

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

SEPTEMBER 15, 2025

Abstract

We use advances in natural language processing to construct new measures of workers' task-level exposure to artificial intelligence (AI) and machine learning from 2010 to 2023, capturing variation across firms, occupations, and time. Tasks with higher AI exposure subsequently experience reduced labor demand. To interpret these patterns, we develop a model that separates direct substitution from indirect reallocative effects of labor-saving technologies. Two variables summarize the impact of AI on within-firm labor demand: the mean exposure of an occupation's tasks, which depresses demand, and the concentration of exposure in a few tasks, which offsets losses by enabling workers to reallocate effort. Using an instrument based on historical university hiring networks, we find causal evidence consistent with these predictions. Despite strong substitution at the task level, overall employment effects are modest, as reduced demand in exposed occupations is offset by productivity-driven increases in labor demand at AI-adopting firms.

*We are grateful to Matthew Akuzawa, 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, Pascual Restrepo (discussant), 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 revived 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). Yet, 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 its 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 (the AI integrators). Using LLMs, we extract detailed information on how these workers apply AI to their firms. While this approach relies on resume data and therefore misses cases where firms outsource AI development, it closely tracks survey-based measures of adoption such as the BTOS.³ Consistent with Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas (2023a), we find that firms with high AI utilization are 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 step 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 specification with a rich set of fixed effects: our most demanding specification includes the

¹For 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.

²See 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.

³Using the Annual Business Survey, Chequer, Herkenhoff, Papanikolaou, Schmidt, and Seegmiller (2025) report that among intensive AI users, 49% of firms conduct AI R&D in-house, rising to 77% when weighting by employment.

interaction of firm by occupation by year and task by year fixed effects, so identification comes from comparing two tasks with different AI exposure performed by the same occupation in the same firm and year, relative to their economy-wide averages. We find that a one standard deviation increase in *task*-level AI exposure in a given firm reduces by 2% the relative demand for skills related to that task. 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 equally exposed to AI. On average, tasks performed by higher-paid occupations (those near the 90th percentile of the wage distribution) are more exposed. Yet substitution at the task level does not mechanically imply reduced labor demand at the occupation level. Two countervailing forces can offset direct substitution. First, as AI takes over some tasks, workers may reallocate effort toward tasks that remain difficult to automate, raising their productivity. Second, if AI adoption boosts firm productivity and growth, firms may expand employment overall, including in occupations with high exposure. These mechanisms highlight why task-level substitution need not translate into lower aggregate 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 the automated filing of expense reports. Crucially, workers optimally allocate their time across tasks. Improvements in labor-saving 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 key 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.⁴ In addition, there is a third channel affecting labor demand across firms that hinges on 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 and growth. 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 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 and the intensity of firm-level spillovers on labor demand. 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 random across firms or occupations. At the firm level, adoption costs may be correlated with underlying growth potential, biasing OLS estimates of the effect of AI on productivity. At the occupation level, labor-saving technologies may be selectively developed for jobs in short supply, creating bias in estimates of the effect of AI exposure on labor demand. To address these concerns, we build on [Babina, Fedyk, He, and Hodson \(2024\)](#) and instrument for a firm’s employment of AI-focused workers with the growth in AI graduates from the universities from which the firm historically hires. The key variation in this strategy comes from pre-existing variation in university hiring networks: firms historically tied to universities whose graduates later enter AI-related occupations face lower adoption costs as they gain exposure to supply-side increases in AI-trained labor. For example, if State Street hires disproportionately from Boston University while BNY Mellon hires more from Wharton, and Boston University produces more AI graduates than Wharton, then the relative cost of adopting AI will fall more for State Street than for BNY Mellon.

Using our instrument, we find that adoption of AI by firms leads to higher growth in revenue, measured productivity, profits, and employment. Notably, our IV estimates are about 30 percent larger than OLS estimates, consistent with selective adoption by large, profitable firms that typically grow more slowly than smaller peers.

We next turn to the demand for workers in affected firms and occupations. To instrument for within-firm variation in the direction of AI adoption, we interact the average exposure of each occupation across all other firms (in the same period) with the predicted firm-level intensity of AI adoption from the university-based instrument. The granularity of our exposure measures allows us to saturate the specification with a rich set of controls, including firm by year and occupation by year

⁴These two sufficient statistics emerge naturally in our framework, but we also show they arise in the [Acemoglu and Restrepo \(2018\)](#) model, underscoring their broader applicability.

fixed effects. Identification therefore comes from comparing differentially exposed occupations within the same firm as well as the same occupation across firms with varying direction and intensity of AI adoption. Consistent with the model, we find that higher mean exposure of an occupation’s tasks to employer-adopted AI applications reduces subsequent employment growth in that occupation–firm pair, while greater concentration of exposure raises it. The IV results mirror the OLS findings qualitatively, but are larger in magnitude: in our preferred specification, a one–standard deviation increase in mean task exposure reduces the occupation’s within-firm employment share by about 14.5 percent over five years, while a similar increase in concentration raises it by about 7.5 percent. Finally, replacing firm–year fixed effects with industry–year fixed effects, we find a strong link in both OLS and IV specifications between overall firm AI adoption intensity and occupational employment growth, consistent with AI-induced firm growth increasing labor demand.

In brief, our empirical results reveal an economically significant effect of AI on firm growth, productivity and labor demand, but also evidence for within firm labor reallocation. These within-firm reallocation patterns 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 across different occupations, job types, and worker earnings levels.

Overall, we find that the overall impact of AI on the composition of labor is limited due to the presence of countervailing forces. Even though the direct substitution effect (mean task exposure) is quantitatively strong, it is counterbalanced by labor-augmenting effects from task reallocation (the concentration of AI task exposures) and by AI-driven increases in firm labor demand. Although these forces differ somewhat across the pay distribution, their net effects are more uniform than one might expect. Substitution effects are strongest for higher-paid occupations, but so too are reallocative gains from concentrated task exposure. As a result, within firms, employment in highly exposed occupations (90th percentile of the pay distribution) falls by about 3.1% relative to the least exposed (bottom percentile). After taking into account the impact of firm growth on labor demand, this effect is mildly reversed, since the jobs that are more exposed to AI are more prevalent in firms that adopt AI and grow faster. 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 overall estimates are muted does not necessarily imply there are no winners and losers among occupations. The most adversely affected occupations are in business, financial, and engineering fields, which see declines of nearly 2% in their employment share over five years.

These losses reflect strong substitution at the task level, only partly offset by complementarities and firm-level increases in labor demand. At the same time, even occupations with little direct task exposure to AI—such as food preparation and serving—also experience employment share declines, since their employers are less likely to adopt AI and thus grow more slowly than AI-adopting firms.

In sum, our results help explain why the aggregate impact of AI on labor reallocation may be hard to detect, even though we find 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 predicted to realized changes, we estimate that about 14% of the variation in occupational employment share growth can be attributed to their AI exposure during our sample period—roughly half from within-firm reallocation away from exposed tasks, and half from differences in average AI use across occupations’ employers.

Our paper contributes to a growing literature on the labor market consequences of technological change, and especially of AI. Closest to us is work linking AI adoption to firm outcomes and labor demand. [Acemoglu et al. \(2022\)](#) find evidence of labor substitution at the establishment level but little impact on employment or wages in exposed occupations. [Babina et al. \(2024\)](#) show that AI adoption correlates with firm growth but does not generate measurable productivity gains or automation. [Acemoglu et al. \(2023a\)](#) emphasize that the positive correlation between AI adoption and firm growth largely reflects selection into which firms adopt advanced technologies. [Gathmann, Grimm, and Winkler \(2024\)](#) contrast AI with robots, arguing that AI reduces demand for abstract tasks while raising demand for certain routine ones. Closest in spirit to our approach, [Aghion, Bunel, Jaravel, Timo Mikaelson, and Sogaard \(ming\)](#) show that the effects of AI on labor demand depend not just on exposure, but also on the type of AI technology. Related work includes [Jiang, Park, Xiao, and Zhang \(2025\)](#), who find that AI use lengthens working hours, and [Aum and Shin \(2025\)](#), who document declines in labor demand for skilled workers in Korea following digital technology investments, including AI.

Our measure of AI adoption ends in 2023 and therefore largely excludes the recent rise of generative AI (GenAI). Despite the short time since GenAI became broadly available in late 2022, a growing literature has begun to study its effects on firms and workers. [Eloundou et al. \(2023\)](#) construct an occupation-level exposure measure and show that most occupations are significantly exposed, with high-wage jobs more affected than low- or middle-wage ones—unlike prior waves of automation. Building on this measure, [Eisfeldt, Schubert, Taska, and Zhang \(2023\)](#) develop a firm-level proxy for workforce exposure to GenAI, arguing that adoption has improved profitability by reducing labor costs. In contrast, [Auer, Köpfer, and Sveda \(2024\)](#) suggest that GenAI may complement high-wage workers and that displacement risks fall more heavily on low-wage jobs. At the worker level, [Humlum and Vestergaard \(2024\)](#) document widespread but uneven GenAI

adoption, concentrated among younger, higher-achieving, and male workers. Finally, using a version of Hulten’s theorem, [Acemoglu \(2024\)](#) argues that GenAI’s aggregate productivity effects will likely be modest. Even though GenAI differs from the first wave of AI technologies in the types of tasks it can perform, we expect the same economic forces we uncover—direct substitution at the task level, reallocation across tasks, and firm-level growth effects—to continue to shape the overall composition of the labor force.

Compared to prior work, our study brings a distinct measurement contribution. Our approach combines firm-level adoption data with occupation–task detail, producing direct, theory-grounded measures of AI exposure. These measures not only capture the direct substitution effect of AI but also identify reallocative effects across tasks. While [Babina et al. \(2024\)](#); [Babina, Fedyk, He, and Hodson \(2023\)](#) also use online resumes to track AI adoption, our work goes further by mapping each AI application adopted by a firm to the specific tasks its occupations perform. This level of granularity moves beyond occupation-level exposure measures ([Webb, 2020](#); [Felten, Raj, and Seamans, 2018](#); [Brynjolfsson, Mitchell, and Rock, 2018](#); [Eloundou et al., 2023](#); [Eisfeldt et al., 2023](#)), which abstract from within-occupation heterogeneity, and beyond job-posting studies ([Acemoglu et al., 2022](#); [Eisfeldt et al., 2023](#)), which capture hiring intentions rather than realized outcomes. By focusing on realized employment changes at the firm–occupation–task level, our analysis helps explain why aggregate effects appear muted even as substitution is evident at the micro level.

More broadly, our work connects to the literature on the labor market effects of labor-substituting technologies. One important stream emphasizes how technology substitutes for routine tasks, making an occupation’s routine task share a key predictor of its exposure to labor-saving innovations over the past several decades ([Autor et al., 2006](#); [Acemoglu and Autor, 2011](#); [Goos, Manning, and Salomons, 2014](#)). A second stream focuses on constructing direct measures of specific labor-saving technologies—such as robots, software, or AI—and tracing their effects on employment and wages ([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](#)). Our paper contributes to both streams by combining a task-based framework with direct measures of AI adoption, allowing us to disentangle substitution effects at the task level from reallocative and productivity-driven effects at the firm level.

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 integrators to identify which applications are developed by specific firms.⁵ 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 data source is Revelio Labs, a workforce analytics provider that compiles a near-universe of LinkedIn profiles and job postings. The resume data contain rich information on workers' educational and employment histories, including universities attended, fields of study, employers, job titles, employment dates, and self-reported job descriptions. The job postings data include firm identifiers, occupational and industry codes (SOC, NAICS), posting and removal dates, seniority, salary (or imputed salary when missing), job location (state, city, MSA, ZIP), full posting text, and—when the firm is public—stock market identifiers such as CUSIP.

Our main analysis covers the period from 2014 to 2023, a period that combines improved resume coverage with the rapid spread of ‘first generation’ AI in the workplace. When constructing our shift-share instrument in Section 3.1, we use the profiles back to 2005. When 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. 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. For the job postings data, we include records from 2010 to 2023, leveraging earlier years to construct our measures of task reallocation in Section 1.4. To reduce the raw dataset of over a billion postings, we randomly sample up to ten postings per year for each occupation–firm pair. We restrict our analysis to publicly traded companies in Compustat which includes data on firm outcomes. For positions of non-AI workers, we require a valid occupation identifier. After applying these restrictions, our dataset includes approximately 58 million LinkedIn profiles and 14 million job postings that are linked to

⁵We use the term ‘AI Integrator’ to refer to specialized workers whose job is to adapt or develop customized AI applications for their firms. Our view aligns with ChatGPT5, which provides the following definition: “The AI Integrator bridges the gap between cutting-edge artificial intelligence technologies and real-world organizational processes. This role focuses on identifying where AI tools can be deployed effectively, adapting them to existing workflows, and ensuring smooth adoption across teams. The AI Integrator combines technical fluency in AI systems with a strong understanding of business needs, user behavior, and change management.”

Compustat firms.

We measure employment at the occupation–firm level using Revelio counts, focusing on active positions and excluding those tagged as AI integrators. A position is considered active in a given year if the recorded start and end dates indicate at least six months of employment during that year. Aggregated to the firm–year level, Revelio-based employment closely tracks Compustat reports. Appendix Figure A.1 shows 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 correlations, 0.76 in levels and 0.56 in five-year changes, indicate that resume data provide a reliable proxy for firm-level employment dynamics.

An important caveat of the Revelio resume data is that they are not a random sample of the U.S. workforce but are drawn from self-reported, publicly posted resumes. Hence, the resume data tends to overrepresent professional occupations, younger cohorts, and digitally engaged workers. Even so, several studies suggest the data provide broad and reliable coverage of the white-collar labor market. [Tambe \(2025\)](#) shows that education, occupation, and age distributions in Revelio closely track the CPS and ACS, particularly for managerial and professional roles. Likewise, [Hershbein and Kahn \(2018\)](#) validate online postings data (from Lightcast, formerly Burning Glass), showing that occupation- and industry-level trends align with the BLS’s Occupational Employment Statistics. Taken together, these findings indicate that online labor market data—while subject to selection biases—are a credible source for analyzing structural change. In fact, Revelio’s overrepresentation of high-skill, digitally intensive sectors such as professional services, tech, and finance makes it especially well-suited to studying the labor market impact of AI, which is disproportionately concentrated in these areas.

1.2 Extracting AI Applications

Our primary source for measuring firm AI adoption is job descriptions extracted from the resumes of AI integrators. Revelio provides especially strong coverage of this population, since employees involved in developing and deploying AI have clear incentives to highlight their expertise and projects on online resumes in order to attract new opportunities.⁶ Because deploying AI typically requires in-house investment, these detailed descriptions provide a unique window into both the level of adoption and the direction of firms’ AI investments. They also capture adoption activities invisible in patent data—another common textual source in the literature—since many AI applications combine relatively standardized algorithms with proprietary data and are thus unlikely to be patentable.

⁶For example, as of mid-2023, [LinkedIn’s website](#) reported that, on LinkedIn, 117 job applications were submitted every second, 8 people were hired every minute, and that “45% of hirers on LinkedIn explicitly use[d] skills data to fill their roles.” Recruiting blogs likewise emphasize that complete LinkedIn profiles improve job search outcomes by increasing engagement from recruiters.

Methodology

To measure firm adoption of AI applications, we analyze job descriptions where workers explicitly report involvement in AI development. We summarize the procedure here and provide full details in Appendix B.1. In the first step, we scan job descriptions in the Revelio resumes for AI-related keywords to flag positions where AI may be adopted. We then use a large language model (Llama 3.1 70B) to extract and refine phrases that describe how AI is used. The model first identifies and cleans relevant phrases, then filters out references to specific AI tools and removes vague mentions that lack a concrete use case. This process yields more than 1 million distinct AI use cases from roughly 500,000 job positions. We classify an AI application as active in firm f during year t if the associated resume position was active for at least six months in that year.⁷

Figure 1 illustrates our procedure with an example from the resume of a JP Morgan employee in an AI-implementing role. The worker described 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.

Our process first flags this resume as belonging to an AI integrator by identifying the term “AI/ML” as AI-related. Next, the LLM extracts specific phrases describing AI use, including “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.” The LLM queries clean and standardize these phrases into concrete AI applications: (i) “Forecast risk and fraud in various lending businesses, including auto, card, and home lending,” and (ii) “Assess credit risk and provide regulatory compliance across different lines of business.”

Validation

A potential limitation of our approach is that it only captures AI applications developed internally. Yet this is not a narrow margin: using data from the Annual Business Survey, [Chequer et al. \(2025\)](#) report that among intensive AI users, 49% conduct R&D in-house, rising to 77% when weighted by employment. Similarly, [Bonney, Breaux, Buffington, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, and Savage \(2024\)](#) find that in the BTOS survey, about half of firms (employment-weighted) trained existing staff to use AI in the past six months, and about 10% hired AI-trained workers. To assess the validity of our adoption measure, we benchmark it against three independent sources:

⁷As a validation check, we benchmark our firm-level AI integrator counts against those of [Babina et al. \(2024\)](#), who use Cognism resumes. Despite differences in data sources and methodology, the firm-level measures are highly correlated.

survey-reported adoption, firms’ hiring demand, and AI patenting. Across all three, our measure closely aligns with these alternative sources.

First, we compare adoption rates from our resume-based measure with survey-reported AI utilization in the Census [Business Trends and Outlook Survey \(BTOS\)](#), which tracks firm AI use in real time ([Bonney et al., 2024](#)). The BTOS reports sector (2-digit NAICS) by firm size adoption rates.⁸ To align with their data, we calculate the probability that a firm in a given sector–size cell has at least one AI worker resume in Revelio as of end-2023. For comparability, we use all Revelio firms (not just Compustat) and, consistent with the BTOS frame of single-unit employers, restrict the top size bin to firms with fewer than 1,000 employees. Panel A of [Figure 2](#) plots BTOS-reported adoption rates against resume-based adoption rates in log scale to highlight variation among low-adoption cells. The correlation between the two measures is strikingly high (0.9 in levels), indicating that our resume-based measure closely tracks survey-reported patterns of AI use.

Second, we benchmark our AI adoption measure against firms’ stated demand for AI workers. [Appendix Figure A.2](#) compares the number of new AI resumes at firm f in year t with the number of AI-related job postings in the same year. We tag postings using the same keyword filters applied to resumes of AI integrators. The figure plots new AI resumes against AI postings, residualized on total resume-based employment, total postings, and year fixed effects. The relationship is strong: the partial correlation is 0.67, indicating that periods when workers report starting new AI positions align closely with periods when firms advertise AI-related jobs.

Last, we explore whether firms that employ AI integrators also patent in AI. Using the [USPTO AI patent database](#) from [Pairolero, Giczy, Torres, et al. \(2025\)](#), [Appendix Figure A.3](#) shows a positive association: firms with higher AI adoption using our measure are more likely to patent in AI, with a partial correlation of 0.36 after controlling for employment and overall patenting. While not all AI adopters file patents each year—as expected—the overlap underscores that our measure tracks a genuine focus towards AI development and adoption.

Which firms adopt AI?

We next examine how our AI adoption measure correlates with firm characteristics over the 2014–2023 period. [Panel B of Figure 2](#) shows that information and professional services have the highest AI utilization rates, followed by finance and insurance, while accommodation, food services, and other services have the lowest. Within each major NAICS sector, AI adoption rises monotonically with firm size.

⁸We 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 \times 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.

Table 1 relates the number of AI applications at firm f in year t to log sales per worker, log sales, log profits (sales minus cost of goods sold), log TFP (estimated following [İmrohoroğlu and Şelale Tüzel \(2014\)](#)), and log average salary from Revelio-imputed compensation. Each specification includes industry (3-digit NAICS) interacted with calendar year fixed effects and controls for the log number of Revelio resumes to address the mechanical link between resume coverage and observed AI applications. The results show a strong positive association between AI adoption and firm performance: a one-standard deviation increase in AI utilization is associated with about 12% higher productivity (sales per worker and TFP), 31% higher sales, 42% higher profits, and 11% higher average pay.

In brief, we see that AI-using firms are larger, more productive, and higher paying than their peers, consistent with survey evidence on advanced technology adoption ([Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo, and Zolas, 2023b](#); [McElheran, Li, Brynjolfsson, Kroff, Dinlersoz, Foster, and Zolas, 2023](#)). While suggestive of economies of scale or the advantages of large data endowments, these correlations may also reflect selection into which firms adopt AI—hence, our empirical design will need to leverage plausibly exogenous variation in AI adoption.

Heterogeneity in AI applications across firms and sectors

A unique advantage of our resume-based approach in measuring AI adoption is that it provides a highly granular view of the direction of AI investments. Much like patents reveal detailed information about new products or processes, employees’ job descriptions offer insight into the different ways firms apply AI.

To illustrate the richness of our measure, we first compare two large publicly traded firms in our sample: JPMorgan Chase (JPMC) and Walmart. Because both report a large number of AI applications, we cluster them using k -means (with $k = 5$) and assign labels with a large language model. The left panel of Table 2 shows several of these clusters. At JPMC, prominent applications include “Predictive Modeling and Financial Forecasting,” relevant for underwriting and asset management, and “Customer Engagement and Personalization,” aimed at improving client service. At Walmart, applications focus instead on “Forecasting, Pricing, and Supply Chain Optimization” and “Process Automation and Operational Efficiency,” reflecting its retail and logistics footprint. Both firms also emphasize fraud detection—“Fraud Detection, AML & Risk Mitigation” at JPMC and “Fraud, Security, and Anomaly Detection” at Walmart—which is not surprising given the massive volume of transactions they process.

We next extend this exercise to all AI applications after 2010, clustering them into 20 categories and assigning descriptive labels with a large language model; Appendix B.2 provides longer descriptions and representative examples. Figure 3 displays a heatmap of the share of applications

by cluster (rows) and NAICS sector (columns), with darker shading indicating higher prevalence; clusters are ordered by their average frequency across industries. Several patterns emerge. The most common cluster across nearly all sectors is “Business Intelligence Insights,” reflecting the broad applicability of predictive analytics. By contrast, certain applications are more sector-specific: “Financial Risk Modeling” in Finance, “Personalized Recommendation Engines” in Retail, “Demand and Sales Forecasting” in Utilities, and “Healthcare Diagnostics and Genomics” in Health. Within Education (which includes higher education researchers), the most prevalent applications are “Healthcare Diagnostics and Genomics,” “Scientific and Industrial Modeling,” “Image and Video Recognition,” and “Autonomous Navigation and Robotics.” This exercise highlights the richness of our resume-based measure of AI adoption, which often provides more detailed insights than available firm surveys.⁹

1.3 Building a Task-level Exposure Measure

We now turn to measuring the extent to which workers’ tasks are exposed to the AI application identified in the previous section. Our starting point is that AI applications that are semantically similar to a given task are more likely to be relevant for that task. In the first part of this section, we describe how we estimate this semantic similarity and illustrate the approach with examples. In the second part, we show that skills tied to exposed tasks subsequently appear less often in firms’ job postings after adoption, providing evidence that our measure captures substitution between AI and labor in performing those tasks.

Estimating the similarity between AI applications and tasks

We define occupations using detailed 6-digit SOC codes and obtain task descriptions from the O*NET database.¹⁰ To measure the semantic similarity between AI applications and worker tasks, we use text embeddings. Embeddings encode the semantic meaning of text as vectors in a high-dimensional space, such that semantically similar documents exhibit high cosine similarity.¹¹

⁹Several 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.

¹⁰While 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 description data from the O*NET 28.3 release (May 2024), published by the U.S. Department of Labor.

¹¹We employ the GTE-Large embeddings model, developed by Alibaba DAMO Academy, which encodes text into a 1,096-dimensional vector and performs well in document similarity benchmarks relative to models of similar scale. Earlier word-level embedding models such as GloVe and word2vec ([Pennington, Socher, and Manning, 2014](#); [Mikolov, Sutskever, Chen, Corrado, and Dean, 2013](#)) have been widely used in economics ([Seegmiller, Papanikolaou, and](#)

Using these embeddings, we represent each of roughly one million cleaned AI applications and each of roughly 20,000 O*NET tasks as 1,096-dimensional vectors. This yields a matrix of cosine similarities $\rho(i, j)$ across all application–task pairs. These similarities serve as the building blocks for constructing our measures of task-level AI exposure.

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

$$\xi_{j,f,t} = \frac{1}{N_{f,t}} \sum_{i=1}^{N_{f,t}} I_{j,i}^{95}. \quad (1)$$

Intuitively, $\xi_{j,f,t}$ captures the fraction of firm f 's AI applications in year t that are deemed relevant to task j . Because most AI application–task pairs are plausibly unrelated, we impose a sparsity restriction: a task is considered exposed to an AI application only if the cosine similarity of their text embeddings lies above the unconditional 95th percentile of the overall distribution of similarity scores. Accordingly, the indicator $I_{j,i}^{95}$ equals one if the similarity $\rho(i, j)$ between task j and application i exceeds this cutoff, and zero otherwise.

Examples

The right panels of Table 2 report examples of O*NET tasks (middle column) and their associated occupations (right column) with the highest average exposure probability to AI applications within each cluster for JPMC and Walmart. For JPMC, applications in fraud detection, predictive modeling, and financial forecasting map most closely to tasks performed by Other Financial Specialists and by Accountants and Auditors. By contrast, in the ‘Customer Engagement and Satisfaction’ cluster, the most exposed task is to ‘monitor customer preferences to determine focus of sales efforts,’ performed by Sales Managers. Walmart’s locus of exposure looks quite different. Within its fraud detection cluster, highly exposed tasks include ‘analyze retail data to identify current or emerging trends in theft or fraud’ (performed by Other Managers) and ‘monitor machines that automatically measure, sort, or inspect products.’ In pricing and supply chain optimization, Purchasing Managers and Wholesale and Retail Buyers emerge as the most exposed, while in process automation and operational efficiency the most exposed tasks are linked to Sales Engineers and Other Engineering Technologists.

These examples help to illustrate the rich heterogeneity in exposure that 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 next exploit this heterogeneity to identify how AI adoption shapes labor demand for

Schmidt, 2023; Kogan et al., 2023; Autor et al., 2024), but they assign fixed vectors to words regardless of context. Recent embedding models instead capture contextual meaning.

exposed tasks.

1.4 AI Exposure and Demand for Tasks

Equation (1) can be interpreted as the probability that task j is exposed to AI applications developed in firm f in year t . We also construct a measure of task-level exposure that accounts for differences in the AI adoption intensity in firm f ,

$$\text{Task-Level AI Exposure}_{j,f,t} = \xi_{j,f,t} \times \log(1 + N_{f,t}), \quad (2)$$

where $N_{f,t}$ denotes the number of AI applications developed in firm f in year t . Armed with a measure of task-level exposure to AI, we next ask what this measure captures: are tasks semantically similar to AI applications *substituted* or *complemented* by Artificial Intelligence? Answering this requires data on how demand for specific tasks has evolved over time.

We proxy for task utilization using the composition of skill requirements in firms’ online job postings. Revelio links job postings to both the firm and the relevant occupation. We parse the job posting texts and identify the required skills in each posting using the [Open Skills API](#) from LightCast.¹² LightCast identifies about 30,000 distinct skills. On average, each posting in our dataset lists 17 distinct skills. LightCast provides textual descriptions of each skill, which allow us to connect these skills to ONET tasks by computing the cosine similarity of their GTE embeddings. To impose sparsity, we link a skill to a task only if the similarity score falls in the top percentile of the task–skill distribution. Under this rule, the typical task has roughly 300 associated skills that could be relevant in performing it.

Next, we measure how intensively each task is demanded within firm–occupation job postings. For each occupation o at firm f , we calculate the fraction $S_{j,o,f,t}$ of all listed skills in the postings that are linked to a given task j . To smooth out year-to-year fluctuations, we compute this share over a rolling five-year window, so that the value for year t reflects postings from $t - 4$ through t . To obtain some intuition, Appendix Table A.7 focuses on how these shares have evolved for the three most AI-exposed tasks for a number of selected occupations. For each of these tasks, we report the DHS change in the fraction $S_{j,o,f,t}$ of all listed skills in the postings that are linked to a given task j , averaged across firms.¹³ Examining the table, we see that these exposed tasks experience significant declines in demand relative to other tasks in the same firm–occupation cell. For instance, Claims Adjusters, Appraisers, Examiners, and Investigators show steep drops in mentions of skills linked to document review and claims investigation. Consistent with the prevalence of business-intelligence AI

¹²A growing literature uses LightCast (formerly Burning Glass) skill data, including [Deming and Kahn \(2018\)](#), [Deming and Noray \(2020\)](#), [Acemoglu et al. \(2022\)](#), and [Braxton and Taska \(2023\)](#).

¹³The DHS change allow is a second-order approximation to the log change, so coefficients can be interpreted in units of percentage changes, but it also accommodates cases where one of the shares is equal to zero ([Davis, Haltiwanger, and Schuh, 1996](#)). By construction, these residualized variables average to zero across all tasks in each occupation–firm–year cell, so the table highlights tasks most affected relative to others within the same occupation.

applications, Industrial Production Managers experience large declines in skills related to planning and tracking operational decisions.

We then study the relationship between task-level AI exposure and subsequent labor demand more systematically using the following specification:

$$\Delta_{DHS} S_{j,o,f,t+5} = \beta \text{Task-Level AI Exposure}_{j,f,t} + \alpha_{o,f,t} + \delta X_{j,o,f,t} + \epsilon_{j,o,f,t}. \quad (3)$$

The dependent variable is the DHS change in task demand intensity between years t and $t + 5$. We control for task by calendar year fixed effects to absorb time-variation in the propensity of particular skills to appear in postings, and, depending on the specification, we add interactions of industry, firm, occupation, and calendar year effects. When not including the full set of firm by occupation by year fixed effects, we also control for the occupation-level mean exposure to AI to isolate within-occupation task reallocation—a control which is consistent with our theoretical framework in the next section. Our most saturated specification includes both task by year and firm by year by occupation fixed effects, hence the coefficient β is identified solely from deviations in task exposure within the same firm–occupation–year cell, relative to the economy-wide trend for that task.¹⁴ We cluster standard errors at the occupation–firm level and scale our task-level exposure measure (2) to unit standard deviation.

Table 3 presents the results. Across all specifications, coefficients on task-level AI exposure are negative and of similar magnitude: greater exposure reduces the intensity with which firms demand skills linked to those tasks. The effects are economically meaningful. In our most saturated specification (column 4), a one–standard deviation increase in task-level AI exposure lowers demand for the associated skills by about 2 percent over five years. Because the dependent variable is a share, this implies reallocation of labor effort away from exposed tasks and toward less-exposed tasks within the same occupation.

In short, tasks highly exposed to AI experience subsequent declines in their relative demand, consistent with AI substituting for human labor in performing them. In the next section, we consider how these task-level effects translate into labor demand for entire occupations.

2 Measuring the Exposure of Jobs to AI

So far, we have constructed a measure of task-level exposure to AI. Because tasks that are semantically similar to AI applications are likely to experience reduced labor demand, this measure primarily captures substitution between AI and labor. Yet workers perform bundles of tasks, and these tasks may be complements. Thus, the fact that some tasks can now be performed by AI does not

¹⁴For example, within JPMC’s Analyst positions in 2018, the comparison comes from whether skills linked to fraud detection modeling declined more than skills linked to client communication, over and above the average change in demand for those skills across all firms.

necessarily imply that the occupation as a whole will experience lower labor demand. To map task-level exposures into shifts in occupational demand, we develop a model following [Acemoglu and Autor \(2011\)](#); [Acemoglu and Restrepo \(2018\)](#); [Caunedo, Jaume, and Keller \(2023\)](#); [Kogan et al. \(2023\)](#). Using the model as a guide, we then construct 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. Labor and capital are substitutes in production, and capital is task-specific. Technological progress shows up as declines in the quality-adjusted price of capital. Within a job, workers allocate time across tasks. Across jobs, they choose optimally given wages and idiosyncratic preferences. The key feature of our model is that we allow technology to hit tasks unevenly, so workers can reorganize their effort across tasks when capital substitutes for labor. To simplify the exposition, we only discuss the key model ingredients and delegate all details to [Appendix A](#).

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. If a technology uniformly improves the capital that is specific to most of the worker’s tasks, then labor demand for that occupation decreases. 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 can potentially increase, since tasks are complements and workers can endogenously allocate their time among tasks.

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,

$$\bar{Y} = \left(\int_f \alpha_f^{\frac{1}{\theta}} Y_f^{\frac{\theta-1}{\theta}} df \right)^{\frac{\theta}{\theta-1}}. \quad (4)$$

Here, θ captures the elasticity of substitution across firms and α_f can represent the number of products produced by the firm—see [Appendix A.3](#). Each firm produces a differentiated good by combining the output of many occupations,

$$Y_f = \left(\int_o Y(o, f)^{\frac{x-1}{x}} \right)^{\frac{x}{x-1}}. \quad (5)$$

Firms make profits because of imperfect competition, reflecting both monopolistic competition in product markets and monopsonistic power in labor markets. 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

$$Y(o, f) = \left(\sum_j \alpha(j)^{\frac{1}{\psi}} y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}} \quad (6)$$

Here, ψ denotes the elasticity of substitution across tasks within a given job. This parameter governs the elasticity of labor demand for each task. The weight $\alpha(j)$ measures the contribution of task j to occupation-level output. Appendix A.3 shows how this structure can arise endogenously in a model where each task is itself composed of many micro-tasks; in that setting, $\alpha(j)$ is simply the measure of micro-tasks assigned to task j . For clarity, we suppress firm and occupation subscripts unless required.

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

$$y(j) = \left(\gamma_j l(j)^{\frac{\nu-1}{\nu}} + (1 - \gamma_j) k(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}. \quad (7)$$

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 ,

$$\Delta \log q(j) = -\varepsilon(j). \quad (8)$$

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 make two decisions about labor supply. They first choose a job, defined as an occupation-firm pair. They then allocate time across the tasks within that job. The effective supply of labor by worker in job (o, f) to task j is given by

$$l(j) = \alpha(j)^\beta h(j)^{1-\beta}, \quad (9)$$

where $h(j)$ denotes the time spent on task j and $\alpha(j)$ is a measure of task importance discussed above. The total number of hours a worker can supply across all J tasks is equal to one. The parameter $\beta \in (0, 1)$ captures the degree of decreasing returns to effort at the task level; a smaller

value of β implies that there is more scope of reallocating effort across tasks. As a result, the optimal allocation of effort to task j is equal to

$$h(j) = \frac{\alpha(j) w(j)^{\frac{1}{\beta}}}{\sum_{j' \in J} \alpha(j') w(j')^{\frac{1}{\beta}}}, \quad (10)$$

where $w(j)$ is the job-specific wage for performing task j .

Workers optimally choose jobs based on the total earnings of that job 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 α_f and a shape parameter ζ . Using the properties of the Fréchet distribution, the measure of workers that choose job (o, f) is given by

$$N(o, f) = \alpha_f \bar{\zeta} W(o, f)^\zeta, \quad (11)$$

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

$$W(o, f) \equiv \sum_{j \in J_o} l(j) w(j) \quad (12)$$

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 approximate the change in labor demand around the symmetric equilibrium in which the labor share is constant across tasks (setting $\gamma_j = \gamma$, $w(j) = w$, and $q(j) = q$ for all $j \in J$). In this case, we obtain a simple equation for employment growth due to changes in technology,

$$\Delta \log N(o, I) \approx \underbrace{\zeta \eta_m m(\varepsilon) + \frac{\zeta}{2\beta} \eta_o^2 C(\varepsilon)}_{\text{Direct effects}} + \underbrace{\zeta \eta_z \Delta_\varepsilon \log Z_f + \Delta \log \alpha_f}_{\text{Firm Spillovers}} + \underbrace{\frac{\zeta \eta_z}{\theta - \chi} \Delta \log \left(\frac{\bar{Y}}{\bar{\zeta}} \right)}_{\text{Aggregate Spillovers}}. \quad (13)$$

Here,

$$m(\varepsilon) \equiv \sum_{j \in J} \frac{\alpha(j)}{\sum_{k \in J} \alpha(k)} \varepsilon(j) \quad (14)$$

denotes the mean improvement of the technology across all tasks;

$$C(\varepsilon) \equiv \sum_{j \in J} \frac{\alpha(j)}{\sum_{k \in J} \alpha(k)} \left(\varepsilon(j) - m(\varepsilon) \right)^2 \quad (15)$$

denotes the concentration in exposure to specific tasks; Z_f denotes the productivity of firm f ; and the last term captures aggregate spillovers.

Equation (13) is the core of our empirical analysis. We now discuss each term in turn. The first two terms in (13) capture the direct effects of technology on the marginal product of labor for job (o, f) . The first component of the direct effect (14) corresponds to the average improvement in technology across all tasks in the job—weighted by the O*NET importance weight of each task. Its impact on labor demand is governed by

$$\eta_m \equiv -\frac{s_k (\nu - \chi)}{\zeta + \nu s_k + \chi (1 - s_k)}, \quad (16)$$

where s_k denotes the capital share (assumed equal across tasks). The sign of η_m depends on whether the elasticity of substitution between capital and labor ν exceeds the elasticity of substitution across occupations within firms χ . A capital improvement specific to task j substitutes directly for labor in that task, with the degree of substitution governed by ν . At the same time, the productivity of the occupation rises in this firm, which raises demand for its output. The elasticity of that demand is χ (Hicks, 1932). The net effect on labor demand depends on the sign of $\nu - \chi$: whether the fall in task-level labor demand from labor-saving technology outweighs the increase in occupation-level demand from higher firm productivity.

The second term in equation (13) rises with the dispersion of technological improvements $\varepsilon(j)$ across tasks. This second-order correction term captures how much the occupation’s average exposure $m(\varepsilon)$ is concentrated in a small number of tasks. Holding the mean exposure constant, labor demand increases with this term: when technology greatly improves the capital used in some tasks but not others, workers reallocate effort toward the unaffected tasks. Two forces drive this concentration effect. First, workers optimally shift time across tasks in response to changes in relative productivity, with the scope for reallocation inversely related to β . Second, log wages are convex in the vector of task-level wages. By Jensen’s inequality, mean-preserving spreads in task prices $\log w(j)$ raise occupation wages and labor demand. Both forces are stronger when productivity gains are concentrated in a subset of tasks. The importance of this term depends on

$$\eta_o = -\frac{s_k \beta (\nu - \psi)}{(1 - \beta) + \beta (\nu s_k + \psi (1 - s_k))}, \quad (17)$$

which captures the impact of a technology improvement specific to task j on the wage of task j relative to the wage paid on other tasks $j' \neq j$.

To build intuition for the role of the second term, consider the case of two tasks $j \in \{1, 2\}$. Figure 4 plots the effect of technology changes on overall labor demand (left panel) and on the time allocated to the first task (right panel). As the average improvement in technology (14) rises, labor demand falls. The size of this decline, however, depends on how the technology improvements are distributed across tasks. When improvements are spread evenly across the two tasks (red line), labor demand falls the most: there is no scope for reallocating effort. By contrast, when improvements are uneven—some tasks more affected than others—workers reallocate their effort toward the less

affected task. In this case, illustrated by the blue and green lines, the concentration term (15) offsets the fall in labor demand. Depending on the magnitude of the technology improvement and the degree of concentration, labor demand can increase even though the occupation is highly exposed to technology.

This simple example highlights a central point of the paper: the effect of advances in labor-saving technologies on occupation labor demand is nuanced. It depends not only on which tasks it substitutes for, but also on how unevenly those substitutions occur. Even if technology now performs some tasks more efficiently than before, demand for the occupation need not decline—even holding firm-level productivity constant. In fact, if the improvements are large enough and highly concentrated in a few tasks, labor demand for the occupation can rise.

The above discussion focuses on within-firm shifts in occupation labor demand. In addition to the direct effect discussed above, the last two terms in (13) that capture spillovers at the firm and aggregate level, respectively. The third term in (13) captures the effect of technology on overall firm labor demand. It has two components, which essentially capture process and product innovation. The first operates through firm productivity, defined as the reciprocal of unit labor cost,

$$Z_f \equiv \left(\int_o P(o, f)^{1-\chi} \right)^{-\frac{1}{1-\chi}}, \quad (18)$$

where $P(o, f)$ is the price paid by firm f for the output of occupation o . The second operates through changes in α_f . The impact of the first component on labor demand is governed by

$$\eta_z \equiv \frac{\theta - \chi}{\chi + s_k(\nu - \chi) + \zeta}. \quad (19)$$

As long as the elasticity of substitution across firm output, θ , exceeds the elasticity of substitution across occupations within the firm, χ , higher firm productivity raises labor demand across all occupations. The second component of the third term captures shifts in labor demand from technologies that affect the overall demand for firm output—for instance, the emergence of new products. The last term in (13) captures economy-wide increases in labor demand (or labor supply) due to technological improvements.

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

$$\Delta \log h(j) \approx \frac{\eta_o}{\beta} (\varepsilon(j) - m(\varepsilon)). \quad (20)$$

We see that the amount of effort allocated to a particular task within an occupation is a function of the task’s technology exposure relative to the average exposure across all tasks. Thus, equation (20) maps directly into our empirical specification in Section 1.4.

In the next section we construct direct empirical proxies for the first three terms of (13). Because

all specifications include calendar-year fixed effects, the last term cannot be estimated. This ‘missing intercept’ implies that our regressions identify only relative shifts in employment across jobs and are silent about level effects. Further, in some specifications we also include firm-by-year fixed effects. These isolate the within-firm component of labor demand—the first two terms—and help address concerns about endogenous adoption of AI across firms.

Last, we emphasize that the mean and concentration statistics we derive are not unique to our baseline CES model. In Appendix A.4, we show that the same estimating equation emerges in the framework of Acemoglu and Restrepo (2018), where technological progress shifts an automation threshold and reassigns micro tasks from labor to capital. In this interpretation, our job-level CES production function (6)–(7) serves as a reduced-form representation of that task reassignment process: CES coefficients γ_j and $1 - \gamma_j$ no longer sum to one but instead reflect the endogenous division of tasks between labor and capital. This formulation highlights an extensive margin of displacement—some micro tasks exit the domain of labor entirely—which reduces wages and labor demand. Nevertheless, mean exposure and concentration remain the relevant sufficient statistics for the relative employment effects within a firm. Incorporating the reassignment mechanism yields predictions that closely resemble those of our baseline model under a higher elasticity of substitution ν , underscoring the generality of our empirical framework.

2.2 Measuring the Exposure of Workers to AI

Our model implies that the impact of a specific technology on firm demand for an occupation can be summarized by three terms: the average improvement in labor-saving technology across tasks (14), the degree of concentration of these improvements in specific tasks (15), and firm-level spillovers. We now construct empirical analogues of these objects.

First, we need to take a stance on the variable driving firm spillovers—the third term of equation (13). In the model, firm growth depends on productivity Z_f and the scope of varieties α_f the firm produces. Both margins plausibly rise with the number $N_{f,t}$ of AI applications developed in firm f in year t . However, this variable is highly skewed: some firms develop thousands of applications while others adopt only a handful (Appendix Table A.1). To ensure that the highly skewed distribution of $N_{f,t}$ does not influence our findings, we measure the intensity of adoption using $\log(1 + N_{f,t})$.

Next, we construct the occupation-level measures of AI exposure. We begin with the weighted average exposure probability across an occupation’s tasks,

$$\mu_{o,f,t} = \sum_j \omega_{o,j} \xi_{j,f,t} \quad (21)$$

The weights $\omega_{o,j}$ correspond to the O*NET task importance scores: each task is assigned a score between 1 and 5 to indicate how central the task is to the job. We rescale these scores so that

the weights sum to one within each occupation. This measure (21) captures the direction of AI innovation in firm f , but not the intensity of adoption. To incorporate intensity, we scale by the firm’s AI adoption intensity,

$$\text{AI Exposure Average}_{o,f,t} = \mu_{o,f,t} \times \log(1 + N_{f,t}) \quad (22)$$

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,

$$c_{o,f,t} = \sum_j \omega_{o,j} (\xi_{j,f,t} - \mu_{o,f,t})^2, \quad (23)$$

and after incorporating differences in adoption across firms we obtain

$$\text{AI Exposure Concentration}_{o,f,t} = c_{o,f,t} \times \log(1 + N_{f,t}). \quad (24)$$

Equations (22) and (24) are the empirical counterparts to equations (14) and (15) in the model. They vary across jobs (o, f) both because firms develop different AI applications that are related to different tasks and because they adopt them with different intensity. Our model implies that an increase in (22) will lead to a decrease in employment *within* the firm, while holding (22) constant, an increase in (24) will lead to higher employment within the firm.

2.3 Stylized Facts about AI Job Exposure

Figure 5 plots the mean AI exposure for each application cluster, averaged across broad occupation groups. Applications in the Business Intelligence Insights category affect a wide range of workers, especially those in Management, Business and Financial, and Computer and Math occupations. By contrast, ‘Cybersecurity and Fraud Detection’ primarily affects Protective Services, which otherwise show little AI exposure. Customer experience automation and marketing optimization are concentrated in Sales occupations, while task and workflow optimization AI applications are most relevant for Production and Office & Administrative occupations. Financial risk modeling primarily impacts Business and Financial roles.

These patterns suggest that white-collar occupations face greater AI exposure overall. Figure 6 confirms this: average exposure $\mu_{o,f,t}$ rises with wage rank, peaks around the 90th percentile, and then declines. Unlike earlier technological shifts that displaced middle-skill workers (Autor et al., 2006; Kogan et al., 2023), AI is most relevant for higher-wage jobs. At first glance, this could suggest that high earners are especially vulnerable. However, Appendix Figure A.5 shows that technologies generating high mean exposure also tend to generate concentrated exposure, which mitigates displacement by enabling task reallocation. Still, the correlation is far from perfect—about 0.67 at the occupation–firm level—so mean and concentration capture distinct margins. For instance,

‘Operational Data Analytics’ affects tasks more narrowly than ‘Customer Experience Automation’ despite similar mean exposure, leading to very different implications for labor demand.

These aggregate comparisons mask important heterogeneity across occupations, as Figure 5 shows. For example, Management Analysts and Financial Risk Specialists show similar mean exposure to ‘Business Intelligence Insights’ applications, but exposure is more concentrated for Management Analysts, making them less likely to be displaced according to our model. Likewise, Compensation, Benefits, and Job Analysis Specialists have lower mean task exposure to AI but highly concentrated exposure, giving them more scope to reallocate effort toward unaffected tasks; our model thus predicts rising labor demand for these jobs. The same logic applies within other clusters. Under ‘Financial Risk Modeling’, Financial Managers face both higher and broader exposure than Brokerage Clerks, implying larger declines in labor demand. Under ‘Customer Experience Automation’, Sales Managers and Telemarketers show similar mean exposure, but the more concentrated exposure of Sales Managers suggests smaller employment losses relative to Telemarketers.

3 The Impact of AI Exposure on Labor Demand

We begin by outlining our identification strategy for estimating the impact of AI exposure on occupation-level labor demand. We then present the estimates and a series of robustness checks. Finally, we study how AI-driven shifts in labor demand translate into reallocation over the sample period.

3.1 Identification Strategy

Our goal is to estimate the effect of AI exposure on firm and labor market outcomes. The challenge is that adoption is endogenous: firms choose whether and how to adopt AI, and those choices are correlated with other drivers of performance and labor market conditions. In this section we outline our strategy to address these concerns. First, we introduce a firm-level instrument that isolates plausibly exogenous variation in the cost of adopting AI. Second, we address the concern that firms endogenously choose *how* to implement AI—for example, by targeting automation at tasks that are especially scarce.

Endogeneity in level of AI adoption

The first challenge is that AI adoption is endogenous. Larger, more productive, and more profitable firms are more likely to adopt, but such firms also tend to grow more slowly (Evans, 1987). This selection bias means that adopters may underperform non-adopters for reasons unrelated to AI. Alternatively, firms with superior management or stronger technological infrastructure may both

adopt AI and realize higher gains, biasing estimates upward. These concerns make clear the need for exogenous variation in AI adoption to identify causal effects.

To address these concerns, we first construct a shift-share instrument that isolates plausibly exogenous variation in firm-level AI adoption costs. The strategy exploits two sources of variation: cross-firm differences in historical university hiring patterns and cross-university differences in the rise of AI-related work. We define the ‘shift’ as the share of a university’s graduates entering AI-related occupations during 2014–2018, measured from resume data spanning a wide set of institutions and cohorts. We define the ‘share’ as the distribution of each firm’s hiring across universities during the pre-AI period (2005–2009), before the diffusion of modern AI tools. These hiring shares are held fixed throughout the analysis to avoid simultaneity with subsequent adoption. Multiplying each university’s AI shift by the firm’s baseline hiring share, and aggregating across feeder institutions, yields a firm-level measure of predicted AI exposure. In essence, the instrument exploits heterogeneity in firms’ pre-existing hiring networks: firms that drew more heavily from universities whose graduates later entered AI-intensive jobs face lower costs of attracting AI talent, and hence lower costs of adoption.

We instrument for the (log) number of AI applications using past hiring practices:

$$\log(1 + N_{f,t})^{IV} \equiv \log\left(1 + N_{f,t}^{total} \times p_{f,t}^{AI}\right). \quad (25)$$

The key assumption is that the number of AI applications within a firm rises with the predicted number of AI integrators. Here, $N_{f,t}^{total}$ is the total number of resumes in active job positions at firm f in year t , and $p_{f,t}^{AI}$ is the predicted probability that a worker is an AI integrator. We construct $p_{f,t}^{AI}$ as

$$p_{f,t}^{AI} = \sum_{u \in \mathcal{U}} w_{u \rightarrow f}^{2005-2009} \times \frac{\sum_{k \neq f} N_{u,k,t}^{AI}}{\sum_{k \neq f} N_{u,k,t}^{total}} \quad (26)$$

Here $N_{u,k,t}^{AI}$ is the actual number of AI integrators employed in firm k in year t who report graduating from university u on their resumes; $N_{u,k,t}^{total}$ is the number of total workers who report graduating from university u . The set \mathcal{U} represents universities represented by workers with active job positions in the resume data from 2005 to 2009. The weight $w_{u \rightarrow f}^{\tau}$ is the average share of workers at firm f during period τ who graduated from university u , computed from the Revelio resume data. When constructing (26), we include all workers in the Revelio data, not just those employed in firms in Compustat.

Figure 7 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, based off their average university employment shares from 2005–2009. 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 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.

A key rationale for the relevance of our instrument is that the scarcity of workers with the skills to develop and deploy AI is often a binding constraint on adoption. Survey evidence supports this view: in the Business Trends and Outlook Survey (BTOS) survey, 3.6 percent of firms (employment-weighted) cite the lack of a skilled workforce as a barrier to adoption; 8.8 percent cite lack of knowledge about AI capabilities; and 4.6 percent cite high costs (Bonney et al., 2024). Evidence from the 2018 ABS survey points in the same direction: among firms testing AI, 36 percent (employment-weighted) report that the lack of available talent hindered adoption, a share comparable to the fraction citing lack of data as a barrier (Chequer et al., 2025).

In Panel A of Appendix Table A.3, 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 above 30. In Panel B, we demonstrate that the predicted share of AI workers based on these historical university links (26) significantly predicts the actual share of AI workers at the firm, with a coefficient of 0.54 and t -statistic of 7.5 (conditional on controls for log resume employment and industry by year fixed effects). Together, these findings confirm that our instrument is strongly correlated with actual AI adoption.

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. The fact that we leave out the focal firm, helps with the first two assumptions; Section 3.4 below includes further robustness checks. 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; this assumption is consistent given the evidence in Appendix Table A.3.

Endogeneity in the direction of AI adoption

Firms not only decide whether to adopt AI, but also how to deploy it: they may target adoption to tasks with the largest expected cost savings—such as those facing labor scarcity or declining productivity. The result is a non-random pattern of AI exposure across occupations. This selective implementation biases estimates of the effect of average AI exposure in a predictable direction: it induces downward bias in the OLS coefficients on both the mean and the concentration of AI exposure. Measurement error introduces a distinct challenge: our firm-occupation measures are often based on limited data on actual implementation. This raises the possibility of attenuation bias.

To address both of these issues, our IV strategy uses AI applications across *all other* firms k to measure technology exposure (1) for task j as

$$\xi_{j,f,t}^{IV} \equiv \frac{1}{N_t - N_{f,t}} \sum_{k \neq f} \left(\sum_{i=1}^{N_{k,t}} I_{j,i}^{95} \right). \quad (27)$$

Next, using (27), we construct the occupation-level instruments for the mean and concentrated exposure measures as

$$\text{AI Exposure Average}_{o,f,t}^{IV} \equiv \left[\sum_j \omega_{o,j} \xi_{j,f,t}^{IV} \right] \times \log(1 + N_{f,t})^{IV} \quad (28)$$

and

$$\text{AI Exposure Concentration}_{o,f,t}^{IV} \equiv \left[\sum_j \omega_{o,j} \left(\xi_{j,f,t}^{IV} - \sum_j \omega_{o,j} \xi_{j,f,t}^{IV} \right)^2 \right] \times \log(1 + N_{f,t})^{IV}. \quad (29)$$

In a nutshell, we instrument for the mean and concentration of technology exposure for a particular occupation in a given firm with the analogous measures computed across all other firms. This approach serves two goals. First, by drawing on variation in AI exposure that is external to the firm, allows us to focus on broader technological diffusion of AI towards certain tasks, rather than any selected targeting of particular occupations for purely firm-specific reasons. Second, by computing the means and concentrations across a broad set of firms, the estimation noise in the mean and concentration decrease, which helps reduce attenuation bias.

3.2 AI Exposure and Firm Growth

We estimate the effect of AI adoption on subsequent firm outcomes using

$$\log Y_{f,t+5} - \log Y_{f,t} = \beta \log(1 + N_{f,t}) + \alpha_{I,t} + \delta X_{f,t} + \epsilon_{f,t}, \quad (30)$$

where the dependent variable is the five-year log growth of outcome $Y_{f,t}$ and the key regressor is $\log(1 + N_{f,t})$, the log of one plus the number of AI applications at firm f in year t . We control for log Compustat employment and the log number of Revelio employees, which account for the mechanical correlation between the number of AI integrators (and hence applications) and firm size or resume coverage. We further include industry-by-year fixed effects, $\alpha_{I,t}$, to absorb industry-specific trends, and cluster standard errors at the firm level. Last, we control for one lag of the dependent variable to account for pre-existing trends in growth rates, and the level of $Y_{f,t}$ to account for decreasing returns to scale.

Our main coefficient of interest is β , which captures the effect of AI adoption intensity on subsequent firm performance. In the model, firm growth depends on productivity Z_f and the scope of varieties α_f the firm produces. Both margins plausibly rise with the number of AI applications developed, so β provides a reduced-form estimate of the impact of AI adoption. Table 4 presents the results: the first four columns report OLS estimates, and the last four report IV estimates using the shift-share instrument introduced above. First-stage F -statistics are large across all IV specifications, alleviating concerns about weak instruments.

Examining the table, we see a strong positive relation between the intensity of AI adoption and subsequent five-year growth in sales, employment, profits, and total factor productivity across specifications. The magnitudes are economically sizeable: focusing on the IV results, a one-standard deviation increase in our AI adoption measure leads to a 9.7 percentage point increase in firm revenue growth and a 7.5 percentage point increase in total factor productivity over the subsequent five years. Notably, profits increase roughly in line with sales, suggesting that markups remain relatively stable. Importantly, we find a positive effect of AI adoption on employment: a one-standard deviation increase in AI adoption leads to a 6.8 percentage point increase in firm employment.

We note that the IV estimates are somewhat larger than the OLS estimates. This pattern is consistent with firms that adopt AI also being those that naturally grow more slowly, which biases OLS coefficients toward zero. Survey evidence (Acemoglu et al., 2023b) supports this view, showing that advanced technology adopters are typically larger and older firms already in a lower-growth phase of their life cycle. Our instrumental variables strategy helps mitigate this downward bias.

In brief, AI adoption intensity is associated with higher growth in revenue, profitability, and employment. The employment gains reflect firm-level spillover effects—consistent with the third term in equation (13)—rather than direct complementarity between AI and worker tasks. To examine the latter, we next compare relative employment across occupations with differing AI

exposure.

3.3 AI Exposure and Occupational Labor Demand

Having established a causal impact of AI on firm productivity and growth, we now turn to the main part of our analysis: labor demand for specific jobs. Our strategy exploits the granularity of the exposure measures in (22) and (24), which capture how exposed particular occupations are to AI applications within a firm.

Our outcome variable is occupation–firm–year employment, measured at the 6-digit SOC level using Revelio resume data. We count the number of active positions in each firm-occupation-year, weighting observations with Revelio’s adjustment factors to correct for occupational coverage bias. A position is considered active in a given year if it was held for at least half the year. Because our model emphasizes the labor-substituting effects of AI at the task level, we exclude 2-digit SOC 15 (Computer and Mathematical occupations), since these are by far the most likely broad occupation group to be tagged as AI implementers, representing about 50% of all AI integrator positions in our data. These workers, even if not directly developing AI, likely implement and maintain the software and hardware necessary for its use—and hence their labor input is plausibly complementary to AI adoption.

Our first specification exploits both within- and between-firm variation:

$$\begin{aligned} \log(\text{Emp}_{f,o,t+5}) - \log(\text{Emp}_{f,o,t}) = & \beta \text{AI Exposure Average}_{f,o,t} + \gamma \text{AI Exposure Concentration}_{f,o,t} \\ & + \delta \log(1 + N_{f,t}) + \Gamma X_{f,o,t} + \alpha_t + \epsilon_{f,t} \end{aligned} \quad (31)$$

This specification includes only calendar year fixed effects and directly controls for firm-level AI adoption, $\log(1 + N_{f,t})$, which in the model reflects shifts in productivity Z_f and scope α_f . Thus, its first three terms correspond to the first three terms in equation (13). The control vector X includes firm employment (log Revelio and Compustat) and lagged growth in occupation-firm employment. Standard errors are clustered at the occupation–firm level, and observations are weighted by each cell’s share of total employment. To aid interpretation, we scale independent variables to unit standard deviation.

Our second specification focuses on within-firm patterns by differencing out firm-level shocks:

$$\begin{aligned} \log(\text{Emp}_{f,o,t+5}) - \log(\text{Emp}_{f,o,t}) = & \beta \text{AI Exposure Average}_{f,o,t} + \gamma \text{AI Exposure Concentration}_{f,o,t} \\ & + \Gamma X_{f,o,t} + \alpha_{f,t} + (\alpha_{o,t}) + \epsilon_{f,t} \end{aligned} \quad (32)$$

Relative to equation (31), this specification adds firm by year fixed effects $\alpha_{f,t}$, which absorb all contemporaneous firm-level shocks. The benefit is that identification now comes from variation in occupational exposure within firms, eliminating the bias from general firm-level selection into AI

adoption. As a variant, we also include 6-digit SOC interacted with calendar year effects, which purge occupation-specific trends by exploiting how exposure differs within a firm relative to the average exposure of that occupation across all firms. The granularity of our exposure measure makes this feasible, unlike existing approaches that rely only on cross-occupation differences (Webb, 2020; Eisfeldt et al., 2023; Eloundou et al., 2023; Brynjolfsson et al., 2018). The tradeoff is that we can no longer identify δ , the coefficient on firm-level spillovers from AI, so the specification isolates only the first two terms in equation (13).

We report the estimated coefficients from (31) and (32) in Table 5. Columns one through four report the OLS estimates; columns five to eight report the IV estimates using the shift-share instrument in Section 3.1. The first-stage F-statistics remain high, alleviating concerns about weak instruments. We next turn to the substantive results, analyzing the mechanisms through which AI exposure shapes labor demand.

First, we examine the effect of mean AI task exposure. Table 5 shows that the estimated coefficient β on mean AI exposure (22) is negative and highly significant in all specifications. The magnitudes are large: a one-standard deviation increase in exposure predicts a 5.7 to 8.7 percent employment decline in the OLS estimates and a 10.4 to 15.6 percent decline in the IV estimates. The magnitude of the estimates is smallest when we include occupation-year fixed effects in the most strict specifications in columns 4 and 8. The larger IV coefficients point to attenuation bias rather than selective adoption as the source of these differences. If firms selectively adopt AI in occupations facing rising labor scarcity or falling productivity, OLS estimates would be biased downward, not toward zero.

We next turn to the concentration of AI exposure across tasks, which the model predicts can offset the negative average effects. Table 5 shows that the estimated coefficient γ on the concentration measure (24) is positive and highly significant in all specifications. A one-standard deviation increase in concentration is associated with a 1.3 to 1.9 percent increase in employment in the OLS estimates, and a 7.5 to 9.1 percent increase in the IV estimates. The larger IV coefficients again suggest attenuation bias from measurement error, which is especially likely for concentration given that firms have a small number of AI applications.

Finally, the estimates of the firm spillover term δ show that a one-standard deviation increase in AI adoption intensity is associated with a 10 to 11 percent increase in employment in the OLS specifications and almost 20 percent increase in the IV specifications. Identifying this coefficient requires dropping firm-year fixed effects (columns 1–2 and 5–6). Consistent with Section 3.2, the IV coefficients exceed the OLS estimates, likely because selective adoption by larger, slower-growing firms biases the OLS coefficients downward.

Overall, the signs of these three effects align with the theoretical predictions of our model. The negative estimate of β confirms AI-labor substitution at the task level. The positive estimate

of γ indicates that within-occupation spillovers across tasks mitigate some of this substitution effect. Finally, the positive estimate of δ shows that firm-wide AI adoption raises employment through productivity growth. The model is silent on which mechanism dominates, as their relative importance depends on parameter values. In Section 3.5, we return to this question and assess how much AI adoption contributed to labor reallocation during our sample period.

3.4 Discussion: Identification Assumptions

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.

Our 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. Next, we implement a series of empirical tests designed to validate this exclusion restriction.

Anticipatory hiring. A first concern is that firms may have shaped their pre-period 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. We can address this concern by removing from the sample the top 50 universities with the highest number of graduates working in AI during the 2014–2018 period.

University prestige and firm quality. A related concern is that association with elite universities could drive both the ‘shift’ and the ‘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. Firms that hire disproportionately from such schools may simply be more capable overall. In this case, 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 also 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.

Reverse causality. Another possibility is that large or technologically forward-looking firms

influenced the AI curricula of their feeder universities. 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.

Geographic clustering. A further concern is that firms located in 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.

Panel A of Appendix Table A.4 simultaneously addresses the concerns discussed above. Specifically, we first remove 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. Last, we 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). Examining Panel A, we see that the IV estimates remain statistically significant and only slightly attenuated relative to our main specification, which helps address the concerns above. Further, the stability of predictions even after removing some of the most important universities for predicting AI use is particularly helpful in providing support for the exogeneity of the pre-period shares (Goldsmith-Pinkham, Sorkin, and Swift, 2020).

Broader technical intensity. Last another concern is that firms hiring from AI-intensive universities may benefit from exposure to a broader supply of technically skilled graduates, beyond those developing AI applications; this possibility serves to confound the estimated effect of AI exposure on firm growth. To address this concern, Panel B of Appendix Table A.4 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 conduct parallel robustness checks for the occupation–firm IV specifications in Appendix Table A.5. In the first four columns, we drop the top 50 universities when constructing the IV, as well as the top 50 AI-using firms and firms in high-tech industries. In the last four columns, we add controls for predicted shares of workers in computer and mathematical occupations (SOC 15) and engineering occupations (SOC 17), both directly and interacted with the leading terms in brackets in equations (28) and (29). Apart from these adjustments, the specifications are identical to 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.

3.5 Overall Effects of AI on Labor Reallocation

We conclude our analysis by assessing the overall impact of AI developments on labor demand implied during our sample period. As we noted in our discussion of equation (13), our empirical design identifies only relative effects across occupations. The aggregate term in equation (13) is absorbed by the calendar-year fixed effects, which capture contemporaneous macroeconomic conditions and policy events—such as the recovery from the 2008–2009 financial crisis and the 2017 Tax Cuts and Jobs Act—and thus isolate the contribution of AI from other shocks.

With this caveat in mind, we focus on the role of AI in driving reallocation across occupations. To quantify these reallocative effects, we compute the expected net marginal impact of AI exposure on labor demand conditional on job (occupation \times firm) characteristics K ,

$$E[\widehat{\beta}\text{AI Exposure Average}^\perp + \widehat{\delta}\text{AI Exposure Concentration}^\perp + \widehat{\gamma}\log(1 + N_{f,t})^\perp \mid K = k]. \quad (33)$$

Equation (33) maps directly into the first three terms of equation (13). We take the estimated coefficients $\widehat{\beta}$, $\widehat{\delta}$, and $\widehat{\gamma}$ from the IV estimates in column (5) of Table 5 and apply them to our exposure measures (22) and (24). 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 category). Because averages are absorbed by controls, these marginal effects are relative to mean employment growth. By construction, integrating across the employment-share weighted distribution of K yields zero aggregate effects. This decomposition highlights which occupations gain and which lose from AI exposure, even though aggregate effects are absorbed by fixed effects.

Our first exercise compares occupations by pay level to measure how AI affects labor demand

at different points in the earnings distribution. For each wage salary percentile k , we compute equation (33). The top panel of Figure 8 reports the results. The red line shows the direct effect of mean AI exposure on employment (the first term in (33)); the green line captures reallocative effects from concentrated exposure (the second term); and the yellow line reflects firm-level productivity effects (the third term). The blue line combines all three to give the total net effect.

Consider first the direct substitution effect—the red line in Figure 8. The negative impact of mean AI exposure on employment shares rises steadily with job salary—consistent with Figure 6, which shows greater exposure among higher-paid occupations. At the top of the earnings distribution, our estimates imply that these occupations lose about 7% of employment share over five years from this direct effect, while lower-paid occupations gain nearly 10% in relative terms.

These direct effects, however, omit two offsetting forces: within-job reallocation of effort and firm-level productivity gains. High-paying occupations are typically located at firms that make intensive use of AI (yellow line) and also face more concentrated exposure across tasks (green line). Together, these factors largely neutralize the substitution effect, leaving the net impact (blue line) modest across most salary levels, with essentially no effect for the middle 60 percent of the pay distribution. At the very top, the net effect is positive at around 3 percent over a 5-year period, despite substantial direct exposure for these occupations. This pattern is consistent with [Acemoglu et al. \(2022\)](#), who document negligible aggregate employment effects of AI during 2010–2018 despite strong firm-level impacts. Our decomposition explains why: AI’s impact is large, but its overall footprint is muted by offsetting forces.

We next restrict attention to within-firm variation in AI exposure. Panel B of Figure 8 shows the results. The direct substitution effect is again strongest for high-paying jobs: at the 90th percentile, within-firm employment falls by about 2.8% over five years. Concentration of exposure across tasks tempers this decline to about 1.5%. At the other end of the distribution, low-wage jobs see modest employment gains of roughly 2%. Similar to our overall analysis, the net employment effects remain small once offsetting forces are taken into account. Substitution pressures are strong, but reallocation within jobs and productivity gains at the firm level substantially mute their impact. Taken together, the within-firm estimates reinforce the main conclusion: AI exposure reshapes employment patterns, but its overall footprint is modest because powerful forces offset each other.

Turning from earnings differences to occupational categories, we next compare across broad occupation groups defined by 2-digit SOC codes. Table 6 summarizes the components of AI-related relative employment effects for each group; employment-weighted totals sum to zero by construction. We see that several groups experience meaningful declines: ‘Business and Financial’ occupations fall by 2.2% on average, ‘Architecture and Engineering’ by 1.3%, and ‘Food Preparation and Serving’ by 4.6%, largely because their employers rarely adopt AI. By contrast, legal occupations expand by 7.6%, reflecting low direct exposure to the first wave of AI technologies in our sample coupled with

sizable firm-level productivity gains.

These patterns underscore that AI’s aggregate effects are modest overall, but mask substantial heterogeneity across occupational groups. To highlight this point, we next turn to detailed variation within groups at the 6-digit occupation level. Figure 9 examines AI impacts at the detailed 6-digit SOC level. Panels A and B plot task-level and firm-level AI effects against realized employment shifts, while Panel C shows their combined impact. Both occupation-specific task exposure and firm-level AI utilization are strongly correlated with actual employment changes. Quantitatively, AI-related measures explain a large share of the variation in employment shifts, with a regression R^2 of about 16 percent, of which at least 45 percent is attributable directly to task exposure.

One concern with this analysis is external validity beyond the Revelio sample, which overrepresents white-collar jobs relative to other occupations. To address this, we replicate the above decompositions, except we now re-scale occupation weights so that each 2-digit occupation category reflects their respective employment shares from the Bureau of Labor Statistics Occupational Employment Statistics (BLS-OES), which better reflect the economy-wide distribution of jobs. Appendix Figure A.6 and Appendix Table A.6 show that the results are substantively similar, with a slightly more positive net reallocation toward high-paying roles and somewhat larger declines for business and engineering occupations.

In sum, our analysis confirms AI significantly drove occupational employment shifts during our sample period, though its overall effects are modest. That said, a final caveat with our analysis is that it is restricted to the sample of publicly traded firms in Compustat. To the extent that large firms are more likely to adopt AI than small firms, and our sample is skewed towards these large firms, our estimates may overstate the overall economic impact of AI during our sample period.

4 Conclusion

Our work examines how AI adoption shapes firm dynamics and labor demand using firm–occupation variation in a novel measure of AI exposure. Three findings stand out. First, adoption is concentrated in large, productive firms and raises sales, profits, and TFP. Second, at the occupational level, higher-wage jobs are more exposed; average exposure reduces employment, while concentrated exposure reallocates labor toward complementary tasks and offsets these declines. Third, firm-wide adoption generates positive employment effects, consistent with productivity-driven increases in labor demand.

Overall, AI substitutes for labor at the task level but its net employment effects are muted by offsetting forces. Highly AI-exposed occupations experience declines in labor demand, yet this effect is mitigated by two offsetting forces. Consistent with our theoretical 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 leads to higher employment growth as firms expand their scale. Taken together, AI task exposure need not correlate with lower labor demand, even when AI is a task-level substitute.

We conclude by quantifying AI’s role in driving labor reallocation over the 2014 to 2023 period. Our estimates imply 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. Notably, AI exposure is greatest at the top of the wage distribution, yet these occupations experience modest increases in relative labor demand due to AI adoption. Our evidence emphasizes the importance of labor reallocation across tasks and firms, with the labor market effects partly shaped by which firms adopt AI. That said, our decomposition of the overall effects pertains to the first wave of AI adoption during the 2014–2023 period and need not generalize to the recent rise of generative AI. Even though the economic mechanisms we identify are likely to be similar, the set of tasks that are exposed to generative AI are likely to be different than those exposed to the AI applications in our sample.

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.
- Hicks, J. (1932). *The Theory of Wages*. Macmillan.
- 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.
- 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 B.1 for further details.

Figure 2: Sector \times 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 \times 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: The role of concentrated technology task exposure on labor demand




Note: This figure uses our model described in Section 2 of the main text and appendix A to study the impact of AI technology shocks for an occupation with two initially identical tasks. The left panel shows the overall occupational labor demand growth, and the right panel the growth in hours for the first task, for different shock configurations. The red line varies $m(\epsilon)$ and assumes both tasks are affected equally; the blue line varies $m(\epsilon)$ but allows technology shocks to be relatively more concentrated in the first task ($\epsilon_1 = 1.5 \times m(\epsilon)$); the green line assumes maximal concentration, with only the first task being exposed. We simulate from the exact solution of the model and assume the occupation is “small” within the firm, so that the shock doesn’t affect the model firm-level productivity index Z_f (which is held fixed at the initial value of 1). The calibration is $\nu = 3$; $\psi = 0.5$; $\chi = 2$; $\zeta = 1$; $\beta = 0.5$.

Figure 5: 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 6: 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 7: 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 3.1 for details.


Figure 8: Impact of AI on employment growth across the pay distribution




Note: This figure implements the decomposition of employment marginal effects from equation (33) in Section 3.5, 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 3.3 of the main text for details.

Figure 9: 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 (33) in Section 3.5. 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 16 percent of realized employment growth, of which at least 45 percent is attributable to direct occupation exposure. See section 3.3 of the main text for details.

Tables

Table 1: Characteristics of AI-using firms

	(1)	(2)	(3)	(4)	(5)
	Log Sales per worker	Log Sales	Log Profit	Log TFP	Log Average Salary
log(1 + AI uses)	0.117*** (6.87)	0.310*** (12.39)	0.415*** (17.48)	0.125*** (10.37)	0.109*** (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 \times 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 $\log(1 + N_{f,t})$ 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 \times 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

	(1)	(2)	(3)	(4)
Task-level AI Exposure	-2.00*** (-19.86)	-2.02*** (-20.09)	-2.09*** (-20.69)	-1.91*** (-19.21)
N	12,341,269	12,341,269	12,341,269	12,337,733
R ²	0.078	0.082	0.11	0.34
Occ Mean Exposure Control	X	X	X	
Firm Size Controls	X	X		
Industry \times Year FE		X		
Firm \times Year FE			X	
Firm \times Occ \times Year FE				X
Task \times Year FE	X	X	X	X

Note: This table includes estimates of Equation (3) from the main text. The unit of analysis is at the task–firm–year 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 2 of the main text. Standard errors are clustered by occupation–firm, with associated t -statistics reported in parentheses. Besides the designated fixed effects, we also control for the average task-level AI exposure of the occupation in specifications without the full complement of firm \times occupation \times 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) Sales	(2) Emp	(3) Profit	(4) TFP	(5) Sales	(6) Emp	(7) Profit	(8) TFP
log(1 + AI applications)	5.57*** (3.82)	3.97*** (3.48)	5.91*** (4.15)	5.32*** (5.80)	9.71*** (3.89)	6.84*** (3.75)	8.22** (3.20)	7.53*** (5.11)
N	12,757	13,225	11,652	6,065	12,282	12,688	11,246	6,035
R-sq	0.14	0.12	0.13	0.25	0.070	0.051	0.027	0.19
F-stat					2,145.5	2,287.3	1,842.7	1,038.7
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 (30). 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 two-stage least squares with the instrument $\log(1 + \text{Number of AI Uses})_{f,t}^{IV}$ IV defined in Section 3.1 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	-8.33*** (-12.31)	-8.72*** (-13.75)	-7.95*** (-12.50)	-5.70*** (-10.24)	-15.6*** (-10.26)	-14.6*** (-10.79)	-14.6*** (-16.49)	-10.4*** (-10.66)
AI Exposure Concentration	1.28** (3.11)	1.67*** (4.43)	1.91*** (4.66)	1.33*** (4.23)	9.07*** (6.54)	7.51*** (5.73)	7.50*** (8.24)	7.46*** (5.38)
log(1 + AI uses)	11.2*** (16.79)	10.3*** (15.66)			19.8*** (18.31)	19.7*** (16.29)		
N	1,454,540	1,454,539	1,454,255	1,454,255	1,452,305	1,452,305	1,452,211	1,452,211
R ²	0.11	0.19	0.50	0.59	0.074	0.060	0.024	-0.0014
F-stat (AI Exposure Average)					971.4	1,016.2	2,292.3	1,142.8
F-stat (AI Exposure Concentration)					1,022.4	1044.1	2164.6	465.3
F-stat (log(1 + AI uses))					1,031.9	1,466.3		
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 (13) 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 3.1 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	-2.14	1.42	0.91	0.20	19.0
Business and Financial	13	-9.79	5.90	1.73	-2.17	18.3
Architecture and Engineering	17	-6.64	2.52	2.83	-1.29	9.16
Science	19	1.58	-0.19	0.58	1.97	2.41
Community and Social Service	21	11.3	-5.53	-0.040	5.78	0.33
Legal	23	10.9	-6.29	2.93	7.58	0.73
Education and Library	25	10.4	-4.96	0.53	5.92	0.94
Arts, Entertainment, Media	27	8.70	-5.05	3.33	6.98	5.10
Healthcare Practitioners	29	4.61	-1.77	-0.060	2.77	1.85
Healthcare Support	31	6.87	-3.25	-0.48	3.14	0.46
Protective Service	33	9.04	-5.39	-1.05	2.60	0.42
Food Preparation and Serving	35	13.3	-6.42	-11.4	-4.56	2.62
Cleaning and Maintenance	37	14.8	-8.35	-4.86	1.60	0.44
Personal Care and Service	39	12.9	-6.43	-5.03	1.48	0.98
Sales and Related	41	1.69	-0.95	-2.48	-1.74	13.5
Office and Administrative	43	3.10	-2.50	0.43	1.03	10.5
Farming, Fishing, and Forestry	45	13.9	-7.56	-6.16	0.20	0.47
Construction and Extraction	47	6.87	-4.40	0.11	2.58	2.04
Installation and Repair	49	3.88	-3.24	-0.30	0.33	2.69
Production	51	5.37	-2.37	-2.03	0.97	3.91
Transportation	53	8.23	-4.17	-4.23	-0.17	4.19

Note: This table shows results from estimating the decomposition (33) 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 3.3 of the text for further details.

A Model Appendix

Here, we provide more details on the model derivation. To facilitate discussion of some extensions that appear later below, we will also allow for additional dimensions of heterogeneity relative to the setup presented in the main text.

A.1 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,

$$\bar{Y} = \left(\int_{\mathcal{F}} \alpha_f^{\frac{1}{\theta}} Y_f^{\frac{\theta-1}{\theta}} df \right)^{\frac{\theta}{\theta-1}}. \quad (\text{A.1})$$

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

$$Y_f = \left(\int_{\mathcal{O}} Y(o, f)^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}. \quad (\text{A.2})$$

where χ reflects the firm specific elasticity of substitution across occupation outputs. Going forward, we will suppress the firm subscript for brevity. 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 $\Theta = \frac{\theta}{\theta-1}$. 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

$$Y(o, f) = \left(\sum_j \alpha_o(j)^{\frac{1}{\psi}} y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}, \quad (\text{A.3})$$

where $\alpha_o(j)$ captures the importance of the task j for occupational output of occupation o . Here, ψ denotes the elasticity of substitution in firm f for occupation o across tasks, which determines the elasticity of labor demand for each task. To simplify the notation, we will suppress the firm subscript and occupation subscripts. Each task j in job (o, f) is produced by a labor input $l(j)$ and a capital input $k(j)$,

$$y(j) = \left(\gamma_j l(j)^{\frac{\nu-1}{\nu}} + (1 - \gamma_j) k(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}. \quad (\text{A.4})$$

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. There is an implicit firm and occupation subscript here that we have suppressed. 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. Below, section A.4 discusses an alternative interpretation of this production function which follows Acemoglu and Restrepo (2018) to develop a CES representation of an additional layer of “micro tasks”. In that setting, the constants in (A.4) need not sum to one and potentially change in response to technology.

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

$$l(j) = \alpha(j)^\beta h(j)^{1-\beta}. \quad (\text{A.5})$$

Here, the parameter $\beta \in (0, 1)$ captures the degree of decreasing returns to effort at the task level. Hence, when $\beta \rightarrow 1$, efficiency units of effort across tasks is exogenously fixed at $\alpha(j)$. The total number of hours a worker can supply across all J tasks is equal to one.

A.2 Equilibrium Conditions

Hours Allocation

Proposition 1 *At the o, f level, the optimal time allocation problem has the solution*

$$h(j) = \frac{\alpha(j)w(j)^{\frac{1}{\beta}}}{\sum_{k \in J} \alpha(k)w(k)^{\frac{1}{\beta}}}. \quad (\text{A.6})$$

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

$$\mathcal{L} = \sum_{j=1}^J w(j)\alpha(j)^\beta h(j)^{1-\beta} dj - \lambda \left(\sum_{j=1}^J h(j) dj - 1 \right). \quad (\text{A.7})$$

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

$$(1 - \beta) w(j)\alpha(j)^\beta h(j)^{-\beta} = \lambda. \quad (\text{A.8})$$

This leads to

$$h(j) = \alpha(j) \left[(1 - \beta) \frac{w(j)}{\lambda} \right]^{\frac{1}{\beta}}. \quad (\text{A.9})$$

Take the sum of both sides,

$$\sum_{j=1}^J h(j) = \sum_{j=1}^J \alpha(j) \left[(1 - \beta) \frac{w(j)}{\lambda} \right]^{\frac{1}{\beta}} = (1 - \beta)^{-\frac{1}{\beta}} \lambda^{\frac{1}{\beta}} \sum_{j=1}^J \alpha(j)w(j)^{\frac{1}{\beta}} = 1. \quad (\text{A.10})$$

Thus,

$$\lambda = (1 - \beta) \left(\sum_{j \in J} \alpha(j) w(j)^{\frac{1}{\beta}} \right)^{\beta}. \quad (\text{A.11})$$

Apply this to (A.9) yields (A.6). ■

Firm Output

Next, we derive the optimality conditions for the firm problem. The firm-level cost minimization problem can be expressed as

$$\min_{Y(o,f)} \int_O P(o, f) Y(o, f) \quad \text{s.t.} \quad Y_f = \left(\int_O Y(o, f)^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}. \quad (\text{A.12})$$

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

$$Y(o, f) = P(o, f)^{-\chi} Z_f^{-\chi} Y_f \quad (\text{A.13})$$

where

$$Z_f \equiv \left(\int_O P(o, f)^{1-\chi} \right)^{-\frac{1}{1-\chi}} = P_f^{-1} \quad (\text{A.14})$$

$P(o, f)$ denotes the marginal cost firm f it pays for the output of occupation o and P_f denotes the marginal cost firm f pays for the next unit of output. These prices need not be the same across firms. Suppose we label the price of firm f goods as \mathcal{P}_f . 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,

$$\begin{aligned} \mathcal{P}_f Y_f &= \Theta \int_O Y(o, f) P(o, f) \\ \mathcal{P}_f Y_f &= \Theta \int_O P(o, f)^{1-\chi} \left(\int_O \alpha_f(o) P(o, f)^{1-\chi} \right)^{\frac{\chi}{1-\chi}} Y_f \\ \mathcal{P}_f &= \Theta \left(\int_O P(o, f)^{1-\chi} \right)^{\frac{1}{1-\chi}} = \Theta Z_f^{-1} \end{aligned} \quad (\text{A.15})$$

Each firm faces the inverse demand curve

$$Y_f = \alpha_f \mathcal{P}_f^{-\theta} \mathcal{P}^{\theta} \bar{Y} \quad (\text{A.16})$$

where

$$\mathcal{P} \equiv \left(\int_{\mathcal{F}} \alpha_f \mathcal{P}_f^{1-\theta} \right)^{\frac{1}{1-\theta}}. \quad (\text{A.17})$$

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

$$Y_f = \alpha_f \mathcal{P}_f^{-\theta} \bar{Y}, \quad (\text{A.18})$$

and given the price above, this implies

$$Y_f = \alpha_f \Theta^{-\theta} Z_f^\theta \bar{Y} \quad (\text{A.19})$$

Occupation Output

Now, consider the occupation's task minimization problem (where we omit o and f notation for brevity)

$$\min_{y(j)} \sum_{j \in J} p(j) y(j) \quad s.t. \quad Y(o, f) = \left(\sum_{j \in J} \alpha(j)^{\frac{1}{\psi}} y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}} \quad (\text{A.20})$$

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.20):

$$y(j) = \alpha(j) p(j)^{-\psi} \left[\sum_{j \in J} \alpha(j) p(j)^{1-\psi} \right]^{\frac{\psi}{1-\psi}} Y(o, f) \quad (\text{A.21})$$

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

$$P(o, f) = \left[\sum_{j \in J} \alpha(j) p(j)^{1-\psi} \right]^{\frac{1}{1-\psi}} \quad (\text{A.22})$$

Using the above combined with equations (A.13) and (A.19), we get

$$y(j) = \alpha_f \alpha(j) p(j)^{-\psi} X(o, f)^{\chi-\psi} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \quad (\text{A.23})$$

where

$$X(o, f) = \left[\sum_{j \in J} \alpha(j) p(j)^{1-\psi} \right]^{-\frac{1}{1-\psi}} = P(o, f)^{-1} \quad (\text{A.24})$$

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

Task Labor and Capital Demand

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

$$\min_{l(j), k(j)} q(j)k(j) + w(j)\mathcal{M}(j)l(j) \quad s.t. \quad y(j) = \left[\gamma_j l(j)^{\frac{\nu-1}{\nu}} + (1 - \gamma_j)k(j)^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}} \quad (\text{A.25})$$

Going forward, we make the re-parameterization $a_j \equiv \gamma_j^\nu$, $b_j \equiv (1 - \gamma_j)^\nu$. After solving (A.25), the per-unit cost of task j equals

$$p(j) = \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{1}{1-\nu}} \quad (\text{A.26})$$

Using equation (A.26), we can rewrite (A.23) as

$$y(j) = \alpha_f \alpha(j) \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{-\psi}{1-\nu}} X(o, f)^{\chi-\psi} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \quad (\text{A.27})$$

Plugging in (A.27) to the CES Hicksian demand for $k(j)$ and $l(j)$ and imposing labor market clearing gives

$$l(j) = \alpha_f \alpha(j) \frac{a_j}{[\mathcal{M}(j)w(j)]^\nu} \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{\nu-\psi}{1-\nu}} \times X(o, f)^{\chi-\psi} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \quad (\text{A.28})$$

$$k(j) = \alpha_f \alpha(j) \frac{b_j}{q(j)^\nu} \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{\nu-\psi}{1-\nu}} \times X(o, f)^{\chi-\psi} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \quad (\text{A.29})$$

Task Labor Supply

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

$$L_o(j) = N(o, f) \alpha(j)^\beta h(j)^{1-\beta}. \quad (\text{A.30})$$

Using the properties of the Fréchet distribution and supposing the measure of workers available to the industry is fixed at \bar{N} , it follows that the expected measure of workers to job o in firm f is equal to

$$N(o, f) = \frac{\bar{N}}{\underbrace{\int_{f' \in \mathcal{F}} \alpha_{f'} \int_{o' \in \mathcal{O}} W(o', f')^\zeta \mathrm{d}o' \mathrm{d}f'}_{\bar{\zeta}}} \alpha_f W(o, f)^\zeta. \quad (\text{A.31})$$

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

$$W(o, f) \equiv \sum_{j \in J_o} \alpha(j)^\beta h(j)^{1-\beta} w(j). \quad (\text{A.32})$$

Given (A.6), the total earnings for that job are equal to

$$W(o, f) = \sum_{j \in J_o} \alpha(j)^\beta \left[\frac{\alpha(j)w(j)^{\frac{1}{\beta}}}{\sum_{k \in J} \alpha(k)w(k)^{\frac{1}{\beta}}} \right]^{1-\beta} w(j) = \frac{\sum_{j \in J} \alpha(j)w(j)^{\frac{1}{\beta}}}{(\sum_{k \in J} \alpha(k)w(k)^{\frac{1}{\beta}})^{1-\beta}} = \left[\sum_{j \in J} \alpha(j)w(j)^{\frac{1}{\beta}} \right]^\beta. \quad (\text{A.33})$$

Notice that as long as hours are flexible, $0 < \beta < 1$, 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

$$\begin{aligned} \alpha(j)^\beta h(j)^{1-\beta} N(o, f) &= \alpha(j)w(j)^{\frac{1-\beta}{\beta}} \left(\sum_{j \in J} \alpha(j)w(j)^{\frac{1}{\beta}} \right)^{\beta-1} \alpha_f \left[\sum_{j \in J} \alpha(j)w(j)^{\frac{1}{\beta}} \right]^{\zeta\beta} \bar{\zeta} \\ &= \alpha_f \alpha(j)w(j)^{\frac{1}{\beta}-1} \left(\sum_{j \in J} \alpha(j)w(j)^{\frac{1}{\beta}} \right)^{\beta-1+\zeta\beta} \bar{\zeta} \end{aligned} \quad (\text{A.34})$$

Replacing the left-hand-side of equation (A.28) with the equation for labor supply for task j yields a system of J equations in J unknowns—the task prices $w(j)$,

$$\begin{aligned} w(j)^{\frac{1}{\beta}} \left(\sum_{j \in J_o} \alpha(j)w(j)^{\frac{1}{\beta}} \right)^{\beta-1+\zeta\beta} \bar{\zeta} &= a_j \mathcal{M}(j)^{-\nu} w(j)^{1-\nu} \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{\nu-\psi}{1-\nu}} \\ &\quad \times X(o, f)^{\chi-\psi} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \end{aligned} \quad (\text{A.35})$$

where we will show below that the ratio of the marginal cost of $l(j)$ to the wage $\mathcal{M}(j)$ is constant. Dividing (A.35) for two different tasks yields

$$1 = \frac{a_j \mathcal{M}(j)^{-\nu} w(j)^{1-\nu-1/\beta} \left(a_j [\mathcal{M}(j)w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{\nu-\psi}{1-\nu}}}{a_k \mathcal{M}(k)^{-\nu_k} w(k)^{1-\nu_k-1/\beta} \left(a_k [\mathcal{M}(k)w(k)]^{1-\nu_k} + b_k q(k)^{1-\nu_k} \right)^{\frac{\nu_k-\psi}{1-\nu_k}}}. \quad (\text{A.36})$$

Since the term in both numerator and denominator is monontonic in the wage, an implication of (A.36) is that within the same occupation-firm combination, two tasks which have the same capital prices will have the same wage.

Using equation (A.26), we can get the following expression for occupation-firm level productivity:

$$X(o, f) = \left[\sum_{j \in J} \alpha(j) \left(a_j [\mathcal{M}(j) w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{1-\psi}{1-\nu}} \right]^{-\frac{1}{1-\psi}}. \quad (\text{A.37})$$

Monopsony wedge $\mathcal{M}(j)$

Here, we derive an important property which holds for $\mathcal{M}(j)$, the ratio of the marginal cost of $l(j)$ to the task j wage $w(j)$.

Proposition 2 *Given our assumptions on labor supply, the wedge between marginal cost of labor and wages for task j is constant across all tasks and satisfies*

$$\mathcal{M}(j) = 1 + \frac{1}{\zeta}. \quad (\text{A.38})$$

Proof. In order to derive the wage markdown for task $\mathcal{M}(j)$, 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 $w(j)$

$$\frac{\partial \log l(j)}{\partial \log w(j)} = \frac{\partial l(j)}{\partial w(j)} \frac{w(j)}{l(j)} = \left(\frac{1}{\beta} - 1 \right) + \left(1 - \frac{1}{\beta} + \zeta \right) h(j) = \left(\frac{1}{\beta} - 1 \right) (1 - h(j)) + \zeta h(j). \quad (\text{A.39})$$

We also need the cross-price terms

$$\frac{\partial \log l(j)}{\partial \log w(k)} = \left(1 - \frac{1}{\beta} + \zeta \right) h(k), \quad (\text{A.40})$$

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.39-A.40), we work with the identity

$$\log l(j) = \log(\alpha_f \bar{\zeta}) + \left(\frac{1}{\beta} - 1 \right) \log w(j) + [\beta - 1 + \zeta \beta] \log \left[\alpha(j) \sum_{j \in J} \exp(\log \alpha(j) + \frac{1}{\beta} \log w(j)) \right]. \quad (\text{A.41})$$

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

$$\frac{\partial}{\partial w(j)} \log \left[\alpha(j) \sum_{j \in J} \exp(\log \alpha(j) + \frac{1}{\beta} \log w(j)) \right] = \frac{1}{\beta} \frac{\alpha(j) w(j)^{1/\beta}}{\sum_{j' \in J} \alpha(j') w(j')^{1/\beta}} = \frac{1}{\beta} h(j). \quad (\text{A.42})$$

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

$$\frac{d \log w(k)}{d \log l(j)} \Big|_{\substack{d \log l(k)=0 \\ k \neq j}} = \mathbf{1}[k = j] \frac{\beta}{1 - \beta} + \frac{\frac{1}{\beta} - 1 - \zeta}{\left(\frac{1}{\beta} - 1 \right) \zeta} h(j), \quad (\text{A.43})$$

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

$$\frac{\partial[W(o, f)N(o, f)]}{\partial w(j)} = l(j)(1 + \zeta). \quad (\text{A.44})$$

To derive equation (A.44), we start with the fact that total wage earnings equals $\bar{\zeta}\alpha_f W(o, f)^{\zeta+1}$. Then by differentiating and using the definition of $l(j)$ from equation (A.34), we get

$$\frac{\partial[W(o, f)N(o, f)]}{\partial w(j)} = (1 + \zeta) \underbrace{\bar{\zeta}\alpha_f W(o, f)^\zeta}_{=N(o, f)} \frac{\alpha(j)^{1-\beta} w(j)^{\frac{1-\beta}{\beta}}}{\left[\sum_{j \in J} \alpha(j) w(j)^{\frac{1}{\beta}}\right]^{1-\beta}} \alpha(j)^\beta = (1 + \zeta)l(j). \quad (\text{A.45})$$

We can then combine these pieces (A.39, A.40, A.44) to compute marginal cost:

$$\left. \frac{\partial[W(o, f)N(o, f)]}{\partial l(j)} \right|_{\substack{d \log l(k)=0 \\ k \neq j}} = \sum_{k=1}^J \left. \frac{\partial[W(o, f)N(o, f)]}{\partial w(k)} \frac{w(k)}{l(j)} \frac{d \log w(k)}{d \log l(j)} \right|_{dl(k)=0} \quad (\text{A.46})$$

$$= (1 + \zeta) \left\{ \frac{\beta w(j)}{1 - \beta} + \sum_{k=1}^J \frac{w(k)l(k)}{l(j)} \frac{\frac{1}{\beta} - 1 - \zeta}{\left(\frac{1}{\beta} - 1\right)\zeta} h(j) \right\} \quad (\text{A.47})$$

$$= (1 + \zeta) \left\{ \frac{\beta w(j)}{1 - \beta} + w(j) \frac{\frac{1}{\beta} - 1 - \zeta}{\left(\frac{1}{\beta} - 1\right)\zeta} \right\} \quad (\text{A.48})$$

$$= \frac{1 + \zeta}{\zeta} w(j).$$

Recalling that $\mathcal{M}(j)$ is the ratio of marginal cost to the wage, and rearranging, we get our desired result in (A.38). ■

A.3 Microfoundations for weights α

While the weights α_f and $\alpha(j)$ are simply treated as reduced form parameters capturing task importance in the main text, this section discusses an alternative interpretation of these parameters which involves an additional layer of aggregation. This creates a connection between our model and frameworks such as Romer (1990) which emphasize the creation of new product varieties as a source of growth.

Firm level shifter α_f as the number of product varieties

A natural interpretation of the firm level shifter is that it captures the scope of production captured by the firm. Concretely, let us suppose that Y_f is itself a CES composite of an interval of N_f

different product varieties

$$\bar{Y} = \left[\int_{f \in \mathcal{F}} \int_0^{N_f} \tilde{\alpha}_f^{1/\theta} Y_{f,v}^{\frac{\theta-1}{\theta}} dv df \right]^{\frac{\theta}{\theta-1}} \quad (\text{A.49})$$

where $\tilde{\alpha}_f$ is an exogenous firm-level taste shifter.¹⁵ Following the usual properties of CES preferences, we can write the aggregate price index for industry output as

$$\mathcal{P} = \left[\int_{f \in \mathcal{F}} \int_0^{N_f} \tilde{\alpha}_f P_{f,v}^{1-\theta} dv df \right]^{\frac{1}{1-\theta}}, \quad (\text{A.50})$$

and the demand curve for variety v for firm f as

$$Y_{f,v} = \tilde{\alpha}_f \left[\frac{P_{f,v}}{\mathcal{P}} \right]^{-\theta} \bar{Y}. \quad (\text{A.51})$$

Now, we introduce a composite firm-level good Y_f , defined as

$$Y_f \equiv \left[\left(\frac{1}{N_f} \right)^{\frac{1}{\theta}} \int_0^{N_f} Y_{f,v}^{\frac{\theta-1}{\theta}} dv \right]^{\frac{\theta}{\theta-1}} \iff N_f^{\frac{1}{\theta}} Y_f^{\frac{\theta-1}{\theta}} = \int_0^{N_f} Y_{f,v}^{\frac{\theta-1}{\theta}} dv, \quad (\text{A.52})$$

so aggregate output is

$$\bar{Y} = \left[\int_{f \in \mathcal{F}} \int_0^{N_f} [\tilde{\alpha}_f N_f]^{1/\theta} Y_{f,v}^{\frac{\theta-1}{\theta}} dv df \right]^{\frac{\theta}{\theta-1}} \equiv \left[\int_{f \in \mathcal{F}} \alpha_f^{1/\theta} Y_f^{\frac{\theta-1}{\theta}} dv df \right]^{\frac{\theta}{\theta-1}} \quad (\text{A.53})$$

with $\alpha_f \equiv \tilde{\alpha}_f N_f$.

Next, let's verify that other components of the firm problem can be written in terms of this composite Y_f . Notice that we can define a firm level price index for this composite good as follows:

$$P_f \equiv \left[\frac{1}{N_f} \int_0^{N_f} P_{f,v}^{1-\theta} dv \right]^{\frac{1}{1-\theta}} \iff \int_0^{N_f} P_{f,v}^{1-\theta} dv = N_f P_f^{1-\theta}. \quad (\text{A.54})$$

Rearranging the identity $\frac{P_{f,v} Y_{f,v}}{\mathcal{P} \bar{Y}} = \tilde{\alpha}_f \left[\frac{P_{f,v}}{\mathcal{P}} \right]^{1-\theta}$ and integrating, we can rewrite the industry price index as

$$\mathcal{P} = \left[\int_f \tilde{\alpha}_f \int_0^{N_f} P_{f,v}^{1-\theta} dv df \right]^{\frac{1}{1-\theta}} = \left[\int_f \alpha_f P_f^{1-\theta} df \right]^{\frac{1}{1-\theta}}, \quad (\text{A.55})$$

which is a single layer index involving the composite price P_f and the new weight $\alpha_f \equiv N_f \tilde{\alpha}_f$. Further note that we can write the demand for a firm's composite good as

$$Y_f = \alpha_f \left[\frac{P_f}{\mathcal{P}} \right]^{-\theta} \bar{Y}. \quad (\text{A.56})$$

¹⁵This argument naturally extends to explicitly allowing the elasticity of substitution to differ across varieties offered by the same firm via an additional nested structure.

All that remains is to verify that the isomorphism holds on the production technology. We will further suppose that the firm has the same production technology for each variety. Given our assumptions above, in order to most efficiently produce one unit of Y_f , the firm will therefore produce the same amount c of each variety, where c satisfies

$$1 = \left[\left(\frac{1}{N_f} \right)^{\frac{1}{\theta}} \int_0^{N_f} c^{\frac{\theta-1}{\theta}} dv \right]^{\frac{\theta}{\theta-1}} = c \left[N_f^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}} = c N_f \quad \iff \quad c = \frac{1}{N_f}. \quad (\text{A.57})$$

Hence, production of one unit of Y_f requires producing measure $1 = N_f \cdot c$ total units across the underlying varieties, regardless of the value of N_f .

Hence we can interpret the taste shifter α_f as partially including a component of firm productivity which includes the “love of variety” across N_f differentiated products it produces. To the extent that N_f varies across firms, this component generates additional dispersion in firm scale conditional on productivity. Further, if AI facilitates greater customization of products and services (see, e.g., [Babina et al., 2024](#), for related evidence), an additional source of spillovers can result from increases in N_f which boost firm labor demand holding Z_f constant.

Further note that the taste shifter α_f symmetrically enters both the labor supply and labor demand blocks in the equations above. A natural reason for this would be if $\alpha_f = N_f$ and jobs are posted at the variety level. In such a case, N_f would also emerge as a multiplicative constant in the Fréchet labor supply block, shifting employment but not impacting per-worker wages.¹⁶ If such an assumption does not hold, N_f would shift firm demand in (A.28) but not in the labor supply equation (A.34), which would introduce an additional multiplicative constant into the labor market clearing condition (A.35).

Task-level $\alpha(j)$ shifter as measure of micro-level tasks

We can derive our specification of task-level labor supply and demand via a bottom up aggregation of micro-level tasks. On the firm side, we can iterate on the argument above to interpret $y(j)$ as a composite of task j type output over an interval of length $\alpha(j)$

$$y(j) = \left[\alpha(j)^{-\frac{1}{\psi}} \int_0^{\alpha(j)} y_j(i)^{\frac{\psi-1}{\psi}} di \right]^{\frac{\psi}{\psi-1}}, \quad (\text{A.58})$$

where $y_j(\tau)$ takes the same form as the production function $y(j)$ from (A.4), with task level labor and capital inputs $l_j(i)$ and $k_j(i)$, respectively.

While $k(j)$ is perfectly divisible across tasks without frictions, there are frictions which prevent workers from fully specializing their time in a single atomistic task. Instead, effective labor supply to each task is

$$l_j(i) = h_j(i)^{1-\beta} \quad (\text{A.59})$$

¹⁶This property follows from the fact that the max of N_f iid Fréchet draws is also Fréchet distributed, but with a scale parameter equal to N_f times the original scale parameter.

Suppose that the worker has already decided to allocate $h(j)$ in total hours to task j . Then, since she is equally productive across all micro-tasks in interval j , the earnings maximizing choice of $h_\tau(j)$ is to divide effort equally among tasks. So, $h_j(i) = h(j)/\alpha(j)$, and therefore $l_j(i) = \left[\frac{h(j)}{\alpha(j)}\right]^{1-\beta}$.

Notice therefore that total efficiency units of labor satisfy $l(j) = h(j)^{1-\beta} \alpha(j)^{\beta-1} \int_0^{\alpha(j)} di = \alpha(j)^\beta h(j)^{1-\beta}$. Likewise, if the firm has total capital $k(j)$ to allocate over the interval, it will optimally allocate $\frac{k(j)}{\alpha(j)}$ to each task within the interval. Because the micro-task level CES production function exhibits constant returns to scale,

$$y_j(i) = \frac{1}{\alpha(j)} \left(\gamma_j l(j)^{\frac{\nu-1}{\nu}} + (1 - \gamma_j) k(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}} \quad (\text{A.60})$$

and

$$\left[\alpha(j)^{-\frac{1}{\psi}} \int_0^{\alpha(j)} y_j(i)^{\frac{\psi-1}{\psi}} di \right]^{\frac{\psi}{\psi-1}} = \left(\gamma_j l(j)^{\frac{\nu-1}{\nu}} + (1 - \gamma_j) k(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}} \underbrace{\left[\alpha(j)^{-\frac{1}{\psi}} \alpha(j)^{\frac{1-\psi}{\psi}} \int_0^{\alpha(j)} di \right]^{\frac{\psi}{\psi-1}}}_{=1} = y(j), \quad (\text{A.61})$$

which exactly corresponds with the production function and aggregate efficiency units of labor functional forms which we assumed.

As was the case above, we could also add constant terms $\tilde{\alpha}(j)$ into the expressions. However, this formulation provides one concrete interpretation for the $\alpha(j)$. It also suggests that it is potentially sensible to impose that these coefficients sum to 1 across tasks performed by the same occupation.

A.4 Alternative formulation of automation

In this subsection, we briefly discuss an alternative interpretation of our model which obtains from applying the notation of automation developed in (Acemoglu and Restrepo, 2018, henceforth ‘‘AR’’). More precisely, we embed a production process analogous to theirs to describe the innermost layer of production, then layer additional CES layers on top of it.

As we develop more formally below, automation in this framework can manifest in reduced form as a shift in the constants a_j and b_j which appear in the equilibrium conditions of the model. (A minor difference is that the coefficients in the CES production function in (A.4) no longer sum to one.) This formulation provides a potential rationale for allowing for a technological change $\varepsilon(j)$ to shift these constants, with

$$\frac{d \log a_j}{d \log \varepsilon(j)} \leq 0 \quad \text{and} \quad \frac{d \log b_j}{d \log \varepsilon(j)} \geq 0, \quad (\text{A.62})$$

respectively. For this reason, we also develop comparative statics to address this case as well.

We now interpret j as an index of the set of tasks performed an occupation (suppressing firm and occupation from the notation for brevity). Specifically, task j output $y(j)$ output is produced via a

CES technology over a unit interval of micro-tasks:

$$y(j) = \left[\int_0^1 y_i(j)^{\frac{\nu-1}{\nu}} di \right]^{\frac{\nu}{\nu-1}}, \quad (\text{A.63})$$

where micro-task-level output is produced with a linear technology

$$y_j(i) = k_j(i) + \phi_j(i)l_j(i). \quad (\text{A.64})$$

Here, $\phi_j(i) > 0$ represents the effectiveness of human labor and capital for a specific micro-task respectively. We assume that tasks are ordered such that $\phi_j(i)$ is strictly increasing in i . As in AR, we will define the automation threshold I_j as frontier of capital's ability to replace human labor, where the idea is that tasks in which capital has the largest relative productivity advantage will be automated first. Thus, given these assumptions, the cost of a micro-task is

$$p_j(i) = \begin{cases} \min\left(q(j), \frac{w(j)}{\phi_j(i)}\right) & \text{if } i \leq I_j, \\ \frac{w(j)}{\phi_j(i)} & \text{if } i > I_j, \end{cases} \quad (\text{A.65})$$

where $\phi_j(i)$ and $\varphi_j(i)$ are both positive.

Analogously with AR, we also define \tilde{I}_j as the point at which capital and labor are equally efficient in production:

$$\frac{w(j)}{q(j)} = \phi_j(\tilde{I}_j) > \phi_j(I_j). \quad (\text{A.66})$$

A key condition in order for shifts in the automation frontier I_j to influence labor demand is that $\tilde{I}_j > I_j$, the predominant case emphasized by AR. That is, for all tasks that AI helps to automate, capital is strictly preferred to human labor. Suppose we assume the opposite case holds. Then, the marginal cost of labor would be lower than capital, and so shifts in the automation frontier I_j would have no effect locally.

Solving for the total output for a given threshold $I_j < \tilde{I}_j$, AR show that

$$y(j) = \left[I_j^{\frac{1}{\nu}} k(j)^{\frac{\nu-1}{\nu}} + \left(\int_{I_j}^1 \phi_j(i)^{\nu-1} di \right)^{\frac{1}{\nu}} l(j)^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}} \quad (\text{A.67})$$

Hence, there is an isomorphic representation of their model as a simple CES index of capital and labor. Given this representation, we can define the following:

$$a_j = \int_{I_j}^1 \phi_j(i)^{\nu-1} di \quad b_j = I_j, \quad (\text{A.68})$$

where $l(j)$ and $k(j)$ are defined by

$$l(j) \equiv \int_{I_j}^1 l_j(i) di, \quad \frac{l_j(i)}{l(j)} = \frac{\phi_j(i)^{\nu-1}}{\int_{I_j}^1 \phi_j(i')^{\nu-1} di'}, \quad \text{and} \quad k(j) \equiv \int_{I_j}^1 k_j(i) di. \quad (\text{A.69})$$

Then, all remaining equilibrium conditions are unchanged with these modified production function parameters.

To consider automation shocks in the spirit of AR, we can allow technology to shift the automation threshold I_j (and in turn, a_j and b_j), while keeping the capital price $q(j)$ constant. If, as in the main text we additionally linearize around a case with symmetric prices and capital shares, we still obtain an estimating equation for the direct effect identical to (13) in the main text, where mean $m(\varepsilon)$ and concentration $C(\varepsilon)$ are defined in the same way as (14) and (15). The only difference is that the coefficients have different economic interpretations. The approximation becomes

$$\Delta \log W(o, f) \approx \eta_m m(\varepsilon) + \frac{1}{2\beta} \eta_o^2 C(\varepsilon) + \text{Spillovers} \quad (\text{A.70})$$

where

$$\eta_m \equiv \frac{\eta_a + (\nu - \chi)\pi}{\zeta + \nu s_k + \chi(1 - s_k)}, \quad \eta_o = \frac{\eta_a + (\nu - \psi)\pi}{\frac{1}{\beta} - 1 + \nu - (\nu - \psi)s_l} \quad (\text{A.71})$$

and

$$\begin{aligned} \eta_a &= -\frac{I\phi(I)^{\nu-1}}{\int_I^1 \phi(i)^{\nu-1} di} < 0 \\ \pi &= I p(j)^{\nu-1} q^{1-\nu} \left(1 - \underbrace{\left(\frac{Mw}{q\phi(I)} \right)^{1-\nu}}_{\geq 1} \right) (1 - \nu)^{-1} < 0 \end{aligned} \quad (\text{A.72})$$

The impact of the automation shocks is driven by two key elasticities. The first is a task substitution term η_a , which is always negative. The second involves the product of $\nu - \chi$ with π , a factor which captures the elasticity of the marginal cost of task j output to $\varepsilon(j)$. η_a is negative given that $\phi(I) > 0$, while π is negative given our assumption that $I < \tilde{I}$ from (A.66). Analogously, the factor η_o which drives then concentration coefficient is influenced by the direct substitution term η_a . Hence, a model in which technology shifts the automation thresholds delivers similar comparative statics as the baseline model (in which q shifts) with a larger elasticity of substitution ν .

Another advantage of this framework is that one can easily allow for both shifts in the CES coefficients and changes in the cost of task-specific capital q . If in addition to the AR effects we allow the cost of capital q to shift, π would shift further downwards by a factor equal to the capital share s_k times the elasticity of q with respect to the technology shock. For additional details about the more general case, which also provides formulas which apply outside of the symmetric case, see section A.6.

A.5 Task wage elasticities: general case

In each of the next sections, we will assume use log-linear approximations to derive the elasticity of task wages with respect to various changes in the economic environment. Our analysis proceeds in several steps. We first illustrate how task-level wages respond to an exogenous shock to labor

demand induced by aggregate output, holding other primitives constant. Next, we consider different shifts induced by technology.

Before proceeding, we introduce some convenient intermediate quantities which will appear in our comparative statics derivations. Define

$$s_l(j) \equiv \frac{a_j M^{1-\nu} w(j)^{1-\nu}}{a_j M^{1-\nu} w(j)^{1-\nu} + b_j q(j)^{1-\nu}}, \quad s_k(j) = 1 - s_l(j) \quad (\text{A.73})$$

as the task-level labor share and capital share of output, respectively. We can also define the occupation and firm-level share of productivity by task.

$$s_p(j) \equiv \frac{\alpha(j)(a_j M^{1-\nu} w(j)^{1-\nu} + b_j q(j)^{1-\nu})^{\frac{1-\psi}{1-\nu}}}{\sum_{k=1}^J \alpha(k)(a_k M^{1-\nu} w(k)^{1-\nu} + b_k q(k)^{1-\nu})^{\frac{1-\psi}{1-\nu_k}}} \quad (\text{A.74})$$

A helpful accounting identity for these derivations is

$$s_p(j)s_l(j) = h(j)\bar{s}_l \quad (\text{A.75})$$

where \bar{s}_l is the weighted average labor share of marginal cost (note that this differs from the labor share of input costs because of monopsony markdowns) which is defined as

$$\bar{s}_l \equiv \sum_{j \in J_o} s_p(j)s_l(j). \quad (\text{A.76})$$

The accounting identity holds as an individual's fraction of hours allocated to task j , $h(j)$, will equal the share of the wage bill coming from task j .

In addition, we define the following composite parameters:

$$\begin{aligned} \Gamma(j) &= \left(\frac{1}{\beta} - 1 + \nu s_k(j) + \psi s_l(j) \right)^{-1} \\ \bar{\Gamma} &= \sum_{j \in J_o} h(j)\Gamma(j) \\ \kappa &= \left(1 - \frac{1}{\beta} + \zeta + (\chi - \psi)\bar{s}_l \right). \end{aligned} \quad (\text{A.77})$$

We will suppress the occupation subscript on the second expression.

Output Shock

Suppose total output is hit with an exogenous shock holding all else constant. Output develops in the following way

$$\tilde{Y} = \bar{Y} e^\varepsilon \quad (\text{A.78})$$

and task prices change

$$w(j) = w(j) e^{\eta_Y(j)\epsilon}, \quad \forall j \quad (\text{A.79})$$

with capital costs staying constant. Starting with equation A.35 and plugging in equation A.37 gives us

$$\begin{aligned} w(j)^{\frac{1}{\beta}} \left(\sum_{k \in J_o} \alpha(k) w(k)^{\frac{1}{\beta}} \right)^{\beta-1+\zeta\beta} \bar{\zeta} &= a_j \mathcal{M}^{-\nu} w(j)^{1-\nu} \left(a_j [\mathcal{M} w(j)]^{1-\nu} + b_j q(j)^{1-\nu} \right)^{\frac{\nu-\psi}{1-\nu}} \\ &\times \left[\sum_{k \in J_o} \alpha(k) \left(a_k [\mathcal{M} w(k)]^{1-\nu} + b_k q(k)^{1-\nu} \right)^{\frac{1-\psi}{1-\nu}} \right]^{-\frac{\chi-\psi}{1-\psi}} Z_f^{\theta-\chi} \Theta^{-\theta} \bar{Y}. \end{aligned} \quad (\text{A.80})$$

Taking equation A.80, plugging in the new wages and output and dividing through by the original gives the following.

$$\begin{aligned} e^{\eta_Y(j)\frac{1}{\beta}\epsilon} \left(\sum_{k \in J_o} h(k) e^{\eta_Y(k)\frac{1}{\beta}\epsilon} \right)^{\beta-1+\zeta\beta} &= e^{\eta_Y(j)(1-\nu)\epsilon} (s_l(j) e^{\eta_Y(j)(1-\nu)\epsilon} + s_k(j))^{\frac{\nu-\psi}{1-\nu}} \\ &\times \left(\sum_{k \in J_o} s_p(k) (s_l(k) e^{\eta_Y(k)(1-\nu)\epsilon} + s_k(k))^{\frac{1-\psi}{1-\nu}} \right)^{-\frac{\chi-\psi}{1-\psi}} e^\epsilon \end{aligned} \quad (\text{A.81})$$

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

$$\eta_z(j) = \frac{-\sum_{k \in J_o} h(k) \eta_z(k) (1 - \frac{1}{\beta} + \zeta + (\chi - \psi) \bar{s}_l)}{\frac{1}{\beta} - 1 + \nu s_k(j) + \psi s_l(j)} \quad (\text{A.82})$$

Plugging in the expressions from A.77 we get

$$\begin{aligned} \eta_Y(j) &\equiv \left(\frac{1}{\beta} - 1 + \nu s_k(j) + \psi s_l(j) \right)^{-1} \left(1 + \sum_{k \in J_o} h(k) \frac{(1 - \frac{1}{\beta} + \zeta + (\chi - \psi) \bar{s}_l)}{\frac{1}{\beta} - 1 + \nu_k s_k(k) + \psi s_l(k)} \right)^{-1} \\ &= \frac{\Gamma(j)}{1 + \kappa \bar{\Gamma}} \end{aligned} \quad (\text{A.83})$$

We can extract two other interesting elasticities from this proof. First, an analogous argument yields that the elasticity of task wages with respect to an exogenous shock to firm level productivity satisfies:

$$\eta_z(j) \equiv (\theta - \chi) \eta_Y(j) = (\theta - \chi) \frac{\Gamma(j)}{1 + \kappa \bar{\Gamma}}. \quad (\text{A.84})$$

Second, if we wanted the elasticity of the overall wage $W(o, f)$ with respect to output shocks, we get the following formula:

$$\bar{\eta}_Y = \sum_{j \in J_o} h(j) \eta_Y(j) = \frac{\bar{\Gamma}}{1 + \kappa \bar{\Gamma}}. \quad (\text{A.85})$$

Technology Improvement Shocks

Next, we turn to our main results on technological improvements. We begin by considering the partial derivative of task-level wages to technological improvements specific to task j , then will build up sufficient statistics. We will allow for capital automation and deepening effects. Motivated by the alternative formulation of technological change from section A.4, we will use elasticities for capital cost, as well as the CES capital/labor shares. Since this is an analysis on the effect on wages, we will have the target elasticity on the wage term. We give below the definition for how each evolves after a shock ε on task j .

For j

$$\tilde{a}_j = a_j e^{\eta_a(j)\varepsilon} \quad \tilde{b}_j = b_j e^{\eta_b(j)\varepsilon} \quad \tilde{q}(j) = q(j) e^{\eta_q(j)\varepsilon} \quad (\text{A.86})$$

For all k

$$\tilde{w}(k) = w(k) e^{\eta(k)\varepsilon} \quad (\text{A.87})$$

If we replace the above into (A.35) for $k \in J_o$, we get two equations, one for the ‘shocked’ task and another equation which is common to all unshocked tasks. After dividing both sides of each equation with equation (A.35) evaluated at the pre-shock equilibrium, we obtain the following at j :

$$e^{\eta(j)\frac{1}{\beta}\varepsilon} \left(\sum_{i=1}^J h(i) e^{\eta(i)\frac{1}{\beta}\varepsilon} \right)^{\beta-1+\zeta\beta} = e^{\eta_a(j)\varepsilon} e^{\eta(j)(1-\nu)\varepsilon} \times \left(s_l(j) e^{\eta_a(j)\varepsilon} e^{\eta(j)(1-\nu)\varepsilon} + s_k(j) e^{\eta_b(j)\varepsilon} e^{\eta_q(j)(1-\nu)\varepsilon} \right)^{\frac{\nu-\psi}{1-\nu}} \left(\frac{\tilde{X}(o, f)}{X(o, f)} \right)^{\chi-\psi} \quad (\text{A.88})$$

as well as the additional equations for other tasks ($k \neq j$)

$$e^{\eta(k)\frac{1}{\beta}\varepsilon} \left(\sum_{i=1}^J h(i) e^{\eta(i)\frac{1}{\beta}\varepsilon} \right)^{\beta-1+\zeta\beta} = e^{\eta(k)(1-\nu)\varepsilon} \left(s_l(k) e^{\eta(k)(1-\nu)\varepsilon} + s_k(k) \right)^{\frac{\nu-\psi}{1-\nu}} \left(\frac{\tilde{X}(o, f)}{X(o, f)} \right)^{\chi-\psi}. \quad (\text{A.89})$$

Then we can write

$$\frac{\tilde{X}(o, f)}{X(o, f)} = \left(s_p(j) \left(e^{\eta_a(j)\varepsilon} s_l(j) e^{\eta(j)(1-\nu)\varepsilon} + e^{\eta_b(j)\varepsilon} s_k(j) e^{\eta_q(j)(1-\nu)\varepsilon} \right)^{\frac{1-\psi}{1-\nu}} + \sum_{k \neq j} s_p(k) \left(s_l(k) e^{\eta(k)(1-\nu)\varepsilon} + s_k(k) \right)^{\frac{1-\psi}{1-\nu}} \right)^{-\frac{1}{1-\psi}}. \quad (\text{A.90})$$

Taking logs of both sides (A.88) and (A.89), differentiating with respect to ε , evaluating it at

$\epsilon = 0$ yields J linear equations. One useful term here is:

$$\frac{\partial}{\partial \epsilon} \log \left(\frac{\tilde{X}(o, f)}{X(o, f)} \right) \Big|_{\epsilon=0} = - \left[s_p(j) \left(\frac{\eta_a(j) s_l(j) + (\eta_b(j) + \eta_q(j)(1 - \nu)) s_k(j)}{1 - \nu} \right) + \sum_{k=1}^J s_p(k) s_l(k) \eta(k) \right] \quad (\text{A.91})$$

Next, let's define the following quantity

$$\pi(j) = \frac{\eta_a(j) s_l(j) + (\eta_b(j) + \eta_q(j)(1 - \nu)) s_k(j)}{1 - \nu}. \quad (\text{A.92})$$

Notice that $\pi(j)$ has a natural economic interpretation: it is the elasticity of the marginal cost of type j output $p(j)$ to $\epsilon(j)$, holding wages fixed. The second sum in (A.91) captures changes in productivity induced by movements in task-level wages.

Collecting these pieces together yields an equation which must hold for each of the tasks:

$$\begin{aligned} \frac{1}{\beta} \eta(k) + \frac{\beta - 1 + \zeta \beta}{\beta} \sum_{i=1}^J h(i) \eta(i) &= \eta(k)(1 - \nu_k) + (\nu_k - \psi) s_l(k) \eta(k) \\ &- (\chi - \psi) \left[s_p(j) \pi(j) + \sum_{k=1}^J s_p(k) s_l(k) \eta(k) \right] \\ &+ \left(\eta_a + (\nu - \psi) \pi(j) \right) \mathbf{1}_{k=j}. \end{aligned} \quad (\text{A.93})$$

In the above equation, the second line corresponds to the $\frac{\partial}{\partial \epsilon} \log \left(\frac{\tilde{X}(o, f)}{X(o, f)} \right) \Big|_{\epsilon=0}$ term. One potential way to solve for η_k goes as follows. First, isolate all η_k terms onto one side as follows:

$$\begin{aligned} \left((1 - \nu_k) + (\nu_k - \psi) s_l(k) - \frac{1}{\beta} \right) \eta_k &= \frac{\beta - 1 + \zeta \beta}{\beta} \sum_{i=1}^J h(i) \eta_i \\ &+ (\chi - \psi) \left[s_p(j) \pi(j) + \sum_{k=1}^J s_p(k) s_l(k) \eta(k) \right] \\ &- \left(\eta_a + (\nu - \psi) \pi(j) \right) \mathbf{1}_{k=j} \end{aligned} \quad (\text{A.94})$$

To make notation easier, define the following:

$$\begin{aligned} A_k &= \frac{\beta - 1 + \zeta \beta}{\beta} h(k) + (\chi - \psi) s_p(k) s_l(k) = h(k) \kappa \\ B &= s_p(j) (\chi - \psi) \pi(j) \\ C &= - \left(\eta_a + (\nu - \psi) \pi(j) \right) \end{aligned} \quad (\text{A.95})$$

With these definitions, rewrite the system of equations as follows:

$$\begin{aligned}
-\eta_k \cdot \Gamma(k)^{-1} &= \sum_{i=1}^J A_i \eta_i + B + C \mathbf{1}_{k=j} \\
-A_k \cdot \eta_k &= A_k \Gamma(k) \sum_{i=1}^J A_i \eta_i + A_k (B + C \mathbf{1}_{k=j}) \Gamma(k)
\end{aligned} \tag{A.96}$$

Sum both sides from $k = 1$ to J

$$\begin{aligned}
-\sum_{k=1}^J A_k \eta_k &= \sum_{k=1}^J A_k \Gamma(k) \left(\sum_{i=1}^J A_i \eta_i \right) + \sum_{k=1}^J A_k (B + C \mathbf{1}_{k=j}) \Gamma(k) \\
&= \left(\sum_{k=1}^J A_k \Gamma(k) \right) \left(\sum_{i=1}^J A_i \eta_i \right) + \sum_{k=1}^J A_k (B + C \mathbf{1}_{k=j}) \Gamma(k) \\
\sum_{k=1}^J A_k \eta_k &= -\frac{\sum_{k=1}^J A_k (B + C \mathbf{1}_{k=j}) \Gamma(k)}{1 + \sum_{k=1}^J \Gamma(k) A_k \Gamma(k)}
\end{aligned} \tag{A.97}$$

Plug in this final equality back into the original equation

$$\begin{aligned}
-\eta_k \cdot \Gamma(k)^{-1} &= -\frac{\sum_{i=1}^J A_i (B + C \mathbf{1}_{i=j}) \Gamma(i)}{1 + \sum_{i=1}^J A_i \Gamma(i)} + (B + C \mathbf{1}_{k=j}) \\
\eta_k &= \Gamma(k) \frac{-(B + C \mathbf{1}_{k=j}) + \sum_{i=1}^J A_i (B + C \mathbf{1}_{i=j}) \Gamma(i) - (B + C \mathbf{1}_{k=j}) \sum_{i=1}^J A_i \Gamma(i)}{\left(1 + \sum_{i=1}^J A_i \Gamma(i)\right)}
\end{aligned} \tag{A.98}$$

If $k \neq j$, the numerator simplifies as follows:

$$\sum_{i=1}^J A_i (B + C \mathbf{1}_{i=j}) \Gamma(i) - B \sum_{i=1}^J A_i \Gamma(i) = C A_j \Gamma(j) \tag{A.99}$$

If $k = j$, then the numerator reduces as follows.

$$\begin{aligned}
\sum_{i=1}^J A_i (B + C \mathbf{1}_{i=j}) \Gamma(i) - (B + C) \sum_{i=1}^J A_i \Gamma(i) &= A_j C \Gamma(j) - C \sum_{i=1}^J A_i \Gamma(i) \\
&= -C \sum_{i \neq j}^J A_i \Gamma(i)
\end{aligned} \tag{A.100}$$

For both, the denominator reduces to

$$1 + \sum_{i=1}^J A_i \Gamma(i) = 1 + \sum_{i=1}^J h(i) \kappa \Gamma(i) = 1 + \kappa \bar{\Gamma} \tag{A.101}$$

Finally, the equations for η_k are defined below.

For j ,

$$\eta_j = \Gamma(j) \frac{-B - C - C \sum_{i \neq j}^J A_i \Gamma(i)}{(1 + \kappa \bar{\Gamma})} \quad (\text{A.102})$$

For $k \neq j$

$$\eta_k = \Gamma(k) \frac{-B + A_j C \Gamma(j)}{(1 + \kappa \bar{\Gamma})} \quad (\text{A.103})$$

Our end goal is to try to find a constant part of the elasticity that does not vary with which task we examine. We want to split the task between a general level elasticity that captures the cross-task elasticity, and a type of own task elasticity that is the difference between the true own task elasticity and the cross-task elasticity. We can start with the latter. Note that if we plug in j to the latter equation, we get the cross-task elasticity equivalent, denoted by $\hat{\eta}_j$. Taking the difference we can get the target own-task elasticity.

$$\begin{aligned} \eta_j - \hat{\eta}_j &= \Gamma(j) \frac{-B - C - C \sum_{i \neq j}^J A_i \Gamma(i)}{(1 + \kappa \bar{\Gamma})} - \Gamma(k) \frac{-B + A_j C \Gamma(j)}{(1 + \kappa \bar{\Gamma})} \\ &= \Gamma(j) \frac{-C - C \sum_{i=1}^J A_i \Gamma(i)}{(1 + \kappa \bar{\Gamma})} \\ &= -\Gamma(j) C \\ &= \left(\eta_a + (\nu - \psi) \pi(j) \right) \Gamma(j) \equiv \eta_o(j) \end{aligned} \quad (\text{A.104})$$

It will be helpful for our analysis to extract the constant term the elasticity expression.

$$\eta_k = \eta_c(j, k) \equiv \eta_c \cdot f(j, k) \quad (\text{A.105})$$

Notice that only the denominator of η_k is not task or shock specific and is thus the only candidate for the invariant cross-task term. Thus let us define

$$\eta_c \equiv \frac{1}{1 + \kappa \bar{\Gamma}} \quad (\text{A.106})$$

and we can rephrase the elasticity of task wages with respect to total output shocks as

$$\eta_Y(j) = \eta_c \Gamma(j) \quad \bar{\eta}_Y = \eta_c \bar{\Gamma} \quad (\text{A.107})$$

Going back to the original problem, our final form for the cross-task elasticity of wages is

$$\eta_c(j, k) = -\eta_c \Gamma(k) \Gamma(j) h(j) \left[\eta_a(j) \kappa + \pi(j) \left(\frac{\bar{s}_l}{s_l(j)} (\chi - \psi) \Gamma(j)^{-1} + (\nu - \psi) \kappa \right) \right]. \quad (\text{A.108})$$

A.6 Occupational wage changes: general case

Next, we will make use of the elasticities which obtain from our loglinearized solutions for wages in order to approximate the changes in wages and employment and wages at the firm-occupation level. While we find that the loglinear approximation is quite accurate for *task-level* wages $w(j)$, occupational wages $W(o, f)$ and employment $N(o, f)$ are a nonlinear functions of $\{w(j)\}_j$. To get a more accurate solution given these nonlinearities, we plug these log-linearized wage changes into a second order approximation of $W(o, f)$, a process we refer to as a “second order correction”.

Hence, in this section we will approximate the log change in occupation level wages as

$$\Delta \log W(o, f) \approx \Delta_1 \log W(o, f) + \Delta_2 \log W(o, f) \quad (\text{A.109})$$

where $\Delta_1 \log W(o, f)$ represents the first order approximation term, and $\Delta_2 \log W(o, f)$ captures the additional second order correction. Next, we characterize each of these terms for the general case, as well as important special cases.

Change in Wages: First Order

Proposition 3 *For a series of shocks ε to task capital costs, the first order approximation of the log change in wages equals*

$$\Delta_1 \log W(o, f) = \sum_{j \in J_o} \left((\eta_a(j) + (\nu - \psi)\pi(j)) \frac{\partial \log w(j)}{\partial \log \bar{Y}} - \frac{\bar{s}_l}{s_l(j)} (\chi - \psi)\pi(j) \frac{\partial \log W(o, f)}{\partial \log \bar{Y}} \right) h(j) \varepsilon(j) \quad (\text{A.110})$$

$$\text{where } \pi(j) = \frac{\eta_a s_l(j) + (\eta_b + \eta_q(1 - \nu))s_k(j)}{1 - \nu}.$$

Proof. To derive the expression for wage growth, we start with the following expression for the log change in occupation wages.

$$\log \frac{W_1(o, f)}{W_0(o, f)} = \beta \log \left(\frac{\sum_{j \in J} \alpha(j) w_1(j)^{\frac{1}{\beta}}}{\sum_{j \in J} \alpha(j) w_0(j)^{\frac{1}{\beta}}} \right) \quad (\text{A.111})$$

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

$$w_1(j) = w_0(j) e^{\eta_o(j) \varepsilon(j) + \sum_k \eta_c(k, j) \varepsilon(k)} \quad (\text{A.112})$$

Hence, we can rewrite the expression for wage growth as

$$\log \frac{W_1(o, f)}{W_0(o, f)} = \beta \log \left(\sum_{j \in J} h(j) e^{\frac{1}{\beta} \left[\eta_o(j) \varepsilon(j) + \sum_k \eta_c(k, j) \varepsilon(k) \right]} \right) \quad (\text{A.113})$$

To get a first-order approximation for the log difference, first take the partial derivatives with respect

to each shock.

$$\begin{aligned}
\frac{\partial}{\partial \varepsilon(j)} \log \frac{W_1(o, f)}{W_0(o, f)} \Big|_{\varepsilon=0} &= h(j)\eta_o(j) + \sum_{k \in J_o} h(k)\eta_c(j, k) \\
&= h(j)\Gamma(j) \left[\eta_a(j)(1 - \eta_c \bar{\Gamma} \kappa) + \pi(j) \left((\nu - \psi) - \bar{\Gamma} \eta_c (\kappa(\nu - \psi) + \frac{\bar{s}_l}{s_l(j)} (\chi - \psi) \Gamma(j)^{-1}) \right) \right] \\
&= h(j)\eta_c \left[\left(\eta_a(j) + (\nu - \psi)\pi(j) \right) \Gamma(j) - \frac{\bar{s}_l}{s_l(j)} (\chi - \psi)\pi(j)\bar{\Gamma} \right]
\end{aligned} \tag{A.114}$$

Plugging this expression into our first order approximation, we get the following:

$$\begin{aligned}
\Delta_1 \log W(o, f) &= \sum_{j \in J_o} \frac{\partial}{\partial \varepsilon(j)} \log \frac{W_1(o, f)}{W_0(o, f)} \Big|_{\varepsilon=0} \varepsilon(j) \\
&= \sum_{j \in J_o} \left[\left(\eta_a(j) + (\nu - \psi)\pi(j) \right) \eta_c \Gamma(j) - \frac{\bar{s}_l}{s_l(j)} (\chi - \psi)\pi(j)\eta_c \bar{\Gamma} \right] h(j)\varepsilon(j) \\
&= \sum_{j \in J_o} \left(\left(\eta_a(j) + (\nu - \psi)\pi(j) \right) \frac{\partial \log w(j)}{\partial \log \bar{Y}} - \frac{\bar{s}_l}{s_l(j)} (\chi - \psi)\pi(j) \frac{\partial \log W(o, f)}{\partial \log \bar{Y}} \right) h(j)\varepsilon(j).
\end{aligned} \tag{A.115}$$

■

The general expression incorporates a number of different pieces, reflecting the fact that capital shares impact the appropriate weights on different components of the expression and we allowed technology to impact capital prices (to capture task-specific capital deepening) as well as the CES parameters a_j and b_j (to capture task substitution). However, the first order term simplifies dramatically if we focus on the symmetric case.

If technology only influences $q(j)$ – the baseline case considered in the paper – the symmetric solution follows immediately as a special case of our general result above.

Corollary 1 *Suppose we are in the symmetric scenario where all tasks have the same initial capital share, wage, and capital cost. Tasks differ by their CES weight $\alpha(j)$. Suppose that AI improvements only lower the cost of capital, so that $\eta_a = \eta_b = 0$ and $\eta_q = -1$. Then*

$$\Delta_1 \log W(o, f) = \eta_m m(\varepsilon) \tag{A.116}$$

where

$$\eta_m \equiv -\frac{s_k(\nu - \chi)}{\zeta + \nu s_k + \chi(1 - s_k)}. \tag{A.117}$$

and

$$m(\varepsilon) = \sum_{j \in J_o} \frac{\alpha(j)}{\sum_{k \in J_o} \alpha(k)} \varepsilon(j) \tag{A.118}$$

Note that this is the expression and setup that we use in the paper.

Proof. To set this problem up, we first need to plug in the task cost elasticities into $\pi(j)$. Thus we

get the new form:

$$\pi = -s_k \quad (\text{A.119})$$

In the scenario where we have symmetric capital shares, notice also that $\eta_Y(j) = \bar{\eta}_Y$, $s_l(j) = \bar{s}_l$, and $h(j) = \frac{\alpha(j)}{\sum \alpha(k)}$. Thus collapsing down the sum and simplifying gives the above expression. ■

Next, we consider the pure task substitution scenario of [Acemoglu and Restrepo \(2018\)](#) in the symmetric case. The key result is that the same sufficient statistic appears as in the baseline model, except the mapping between the coefficient η_m and the model parameters is different.

Corollary 2 *Suppose we are in the symmetric scenario where each task has the same initial capital share, wage, and capital cost. Tasks differ by their CES weight $\alpha(j)$ and the parameters a_j and b_j are defined as in section A.4. Suppose that $\eta_q = 0$ and AI shifts the automation threshold from I to $I \exp(\epsilon)$, with $\phi(I) < \frac{Mw}{q}$. Then,*

$$\eta_a = -\frac{I\phi(I)^{\nu-1}}{\int_I^1 \phi(i)^{\nu-1} di} \quad \eta_b = 1 \quad (\text{A.120})$$

The first order approximation term of the change in wages simplifies to the following:

$$\Delta_1 \log W(o, f) = \eta_m m(\epsilon) \quad (\text{A.121})$$

where

$$\eta_m \equiv \frac{\eta_a + (\nu - \chi)\pi}{\zeta + \nu s_k + \chi(1 - s_k)}, \quad (\text{A.122})$$

$\pi < 0$, and

$$m(\epsilon) = \sum_{j \in J_o} \frac{\alpha(j)}{\sum_{k \in J_o} \alpha(k)} \epsilon(j) \quad (\text{A.123})$$

Proof. From section A.4, we have the following expressions for the a_j and b_j .

$$a = \int_I^1 \phi(i)^{\nu-1} di \quad b = I \quad (\text{A.124})$$

Their respective elasticities with a respect to a shock in I look as follows

$$\eta_a = -\frac{I\phi(I)^{\nu-1}}{\int_I^1 \phi(i)^{\nu-1} di} \quad \eta_b = 1. \quad (\text{A.125})$$

From proposition 3, we know that in the symmetric case the mean term of log wage change looks like

$$\eta_m \equiv \frac{\eta_a + (\nu - \chi)\pi}{\zeta + \nu s_k + \chi(1 - s_k)} \quad (\text{A.126})$$

We can use our earlier definition for a and b , plug them into the above as well as π , s_k , and s_l to get

the following derivation:

$$\begin{aligned}
\pi &= (1 - \nu)^{-1} \frac{\eta_a \left(\int_I^1 \phi(i)^{\nu-1} di \right) (Mw)^{1-\nu} + I q^{1-\nu}}{\left(\int_I^1 \phi(i)^{\nu-1} di \right) (Mw)^{1-\nu} + I q^{1-\nu}} \\
&= (1 - \nu)^{-1} I p(j)^{\nu-1} \left(- \left(\frac{Mw}{\phi(I)} \right)^{1-\nu} + q^{1-\nu} \right) \\
&= I p(j)^{\nu-1} q^{1-\nu} \underbrace{\left(1 - \left(\frac{Mw}{q\phi(I)} \right)^{1-\nu} \right)}_{\geq 1} (1 - \nu)^{-1} < 0
\end{aligned} \tag{A.127}$$

In the above expression, the first three terms are all trivially positive. From the assumption made in equation A.66, we see that $\frac{Mw}{q\phi(I)}$ is greater than one. Thus for any value of ν , π is negative. Plugging this back into our expression for η_m , we see that both terms in the parenthesis are negative, and the term out front is strictly positive. Thus treating AI as an automation shock has a strictly negative mean effect on wages. ■

Change in Occupation Wages: Second Order Corrections

Next, we derive the second order correction and concentration terms for firm-occupation level wages.

Proposition 4 *For a series of shocks ε to tasks in occupation o , the second order approximation term for the the log change in wages is equal to*

$$\Delta_2 \log W(o, f) = \frac{1}{2\beta} C(\Omega(j)) \tag{A.128}$$

where

$$\Omega(j) = -(\Gamma(j) - \bar{\Gamma}) \left(\sum_{k \in J_o} \left[\eta_a(k) \kappa + \pi(k) \left(\frac{\bar{s}_l}{s_l(k)} (\chi - \psi) \Gamma(k)^{-1} + (\nu - \psi) \kappa \right) \right] \eta_c h(k) \Gamma(k) \varepsilon(k) \right) + \eta_o(j) \varepsilon(j) \tag{A.129}$$

and

$$C(x) \equiv \sum_{j \in J} h(j) (x(j) - m(x))^2 \tag{A.130}$$

Proof. We start with the same change in wealth equation derived in proposition 3.

$$\log \frac{W_1(o, f)}{W_0(o, f)} = \beta \log \left(\sum_{j \in J} h(j) e^{\frac{1}{\beta} \left[\eta_o(j) \varepsilon(j) + \sum_k \eta_c(k, j) \varepsilon(k) \right]} \right) \tag{A.131}$$

To get a second-order approximation for the log difference, take the second order partial derivatives with respect to each shock.

$$\begin{aligned}
\frac{\partial^2}{\partial \varepsilon(j) \partial \varepsilon(k)} \log \frac{W_1(o, f)}{W_0(o, f)} &= \frac{1}{\beta} \left[\sum_{i \in J_o} h(i) \eta_c(j, i) \eta_c(k, i) - \left(\sum_{i \in J_o} h(i) \eta_c(j, i) \right) \left(\sum_{i \in J_o} h(i) \eta_c(k, i) \right) \right. \\
&\quad + h(j) \eta_o(j) \left(\eta_c(k, j) - \sum_{i \in J_o} h(i) \eta_c(k, i) \right) + h(k) \eta_o(k) \left(\eta_c(j, k) - \sum_{i \in J_o} h(i) \eta_c(j, i) \right) \\
&\quad \left. - h(j) h(k) \eta_o(j) \eta_o(k) \right] \\
\frac{\partial^2}{\partial \varepsilon(j)^2} \log \frac{W_1(o, f)}{W_0(o, f)} &= \frac{1}{\beta} \left[\sum_{i \in J_o} h(i) \eta_c(j, i)^2 - \left(\sum_{i \in J_o} h(i) \eta_c(j, i) \right)^2 \right. \\
&\quad + 2h(j) \eta_o(j) \left(\eta_c(j, j) - \sum_{i \in J_o} h(i) \eta_c(j, i) \right) \\
&\quad \left. + h(j) \eta_o(j)^2 - h(j)^2 \eta_o(j)^2 \right]
\end{aligned} \tag{A.132}$$

For ease of simplification, suppose we condense the formula for the cross-task elasticity as follows.

$$\eta_c(j, k) = -\eta_c \Gamma(k) \Gamma(j) h(j) \left[\eta_a(j) \kappa + \pi(j) \left(\frac{\bar{s}_l}{s_l(j)} (\chi - \psi) \Gamma(j)^{-1} + (\nu - \psi) \kappa \right) \right] \equiv \tilde{\eta}_c(j) \Gamma(j) \Gamma(k) \tag{A.133}$$

Plugging this into our second order approximation, we get the following.

$$\begin{aligned}
\frac{1}{2} \sum_{j \in J_o} \sum_{k \in J_o} \frac{\partial^2}{\partial \varepsilon(j) \partial \varepsilon(k)} \log \frac{W_1(o, f)}{W_0(o, f)} \Big|_{\varepsilon=0} \varepsilon(j) \varepsilon(k) &= \frac{1}{2\beta} \sum_{j \in J_o} h(j) \left((\Gamma(j) - \bar{\Gamma}) \left(\sum_{k \in J_o} \tilde{\eta}_c(k) \Gamma(k) \varepsilon(k) \right) + \eta_o(j) \varepsilon(j) \right)^2 \\
&\quad - \left(\sum_{j \in J_o} h(j) \eta_o(j) \varepsilon(j) \right)^2
\end{aligned} \tag{A.134}$$

To further break down the above approximation, start by defining the following two quantities

$$m(x) = \sum_{j \in J} h(j) x(j) \tag{A.135}$$

$$C(x) \equiv \sum_{j \in J} h(j) (x(j) - m(x))^2 = \sum_{j \in J} h(j) x(j)^2 - m(x)^2 \tag{A.136}$$

The above two quantities are the mean and concentration of the variable x weighted by hours. Next, let

$$\begin{aligned}\Omega(j) &= (\Gamma(j) - \bar{\Gamma}) \left(\sum_{k \in J_o} \tilde{\eta}_c(k) \Gamma(k) \varepsilon(k) \right) + \eta_o(j) \varepsilon(j) \\ &= -(\Gamma(j) - \bar{\Gamma}) \left(\sum_{k \in J_o} \left[\eta_a(k) \kappa + \pi(k) \left(\frac{\bar{s}_l}{s_l(k)} (\chi - \psi) \Gamma(k)^{-1} + (\nu - \psi) \kappa \right) \right] \eta_c h(k) \Gamma(k) \varepsilon(k) \right) + \eta_o(j) \varepsilon(j)\end{aligned}\tag{A.137}$$

Notice that

$$m(\Omega(j)) = \sum_{j \in J_o} h(j) \left[(\Gamma(j) - \bar{\Gamma}) \left(\sum_{k \in J_o} \tilde{\eta}_c(k) \Gamma(k) \varepsilon(k) \right) + \eta_o(j) \varepsilon(j) \right] = \sum_{j \in J_o} h(j) \eta_o(j) \varepsilon(j)\tag{A.138}$$

Thus we can rewrite as:

$$\begin{aligned}\frac{1}{2} \sum_{j \in J_o} \sum_{k \in J_o} \frac{\partial^2}{\partial \varepsilon(j) \partial \varepsilon(k)} \log \frac{W_1(o, f)}{W_0(o, f)} \Big|_{\varepsilon=0} &= \frac{1}{2\beta} \sum_{j \in J_o} h(j) \Omega(j)^2 - \frac{1}{2\beta} m(\Omega(j))^2 \\ &= \frac{1}{2\beta} C(\Omega(j))\end{aligned}\tag{A.139}$$

■

Corollary 3 *Suppose we are in the symmetric scenario where each task has the same initial capital share, wage, and capital cost. Tasks differ by their CES weight $\alpha(j)$. The second order approximation term of the change in wages simplifies to the following:*

$$\Delta_2 \log W(o, f) = \frac{1}{2\beta} \eta_o^2 C(\varepsilon).\tag{A.140}$$

where

$$\eta_o = \frac{\eta_a + (\nu - \psi) \pi}{\frac{1}{\beta} - 1 + \nu s_k + \psi s_l}\tag{A.141}$$

and

$$C(x) \equiv \sum_{j \in J} h(j) (x(j) - m(x))^2\tag{A.142}$$

Proof. Notice that when capital shares are symmetric, the first term of $\Omega(j)$ falls out entirely leaving us with just the $\eta_o(j)$, which is now a constant and can be pulled out of the concentration function. We know from equation [A.104](#) that

$$\eta_o = \left(\eta_a + (\nu - \psi) \pi \right) \Gamma.\tag{A.143}$$

The second order term is positive regardless of the sign of η_o . ■

B Data Appendix

Here, we elaborate on the details of our empirical analysis.

B.1 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).¹⁷ 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.

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.

¹⁷Upon 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.

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 outputting 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.

B.2 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 5, 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 $\log(1 + \text{Newly Added AI Resumes}_{f,t})$ against $\log(1 + \text{AI-Related Job Postings}_{f,t})$. 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 by 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 $\log(1 + N_{f,t})$, 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 $\log(1 + \text{AI Workers (BFHH)}_{f,t})$, the log of 1 plus the number of AI workers in Babina et al. (2024) against Revelio-implied AI workers $\log(1 + \text{AI Workers}_{f,t})$. 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, bottom axis) and the average mean concentration in AI exposure (orange bar, top axis) 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 (33) in Section 3.5, 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 8 and section 3.3 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 3.5 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 (??) 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 3.5 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 (??) 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: Relevance tests for university network IV**Panel A: Lagged university \times firm shares predict future university \times firm shares**

	(1)
	Average Share (2014-2018)
Average Share (2005-2009)	0.480*** (34.41)
N	861524
R-sq (within)	0.117
Firm FE	X
University FE	X

Variation: university \times firm**Panel B: Shift-share for firm AI worker share predicts actual AI worker share**

	(1)
	Actual AI Worker Share
Predicted AI Worker Share	0.537*** (7.46)
N	16560
R-sq (within)	0.0433
Revelio Emp Control	X
Ind \times Year FE	X

Variation: firm \times 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-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 \times year fixed effects. T-stats from standard errors clustered by firm are in parentheses.

Table A.4: IV robustness tests

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

	IV (Drop Top 50 AI Firms/Universities+Tech Industry)			
	(1)	(2)	(3)	(4)
	Sales	Emp	Profit	TFP
log(1 + AI applications)	8.21*	7.41**	8.16*	4.69*
	(2.54)	(3.06)	(2.46)	(2.52)
N	9,458	9,847	8,507	4,256
R-sq	0.084	0.050	0.034	0.17
Controls	X	X	X	X
Ind × 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)	(2)	(3)	(4)
	Sales	Emp	Profit	TFP
log(1 + AI applications)	9.57***	6.64***	8.29**	7.75***
	(3.83)	(3.64)	(3.22)	(5.25)
N	12,282	12,688	11,246	6,035
R-sq	0.070	0.051	0.027	0.18
Controls	X	X	X	X
Shift-Share Controls	X	X	X	X
Ind × 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 × 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.5: AI exposure and occupational employment growth (IV Robustness)

Dependent variable: $100 \times$ 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	-18.8*** (-7.64)	-16.5*** (-9.04)	-16.8*** (-10.80)	-9.25*** (-6.03)	-14.9*** (-5.18)	-12.9*** (-5.13)	-14.9*** (-8.21)	-10.0*** (-9.27)
AI Exposure Concentration	11.3*** (6.64)	6.96*** (4.48)	7.39*** (5.65)	4.42* (2.33)	12.4*** (4.63)	11.5*** (4.42)	13.7*** (6.95)	9.23*** (6.32)
$\log(1 + \text{AI uses})$	11.0*** (5.33)	7.62*** (4.69)			19.8*** (13.33)	18.9*** (12.06)		
N	1,084,376	1,084,376	1,084,302	1,084,302	1,452,305	1,452,305	1,452,211	1,452,211
R ²	-0.0080	0.014	-0.0025	-0.0021	0.023	0.013	-0.032	-0.012
F-stat (AI Exposure Average)	1016.5	1,234.8	2,274.5	1,150.6	405.1	393.7	958.2	1,121.2
F-stat (AI Exposure Concentration)	741.9	861.5	2107.6	543.4	168.7	126.0	294.2	382.9
F-stat ($\log(1 + \text{AI uses})$)	638.5	793.8			1,386.5	1,929.6		
Controls	X	X	X	X	X	X	X	X
Year FE	X				X			
Industry \times Year FE		X				X		
Firm \times Year FE			X	X			X	X
Occ \times 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 (13) 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 \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 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 $\sigma_{o,t}^2$ (respectively defined in equations (??) and (??) 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.6: 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	-7.33	4.21	2.77	-0.35	5.37
Business and Financial	13	-15.7	8.95	3.80	-2.98	5.69
Architecture and Engineering	17	-11.7	5.27	4.66	-1.78	1.87
Science	19	-3.67	2.64	2.58	1.55	0.85
Community and Social Service	21	6.40	-2.85	1.77	5.32	1.49
Legal	23	5.86	-3.53	4.83	7.16	0.82
Education and Library	25	4.94	-2.12	2.48	5.31	6.08
Arts, Entertainment, Media	27	3.27	-2.19	5.32	6.40	1.31
Healthcare Practitioners	29	0.18	0.68	1.43	2.29	5.93
Healthcare Support	31	2.42	-0.80	1.02	2.64	2.91
Protective Service	33	4.45	-2.85	0.60	2.20	2.43
Food Preparation and Serving	35	7.64	-3.56	-9.60	-5.51	9.29
Cleaning and Maintenance	37	9.71	-5.64	-3.13	0.93	3.21
Personal Care and Service	39	7.57	-3.61	-3.18	0.79	3.16
Sales and Related	41	-3.44	1.80	-0.73	-2.37	11.1
Office and Administrative	43	-2.25	0.34	2.28	0.36	16.3
Farming, Fishing, and Forestry	45	8.56	-4.71	-4.22	-0.36	0.35
Construction and Extraction	47	1.99	-1.71	1.91	2.19	4.13
Installation and Repair	49	-0.81	-0.63	1.38	-0.069	4.03
Production	51	0.77	0.18	-0.41	0.54	6.62
Transportation	53	3.14	-1.48	-2.58	-0.92	7.05

Note: This table shows results from estimating the decomposition (33) for broad 2-digit SOC occupation groups, except occupations within a 2-digit SOC occupation category are re-weighted so that the 2-digit groups reflect their yearly employment shares in the BLS-OES data. See notes to Table 6 and Section 3.3 of the text for further details.

Table A.7: Changes in Demand for 3 Most Exposed Tasks for Selected Occupations

Occupation	Task	Change
Industrial Production Managers	• Develop or implement production tracking or quality control systems, analyzing production, quality control, maintenance, or other operational reports to detect production problems.	-26.6
	• Prepare reports on operations and system productivity or efficiency.	-24.8
	• Collect and analyze production samples to evaluate quality.	-5.9
Claims Adjusters, Appraisers, Examiners, and Investigators	• 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.	-32.1
	• Analyze information gathered by investigation and report findings and recommendations.	-22.7
	• Prepare reports to be submitted to company’s data processing department.	-25.2
Human Resources Workers	• Analyze employment-related data and prepare required reports.	-19.2
	• Evaluate selection or testing techniques by conducting research or follow-up activities and conferring with management or supervisory personnel.	-25.7
	• Evaluate recruitment or selection criteria to ensure conformance to professional, statistical, or testing standards, recommending revisions, as needed.	-10.3
Accountants and Auditors	• Examine and evaluate financial and information systems, recommending controls to ensure system reliability and data integrity.	-17.0
	• Collect and analyze data to detect deficient controls, duplicated effort, extravagance, fraud, or non-compliance with laws, regulations, and management policies.	-23.7
	• Develop, implement, modify, and document recordkeeping and accounting systems, making use of current computer technology.	-23.2
Financial and Investment Analysts	• Analyze financial or operational performance of companies facing financial difficulties to identify or recommend remedies.	-21.0
	• Inform investment decisions by analyzing financial information to forecast business, industry, or economic conditions.	-13.6
	• Evaluate capital needs of clients and assess market conditions to inform structuring of financial packages.	-14.9
Insurance Claims and Policy Processing Clerks	• Review and verify data, such as age, name, address, and principal sum and value of property, on insurance applications and policies.	-21.3
	• Organize or work with detailed office or warehouse records, using computers to enter, access, search or retrieve data.	-47.4
	• Enter insurance- and claims-related information into database systems.	-5.5

Note: Table lists examples of the most AI-exposed tasks for selected occupations along with the averages in the subsequent changes in demand for skills related to these tasks. For each of the selected occupations, we isolate the three most AI-exposed tasks (based on the task AI exposure probability $\xi_{o,f,t}$), then report the average of 100 times the 5-year [Davis et al. \(1996\)](#) change in the share of the skills in firms’ job postings for that occupation that are related to that task (see main text and notes to [Table 3](#) for more details on the variable’s construction). Specifically, we average the [Davis et al. \(1996\)](#) change across all firm–years for each task to get the average outcome at the task level. We then subtract off the occupation-level mean across tasks (weighted by O*NET task-importance), so that the total change sums to zero across all tasks within the occupation.