\frac{1}{\chi-1}}, \tag{19}$$

where $X(o,f)$ is the productivity of occupation o employed in firm f. If a given technology affects the productivity of multiple occupations in firm f at the same time, then this is an additional force that will affect wage earnings that is not captured by the two elasticities η_o and η_c above. After log-linearizing around the symmetric equilibrium, we can calculate the elasticity of task prices $w(j)$ to an increase in firm productivity,

$$\eta_z \equiv \frac{\partial \log w(j)}{\partial \log Z_f} = \frac{\theta - \chi}{s_k\, \nu + s_l\, \chi + \zeta} \tag{20}$$

As long as the elasticity of substitution across firm products is greater than the elasticity of substitution across occupations within a given firm—a reasonable assumption—then an increase in firm productivity leads to an increase in the wage of each task, which directly translates into a corresponding increase in employment

$$\Delta_z \log N(o,f) \approx \zeta\, \eta_z\, \Delta_z \log Z_f \tag{21}$$

Lastly, the model has a direct implication for how task-level hours respond. In particular, hours growth is approximately proportional to the task-specific cost shock minus the occupational average cost shock:

$$\Delta \log h(j) \approx \frac{\eta_o - \eta_c}{\beta} \left(\varepsilon(j) - m(\varepsilon) \right) \tag{22}$$

Putting everything together, the effect of technology improvements on firm labor demand can be summarized by equations (16) and (21). The first equation derives the impact of a given technology ε on labor demand holding firm productivity constant, whereas the second equation illustrates the effect of changes in firm productivity while holding everything else constant. Furthermore, equation (22) shows that at the task-level, the amount of effort allocated to a particular task within an occupation should be affected by how exposed the task is, relative to the average exposure across all tasks.

2.2 Measuring the Exposure of Workers to AI

Our model in the previous section implies that the impact of a specific technology on firm labor demand for a specific occupation can be summarized by the two statistics in equations (17) and (18): the average improvement in the labor-saving technology across the tasks performed by workers in that job, and the concentration in these improvements across worker tasks. We next construct empirical analogues of these objects in the data.

Occupation-Level Exposure to AI: Mean vs. Concentration

We next construct direct empirical analogues of our two exposure measures (17) and (18). We first we compute the weighted average exposure probability across an occupation's tasks,

$$\mu_{o,f,t} = \sum_j \omega_{o,j} \text{ Exposure Probability}_{j,f,t} \tag{23}$$

The weights $\omega_{o,j}$ are obtained from the O*NET task importance scores. O*NET assigns a score between 1 and 5 on each task performed by a specific occupation to indicate how central the task is as part of the job. We take these weights $\omega_{o,j}$ from these importance scores and rescale them so that they sum to one within each occupation.

However, this mean exposure measure does not take into account differences across firms in the intensity of AI adoption. To take into account that some firms may utilize AI applications more intensively than others, we adjust these measures with an estimate of the AI intensity in firm f,

$$\text{AI Exposure Average}_{o,f,t} = \mu_{o,f,t} \times \log(1 + N_{f,t}) \tag{24}$$

where $N_{f,t}$ is the number of AI applications in use at firm f during year t.

Similarly, we construct an analogous version for our concentration of AI exposure across tasks as,

$$c_{o,f,t} = \sum_j \omega_{o,j} \left(\text{Exposure Probability}_{j,f,t} - \mu_{o,f,t}\right)^2, \tag{25}$$

and after incorporating differences in adoption across firms we obtain

$$\text{AI Exposure Concentration}_{o,f,t} = c_{o,f,t} \times \log(1 + N_{f,t}). \tag{26}$$

Equations (24) and (26) are the empirical equivalents of equations (17) and (18) in the model. Importantly, they vary across firms not only because different firms develop different AI applications, but also because different firms adopt AI more intensively than others—a proxy for the size of the ε shocks in the model.[14]

[14] We multiply by the log of one plus, rather than incorporating $N_{f,t}$ directly into the definition of (1), because the

Our model implies that an increase in (24) will lead to a decrease in employment *within* the firm, while holding (24) constant, an increase in (26) will lead to higher employment within the firm. The first prediction stems from the direct task-level substitution of labor with AI, while the latter comes from cross-task productivity spillovers generated by the cost improvements that come from AI applications targeted towards the occupation's tasks.

2.3 Stylized facts about AI job exposure

Figure 4 plots the mean AI exposure for each application cluster, averaged across broad occupation groups. AI applications that fall under the 'Business Intelligence Insights' category tend to affect most workers, but especially those in the 'Management', 'Business and Financial', and 'Computer and Math' occupation categories. AI applications differ in their breadth of exposure across occupations: applications in the 'Business Intelligence Insights' affect a broad range of occupations, whereas 'Cybersecurity and fraud Detection' applications are primarily related to Protective Services occupations, which otherwise have low exposures to AI applications. Likewise, AI applications related to customer experience automation as well as marketing and ad optimization are highly related to tasks performed by sales occupations. Echoing our examples from Table 2, task and workflow optimization applications are highly relevant for Production as well as Office and Administrative occupations, while Business and Financial occupations are most exposed to financial risk modeling AI applications.

This discussion suggests that white collar occupations have greater exposure to AI. Indeed, this appears to be the case: in Figure 5, we rank occupation–firm pairs in the Revelio data based on their average salary as imputed by Revelio. We then plot the average of $\mu_{o,f,t}$ across employment-weighted salary percentile ranks. Consistent with Webb (2020), we find that AI exposure rises with wage rank up to approximately the 90th percentile, after which it declines. This pattern contrasts with previous technological shifts, which primarily affected middle-skill labor (Autor et al., 2006; Kogan et al., 2023). Based on this result, one might expect that higher paid-occupations are more likely to be displaced by AI; however this discussion ignores the fact that exposures for these occupations can be highly concentrated.

In Appendix Figure A.5, we compute the average value of our two exposure measures (24) and (26) across occupations, weighted by employment. Examining the figure, several facts stand out. First, there is considerable dispersion in the degree mean AI task exposure across these clusters. AI Applications that fall under the 'Business Intelligence Insights', 'Resource & Performance Optimization', 'AI Solution Consulting', and 'Task & Workflow Automation' cluster are the ones

distribution of the number $N_{f,t}$ of AI applications is highly skewed within firms—some firms have AI applications in the several thousands while others have very few. Our specification implicitly takes a stand on the mapping between the number of AI applications and the decline ε in the quality-adjusted price of labor-saving technologies. See Appendix Table A.1 which summarizes the distribution of these variables in our sample.

that are, on average, most closely related to worker tasks. By contrast, the mean exposure of AI applications in the 'Image & Video Recognition', 'Healthcare Diagnostics & Genomic', 'Scientific & Industrial Modeling', or 'Autonomous Navigation & Robotics'. Thus, all else equal, we expect the first group of AI applications to be more displacive of workers than the second group. Second, the mean level of AI exposure is positively correlated with the concentration measure, but the correlation is not perfect. For instance, we see that applications in the 'Operational Data Analytics' category tend to affect worker tasks in a more concentrated way than applications in the 'Customer Experience Automation' category, even though on average, worker tasks are similarly exposed to both sets of applications. Thus, our model would imply that AI applications in the second category will lead to a greater decline in occupational demand than applications in the first category.

Naturally, these comparisons obscure significant heterogeneity in exposures across occupations, as Figure 4 suggests. Importantly, occupations tend to differ both in the mean exposure, but also in the concentration of exposure to each of these AI applications. For example, Management Analysts and Financial Risk Specialists are two occupations that have similar mean exposure to AI applications that fall under the 'Business Intelligence Insights' category. However, the exposure of Management Analysts is significantly more concentrated in specific tasks than the exposure of Financial Risk Specialists, which implies that Management Analysts are less likely to be displaced than Financial Risk Specialists from these AI applications. Along these lines, Compensation, Benefits, and Job Analysis Specialists have a significantly lower mean exposure to 'Business Intelligence Insights' applications, but their exposure is highly concentrated in specific tasks, which suggests that they have greater scope of reallocating their effort towards less affected tasks.

Similarly, the fact that AI applications under the 'Financial Risk Modeling' category are mostly related to worker tasks in Business and Financial occupations masks considerable differences: the tasks performed by Financial Managers have both higher mean, but also less concentrated, exposure to these AI applications than the tasks of Brokerage Clerks. Thus, ceteris paribus, we would expect that AI applications under this category to decrease the occupational demand for Financial Managers relative to Brokerage Clerks, as the former face both a higher, but also broader task exposure to these AI applications. Along these lines, Sales Managers and Telemarketers both face similar mean exposure to AI applications under the 'Customer Experience Automation' category, but Sales Managers face a significantly more concentrated exposure and therefore should experience a smaller decrease in labor demand than Telemarketers.

3 The Impact of AI Exposure on Labor Demand

We begin by detailing our identification strategy that allows us to estimate the impact of AI exposure on occupation labor demand. After presenting our estimates and discussing different robustness checks, we then examine the impact of our estimates on labor reallocation over our sample period.

3.1 Empirical Strategy

Our objective is to estimate the impact of AI exposure on firm-level and labor market outcomes. However, because firms endogenously choose whether and how to adopt new technologies, including AI, it is difficult to disentangle the effect of AI exposure from other correlated drivers of firm performance. For example, larger, more productive, and more profitable firms may be more likely to adopt AI. Since such firms tend to follow lower growth trajectories (Evans, 1987), this raises concerns about selection bias: firms with characteristics that make them more likely to adopt AI may grow more slowly for reasons unrelated to adoption itself. Additionally, unobserved heterogeneity may upward bias estimates if firms with better management practices or more advanced technological infrastructure are both more inclined to adopt AI and better positioned to realize its gains.

Taken together, these issues highlight the need for an instrumental variables approach to isolate exogenous variation in AI adoption to establish a causal relationship with firm outcomes. To address these issues, we implement a Bartik shift-share instrument to predict AI exposure at the firm-level.[15] Our identification strategy exploits cross-firm variation in historical university hiring patterns and cross-university variation in the rise of AI-related work to construct a shift-share instrument for firm-level AI exposure. Specifically, we interact university-level measures of AI intensity—defined as the share of graduates who enter AI-related occupations in the 2014–2018 period—with firm-specific hiring shares from a pre-specified baseline period (2005–2009), before the widespread diffusion of modern AI tools. The shifts are measured as the share of a university's graduates who enter AI-related occupations during the 2014–2018 period, based on resume data covering a wide range of institutions and graduating cohorts. The shares capture each firm's historical tendency to hire from different universities and are held fixed throughout the analysis period to avoid simultaneity bias with firm-level hiring or technology adoption. By multiplying the AI shift for each university with the corresponding firm-specific hiring share and aggregating across all feeder institutions, we obtain a firm-level measure of predicted AI exposure. This shift-share construct serves as an instrument for actual AI employment at the firm level, allowing us to estimate the causal effects of AI adoption on firm outcomes.

One of the key rationales for the relevance of our instrument is that the scarcity of available workers with the relevant skills and expertise to develop and deploy AI can often be a barrier to its adoption. Consistent with this conjecture, the BTOS survey asks firms directly about reasons why firms plan not to use AI. On an employment-weighted basis, while three quarters of firms say that AI is not applicable to their business, Bonney et al. (2024) report that 3.6 percent of firms

[15] A Bartik instrument can be understood as a weighted average of a common set of shocks (shifts) with weights that reflect heterogeneity in firms' exposure to those shocks (shares).

specifically identify the lack of a skilled workforce as a barrier preventing adoption. In addition, 4.6 percent of firms indicate that AI is too expensive and 8.8 percent of firms identify the lack of knowledge on the capabilities of AI as a significant barrier, two additional factors that are likely impacted the ease with which a firm can find workers with the relevant expertise. Analyzing data from the 2018 ABS survey, Chequer et al. (2025) report that, among the firms who are testing AI, 36 percent of firms (weighted by employment) report that the lack of available talent adversely affected the adoption of AI, a magnitude comparable to the fraction of firms saying that lack of available data was a barrier to adoption.

3.2 Instrumental Variables

Concerns about identification differ when measuring AI exposure within firms at the occupation level, rather than at the firm level overall. First, selection bias may arise if firms selectively implement AI in occupations or tasks where the expected cost savings are highest; for example, where firms anticipate increased labor scarcity or diminished labor productivity. This leads to a non-random pattern of AI adoption across occupations. Such targeted implementation would tend to bias estimates of the effect of average AI exposure in a predictable way: it would induce a downward bias in the OLS coefficients on both the mean and the concentration of AI exposure. After instrumenting, we would expect the IV coefficient on average AI exposure to become more negative, and the coefficient on concentration to become more positive, relative to their OLS estimates.

Second, measurement error may attenuate estimated effects. Since occupational-level AI exposure within firms is often based on limited data on actual implementation, there is potential for attenuation bias, especially given the highly granular nature of our measurement and the relatively few observed AI use cases within each firm. To address these issues, our IV strategy exploits variation in the mean and concentration of occupational exposure to AI across *all* firms to predict exposure within a given firm. By drawing on external variation in AI exposure across many firms and applications, this approach helps mitigate both measurement error and selection bias from targeted occupational adoption.

The key variation in our IV strategy comes from labor market exposure to AI through persistent university hiring networks. Firms that historically hired from universities whose graduates later entered AI-related occupations are more likely to adopt AI themselves. This reflects an external supply-side channel of AI diffusion: firms become more exposed to AI adoption not through endogenous internal capabilities, but through their established access to AI-related labor via university pipelines. By using pre-AI hiring patterns and post-AI university-level intensity, we isolate plausibly exogenous variation in firm-level AI adoption.

We predict the number of AI workers at the firm as follows:

$$N_{f,t}^{AI,pred} = N_{f,t}^{total} \times p_{f,t}^{AI} \tag{27}$$

Here $N_{f,t}^{total}$ is the total number of resumes in active job positions at firm f at time t, while $p_{f,t}^{AI}$ is the predicted probability of a given worker implementing AI at the firm. Our estimate of $p_{f,t}^{AI}$ is

$$p_{f,t}^{AI} = \sum_{u \in \mathcal{U}} w_{u \to f}^{2005-2009} \times \frac{N_{u,t}^{AI}}{N_{u,t}^{total}} \tag{28}$$

Here $N_{u,t}^{AI}$ is the actual number of AI workers in year t who report graduating form university u on their resumes, and $N_{u,t}^{total}$ is the number of total workers in year t who report graduating from university u; $\mathcal{U}$ is the set of all universities represented by workers with active job positions in the resume data from 2005-2009. We further denote by $w_{u \to f}^{2005-2009}$ the average share of all workers employed at firm f from 2005-2009 and who graduated from university u. We compute $w_{u \to f}^{2005-2009}$ using job positions from Revelio resumes over the 2005-2009 time period.

Figure 6 illustrates the core idea behind our identification strategy by comparing the top ten universities in the historical hiring networks of BNY Mellon and State Street.[16] Both firms are among the largest global custodians and operate in similar segments of the financial services industry, yet their university hiring patterns differ sharply. BNY Mellon, headquartered in Pennsylvania, primarily hires from nearby institutions in Pennsylvania, New York, and New Jersey. In contrast, State Street, based in Massachusetts, draws heavily from universities in Massachusetts and other New England states. These patterns are based on resume data from 2005–2009 and reflect persistent pre-AI labor supply linkages across firms.

Our empirical strategy leverages these differences to construct a shift-share instrument for firm-level AI exposure. Specifically, we interact each firm's historical hiring shares across universities (the shares) with university-level AI intensity in the 2014–2018 period (the shifts), defined as the share of graduates from each university who enter AI-related occupations. If the universities supplying BNY Mellon and State Street differ in their graduates' tendency to adopt AI in the later period (2014-2018), then the two firms will have systematically different predicted exposure to AI—even though they are observationally similar in size, industry, and product market. This variation in predicted AI exposure, arising from stable differences in hiring networks and time-varying AI adoption at the university level, provides the key source of identifying variation in our instrumental variables design.

[16] This is based off their average university employment shares from 2005-2009.

We define our instrument for the number of AI applications, $\log(1+N_{f,t})$, as:

$$\log(1+N_{f,t})^{IV} \equiv \log(1+N_{f,t}^{AI,pred}) \tag{29}$$

which relies on the plausible assumption that the number of AI use cases within a firm increases with the predicted number of workers capable of implementing AI. This predicted AI employment measure captures exogenous variation in AI adoption driven by pre-existing labor market networks, rather than by endogenous firm-level decisions about technology use.

Next, we construct the occupation-level instruments. We let N_t denote the number of AI applications extracted across *all* firms in year t, and we define

$$\text{Exposure Probability}_{j,t} = \frac{1}{N_t}\sum_{i=1}^{N_{f,t}} \text{Above 95th percentile indicator}_{j,i} \tag{30}$$

and

$$\mu_{o,t} = \sum_j \omega_{o,j}\text{Exposure Probability}_{j,t} \tag{31}$$

$$c_{o,t} = \sum_j \omega_{o,j}\left(\text{Exposure Probability}_{j,t} - \mu_{o,t}\right)^2 \tag{32}$$

which are respectively the task-level exposure probabilities, and the means and concentrations of the exposure probabilities across all firms in year t. Finally, we arrive at our instruments for the mean and concentration of occupational task exposure to AI:

$$\text{AI Exposure Average}\ {}^{IV}_{o,f,t} \equiv \mu_{o,t} \times \log(1+\text{Predicted AI Employees})_{f,t} \tag{33}$$

$$\text{AI Exposure Concentration}^{IV}_{o,f,t} \equiv c_{o,t} \times \log(1+\text{Predicted AI Employees})_{f,t} \tag{34}$$

The logic behind these instruments is that the average and concentration of AI exposure across all firms within an occupation instead reflects broader technological diffusion of AI towards certain tasks, rather than any selected targeting of particular occupations for purely firm-specific reasons. Additionally, by computing the means and concentrations across a broad set of firms, the estimation noise in the mean and concentration decrease considerably, which helps to alleviate attenuation concerns. As we demonstrate in Section 3.4 below, this attenuation appears especially impactful for the coefficient estimates on the AI exposure concentration.

Our instrumental variables strategy rests on three key identification assumptions. First, we assume that the university-level propensity of graduates to work in AI is exogenous to the firm—that is, firms do not directly influence whether graduates from a particular university enter AI-related occupations. Second, we assume that a firm's historical hiring shares across universities, measured during the pre-AI period (2005–2009), are uncorrelated with its subsequent growth (2014–2018), except through their impact on AI adoption. Third, we assume that university–firm hiring relationships are sufficiently stable over time, such that a firm's historical hiring patterns are predictive of its post-period hiring behavior.

In Panel A of Appendix Table A.4, we show that university firm hiring relationships are highly persistent: the average share of a firm f's employment coming from university u in 2005-2009 strongly predicts the corresponding share in the 2014-2018 period, even after including firm and university fixed effects. The estimated coefficient is approximately 0.5, with a t-statistic of 27. In Panel B, we demonstrate that the predicted share of AI workers based on these historical university links, $p^{AI}_{f,t}$, significantly predicts the actual share of AI workers at the firm, with a coefficient of 0.58 and t-statistic of 7.5 (conditional on controls for log resume employment and industry $\times$ year fixed effects). Together, these findings support both the exclusion and relevance assumptions, and confirm that our instrument, $\log(1+N_{f,t})^{IV}$, is strongly correlated with actual AI utilization. Accordingly, we show in Section 3.3 that the full instrument $\log(1+N_{f,t})^{IV}$ has a high F-stat for predicting actual AI utilization. Further robustness checks are reported in Section 3.5.

3.3 AI Exposure and Firm Growth

We begin by estimating the effect of current AI use on subsequent firm outcomes. Specifically, we estimate regressions of the form:

$$\log Y_{f,t+5} - \log Y_{f,t} = \beta \log(1+N_{f,t}) + \alpha_{Ind,t} + \delta X_{f,t} + \epsilon_{f,t}, \tag{35}$$

where the dependent variable is the five-year log growth in a firm outcome $Y_{f,t}$, and the key independent variable is $\log(1+N_{f,t})$, the log of one plus the number of AI-related job titles at firm f in year t.

Our baseline controls include log Compustat employment and the log number of resumes from Revelio at the firm-year level. These controls adjust for the mechanical relationship between AI use and firm size, as both larger firms and firms with more complete resume coverage are likely to report more AI roles.[17] Finally, we include industry-by-year fixed effects ($\alpha_{Ind,t}$) to absorb industry-specific time trends, and we cluster standard errors at the firm level.

We present estimates of equation (35) in Table 4. Across all specifications, we find a strong positive relationship between the number of AI applications at firm f in year t and subsequent five-year growth in sales, employment, profits, and total factor productivity. Columns 1–4 report OLS estimates, while columns 5–8 present IV estimates, as defined in equation (29)

[17] Our model implies that a firm's productivity index Z_f increases in the average occupational productivity $X(o,f)$, or equivalently, the inverse cost of occupational output. Because AI use affects exposure across occupations, it raises Z_f. We therefore use the uninteracted firm-level measure $\log(1+N_{f,t})$ to proxy for $\Delta \log Z_f$. To avoid confounding due to firm size or resume coverage, we always control for the log number of Revelio employees at the firm in year t.

First, the coefficient estimates are strongly positive and statistically significant across both OLS and IV specifications. Focusing on the IV results, a one–standard deviation increase in our AI adoption measure is associated with a 9.6 percentage point increase in firm revenue growth and a 7.4 percentage point increase in total factor productivity over the subsequent five years. These estimates support our interpretation of the number of AI applications as a shifter of the firm productivity index Z_f in the model: by improving cost efficiency, AI use enables firms to expand and become more productive. Notably, profits increase roughly in line with sales, suggesting that markups remain relatively stable.

We note that our IV estimates are somewhat larger in magnitude than our OLS estimates. This pattern would arise if firms which tend to adopt AI also are the types of firms which tend to grow more slowly, as it would tend to bias the OLS estimates towards zero. This possibility is consistent with survey evidence (Acemoglu et al., 2023b) that shows that advanced technology adopters tend to be larger and older to begin with, meaning they are likely to naturally be in the "low–growth" phase of their life-cycle. Our instrumental variables strategy therefore tends to mitigate this downward bias. Also of note are the large F-statistics on the instrument in each of the IV specifications, alleviating any potential concerns related to weak instruments.

Notably, we find a positive effect of AI adoption on employment: AI adoption leads to a 6.7 percentage point increase in firm employment. Here, we emphasize that the fact that firm employment increases should not be interpreted as evidence that AI is directly complementary to worker tasks. Through the lens of the model, the impact of firm productivity on labor demand can be large when the elasticity of substitution between firms (θ in the model) is much higher than the elasticity of substitution between occupations within firms (χ in the model), generating positive net labor demand effects even when the task-level AI–labor substitution ν is large. To infer the degree of labor substitution, we need to examine more granular evidence at the occupation–firm level, which is the focus of the next section.

3.4 AI Exposure and Occupational Labor Demand

Having established that AI has a causal impact on firm productivity and growth, we next turn to the main part of our analysis that focuses on labor demand for specific jobs. To do so, we leverage the granularity of our exposure measures in (24) and (26) that establish which occupations are more or less exposed to AI applications within a given firm.

Our main outcome variable is employment of a particular occupation in a specific firm at a point in time. We measure employment at the 6-digit SOC occupation level within a firm using the Revelio resume data, which allows us to keep track of how many positions were active at the given firm and occupation each year. We use the weights provided by Revelio to reweigh observations to account for the skew in the data towards certain occupations. We consider a position active in a given year if it was in place for at least half of the year. Our model focuses on the labor-substituting effects of AI at the task level. Therefore, when measuring employment and the occupation–firm level, we exclude occupations from the 2-digit SOC code 15 (Computer and Mathematical occupations) because these are by far the most likely broad occupation group to be tagged as AI implementers, representing about 50% of all AI positions in our data. Moreover, even if not directly running the AI models, workers in these occupations are also those who would implement and maintain the software and hardware necessary to use AI at the firm.

Our first specification exploits both within- and between-firm variation. Specifically, we estimate

$$\log(Emp_{f,o,t+5}) - \log(Emp_{f,o,t}) = \beta \,\text{AI Exposure Average}_{f,o,t} + \delta \text{AI Exposure Concentration}_{f,o,t} + \gamma \log(1 + N_{f,t}) + \Gamma X_{f,o,t} + \alpha_t \epsilon_{f,t} \tag{36}$$

In this specification, we include only calendar year fixed effects and directly control for firm-level AI utilization $\log(1+N)_{f,t}$, which allows us to examine the direct effects of shifting firm productivity Z_f in the model. In this specification, the vector of controls $X_{f,o,t}$ includes the log of firm employment (based on Revelio and Compustat); the growth in employment over the last year for that particular occupation in that specific firm; and the market value of innovation to book assets from Kogan, Papanikolaou, Seru, and Stoffman (2017). We cluster standard errors by occupation–firm, as this is the level at which our main measures vary, and we weight observations by the share of yearly total employment in the occupation–firm cell. We scale the independent variables to have unit standard deviation, and we multiply the dependent variable by 100 so that coefficients read as the percent increase in employment growth for a standard deviation change in exposure.

Our second specification focuses on within-firm patterns, by differencing out shocks to the firm,

$$\log(Emp_{f,o,t+5}) - \log(Emp_{f,o,t}) = \beta \,\text{AI Exposure Average}_{f,o,t} + \delta \text{AI Exposure Concentration}_{f,o,t} + \Gamma X_{f,o,t} + \alpha_{f,t} + (\alpha_{o,t}) + \epsilon_{f,t} \tag{37}$$

Specifically, in equation (37) we include the interaction of firm–year fixed effects $\alpha_{f,t}$ to the specification in (36). The cost of doing so is that we can no longer identify the coefficient γ capturing the impact of firm productivity improvements due to AI on labor demand. The advantage, however, is that by leveraging cross-occupation and within-firm variation, this specification nets out all endogeneity driven by general firm-level selection into AI adoption, and isolates the relative demand effects stemming from differences in occupational exposure within the firm. As a variant of this specification, we also include granular occupation (6-digit SOC) interacted with calendar year fixed effects $\alpha_{o,t}$. This variant further absorbs occupation–specific trends by exploiting variation in the degree to which a particular occupation is more exposed to AI within a given firm compared to the typical exposure of that occupation to AI among all firms. This specification is feasible due to the granular nature of our measure compared to existing measures that identify only cross-occupation differences in exposure (Webb, 2020; Eisfeldt et al., 2023; Eloundou et al., 2023; Brynjolfsson et al., 2018).

We report our estimates in Table 5. In columns 1 through 4 we report the estimated coefficients from (36) and (37) using OLS. In columns 5 through 8 we estimate the same specifications through two-stage least squares using the instrumental variables described in Section 3.2.

Confirming the predictions of the model, the estimated coefficients on $\widehat{\beta}$ on our mean AI exposure measure (24) is negatively and highly significant, while the estimated coefficient $\widehat{\delta}$ on the concentration measure (24) is also positive and highly significant across specifications. When we do not include firm–year fixed effects in columns 1-2 and 5-6, we can separately estimate the impact of firm-level AI induced productivity gains, which also has highly significantly positive relationship with employment growth. Coefficient estimates are reasonably stable when comparing within OLS or within IV specifications, though they are unsurprisingly the smallest when we include occupation–year fixed effects in the most strict specifications in columns 4 and 8. Instrumented specifications again have F-statistics that are far beyond typical weak instruments thresholds, which diminishes any such concerns.

In terms of magnitudes, there are several points worth noting. The range of coefficient estimates on the average AI exposure double from between -5.5 and -8 percent for the OLS specifications, to -14 to -15 percent for the IV specifications. This pattern is even more stark for the concentration of AI exposure, which increases by about 5-6 times when comparing IV to OLS specifications, going from a significantly positive but small 1.2 to 1.9 percent in OLS specifications, versus a large and highly significant positive 7.3 to 8.8 percent across IV specifications. If the primary source of endogeneity were selective adoption within firms towards occupations and tasks where labor scarcity is expected to increase or labor productivity to decline, we would expect to see a positive bias in both OLS coefficients: the coefficient on our mean AI exposure measure (24) would become less negative and go towards zero after being instrumented, while the coefficient on the concentration measure (24) would become more positive. Instead, both coefficient estimates move away from zero and in opposite directions after being instrumented, which is more consistent with measurement noise causing attenuation in our estimates (especially for our exposure concentration measure). While we cannot rule out that there is some targeted AI adoption within the firm, it doesn't seem to be the primary source of bias in the OLS.

The negative effect stemming from average AI exposure is directly in line with the model's predictions under the assumption of labor–AI substitution at the task level. The positive coefficient on the concentration measure (24) suggests that there are indeed productivity spillovers across tasks within an occupation which mitigate some of this negative direct substitution effect. The model allows for either force to dominate the other depending on the parameter mix, most especially the various elasticities of substitution. In the data we find that this direct substitution effect appears to be quantitatively larger. On the other hand, the firm-level productivity effect of AI use raises employment in general. Thus whether AI use on net erodes or increases relative employment for particular groups is an empirical question that depends on the empirical distribution of each measure, as well as the relative coefficient estimates.

Last, we also use our instrument to revisit our findings on the role of AI exposure on the demand for particular tasks—equation (5). As we see in Appendix Table A.3, the magnitudes of the estimated slope coefficients are very similar between the OLS and IV specifications.

3.5 Robustness: Validity of the IV

Our identification strategy rests on the assumption that firm-level exposure to cross-university variation in AI intensity—measured from resume-based AI graduates—is exogenous to firm-specific unobservables that may also affect outcomes. Our strategy exploits a common set of university-level shocks (the "shifts") and interacts them with firm-specific historical hiring shares (the "shares"), measured in the pre-period. Because identification comes from cross-firm differences in how these common shocks are mediated by historical hiring patterns, our key assumption is that the firm's hiring network structure—measured prior to the rise of any widespread AI use—is uncorrelated with unobserved shocks to firm outcomes. In this section, we discuss potential threats to this assumption and describe empirical strategies to assess and address them.

This instrument relies on the assumption that a firm's historical hiring shares from 2005–2009 affect outcomes during the 2014–2023 period only through differential exposure to AI-trained labor. In Appendix Table A.5, we implement a series of empirical tests designed to validate this identifying assumption.

First, one concern is that firms may have endogenously formed their university hiring networks in anticipation of future AI needs. Firms with early AI ambitions might have disproportionately hired from institutions that were already beginning to ramp up AI-related training during the pre-period, violating the assumption that pre-period hiring shares are orthogonal to future firm-specific shocks. To address this, Panel A of Appendix Table A.5 presents results from an exclusion strategy that removes from the sample the top 50 universities with the highest number of graduates working in AI during the 2014–2018 period. These include Stanford, UC Berkeley, MIT, Carnegie Mellon, Georgia Tech, the University of Washington, most Ivy League schools (including Harvard), and other major producers of AI talent, together accounting for roughly half of all AI worker-years in the data. In addition, we exclude the top 50 firms that collectively employed the most AI workers over the same period, including all members of the "Magnificent Seven" (Amazon, Apple, Google, Meta, Microsoft, Nvidia, Tesla), as well as Boeing, Lockheed Martin, Raytheon, PayPal, Twitter (pre-privatization), eBay, Netflix, Dell, Intel, and HP. We further drop all firms in "high-tech" industries that were most likely to be closely tied to early AI diffusion—defined as 2-digit NAICS 51 (Information), 54 (Professional Services), and 3-digit NAICS 334 (Computer Hardware Manufacturing).

Second, a related threat is that endogenous firm association with elite universities may mechanically drive both the "shift" and "share" components of the instrument. For example, prestigious institutions like MIT, Stanford, and Harvard are highly represented in both AI talent production and high-status hiring pipelines. If firms that hire disproportionately from such schools are simply more capable to begin with, then our instrument may capture selection on firm quality rather than exogenous AI exposure. Since the top AI-producing universities tend to also be highly-prestigious in other ways, exclusion of the top 50 AI-producing universities as described above speaks to this concern. Again, we find that the IV estimates persist despite dropping these universities, which supports the view that results are not being driven by firm quality or prestige effects.

Third, another possible concern is reverse causality: that large or technologically forward-looking firms may have influenced the AI curricula or graduation output of the universities from which they hire. In this case, the shift may be endogenous to firm behavior, especially if those firms represent a significant share of AI demand. Although our identification strategy does not rely on the exogeneity of the shift, such feedback could still bias the estimated relationship if it contaminates the pre-period hiring shares. By excluding the top 50 AI firms and top 50 AI-producing universities from the shift-share calculation, we minimize the likelihood that firms shaped the curriculum of their feeder institutions in ways that could violate the exclusion restriction. The continued strength and significance of the IV after these removals indicates that reverse causality of this sort is unlikely to drive the results.

Fourth, we consider the possibility that firms located in geographic technology hubs may both hire from nearby AI-intensive universities and benefit from regional agglomeration effects, such as knowledge spillovers or local infrastructure. In this scenario, the instrument could reflect geographic clustering rather than differential AI exposure. While we do not directly estimate spatial spillovers, the exclusion of universities and firms in well-known tech regions—e.g., Stanford (Bay Area), MIT (Boston), and firms in NAICS 51 and 54—removes many of the institutions most susceptible to this channel.

We report the findings of dropping firms and universities in Panel A of Appendix Table A.5, which simultaneously addresses the concerns discussed above. Despite these exclusions, the IV estimates remain statistically significant and only slightly attenuated relative to our main specification, suggesting that our identification is not driven by forward-looking firms selectively forming AI-linked hiring networks, or endogenous firm-university associations based on prestige. Since our identification assumption rests on the exogeneity of the pre-period shares, the stability of predictions even after removing some of the most important universities for predicting AI use is particularly helpful for the identifying assumption (Goldsmith-Pinkham, Sorkin, and Swift, 2020).

Fifth, another concern is that firms hiring from AI-intensive universities may benefit from exposure to a broader supply of technically skilled graduates, beyond those working in AI specifically, thereby confounding the estimated effect of AI exposure on firm growth. To address this, Panel B of Appendix Table A.5 augments the IV model with additional shift-share controls that predict the firm's exposure to technical labor more broadly. Specifically, we include predicted hiring shares for computer and mathematical occupations (2-digit SOC 15) and engineering occupations (SOC 17), based on the same university hiring network. These controls account for broader exposure to technical talent unrelated to AI. Their inclusion has no meaningful effect on the IV coefficients, suggesting that the main instrument captures AI-specific exposure rather than general technical intensity.

We repeat similar robustness checks for the occupation-firm IV specifications in Appendix Table A.6. In the first four columns, we again drop the top 50 universities when constructing the IV, and we drop the top 50 AI-using firms plus "high-tech" industries from the sample. In the last four columns of the table, we add controls for predicted shares of workers in computer/mathematical occupations and engineering occupations (both by themselves and interacted with $\mu_{o,t}$ and $\sigma^2_{o,t}$ defined in equations (31) and (32), respectively). The specifications are otherwise the exact same as in the last four columns of Table 5. Again, we find similar conclusions as with our baseline estimates, with mildly attenuated (but still significant) effects when we drop the top AI-using firms from the sample in the first four columns.

3.6 Aggregate Effects

Assessing the aggregate impact of AI technology developments on labor demand directly from our estimates is challenging due to the inclusion of calendar year fixed effects in our empirical specifications. These fixed effects control for contemporaneous macroeconomic conditions and policy events, such as the economic recovery from the 2008–2009 financial crisis and the enactment of the Tax Cuts and Jobs Act of 2017, and therefore allow us to isolate the impact of AI on labor demand from other confounding influences. However, our estimates remain useful for evaluating AI's role in reallocating employment across occupations.

To quantify this reallocative effect, we calculate the expected net marginal effect of AI exposure on labor demand conditional on occupational characteristics K, such as average occupation wage percentile or broad occupational categories,

$$E\left[\, \widehat{\beta}\text{AI Exposure Average}^{\perp} + \widehat{\delta}\text{AI Exposure Concentration}^{\perp} + \widehat{\gamma}\log(1 + N_{f,t})^{\perp} \mid K = k\right] \tag{38}$$

Here, the estimated coefficients $\widehat{\beta}$, $\widehat{\delta}$, and $\widehat{\gamma}$ are taken from the IV estimates in column (5) of Table 5. The notation $\perp$ denotes that variables have been orthogonalized with respect to all non-AI controls, and K is the characteristic of interest (salary rank or occupation broad category). Because the average of each variable has been netted out via controls, these marginal effects on employment are relative to the average employment growth, so the total effects integrated across the entire employment share–weighted distribution of K integrate to zero by construction.

We begin by examining how AI's impacts differ across occupations by earnings. To do this, we calculate expected marginal effects at each salary percentile. Specifically, we set k to represent the salary percentile of an occupation and compute equation (38) separately for each percentile. The top panel of Figure 7 illustrates these calculations. The red line shows the direct marginal effect of mean AI exposure on employment (the first term in (38)), the green line represents reallocative effects (the second term due to concentrated exposure), and the yellow line captures the impact of firm productivity (the third term). The total net effect, combining all three, is shown by the blue line.

Consider first the direct substitution effect—the red line in Figure 7. The negative impact of mean AI exposure increases steadily with salary. This aligns with Figure 5, where higher-paid occupations have greater AI exposure. Quantitatively, occupations at the highest salary percentiles are expected to see an employment share reduction of about 7% over five years due to this direct effect. Conversely, lower-paid occupations might increase by around 9%.

However, these direct effects miss two key offsetting factors: within-job reallocation of effort and productivity effects. High-paying occupations are typically at firms heavily utilizing AI (yellow line) and also experience greater concentration in AI exposure (green line). These two effects neutralize the direct substitution effect, resulting in a significantly more muted net impact (blue line) across most salary levels. Interestingly, at the highest pay levels, the net employment effect is slightly positive, despite substantial direct AI exposure. Our aggregate results align with Acemoglu et al. (2022), who find negligible aggregate employment impacts from AI during 2010–2018 despite clear effects at individual firms. Our decomposition illustrates why detecting aggregate AI employment effects is challenging: the overall effect can be small, not because the impact of AI is small but rather because it results in several offsetting forces.

This decomposition is based on estimates from equation (36), using both within- and between-firm variations in AI exposure. Results remain robust when focusing exclusively on within-firm variation, as we see in panel B of Figure 7. The within-firm analysis reveals a stronger direct substitution effect on high-paying positions. At about the 90th salary percentile, within-firm employment falls by around 2.6% due to direct substitution, moderated to 1.4% after accounting for the fact that AI exposure for these occupations is concentrated in specific tasks. Meanwhile, lower-paid positions see modest employment gains around 1.8%. Overall, while strong substitution effects exist, their net employment impacts remain quantitatively small both within firms and in aggregate.

We next explore how AI effects vary across broad occupation categories (defined by 2-digit

SOC codes). Table 6 summarizes each AI-related employment component for these groups. By construction, total employment-weighted effects sum to zero within each component. Examining Table 6, several occupational groups experience significant employment declines due to AI. "Business and Financial" occupations face an average 1.7% decline, and “Architecture and Engineering” occupations lose 1.5%. "Food Preparation and Serving" occupations provide another instructive example, experiencing a net 4.5% employment reduction primarily because their employers rarely utilize AI. Conversely, legal occupations see a robust 7.1% employment increase due to low direct AI exposure and substantial firm-level productivity benefits. Yet, recent developments in generative AI, such as large language models, may soon impact even legal occupations more directly (Eloundou et al., 2023). Overall, AI’s aggregate impacts remain moderate, though with greater variation across occupational groups than salary levels.

However, these aggregate group effects obscure significant within-group variation. To capture this, Figure 8 examines AI impacts at the detailed (6-digit SOC) occupational level. Panels A and B separately plot task-level and firm-level AI effects against realized employment shifts, while Panel C plots their combined effect. Clearly, occupation-specific task exposure and firm AI utilization are both critical factors, strongly correlated with actual employment shifts. Specifically, AI-related effects significantly explain actual employment share changes, as indicated by a regression R^2 of approximately 16 percent, with roughly 28 percent attributed directly to occupation-specific task exposure.

Finally, recognizing our data primarily reflects professional roles, we reweight occupations using Bureau of Labor Statistics Occupational Employment Statistics (BLS-OES) data to mirror broader U.S. employment. These re-weighted analyses (Appendix Figure A.6 and Appendix Table A.7) yield substantively similar conclusions, albeit with slightly more positive net reallocation toward high-paying roles and more pronounced declines for business and engineering occupations.

In sum, our analysis confirms AI significantly drives occupational employment shifts. However, given our focus on large, publicly traded firms that utilize AI more intensively, these estimates may somewhat overstate AI’s broader economic impact.

4 Conclusion

This paper examines the impact of artificial intelligence adoption on firm dynamics and labor demand, leveraging firm-occupation level variation in AI exposure. We document three primary empirical patterns. First, AI adoption is concentrated in larger, more productive firms, which tend to have distinct growth trajectories. Instrumental variable estimates confirm that AI adoption leads to higher firm-level sales, profits, and total factor productivity. Second, at the occupational level, AI exposure is concentrated in higher-wage positions, with employment effects that depend on the concentration of exposure across tasks. Higher average exposure reduces within-firm employment, while greater concentration in exposure mitigates these declines by reallocating labor toward complementary tasks. Third, firm-wide AI adoption generates positive employment effects, consistent with AI-driven productivity gains increasing aggregate labor demand.

The results suggest that while AI substitutes for labor at the task level, its net employment effects are shaped by offsetting forces. Highly AI-exposed occupations experience declines in labor demand, yet within-occupation task reallocation and firm-wide AI-driven growth help sustain overall employment levels. Consistent with our conceptual framework, we find that occupations with high concentration in AI exposure experience relatively higher employment growth, underscoring the importance of within-occupation task complementarities. Across firms, AI adoption is associated with higher employment growth in AI-intensive firms, indicating that firms integrating AI more effectively also expand their workforce.

Despite these countervailing forces, the labor market effects of AI remain nontrivial. Our estimates suggest that AI exposure accounts for roughly 16 percent of the variation in occupational employment growth among publicly traded firms, with just over half of this effect stemming from task-level substitution and the remainder from firm-wide adoption. However, at the top of the wage distribution, where AI exposure is most pronounced, we find limited net employment effects due to offsetting firm growth and reallocation within occupations. These results highlight why aggregate employment impacts of AI may be difficult to detect, even as task-level substitution is significant.

Taken together, these findings provide new evidence on the labor market implications of AI adoption. Rather than leading to broad-based job losses, AI appears to be reallocating labor across tasks and firms, with the magnitude of displacement effects contingent on the structure of AI adoption at the firm level. Future research should further investigate how firms adjust their workforce composition in response to AI, how skill demands evolve, and how AI-induced productivity gains translate into wage and employment dynamics over time.

References

Acemoglu, D. (2024, May). The simple macroeconomics of ai. Working Paper 32487, National Bureau of Economic Research.

Acemoglu, D., G. Anderson, D. Beede, C. Buffington, E. Childress, E. Dinlersoz, L. Foster, N. Goldschlag, J. Haltiwanger, Z. Kroff, P. Restrepo, and N. Zolas (2023a, May). Advanced technology adoption: Selection or causal effects? *AEA Papers and Proceedings 113*, 210–14.

Acemoglu, D., G. Anderson, D. Beede, C. Buffington, E. Childress, E. Dinlersoz, L. Foster, N. Goldschlag, J. Haltiwanger, Z. Kroff, P. Restrepo, and N. Zolas (2023b, May). Advanced technology adoption: Selection or causal effects? *AEA Papers and Proceedings 113*, 210–14.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pp. 1043–1171.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics 40*(S1), S293–S340.

Acemoglu, D. and P. Restrepo (2018, June). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review 108*(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2021, June). Tasks, automation, and the rise in us wage inequality. Working Paper 28920, National Bureau of Economic Research.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2020). What are the labor and product market effects of automation?: New evidence from france. Technical report, Cato Institute.

Aghion, P., S. Bunel, X. Jaravel, A. R. Timo Mikaelsen, and J. Sogaard (Forthcoming). How different uses of ai shape labor demand: Evidence from France. *American Economic Association: Papers and Proceedings*.

Auer, R., D. Köpfer, and J. Sveda (2024, September). The rise of generative AI: modelling exposure, substitution and inequality effects on the US labour market. BIS Working Papers 1207, Bank for International Settlements.

Aum, S. and Y. Shin (2025, February). The labor market impact of digital technologies. Working Paper 33469, National Bureau of Economic Research.

Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024, 03). New frontiers: The origins and content of new work, 1940–2018*. *The Quarterly Journal of Economics 139*(3), 1399–1465.

Autor, D. H. and D. Dorn (2013, August). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–97.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006, May). The Polarization of the U.S. Labor Market. *American Economic Review 96*(2), 189–194.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics 151*, 103745.

Babina, T., A. Fedyk, A. X. He, and J. Hodson (2023, June). Firm investments in artificial intelligence technologies and changes in workforce composition. Working Paper 31325, National Bureau of Economic Research.

Benmelech, E. and M. Zator (2022, January). Robots and firm investment. Working Paper 29676, National Bureau of Economic Research.

Bonfiglioli, A., R. Crinò, H. Fadinger, and G. Gancia (2020, November). Robot Imports and Firm-Level Outcomes. Crc tr 224 discussion paper series, University of Bonn and University of Mannheim, Germany.

Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, and K. Savage (2024, April). Tracking firm use of ai in real time: A snapshot from the business trends and outlook survey. Working Paper 32319, National Bureau of Economic Research.

Braxton, J. C. and B. Taska (2023, February). Technological change and the consequences of job loss. *American Economic Review 113*(2), 279–316.

Brynjolfsson, E., T. Mitchell, and D. Rock (2018, May). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings 108*, 43–47.

Caunedo, J., D. Jaume, and E. Keller (2023, June). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–85.

Chequer, M., K. Herkenhoff, D. Papanikolaou, L. Schmidt, and B. Seegmiller (2025). A model of ai and labor markets: Theory meets resume data. Technical report, University of Minnesota.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021, 05). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association 19*(6), 3104–3153.

Davis, S. J., J. C. Haltiwanger, and S. Schuh (1996). *Job Creation and Destruction.* MIT Press.

de Souza, G. and H. Li (2023, November). Robots, Tools, and Jobs: Evidence from Brazilian Labor Markets. Working Paper Series WP 2023-42, Federal Reserve Bank of Chicago.

Dechezleprêtre, A., D. Hémous, M. Olsen, and C. Zanella (2021). Induced automation: evidence from firm-level patent data. *University of Zurich, Department of Economics, Working Paper* (384).

Deming, D. and L. B. Kahn (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics 36*(S1), S337–S369.

Deming, D. J. and K. Noray (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Eisfeldt, A., G. Schubert, B. Taska, and M. B. Zhang (2023). Generative ai and firm values. Working paper, U.S. Census Bureau.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2023). Gpts are gpts: An early look at the labor market impact potential of large language models.

Evans, D. S. (1987). Tests of alternative theories of firm growth. *Journal of Political Economy 95*(4), 657–674.

Felten, E. W., M. Raj, and R. Seamans (2018, May). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings 108*, 54–57.

Gathmann, C., F. Grimm, and E. Winkler (2024, December). AI, Task Changes in Jobs, and Worker Reallocation. IZA Discussion Papers 17554, Institute of Labor Economics (IZA).

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *Quarterly Journal of Economics 135*(4), 1793–1868.

Goos, M., A. Manning, and A. Salomons (2014, August). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review 104*(8), 2509–26.

Hershbein, B. and L. B. Kahn (2018, July). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review 108*(7), 1737–72.

Humlum, A. (2019). Robot adoption and labor market dynamics. *Princeton University.*

Humlum, A. and E. Vestergaard (2024, May). The Adoption of ChatGPT. IZA Discussion Papers 16992, Institute of Labor Economics (IZA).

Jiang, W., J. Park, R. J. Xiao, and S. Zhang (2025, February). Ai and the extended workday: Productivity, contracting efficiency, and distribution of rents. Working Paper 33536, National Bureau of Economic Research.

Jiang, W., Y. Tang, R. J. Xiao, and V. Yao (2021). Surviving the fintech disruption. Technical report, National Bureau of Economic Research.

Koch, M., I. Manuylov, and M. Smolka (2021, 01). Robots and Firms. *The Economic Journal 131*(638), 2553–2584.

Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2023, November). Technology and labor displacement: Evidence from linking patents with worker-level data. Working Paper 31846, National Bureau of Economic Research.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017, 03). Technological Innovation, Resource Allocation, and Growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Mann, K. and L. Püttmann (2023, 05). Benign Effects of Automation: New Evidence from Patent Texts. *The Review of Economics and Statistics 105*(3), 562–579.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. S. Foster, and N. Zolas (2023, October). Ai adoption in america: Who, what, and where. Working Paper 31788, National Bureau of Economic Research.

Mikolov, T., I. Sutskever, K. Chen, G. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. *CoRR abs/1310.4546*.

Pairolero, N. A., A. V. Giczy, G. Torres, et al. (2025). The artificial intelligence patent dataset (aipd) 2023 update. *Journal of Technology Transfer*.

Pennington, J., R. Socher, and C. D. Manning (2014). Glove: Global vectors for word representation. In *EMNLP*.

Seegmiller, B., D. Papanikolaou, and L. D. Schmidt (2023). Measuring document similarity with weighted averages of word embeddings. *Explorations in Economic History 87*(C), S0014498322000729.

Tambe, P. (2025, February). Reskilling the workforce for ai: Domain knowledge and algorithmic expertise. Technical report. Available at SSRN: https://ssrn.com/abstract=3776492.

Webb, M. (2020). The impact of artificial intelligence on the labor market.

İmrohoroğlu, A. and Şelale Tüzel (2014). Firm-level productivity, risk, and return. *Management Science 60*(8), 2073–2090.

Figures

Figure 1: Illustration of process for identifying AI applications from resumes and exposed occupation tasks

Example resume job description of a worker employed at JP Morgan:

> Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. AI/ML model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments.

Step 1: Identify AI-related terms (if any)

"Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses. **AI/ML** model delivery in public cloud, private cloud and on prem. managing credit risk deployment services platform with continuous delivery, development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

Step 2: Use large language models to extract the phrases likely to contain specific AI applications

"**Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses.** AI/ML model delivery in public cloud, private cloud and on prem. Managing credit risk deployment services platform with continuous delivery, **development and deployment of quantitative risk models that serve regulatory and credit risk assessments.**"

Step 3: Use large language models to clean the extracted AI applications

Extracted phrase: "Technology delivery lead for risk and fraud forecasting models in auto, card, and home lending businesses."

> Cleaned AI application: **"Forecast risk and fraud in various lending businesses, including auto, card, and home lending."**

Extracted phrase: "Development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

> Cleaned AI application: **"Assess credit risk and provide regulatory compliance across different lines of business."**

Step 4: Use GTE sentence embeddings to measure textual similarity to identify highly exposed tasks

AI application: "Forecast risk and fraud in various lending businesses, including auto, card, and home lending."

> Most exposed O*NET occupation task by cosine similarity: **"Prepare reports that include the degree of risk involved in extending credit or lending money."** (Credit Analysts, SOC code = 132041)

AI application: "Development and deployment of quantitative risk models that serve regulatory and credit risk assessments."

> Most exposed O*NET occupation task by cosine similarity: **"Analyze credit data and financial statements to determine the degree of risk involved in extending credit or lending money."** (Credit Analysts, SOC code = 132041)

Note: This figure shows an example of the process for identifying AI applications from online resumes and linking with exposed job tasks. See section 1.2 in the main text and appendix A.2 for further details.

Figure 2: Sector × size AI utilization rates

Panel A: Comparison with statistics from Census BTOS survey

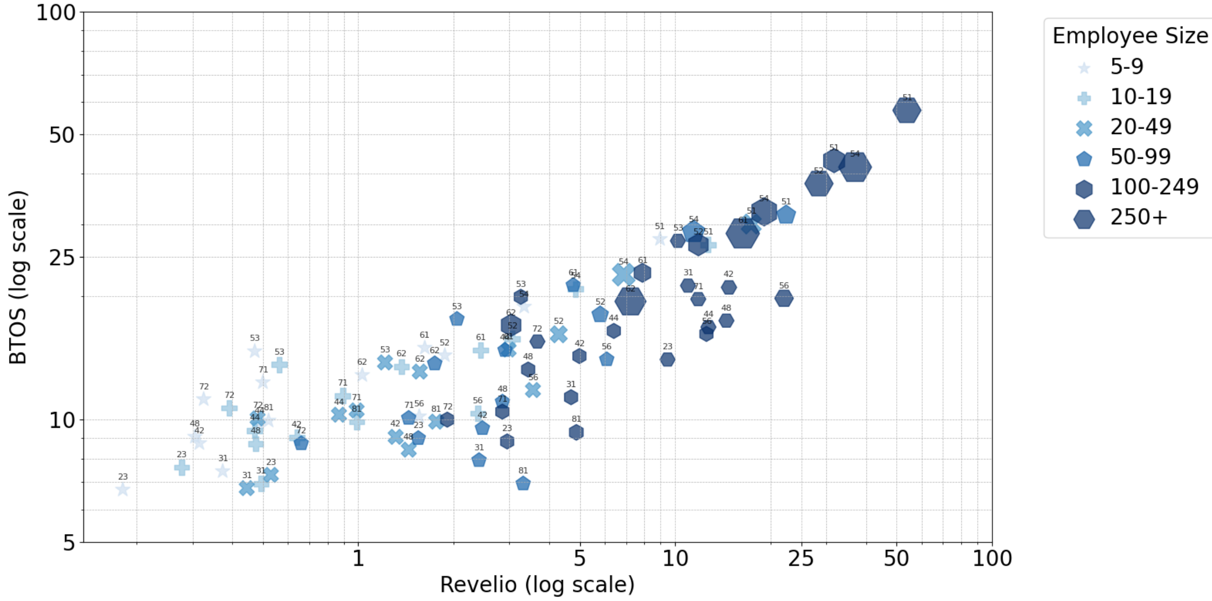

Panel B: Detailed resume-based AI utilization rates

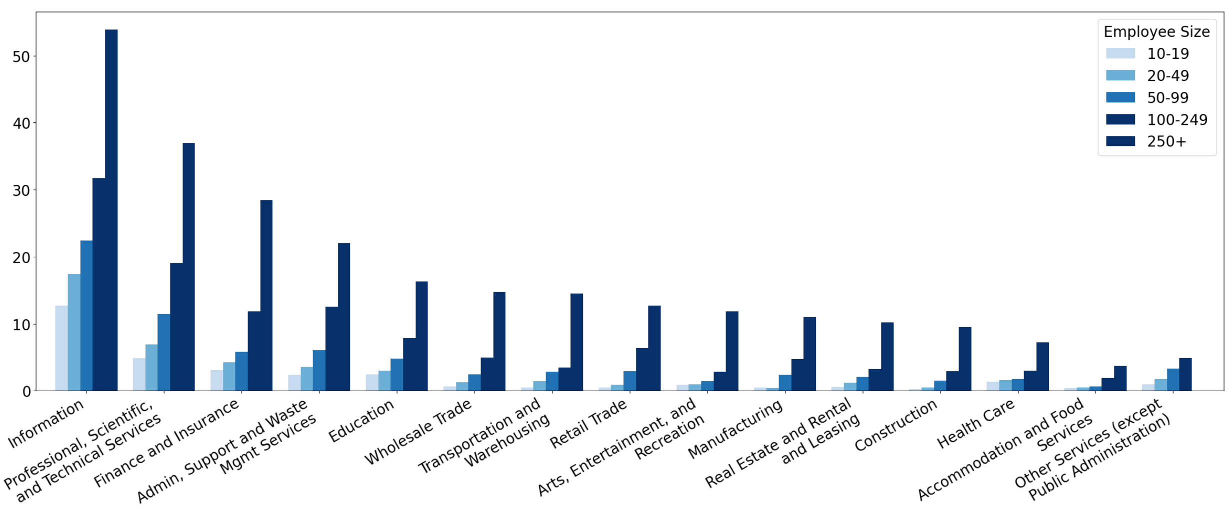

Note: Panel A of this figure plots a scatter plot of AI-implied utilization rates by sector × firm size bins, as implied by the share of firms who report using artificial intelligence technologies in the BTOS survey (y-axis), against the share of firms with at least one AI resume in the Revelio data in the year 2023 (x-axis). We use Revelio NAICS codes and firm identifiers, and we restrict to firms with less than 1,000 employees. The correlation in levels is approximately 0.90. The plot legend indicates firm size bins, and the labels on the points in the graph denote the sector (2-digit NAICS code). Panel B offers a detailed breakdown of the resume-based AI utilization rates from panel A.

Figure 3: Composition of AI application type by sector

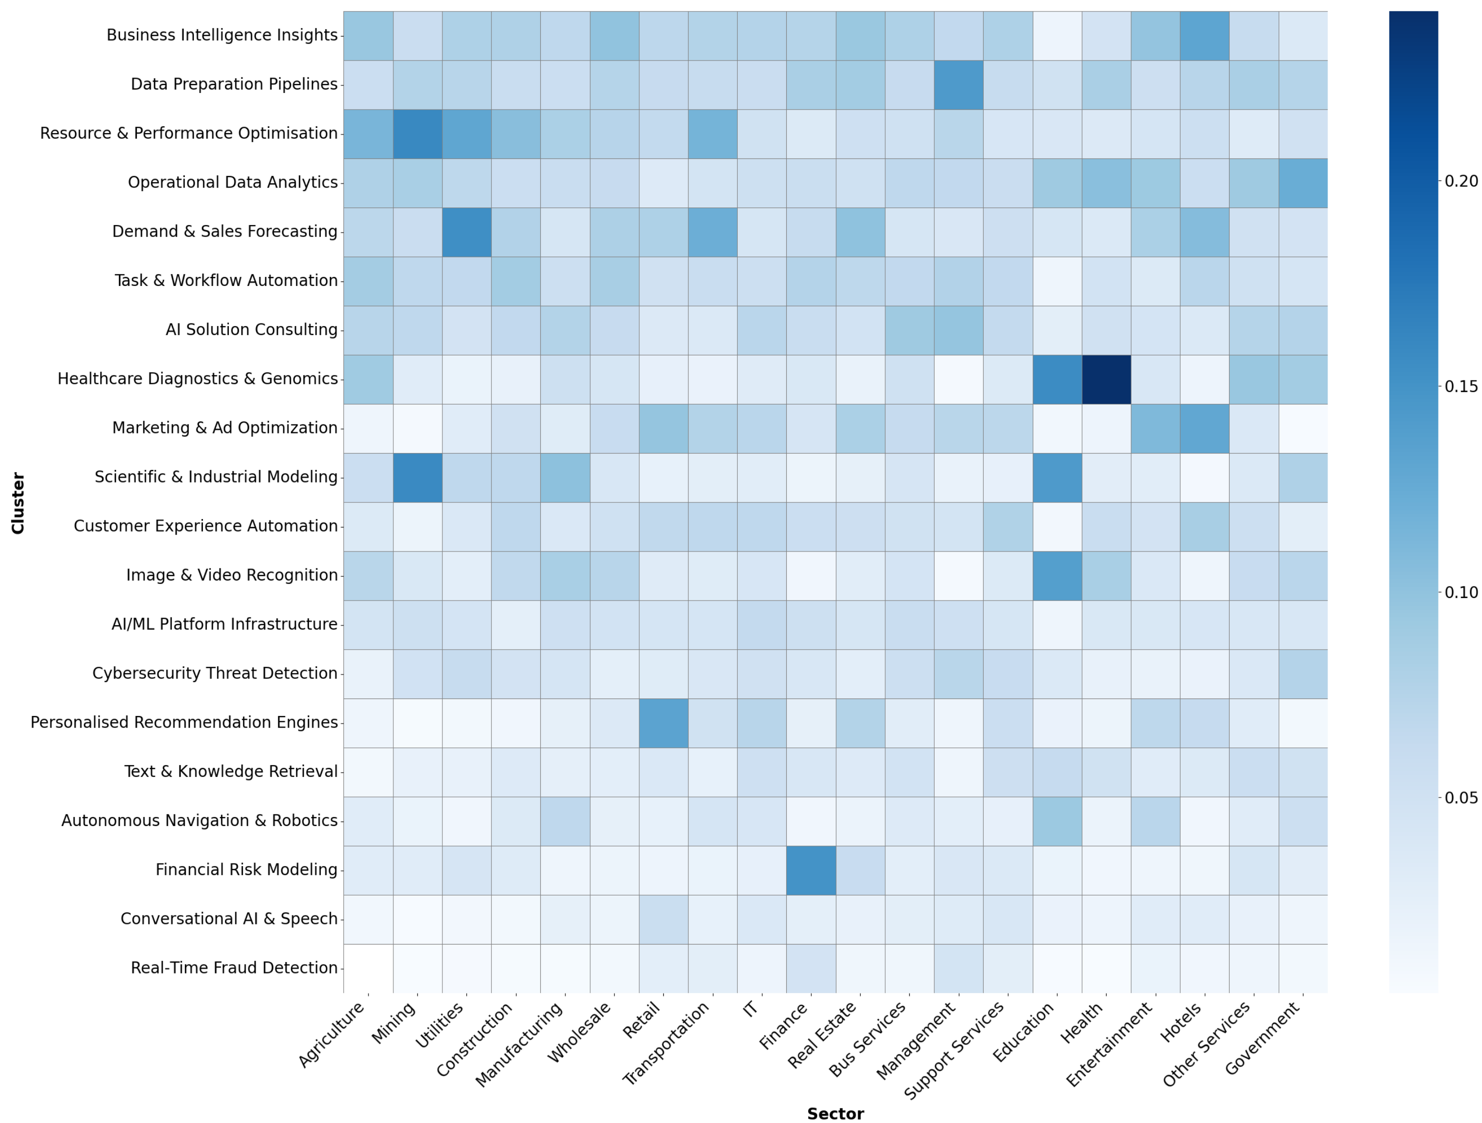

Note: This heatmap characterizes the distribution of categories of AI applications (in rows) within major NAICS industry sectors (in columns). Rows correspond with 20 categories of AI applications obtained by first running k-means clustering algorithm, then using a LLM to label each cluster. The color of cell shading indicates the fraction of AI applications which fall in a given category within each sector. Darker shading indicates with a higher share of applications in that category. Statistics are calculated using the full Revelio sample. Clusters are ordered vertically in descending order by their average frequency.

Figure 4: Occupational AI exposure by AI application type

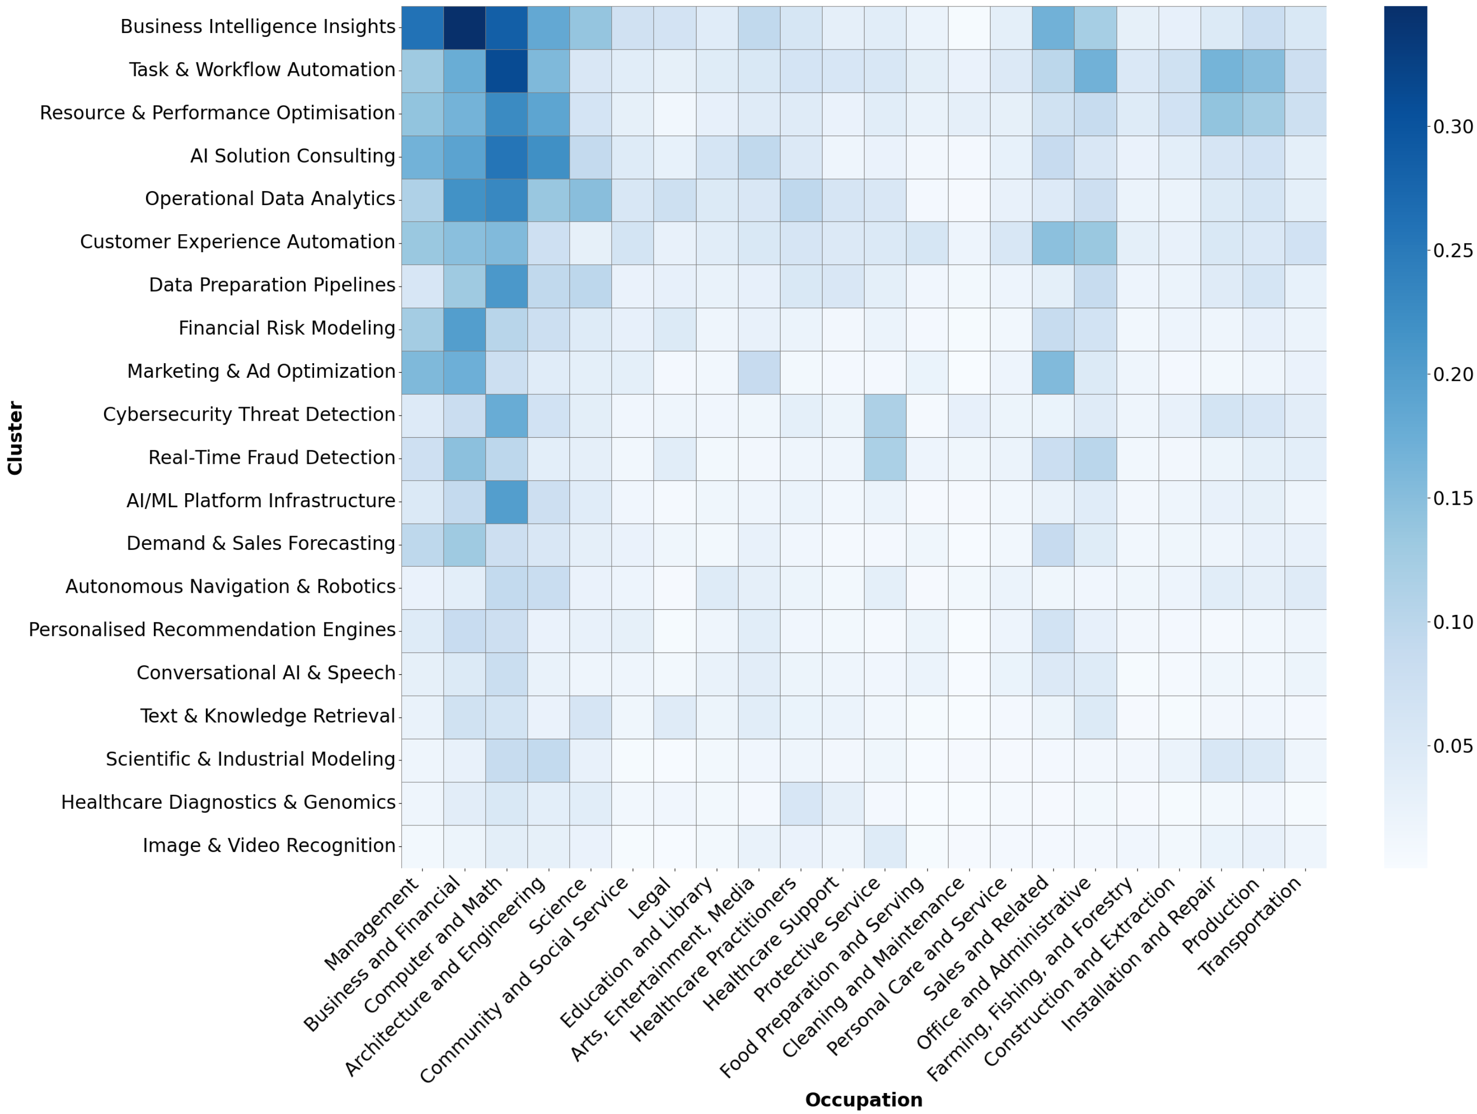

Note: This heatmap characterizes the average fraction of tasks exposed per AI application for different major SOC occupation groups (in columns) in different categories of AI applications (in rows). Rows correspond with 20 categories of AI applications obtained by first running k-means clustering algorithm, then using a LLM to label each cluster. The color of cell shading indicates the average fraction of exposed tasks per AI application in a given category for each group of occupations. Darker shading indicates a higher average exposure probability in that category. Statistics are calculated using the full Revelio sample. Clusters are ordered vertically in descending order by their average exposure probability. Note that the Computer and Math occupations are excluded from our reallocation analyses.

Figure 5: AI Exposure Probability by Salary Rank

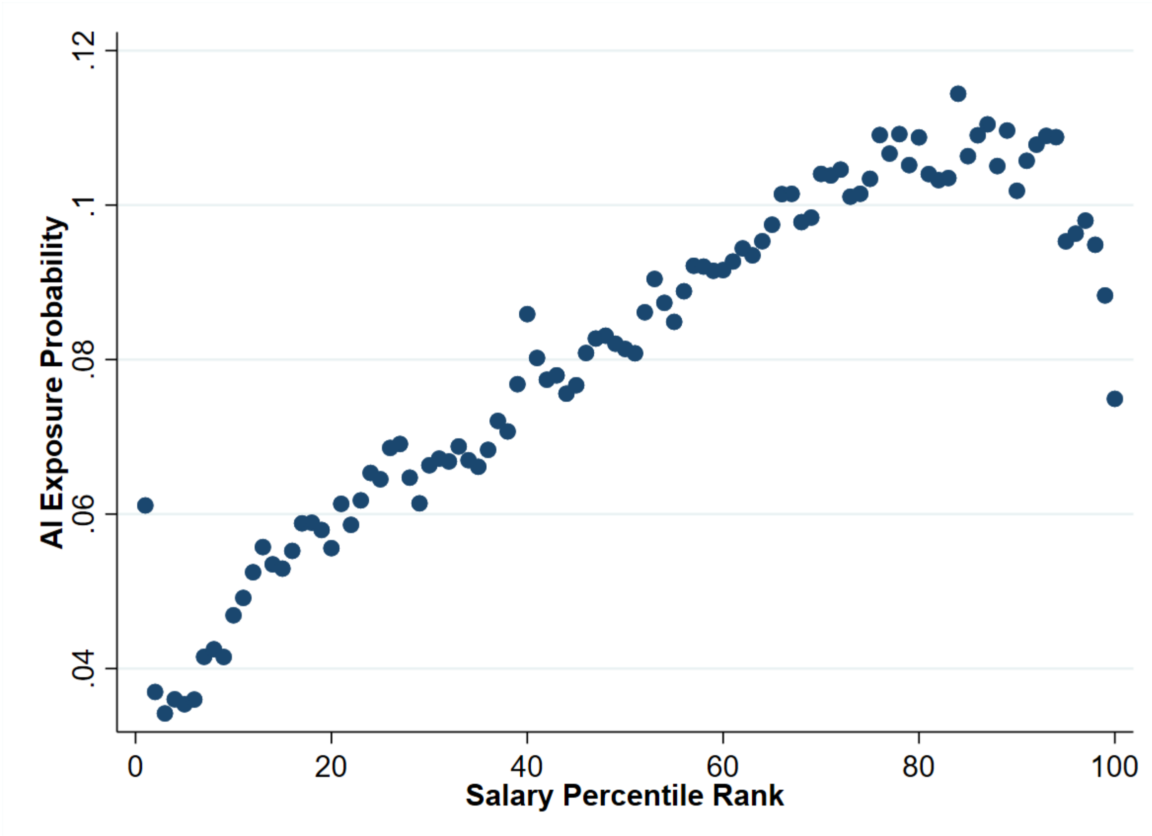

Note: This figure plots the average task-level probability of exposure to a given AI application by salary rank. We use imputed salaries for each job position to compute the ranks. See Section 2.2 for details.

Figure 6: Examples of pre-existing hiring networks

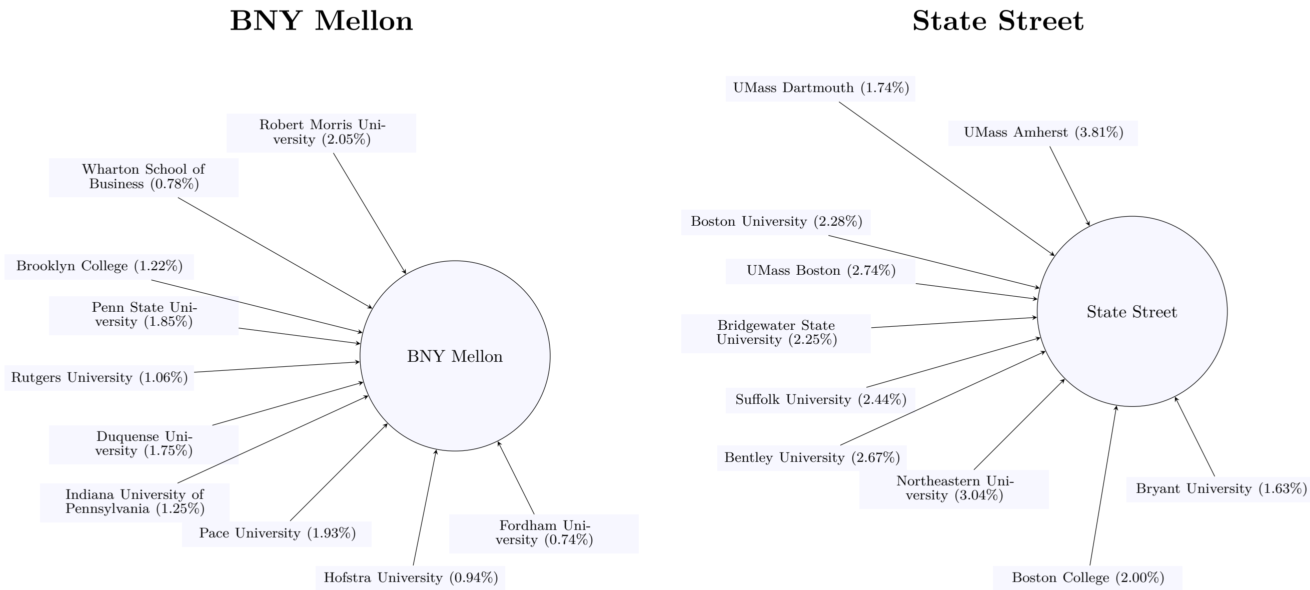

Note: This figure shows the average share of employment coming from the top 10 universities in BNY Mellon's and State Street's university hiring networks. Each percentage in parentheses corresponds to the average share of the given firm's workers during 2005-2009 who graduated from the given university. See section 3.2 for details.

Figure 7: Impact of AI on employment growth across the pay distribution

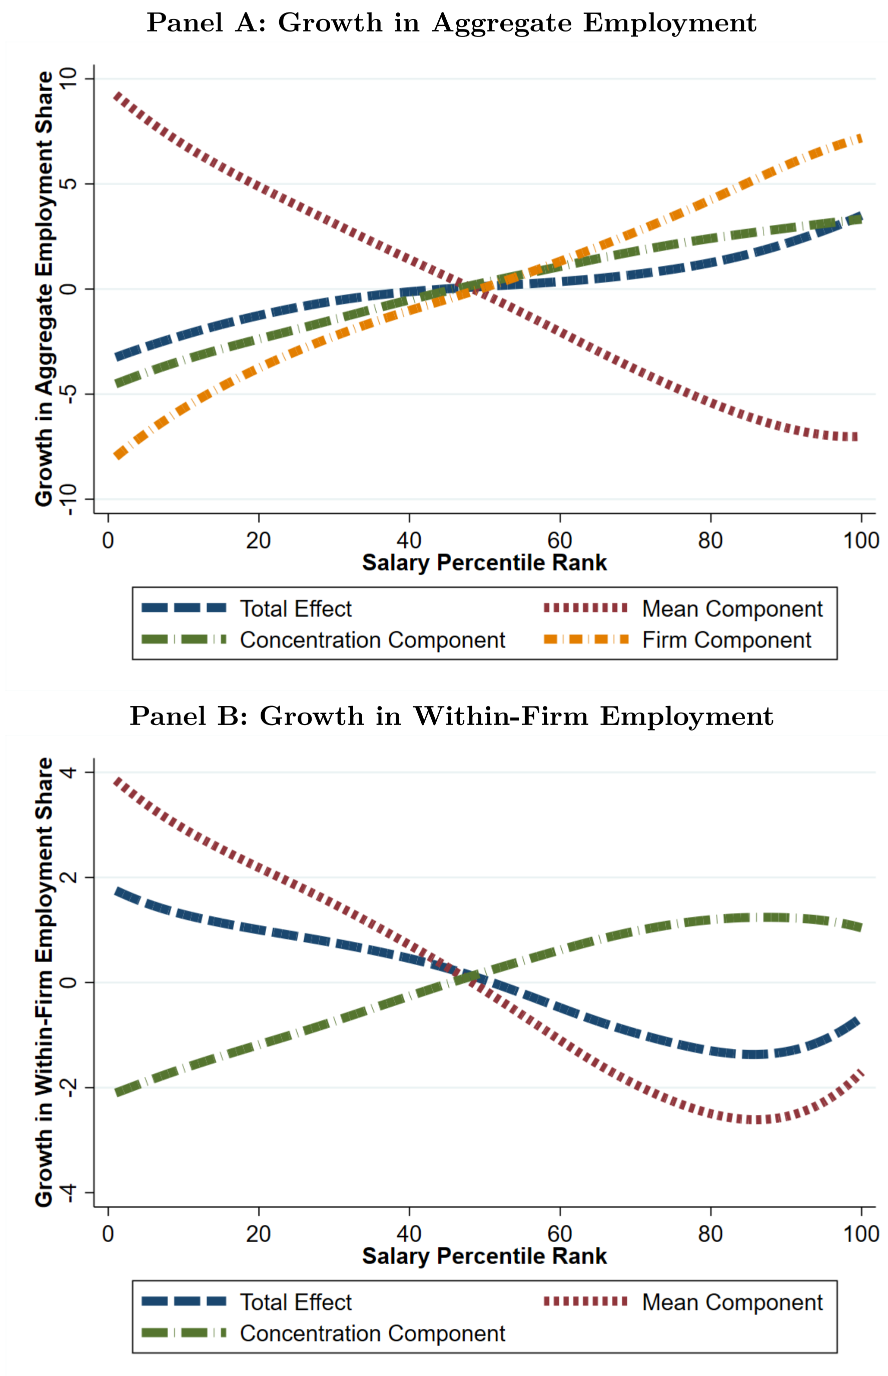

Note: This figure implements the decomposition of employment marginal effects from equation (38) in Section 3.6, where we compute the expected impact of the different components of exposure to artificial intelligence on changes in employment share at each points in the salary percentile distribution. Plots are lowess-smoothed to enhance readability. In Panel A, we plot the impacts on employment shares in the aggregate, while in Panel B, we look purely at within-firm reallocation. In red, we plot the impact of direct task-level substitution driven by our measure of average AI exposure; in green we plot the impact of across-task productivity spillovers driven by the variance of AI exposure within the occupation. In Panel A we also show the effect of firm-level AI use in yellow. The total net effect is in blue. See section 3.4 of the main text for details.

Figure 8: Actual growth in employment shares relative to AI-implied growth

A. Effect from Direct Task Exposure to AI (mean and dispersion) | B. Effect from Employers' AI Use (firm productivity) | C. Overall AI-Implied Change (total effect)

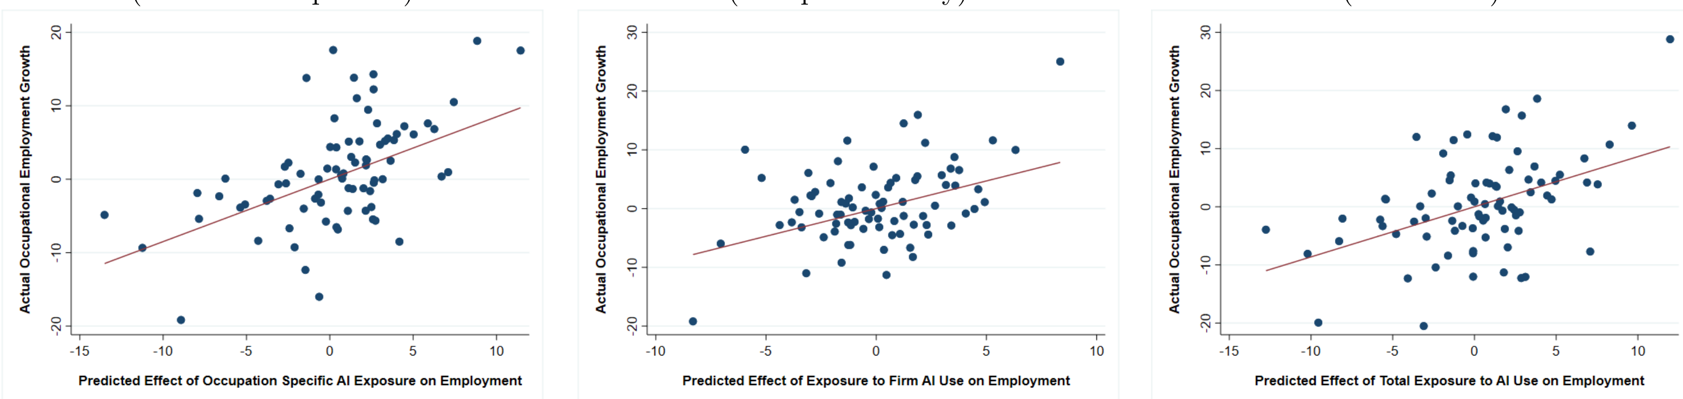

Note: This figure plots residualized binscatter plots of actual occupational employment growth against predicted occupational employment growth implied by our decomposition of employment marginal effects from equation (38) in Section 3.6. We implement the decomposition at the 6-digit SOC occupation level. In Panel A, we plot the relationship between actual occupational employment growth and the total effect of direct occupation task exposure to AI (including mean and variance of task exposure), after netting out the firm-level component. In Panel B, we do the opposite exercise by plotting the partial relationship between actual occupational employment growth and the occupational average exposure to firm-level AI use, after netting out the effects of direct task exposure. The two components taken together can explain 14 percent of realized employment growth, of which 51 percent is attributable to direct occupation exposure. See section 3.4 of the main text for details.

Tables

Table 1: Characteristics of AI-using firms

	(1) Log Sales per worker	(2) Log Sales	(3) Log Profit	(4) Log TFP	(5) Log Average Salary
$\log(1 + \text{AI uses})$	0.117***	0.310***	0.415***	0.125***	0.109***
	(6.87)	(12.39)	(17.48)	(10.37)	(18.81)
N	33541	36227	33309	17034	38211
R-sq	0.345	0.644	0.614	0.181	0.427
Revelio Emp Control	X	X	X	X	X
Ind × Year FE	X	X	X	X	X

Note: This table shows regression coefficients of the logs of sales per worker, total sales, profits (defined as sales minus cost of goods sold), revenue total factor productivity, and log average Revelio salary on firm-level AI utilization $\log(1 + \mathrm{N}_{f,t})$ defined in the main text. As controls, we include the log of total employment based on Revelio resume counts in the given year and 3-digit NAICS industry × year fixed effects. We cluster standard errors by firm and report t-statistics in parenthesis. The sample period spans 2014-2023.

Table 2: AI applications and most exposed tasks at JPMorgan Chase and Walmart

Panel A: Examples from JPMorgan Chase		
Application Cluster	Highly Exposed Tasks	Associated Occupation
Fraud Detection, AML & Risk Mitigation	Collect and analyze data to detect deficient controls, duplicated effort, extravagance, fraud, or non-compliance with laws, regulations, and management policies.	Accountants and Auditors
	Research or evaluate new technologies for use in fraud detection systems.	Other Financial Specialists
Predictive Modeling & Financial Forecasting	Consult financial literature to ensure use of the latest models or statistical techniques.	Other Financial Specialists
	Research or develop analytical tools to address issues such as portfolio construction or optimization, performance measurement, attribution, profit and loss measurement, or pricing models.	Other Financial Specialists
Customer Engagement & Personalization	Monitor customer preferences to determine focus of sales efforts.	Sales Managers
	Identify interested and qualified customers to provide them with additional information.	Models, Demonstrators, and Product Promoters
Panel B: Examples from Walmart		
Application Cluster	Highly Exposed Task	Associated Occupation
Forecasting, Pricing, and Supply Chain Optimization	Analyze market and delivery systems to assess present and future material availability.	Purchasing Managers
	Monitor and analyze sales records, trends, or economic conditions to anticipate consumer buying patterns, company sales, and needed inventory.	Wholesale and Retail Buyers, Except Farm Products
Process Automation and Operational Efficiency	Plan and modify product configurations to meet customer needs.	Sales Engineers
	Monitor and adjust production processes or equipment for quality and productivity.	Other Engineering Technologists & Technicians, Except Drafters
Fraud, Security, and Anomaly Detection	Analyze retail data to identify current or emerging trends in theft or fraud.	Other Managers
	Monitor machines that automatically measure, sort, or inspect products.	Inspectors, Testers, Sorters, Samplers, and Weighers

Note: This table provides some examples of AI applications used by two large firms, JPMorgan Chase and Walmart, along with closely related tasks (middle column) and their associated occupations (right column). Specifically, we cluster the set of AI applications for each firm into 5 clusters using a k-means algorithm, then use a LLM to supply labels to each of the clusters, which are reported in the left column. Three example clusters are shown for each firm. The labels for the other two clusters for JPMorgan Chase are 'Data Engineering & Analytics Infrastructure' and 'Automation & Workflow Optimization'. For Walmart, they are 'Personalization, Recommendations, and Enhanced Search' and 'Data Pipelines, Integration, and Big Data Infrastructure'.

Table 3: AI task Exposure and Labor Demand

Dependent Variable: 100 × 5-year Davis et al. (1996) change in share of job posting skills related to task

	OLS		
	(1)	(2)	(3)
Task-level AI Exposure	-4.78***	-4.75***	-4.80***
	(-13.46)	(-13.97)	(-14.14)
N	13,241,933	13,241,933	13,238,128
R^2	0.073	0.10	0.32
F-stat			
Task Importance Control	X	X	X
Mean Task Exposure Control	X	X	
Firm × Year FE	X	X	
Occ × Year FE		X	
Firm × Occ × Year FE			X

Note: This table includes estimates of Equation (5) from the main text. The unit of analysis is at the task–firm level, and the dependent variable is the Davis et al. (1996) change in the share of job posting skills demanded that are textually linked to the give task when comparing across the firms' job postings for a given occupation in the next 5 years versus the previous 5 years. The "Task-AI Exposure" is defined in equation 3 of the main text. Standard errors are clustered by occupation-firm, with associated *t*-statistics reported in parentheses. Besides the designated fixed and also the average task-level AI exposure of the occupation in specifications without the full complement of firm × occupation × year fixed effects. See section 1.4 of the main text for details.

Table 4: The impact of firm-level AI use on firm growth rates

Dependent variable: 100× 5-year growth rate in the firm outcome designated in each column

	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Sales	Emp	Profit	TFP	Sales	Emp	Profit	TFP
log(1 + AI uses)	5.57***	3.97***	5.91***	5.32***	9.64***	6.73***	8.14**	7.44***
	(3.82)	(3.48)	(4.15)	(5.80)	(3.87)	(3.70)	(3.18)	(5.05)
N	12757	13225	11652	6065	12282	12688	11246	6035
R-sq	0.14	0.12	0.13	0.25	0.070	0.051	0.027	0.19
F-stat					2221.2	2369.4	1905.5	1080.5
Controls	X	X	X	X	X	X	X	X
Ind × Year FE	X	X	X	X	X	X	X	X

Note: This table shows results from estimating Equation (35). The dependent variable is the 5-year forward growth rate in the designated firm outcome. In the last four columns, we estimate the specification using twos-stage least squares with the instrument $\log(1 + \text{Number of AI Uses})^{IV}_{f,t}$ IV defined in Section 3.2 of the main text, with corresponding IV F-statistics from the first-stage regression reported in the table. Controls include a lagged one-year growth rate and level of the dependent variable; the logs of total employment both based on Revelio resume counts and Compustat employment counts; and 3-digit NAICS × year fixed effects. We cluster standard errors by firm and report corresponding t-statistics in parentheses.

Table 5: AI exposure and occupational employment growth (5-year horizon)

Dependent variable: 100× 5-year growth rate in the occupation–firm employment

	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Exposure Average	-8.02***	-8.34***	-7.73***	-5.50***	-15.0***	-14.1***	-14.5***	-10.4***
	(-12.80)	(-14.35)	(-12.87)	(-10.47)	(-10.20)	(-10.69)	(-16.94)	(-11.37)
AI Exposure Concentration	1.21**	1.50***	1.87***	1.42***	8.79***	7.25***	7.46***	7.33***
	(3.24)	(4.37)	(4.94)	(4.87)	(6.38)	(5.56)	(8.34)	(5.74)
$\log(1 + \text{AI uses})$	12.0***	11.3***			18.2***	17.2***		
	(20.44)	(19.28)			(17.27)	(14.92)		
N	2037346	2037345	2035179	2035179	2017628	2017628	2016990	2016990
R^2	0.11	0.18	0.54	0.60	0.088	0.079	0.018	-0.0017
F-stat (AI Exposure Average)					1066.3	1126.0	2490.7	1368.5
F-stat (AI Exposure Concentration)					1025.2	1074.9	2273.3	594.0
F-stat ($\log(1 + \text{AI uses})$)					1748.9	1725.7		
Controls	X	X	X	X	X	X	X	X
Year FE	X				X			
Industry × Year FE		X				X		
Firm × Year FE			X	X			X	X
Occ × Year FE				X				X

Note: This table shows regression estimates of Equation (16) from the main text. Columns (1) through (4) correspond to the OLS estimates. Columns (5) through 8 correspond to the two-stage least squares with the set of instruments described in Section 3.2 of the main text. We include associated F-statistics for each instrumented variable in the table. In addition to the designated fixed effects, all specifications include a control for the lagged one-year employment growth; specifications with only year fixed effects; the logs of total employment both based on both Revelio resume counts and Compustat employment counts. Observations are weighted by the yearly occupation-firm cell's share of employment. T-stats based on standard errors clustered by occupation–firm are in parentheses.

Table 6: Impact of AI on relative employment growth by occupation group

	2-digit SOC	Mean Component	Concentration Component	Firm Component	Total	% of Emp
Management	11	-2.27	1.45	1.15	0.34	19.0
Business and Financial	13	-9.57	5.65	2.26	-1.65	17.6
Architecture and Engineering	17	-6.27	2.46	2.30	-1.52	9.10
Science	19	0.85	0.031	1.43	2.31	2.36
Community and Social Service	21	10.4	-5.14	-0.15	5.07	0.33
Legal	23	9.65	-5.68	3.13	7.10	0.71
Education and Library	25	9.00	-4.61	0.32	4.71	1.00
Arts, Entertainment, Media	27	7.55	-4.48	3.07	6.15	5.38
Healthcare Practitioners	29	4.93	-1.98	-1.44	1.51	1.92
Healthcare Support	31	6.93	-3.13	-1.41	2.39	0.47
Protective Service	33	9.01	-5.25	-2.03	1.73	0.43
Food Preparation and Serving	35	12.6	-6.25	-10.8	-4.46	2.75
Cleaning and Maintenance	37	14.0	-7.92	-4.74	1.33	0.46
Personal Care and Service	39	12.1	-6.18	-4.95	0.96	1.09
Sales and Related	41	1.66	-0.74	-2.51	-1.59	13.3
Office and Administrative	43	2.82	-2.22	0.12	0.71	10.6
Farming, Fishing, and Forestry	45	12.8	-7.04	-4.80	1.00	0.46
Construction and Extraction	47	6.10	-4.03	-0.068	2.01	2.07
Installation and Repair	49	3.76	-3.06	-0.78	-0.078	2.72
Production	51	5.52	-2.38	-2.68	0.46	3.94
Transportation	53	7.86	-4.08	-3.79	-0.017	4.26

Note: This table shows results from estimating the decomposition (38) for broad 2-digit SOC occupation groups. The column "Mean Component" provides the average relative employment growth impact of the average task-level exposure to AI within the occupation group, while the "Variance Component" shows the impact of the variance in task-level exposure to AI. The "Firm-Component" column gives the employment impact of the occupation groups' average exposure to firm-level AI use. Effects are expressed relative to the aggregate average growth rate, so that the total employment-weighted effect sums to 0. See Section 3.4 of the text for further details.

A Appendix

Section A.1 contains the derivations of the model.

A.1 Model Appendix

Here, we provide more details on the model derivation.

Optimal Hours Allocation

The optimal time allocation problem has the solution

$$h(j) = \frac{w(j)^{\frac{1}{\beta}}}{\sum_{j \in J} w(j)^{\frac{1}{\beta}}}. \tag{A.1}$$

To see this, note that each worker solves the following optimization problem, taking into account the constraint on hours,

$$\mathcal{L} = \sum_{j=1}^{J} w(j) h(j)^{1-\beta}\, dj - \lambda \left(\sum_{j=1}^{J} h(j)\, dj - 1 \right). \tag{A.2}$$

The first-order condition with respect to devoted to task j $h(j)$ is

$$(1-\beta)\, w(j)\, h(j)^{-\beta} = \lambda. \tag{A.3}$$

This leads to

$$h(j) = \left[(1-\beta) \frac{w(j)}{\lambda} \right]^{\frac{1}{\beta}}. \tag{A.4}$$

Take sum on both sides,

$$\sum_{j=1}^{J} h(j) = \sum_{j=1}^{J} \left[(1-\beta) \frac{w(j)}{\lambda} \right]^{\frac{1}{\beta}} dj = (1-\beta)^{-\frac{1}{\beta}} \lambda^{\frac{1}{\beta}} \sum_{j=1}^{J} w(j)^{\frac{1}{\beta}} = 1. \tag{A.5}$$

Thus,

$$\lambda = (1-\beta) \left(\sum_{j \in J} w(j)^{\frac{1}{\beta}} \right). \tag{A.6}$$

Apply this to (A.4) yields (A.1).

Equilibrium First-Order Conditions for Labor and Capital Demand

The firm-level cost minimization problem can be expressed as

$$\min_{Y_f(o)} \int_{\mathcal{O}} P(o,f) Y_f(o) \quad \text{s.t. } Y_f = \left(\int_{\mathcal{O}} Y_f(o)^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}. \tag{A.7}$$

Given the above, the labor demand of firm f for occupation o is equal to

$$Y(o,f) = P(o,f)^{-\chi}\, Z_f^{-\chi}\, Y_f \tag{A.8}$$

where

$$Z_f \equiv \left(\int_O P(o,f)^{1-\chi}\right)^{-\frac{1}{1-\chi}} \tag{A.9}$$

and $P(o,f)$ denotes the marginal cost firm f it pays for the output of occupation o. These prices need not be the same across firms. To simplify notation, in what follows we will suppress the firm subscripts. Firms make profits because of imperfect competition, reflecting both pricing power in product markets and monopsony power in labor markets. Denote their markup over marginal cost Z_f^{-1} by $\Theta = \frac{\theta}{\theta-1} > 1$. Since the firm has monopsony power in the labor market, its marginal cost will exceed its average cost, as we discuss further below. As a result,

$$\begin{aligned}
P_f Y_f &= \Theta \int_O Y_f(o) P(o,f) \\
P_f Y_f &= \Theta \int_O P(o,f)^{1-\chi} \left(\int_O P(o,f)^{1-\chi}\right)^{\frac{\chi}{1-\chi}} Y_f \\
P_f &= \Theta \left(\int_O P(o,f)^{1-\chi}\right)^{\frac{1}{1-\chi}} = \Theta\, Z_f^{-1}
\end{aligned} \tag{A.10}$$

Each firm faces the inverse demand curve

$$Y_f = P_f^{-\theta}\, P^{\theta}\, \bar{Y} \tag{A.11}$$

where

$$P \equiv \left(\int_{\mathcal{F}} P_f^{1-\theta}\right)^{\frac{1}{1-\theta}}. \tag{A.12}$$

Without loss of generality, we can normalize the aggregate price index $P = 1$, which implies

$$Y_f = P_f^{-\theta}\, \bar{Y}, \tag{A.13}$$

and given the price above, this implies

$$Y_f = \Theta^{-\theta}\, Z_f^{\theta}\, \bar{Y} \tag{A.14}$$

Now, consider the occupation's task minimization problem

$$\min_{y(j)} \sum_{j\in J} p(j)\, y(j) \qquad s.t. \;\; Y(o,f) = \left(\sum_{j\in J} y(j)^{\frac{\psi-1}{\psi}}\right)^{\frac{\psi}{\psi-1}} \tag{A.15}$$

Here $p(j)$ is the marginal cost index of producing task j output $y(j)$ after optimal input choices have been made within task j. Due to monopsony power, the marginal cost $p(j)$ will exceed the average cost of producing $y(j)$ given that the firm will internalize that hiring a marginal worker will require paying higher wages to additional, inframarginal workers. Our CES structure admits the following Hicksian demand for $y(j)$ from the FOC for problem (A.15):

$$y(j) = p(j)^{-\psi} \left[\sum_{j \in J} p(j)^{1-\psi}\right]^{\frac{\psi}{1-\psi}} Y(o,f) \tag{A.16}$$

Here, $P(o,f)$ is the marginal cost of occupation o's output.

$$P(o,f) = \left[\sum_{j \in J} p(j)^{1-\psi}\right]^{\frac{1}{1-\psi}} \tag{A.17}$$

Using the above combined with equations (A.8) and (A.14), we get

$$y(j) = \frac{1}{p(j)^{\psi}} X(o)^{\chi-\psi}\, Z_f^{\theta-\chi} \Theta^{-\theta}\, \bar{Y}. \tag{A.18}$$

where

$$X(o,f) \quad = \quad \left[\sum_{j \in J} p(j)^{1-\psi}\right]^{-\frac{1}{1-\psi}} = P(o,f)^{-1} \tag{A.19}$$

In the above, $X(o,f)$ is the productivity (the inverse of the unit cost) of productivity O and Z_f is the productivity of firm f.

The factor allocation associated with the monopsonistic cost minimization problem is isomorphic to the solution to a perfectly competitive firm which faces a wedge between the marginal cost of labor and the wage. We denote this wedge by $\mathcal{M}(j)$, which we compute below. In deriving comparative statics, we assume that the firm treats $\mathcal{M}(j)$ as constant when choosing its factor allocations for simplicity and analytical tractability. The cost minimization problem within task j is

$$\min_{l(j),k(j)} q(j)k(j) + w(j)\mathcal{M}(j)l(j) \qquad s.t. \;\; y(j) = \left[(1-\gamma_j)k(j)^{\frac{\nu-1}{\nu}} + \gamma_j\, l(j)^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}} \tag{A.20}$$

Going forward, we make the re-parameterization $a_j \equiv \gamma_j^{\nu}$, $b_j \equiv (1-\gamma_j)^{\nu}$. After solving (A.20), the per-unit cost of task j equals

$$p(j) = \Big(a_j \left[\mathcal{M}(j)w(j)\right]^{1-\nu} + b_j\; q(j)^{1-\nu}\Big)^{\frac{1}{1-\nu}} \tag{A.21}$$

Using equation (A.21), we can rewrite (A.18) as

$$y(j) = \Big(a_j \left[\mathcal{M}(j)w(j)\right]^{1-\nu} + b_j\; q(j)^{1-\nu}\Big)^{\frac{-\psi}{1-\nu}} X(o,f)^{\chi-\psi}\, Z_f^{\theta-\chi} \Theta^{-\theta}\, \bar{Y}. \tag{A.22}$$

Plugging in (A.22) to the CES Hicksian demand for $k(j)$ and $l(j)$ and imposing labor market clearing gives

$$l(j) = \frac{a_j}{[\mathcal{M}(j)w(j)]^{\nu}} \left(a_j[\mathcal{M}(j)w(j)]^{1-\nu} + b_j\, q(j)^{1-\nu}\right)^{\frac{\nu-\psi}{1-\nu}} X(o)^{\chi-\psi}\, Z_f^{\theta-\chi}\Theta^{-\theta}\,\bar{Y}. \quad \text{(A.23)}$$

$$k(j) = \frac{\beta_j}{q(j)^{\nu}} \left(a_j[\mathcal{M}(j)w(j)]^{1-\nu} + b_j\, q(j)^{1-\nu}\right)^{\frac{\nu-\psi}{1-\nu}} X(o)^{\chi-\psi}\, Z_f^{\theta-\chi}\Theta^{-\theta}\,\bar{Y}. \quad \text{(A.24)}$$

Labor Market Clearing

If there are $N(o,f)$ workers in a occupation–firm pair (o,f) then the total supply of

$$L_o(j) = N(o,f)\, h^{1-\beta}(j). \quad \text{(A.25)}$$

Using the properties of the Fréchet distribution, it follows that the expected measure of workers to job o in firm f is equal to

$$N(o,f) = \underbrace{\frac{1}{\int_{f'\in\mathcal{F}}\int_{o'\in\mathcal{O}} W(o,f)^{\frac{\zeta}{1+\zeta}}}}_{\bar{\zeta}} W(o,f)^{\zeta}. \quad \text{(A.26)}$$

where $W(o,f)$ is the total earnings on the job,

$$W(o,f) \equiv \sum_{j\in J_o} h(j)^{1-b} w(j). \quad \text{(A.27)}$$

Given (A.1), the total earnings for that job are equal to

$$W(o,f) = \sum_{j\in J_o} h(j)^{1-\beta} w(j) = \frac{\sum_{j\in J} w(j)^{\frac{1}{\beta}}}{(\sum_{j\in J} w(j)^{\frac{1}{\beta}})^{1-\beta}} = \left[\sum_{j\in J} w(j)^{\frac{1}{\beta}}\right]^{\beta}. \quad \text{(A.28)}$$

Notice that as long as hours are flexible, $0 < b < 1$, then the occupation level wage is convex in the task prices. Put differently, because the worker can reallocate hours, she benefits from a mean-preserving spread in $w(j)$.

So, the total labor supply for task j is equal to

$$\begin{aligned} h(j)^{1-\beta}\, N(o,f) &= w(j)^{\frac{1-\beta}{\beta}} \left(\sum_{j\in J} w(j)^{\frac{1}{\beta}}\right)^{\beta-1} \left[\sum_{j\in J} w(j)^{\frac{1}{\beta}}\right]^{\zeta\,\beta} \bar{\zeta}. \\ &= w(j)^{\frac{1}{\beta}-1} \left(\sum_{j\in J} w(j)^{\frac{1}{\beta}}\right)^{\beta-1+\zeta\,\beta} \bar{\zeta} \end{aligned} \quad \text{(A.29)}$$

Replacing the left-hand-side of equation (A.23) with the equation for labor supply for task j

yields a system of J equations in J unkowns—the task prices $w(j)$,

$$w(j)^{\frac{1}{\beta}}\left(\sum_{j\in J_o} w(j)^{\frac{1}{\beta}}\right)^{\beta-1+\zeta\,\beta}\bar{\zeta} = a_j\,\mathcal{M}(j)^{-\nu}\,w(j)^{1-\nu}\left(a_j[\mathcal{M}(j)w(j)]^{1-\nu}+b_j\,q(j)^{1-\nu}\right)^{\frac{\nu-\psi}{1-\nu}}\;X(o,f)^{\chi-\psi}\,Z_f^{\theta-\chi}\Theta^{-\theta}\,\bar{Y}. \tag{A.30}$$

where we will show below that the labor wedge $\mathcal{M}(j)$ is constant.

Using equation (A.21), we can write

$$X(o,f) \quad = \quad \left[\sum_{j\in J}\left(a_j\,[\mathcal{M}(j)\,w(j)]^{1-\nu}+b_j\;q(j)^{1-\nu}\right)^{\frac{1-\psi}{1-\nu}}\right]^{-\frac{1}{1-\psi}} \tag{A.31}$$

Wage markdown

Firms are facing an upward labor supply curve and are monopsonists in the labor market. Next, we derive $\mathcal{M}(j)$ the wedge between the marginal cost of type j labor and the wage $w(j)$. We will show below that the labor wedge is a constant equal to $\mathcal{M}(j)=1+1/\zeta$, a familiar expression in the monopsony literature which obtains when the labor supply curve has a constant elasticity. The proportional difference between marginal cost and the wage equals $\frac{1}{\zeta}$, the elasticity of the inverse labor supply curve, so the firm optimally marks down wages below marginal cost by a factor equal to $1/\mathcal{M}(j)$.

Derivation: In order to derive $\mathcal{M}(j)$, we require several building blocks. First, we need to know how the quantity of task j labor $l(j)$ changes with respect to its own price $w(j)$

$$\frac{\partial\log l(j)}{\partial\log w(j)}=\frac{\partial l(j)}{\partial w(j)}\frac{w(j)}{l(j)}=\left(\frac{1}{\beta}-1\right)+\left(1-\frac{1}{\beta}+\zeta\right)h(j)=\left(\frac{1}{\beta}-1\right)(1-h(j))+\zeta h(j). \tag{A.32}$$

We also need the cross-price terms

$$\frac{\partial\log l(j)}{\partial\log w(k)}=\left(1-\frac{1}{\beta}+\zeta\right)h(k), \tag{A.33}$$

which has a sign which depends on whether the between task substitution effect $(1-1/\beta<0)$ dominates the induced increase in the number of workers from higher total wages (ζ).

To derive these equations (A.32-A.33), we work with the identity

$$\log l(j)=\log\bar{\zeta}+(1-1/\beta)\log w(j)+[\beta-1+\zeta\beta]\log\sum_{j\in J}\exp(\tfrac{1}{\beta}\log w(j)). \tag{A.34}$$

It is straightforward that differentiating the above equation yields the desired results , since

$$\frac{\partial}{\partial w(j)}\log\sum_{j\in J}\exp(\tfrac{1}{\beta}\log w(j))=\frac{1}{\beta}\frac{w(j)^{1/\beta}}{\sum_{j'\in J}w(j)^{1/\beta}}. \tag{A.35}$$

Next, we need to understand the set of wage changes which allows the firm to increase $l(j)$ while holding the quantity of labor in all other tasks fixed. In elasticity form, the requisite wage changes are

$$\frac{d\log w(k)}{d\log l(j)}\Big|_{\substack{d\log l(k)=0\\ k\neq j}} = \mathbf{1}[k=j]\frac{\beta}{1-\beta} + \frac{\frac{1}{\beta}-1-\zeta}{\left(\frac{1}{\beta}-1\right)\zeta}h(j), \tag{A.36}$$

an expression we obtain by inverting the Jacobian matrix capturing the set of elasticities of task quantities with respect to task prices.

Finally, we need to understand how total costs change with each of the task-level wages

$$\frac{\partial[W(o,f)N(o,f)]}{\partial w(j)} = l(j)(1+\zeta). \tag{A.37}$$

To derive equation (A.37), we start with the fact that total wage earnings equals $\bar{\zeta}W(o,f)^{\zeta+1}$. Then by differentiating and using the definition of $l(j)$ from equation (A.29), we get

$$\frac{\partial[W(o,f)N(o,f)]}{\partial w(j)} = (1+\zeta)\underbrace{\bar{\zeta}W(o,f)^{\zeta}}_{=N(o,f)}\frac{w(j)^{\frac{1}{\beta}-1}}{\left[\sum_{j\in J}w(j)^{\frac{1}{\beta}}\right]^{\beta-1}} = (1+\zeta)l(j). \tag{A.38}$$

We can then combine these pieces (A.32, A.33, A.37) to compute marginal cost:

$$\begin{aligned}\frac{\partial[W(o,f)N(o,f)]}{\partial l(j)}\Big|_{\substack{d\log l(k)=0\\ k\neq j}} &= \sum_{k=1}^{J}\frac{\partial[W(o,f)N(o,f)]}{\partial w(k)}\frac{w(k)}{l(j)}\frac{d\log w(k)}{d\log l(j)}\Big|_{dl(k)=0} \\ &= (1+\zeta)\left\{\frac{\beta\, w(j)}{1-\beta} + \sum_{k=1}^{J}\frac{w(k)l(k)}{l(j)}\frac{\frac{1}{\beta}-1-\zeta}{\left(\frac{1}{\beta}-1\right)\zeta}h(j)\right\}.\end{aligned} \tag{A.39}$$

Recalling that $\mathcal{M}(j)$ is the ratio of marginal cost to the wage, and rearranging, we get that

$$\mathcal{M}(j) = \frac{(1+\zeta)}{\frac{1}{\beta}-1}\frac{\zeta[s(j)-h(j)]+(\frac{1}{\beta}-1)h(j)}{\zeta s(j)} = M \equiv 1+\frac{1}{\zeta}, \tag{A.40}$$

where $s(j)\equiv w(j)l(j)/[W(o,f)\,N(o,f)]$ is the cost share in task j. To obtain the final expression, we note that $s(j)=h(j)$.

A Log-linear Approximation

To derive an approximate solution, let us focus on the the symmetric steady state with $a_j = a$, $\beta_j=\beta$, $w(j)=w$ and $q(j)=q$. Suppose $j=1$ gets shocked and the other tasks are not. Given the definitions of the elasticities, we can write

$$q(1) = q\,e^{-\epsilon}, \qquad q(j)=q, \qquad j\geq 2 \tag{A.41}$$

$$w(1) = w\,e^{\eta_o\,\epsilon}, \qquad w(j) = w\,e^{\eta_c\,\epsilon}, \qquad j\geq 2. \tag{A.42}$$

If we replace the above into (A.30) for $j = 1$ and $j = 2..J$, we get two equations, one for the 'shocked' task and another equation which is common to all unshocked tasks, for $j = 1$. After dividing both sides of each equation with equation (A.30) evaluated at the pre-shock equilibrium, we obtain

At $j = 1$:

$$e^{\eta_o \frac{1}{\beta}\epsilon} \left(\frac{1}{J} e^{\eta_o \frac{1}{\beta}\epsilon} + \frac{J-1}{J}\, e^{\eta_c \frac{1}{\beta}\epsilon}\right)^{\beta - 1 + \zeta\,\beta} = e^{\eta_o\,(1-\nu)\epsilon} \left(s_l\, e^{\eta_o\,(1-\nu)\epsilon} + (1-s_l) e^{-(1-\nu)\,\epsilon}\right)^{\frac{\nu-\psi}{1-\nu}} \left(\frac{\tilde{X}(o,f)}{X(o,f)}\right)^{\chi-\psi} \tag{A.43}$$

at $j \neq 1$

$$e^{\eta_c \frac{1}{\beta}\epsilon} \left(\frac{1}{J} e^{\eta_o \frac{1}{\beta}\epsilon} + \frac{J-1}{J}\, e^{\eta_c \frac{1}{\beta}\epsilon}\right)^{\beta - 1 + \zeta\,\beta} = e^{\eta_c\,(1-\nu)\epsilon} \left(s_l\, e^{\eta_c\,(1-\nu)\epsilon} + 1 - s_l\right)^{\frac{\nu-\psi}{1-\nu}} \left(\frac{\tilde{X}(o,f)}{X(o,f)}\right)^{\chi-\psi} \tag{A.44}$$

where

$$\frac{\tilde{X}(o,f)}{X(o,f)} = \left[\frac{J-1}{J}\left(s_l\, e^{\eta_c\,(1-\nu)\epsilon} + 1 - s_l\right)^{\frac{1-\psi}{1-\nu}} + \frac{1}{J}\left(s_l\, e^{\eta_o\,(1-\nu)\,\epsilon} + (1-s_l) e^{-(1-\nu)\,\epsilon}\right)^{\frac{1-\psi}{1-\nu}}\right]^{-\frac{1}{1-\psi}} \tag{A.45}$$

Taking logs of both sides (A.43) and (A.44), differentiating with respect to ϵ, evaluating it at $\epsilon = 0$ yields two linear equations in two unknowns, η_o and η_c. Define

$$s_l \equiv \frac{a\,M^{1-\nu}\,w^{1-\nu}}{a\,M^{1-\nu}\,w^{1-\nu} + b\,q^{1-\nu}}, \qquad s_k = 1 - s_l \tag{A.46}$$

is the task-level labor share and capital share of output, respectively. The solution to these equations are

$$\eta_c = -\frac{s_k}{J}\,\frac{(\nu-\chi)(1-\beta) - \beta\Big(\nu\,(\chi-\psi) + \zeta(\nu-\psi)\Big)}{\Big(s_k\,\nu + s_l\,\chi + \zeta\Big)\Big(1 - \beta\,(1 - s_k\,\nu - \psi\,s_l)\Big)} \tag{A.47}$$

as long as $\nu \geq 1$, the above expression is decreasing in ψ. Cross-task spillovers become positive if the within-occupation task complementarity is strong enough, that is the parameter ψ is less than

$$\bar{\psi} = \frac{((\chi+\zeta+1)\,\nu - \chi)\,\beta + \chi - \nu}{\beta\,(\nu+\zeta)} \tag{A.48}$$

The own-task elasticity, which captures the impact of innovation $\varepsilon(j)$ on $w(j)$, is given by

$$\eta_o = -\frac{s_k}{J}\,\frac{(\nu-\chi)(1-\beta) + \beta\left((J-1)(\nu-\psi)\zeta + \nu(\psi-\chi) + J\,(\nu-\psi)(s_k\,\nu + s_l\,\chi)\right)}{\Big(s_k\,\nu + s_l\,\chi + \zeta\Big)\Big(1 - \beta\,(1 - s_k\,\nu - \psi\,s_l)\Big)}. \tag{A.49}$$

A sufficient, but not necessary, condition for η_o to be negative is that the number of tasks is sufficiently large

$$J \geq \frac{\nu + \zeta}{s_l\,\chi + s_k\,\nu + \zeta} \tag{A.50}$$

and $\nu \geq \psi$.

Last, to derive the expression for wage growth,

$$\log W_1(o,f) - \log W_0(o,f) = b\,\log\left[\sum_{j\in J} w_1(j)^{\frac{1}{b}}\right] - \log\left[\sum_{j\in J} w_0(j)^{\frac{1}{b}}\right] \tag{A.51}$$

We can consider a technology shock that affects the job (o,f) described by a vector of $\varepsilon_1 ... \varepsilon_J$. Thus, we can write

$$w_1(j) = w_0(j) e^{\eta_o\,\varepsilon(j) + \eta_c \sum_{j' \neq j} \varepsilon(j')}. \tag{A.52}$$

Plugging the above into the change in wages, and approximating around $\varepsilon(j) = 0$ for all j, we obtain Approximating the above using a second order expansion, we get that

$$\begin{aligned} \log W_1(o,f) - \log W_0(o,f) \approx{}& +\eta_c\,(J-1)\,\frac{1}{J}\sum_j \varepsilon(j) + \eta_o \frac{1}{J}\sum_j \varepsilon(j) + \\ & + \frac{b}{2}\left(\frac{\eta_o + \eta_c}{b}\right)^2 \frac{J-1}{J^2}\,\left(\sum_j \varepsilon(j)^2 - \frac{2}{J-1}\sum_j\sum_{j'>j}\varepsilon(j)\varepsilon(j')\right) \end{aligned} \tag{A.53}$$

Define

$$m(\varepsilon) = \frac{1}{J}\sum_j \varepsilon(j) \tag{A.54}$$

$$\begin{aligned} C(\varepsilon) \equiv{}& \frac{1}{J}\sum\left(\varepsilon(j) - m(\varepsilon)\right)^2 \\ ={}& \frac{J-1}{J^2}\left(\sum_j \varepsilon(j)^2 - \frac{2}{J-1}\sum_j\sum_{j'>j}\varepsilon(j)\varepsilon(j')\right) \end{aligned} \tag{A.55}$$

Replacing (A.54) and (A.55) into (A.53) yields the expression for wage growth in the text. The expression for employment growth follows directly from (A.26).

The last step consists of deriving the elasticity to a shock to firm productivity. In the symmetric steady state with $a_j = a$, $\beta_j = \beta$, and $q(j) = q$,

$$J^{\beta-1+\zeta\,\beta}\,w^{1+\zeta}\bar{\zeta} = a\,w^{1-\nu}\left(a\,w^{1-\nu} + b\,q^{1-\nu}\right)^{\frac{\nu-\psi}{1-\nu}}\;X(o,f)^{\chi-\psi}\,Z_f^{\theta-\chi}\Theta^{-\theta}\,\bar{Y}. \tag{A.56}$$

where

$$X(o,f) \;=\; \left[J\,\left(a\,w^{1-\nu} + b\;q^{1-\nu}\right)^{\frac{1-\psi}{1-\nu}}\right]^{-\frac{1}{1-\psi}} \tag{A.57}$$

Plugging the last to the second last equation

$$J^{\beta-1+\zeta\,\beta}\,w^{1+\zeta}\bar{\zeta} = a\,w^{1-\nu}J^{\frac{\psi-\chi}{1-\psi}}\left(a\,w^{1-\nu}+b\,q^{1-\nu}\right)^{\frac{\nu-\chi}{1-\nu}}\;Z_f^{\theta-\chi}\Theta^{-\theta}\,\bar{Y}. \tag{A.58}$$

Suppose productivity of the firm improves

$$\tilde{Z}_f = Z_f\,e^{\epsilon} \tag{A.59}$$

then task prices will change

$$w(j) = w\,e^{\eta_z\,\epsilon}, \qquad \forall j \tag{A.60}$$

and (A.58) becomes

$$J^{\beta-1+\zeta\,\beta}\,w^{1+\zeta}e^{(1+\zeta)\eta_z\,\epsilon}\bar{\zeta} = a\,w^{1-\nu}e^{(1-\nu)\eta_z\,\epsilon}J^{\frac{\psi-\chi}{1-\psi}}\left(a\,w^{1-\nu}e^{(1-\nu)\eta_z\,\epsilon}+b\,q^{1-\nu}\right)^{\frac{\nu-\chi}{1-\nu}}\;Z_f^{\theta-\chi}e^{(\theta-\chi)\epsilon}\Theta^{-\theta}\,\bar{Y}. \tag{A.61}$$

Divide both sides of (A.61) with the corresponding sides of (A.58) gives

$$e^{(1+\zeta)\eta_z\,\epsilon} = e^{(1-\nu)\eta_z\,\epsilon}\left(s_l\,e^{(1-\nu)\eta_z\,\epsilon}+s_k\right)^{\frac{\nu-\chi}{1-\nu}}\;e^{(\theta-\chi)\epsilon} \tag{A.62}$$

Taking logs of both sides, differentiating with respect to ϵ, evaluating at $\epsilon = 0$ and solving for η_z, we get

$$\eta_z = \frac{\theta-\chi}{s_k\,\nu+s_l\,\chi+\zeta} \tag{A.63}$$

and in terms of wage earnings

$$\begin{aligned} \log W_1(o,f) - \log W_0(o,f) &= \beta\,\log\left[J\,w^{\frac{1}{\beta}}\,e^{\frac{1}{\beta}\,\eta_z\,\epsilon}\right] - \beta\log\left[J\,w^{\frac{1}{\beta}}\right] \\ &= \eta_z\,\epsilon. \end{aligned} \tag{A.64}$$

If all task prices go up by the same amount, there is no reallocation, and wages just go up by the same amount.

Task Reallocation

We now derive the elasticity of task-j hours allocation $h(j)$ to AI-driven cost changes. Applying (A.52) to (A.1) for some task j around the symmetric initial equilibrium, we have

$$
\begin{aligned}
\Delta \log h(j) &= \frac{1}{\beta}\left(\eta_o \varepsilon(j) + \eta_c \sum_{j' \neq j} \varepsilon(j')\right) - \log\left(\sum_k \exp\Big(\frac{\eta_o}{b}\varepsilon(k) + \frac{\eta_c}{\beta}\sum_{k' \neq k}\varepsilon(k')\Big)\right) \\
&\approx \frac{1}{\beta}\left[\eta_o \varepsilon(j) + \eta_c \sum_{j' \neq j}\varepsilon(j') - \frac{1}{J}\Big(\sum_k \eta_o \varepsilon(k) + \eta_c \sum_k \sum_{k' \neq k}\varepsilon(k')\Big)\right] \\
&= \frac{1}{\beta}\left[\eta_o\left(\varepsilon(j) - m(\varepsilon)\right) + \eta_c\left(Jm(\varepsilon) - \varepsilon(j) - (J-1)m(\varepsilon)\right)\right]
\end{aligned} \tag{A.65}
$$

which simplifies to

$$
\Delta \log h(j) \approx \frac{\eta_o - \eta_c}{\beta}\left(\varepsilon(j) - m(\varepsilon)\right) \tag{A.66}
$$

Thus task-specific hours reallocation is proportional to the task-level AI capital cost shock relative to the average shock across all tasks within the occupation.

A.2 Extracting Firm-Level AI Applications

To identify potential AI applications at the firm level, we first impose a filter on the text in workers' job description. After converting job descriptions to lower case, require that the description includes at least one of the following strings: " ai "; "artificial intelligence"; "automl"; "computer vision"; "convolutional neural net"; "deep learning"; "genai"; "generative adversarial net"; "generative ai"; "generative artificial intelligence"; "generative pre trained transformer"; "generative pretrained transformer"; "gradient boost"; "hugging face"; "keras"; "large language model"; "lightgbm"; " llm "; "lstm"; "machine learning"; " ml "; "mlflow"; "natural language processing"; "neural net"; " nlp "; "prompt engineer"; "pytorch"; "recurrent neural net"; "reinforcement learn"; "reinforcement learning"; "rnn"; "tensorflow"; "xgboost". We further require AI positions to come from jobs with 2-digit SOC code between 11 through 19 (professional occupations).[18] Upon reading many examples, we find that AI-tagged positions coming from these occupations are nearly exclusively direct implementers of AI, while this is occasionally not the case for the non-professional occupations.

This results in 561,974 distinct job positions which describe implementing artificial intelligence in at least one application. We consider the position to be active at a firm in a specific year if the position is current for at least a 6 month window within the given year. We next apply a series of filters using large language models to read these descriptions of AI positions to extract and clean the phrases which describe specific ways in which AI is being applied. To do this, we use the Llama 3.1 70B model created by Meta. We access the model using an API provided by DeepInfra. We set the temperature parameter to zero in all prompts in order minimize any potential variability in responses to the exact same query.

[18] Upon examination, codes beginning with 4-digit SOC 19-30 had a lot of false positives, so we also excluded any potential AI resumes with this code.

Step-1 LLM Filter: Identifying and cleaning AI-related phrases

Our first-step LLM filter extracts the specific raw phrases in a job description which describe using AI, as well as an LLM-generated summary of the AI application. The prompt instructs the LLM to follow a four-step process in order to guide its "reasoning". The steps are as follows: 1), filtering out the tasks in the task which are unrelated to applications of AI (including discarding descriptions of hardware related to AI rather than the specific use of AI); 2), generate a list of applications identified from the first step; 3), audit answers to ensure that the AI application is clearly specified; and 4), reread the original text to make sure no AI applications were missed in the original reading. Finally, the LLM is asked to report to the user the key applications filtered from the text; the original raw text that generated the specific key application; and finally, the final answer which is the cleaned AI applications.

Our specific LLM prompt for this step is as follows:

*Your current task is to review the following descriptions of job duties being performed by employees of the same company and summarize each of the applications of AI that you see being performed. The goal is to produce an itemized list, where each item corresponds with a different use case for artificial intelligence methods being described. For each application, please describe, in a few sentences based ONLY on the resume descriptions, what functions AI tools are being applied to perform (it is important not to make predictions unless a use case is described in the text). Your answers should be focused on which tasks these AI tools are being used to perform, rather than on which tools are being used. In other words, I only want you to summarize instances in which these employees describe using AI to perform a specific function or solve a particular problem. I am looking for descriptions of the tasks and functions that *the AI tools themselves are performing*, rather than just the responsibilities or activities of the employees who are working with those tools.*

To organize your efforts, I suggest you follow a four-step process. In the first step, please filter out descriptions of tasks which are unrelated to applications of artificial intelligence. If a description does not refer to how an artificial intelligence method is being used (e.g., because it describes development of hardware or other infrastructure related to AI deployment), please disregard the information. In the second step, produce your temporary itemized list from the filtered text. Now let's start the third step: Think aloud. Please audit your answers according to the original text. Sometimes, a task is clearly AI-related, but the specific application is not really specified. An example would be an employee mentioning that they are maintaining data infrastructure or deploying algorithms without saying anything about which data they are using or what the purpose of the underlying algorithms are. When reviewing your preliminary set of bullets, feel free to discard items which fall into this category of not specifying an actual application. For fourth step, please provide your final answer to improve your previous answers. Before finalizing your answer, please also reread the original body of text and identify any additional applications, if any, which were not included in the original list. Extract key applications from the following text document. Please output ONLY as a JSON list (Do not include "“' and anything else). The JSON should represent a table with three columns:

(1) The first column, labeled 'Key Application', should contain concise summaries or key insights extracted from the text.

(2) The second column, labeled 'Raw Excerpt', should include the corresponding raw excerpts from the text that support each key point.

(3) The third column, labeled 'Final Answer', should include your final answer.

< INSERT JOB DESCRIPTION HERE >

END PROMPT

This prompt does not require that a given position can only use AI for one purpose. Accordingly, out of the 561,974 distinct AI positions, this first prompt identifies 1,365,190 distinct applications of AI.

Step-2 LLM Filter: Removing uninformative text

In the second step we feed in the LLM's final response (the third output from the step-1 query) as input for the query. Responses from the first-step query typically take the form "AI tools are being used to..." or "using NLP to..." followed by the actual application. Because we don't want our textual representations of documents to be biased by these generic and uninformative phrases about the particular AI techniques–rather we want to highlight the specific application, not the particular tool being used to accomplish the application–we devise a prompt designed to filter such language from the text. The prompt allows for deleting an AI application entirely if the description of its use is still too vague to offer a clearly-defined specific application. After following this step, we have 1,096,725 filtered AI applications remaining. The second prompt is as follows:

The excerpt below describes how an artificial intelligence technology is being applied. Assume that it is already known that the excerpt refers to a use of artificial intelligence; the reader only wants to know the specific final application. Therefore, all references to any type of AI tool (e.g. natural language processing, machine learning, computer vision, generative AI, or any specific AI/ML algorithm) are redundant and should be stripped from the text. If the text only contains reference to an AI tool and without a clearly specified application, you should return 'N/A' when you filter the text.

For reference, here are a few examples of correctly applied filters:

-'AI tools are being used to measure text similarity in educational settings using NLP' should

become ‘Measure text similarity in educational settings’

-‘Machine learning is being applied to perform tasks related to database analysis and firmware/software development for embedded environments’ should become ‘Perform tasks related to database analysis and firmware/software development for embedded environments’

-‘AI-powered chatbots are being used to provide customers with quick solutions and answers using natural language processing capabilities.’ should become ‘Provide customers with quick solutions and answers.’

-‘Analyzing customer reviews using NLP to understand customer needs and wants’ should become ‘Analyze customer reviews to understand customer needs and wants’

-‘AI tool is being used to deploy computer vision model’ should become ‘N/A”, because computer vision models themselves are an AI tool, and the exact use of computer vision is not specified.’

With this in mind, please filter the following excerpt describing an AI application. < STEP 1 LLM OUTPUT HERE >

END PROMPT

Step 3 LLM Filter: Small refinements on step 2

Upon inspection of the LLM output in the second step, we found a few specific phrases which were more likely to be associated with some remaining uninformative text that occasionally bypassed the filter. Accordingly, for the final step we first identify a small subset the AI-related applications with the specific keywords *‘data analysis’, ‘text analysis’, ‘predictive analytics’, ‘visualization’, ‘predictive analysis’*. Because there are a much smaller set of texts to consider in this step, we use the more expensive but higher-performing GPT-4o model using the OpenAI API. The prompt is: *The excerpt below describes how an artificial intelligence technology is being applied. Please determine if ithe application is very specific. If yes, please summarize the application (without outputing anything else). All references to any type of AI tool (e.g. natural language processing, machine learning, computer vision, generative AI, or any specific AI/ML algorithm) are redundant and should be stripped from the text. Otherwise, respond ‘N/A’. Here are some examples:*

-‘Predictive Analytics’ should be ‘N/A’ as it is very broad;

-‘Data Visualization’ should be ’N/A’ as it is very broad;

-‘AI-driven NFT Collection Visualization’ should be kept as it is a very specific application.

-*'Perform exploratory data analysis for invoice anomalies' should be 'invoice anomalies'*

-*'Provide self-service data access and custom visualization interfaces for the oceanic team' should be 'custom visualization interfaces for the oceanic team' as this is a specific application.*

With this in mind, please filter the application: <INSERT FILTERED APPLICATION HERE>

END PROMPT

Upon application of this final filter, we are left with 1,056,068 distinct AI applications in our final sample.

A.3 Categorizing AI Applications

In this section, we provide additional detail about the 20 categories of AI applications that are extracted for use in Figures 3 and 4, as well as examples discussed in the main text. Given our list of cleaned AI applications, we first run a k-means algorithm which partitions the set of applications into 20 clusters. We then feed a csv with a list of 500 randomly-selected clusters into OpenAI's o3 model and ask the model to produce short labels which summarize the nature of the applications and allow a reader to distinguish between clusters.

Here, we produce additional output from the LLM which provides additional detail about each of the 20 clusters. In particular, we also asked the model to provide a one paragraph summary of each cluster. This information about each of the clusters, which is almost entirely AI-generated and only lightly edited, is reproduced here.

1. **Real-Time Fraud Detection**
   Tools in this cluster watch payment and account activity to spot fraud the moment it happens. They compare each new transaction with past patterns of behaviour and device location to flag anything unusual. When risk looks high they can block the payment or ask for extra verification. Banks and merchants use them to cut charge-backs and stolen-card losses. The main value is faster, more accurate fraud decision-making without slowing honest customers.

2. **Task & Workflow Automation**
   These systems act like tireless digital assistants that carry out repetitive business tasks end-to-end. They read requests from email or forms, log into internal tools, and finish the workflow with little or no human help. Common uses include rolling out new software versions, filling in benefit forms, and running nightly test suites. Managers get dashboards that show time saved and errors prevented. The appeal is lower labour cost and faster turnaround for routine work.

3. **Financial Risk Modeling**
   Applications here help lenders judge how likely a borrower is to repay. They study credit history, bank activity, job data, and the wider economy to build a risk score for each loan. Loan officers use the score and explanation notes to approve, reject, or set interest rates. The models must follow strict banking rules and give clear reasons. Better scoring means fewer bad loans and more fair access to credit.

4. **Demand & Sales Forecasting**
   These tools predict future demand for products, traffic, or revenue so companies can plan ahead. They combine sales history, upcoming promotions, holidays, and even weather to create weekly or daily forecasts. Planners see charts with confidence ranges and can adjust the forecast if they know something the model does not. The predictions drive stock orders, staffing, and pricing decisions. Accurate forecasts reduce out-of-stock losses and waste.

5. **Autonomous Navigation & Robotics**
   Systems here let robots, drones, or driverless vehicles understand their surroundings and move safely. Cameras, lasers, and GPS build a map; planning software chooses a path while avoiding people and obstacles. Industry examples include warehouse forklifts, delivery robots, and inspection drones. Much testing happens in virtual simulators before real-world trials. Key goals are safety certification and reliable behaviour in changing conditions.

6. **Marketing & Ad Optimization**
   Marketing optimisers decide how to spend ad money and which message to show each visitor. They watch clicks, purchases, and ad prices in real time, then adjust bids, budgets, and creative versions. Some systems even write new headlines or design images on the fly. Dashboards report extra sales generated versus cost. Better targeting boosts return on advertising spend and cuts wasted impressions.

7. **Text & Knowledge Retrieval**
   These platforms make it easy to search large collections of documents or get concise answers. They break text into passages, store numerical fingerprints, and fetch the parts most similar to a user query. Chatbots use the retrieved text to form grounded answers, reducing hallucination. Security features mask personal data and control who can access which documents. The benefit is faster, more accurate knowledge lookup for employees and customers.

8. **Scientific & Industrial Modeling**
   Models in this group replace slow physics simulations with fast learned estimates. Engineers feed them sensor logs or design files, and the models predict outcomes like fluid flow or material stress. This speeds up design cycles and lets teams test more options digitally before building prototypes. Outputs feed digital twins that guide real-time adjustments on the factory floor. Trust comes from careful comparison with lab or field measurements.

9. **Business Intelligence Insights**
   Business-intelligence tools here scan dashboards and databases to find trends and root causes without manual digging. They alert users when a key metric shifts and explain the main drivers by region, channel, or customer type. Natural-language summaries make insights readable for non-analysts. Teams act faster because they see problems as they emerge, not at month-end. The main gain is time saved and decisions based on data, not hunches.

10. **Conversational AI & Speech**
    These applications listen, speak, and hold natural conversations. Speech recognition turns audio into text; other models detect intent and choose the right reply. A synthetic voice then speaks the answer, matching brand tone and emotion. Use cases range from call-centre assistants to voice control in cars. Success is measured by lower call times, higher customer satisfaction, and fewer misunderstandings.

11. **Healthcare Diagnostics & Genomics**
    Medical AI here helps doctors spot disease earlier and tailor treatment. Image models highlight suspicious areas on scans; genomic tools flag gene variants that raise risk. Results show inside a clinical dashboard with explanations and uncertainty levels. Systems follow privacy laws and go through health-authority approval. Outcome is faster, more accurate diagnosis and more personalised care.

12. **Cybersecurity Threat Detection**
    Cyber-defence engines watch network traffic and device logs for signs of attack. They learn patterns of normal behaviour and raise an alert when something looks off, like odd logins or data spikes. Integration with response playbooks lets teams isolate a machine automatically to stop spread. Dashboards trace how the attack unfolded in plain language. Benefit: fewer false alarms and quicker containment of real threats.

13. **Operational Data Analytics**
    Operational analytics focus on keeping factories, delivery fleets, or IT systems running smoothly. They analyse sensor feeds and log files to catch glitches before they cause downtime. Root-cause tools suggest the most likely fix and estimate impact if nothing is done. Alerts arrive on phones or control-room screens in near real time. Firms save money by preventing breakdowns and using resources more efficiently.

14. **Customer Experience Automation**
    Customer-experience platforms automate support across chat, email, and phone. They route questions, draft helpful replies, and predict when a customer might leave. Agents get suggested answers and next-best offers on screen, saving typing time. Sentiment tracking flags frustrated callers so a human can jump in. The goal is faster resolution and happier, loyal customers.

15. **Image & Video Recognition**

Computer-vision tools here recognise objects, scenes, and people in images or video. Factories use them to spot defects; streaming sites to moderate content; retailers to study foot traffic. Edge devices run lightweight models for low latency, while heavier analysis happens in the cloud. Active-learning workflows keep accuracy high by retraining on new examples. Payoff comes from automated monitoring and insight that was impossible at human speed.

16. **Resource & Performance Optimisation**
    Optimisation engines act like digital planners that constantly search for a better schedule or setting. They juggle many constraints – staff skills, machine capacity, delivery times – to propose the best plan at that moment. When conditions change, the engine reruns and updates the plan automatically. Interfaces let managers explore trade-offs, like cost versus speed. Benefits include higher throughput, lower energy bills, or shorter wait times.

17. **AI/ML Platform Infrastructure**
    Platform tools give data-science teams an organised way to build, track, and serve models. They store features, run training jobs, and push models to live endpoints with version control. Monitoring dashboards show drift and accuracy over time. Governance modules record who approved each model and when. The result is faster, safer deployment of machine-learning projects across a company.

18. **Data Preparation Pipelines**
    Data-preparation services clean and transform raw feeds so models can learn from them. They detect bad rows, fill gaps, standardise units, and link records that refer to the same entity. Visual lineage graphs show how each column was created. Built-in tests warn if upstream systems change and break assumptions. Outcome: higher-quality datasets delivered in hours instead of weeks.

19. **AI Solution Consulting**
    Consultancies here help clients turn ideas into working AI products. They run discovery workshops, build proof-of-concepts, and guide deployment in production. Teams mix industry specialists with data engineers to ensure solutions fit business reality. Knowledge transfer and change-management plans aim for long-term adoption. Value lies in reduced risk and faster time to measurable return on investment.

20. **Personalised Recommendation Engines**
    Recommendation engines suggest the next item a user is likely to enjoy or buy. They learn from past clicks, ratings, and viewing history to rank options for each person. Mixing in novelty ensures users discover new content, not just repeats. Live experiments track uplift in engagement or revenue. Better recommendations boost satisfaction and increase sales or watch time.

Appendix Figures and Tables

Figure A.1: Comparison of Revelio and Compustat Employment (Binscatters)

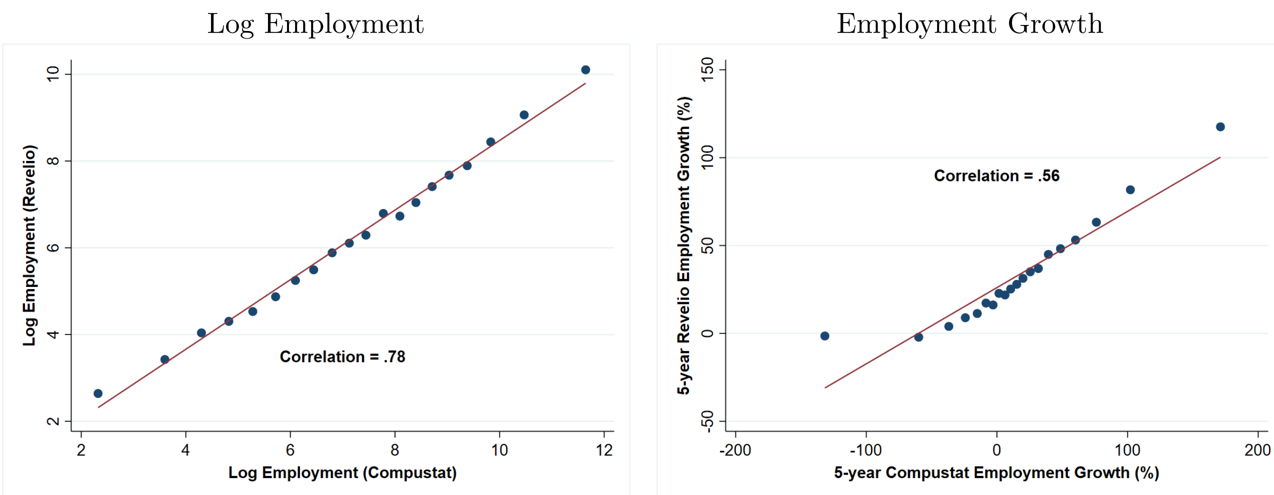

Note: This figure plots binscatters of log Revelio employment against log compustat employment (left) and 5-year Revelio employment growth against Compustat 5-year employment growth (right). Variable correlations are 0.76 (log employment) and 0.56 (employment growth). The sample spans 2014-2023.

Figure A.2: AI-related job postings versus new AI resumes (binscatters)

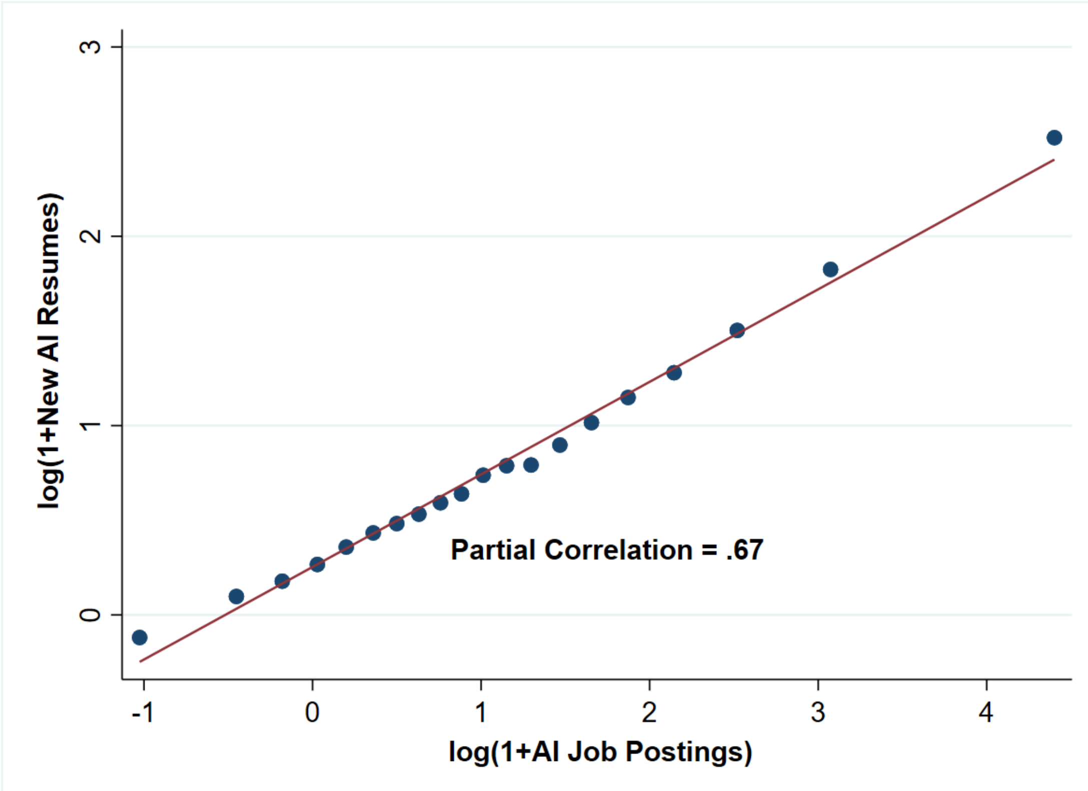

Note: This figure represents a residualized binscatter of log(1 + Newly Added AI Resumes$_{f,t}$) against log(1 + AI-Related Job Postings$_{f,t}$). A resume is tagged as a newly-added AI resume in year t if an AI position began at the firm in year t. AI-Related Job Postings$_{f,t}$ are the count of jobs posted by the firm f in year t which contain the same AI keywords used to tag AI resumes. Controls include the log of total job postings bu the firm in that year and the log of total resume-implied employment in the Revelio data. The partial correlation between the two series is 0.67. The sample spans 2014-2023.

Figure A.3: AI-related patenting versus AI resumes (binscatters)

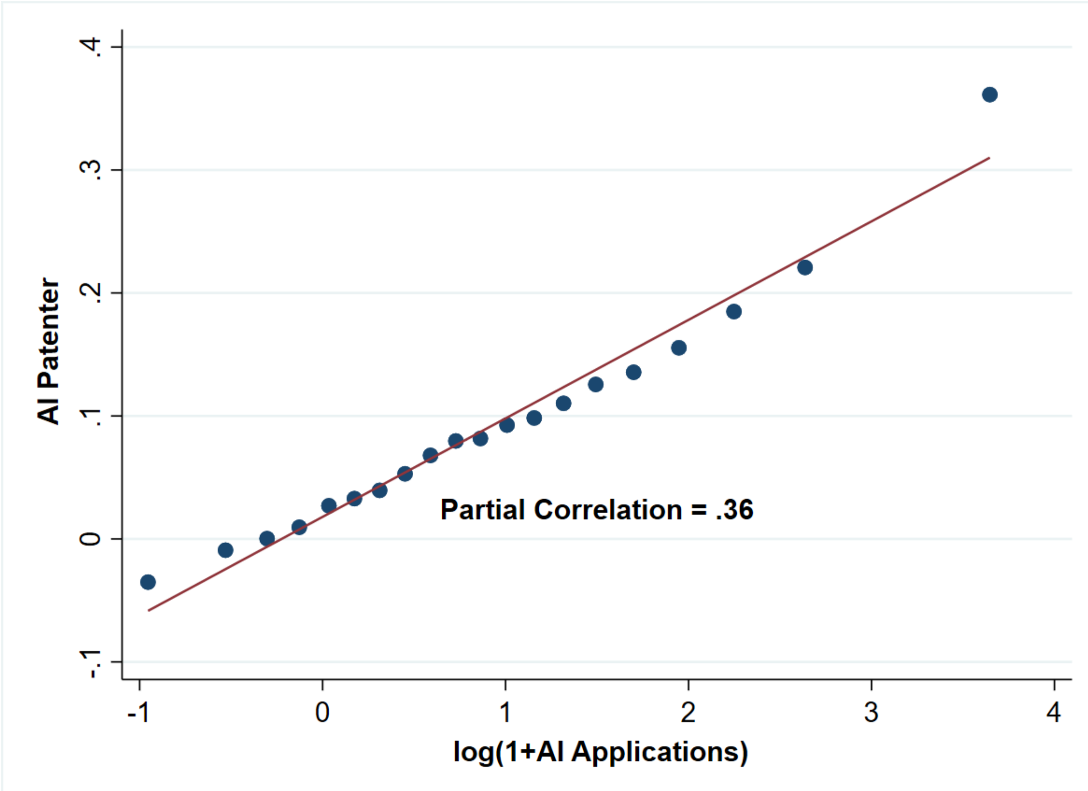

Note: This figure represents a residualized binscatter of the probability of AI-related patenting (based off the AI patent database from Pairolero et al. (2025)) against $\log(1 + \mathrm{N}_{f,t})$, the log of one plus the number of distinct AI applications at firm f in year t. Controls include the log of total resume-implied employment in the Revelio data and an indicator for non-AI patenting status. The partial correlation between the two series is 0.36. The sample spans 2014-2023.

Figure A.4: Comparison of Revelio AI-related employees with Babina et al. (2024) (binscatters)

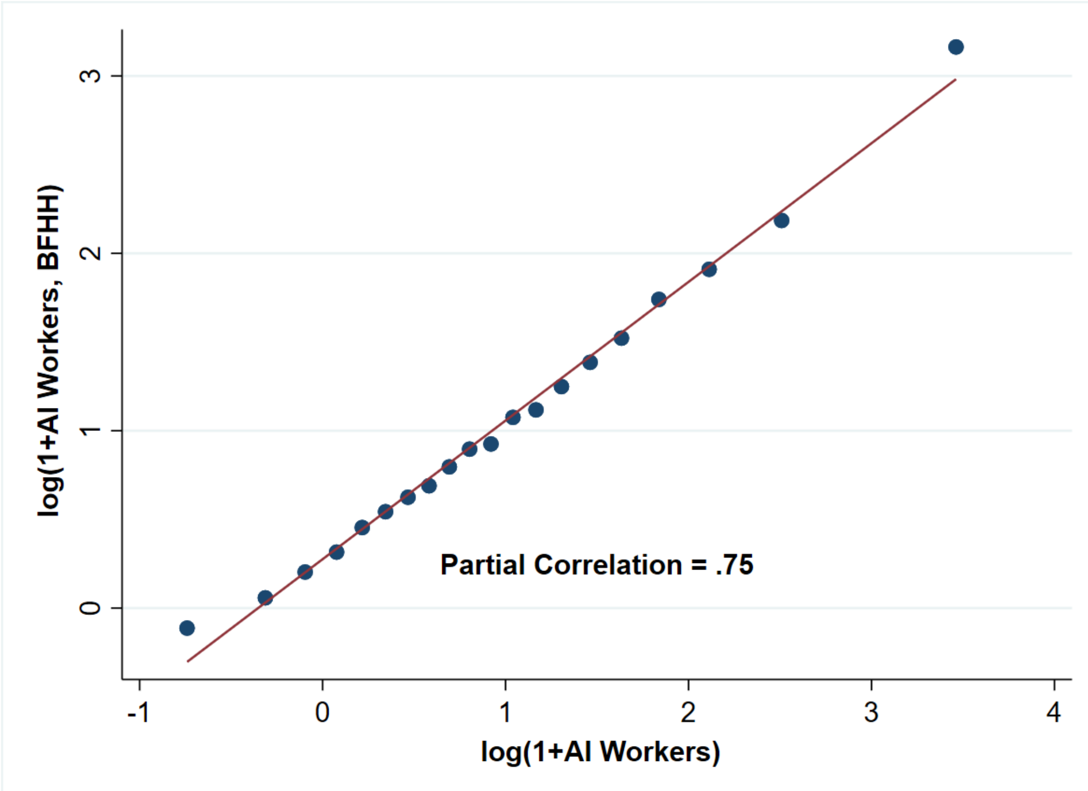

Note: This figure represents a residualized binscatter of log(1 + AI Workers $(\text{BFHH})_{f,t}$), the log of 1 plus the number of AI workers in Babina et al. (2024) against Revelio–implied AI workers log(1 + AI $\text{Workers}_{f,t}$). The figure controls for the log of Revelio resume-implied employment and Cognism resume-implied employment from Babina et al. (2024). The partial correlation between the two variables is 0.75. The sample period spans 2014-2018.

Figure A.5: AI Applications: Mean AI Exposure and Concentration

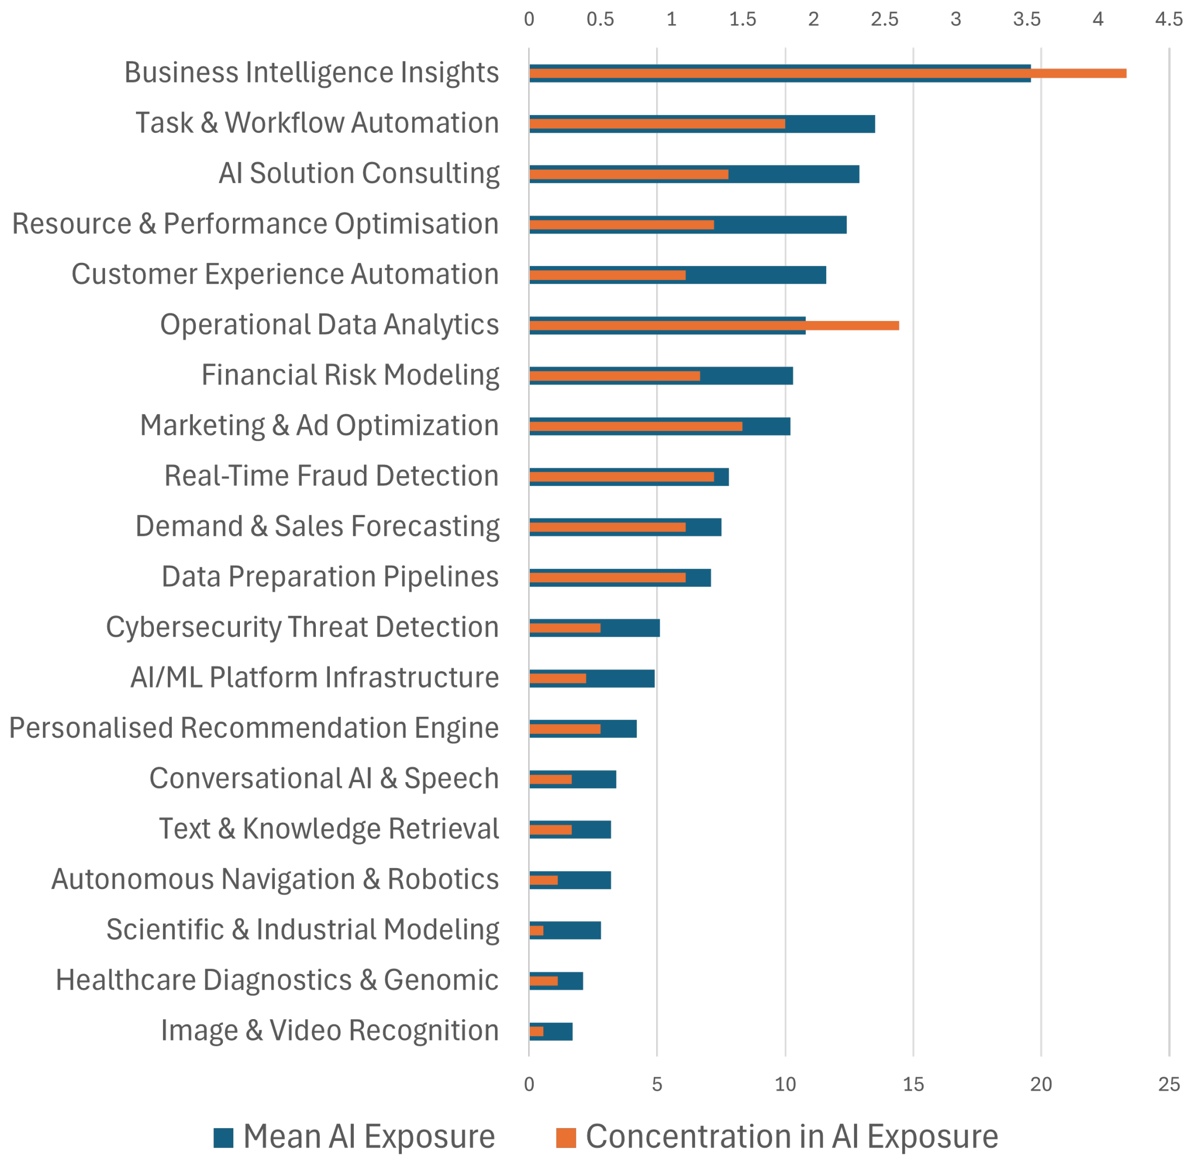

Note: The figure reports the average mean AI exposure (blue bar) and the average mean concentration in AI exposure (orange bar) across occupations (at the SOC6 level) for the 20 different clusters of AI applications. The averages are weighted by occupation employment counts.

Table A.1: Descriptive Statistics

Variable	Mean	SD	p5	p25	p50	p75	p95
5 yr % employment growth, annualized	2.8	7.7	-7.7	-1.2	2.2	6.1	16.0
# Workers	40.5	345.0	1.0	2.0	4.4	14.7	121.4
# AI Applications	368.6	1121.2	0.0	5.0	33.0	181.0	1636.0
log(1 + # AI Applications)	3.52	2.36	0.00	1.79	3.53	5.20	7.40
AI Exposure Average	0.45	0.43	0.00	0.08	0.34	0.69	1.27
AI Exposure Concentration	0.05	0.05	0.00	0.01	0.03	0.07	0.14

Figure A.6: Impact of AI on employment growth across the pay distribution (re-weighting to reflect BLS-OES 2-digit SOC shares)

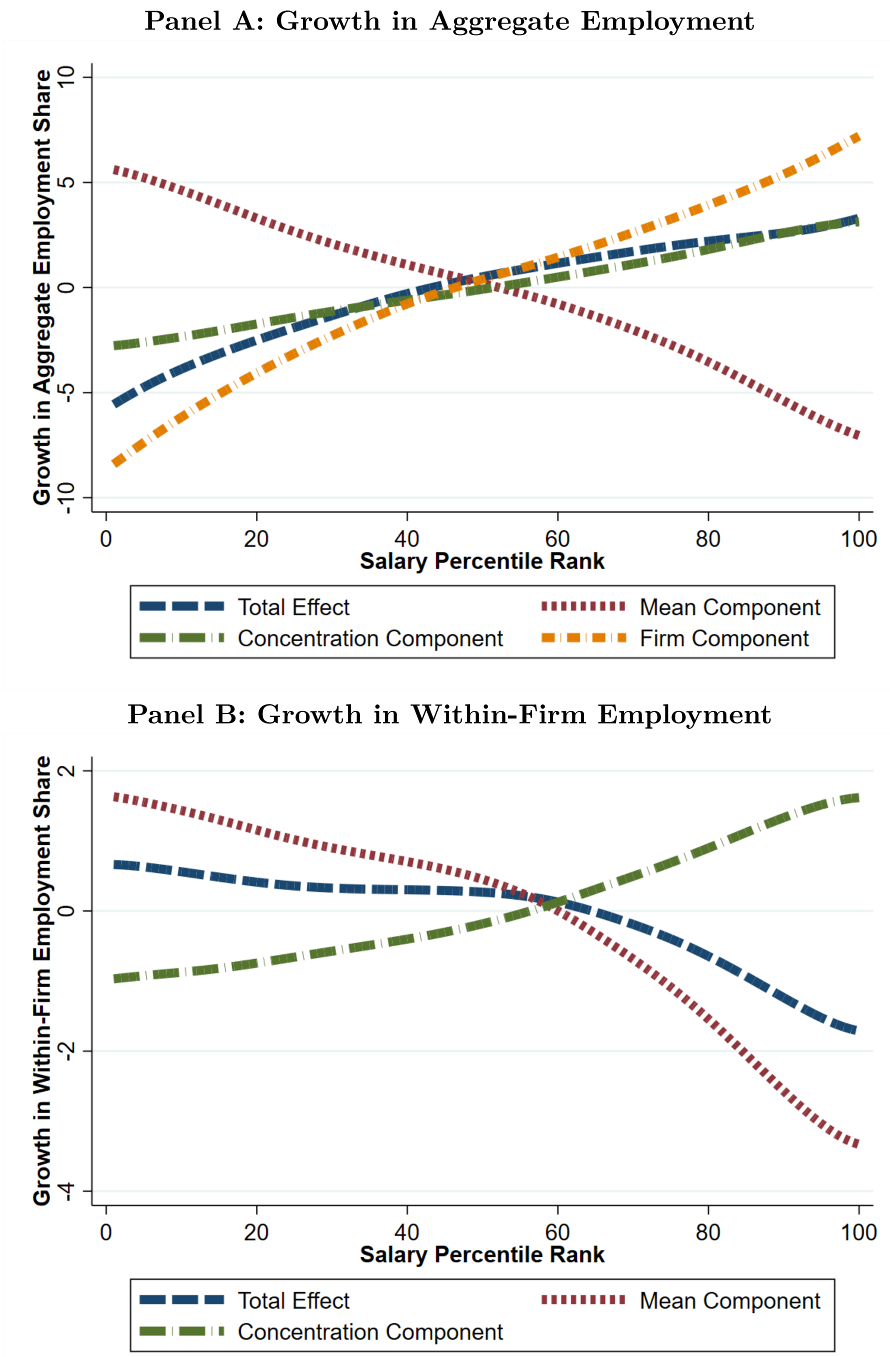

Note: This figure implements the decomposition of employment marginal effects from equation (38) in Section 3.6, except occupations within a 2-digit SOC occupation category are re-weighted to reflect their yearly employment shares in the BLS-OES data. See the notes to Figure 7 and section 3.4 of the main text for details.

Figure A.7: Impact of AI on employment growth across Occupation-Task Level AI Exposure Probability Percentile Rank

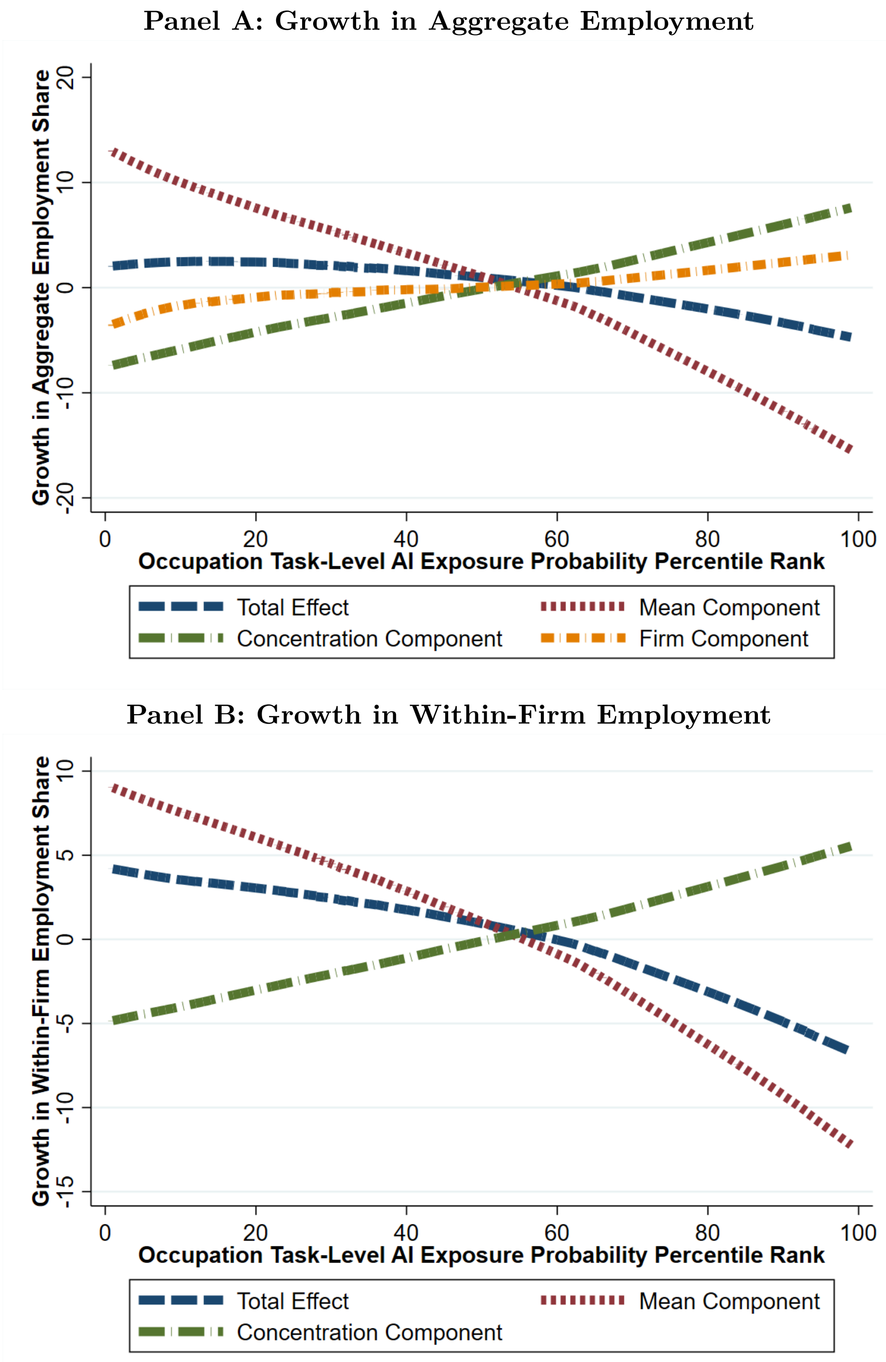

Note: This figure implements the decomposition of employment marginal effects described in section 3.6 of the main text, except we now compute average marginal effects by ranking occupations based on their cross-sectional percentile rank for average task exposure probability (defined in equation (31) of the main text).

Figure A.8: Impact of AI on employment growth across Occupation-Task Level AI Exposure Probability Percentile Rank (re-weighting to reflect BLS-OES 2-digit SOC shares)

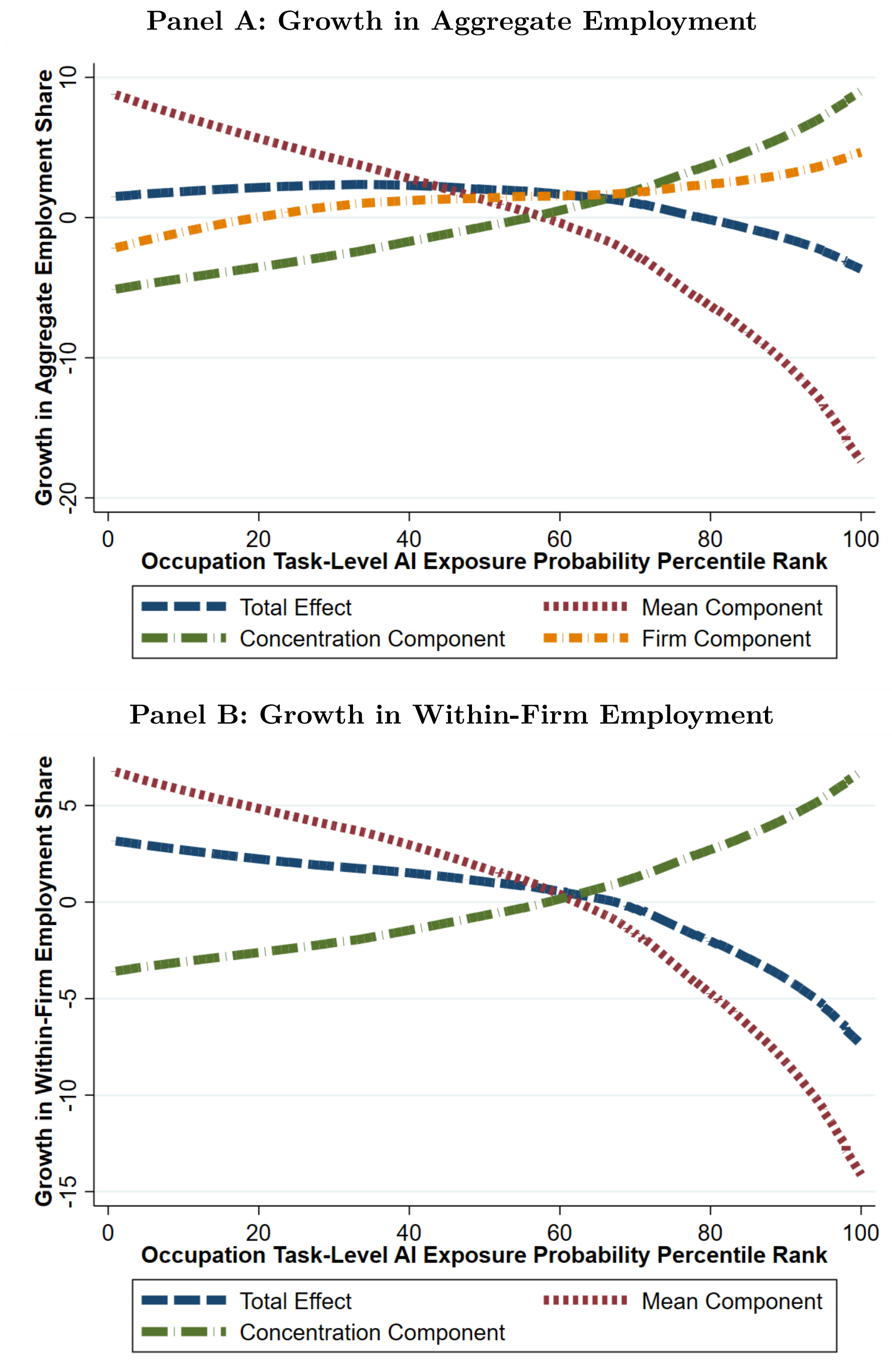

Note: This figure implements the decomposition of employment marginal effects described in section 3.6 of the main text, except we now compute average marginal effects by ranking occupations based on their cross-sectional percentile rank for average task exposure probability (defined in equation (31) of the main text), and occupations within a 2-digit SOC occupation category are re-weighted to reflect their yearly employment shares in the BLS-OES data.

Table A.2: Top and bottom 25 occupations by average AI exposure

25 Most Exposed Occupations	25 Least Exposed Occupations
Occupation	Occupation
Market Research Analysts and Marketing Specialists	Tire Builders
Management Analysts	Terrazzo Workers and Finishers
Logisticians	Tire Repairers and Changers
Computer Hardware Engineers	Tree Trimmers and Pruners
Financial Specialists	Bartenders
Computer and Information Systems Managers	Helpers–Carpenters
Sales Engineers	Dishwashers
Financial Risk Specialists	Food Preparation Workers
Transportation, Storage, and Distribution Managers	Maids and Housekeeping Cleaners
Industrial Engineers	Aircraft Service Attendants
Life, Physical, and Social Science Technicians	Animal Trainers
Aerospace Engineers	Actors
Materials Engineers	Ophthalmic Laboratory Technicians
Sales Managers	Gambling Dealers
Sales Representatives of Services	Cooks, Private Household
Credit Analysts	Janitors and Cleaners
Cost Estimators	Childcare Workers
Advertising and Promotions Managers	Food Servers, Nonrestaurant
Marketing Managers	Mechanical Door Repairers
Chemical Engineers	Cooks, Restaurant
Electrical Engineers	Judicial Law Clerks
Purchasing Agents	Insurance Appraisers, Auto Damage
Purchasing Managers	Makeup Artists, Theatrical and Performance
Production, Planning, and Expediting Clerks	Flight Attendants
Bioengineers and Biomedical Engineers	Home Health Aides

Note: This table details the top and bottom 25 occupations ranked by average AI exposure, as determined by the measurement process described in Section 2.2. These rankings are based on the computed AI exposure scores, which leverage task-level similarities between AI applications and occupational descriptions. See Section 2.2 for more details.

Table A.3: AI task Exposure and Labor Demand: OLS vs IV

Dependent Variable: 100 × 5-year Davis et al. (1996) change in share of job posting skills related to task

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Task-level AI Exposure	-4.78***	-4.75***	-4.80***	-4.71***	-4.24***	-4.82***
	(-13.46)	(-13.97)	(-14.14)	(-9.88)	(-11.07)	(-12.21)
N	13241933	13241933	13238128	13213696	13213696	13209898
R^2	0.073	0.10	0.32	0.0073	0.0033	0.0038
F-stat				13234.1	24667.0	28490.2
Task Importance Control	X	X	X	X	X	X
Mean Task Exposure Control	X	X		X	X	
Firm × Year FE	X	X		X	X	
Occ × Year FE		X			X	
Firm × Occ × Year FE			X			X

Note: This table includes estimates of Equation (5) from the main text. The unit of analysis is at the task–firm level, and the dependent variable is the Davis et al. (1996) change in the share of job posting skills demanded that are textually linked to the give task when comparing across the firms' job postings for a given occupation in the next 5 years versus the previous 5 years. The "Task-AI Exposure" is defined in equation 3 of the main text. In Columns (1)-(3), we estimate the specification using OLS, and in Columns (4)-(6), we use the instrument defined in main text Equation (4), with associated F-statistics reported below. Standard errors are clustered by occupation-firm, with associated t-statistics reported in parentheses. Besides the designated fixed and also the average task-level AI exposure of the occupation in specifications without the full complement of firm × occupation × year fixed effects. See section 1.4 of the main text for details.

Table A.4: Relevance tests for university network IV

Panel A: Lagged university × firm shares predict future university × firm shares:

	(1) Average Share (2014-2018)
Average Share (2005-2009)	0.482***
	(26.98)
N	864502
R-sq (within)	0.118
Firm FE	X
University FE	X

Variation: university × firm

Panel B: Shift-share for firm AI worker share predicts actual AI worker share:

	(1) Actual AI Worker Share
Predicted AI Worker Share	0.582***
	(7.46)
N	16560
R-sq (within)	0.0436
Revelio Emp Control	X
Ind × Year FE	X

Variation: firm × year

Note: Panel A of this table regresses the firm f average share of non-AI employment coming from university u from 2014-18 on the same shares from 2005-2009. We restrict to universities that were observable in 2005-09. Panel A also includes university and firm fixed effects; t-stats from standard errors clustered by university and firm are in parentheses. Panel B of this table regresses the actual share of AI workers at firm f in year t on the predicted AI share based off university firm hiring networks, plus controls for log resume-based employment and 3-digit NAICS industry × year fixed effects. T-stats from standard errors clustered by firm are in parentheses.

Table A.5: IV robustness tests

Panel A: Exclude top 50 AI employers/universities and tech industries

	IV (Drop Top 50 AI Firms/Universities+Tech Industry)			
	(1) Sales	(2) Emp	(3) Profit	(4) TFP
log(1 + AI uses)	8.21*	7.36**	8.17*	4.55*
	(2.56)	(3.04)	(2.48)	(2.46)
N	9458	9847	8507	4256
R-sq	0.084	0.050	0.034	0.17
Controls	X	X	X	X
Ind $\times$ Year FE	X	X	X	X

Panel B: Add shift-share controls for predicted share of employees in computer science and engineering occupations

	IV (Add shift-share controls)			
	(1) Sales	(2) Emp	(3) Profit	(4) TFP
log(1 + AI uses)	9.51***	6.54***	8.21**	7.60***
	(3.82)	(3.60)	(3.21)	(5.17)
N	12282	12688	11246	6035
R-sq	0.070	0.051	0.027	0.18
Controls	X	X	X	X
Shift-Share Controls	X	X	X	X
Ind $\times$ Year FE	X	X	X	X

Note: This table presents robustness tests for IV regressions of firm-level growth rates on AI utilization. In the top panel, we construct the 2005-2009 pre-period university $\times$ firm hiring network after dropping all workers who graduated from any of the top 50 universities by number of graduates who specialize in AI during the 2014-2018 period; we also drop the top 50 firms by total employment of AI workers during the same time interval; finally, we also exclude tech industries (2-digit NAICS = 51 or 54, or 3-digit NAICS = 334). In panel B, we add shift-share controls for the university network-predicted share of workers in computer/mathematical or engineering occupations (2-digit soc codes 15 or 17). Sample period spans 2014-2023, and we report in parentheses t-stats from standard errors clustered by firm.

Table A.6: AI exposure and occupational employment growth (IV Robustness)

Dependent variable: 100× 5-year growth rate in the occupation–firm employment

	IV (Drop Univ/Firm/Tech)				IV (Shift-Share Controls)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Exposure Average	-17.0***	-14.7***	-16.1***	-8.98***	-14.2***	-12.9***	-14.8***	-10.3***
	(-7.23)	(-8.01)	(-11.56)	(-6.54)	(-5.50)	(-5.70)	(-9.10)	(-10.16)
AI Exposure Concentration	10.2***	6.19***	7.07***	4.89**	12.7***	12.1***	13.4***	9.16***
	(6.01)	(3.94)	(5.87)	(2.95)	(5.03)	(4.95)	(7.32)	(6.77)
log(1 + AI uses)	9.56***	5.93***			17.6***	16.6***		
	(4.84)	(3.81)			(12.71)	(11.78)		
N	1636223	1636223	1635606	1635606	2017628	2017628	2016990	2016990
R^2	0.036	0.043	-0.0032	-0.0026	0.040	0.034	-0.029	-0.011
F-stat (AI Exposure Average)	1058.5	1143.0	2456.7	1335.2	465.0	453.9	1136.7	1394.9
F-stat (AI Exposure Concentration)	821.9	913.1	2293.9	666.6	157.4	123.3	332.6	501.0
F-stat (log(1 + AI uses))	915.8	905.7			2074.3	2196.1		
Controls	X	X	X	X	X	X	X	X
Year FE	X				X			
Industry × Year FE		X				X		
Firm × Year FE			X	X			X	X
Occ × Year FE				X				X
Drop Firm/Univ/Tech	X	X	X	X				
Shift-Share Controls					X	X	X	X

Note: This table shows 2SLS regression estimates of Equation (16) from the main text. In columns (1) through (4), we repeat the same specifications as in the final 4 columns of Table 5, except we construct the 2005-2009 pre-period university × firm hiring network after dropping all workers who graduated from any of the top 50 universities by number of graduates who specialize in AI during the 2014-2018 period; we also drop the top 50 firms by total employment of AI workers during the same time interval; finally, we also exclude tech industries (2-digit NAICS = 51 or 54, or 3-digit NAICS = 334). In columns (5) through (8), we again repeat the specifications in the final three columns of Table 5, except we add shift-share controls for the university network-predicted share of workers in computer/mathematical or engineering occupations (2-digit soc codes 15 or 17), both by themselves and also interacted with the occupational means and variances of task exposure $\mu_{o,t}$ and $\sigma^2_{o,t}$ (respectively defined in equations (31) and (32) of the main text). Otherwise, specifications are the exact same as the final 4 columns of Table 5; see notes under that table for further details. Observations are weighted by the yearly occupation-firm cell's share of employment. T-stats based on standard errors clustered by occupation–firm are in parentheses.

Table A.7: Impact of AI on relative employment growth by occupation group (re-weighting to reflect BLS-OES 2-digit SOC shares)

	2-digit SOC	Mean Component	Concentration Component	Firm Component	Total	% of Emp
Management	11	-7.09	4.08	3.09	0.078	5.28
Business and Financial	13	-15.0	8.50	4.33	-2.20	5.38
Architecture and Engineering	17	-11.2	5.14	4.27	-1.81	1.82
Science	19	-3.93	2.66	3.41	2.14	0.82
Community and Social Service	21	5.69	-2.59	1.79	4.89	1.47
Legal	23	4.98	-3.10	5.07	6.95	0.79
Education and Library	25	4.16	-1.99	2.31	4.48	6.34
Arts, Entertainment, Media	27	2.63	-1.82	5.08	5.88	1.35
Healthcare Practitioners	29	0.47	0.47	0.40	1.35	6.02
Healthcare Support	31	2.42	-0.67	0.45	2.21	2.95
Protective Service	33	4.40	-2.72	-0.11	1.57	2.46
Food Preparation and Serving	35	7.20	-3.49	-8.76	-5.05	9.58
Cleaning and Maintenance	37	8.97	-5.26	-2.77	0.94	3.28
Personal Care and Service	39	7.06	-3.50	-2.97	0.60	3.46
Sales and Related	41	-3.31	1.92	-0.57	-1.96	10.8
Office and Administrative	43	-2.30	0.50	2.11	0.31	16.2
Farming, Fishing, and Forestry	45	7.77	-4.32	-2.80	0.65	0.33
Construction and Extraction	47	1.39	-1.42	1.87	1.84	4.13
Installation and Repair	49	-0.92	-0.48	1.12	-0.28	4.00
Production	51	0.81	0.19	-0.77	0.24	6.55
Transportation	53	2.67	-1.38	-1.81	-0.51	7.04

Note: This table shows results from estimating the decomposition (38) for broad 2-digit SOC occupation groups, except occupations within a 2-digit SOC occupation category are re-weighted to reflect their yearly employment shares in the BLS-OES data. See notes to Table 6 and Section 3.4 of the text for further details.

Table A.8: Changes in Demand for 5 Most Exposed Tasks for Selected Occupations

Occupation	Task	Change
Industrial Production Managers	• Review plans and confer with research or support staff to develop new products or processes.	-23.7
	• Develop or implement production tracking or quality control systems, analyzing production, quality control, maintenance, or other operational reports to detect production problems.	-18.1
	• Prepare reports on operations and system productivity or efficiency.	-13.2
	• Monitor development of new products to help identify possible problems for mass production.	-11.3
	• Collect and analyze production samples to evaluate quality.	-4.6
Claims Adjusters, Appraisers, Examiners, and Investigators	• Investigate, evaluate, and settle claims, applying technical knowledge and human relations skills to effect fair and prompt disposal of cases and to contribute to a reduced loss ratio.	-33.6
	• Conduct detailed bill reviews to implement sound litigation management and expense control.	-30.5
	• Prepare reports to be submitted to company's data processing department.	-14.6
	• Verify and analyze data used in settling claims to ensure that claims are valid and that settlements are made according to company practices and procedures.	-7.6
	• Analyze information gathered by investigation and report findings and recommendations.	-7.0
Human Resources Workers	• Advise management on organizing, preparing, or implementing recruiting or retention programs.	-40.5
	• Prepare or maintain employment records related to events, such as hiring, termination, leaves, transfers, or promotions, using human resources management system software.	-28.8
	• Evaluate selection or testing techniques by conducting research or follow-up activities and conferring with management or supervisory personnel.	-14.7
	• Analyze employment-related data and prepare required reports.	-11.1
	• Evaluate recruitment or selection criteria to ensure conformance to professional, statistical, or testing standards, recommending revisions, as needed.	-1.1
Accountants and Auditors	• Develop, implement, modify, and document recordkeeping and accounting systems, making use of current computer technology.	-19.2
	• Develop, maintain, or analyze budgets, preparing periodic reports that compare budgeted costs to actual costs.	-17.0
	• Collect and analyze data to detect deficient controls, duplicated effort, extravagance, fraud, or non-compliance with laws, regulations, and management policies.	-13.8
	• Examine and evaluate financial and information systems, recommending controls to ensure system reliability and data integrity.	-7.8
	• Analyze business operations, trends, costs, revenues, financial commitments, and obligations to project future revenues and expenses or to provide advice.	-9.0
Credit Analysts	• Consult with customers to resolve complaints and verify financial and credit transactions.	-40.4
	• Analyze financial data, such as income growth, quality of management, and market share to determine expected profitability of loans.	-14.1
	• Generate financial ratios, using computer programs, to evaluate customers' financial status.	-1.3
	• Review individual or commercial customer files to identify and select delinquent accounts for collection.	0.9

	• Evaluate customer records and recommend payment plans, based on earnings, savings data, payment history, and purchase activity.	-1.6
Financial and Investment Analysts	• Analyze financial or operational performance of companies facing financial difficulties to identify or recommend remedies.	-11.2
	• Monitor developments in the fields of industrial technology, business, finance, and economic theory.	-9.9
	• Inform investment decisions by analyzing financial information to forecast business, industry, or economic conditions.	-7.2
	• Evaluate capital needs of clients and assess market conditions to inform structuring of financial packages.	-8.4
	• Prepare plans of action for investment, using financial analyses.	-8.3
Bookkeeping, Accounting, and Auditing Clerks	• Prepare purchase orders and expense reports.	-34.0
	• Operate computers programmed with accounting software to record, store, and analyze information.	-31.1
	• Classify, record, and summarize numerical and financial data to compile and keep financial records, using journals and ledgers or computers.	-27.0
	• Compile statistical, financial, accounting, or auditing reports and tables pertaining to such matters as cash receipts, expenditures, accounts payable and receivable, and profits and losses.	-22.2
	• Compile budget data and documents, based on estimated revenues and expenses and previous budgets.	-13.7
Insurance Claims and Policy Processing Clerks	• Organize or work with detailed office or warehouse records, using computers to enter, access, search or retrieve data.	-28.0
	• Transcribe data to worksheets, and enter data into computer for use in preparing documents and adjusting accounts.	-26.5
	• Review and verify data, such as age, name, address, and principal sum and value of property, on insurance applications and policies.	-20.1
	• Process and record new insurance policies and claims.	-9.2
	• Enter insurance- and claims-related information into database systems.	-7.3

Note: Table lists the five examples of the most exposed tasks for selected occupations along with the averages in the subsequent changes in demand for skills related to these tasks over 5 year intervals in our sample period defined in equation (5). For each of the selected occupations, we isolate the five most AI-exposed tasks, then report a measure of the dependent variable from equation (5) which is residualized with respect to the controls and the occupation-firm fixed effects then averaged (weighted by employment in each firm-occupation-year cell) across firm-years. As a result of this residualization, this variable averages to zero across all tasks within each firm-occupation-year cell. We perform an analogous construction for residualized task-level exposure in order to identify the 5 most exposed tasks for each of these occupations.