Canaries in the Gold Mine: Early Productivity Gains from Artificial Intelligence Creating Organization Capital

Tania Babina[a], Alex He[b], Renhao Jiang[c]

[a]*University of Maryland, College Park, NBER, and CEPR*
[b]*University of Maryland, College Park*
[c]*University of Maryland, College Park*

Abstract

Using a new measure of AI investments based on AI-skilled hiring—capturing early machine learning through recent generative and agentic AI jobs—we show that firms' AI investments are associated with emerging productivity growth, in contrast to no measurable effects during the previous decade. These productivity gains stem from AI-driven accumulation of organization capital: firm-specific knowledge acquired through learning and enabling firms to operate more efficiently. Using novel measure of organization capital based on workers' job descriptions, we document that the productivity gains are driven by AI-skilled jobs that build organization capital. Overall, productivity gains from AI investments come from a new form of AI-generated organization capital.

Keywords: Artificial Intelligence, productivity, AI investments, employment, intangible capital, organization capital, workforce composition
JEL: D22, E22, J23, J24, L11, O33

⋆First draft: February 26, 2026. This draft: August 17, 2026. Emails: tbabina@umd.edu (Tania Babina), axhe@umd.edu (Alex He), rjiang07@umd.edu (Renhao Jiang). We are grateful for discussions and helpful comments from Katharine Abraham, Daron Acemoglu, David Autor, Martin Beraja, Lee Branstetter, Nicolas Bloom, Timothy Bresnahan, Erik Brynjolfsson, David Byrne, Nicolas Crouzet, David Deming, Joshua Gans, Danial Gross, Avi Goldfarb, Kyle Herkenhoff, Anders Humlum, Xavier Jaravel, Charles Jones, Benjamin Jones, Sophia Kazinnik, Kai Li, Joseba Martinez, Daniel Rock, Rafaella Sadun, Gregor Schubert, Lawrence Schmidt, Bryan Seegmiller, Chad Syverson, Neil Thompson, John Van Reenan, Ben Zhang, and seminar participants at the Baruch AI in Finance Workshop, Drexel University, the OECD Virtual Conference on the Competition in the AI landscape, the UBS New Frontiers in Quant Finance Conference, the University of Maryland, the University of North Carolina Alumni Conference, the University of Sydney, the University of Technology Sydney, the UBC Summer Finance Conference, and the Workshop on Emerging Research on AI, Labor, and the Future of Work at the Economic Innovation Group. We thank Bryan Seegmiller for his suggestions and Yueran Ma for her code on matching firms in Revelio and Compustat. We thank Yaru Su for excellent research assistance. All errors are our own.

Technological change is a major source of increased economic growth (Aghion and Howitt, 1992; Kogan et al., 2017). While many major technologies drive growth through increased productivity, the evidence on productivity of the recent technological wave, Artificial Intelligence (AI), is mixed. The micro-oriented studies that focus on quantifying the impact of adoption of AI tools for specific tasks typically find large productivity gains. However, evidence from more aggregated data—across many AI use cases, firms, or sectors—has not found consistent productivity improvements. The debate echoes Solow's famous observation about an earlier technological revolution: *"You can see the computer age everywhere but in the productivity statistics."* In this paper, we revisit this question using new measures of firm investments in AI, spanning the early machine learning wave and the recent wave of generative and agentic AI.

We document that firms investing more in AI experience significantly faster growth in firm-level productivity between 2018 and 2024—a marked departure from the earlier decade when we find no such gains. These gains are not driven by the recent generative AI advances and instead extend across a broad spectrum of AI technologies. A new measure of organization capital—firm-specific knowledge that enables more efficient production—points to this capital as an important mechanism: the productivity gains are concentrated in AI investments that build organization capital, and its slow accumulation helps explain the long lag in AI's productivity benefits. Overall, the key contribution of this paper is to show that firms realize productivity gains from AI investments through building a new form of AI organization capital.

Measuring firm-level AI investment is challenging because firms do not systematically disclose it. We build on and extend the human-capital-based approach of Babina et al. (2024), which proxies for a firm's AI investments with its employment of AI-skilled workers—a core input into firms' internal AI development and implementation. We combine two large datasets: the near-universe of U.S. online job postings from Lightcast and individual employment histories from Revelio Labs resume data covering roughly two-thirds of the U.S. workforce. From the job postings data, we identify AI-relevant skills in a data-driven way, based on their co-occurrence with core AI skills (e.g., artificial intelligence, machine learning): their AI-relatedness score, which ranges from 0 (for skills that never appear in the same job postings as core AI skills) to 1 (skills that always appear with core AI skills). A key advantage of this approach is that it automatically captures the evolving vocabulary of successive AI waves: for example, a recent skill "Agentic Systems" has an AI-relatedness score of 0.98, even though it did not exist when the original methodology was developed. This property allows us to capture firm investments in early wave of commercial AI, such as machine learning, and recent advances in generative and agentic AI.

We search for highly AI-related skills in worker resume data to classify each job spell as either AI-related or not. The resulting AI measure has a low false-positive rate of 2%, and—as intended—it mostly captures workers engaged in firms' internal AI development, deployment and management rather than users of outsourced AI tools.[1] In other words, if a firm adopts AI by purchasing external AI software (e.g., off-the-shelf generative AI tools) rather than hiring internal AI developers, this activity should not be reflected in our measure. Our AI measure is also different from AI labor exposure measures, which capture AI's *potential* to do certain tasks instead of actual AI adoption. Instead, our approach allows us to estimate the net effects on firm outcomes across many internal AI applications within a firm, not just those targeting labor replacement (Babina, 2026). We aggregate the AI measure to the firm-year level and link to the firm-level data of U.S. public firms.

The data reveal a rapid and broad-based increase in AI investments. By 2024, U.S. public firms employ more than 300,000 AI-skilled workers, about 2.1% of their U.S. workforce. The AI share rises in nearly every sector, consistent with AI being a general purpose technology. The composition of AI activity is also shifting: jobs tied to first-wave technologies (machine learning, natural language processing, computer vision) decelerate late in the sample, while jobs tied to second-wave technologies (generative and agentic AI) accelerate sharply since 2022.

Using long-differences specification, we examine how *changes* in firms' productivity are associated with *changes* in their AI investments, while controlling for industry fixed effects and a rich set of firm, industry, and local labor market controls. By taking first differences in both independent and dependent variables, the long-differences specification allows us to address the concern that AI-investing firms are more productive. Two facts emerge. First, using our new AI data we revisit and confirm the findings in Babina et al. (2024) of no significant correlation between firm AI investments and productivity growth during the rise of AI adoption from 2010 to 2018. Second, over 2018–2024 the relationship turns strongly positive: a one-standard-deviation increase in the growth of a firm's AI worker share is associated with a 6.2 percentage point increase in sales per worker and a 7.5 percentage point increase in revenue total factor productivity (TFP) over the six-year period—roughly one percent of additional productivity growth per year. Notably, productivity growth is positively related to first-wave and second-wave AI investments, indicating that the payoff extends across a broad spectrum of AI technologies rather than

[1] Chequer et al. (2026) link Census firm surveys to worker tax records and find that about half of firms using AI intensively conduct in-house AI R&D, and that talent is a key input to adoption, suggesting that successful deployment of AI depends on internal investments and specialized talent rather than solely on access to off-the-shelf AI tools.

being confined to the recent generative AI wave.[2]

The timing of these effects supports a causal interpretation. Firms that increased AI investment over 2018–2024 display no differential productivity growth in any pre-period window going back to 2010, and the estimated effects build monotonically year by year after 2018. Regressions of ex-post productivity growth on ex-ante AI investment confirm that AI hiring precedes productivity gains rather than the reverse.

We conduct several robustness tests on our key productivity results. In brief, our estimates are similar under an alternative TFP measure following Imrohoroğlu and Tüzel (2014) that corrects for the simultaneity of input choices and when the sample is restricted to firms present since 2010. We show that the results are not mechanical consequences of acquisitions, nor driven by a handful of large technology-oriented firms, nor explained by keyword retrofitting in online resumes after the release of ChatGPT. Finally, we address a series of potential confounders—pre-existing productivity trajectories, industry-level shocks, pandemic exposure, and concurrent investments in non-AI IT technologies—none of which accounts for the estimated effects.

To further address remaining concerns about unobserved firm-level shocks, we instrument firms' AI investments with their pre-2010 hiring ties to universities with historically strong AI research—the ties that were formed before the commercial rise of AI—following Babina et al. (2024). The instrument has a strong first stage, and the IV estimates confirm a significant increase in productivity of AI-investing firms between 2018 and 2024.

We next examine how these productivity gains relate to firm growth: whether firms generate the same sales with fewer workers or sell more with the same employee headcount. Unlike in the earlier period when AI-investing firms expanded both sales and employment at the same rate, over 2018–2024 AI investments are associated with sales growth of 7 percentage points but essentially zero employment growth—a pattern that holds both in Compustat employment and in resume-based job counts that exclude AI workers. AI-investing firms, in other words, have shifted from the labor-intensive expansion of the early AI-adoption cycle toward jobless growth in the later period.

We then turn to the mechanism. The economics of general purpose technologies emphasizes that new technologies raise productivity only in combination with complementary investments in organizational change, worker skills, and redesigned processes (Bresnahan and Trajtenberg, 1995; Brynjolfsson et al., 2002b; Bartel et al., 2007). Consistent with this view, we find that AI-investing firms build intangible (e.g., SG&A and R&D) rather than

[2]These results also suggest that the relationship between firm AI investments and productivity growth is not an artifact of workers relabeling their jobs with fashionable AI terminology terms or putting terms like "AI" in their resumes (since these two waves do not include more general AI search terms like "AI" or "artificial intelligence").

tangible capital (PP&E). The intangible buildup is mirrored in the workforce: AI-investing firms raise the employment of workers with advanced degrees and in senior roles (while reducing the shares of workers without college degrees or in junior roles), and see increase in average wages. (All of these workforce measures exclude AI-related workers, who tend to be more educated, senior and higher-earning employees.) Because intangible capital is produced by and embodied in skilled labor, this reallocation toward high-human-capital workers is itself evidence of intangible capital deepening.

Within intangibles as a mechanism for higher productivity from AI investments, we zero in on the growth in organization capital—firm-specific, durable knowledge that allows a firm to convert inputs into output more efficiently (Prescott and Visscher, 1980), accumulated through learning-by-doing (Atkeson and Kehoe, 2005) and built by employees (Eisfeldt and Papanikolaou, 2013a). AI technologies are, at their core, learning machines built by workers (Agrawal et al., 2022). Firms' internal AI investments create organization capital by accelerating the firm's learning from its own operations: models trained on firm-specific data and embedded in firm-specific workflows are themselves organization capital, built and maintained by skilled AI workers collaborating with colleagues who hold deep knowledge of the firm's processes, business problems, and customers.

Appendix D formalizes this account in a model extending Atkeson and Kehoe (2005), and two predictions follow from it. First, adjustment is slow even though the model contains no adjustment costs or implementation lags. Because organization capital is a stock and AI changes only the flow into it, the response builds as the pre-existing stock, all of it accumulated at the old conversion rate, is retired—which for plausible depreciation rates implies a half-life exceeding five years, and accounts for the long lags in our estimates. Second, and most distinctively, the effect should be concentrated in firms that begin with little organization capital relative to their size, since a firm that has already codified much of what it does has less left to convert. This runs against the natural presumption that AI benefits firms already rich in organizational capability.

The evidence lines up with this mechanism along every margin we can measure. First, using the standard capitalized-SG&A stock of organization capital (Eisfeldt and Papanikolaou, 2013a), AI-investing firms significantly increase their organization capital over 2018–2024, but not over 2010–2018. Second, we introduce a new, text-based measure of organization capital: an LLM reads each worker's job title and description and scores whether the job builds firm-specific knowledge, systems, and processes. For example, an employee who trains and maintains large language models of customer shopping behavior to improve the firm's sales processes counts as organization-capital-building, whereas a worker who simply uses a chatbot for day-to-day tasks does not. AI-related jobs are almost twice as likely as non-AI jobs to be organization-capital-building (58% versus

31%). Third, and most tellingly, the productivity gains are concentrated exactly where the mechanism predicts: growth in organization-capital-related AI jobs strongly predicts productivity growth, while growth in other AI jobs does not, and organization-capital-related AI jobs also predict faster growth in the firm's organization capital stock (measured with capitalized-SG&A). Finally, the effects are concentrated exactly where the model predicts. Splitting firms at the median of organization capital scaled by tangible capital or by total assets, and holding size fixed, the response is two to three times larger among firms that began with little organization capital relative to their size—for both organization capital growth and productivity growth—and is statistically indistinguishable from zero among firms that began organization-capital-rich.

We also examine, and do not find support for, two leading alternative explanations. The first is task automation and labor replacement: if AI raised measured productivity mainly by shedding workers, we would expect employment declines overall and specifically among AI-exposed workers, who are typically high-skilled and highly educated. We find the opposite: employment of highly educated workers rises, and there is no decline in employment in AI-exposed occupations at AI-investing firms. Our measure, which captures internal AI development and implementation rather than exposure to external AI tools, thus isolates a complementary margin to the displacement effects documented elsewhere. The second alternative is market power: if AI-investing firms simply raised prices, revenue-based productivity would rise without any real output gains. But the productivity gains survive aggregation to the industry level, which is inconsistent with pure business stealing, and industries with greater AI investment show no increase in producer prices. The evidence therefore points to an outward shift in the productive frontier rather than a redistribution of rents.

Overall, we find that recent firm investments in internal AI capabilities are associated with large increases in firm-level productivity on the order of 1% per year using a large sample of firms across all sectors of the US economy and AI use cases. These increases in productivity show up in the data a decade after the start of commercial interest in AI technologies. We provide new evidence linking those delayed gains to the slow buildout of organization capital: durable, firm-specific knowledge embodied in systems, routines, and trained AI models. A model of this process shows that the delay requires no adjustment costs to explain—it follows from organization capital being a stock that AI changes only at the margin. Firm investments in AI that build organization capital drive productivity increases while non-organization-capital AI efforts are not associated with productivity growth. The firms at the frontier of AI investments thus serve as canaries—not in a coal mine, warning of danger, but in a gold mine, offering an early signal of the gains that broader AI diffusion may bring.

One key contribution of this paper is to introduce a new firm-level measure of AI investment that captures both earlier (machine learning, computer vision, natural language processing) and more recent AI waves (generative AI, agentic AI).[3] Our measure of internal firm AI investments are complementary to measures of firms' labor exposure to AI (e.g., Felten et al. 2018; Acemoglu et al. 2022; Webb 2020; Eloundou et al. 2023; Eisfeldt et al. 2025b; Jiang et al. 2025) and the use of external AI tools by workers (e.g., Humlum and Vestergaard 2025; Blank et al. 2025; Bick et al. 2026; Bonney et al. 2026; Iscenko et al. 2026). Methodologically, our human-capital-based measure complements AI measures using surveys (e.g., McElheran et al. 2024; Baslandze et al. 2026; Yotzov et al. 2026), patents (e.g., Kalyani et al. 2025; Jiang et al. 2025), and financial disclosures (e.g., Cao et al. 2025; Barrios et al. 2025).[4] Perhaps closest to our AI measures are those of Hampole et al. (2025), who build time-varying task-level exposure to AI applications actually used by the firms in their sample. Complementing this work, we capture firms' internal AI capabilities—irrespective of their AI labor exposure properties—and offer new measures that span successive waves of AI technologies, including the more recent generative and agentic AI.

The paper contributes to the growing literature on the economic impact of AI technologies on firm performance. While there is ample evidence that firm AI efforts improve firm performance, the mechanisms depend on the type of AI activity and application. Babina et al. (2024) examine the earlier period from 2010–2018 and show AI-investing firms grow sales and market valuations through increased product innovation with no change in productivity, in contrast to increased productivity through organization capital from AI investments in this paper. Rock (2019) shows that firm valuations rose after the emergence of AI for firms that employed workers with skills complementary to AI. Jiang et al. (2025) shows that workers more exposed to AI work longer hours and are more productive from using AI tools, and argue that these effects have positive effects on firm performance. Adams et al. (2026) find that firms adopting AI pricing experience faster growth in sales and markups due to improved demand estimation and price setting.[5] Blumenfeld et al. (2026) show that firms with a higher software engineering payroll share have stock returns

[3]We will publicly share our new firm-level AI measures and all relevant materials (e.g., AI search terms) upon publication of this paper.

[4]Babina (2026) describes strengths and weaknesses of these different datasets.

[5]In sectors where the value of data and AI is especially high, such as asset management, AI has large benefits. Bonelli (2026) shows that data-driven venture capital firms become better at identifying startups that resemble prior successes, but less likely to back more novel outliers. Cao et al. (2024) build an AI analyst that beats the majority of human equity analysts in stock price forecasts, though humans retain an advantage where institutional knowledge matters and close the gap once their employers build in-house AI capabilities. Chen et al. (2025) document that AI-oriented hedge funds outperformed earlier in the adoption cycle, but that this differential declined substantially after 2018 as adoption broadened. Sheng et al. (2026) find that hedge funds more reliant on generative AI earn higher abnormal returns in the post-2022 period.

that respond more strongly to AI news, implying market-expected productivity gains concentrated in software engineering. Eisfeldt et al. (2025b) show that firms with more AI-exposed workforces experienced larger positive valuation responses following ChatGPT's release, implying investors expected sizable future gains from reductions in labor expenses. Consistent with that, Hampole et al. (2025) find that AI applications that substitute for labor tasks improve firm performance. We complement this work by examining a distinct channel through which AI drives firm productivity: the creation of organization capital.

More specifically, this paper contributes to the debate on whether recent advances in AI have led to higher productivity.[6] Micro-oriented task-level experiments focusing on generative AI often find immediate time savings or increased task productivity in specific settings, such as customer support and writing (e.g., Brynjolfsson et al. 2025b; Noy and Zhang 2023), but can also generate performance losses outside the model's capability frontier (Dell'Acqua et al., 2023). Relative to these task-based focused studies, this paper's key difference is scope and horizon: it tracks the productivity of a broad sample of firms and use cases over many years rather than short-horizon task performance.

Indeed, the productivity effects of AI investments are harder to assess in more aggregated data than in tightly defined task settings. Evidence at the firm level, largely from surveys of AI use, is mixed. In a 2018 survey of AI use by German firms, Czarnitzki et al. (2023) find that AI adoption is associated with higher levels of (contemporaneously measured) productivity. McElheran et al. (2025) measure firms' AI use in a survey of U.S. manufacturing establishments and find J-curve dynamics: firms that use more AI in 2017 experience lower productivity growth over 2012–2017, but higher productivity growth over 2017–2021. In more recent survey evidence, Aldasoro et al. (2026) find only a modest association between AI use and productivity during 2019–2024 in Germany, while surveys of firm executives in late 2025 and early 2026 find no economically significant productivity effects from AI, both in a multi-country sample of firms (Yotzov et al. 2026) and in the U.S. (Baslandze et al. 2026).[7] Relative to the survey evidence, this paper makes several contributions to the literature on AI and productivity. First, it measures firms' internal AI development and implementation intensity directly via their workforce. Second, these data permit a long panel structure over 2010–2024 (which is not available in surveys), and consistent measurement of AI technologies over time (it is hard to ask what "AI" is consistently when it is emerging and changing): using this long panel, we document an

[6] See a review of this debate in Imas (2025).

[7] The lack of consistent signs of aggregate productivity improvements from AI are consistent with Brynjolfsson et al. (2019) who argue that AI benefits are likely to be delayed by necessary co-investments in complementary innovations, adjustment costs, organizational changes, and new skills. Baily et al. (2025) reach a similar conclusion from the history of general-purpose technologies, arguing that generative AI will likely raise the level of productivity but only slowly.

inflection toward measurable productivity growth after 2018, but not before. Finally, to the best of our knowledge, we are the first to show that firms' internal AI investments help build organization capital, which is a core source of firm-level productivity (Atkeson and Kehoe 2005). This finding helps reconcile why macro productivity statistics may have looked flat early on even as micro studies reported large task-level gains. While AI tools for specific tasks are fast to learn and utilize, organization capital is relatively slow to build and hence its productivity effects are delayed.

Finally, we contribute to the growing literature on intangible capital more generally and organization capital in particular. We make three key contributions to this literature. First, we offer a new way of measuring firms' organization capital using worker job descriptions in their resumes, guided by the view that organization capital is partially embodied in employees and organizational routines. This new approach complements seminal papers on the measurement of intangible and organization capital (e.g., Lev and Radhakrishnan 2005; Eisfeldt and Papanikolaou 2013a; Peters and Taylor 2017) and emerging work that measures intangible capital in new, creative ways. Eisfeldt et al. (2025a) and Charoenwong et al. (2026) build granular measures of firm-level intangible capital, including organization capital, using the text of firms' financial disclosures. He et al. (2024) build measures of investments in customer capital. Cai et al. (2026) measure organization capital from employee reviews of their employers on Glassdoor. Fahlenbrach et al. (2026) also use resume data and measure startup organization capital intensity based on employee positions (e.g., product managers). Second, we show that investment in new technologies like AI is a key contributor to organization capital and drives the positive relationship between organization capital and firm productivity in our setting. This contributes to the literature on the growing economic importance of organization capital (e.g., Eisfeldt and Papanikolaou 2013a), and, more generally, of intangible capital (e.g., Peters and Taylor 2017; Crouzet et al. 2022; Falato et al. (2022); Crouzet and Eberly 2023). Third, we offer a conceptual framework for how a new technology builds organization capital. We extend Atkeson and Kehoe (2005) to distinguish a firm's idiosyncratic productivity from the durable, documented stock it accumulates as a by-product of operations, and model AI as a technology raising the rate at which the first is converted into the second.

1. Data

Following Babina et al. (2024), we proxy for firm-level AI investment using AI-related workers. We construct this measure by using two complementary datasets: job postings from Lightcast and individual worker resumes from Revelio Labs. We link the resulting firm-level AI measures to firm productivity outcomes for U.S. public firms in Compustat, and construct firm-level measures of labor composition and organization capital. This

section describes each key data source.

1.1. Data on Job Postings from Lightcast

We start our analysis with U.S. job postings data from Lightcast (which was created after the merger of Emsi and Burning Glass Technologies—the job postings data used in Babina et al. (2024)). We use the full set of job postings over 2010-2025 to identify AI-specific skills from job demand reflected in job postings. The dataset contains approximately 456 million postings and is constructed from a large set of sources, including more than 64,000 online job boards and company websites. Lightcast aggregates job postings, parses them into a standardized machine-readable format, and applies a deduplication algorithm to mitigate double counting of vacancies posted across multiple boards. This dataset covers a near-universe of job vacancies posted in the U.S., both online and offline, and its representativeness at the occupation level has remained stable over time.

Most importantly for our purposes, the data provide a detailed and standardized list of skills required for each job posting. There are 32,309 unique skills. We use these skill requirements to identify AI-relevant skills, as described in Section 2. We restrict attention to job postings with non-missing employer information (approximately 84% of postings) and at least one standardized required skill (approximately 96% of postings), and exclude internships. This yields a final sample of 367 million job postings.

1.2. Data on Employment Profiles from Revelio

We obtain individual-level employment history data from Revelio Labs. Revelio constructs these data from publicly available online resumes on professional networking websites (e.g., LinkedIn).[8] For each resume, we observe job start and end dates, job titles and descriptions, employer names, employee locations, and educational background, including schools and majors. Although online professional profiles tend to over-represent white-collar workers, the data cover approximately two-thirds of the U.S. workforce as of 2024 and provide a broadly representative breakdown by industry.[9] Revelio uses machine-learning methods to assign seniority levels and occupational classifications to each job spell. The detailed job-spell information yields a job-worker-firm-year panel and allows us to construct consistent firm-year workforce measures.

[8]While Babina et al. (2024) use Cognism as a source of resume data, we use Revelio Labs in this paper. Both datasets draw on online resumes and provide comparable information on workers' employment histories and characteristics, allowing us to construct measures conceptually consistent with those in Babina et al. (2024). The added advantages of Revelio data are its standardized structure and extended time series coverage (Cognism data ends in 2021).

[9]See Revelio Labs, *Revelio Public Labor Statistics (RPLS)*, `https://www.reveliolabs.com/public-labor-statistics/`.

We match employers in Revelio to Compustat firms by leveraging Revelio's time-varying subsidiary structure file as well as a crosswalk from Revelio company identifiers to common firm identifiers (e.g., GVKEY and CUSIP).We implement this linkage to Compustat dynamically to account for acquisitions and divestitures over time. Appendix A.1 describes the linking procedure in detail. Of the 1.07 billion U.S.-based person-firm-year employment records between 2010 and 2024, 211 million (20%) are matched to U.S. public firms (rather than private firms or public-sector employers), covering 30.5 million distinct employees. This match rate is consistent with i) publicly listed U.S. firms accounting for approximately 26% of the overall U.S. employment (Davis et al. 2006), and ii) Revelio data covering approximately two-thirds of the U.S. workforce. For our focal 2018–2024 period, the sample contains 111 million person-firm-years and 23.4 million distinct employees.

1.3. Additional Data Sources

We use several additional data sources. Firm-level outcome and control variables, such as sales, employment, SG&A expenses, come from Compustat. We also collect commuting-zone-level wage and education data from the Census American Community Surveys (ACS), industry-level wages and employment data from the Census Quarterly Workforce Indicators (QWI), and industry-level producer price indices from the Bureau of Labor Statistics (BLS).

2. Measurement and Descriptive Evidence

2.1. Measuring and Validating Firms' Internal AI Investments

This section summarizes the construction of firm-level AI measures and discusses their empirical properties. Appendix A.2 provides a detailed description of their construction, validation, and descriptive properties.

To construct firm-level measures of AI investments, we build upon and extend the methodology developed by Babina et al. (2024). The measure of firms' internal AI investments is based on their intensity of AI-skilled hiring, since AI-skilled labor is a key input to AI development and implementation within a firm. Other inputs to AI, such as data and compute, are complementary to AI-skilled labor, so our human-capital-based measure allows us to capture the relative intensity of AI investments across firms.

Our research question—how firms' productivity changes as they *invest* in internal AI capabilities—requires data that capture internal AI development and implementation across a broad set of firms and use cases, and that reflect actual AI capabilities rather than perceptions of them.[10] Our methodology is reproducible, inexpensive to implement at

[10]Specifically, the data should cover firms that are not just producers of AI but also AI users; capture

scale, and data-driven, allowing it to capture evolving AI job terms across successive waves of AI. It is not intended to capture workers who use AI tools for productivity without working on internal AI development (e.g., chatbot users), and we show below that such workers are rare in our sample.

In a nutshell, we use a data-driven approach to generate highly AI-relevant search terms using job postings data and then search for those terms using information workers provide in their resumes to identify AI-skilled workers. The methodology circumvents the need to pre-specify lists of AI keywords, which may suffer from both Type I (incorrectly labeling tangentially related jobs as AI jobs) and Type II (missing actual AI jobs) errors. The construction proceeds in four steps: (i) we identify AI-relevant skills in job postings via their co-occurrence with core AI skills; (ii) we expand these skills into resume search terms using rule-based methods and LLM; (iii) we screen the candidate terms for false positives; and (iv) we search the validated terms in workers' job titles and descriptions and aggregate to the firm-year level.

Step 1: AI-relevant skills from job postings. We look at the required skills listed in job postings and measure their co-occurrence with four core AI skills: *artificial intelligence*, *machine learning*, *natural language processing*, and *computer vision*. Specifically, for each required skill s, we compute the following measure:

$$w_s^{AI} = \frac{\text{\# of jobs requiring skill } s \text{ and (ML, NLP, CV, or AI in required skills or in job title)}}{\text{\# of jobs requiring skill } s} \tag{1}$$

This measure captures how close each skill is to the core AI technologies. For instance, in our job postings data available through 2025, "Agentic Systems" has a co-occurrence value of 0.98, indicating that 98% of postings requiring it also list at least one core AI skill; "Reinforcement Learning from Human Feedback (RLHF)" scores 0.97 and "Large Language Modeling" 0.93. By contrast, "Microsoft Office" has a value of only 0.003. Appendix Table A.7 reports the 50 skills with the highest AI-relatedness scores and confirms that the co-occurrence methodology of Babina et al. (2024) captures both older AI concepts (e.g., "Doc2vec", "Standard ML Library") and concepts emerging in recent years (e.g., "AgentGPT", "AI-Generated Code"). We then form an initial candidate set of 335 AI-related skills consisting of skills with AI-relatedness scores of at least 0.6 that appear in at least 50 job postings.

Step 2: Expansion into resume search terms. Because resume text is unstructured and

internal development rather than outsourcing; have broad firm and industry coverage; span comprehensive AI use cases; measure actual capabilities rather than perceptions; and reflect all AI efforts, whether successful or not. See Babina (2026) for a discussion of these requirements.

workers may refer to the same skill in multiple ways, we expand each seed skill into the search terms it plausibly takes in resume text (e.g., abbreviations, singular/plural forms, and hyphenated, separated, or merged spellings) and define three matching methods that vary in how strictly they require each candidate keyword to appear in the text.[11] This rule-based expansion yields 433 unique search terms. We then augment this list using an LLM to generate additional closely related terms (including alternative spellings and synonyms, like *TFLite* and *TF-Serving* for *TensorFlow*) and assess their false positive risk. Before adding any term to the keyword list, we run a web search and inspect the top results to judge whether the term is unambiguously AI-related. This LLM-assisted step adds 89 search terms.

Step 3: False-positive screen. We flag terms for which most of the term's matched jobs do not contain core AI terms such as AI, ML, and NLP, since these terms may carry a non-AI meaning in resume text. For each flagged term, we collect all matched jobs and use an LLM to classify whether each position is genuinely AI-related based on its title and description. We then validate the LLM classifications through independent human review of random subsamples.[12] Terms with high false-positive rates in both reviews are removed (e.g., "RL" often refers to "Ralph Lauren" rather than "Reinforcement Learning" in resumes). After this step, the procedure yields the AI skill dictionary of 510 terms.[13]

Step 4: Identifying AI-related jobs and constructing firm-level measures. We classify each worker's job spell as AI-related using the data-driven AI keyword dictionary of terms described above. A job spell is a position at a particular company, with start and end dates, a job title, and a description. Specifically, we label a job as AI-related when its job title or description contains at least one of the 510 AI keywords (see Appendix A.2.5 for details). Because our empirical analyses focus on U.S. public firms, our baseline firm-level AI measure uses U.S.-based employees, whose resume coverage is more complete than international employees and for whom our English-language keyword list is best suited (see Appendix A.1.7 for extensive analyses on data coverage issues). For this sample, we compute the share of AI-related jobs out of all jobs for each firm-year.

We also construct AI measures by AI technology category. Specifically, we use an LLM to assign each AI keyword to one of six categories: general AI, machine learning, natural language processing, computer vision, generative AI, and agentic AI. A job is classified into a category if its title or description contains at least one of that category's

[11] These include: i) case-sensitive word-boundary matches for ambiguous abbreviations, ii) case-insensitive word-boundary matches, and iii) case-insensitive substring matches for longer phrases. See Appendix A.2.2 for details of search term construction and these matching rules.

[12] Across the entire random sample, the human and the LLM classifications agree 91% of the time, suggesting this prompt can be used to separate truly AI-related from false positives.

[13] We can share this list of AI-relevant terms and our code upon request.

keywords. This classification is not mutually exclusive, and a single job can fall into multiple categories. We further define Wave I AI jobs as those matching machine learning, natural language processing, or computer vision terms, and Wave II AI jobs as those matching generative AI or agentic AI terms.[14] Appendix A.2.4 describes the classification and provides examples of jobs tagged by our methodology for each AI technology category. While AI use cases are highly heterogeneous, a common pattern is that a worker uses some AI to learn about some aspect of a firm's business and to improve some aspect of what firm does. For example, a job in "machine learning" category analyzes data and builds a random forest model to identify factors correlated with high customer satisfaction, which informs future resource allocation and brand marketing strategy. A job tagged in "computer vision" leads AI-driven process improvements and applies computer vision models to automate and enhance refund approval workflows, reducing processing time and human error. Job tagged as "agentic AI" integrates LLM-based summarization models for generating investment insights and market summaries.

As a starting point of validating our main AI measures, we find low false-positive rates of AI-related jobs of around 2%. More generally—in the absence of comprehensive data on firms' AI investments—we validate our AI measures by examining the characteristics of jobs that our methodology identifies as AI-relevant.[15] Consistent with AI-skilled talent being highly skilled (Gofman and Jin 2024) and highly-paid (Akcigit et al. 2026), we find that AI-related workers are significantly more educated than non-AI workers (Appendix Figure A.15), are more likely to be in senior roles (Appendix Figure A.16), and have much higher annual wages (Appendix Figure A.17). Appendix Table A.11 and Table A.12 list the AI keywords that identify the most AI jobs, both overall and within each AI subcategory: unsurprisingly, "generative ai" and "agentic ai" pick the most jobs for generative and agentic AI job categories.

Appendix Tables A.14 through A.17 list the job titles and occupations that account for the most AI jobs. They make clear that titles and occupations alone are a poor way to identify AI-skilled talent: "Software Developers" is the second most common occupation containing AI jobs, yet only 5.5% of Software Developers are tagged as AI-related. Employee job descriptions instead offer the richer context needed to identify AI-relevant jobs.

[14]For analyses requiring mutually exclusive groups, we define three groups: jobs that match Wave I but not Wave II, Wave II but not Wave I, and both.

[15]Validating data on internal AI investments against other firm-level AI data is conceptually problematic (Babina 2026): AI patents capture AI *outputs* by AI inventors rather than the *inputs* and downstream use we study; disclosure-based measures reflect strategic release decisions and investor demand for information; and surveys depend on design, response rates, and respondent qualifications, and rarely have a historical panel structure.

We next examine what kinds of jobs employees hold in our AI-related sample by classifying AI jobs into functional categories (see Appendix A.3 for details). From Appendix Tables A.25 and A.26, about 75% of AI jobs are involved in firms' internal AI development and implementation or in the management, coordination, and support of these efforts: roughly half develop and deploy AI internally, 23% manage, oversee, or govern these efforts, and 7.4% sell and support AI products or advise clients on AI strategy. The remaining 25% contain legitimate AI-related keywords and serve some AI-related function but are not focused exclusively on AI implementation: 17.2% are general digital or technical roles with AI tooling, and 7.5% are incidental or low-confidence mentions. Overall, our resume-based measure primarily captures jobs that contribute to internal AI development and deployment as well as AI product management and monetization, rather than workers merely using AI tools: only 0.9% of AI-related jobs are classified as *"Non-AI role using AI tools as productivity aid"*.

Appendix Figure B5 provides a complementary view using Revelio's seven broad job functions. Most AI jobs are in the "Engineer" and "Scientist" functions, but the share outside these two rose from 25% in 2010 to 33% in 2024, with most of the increase after 2018 (Appendix Figure B6), consistent with firms' AI activity maturing from research toward commercialization and deployment across business units.

Figures A.19 and A.20, together with Appendix Tables A.18–A.23, report the firms with the highest AI shares and AI job counts within each industry; the heaviest AI employers are well-known AI investors such as NVIDIA, c3.ai, Amazon, and J.P. Morgan Chase. The time-series figures of firm-level AI-jobs show that these efforts are generally increasing, but typically very gradually, consistent with AI-investments being a slow-moving process.

We further validate our new AI measures in two ways. First, we compare our firm-level measure, based on Revelio Labs resumes, to the measure in Babina et al. (2024), based on Cognism resumes, over 2010–2021, when both are available and before public awareness of AI rose following the release of ChatGPT. The two are highly correlated: the correlation between firm-level AI job counts in our regression sample is 70 to 90% depending on the year (Appendix Table A.24 and Appendix A.2.8). Second, we benchmark our measure against firms' advertised demand for AI workers, comparing new AI-related hires in the resume data with the number of the firm's job postings listing the same AI-related skills. The two series track each other closely within firm-years, with a correlation of 0.89 (Appendix Figure A.21 and Appendix A.2.9). These high correlations with historical resume and job postings datasets also address the concern that workers might relabel past positions using AI terminology that became fashionable only later (Bloom et al. 2026a).

2.2. Measures of Organization Capital

A central mechanism we examine is that AI investments raise productivity by helping firms build organization capital (OC)—durable, intangible, firm-specific knowledge acquired through learning-by-doing and helping firms operate more efficiently. We construct two complementary firm-level measures of organization capital.

Our first measure is a standard accounting-based stock of OC, constructed by capitalizing a firm's past selling, general, and administrative expenses (SG&A)—the expense category that captures spending on firm-specific knowledge, training, and business-process development (Eisfeldt and Papanikolaou, 2013a). Specifically, we apply the perpetual inventory method to firms' deflated SG&A expenses, cumulating these flows into a capital stock with a 20% annual depreciation rate (and 15% for robustness). We then normalize organization capital by property, plant, and equipment (PP&E) or total assets.

Our second measure introduces a novel method for measuring organization capital building capabilities in firms based on the text of job descriptions in workers' resumes (see Appendix A.4.3 for details). Guided by the view that OC is partly built by workers and embodied in firms' operating routines (Lev and Radhakrishnan, 2005; Eisfeldt and Papanikolaou, 2013a), we use worker job descriptions to identify jobs that are related to building such capital. Job descriptions detail the tasks workers perform, the routines, systems, and processes they build and maintain, and the firm-specific knowledge they develop and help improve operations, making them a natural source for measuring the organization-capital-building content of jobs. We classify *all* AI jobs and a *random sample* of non-AI jobs according to whether the job contributes to building organization capital. We score each job in the two samples with an LLM that reads the worker's job title and description and returns an organization-capital score between 0 and 1. We define a job as OC-related if its score is at least 0.6, where 0.6 is the boundary between mixed evidence (0.4–0.6) and substantial evidence (0.6–0.85) on building OC. For AI-related jobs, we aggregate this job-level classification to the firm-year level by computing two shares of total jobs—those that are related and unrelated to building organization capital.[16]

We further move beyond *whether* a job builds organization capital to *how* it does so by assigning each job a standardized organization capital category. From the job descriptions data, the LLM first produces a set of categories that build or do not build OC, then assigns each job to its nearest category. Appendix Table A.30 defines each OC category.

[16]One concern is that our OC classification simply picks up workers who write higher-quality job descriptions. Appendix A.4.3 shows this is not the case: organization-capital jobs have longer descriptions, but they use more repetitive vocabulary, cite fewer specific numbers and tools, and contain spelling errors at similar rates, and among descriptions of similar length the two groups are indistinguishable across all writing quality measures.

In the absence of granular data on organization capital, we proceed by validating our new OC measures in several steps. First, in Appendix Table A.32, we dig into examples of AI-related job descriptions that do and do not build OC. Examples of AI OC-building jobs include an employee who trains and maintains large language models for understanding customer shopping behavior and uses them to improve firm sales processes, or an AI product manager who redesigns business processes and workflows using AI methods, driving productivity improvements. In contrast, a worker who simply uses a chatbot for day-to-day tasks or advises clients on their AI strategy is not classified as OC-related. Overall, OC-building AI jobs create systems a firm continues to use, while other AI jobs apply similar technical skills to one-time analyses, general research, or client work.

Second, we examine word clouds for job descriptions of jobs categorized as OC-building in AI and non-AI-related jobs in Appendix Figure A.22. Among AI-related jobs, both—OC and non OC-building jobs—have *data*, *ai*, and *model* as common words, but OC-related positions emphasize *team*, *lead*, *product*, *business*, *design* and *process*, indicating leading teams and building business products and processes, whereas non-OC-related jobs also list *develop*, *machine* and *learning*, emphasizing technical model development. Among non-AI-related jobs, OC-related positions are led by *team*, *process*, *lead*, *management*, and *business*, highlighting operational and managerial routines, whereas non-OC-related jobs are led by *customer*, *service*, *sale*, *provide*, *team*, *product*, *ensure*, emphasizing customer-facing service, sales, and support work.

Third, to illustrate how AI accelerates the accumulation of OC, Appendix A.4.4 presents case studies of eight firms across different industries and the OC-building job categories in Table 2. How AI builds OC is largely common across them: a model is trained on records the firm was already producing—research documents, meeting transcripts, transactions, sensor readings, product and project data—turning firm-specific data into a durable asset that can be queried, simulated against, or acted on.

Fourth, consistent with prior work that argues that workers in management roles are likely to build organization capital (Eisfeldt and Papanikolaou, 2013a; Dessein and Prat, 2022), we find that 79% of AI-related workers in management roles are classified into organization-capital-building roles (Panel A of Appendix Table A.31); same table shows that only 22.6% of AI workers in "Support for AI products/services" build OC, consistent with customer service roles being more about serving customers.

Finally, in Section 4.2, we formally test and show that the increase in AI jobs that build OC is associated with growth in organization capital constructed with capitalized SG&A expenses (Eisfeldt and Papanikolaou 2013a).

2.3. *Measures of Productivity*

We consider two measures of productivity as dependent variables: log sales per worker and a standard measure of revenue TFP. Revenue TFP is the residual from regressing log real sales on log employment and log capital stock, controlling for firm fixed effects and year fixed effects. The regression is estimated using OLS separately for each industry (see Appendix A.5 for details). The TFP measure is specific to Cobb-Douglas production functions, while sales per worker measures labor productivity for more general production functions.

We construct an alternative measure of firm-level TFP using the control function approach pioneered by Olley and Pakes (1996), following the procedure in Imrohoroğlu and Tüzel (2014) (see Appendix A.5 for details). Because the procedure imposes many sample restrictions and excludes industries where physical inputs are poorly defined (e.g., finance), it covers a small subset of firms and is used only in robustness checks.

2.4. *Measures of Firm-Level Labor Composition*

To construct firm-level measures of labor composition, we use worker-job-firm-year data from Revelio Labs. We calculate total employment as well as employee counts and shares by job seniority and highest education level for each firm-year.[17] For wages, we use Revelio's imputed salary for each job to measure firm-by-year average wage, separately for all jobs and for all jobs excluding AI jobs, computed as the total wage bill divided by the number of jobs with a salary.[18]

2.5. *Descriptive Evidence*

AI Measures. Most of our empirical facts below and the key regression sample are based on U.S.-based employees of U.S. public firms (due to productivity data not available for non-public firms). Figure 1 shows the number (in panel (a)) and the share (in panel (b)) of workers who are AI-skilled from 2010 to 2024. Both the level and share are low at the start of the sample in 2010—consistent with limited commercial interest in AI in the early 2010s—before accelerating in the middle of the decade and continuing to rise in recent years. In 2024, the number of AI skilled jobs is greater than 300,000 and the share is about 2.1%.[19] Because AI workers are highly paid, firms' investment in AI talent is larger than

[17] Seniority uses Revelio's classification, binned into low (entry-level and junior), medium (mid-level), and high (senior and executive). Education is the worker's highest reported degree: no college degree, college degree (bachelor's or associate's), master's degree (including MBA), or doctoral degree.

[18] Revelio does not observe actual pay: these salaries are imputed from worker and job characteristics such as job title, years of experience, seniority, employer, and location.

[19] For perspective, medical doctors—one of the most common U.S. occupations—account for 0.5% of jobs (U.S. Bureau of Labor Statistics 2025). While the two are not directly comparable, this benchmark highlights the rapid rise of AI-related employment.

headcounts indicate: the wage-weighted AI share—the fraction of the total wage bill of U.S. public firms paid to AI-related workers—reaches 3.3% in 2024, well above the 2.1% headcount share (Appendix Figure A.18).

We next examine our data by AI technology subcategory: general AI, machine learning, natural language processing, computer vision, generative AI, agentic AI, and other. We first examine the most frequent words in job titles and descriptions for each AI subcategory. Figure 2 shows that, despite category-specific terms, *data* is common across all categories, consistent with its central role in AI applications. Figure 3 presents AI job shares by major subcategory, with the underlying data reported in Appendix Table A.10. In 2024, general AI (1.5% of jobs) and machine learning (0.8%) are the most common subcategories, while the earlier NLP and CV technologies remain rare (0.13% and 0.06%). Generative AI starts near zero prior to 2020 and rises quickly to 0.45% by 2024, and agentic AI begins increasing rapidly in 2023–2024 from very low levels. Figure 4 aggregates these categories into earlier "Wave I" AI jobs (ML, NLP, or CV) and more recent "Wave II" jobs (generative or agentic AI), making the deceleration of the former and the acceleration of the latter more apparent. These technologies are highly interconnected: 58% of agentic AI jobs also contain generative AI keywords, 40% of generative AI jobs mention ML, NLP, or CV keywords, and firms doing more Wave I AI also do more Wave II AI, with an employment-weighted cross-firm correlation of 0.82 in 2024.

The increase in AI jobs displays an intuitive distribution across industries. Figure 5 plots the share of AI-related workers of U.S. public firms in each 2-digit NAICS sector, separately for the years 2010, 2018 and 2024. The figure highlights that the share of AI-skilled jobs is highest in the "Information" sector, growing from 0.07% in 2010 to 5.7% in 2024. However, almost all sectors see large and monotonic increases over time, consistent with the idea that AI is a general purpose technology.

While our core empirical analysis focuses on U.S. public firms, we also rebuild these figures for U.S.-based jobs outside of U.S. public firms. The patterns are similar—acceleration in AI jobs, deceleration in older AI technologies alongside acceleration in newer ones, and increases across sectors—but the levels are markedly lower, with an AI job share of about 1% in 2024, half that of public firms (Appendix Figures B1, B2, and B3). Even among public firms, larger firms invest more: the AI job share rises monotonically across small, medium, and large firm terciles (Appendix Figure B4).

We also examine where the new AI workforce comes from. Classifying workers by their employment and education histories prior to their first AI-related job at a U.S. public firm, we find that roughly 70% throughout 2010–2024 are experienced U.S. workers switching from non-AI roles, while recent U.S. university graduates are the second-largest and fastest-growing group, doubling from 13% in 2010 to 26% in 2024 (Appendix Figure B7). Firms

thus build their AI workforce primarily by drawing on workers already in the U.S. labor market. This complements Baviskar et al. (2026), who focus on the elite segment of AI talent trained by frontier academic scientists at U.S. universities.

There is substantial heterogeneity in the growth in AI-skilled labor across individual firms, which provides the variation needed to examine the relationship between AI investments and firm productivity. Table 1 reports the cross-sectional distribution of AI-job shares in each year. The distribution is dispersed and highly right-skewed. In 2024, the average firm has an AI job share of 1.46%, with a standard deviation of 3.71%—more than twice the mean.[20] The median firm's AI share is only 0.33%, while firms at the 90th and 99th percentiles reach 3.64% and 18.18%, respectively. This heterogeneity widens over our focal period of 2018–2024 as can also be seen in Appendix Figure A.14.

Organization Capital Measures. Next, we provide stylized facts about our novel measure of organization capital constructed from job descriptions. We first examine how jobs that build organization capital differ between AI and non-AI jobs. In the cross section, AI jobs are considerably more OC-intensive: 58% of AI jobs are OC-building versus 31% of non-AI jobs (Appendix Table A.29). We also find that the share of AI workers classified as OC-building is higher among non-tech firms than among tech firms, consistent with tech firms using AI primarily to build products and services for sale, while non-tech firms use AI to build capabilities internally (Appendix Table A.27). These averages, however, mask a pronounced time trend. Figure 7 shows that the share of AI jobs building OC rises from about one half in 2010 to roughly two thirds in 2024 (panel (a)), both within and outside the tech sector (panel (b)), while the corresponding share for non-AI jobs is flat at about one third throughout the time period (panels (c) and (d)). The stability of the non-AI series indicates that the AI trend reflects a genuine change in the nature of AI work rather than a more general trend of jobs becoming more OC-building. The shift also tracks across AI vintages: 57% of earlier Wave I jobs (ML, NLP, CV) build organization capital versus 66% of recent Wave II jobs (generative and agentic AI), with agentic AI having the highest share of 71% across all AI subtechnologies.

.

The granularity of our data allows us to document how jobs build organization capital. Table 2 reports the distribution of AI and non-AI jobs across categories that do and do not build organization capital. AI jobs concentrate in codifying data, models, and proprietary knowledge into reusable technical systems—data platforms (9.1% of all AI jobs), proprietary-knowledge generative-AI/NLP (6.9%), reusable AI/ML platforms (4.8%), and

[20]The firm-level average AI share in 2024 (1.46%) is lower than the employment-weighted share across all jobs in those firms (2.09%, Figure 1) because larger firms have higher AI shares (Appendix Figure B4).

firm-specific predictive and decision models (2.3%). Non-AI organization capital jobs are more operational, led by hiring and talent development (5.8%) and manufacturing and process systems (3.5%). Both rank building cross-functional operating routines among their top categories, highlighting within-firm coordination as a source of organization capital. Beyond these categories, AI workers build durable capabilities across product, marketing, infrastructure, finance, customer support, security, and R&D systems, each accounting for 1–5% of AI jobs.

Among jobs that do not build organization capital, AI jobs concentrate in portable analytics and one-off models (12.6%), while non-AI jobs are dominated by standardized operations and support (37.9%). Appendix Tables A.32 and A.33 provide examples and show that the LLM requires sufficient evidence of durable organizational capabilities rather than treating AI work as organization-capital-building by default.

Figure 8 decomposes AI jobs by organization capital category through which they build firm-specific assets over time. Essentially every category rises between 2010 and 2024. Firms are not scaling one kind of organization-capital work but accumulating many types of firm-specific assets at once—data platforms, reusable AI/ML platforms, product and design systems, and cross-functional operating routines, among others. The composition also tilts toward codifying proprietary knowledge, with proprietary-knowledge generative-AI and NLP systems growing from negligible in 2010 to among the largest categories by 2024. The breadth of this accumulation indicates that AI work builds a wide range of organizational assets.

3. Main Results: Firm AI Investments and Productivity Growth

3.1. Long-Differences Specification

We begin by examining whether firms that invest more in AI experience faster productivity growth. As is standard in settings with slow-moving processes such as technological change (Acemoglu and Restrepo, 2020), our primary specification is a long-differences regression of changes in firm productivity on changes in AI investments, proxied by the share of AI workers. This strategy suits our setting because AI investments accumulate gradually and their effects are not immediate (Babina et al. 2024). By differencing both independent and dependent variables, the specification allows us to address the concern that AI-investing firms are more productive, since we examine how *changes* in firm productivity relate to *changes* in AI investments. For example, for the sample period of 2018–2024, we estimate the following regression:

$$\Delta Y_{i,[2018,2024]} = \alpha + \beta \Delta \textit{Share AI Jobs}_{i,[2018,2024]} + \textit{Controls}'_{i,2018}\gamma + \textit{Industry FE} + \epsilon_i, \quad (2)$$

where the dependent variable, ΔY_i, is the change in a firm productivity over the sample period. The main independent variable, Δ*Share AI Jobs*$_i$, is the change in the share of AI jobs at firm i, standardized to have a mean of zero and a standard deviation of one. *Industry FE* denotes 2-digit NAICS industry sector fixed effects. In some specifications, we add a rich set of baseline controls measured at the start of the sample period: firm characteristics (log sales, cash/assets, R&D/sales, markup computed following De Loecker et al. (2020)), and the log of the firm's total number of workers in Revelio, log industry wage, and characteristics of the commuting zones where the firms are located.[21] To account for differences in the precision of the AI measurement across firms, the regressions are weighted by the firm's number of jobs in Revelio at the start of the period.[22] Standard errors are clustered at the 5-digit NAICS industry level throughout.

Appendix Table B1 details the construction of our regression sample. We start with 4,022 Compustat firms that are matched to the Revelio resume data and have a coverage ratio (number of workers in Revelio over Compustat employment) between 0.05 and 3. Among these firms, 3,270 are U.S. headquartered, 2,748 have at least 50 U.S. jobs in the Revelio data, 2,262 have more than 50% of jobs in the U.S., and 2,215 have Compustat employment of at least 50. Appendix A.1.7 provides extensive analysis highlighting data issues that motivate these data filters. Given that our focus is the *use* and productivity effects of AI by a wide range of firms rather than the *invention* of new AI technologies by AI-producing sectors, we exclude firms in the tech sectors and semiconductor manufacturing, leaving 1,894 firms.[23] After dropping firms missing 2018 or 2024 data or control variables, the final regression sample contains 1,321 firms in 2024; Appendix Table B2 reports summary statistics.

Before turning to the regression estimates, Appendix Figure B8 presents binned scatterplots of productivity growth against AI investments, conditional on industry fixed effects and baseline controls, with each dot representing roughly the same number of firms. The pattern is clear, and, because the plots impose no functional form, not an artifact of outliers: AI investments are unrelated to productivity growth from 2010 to 2018 but positively related from 2018 to 2024.

Table 3 presents the main results. In Panel A, we first revisit and confirm the findings

[21] Commuting-zone controls are log average wage, log population, exposure to Chinese import competition (Autor et al., 2013), and the shares of foreign-born, routine, manufacturing, finance, IT-occupation, college-educated, and female workers. When firms span multiple commuting zones, we use the average across commuting zones weighted by the firm's employees in each commuting zone in the Revelio data.

[22] Since the numbers of worker resumes are correlated with firm size, this weighting scheme also roughly weights firms in accordance to their contribution to the economy.

[23] We exclude firms in two 2-digit NAICS tech sectors: 51 ("Information") or 54 ("Professional, Scientific, and Technical Services"), and in semiconductor manufacturing (NAICS 334413).

in Babina et al. (2024) of no significant correlation between firm AI investments and productivity growth during the rise of AI adoption from 2010 to 2018. Both sales per worker and revenue TFP show small, statistically insignificant coefficients with or without controls: with controls, the coefficient is 0.001 for sales per worker in column 2 and 0.012 for revenue TFP in column 4. Recovering this null relationship in Babina et al. (2024) is a meaningful check, since we use new sources of job postings and resume data and a substantially updated methodology for identifying AI-related jobs.

We then examine whether this relationship changed since 2018. Panel B shows that AI investments are associated with significant and economically meaningful productivity gains in the later period from 2018 to 2024. With all controls included, a one-standard-deviation increase in the share of AI jobs is associated with a 6.2 percentage point increase in sales per worker (column 2) and a 7.5 percentage point increase in revenue TFP (column 4) over the 6-year period—equivalent to about a 1% increase in productivity per year. These are economically meaningful effects: they correspond to 0.14 of a standard deviation in sales-per-worker growth and 0.19 of a standard deviation in TFP growth, indicating that the AI-worker share accounts for a substantial portion of the variation in firm productivity.

Higher productivity can arise either from producing existing output with fewer inputs or from growing output with inputs held fixed. To tease these apart, we turn from productivity to firm growth and ask whether AI-investing firms generate the same sales with fewer workers or more sales with the same headcount. Appendix Table B3 reports the results for log sales (columns 1 and 2) and log employment (columns 3 and 4). Panel A shows that in the earlier period from 2010 to 2018, AI investments increase sales and employment simultaneously, confirming the finding in Babina et al. (2024).[24] In contrast, Panel B shows that from 2018 to 2024 firm AI investments are associated with growth in firm sales but not employment: a one-standard-deviation increase in AI job share change is associated with 6.7 percentage points of additional sales growth, comparable in magnitude to the earlier period. However, AI's relationship with employment growth is close to zero. We confirm this using total Revelio employment excluding AI-related workers, which nets out growth in the AI workforce itself. Appendix Table B5 shows the same pattern as Compustat employment: growth in the number of all workers and non-AI workers is unrelated to AI investments from 2018 to 2024. The persistence of sales growth alongside zero employment growth suggests that AI-investing firms have shifted from labor-intensive expansion toward jobless growth, with sales gains now coming from rising

[24]The magnitudes are smaller than in Babina et al. (2024) due to differences in the coverage of public firms in Revelio and Cognism resume datasets. In Appendix Table B4, we show that restricting the sample to firms in Revelio data to those that also exist in Cognism makes the coefficients larger and closer to the ones in Babina et al. (2024).

labor productivity—more output per worker—rather than from added headcount.

Given AI's positive productivity effects in the recent period, we next ask whether they are broad-based across industries or concentrated in a few sectors. Appendix Table B6 estimates the relationship between AI investments and productivity separately for four non-tech sectors, Manufacturing, Wholesale & Retail, Finance, and Other, as well as the tech sector (which we exclude from the baseline sample). The increase in productivity is present across all industry sectors, consistent with the characterization of AI as a general purpose technology that diffuses across a wide range of economic activities (Goldfarb et al., 2023).

In Table 4, we examine whether the earlier wave of AI technologies in the 2010s (Wave I: machine learning, natural language processing, and computer vision) or the more recent wave in the 2020s (Wave II: generativeand agentic AI) contributes to productivity growth. We regress productivity growth separately on the growth in the share of workers with Wave I AI skills, Wave II AI skills, or both. We estimate separate rather than joint regressions because the two measures are highly correlated—the employment-weighted correlation between the 2018–2024 changes in firms' Wave I and Wave II AI shares is 0.77—in part because many Wave II AI jobs build on Wave I technologies, as discussed in Section 2.5. This pattern is consistent with the organizational technology ladder in Schubert (2025), in which capabilities built for an earlier technology speed the adoption of the next one.

We find that both waves of AI technologies are positively associated with productivity growth over 2018–2024, so the productivity benefits of AI investment are not confined to recent generative and agentic AI advances but extend across established and emerging AI capabilities. Because neither wave's search terms include generic labels such as "AI," these results also show that the effects are not driven by workers inserting fashionable terms like "AI" into their job descriptions. The Wave II result may seem surprising, since generative and agentic AI are recent enough that one might expect little productivity effect to have materialized yet. Two considerations help reconcile this finding. First, the two waves differ in what they are used for: Wave II jobs are considerably more likely to build organization capital than Wave I jobs, which we show in Section 4.2 is the central mechanism underlying AI-fueled productivity growth. Second, the waves are not independent. The data platforms, pipelines, and instrumentation built during the Wave I period provide the raw material on which generative models are trained. Thus, firms that invested early enter the Wave II period with this foundation already in place and can realize productivity gains more quickly than they otherwise would.

3.2. Dynamic Effects

Because firms build AI capabilities gradually and the productivity effects are uncertain and long-lagged, the timing of those effects is difficult to pin down. We therefore estimate

a series of cross-sectional regressions of $\Delta Y_i^{T \to T'}$, for our two productivity measures, on our main independent variable $\Delta \textit{Share AI Jobs}_i^{2018 \to 2024}$ (2018–2024 change in AI job share), for varying start and end dates T and T', controlling for 2-digit NAICS sector fixed effects and weighting by the firm's U.S. Revelio employment at the start of each period.

Figure 6 plots the resulting coefficients. To the left of the dashed vertical line, the dependent variables are changes in productivity from year T to 2018, for $T = 2010, \ldots, 2017$; all coefficients are close to zero for both log sales per worker (panel (a)) and revenue TFP (panel (b)), so firms that subsequently expanded their AI workforce were not already on differential productivity trajectories. To the right, the dependent variable is the change in productivity from 2018 to year T', for $T' = 2019, \ldots, 2024$. The coefficients increase monotonically over time, suggesting that the productivity gains from AI investments build up gradually over 2018–2024.

The pattern is not specific to 2018—the starting year we pick in this paper based on the ending year in Babina et al. (2024) where we found no correlation between AI investments and productivity growth. Appendix Figure B9 repeats the exercise with 2017 and 2019 as base years, redefining the treatment window accordingly, and finds the same result in both: flat coefficients on productivity before the base year and monotonically rising coefficients after. The timing we document therefore reflects the data rather than the particular base year chosen.

To further examine how AI investments affect subsequent productivity growth, we regress ex-post productivity growth from year t to $t+4$ on ex-ante AI investments from year $t-4$ to t in Appendix Table B7. First, we find that AI-investments in preceding years are associated with subsequent growth in firm productivity, confirming that the effects from AI investments occur with a lag. This also helps rule out reverse causality from firm productivity growth to AI investment. Second, the effects are small and insignificant in earlier periods and grow over time toward the end of our sample period, consistent with the returns to AI investment accumulating over time as firms build up the complementary intangible capital needed to realize the technology's benefits.

We then reverse the ordering to assess whether the relationship instead reflects reverse causality. If more productive firms simply invest more in AI because they have the slack resources, managerial capacity, or data infrastructure to do so, we would expect past productivity growth to predict future AI investment. Appendix Table B8 therefore regresses productivity growth from year $t-4$ to t on AI investment from year t to $t+4$. The estimated coefficients are small and statistically indistinguishable from zero across specifications. Firms that experienced faster productivity growth in the preceding four years did not subsequently invest more in AI.

Taken together, these three exercises describe a consistent timing pattern. Firms that

went on to expand their internal AI investments show no differential productivity growth in the years before 2018, so the relationship we document is not the continuation of a pre-existing trend. Following 2018, the estimated effects rise monotonically through the end of the sample, and the lead-lag specifications confirm that AI investment predicts productivity growth over the subsequent four years, while the reverse relationship—past productivity growth predicting later AI investment—is absent. The evidence thus points to a lagged and gradually accumulating return to AI investment, and is difficult to reconcile with the simplest accounts in which productive firms simply select into AI adoption.

3.3. Robustness

We conduct several robustness tests on our key productivity results.

First, we estimate TFP following Imrohoroğlu and Tüzel (2014), who apply the control function approach of Olley and Pakes (1996) to Compustat data. The central problem in estimating firm-level production functions is that inputs are chosen in response to productivity shocks the firm observes but the econometrician does not, which biases OLS estimates of the input elasticities. The control function approach addresses this simultaneity by using a firm's investment demand, which is monotonically increasing in unobserved productivity conditional on capital, to proxy for the productivity shock and recover consistent estimates of the capital and labor coefficients. Firm-level TFP is then computed as the residual from the estimated production function. We do not use this measure in our baseline because it imposes additional data requirements—most notably on capital and investment series—that substantially reduce the sample size. Appendix Table B9 shows that the results using this measure are similar to our baseline. To confirm that the difference in samples is not driving our conclusions, the last four columns of the table re-estimate our baseline productivity measures on the restricted sample for which the control function estimate is available; the coefficients are similar in magnitude and significance.

Second, one potential concern is that the difference between the 2010–2018 and 2018–2024 periods is driven by firms that entered the sample after 2010, which may differ systematically from incumbent firms adopting AI in the earlier period. Appendix Table B10 restricts the sample to firms that were already present in 2010. The results are similar, indicating that the higher productivity observed in the later period does not arise from differences in firm entry or shifts in sample composition relative to the null findings in the earlier period.

Third, we consider several alternative measures of AI investment. In Appendix Table B11, we use the share of the wage bill that goes to AI-related workers to account for heterogeneity in the quality of AI workers: to the extent that wages reflect differences in skill, this measure weights AI workers by their pay rather than counting all AI jobs equally.

Appendix Table B12 uses the global share of AI jobs rather than the U.S. share. Since our baseline measure is a share of domestic employment, firms that offshore non-AI functions could mechanically raise their measured AI shares without any change in AI investment, and computing the share over the firm's worldwide workforce eliminates this composition effect. Both of these measures yield results comparable to our baseline estimates.

A separate concern is that workers may retrofit AI keywords into their job descriptions, particularly following the release of ChatGPT (Bloom et al., 2026b) or powerful Anthropic models and agentic AI tools in late 2025 and early 2026. This would introduce measurement error into our AI measure and attenuate our estimates. Our results comparing Wave I and Wave II AI technologies alleviate this concern: neither wave includes generic terms such as "artificial intelligence," and we find positive effects for Wave I technologies, which predate the generative AI vocabulary. Reassuringly, the measure in Babina et al. (2024), constructed from pre-2022 resume data, is strongly correlated with ours and yields similar results. To address the concern further, we use an earlier vintage of the Revelio Labs data from October 2025 and find nearly identical results in Appendix Table B13. A related concern is that our skill list is compiled from pooled Lightcast data in 2010–2025, so recent AI vocabulary could be applied retroactively to earlier resumes. In Appendix Table B14, we re-identify AI jobs in each year using only the skills observed in Lightcast job postings in that year, and the estimates are essentially unchanged.

Fourth, a concern is that AI-investing firms may engage in more acquisitions, in which case productivity growth alongside flat employment could arise mechanically from acquiring asset-light or intangible-intensive targets. We first show in Appendix Table B15 that our results are unchanged when we exclude jobs at acquired subsidiaries from the AI investment measure. In addition, because externally purchased intangibles are capitalized on the balance sheet, we examine the relationship between AI investments and firms' balance-sheet intangibles (both in total and the acquired component) in Appendix Table B16, and find no significant relationship. This suggests that the productivity gains we document reflect organic improvements within firms rather than the mechanical effect of absorbing intangible-intensive targets.

Fifth, while our baseline specification already excludes firms in the tech sector and semiconductor manufacturing, we further exclude the remaining Magnificent 7 firms (Apple and Tesla) and find similar results (Appendix Table B17), indicating that our results are not driven by a handful of large, technology-oriented firms.[25]

Finally, while the results in Section 3.2 suggest that reverse causality is unlikely, con-

[25]Five of the Magnificent 7 firms are already excluded from our baseline analysis: Microsoft, Alphabet, and Meta are in the tech sector; Nvidia is in the semiconductor manufacturing industry; and Amazon has over half of its jobs outside the U.S. in 2024.

temporaneous unobserved shocks may still drive both AI investments and productivity growth. We address several such possibilities. (i) AI-investing firms may experience positive productivity shocks or already be on a higher growth trajectory. In Appendix Table B18, we control for initial levels of log sales per worker and revenue TFP, as well as state fixed effects to absorb state-level shocks, and find similar results. We further confirm in Appendix Table B19 that our results are robust to controlling for pre-trends in sales, employment, and TFP over 2010–2018, providing additional assurance that they are not driven by differential pre-existing trajectories among AI-investing firms. (ii) We control for detailed industry fixed effects in Appendix Table B20 to absorb industry-level productivity shocks. Comparing firms within the same industry also holds fixed differences in technology, input use, and market conditions that vary across industries, so that measured productivity better reflects firm-specific efficiency rather than structural industry characteristics. The results are nearly identical even with granular 5-digit NAICS fixed effects. (iii) Because our sample period spans the pandemic, we show in Appendix Table B21 that our results are robust to controlling for firms' Covid exposure (Hassan et al., 2023) and for the share of remote-suitable jobs (Dingel and Neiman, 2020). (iv) AI investments may be correlated with investments in related technologies such as IT. We measure firms' IT investments using the share of jobs in computer- and IT-related occupations excluding AI jobs. Appendix Table B22 shows that non-AI IT investments are not significantly associated with productivity growth over our sample period, and controlling for them does not alter the estimated effects of AI investments.[26] These additional tests address the most likely confounders, but they cannot rule out unobserved shocks entirely. We therefore turn to an instrumental variable strategy next.

3.4. Instrumental Variable Estimates

We instrument for firm-level variation in AI investment based on each firm's differential exposure to the supply of AI-trained graduates from U.S. universities following Babina et al. (2024). Specifically, we instrument firm-level AI investment using variation in firms' ex-ante exposure to the supply of AI talent from universities with strong AI research historically. Because AI-skilled labor is a key constraint on firm investment in AI, firms with greater (historical) hiring ties to these universities, which have produced more AI-trained graduates in recent years, have better access to AI-skilled talent and are better positioned to invest in AI. This instrument and its variants have been used in published (e.g., Krupa and Mullaney 2026; McLemore and Mihov 2026) and ongoing work (e.g., Hampole et al. 2025).

[26]Because we find consistent results in manufacturing and non-manufacturing sectors, it is unlikely that our estimates are driven by robot adoption, which is concentrated in manufacturing.

A full description of the data, methodology, and validation exercises for the instrument are provided in Appendix C. We next summarize the key steps for its construction. We first define AI-strong universities in the U.S. as those with a high share of AI researchers prior to 2010 (before the rise of commercial interest in AI in the 2010s) in the Open Academic Graph (OAGv2) publications data. This definition of AI-strong universities comes from Babina et al. (2024), which provides a detailed description of the methodology. Using our Revelio resume data, we then construct the historical firm-university hiring networks. Our resume data account for approximately 41% of all degrees conferred by all universities in our sample in 2010 according to NSF Integrated Postsecondary Education Data System (IPEDS), with a strong correlation (0.67) between new graduates in our sample and total IPEDS degrees. Validating our classification, the growth in AI-trained graduates during the 2010s was substantially larger at AI-strong universities compared to their non-AI-strong counterparts (Appendix Figure C.2). For our instrumental variable strategy to work, firms must recruit from differentiated and stable sets of universities. Both conditions hold: firms tend to draw talent from a concentrated group of schools,[27] and their pre-2010 hiring networks are highly predictive of where they recruit after 2018 (Appendix Table C.3).

We define the instrument as:

$$IV_i = \sum_u s_{iu}^{Pre2010} AIstrong_u,$$

where $s_{iu}^{Pre2010}$ is the share of STEM workers in firm i between 2005 and 2010 who had their most recent degree (undergraduate or graduate) from university u,[28] and $AIstrong_u$ equals one if university u is an AI-strong university based on pre-2010 publications.

The exclusion restriction requires that a firm's hiring shares between 2005 and 2010 affect outcomes over 2018–2024 only through its exposure to the supply of AI-trained labor. Equivalently, a firm's historical hiring network measured well before AI adoption became widespread must be uncorrelated with unobserved shocks to firm outcomes during our sample period. One potential threat is that ties to AI-strong universities also proxy for exposure to technically skilled graduates more generally, or for connections to elite universities that produce both AI talent and other high-skilled labor. In that case, the instrument would capture access to skilled labor broadly rather than to AI-trained workers specifically. To address this concern, we control for firms' historical links to computer-

[27]The average firm sources 19% of its fresh graduates from a single dominant university, 47% from its top five universities, and 61% from its top ten. By comparison, the single largest university accounts for only 1.6% of the nationwide pool of fresh graduates, the five largest for 7.1%, and the ten largest for 12.9%.

[28]STEM workers are workers with STEM majors in college, including biology, chemistry, engineering, information technology, mathematics, physics, and statistics.

science-strong and top-20 universities: $\sum_u s_{iu}^{Pre2010} CSstrong_u$ and $\sum_u s_{iu}^{Pre2010} Top20_u$, where $CSstrong_u$ is the average share of (non-AI) CS researchers at university u between 2005 and 2009, and $Top20_u$ equals one if a university is among the top 20 universities ranked by the U.S. News & World Report. A related concern is geographic: firms located in technology hubs may both hire from nearby AI-intensive universities and benefit from regional agglomeration effects. Our analysis already excludes firms in the tech sector, which are most susceptible to this channel, and we include state fixed effects in some specifications to further address this concern.

We conduct several additional analyses in Appendix C to address other threats to the identifying assumption. We find no evidence that firms increased their hiring from AI-strong universities during the pre-period in anticipation of future AI needs. We also show that firm connections to AI-strong universities do not predict higher productivity growth before 2010 or over 2010–2018, consistent with the absence of pre-trends and the lagged effects of AI investment documented in Section 3.2.

Appendix Table C.4 presents the first-stage results.[29] Across all specifications, we include our standard controls (industry fixed effects and a set of baseline controls at the firm, industry, and commuting-zone levels) and controls for exposure to CS-strong and top-20 universities. We then progressively introduce (i) state fixed effects, which absorb local labor market conditions and regional agglomeration effects, and (ii) pre-period growth in firm sales, employment, and TFP from 2002 to 2010, which helps account for unobserved firm-level attributes that could jointly influence growth and productivity trajectories and the propensity to hire AI talent. The instrument has a strong first stage with F-statistics above 10 for all specifications, suggesting that firms' pre-2010 exposure to AI-strong universities strongly predicts AI investment after 2018.

Table 5 presents the 2SLS estimates. We include baseline controls in all columns and sequentially add state fixed effects and pre-trends in growth and productivity. The results show a robust and significant effect of AI investments on log sales per worker and revenue TFP. When all controls are included, a one-standard-deviation increase in AI investments leads to an increase in log sales per worker of 14.8%, and an increase in TFP of 11.4%.[30] These results support causal interpretation of firms' internal AI investments leading to increased productivity.

[29]The number of observations is smaller than OLS regressions because we only keep firms for which over half of the STEM workers between 2005 and 2010 have non-missing university information and had the most recent degree from one of the 659 universities with publications data.

[30]The magnitudes are larger than OLS estimates, which may be due to measurement error, or a negative bias from firms with lower projected productivity growth investing in AI.

4. Mechanisms

In this section, we explore potential mechanisms for the positive relationship between AI investments and productivity growth in recent years.

4.1. Firm AI Investments: Tangible versus Intangible Capital

We start by exploring whether AI investments are associated with more investments in intangible or tangible capital. This is a natural starting point because investments in new technologies can either raise labor productivity by increasing tangible capital (i.e., capital deepening), or increasing intangible capital, or both. Adopting general-purpose technologies typically involves not only investment in physical capital, but also complementary innovation and restructuring (Bresnahan and Trajtenberg, 1995). The IT literature makes this point directly: computers and digital technologies generate value when combined with complementary investments in organization capital, worker skills, and redesigned business processes (Brynjolfsson et al., 2002b; Bartel et al., 2007). AI could, therefore, affect both margins: it may raise tangible investment like computing and data storage infrastructure and related equipment, but much of the firm-specific investment needed to make AI productive—experimentation, data curation, training, workflow redesign, and integration—is intangible (Corrado et al., 2009). Examining whether AI-investing firms expand investments in intangible relative to tangible capital therefore provides a first-order test of whether AI investment operates mainly through capital deepening or through a broader transformation of organization and knowledge capital.

We examine the relationship between firm AI investments and firm tangible (property, plant and equipment (PP&E)) and intangible (SG&A and R&D) investments during our focal period from 2018 to 2024. Appendix Table B23 shows evidence consistent with AI-investing firms investing more in intangibles but not in tangible capital. Columns 1 and 2 show that AI-investing firms experience significant increases in SG&A on the order of 5%. The increase in R&D expenditures is half of that: 2.5% when all controls are included in column 4.However, columns 5 and 6 show no increase in tangible investments in the PP&E: when all controls are included, the coefficient is -0.7% and statistically insignificant.

While accounting measures, such as R&D and SG&A, have proven valuable for quantifying intangible investment, they capture the flow of investment spending, while the resulting intangible capital takes the form of knowledge and capabilities held by and built by workers. Because intangible capital is produced and embodied by skilled labor, a firm's worker composition—the prevalence of highly educated, experienced, and highly paid employees—offers a complementary measure that reflects the stock of knowledge and capabilities residing in the firm's workforce. In Table B24, we repeat the main specifications

from 2018 to 2024 where the dependent variables are changes in the shares of workers by education and seniority levels as well as change in average wage. We exclude AI-related workers when computing these outcomes to avoid mechanical relationship since AI-related jobs tend to be more senior and more educated. We find that AI-investing firms expand their employment of workers with advanced degrees (masters or PhDs) and in more senior roles. By contrast, the coefficient for shares of workers without college degrees or low seniority levels is negative and significant.[31] These patterns align with recent evidence that AI adoption disproportionately displaces junior workers (e.g., Brynjolfsson et al. (2025a); Hosseini Maasoum and Lichtinger (2025); Kim and Koning (2026)). In Appendix Table B26, we use average wages from the Revelio data as another proxy for worker skills and find consistent results: higher average wage growth in firms that invest more in AI.

The rise in spending on SG&A and R&D coupled with an increase in hiring of high-skilled workers suggests that AI-investing firms undergo meaningful shifts in investments towards intangible capital in the recent period. This finding contrasts with previous technological advances, which spurred tangible investment and capital deepening among firms adopting new technologies. These results warrant a caveat: capital expenditures on AI infrastructure, such as data centers, are certainly occurring—but they are concentrated upstream among hyperscalers, the large cloud service providers that own and lease global data-center infrastructure and specialized AI hardware. Most non-technology firms do not build frontier-scale compute infrastructure themselves; instead, they rent AI infrastructure capacity.[32] For the non-tech firms in our sample, AI investments therefore appear as a buildup of intangible capital rather than as large increases in tangible capital.

4.2. Firm AI Investments and Organization Capital

An important component of intangible capital is organization capital (OC) (Peters and Taylor, 2017). Given that firm AI investments are associated with a significant increase in SG&A—a key indicator of investment in organization capital—we next examine the role of organization capital in linking AI investments to productivity growth.

Organization capital is intangible, durable, and firm-specific knowledge (Prescott and Visscher, 1980). Firms accumulate it through learning-by-doing, and as they do, they operate more efficiently and become more productive (Atkeson and Kehoe 2005). It is built

[31] Appendix Table B25 examines employment levels and shows that the total employment of highly-educated and senior workers also significantly increases, whereas the employment of workers without college degrees and low seniority declines as firms invest in AI.

[32] Consistent with this view, the FTC identifies Alphabet, Amazon, and Microsoft as the largest cloud service providers involved in major generative-AI partnerships, while McKinsey estimates that hyperscalers will account for roughly 70 percent of forecast U.S. data-center capacity through owned or leased infrastructure (Federal Trade Commission, 2025; Arora et al., 2025).

by workers (Eisfeldt and Papanikolaou, 2013a) and resides in a firm's operating routines (Crouzet et al., 2022). Typically, firms take a long time to build up organization capital and to observe the resulting productivity improvements (Atkeson and Kehoe, 2005). But the faster a firm able to learn, the more quickly it accumulates organization capital and realizes productivity gains.

Organization capital is distinct from related concepts, such as R&D, IT, customer capital, product, process, architectural innovation, or co-invention; while absorptive capacity and architectural competence are specific sup-types of OC. Appendix A.4.1 provides details on the comparisons to these related concepts and discusses core features that define OC. Specifically, OC is *built rather than purchased*, accumulating gradually through firms' learning from their operations rather than being bought on the market. OC is *embedded rather than separable*, residing in platforms, systems, routines, workflows, management practices and coordination capabilities, which distinguishes it from the inputs to intangible capital (e.g., R&D expense) and the outputs they generate (e.g., patents). OC is *durable rather than ephemeral*, representing repeatable useful knowledge rather than one-off outputs or knowledge. Finally, OC is *organization-specific rather than general*, which cuts two ways. Within the firm, the asset is non-rival in the sense of Jones and Tonetti (2020): an established routine or system serves many employees, units, and tasks at once without depletion, so it scales internally at little additional cost. Outside the firm it travels poorly, since routines tuned to one organization lose much of their value elsewhere, making it far less portable and marketable than patents or data.

What is special about AI technologies is that they are fundamentally "learning machines" (Turing 1950). A firm's AI investments can therefore create new organization capital, accelerating learning from its own operations: discovering how to improve workflows, respond to customer needs, and redesign processes and product development. For firms implementing AI internally, models trained on firm-specific data and embedded in firm-specific systems and workflows are themselves part of OC, since their value depends on the organization in which they are deployed. These models are trained and maintained by highly skilled AI workers who collaborate with colleagues with domain expertise and deep knowledge of the firm's workflows, business problems, and customer needs. Thus, AI investments can create organization capital by accelerating the generation of firm-specific productive knowledge—embodied in organizational routines—that enables higher productivity. Consistent with AI technologies being used for the creation of OC, AI-related jobs are almost twice as likely to be classified as OC-related than non-AI-related jobs (see Section 2.5).

To organize these ideas and derive testable predictions, Appendix D develops a model of AI and organization capital that extends the firm life-cycle framework of Atkeson and

Kehoe (2005). The key state variable, *organization capital*, is the durable asset a firm accumulates by running its operations and learning from them: the procedures, structured data, trained models, and codified workflows that the firm continues to use. Because organization capital is built from the firm's own operations, the flow into this stock is proportional to the scale of those operations: what the firm does gets recorded, documented, and used to train models. Organization capital can therefore accumulate only as quickly as the firm generates and learns from its own activity. AI enters our model as an increase in the rate at which a firm's operations are converted into organization capital. In this sense, AI helps make firm's existing activity to be more productive as a source of organizational knowledge. Code repositories, service transcripts, clinical documentation—records of what the firm has done—are used to train AI models and become an operational asset, and the approaches, conventions, and firm-specific knowledge of a particular employee become a firm-specific asset available to many workers within the firm (Singhal et al., 2023; Brynjolfsson et al., 2025b) and help improve their productivity.

The theoretical framework yields two predictions that we can test in the data. First, and central to the timing patterns we document, the adjustment is slow *without any adjustment costs or implementation lags in the model*. Because organization capital is a stock and AI only affects the flow into it, its accumulation remains gradual even when AI accelerates its buildup. Second, the gains should be largest for firms that start with relatively little organization capital for their size. Think of organization capital as a firm's operating playbook: firms with a thin playbook have many processes, routines, and sources of knowledge left to codify and improve, whereas firms with an already extensive playbook have less left to add. As a result, new additions can be large relative to the existing stock for organization-capital-poor firms but only incremental for firms that have already accumulated substantial OC.

To illustrate how AI accelerates the accumulation of organization capital concretely, Appendix A.4.4 presents case studies of eight firms spanning finance, retail, manufacturing, technology, and healthcare, and covering seventeen of the organization-capital-building job categories in Table 2. How AI performs the conversion is largely common across them: a model is trained on records the firm was already producing—research documents, meeting transcripts, sensor readings, product and project data—turning material that had merely documented past work into an asset that can be queried, simulated against, or acted on. What it converts differs sharply. At Morgan Stanley, it is professional judgment: an assistant built on the firm's own research library, and a tool that takes notes during advisor meetings and writes them into the CRM, capturing a record that previously left no trace. At Procter & Gamble, it is physical process knowledge: plant telemetry becomes a digital twin against which process changes are trialed without halting the line. At Home

Depot, it is the expertise of store associates—knowledge that had been available only in one aisle of one store at a time—assembled into a knowledge base available everywhere at once. In each case, the raw material is the firm's own operational record rather than anything purchased or generally available; the conversion is performed by the firm's own AI workers, whose roles are precisely what our measure captures; and the resulting asset is useful to many workers inside the firm.

We now examine more systematically whether the model's predictions are consistent with our data. The first prediction concerns timing, and it matches what we find in Figure 6: the effect of AI investment on productivity does not appear all at once but builds gradually over the years following the investment. Because organization capital is a stock, a firm that raises its rate of learning from operations does not see the gain immediately: for plausible depreciation rates, the model implies that productivity gains take more than five years to reach half their eventual size. Two features of our setting reinforce this pattern. Most firms began investing intensively in AI only after 2015, so much of the early period contains little treatment to respond to; and the composition of AI jobs has shifted over time toward roles that are more likely to build organization capital (Figure 7). Both help explain why we find significant growth in productivity in AI-investing firms during 2018–2024, but not during 2010–2018 (Table 3).

Second, in Table 6, we examine whether firms' AI investments are associated with growth in organization capital constructed with capitalized SG&A expenses (Eisfeldt and Papanikolaou 2013a). In the earlier period of 2010–2018 (when we find no association between AI investments and productivity growth), there is no significant build up of organization capital (Panel A). However, during the recent period from 2018 to 2024, we find a significant growth in organization capital in firms that grew their AI investments more (Panel B). Applying the same lead-lag specification used above for productivity, we find the same timing pattern for organization capital in Appendix Table B27: past AI investment (from year $t-4$ to t) predicts future organization capital growth (from t to $t+4$). Crucially, Appendix Table B28 shows that the reverse relationship is absent—past organization capital growth is not correlated with subsequent AI investment—so the association does not appear to be driven by organizationally capable firms selecting into investing more in AI. The direction of predictability thus runs from AI investment to organization capital, mirroring what we find for productivity, and confirming the predictions of our model: firm productivity growth becomes visible when AI investments help increase the stock of organization capital substantially. Hence, these results suggest a mechanism for the long lag between AI investment and measurable productivity growth:

the intermediate step of accumulating organization capital is itself slow.[33]

Third, to further tighten the link between AI, organization capital, and productivity, we split AI jobs into (i) organization-capital-building-related and (ii) non-organization-capital-building-related components. If AI investments raises productivity by accelerating the accumulation of organization capital, the productivity gains should come from the first (i) rather than the second (ii) component—a sharper test than the aggregate relationship between firm overall AI investments and productivity, since it holds AI investment fixed and varies only what that investment is used for. Table 7 confirms this: growth in the share of organization-capital-building-related AI jobs is strongly and positively associated with productivity growth.[34] In columns 3 and 6 of Table 7, we additionally include the growth in share of non-organization-capital-related AI jobs. The coefficient on organization-capital-related AI jobs remains essentially unchanged. However, non-organization-capital-related AI jobs exhibit no significant relationship with productivity growth, and this null result also holds when we do not include organization-capital-related AI jobs as independent variable. This contrast suggests that the productivity gains from firms' investments in AI operate primarily through the organization capital channel.

To provide additional evidence on this mechanism, Appendix Table B31 shows that organization-capital-related AI jobs are associated with higher growth in organization capital, proxied by capitalized SG&A expenses. A potential concern is that firms already rich in organization capital both hire more organization-capital-related AI workers and enjoy faster productivity growth for reasons unrelated to AI. However, our estimates are virtually unchanged when we control for firms' initial level of organization capital (see Appendix Table B32 for productivity and Appendix Table B33 for organization capital).

Finally, the model's second prediction is sharper than an average effect, and not obvious ex ante. The speed at which organization capital and productivity grow following AI investments depends on the firm's state only through one statistic: its initial stock of organization capital relative to its size. In particular, firms with a low stock of organization capital relative to size should respond both more and faster. We test this by splitting firms at the median of organization capital scaled by PP&E, measured before the focal period, and interacting AI investment with the two groups. Because the model's prediction concerns organization capital relative to size, and because size itself is one of the outcomes,

[33] In Appendix Table B29, the IV estimates support a causal interpretation of the effect of AI investment on organization capital growth.

[34] This result is not trivial since only about 50% of all AI jobs in our samples are classified as organization-capital-related under the baseline 0.6 threshold. Table 7 has fewer observations than the baseline results Table 3 because the AI jobs in these firm-years—typically only one or two—have empty job descriptions. In Appendix Table B30, we find similar results with alternative thresholds of 0.5 and 0.7 for organization-capital-related jobs.

we include size-tercile fixed effects so the comparison holds scale fixed. Table 8 reports productivity growth and Appendix Table B34 reports organization capital growth. The pattern matches the prediction on both margins. For productivity, below-median firms show coefficients of 0.092 for sales per worker and 0.105 for TFP, both significant, while the above-median coefficients are about half as large and statistically insignificant. For organization capital, the coefficient for below-median firms is roughly two to three times larger than for above-median firms.

Taken together, the evidence points to organization capital as the central mechanism linking firm AI investments to productivity growth. AI-investing firms build organization capital along every margin we measure: they expand SG&A and their organization capital stock. Moreover, the productivity gains accumulate gradually and are concentrated precisely where the mechanism predicts: in AI jobs related to building organization capital, rather than in other AI jobs. We also find that firms with a low initial stock of organization capital benefit more, consistent with the model's prediction. These patterns suggest that AI raises productivity by generating firm-specific knowledge—embodied in systems, routines, and firm-trained models—that enables firms to learn faster and operate more efficiently.

4.3. Task Automation and Labor Replacement

A common argument holds that AI can substitute for certain worker tasks, potentially displacing some jobs. Consistent with this view, Brynjolfsson et al. (2025a) show that entry-level workers in occupations more exposed to AI experience lower employment growth, and when firms automate such occupations, their measured productivity can rise (Hampole et al., 2025). If this channel were driving our findings, we would expect firm employment to be negatively correlated with AI investments, which we do not find to be the case. Moreover, under this mechanism, we would expect AI-investing firms to shed the workers most exposed to AI—who, according to standard exposure measures—tend to be highly skilled, technical, educated, and higher-paid. Our results point in the opposite direction: AI investments are associated with an increase in both the employment levels and shares of highly educated workers (Tables B24 and B25) as well as average wages (Table B26), and in unreported results, we find that the share of STEM workers also increase in AI-investing firms.

These findings do not rule out labor-replacing effects of AI in a different setting: instead, they suggest that this mechanism does not drive productivity effects in our setting. First, our measure captures the total effect of all AI applications, not only the labor-replacing ones; second, we focus on internal firm AI development and implementation, whereas AI labor replacement effects might be concentrated in the uses of outsourced AI tools like chatbots. Even on that margin, however, large-scale usage data suggest that adoption at work remains shallow and predominantly collaborative: AI is used for roughly a fifth

of tasks in the median occupation with any AI use, and end-to-end task automation is limited in scope (Iscenko et al., 2026). Overall, the evidence suggests that our firm-level productivity effects are not explained by AI-investing firms shedding employment.

A related possibility is that AI—and generative AI in particular—replaces some tasks within jobs and reallocates workers toward other activities, thereby raising their productivity (e.g., Noy and Zhang 2023). This channel also appears unlikely to explain our results. Demirer et al. (2026) show that large task-level gains from AI coding tools attenuate sharply before reaching completed projects and releases. We find that generative and agentic AI are not driving the productivity increases we document (see Table 4). Moreover, only a small share of AI-related jobs in our sample involve workers using AI tools to help with job tasks (around 1%), whereas the majority are devoted to building organization capital. This suggests that the AI investment captured by our measure centers not on substituting for existing tasks, but on creating new organizational structures, processes, and complementary capabilities—consistent with organization capital accumulation, rather than task replacement, being the primary channel through which firm AI investments raise productivity in our setting.

4.4. Firm AI Investments and Price Increases

Another alternative explanation for our findings is that AI increases firms' market power. If AI-investing firms gain market power and raise prices, revenue-based productivity measures will rise even absent any growth in real output. Although we lack firm-level data on quantities and prices, we can assess this channel using industry-level data. If the firm-level productivity gains reflect transfers from competitors or consumers, rather than an expansion of the productive frontier, they should not aggregate into industry-level productivity improvements. Moreover, if AI-investing firms grow sales by raising prices (and not output), we should also observe higher industry-level prices.

In particular, we estimate the following long-differences specification at the 5-digit NAICS industry level:

$$\Delta \ln y_{j,[2018,2024]} = \gamma \, \Delta\text{Share AI Jobs}_{j,[2018,2024]} + \text{IndustrySectorFE} + \varepsilon_j, \tag{3}$$

where $\Delta \ln y_{j,[2018,2024]}$ denotes the change in industry-level productivity for all Compustat firms in industry j—including both entrants after 2018 and firms that exited before 2024—and $\Delta\text{Share AI Jobs}_{j,[2018,2024]}$ is the corresponding change in the share of AI workers among Compustat firms in industry j. We also examine changes in producer price indices (PPI) from Bureau of Labor Statistics (BLS). As in the firm-level specifications, regressions are weighted by total industry resume counts in 2018.

Table B35 shows that the industry-level results closely mirror the firm-level patterns:

industries with higher AI investments experience significantly higher productivity growth between 2018 and 2024.[35] In the last two columns, we test the price increase channel directly using PPI. Industries that invest heavily in AI exhibit little change in price indices from 2018 to 2024, indicating that the measured productivity gains reflect increases in quantities produced rather than higher prices.[36]

In sum, the industry-level evidence is inconsistent with the market-power explanation. The productivity gains associated with AI investments do not vanish upon aggregation, as they would if they reflected transfers across firms, and they are not accompanied by rising prices, as they would be if AI-investing firms were simply charging more. Instead, the evidence indicates that AI investments expand real output: the productivity growth we document reflects an outward shift in the productive frontier rather than a redistribution of rents.

5. Conclusion

Whether artificial intelligence will raise productivity is perhaps the central question about its economic consequences. This paper provides early firm-level evidence that the productivity payoff has begun to arrive. Using new measures of firms' internal AI investments through 2024—covering both the first wave of machine learning, natural language processing, and computer vision and the recent wave of generative and agentic AI—we show that U.S. public firms investing more in AI experienced significantly faster growth in labor productivity and TFP between 2018 and 2024, in sharp contrast to the absence of such gains over a decade prior to that. The evidence is consistent with organization capital being an important mechanism. AI-investing firms expand SG&A and R&D rather than physical capital, shift toward more educated and senior workers, and accumulate more organization capital. The productivity gains are concentrated in AI jobs focused on developing firm-specific knowledge, systems, and processes—the core signs of knowledge accumulation and build up of organization capital.

Our findings suggest cautious optimism. The AI-investing firms in our sample are early and intensive AI investors whose experience offers an initial signal of AI's broader potential. Indeed, our evidence on increased productivity is based on U.S. public firms that has been investing in AI at about twice higher rates than other U.S. entities, suggesting that our focal firms are at the frontier of AI investments and thus serve as canaries—not in

[35] We find no productivity effects at the industry level in the earlier period from 2010 to 2018.

[36] The industry AI share in Table B35 is constructed from publicly listed firms alone. In Appendix Table B36, we reconstruct the AI share using all public and private firms in the Revelio data and confirm that price indices move little with this broader measure of AI investment. We also find that industry-level total employment follows a similar pattern to the firm-level results.

a coal mine, warning of danger, but in a gold mine, offering an early signal of the gains that broader AI diffusion may bring. If the history of general-purpose technologies is a guide, the productivity gains may grow as more firms accumulate organization capital. Understanding how AI-driven productivity gains evolve and how they are distributed across firms remains an important direction for future research.

References

Abis, S., Veldkamp, L., 2024. The changing economics of knowledge production. Review of Financial Studies 37, 89–118.

Abraham, K.G., Hershbein, B., Houseman, S.N., Truesdale, B.C., 2024. The independent contractor workforce: New evidence on its size and composition and ways to improve its measurement in household surveys. ILR Review 77, 336–365.

Abraham, K.G., Houseman, S.N., Truesdale, B.C., 2026a. Are the self-employed really their own boss?: Evidence on the degree of autonomy and control the self-employed have over their work .

Abraham, K.G., Houseman, S.N., Truesdale, B.C., 2026b. The" missing self-employed": How better data change the picture of self-employment in the united states .

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

Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from U.S. labor markets. Journal of Political Economy 128, 2188–2244.

Adams, J.J., Fang, M., Liu, Z., Wang, Y., 2026. The Rise of AI Pricing: Trends, Driving Forces, and Implications for Firm Performance. Journal of Monetary Economics 157: 103875.

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60, 323–351.

Agrawal, A., Gans, J., Goldfarb, A., 2022. Prediction machines, updated and expanded: The simple economics of artificial intelligence. Harvard Business Press.

Akcigit, U., Chikis, C.A., Dinlersoz, E., Goldschlag, N., 2026. Attention (and Money) is all you need: Why universities are struggling to keep AI talent. Technical Report. National Bureau of Economic Research.

Aldasoro, I., Gambacorta, L., Pál, R., Revoltella, D., Weiss, C., Wolski, M., 2026. AI Adoption, Productivity and Employment: Evidence from European Firms. BIS Working Papers 1325. Bank for International Settlements.

Argote, L., Beckman, S.L., Epple, D., 1990. The persistence and transfer of learning in industrial settings. Management Science 36, 140–154.

Arora, C., Sorel, M., Sachdeva, P., Bhan, A., He, J., Shaw, N., Garg, R., Ravishankar, S., 2025. The next big shifts in ai workloads and hyperscaler strategies. Estimates that hyperscalers will capture about 70 percent of forecast capacity in the U.S. market through owned or leased options.

Atkeson, A., Kehoe, P.J., 2005. Modeling and measuring organization capital. Journal of Political Economy 113, 1026–1053.

Autor, D., Dorn, D., Hanson, G.H., 2013. The china syndrome: Local labor market effects of import competition in the united states. American Economic Review 103, 2121–68.

Babina, T., 2026. Understanding firms' ai efforts and their economic impact. The Review of Corporate Finance Studies , cfag027.

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

Baily, M.N., Byrne, D.M., Kane, A.T., Soto, P.E., 2025. Generative AI at the Crossroads: Light Bulb, Dynamo, or Microscope? Finance and Economics Discussion Series 2025-053. Board of Governors of the Federal Reserve System.

Barrios, J.M., Campbell, J.L., Johnson, R.G., Liu, C., 2025. Artificially Intelligent or Artificially Inflated? Determinants and Informativeness of Corporate AI Disclosures. Working paper.

Bartel, A., Ichniowski, C., Shaw, K., 2007. How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. Quarterly Journal of Economics 122, 1721–1758.

Baslandze, S., Edwards, Z., Graham, J.R., McClure, T., Sparks, M., Meyer, B., Waddell, S.R., Weitz, D., 2026. Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. Journal of Finance: Insights and Perspectives Forthcoming.

Baviskar, S., Branstetter, L., Drayton, C., Fujiy, B.C., Hovy, E., Tambe, P., Wu, L., 2026. Elite talent, invention, and firm productivity in the age of ai .

Benkard, C.L., 2000. Learning and forgetting: The dynamics of aircraft production. American Economic Review 90, 1034–1054.

Benton Institute for Broadband and Society, 2026. Verizon taps AI agents for network automation. June 12. https://www.benton.org/headlines/verizon-taps-ai-agents-network-automation.

Beraja, M., Talamàs, E., 2026. The value of organizational learning technologies. Working Paper .

Bick, A., Blandin, A., Deming, D.J., Fuchs-Schündeln, N., Jessen, J., 2026. Mind the gap: ai adoption in europe and the us. Working Paper, Federal Reserve Bank of St. Louis.

Blank, M., Schubert, G., Zhang, B., 2025. The Household Impact of Generative AI: Evidence from Internet Browsing Behavior. Working Paper. UCLA Anderson School of Management and UC Irvine Merage School of Business. Presented at the 2025 American Economic Association (AEA) Annual Meeting.

Bloom, N., Moore, G., Simon, L.K., Wilkie-Rogers, C., 2026a. Time Travel on Professional Profiles. Working Paper 35546. National Bureau of Economic Research. URL: http://www.nber.org/papers/w35546, doi:10.3386/w35546.

Bloom, N., Moore, G., Simon, L.K., Wilkie-Rogers, C., 2026b. Time Travel on Professional Profiles.

Blumenfeld, A., Hazell, J., Lian, C., Schaab, A., 2026. The macroeconomic effect of ai: Sizing the software engineering channel. Working Paper .

Bonelli, M., 2026. Data-Driven Investors. Review of Financial Studies 39 (7): 1909–1969.

Bonney, K., Breaux, C.L., Dinlersoz, E., Foster, L.S., Haltiwanger, J.C., Pande, A.A., 2026. The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks.

Bresnahan, T.F., Brynjolfsson, E., Hitt, L.M., 2002. Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. Quarterly Journal of Economics 117, 339–376.

Bresnahan, T.F., Greenstein, S., 1996. Technical progress and co-invention in computing and in the use of computers. Brookings Papers on Economic Activity. Microeconomics , 1–77.

Bresnahan, T.F., Trajtenberg, M., 1995. General purpose technologies: 'engines of growth'? Journal of Econometrics 65, 83–108.

Brynjolfsson, E., Chandar, B., Chen, R., 2025a. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper.

Brynjolfsson, E., Hitt, L.M., 2000. Beyond computation: Information technology, organizational transformation and business performance. Journal of Economic Perspectives 14, 23–48. doi:10.1257/jep.14.4.23.

Brynjolfsson, E., Hitt, L.M., Yang, S., 2002a. Intangible assets: Computers and organizational capital. Brookings Papers on Economic Activity 2002, 137–199. doi:10.1353/eca.2002.0003.

Brynjolfsson, E., Hitt, L.M., Yang, S., 2002b. Intangible assets: How the interaction of computers and organizational structure affects stock market valuations. Brookings Papers on Economic Activity 2002, 137–198.

Brynjolfsson, E., Li, D., Raymond, L.R., 2025b. Generative AI at work. Quarterly Journal of Economics 140, 889–942.

Brynjolfsson, E., Rock, D., Syverson, C., 2019. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics, in: Agrawal, A., Gans, J., Goldfarb, A. (Eds.), The Economics of Artificial Intelligence: An Agenda. University of Chicago Press, Chicago, pp. 23–57.

Cai, K., 2025. Verizon says Google AI for customer service agents has led to sales jump. Reuters, April 9. URL: https://money.usnews.com/investing/news/articles/2025-04-09/verizon-says-google-ai-for-customer-service-agents-has-led-to-sales-jump.

Cai, W., Prat, A., Yu, J., 2026. Measuring organizational capital. Journal of Accounting Research .

Cao, S., Jiang, W., Wang, J., Yang, B., 2024. From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses. Journal of Financial Economics 160: 103910.

Cao, Y., Liu, M., Qiu, J., Zhao, R., 2025. Determinants and Information Content of Voluntary Artificial Intelligence Disclosure. Working paper.

Challamel, B., 2023. Moderna, powered by AI: Our journey to becoming a real-time AI organization. Moderna Blogs, November 29. https://www.modernatx.com/media-center/all-media/blogs/moderna-powered-AI.

Charoenwong, B., Kwan, A., Traina, J., 2026. The composition of intangible capital: Evidence from llm classification of corporate disclosures. Working paper.

Chawla, M., Friedman, H., Kim-Gina, J., 2025. The real effects of dissemination of corporate disclosures on innovation. UCLA Electronic Thesis/Dissertation Repository Dissertation.

Chen, S., Sialm, C., Xu, D.X., 2025. The Growth and Performance of Artificial Intelligence in Asset Management. Working paper.

Chequer, M., Herkenhoff, K., Papanikolaou, D., Rothbaum, J., Schmidt, L., Seegmiller, B., 2026. Labor, despite ai. Working Paper .

Cockburn, I.M., Henderson, R.M., 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. The Journal of Industrial Economics 46, 157–182. doi:10.1111/1467-6451.00067.

Cohen, W.M., Klepper, S., 1996. Firm size and the nature of innovation within industries: The case of process and product R&D. The Review of Economics and Statistics 78, 232–243.

Cohen, W.M., Levinthal, D.A., 1989. Innovation and learning: The two faces of R&D. The Economic Journal 99, 569–596. doi:10.2307/2233763.

Cohen, W.M., Levinthal, D.A., 1990. Absorptive capacity: A new perspective on learning and innovation. Administrative Science Quarterly 35, 128–152. doi:10.2307/2393553.

Corrado, C., Hulten, C., Sichel, D., 2009. Intangible capital and u.s. economic growth. Review of Income and Wealth 55, 661–685.

Crouzet, N., Eberly, J., 2021. Intangibles, markups, and the measurement of productivity growth. Journal of Monetary Economics 124, S92–S109. doi:10.1016/j.jmoneco.2021.09.005.

Crouzet, N., Eberly, J.C., 2019. Understanding weak capital investment: The role of market concentration and intangibles. NBER Working Paper w25869 .

Crouzet, N., Eberly, J.C., 2023. Rents and intangible capital: A q+ framework. Journal of Finance 78, 1873–1916.

Crouzet, N., Eberly, J.C., Eisfeldt, A.L., Papanikolaou, D., 2022. The economics of intangible capital. Journal of Economic Perspectives 36, 29–52.

Cullen, Z., Li, D., Li, X., 2026. Labor as capital: AI and the ownership of expertise. Working paper.

CVS Health, 2026a. Aetna reduces claims processing time by more than 20% with AI to improve care experience. Press release, May 26. https://www.cvshealth.com/news/innovation/aetna-reduces-claims-processing-time-by-more-than-20-percent-with-ai-to-improve-care-html.

CVS Health, 2026b. First quarter 2026 results. Form 8-K, Exhibit 99.1, U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/0000064803/000006480326000051/cvs_ex99x1q1-26.htm.

Czarnitzki, D., Fernández, G.P., Rammer, C., 2023. Artificial Intelligence and Firm-Level Productivity. Journal of Economic Behavior & Organization 211, 188–205.

Davis, S.J., Haltiwanger, J., Jarmin, R., Miranda, J., Foote, C., Nagypal, E., 2006. Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. NBER Macroeconomics Annual 21, 107–179.

Davis, S.J., Haltiwanger, J.C., Schuh, S., 1996. Job Creation and Destruction. MIT Press, Cambridge, MA.

De Loecker, J., Eeckhout, J., Unger, G., 2020. The rise of market power and the macroeconomic implications. Quarterly Journal of Economics 135, 561–644.

Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F., Lakhani, K.R., 2023. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Technical Report Working Paper 24-013. Harvard Business School. Working paper.

Demirer, M., Musolff, L., Yang, L., 2026. Writing code vs. shipping code: Productivity effects across generations of ai coding tools. NBER Working Paper 35275 .

Dessein, W., Prat, A., 2022. Organizational capital, corporate leadership, and firm dynamics. Journal of Political Economy 130, 1477–1539. doi:10.1086/718985.

Dingel, J.I., Neiman, B., 2020. How many jobs can be done at home? Journal of Public Economics 189, 104235.

Eisfeldt, A.L., Hartman-Glaser, B., Kim, E.T., Lee, K.B., 2025a. Intangible intensity. First draft.

Eisfeldt, A.L., Papanikolaou, D., 2013a. Organization capital and the cross-section of expected returns. The Journal of Finance 68, 1365–1406.

Eisfeldt, A.L., Papanikolaou, D., 2013b. Organization capital and the cross-section of expected returns. Journal of Finance 68, 1365–1406.

Eisfeldt, A.L., Schubert, G., Zhang, M.B., 2025b. Generative AI and firm values. Journal of Finance forthcoming.

Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2023. GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv:2303.10130 .

Fahlenbrach, R., Sanz, L., Stulz, R.M., 2026. Organization Capital, Large Startups, and the Dearth of IPOs. ECGI Finance Working Paper 1140/2026. European Corporate Governance Institute. Dice Center WP 2026-10; Fisher College of Business WP 2026-03-010.

Falato, A., Kadyrzhanova, D., Sim, J.W., Steri, R., 2022. Rising intangible capital, shrinking debt capacity, and the u.s. corporate savings glut. Journal of Finance 77, 2799–2852.

Federal Trade Commission, 2025. Ftc issues staff report on ai partnerships & investments study. Press release. Discusses partnerships between Alphabet, Amazon, Microsoft and major generative AI developers.

Felten, E.W., Raj, M., Seamans, R., 2018. A method to link advances in artificial intelligence to occupational abilities. AEA Papers and Proceedings 108, 1–4.

Fortune, 2026. How Expedia's CTO is using AI to transform work for 17,000 employees—

and travel for millions. Fortune, January 14. https://www.fortune.com/2026/01/14/how-expedias-cto-is-using-ai-to-transform-work-for-17000-employees-and-travel-for-mil

Foster, L., Haltiwanger, J., Syverson, C., 2016. The slow growth of new plants: Learning about demand? Economica 83, 91–129.

Gofman, M., Jin, Z., 2024. Artificial intelligence, education, and entrepreneurship. The Journal of Finance 79, 631–667.

Goldfarb, A., Taska, B., Teodoridis, F., 2023. Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. Research Policy 52, 104653.

Google Cloud, 2025. Google Cloud and Verizon drive customer experience improvements for Verizon customers with Gemini integration. Press release, April 9. https://www.googlecloudpresscorner.com/2025-04-09-Google-Cloud-and-Verizon-Drive-Customer-Experience-Improvements-for-Verizo

Grant, R.M., 1996. Toward a knowledge-based theory of the firm. Strategic Management Journal 17, 109–122.

Guardian, 2018. Revealed: Google's 'two-tier' workforce training document.

Hampole, M., Papanikolaou, D., Schmidt, L.D., Seegmiller, B., 2025. Artificial intelligence and the labor market. Technical Report. National Bureau of Economic Research.

Hassan, T.A., Hollander, S., van Lent, L., Schwedeler, M., Tahoun, A., 2023. Firm-Level Exposure to Epidemic Diseases: COVID-19, SARS, and H1N1. The Review of Financial Studies 36, 4919–4964.

He, B., Mostrom, L.I., Sufi, A., 2024. Investing in customer capital. Technical Report. National Bureau of Economic Research.

Henderson, R.M., Clark, K.B., 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. Administrative Science Quarterly 35, 9–30. doi:10.2307/2393549.

Henderson, R.M., Cockburn, I.M., 1994. Measuring competence? exploring firm effects in pharmaceutical research. Strategic Management Journal 15, 63–84. doi:10.1002/smj.4250150906.

High, P., 2026. How Expedia is reinventing travel through AI and agentic design. Forbes, May 29. https://www.forbes.com/sites/peterhigh/2026/05/29/how-expedia-is-reinventing-travel-through-ai-and-agentic-design/.

Hosseini Maasoum, S.M., Lichtinger, G., 2025. Generative ai as seniority-biased technological change: Evidence from u.s. resume and job posting data. SSRN Working Paper.

Humlum, A., Vestergaard, E., 2025. Large Language Models, Small Labor Market Effects. NBER Working Paper 33777. National Bureau of Economic Research.

Imas, A., 2025. What is the impact of ai on productivity? https://aleximas.substack.com/p/what-is-the-impact-of-ai-on-productivity. Substack post.

Imrohoroğlu, A., Tüzel, , 2014. Firm-Level Productivity, Risk, and Return. Management

Science 60, 2073–2090.

Iscenko, Z., Strand, S., Chen, Y., Aimard, G., Codreanu, M., Sampathkumar, V., Imas, A., Jacobs, J., Muiruri, E., Mateos-Garcia, J., Ng, J.J., Javed, S., Martin, J., Ajmeri, O., Calin, D., Kim, A., Curto Millet, F., Manyika, J., 2026. Google's ai & economy atlas v1.0: Mapping gemini usage in the economy. Working Paper .

Jiang, W., Park, J., Xiao, R.J., Zhang, S., 2025. AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents. NBER Working Paper No. 33536.

Jones, C.I., Tonetti, C., 2020. Nonrivalry and the economics of data. American Economic Review 110, 2819–58.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., Tahoun, A., 2025. The diffusion of new technologies. Quarterly Journal of Economics 140, 1299–1365.

Kim, H., Koning, R., 2026. Ai-native firms. Working Paper .

Kitishian, D., 2025. Nike's AI strategy: Analysis of dominance in athletic apparel, footwear. Klover.AI research report. URL: https://www.klover.ai/nike-ai-strategy-analysis-of-dominance-in-athletic-apparel-footwear/. secondary source; accessed August 10, 2026.

Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N., 2017. Technological innovation, resource allocation, and growth. Quarterly Journal of Economics 132, 665–712.

Krupa, T.J., Mullaney, M.S., 2026. The use of artificial intelligence in decision-making: Evidence from the effectiveness of corporate tax strategies. Review of Accounting Studies 31, 704–744.

Lev, B., Radhakrishnan, S., 2005. The valuation of organization capital, in: Corrado, C., Haltiwanger, J., Sichel, D. (Eds.), Measuring Capital in the New Economy. University of Chicago Press, Chicago, IL, pp. 73–110.

March, J.G., 1991. Exploration and exploitation in organizational learning. Organization Science 2, 71–87.

McElheran, K., Li, J.F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., Zolas, N., 2024. AI Adoption in America: Who, What, and Where. Journal of Economics & Management Strategy 33 (2): 375–415.

McElheran, K., Yang, M.J., Brynjolfsson, E., Kroff, Z., 2025. The rise of industrial ai in america: Microfoundations of the productivity j-curve(s). SSRN Working Paper No. 5036270.

McLemore, P., Mihov, A., 2026. AI and operational losses: Evidence from U.S. bank holding companies. The Review of Corporate Finance Studies , cfag003.

MIT Sloan Management Review, . How Procter & Gamble uses AI to unlock new insights from data. MIT Sloan Management Review. https://sloanreview.mit.edu/article/how-procter-gamble-uses-ai-to-unlock-new-insights-from-data/.

Mitsubishi UFJ Financial Group, 2024. Who we are. https://www.tradepayables.com/mufg/who-we-are.html. Accessed: 2026-06-10.

New York Times, 2019. Google's shadow work force: Temps who outnumber full-time

employees.

NIKE, Inc., 2024. Creating the unreal: How Nike made its wildest Air footwear yet. Nike Newsroom. URL: https://about.nike.com/en/stories/nike-design-athlete-imagined-revolution. accessed August 10, 2026.

Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. Organization Science 5, 14–37.

Noy, S., Zhang, W., 2023. Experimental evidence on the productivity effects of generative artificial intelligence. Science 381, 187–192.

OECD, Eurostat, 2018. Oslo Manual 2018: Guidelines for Collecting, Reporting and Using Data on Innovation. OECD Publishing. doi:10.1787/9789264304604-en.

Olavsrud, T., 2026. The gen AI helping Aetna review millions of medical records. CIO, July 31. https://www.cio.com/article/4202752/how-gen-ai-helps-aetna-review-millions-of-medical-records.html.

Olley, G.S., Pakes, A., 1996. The Dynamics of Productivity in the Telecommunications Equipment Industry. Econometrica 64, 1263–1297.

OpenAI, 2024. Morgan stanley uses AI evals to shape the future of financial services. Press release, June 26. https://openai.com/index/morgan-stanley/.

Peters, R., Taylor, L.A., 2017. Intangible capital and the investment-q relation. Journal of Financial Economics 123, 251–272.

Polanyi, M., 1966. The Tacit Dimension. Routledge and Kegan Paul, London.

Prescott, E.C., Visscher, M., 1980. Organization capital. Journal of Political Economy 88, 446–461.

Roche Holding AG, 2024. Annual Report 2024. https://assets.roche.com/f/174029/x/74e4b5f7dc/roche-annual-report-2024.pdf. Accessed: June 13, 2026.

Rock, D., 2019. Engineering value: The returns to technological talent and investments in artificial intelligence. Working paper .

Schubert, G., 2025. Organizational technology ladders: Remote work and generative ai adoption. Working Paper .

Sheng, J., Sun, Z., Yang, B., Zhang, A.L., 2026. Generative AI and Asset Management. Review of Financial Studies, advance article.

Singhal, K., Azizi, S., Tu, T., Mahdavi, S.S., Wei, J., Chung, H.W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., Payne, P., Seneviratne, M., Gamble, P., Kelly, C., Babiker, A., Schärli, N., Chowdhery, A., Mansfield, P., Demner-Fushman, D., Agüera y Arcas, B., Webster, D., Corrado, G.S., Matias, Y., Chou, K., Gottweis, J., Tomasev, N., Liu, Y., Rajkomar, A., Barral, J., Semturs, C., Karthikesalingam, A., Natarajan, V., 2023. Large language models encode clinical knowledge. Nature 620, 172–180.

Skift, 2026. Expedia's AI timeline runs through its talent strategy. Skift, July 20. https://skift.com/2026/07/20/expedia-cto-ramana-thumu-ai-talent-strategy/.

Sohel, 2026. How Nike is using AI: Strategy, tools, real examples. Dilogs.ai Blog. URL: https://www.dilogs.ai/blog/how-nike-is-using-ai/. secondary source; accessed

August 10, 2026.

Son, H., 2026. Morgan stanley will soon open its trillion-dollar wealth management funnel to AI agents. CNBC, June 3. https://www.cnbc.com/2026/06/03/ai-agents-morgan-stanley-wealth-management-funnel.html.

Team DigitalDefynd, 2026. 10 ways Nike is using AI [case study]. DigitalDefynd. URL: https://digitaldefynd.com/IQ/ways-nike-use-ai/. secondary source; accessed August 10, 2026.

The Home Depot, 2025a. The Home Depot introduces Magic Apron, a suite of advanced AI tools to help customers with their home improvement needs. Press release, March 6. URL: https://ir.homedepot.com/news-releases/2025/03-06-2025-130241718.

The Home Depot, 2025b. Unveiling Magic Apron: The Home Depot's smartest tool yet. Corporate newsroom, March 6. https://corporate.homedepot.com/news/company/unveiling-magic-apron-home-depots-smartest-tool-yet.

The Home Depot, 2026a. AI at The Home Depot: A look at the future of home improvement. Corporate newsroom, March 30. https://corporate.homedepot.com/news/company/ai-home-depot-look-future-home-improvement.

The Home Depot, 2026b. The Home Depot and Google Cloud launch agentic AI tools to help customers and associates bring projects from 'how-to' to 'done'. Press release, January 11. https://corporate.homedepot.com/news/partnerships/home-depot-and-google-cloud-launch-agentic-ai-tools-help-customers-and-associates.

Turing, A.M., 1950. Computing machinery and intelligence. Mind 59, 433–460.

U.S. Bureau of Labor Statistics, 2025. Overview of may 2024 occupational employment and wages. Occupational Employment and Wage Statistics overview.

Veldkamp, L., Chung, C., 2024. Data and the aggregate economy. Journal of Economic Literature 62, 458–484.

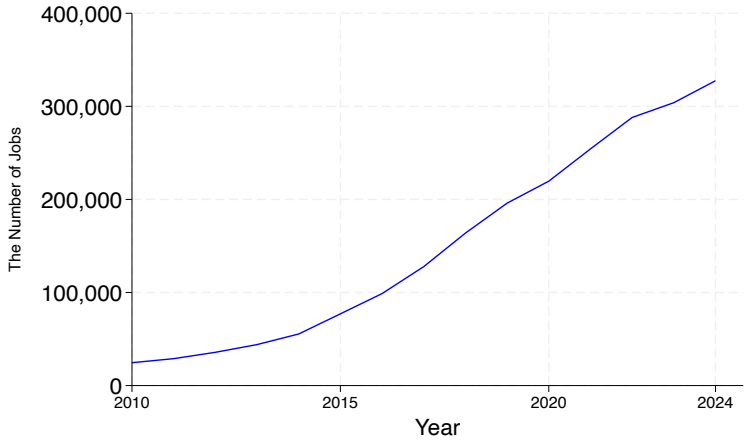
Webb, M., 2020. The impact of artificial intelligence on the labor market. Working Paper .

World Economic Forum, 2021. Factories of the future find growth and sustainability through digitalization. Press release, March 16. https://www.weforum.org/press/2021/03/factories-of-the-future-find-growth-and-sustainability-through-digitalization/.

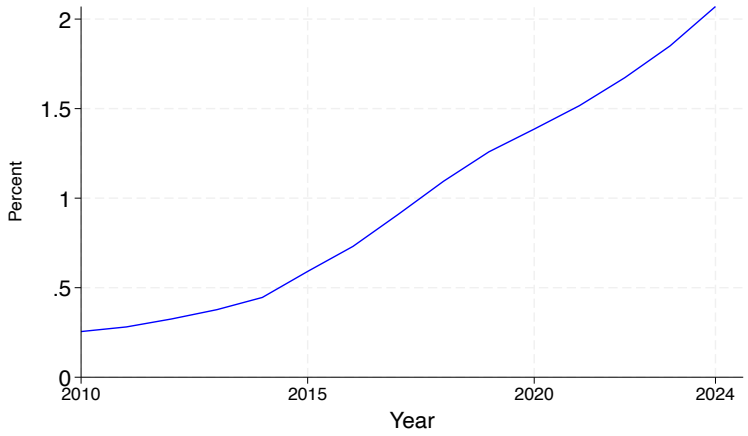
World Economic Forum, 2023. Against economic headwinds, 18 manufacturing lighthouses show how to boost productivity and sustainability by scaling advanced technologies across networks. Press release, January 10. https://www.weforum.org/press/2023/01/against-economic-headwinds-18-manufacturing-lighthouses-show-how-to-boost-productivit

Yotzov, I., Barrero, J.M., Bloom, N., Bunn, P., Davis, S.J., Foster, K.M., Jalca, A., Meyer, B.H., Mizen, P., Navarrete, M.A., Smietanka, P., Thwaites, G., Wang, B.Z., 2026. Firm Data on AI. NBER Working Paper 34836. National Bureau of Economic Research.

Figure 1. Time Series of AI Jobs



(a) Number of AI Jobs

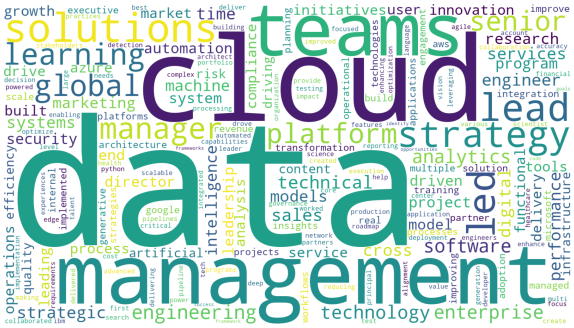


(b) Share of AI Jobs

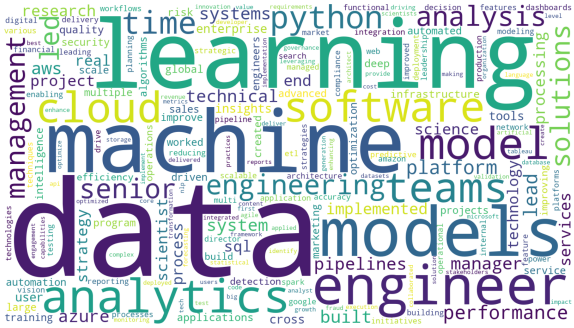
This figure plots annual AI-related U.S. jobs in our matched public-firm sample from 2010 to 2024. Panel (a) reports the total number of AI-related U.S. jobs. Panel (b) reports the aggregate AI job share, defined as AI-related U.S. jobs divided by total U.S. jobs in the same year. AI-related jobs are identified using keyword matches in job titles or job descriptions.

Figure 2. Word Clouds for AI subcategories

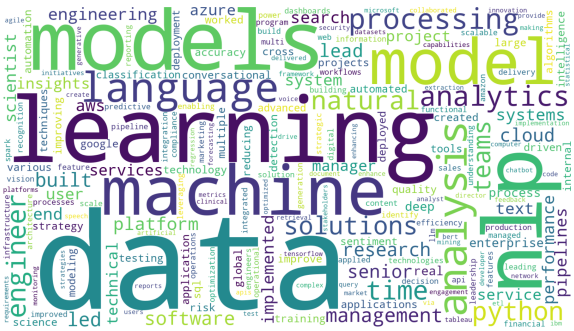
(a) AI



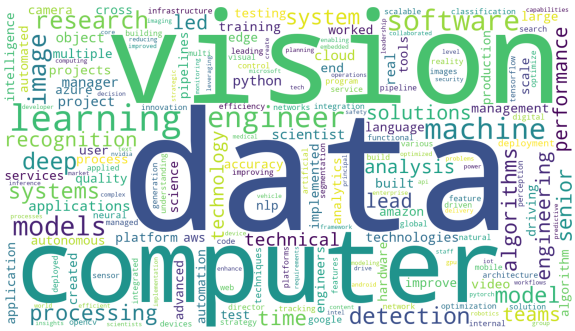
(b) ML



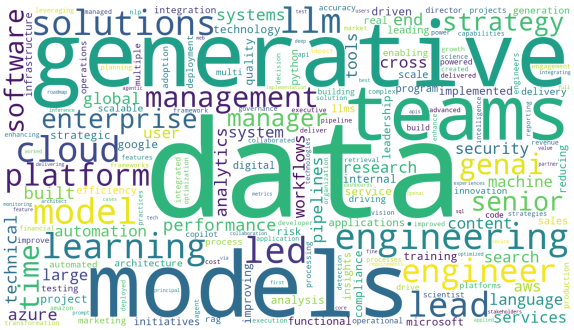
(c) NLP



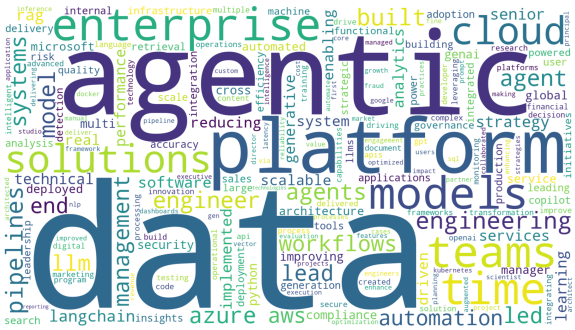
(d) CV



(e) Gen AI

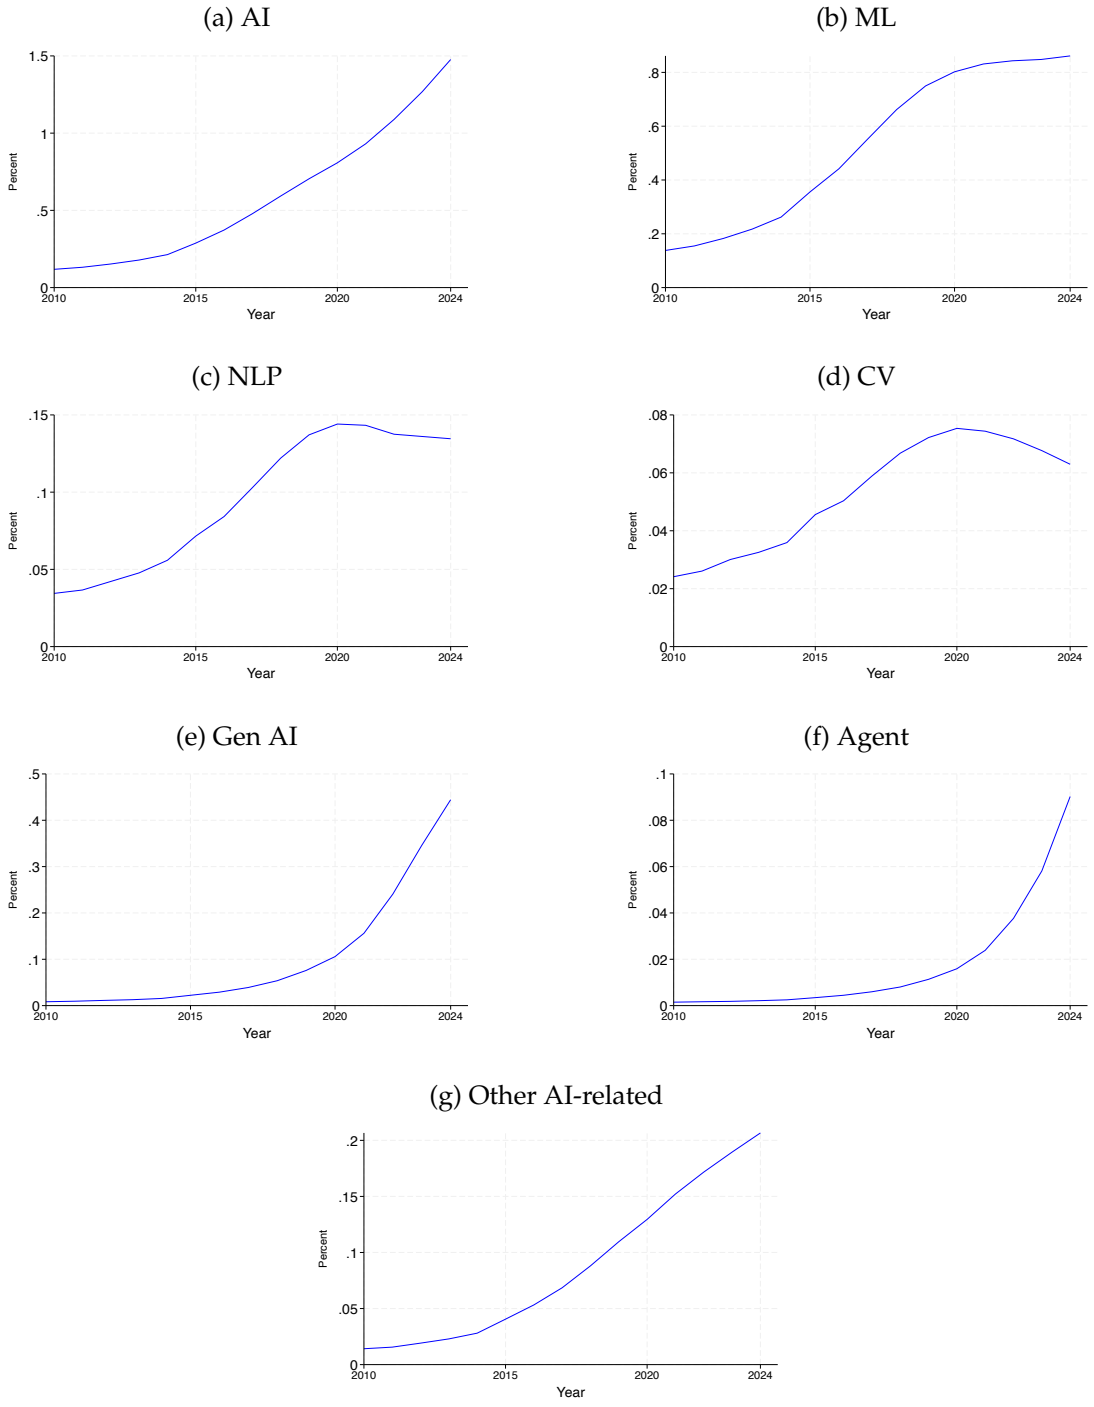


(f) Agentic AI



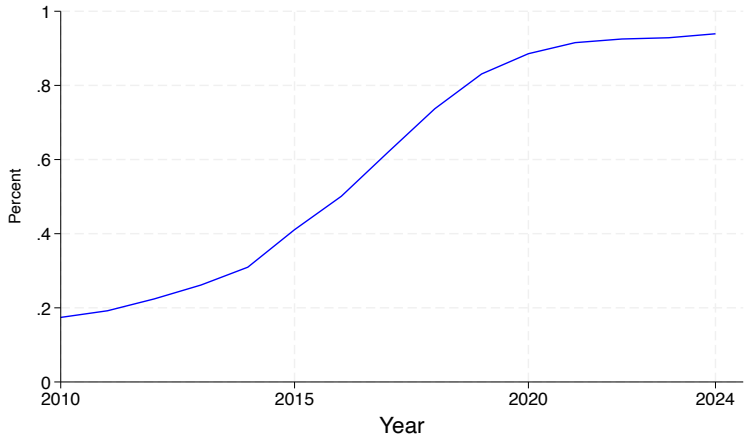
This figure plots word clouds for six AI technology categories: general AI, machine learning (ML), natural language processing (NLP), computer vision (CV), generative AI, and agentic AI. The word clouds are constructed from the job titles and job descriptions of U.S. jobs in our matched U.S. public-firm sample. For each category, words are extracted from positions classified as related to that technology based on keyword matches in either the job title or the job description. The terms "AI" and "artificial intelligence" are excluded from the word clouds, as they appear in the vast majority of jobs in every category and would otherwise dominate all panels. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure 3. Share of AI Jobs by Category

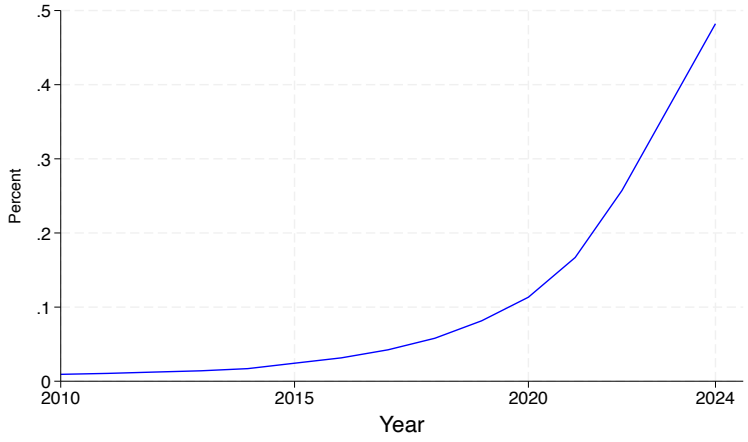


This figure plots annual AI-related U.S. job shares by AI technology category in our matched public-firm sample from 2010 to 2024. Each panel reports the aggregate job share for one category, defined as AI-related U.S. jobs in that category divided by total U.S. jobs in the same year. AI-related jobs are identified using keyword matches in job titles or job descriptions. Categories are not mutually exclusive, so a job can appear in multiple panels if it matches keywords from multiple AI technology categories. The "Other AI-related" category includes AI-related jobs that are not assigned to the six named technology categories.

Figure 4. Share of AI Jobs: Wave I vs. Wave II Technologies



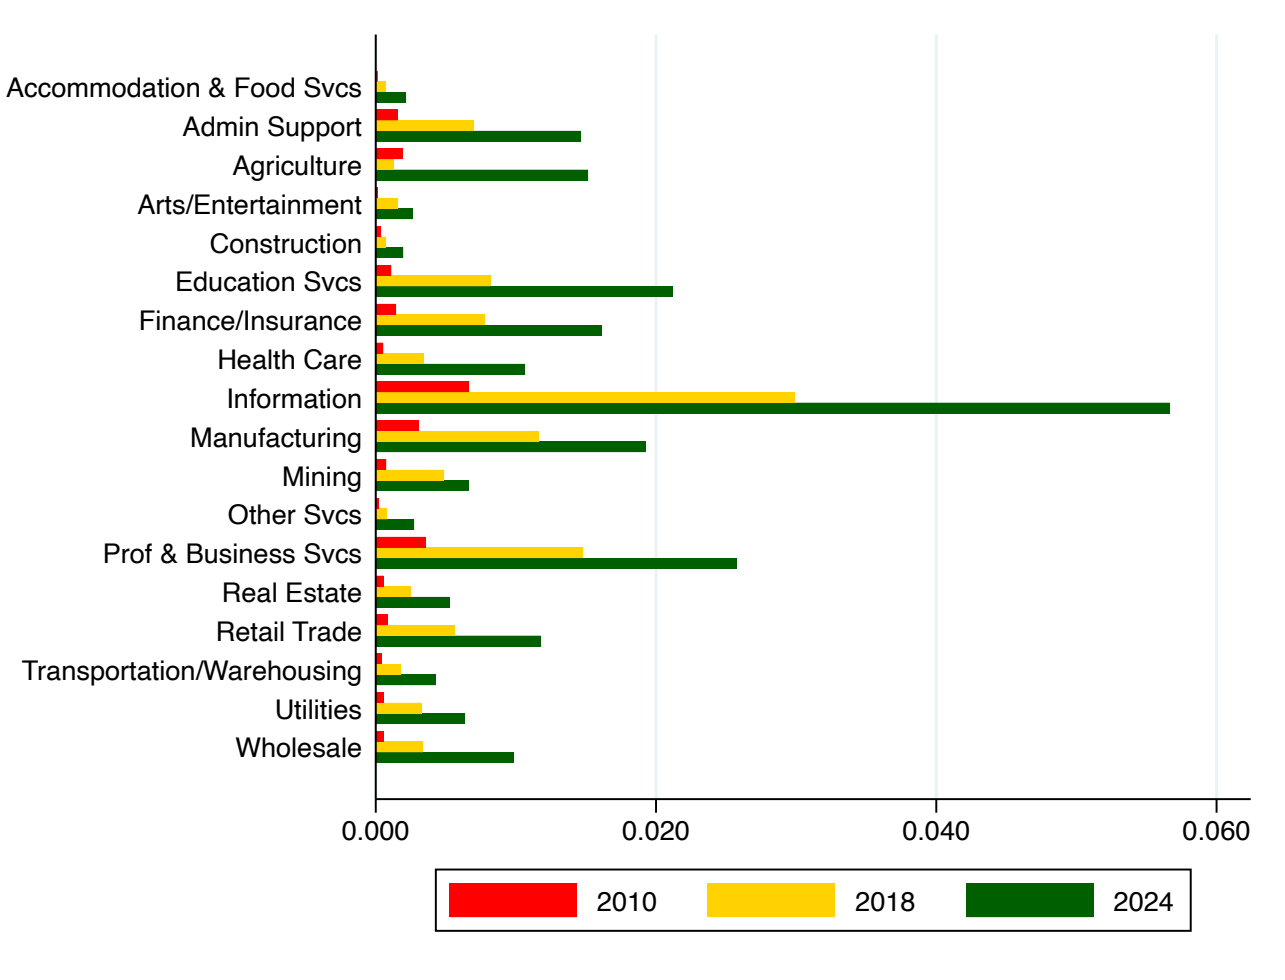
(a) Wave I



(b) Wave II

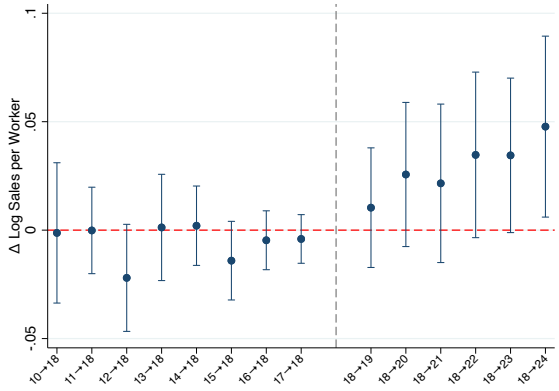
This figure plots the share of AI-related jobs among U.S. public firms from 2010 to 2024, separately for Wave I and Wave II AI technologies. Panel (a) reports the share of "Wave I" AI jobs, defined as jobs containing at least one of the NLP, ML, or CV keywords. Panel (b) reports the share of "Wave II" AI jobs, defined as jobs containing at least one of the Gen AI or agentic AI keywords.

Figure 5. Share of AI Jobs by Industry Sector

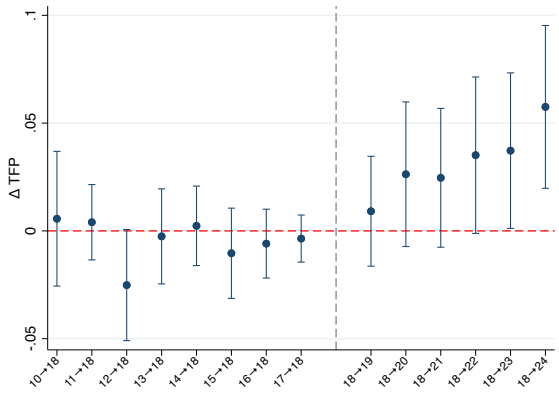


This figure presents the share of AI jobs at the industry level, based on the sample of U.S. public firms. For each sector (based on NAICS-2 digit industry codes), we compute the fraction of AI-related employees in the resume data. The statistics are computed across all public firms in each sector in three years: 2010, 2018, and 2024. In each of these years, we sum AI jobs within each industry and divide it by total jobs in that industry.

Figure 6. AI Investments and Productivity Changes Over Time



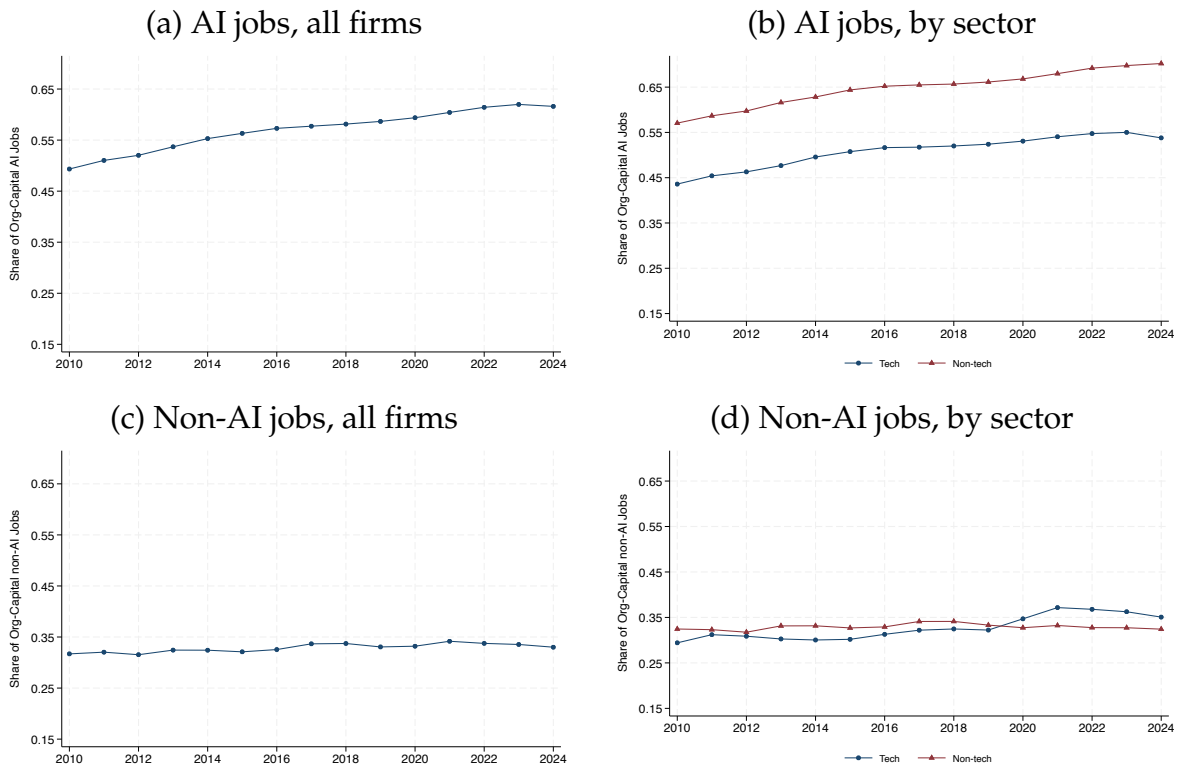
(a) Log Sales per Worker



(b) TFP

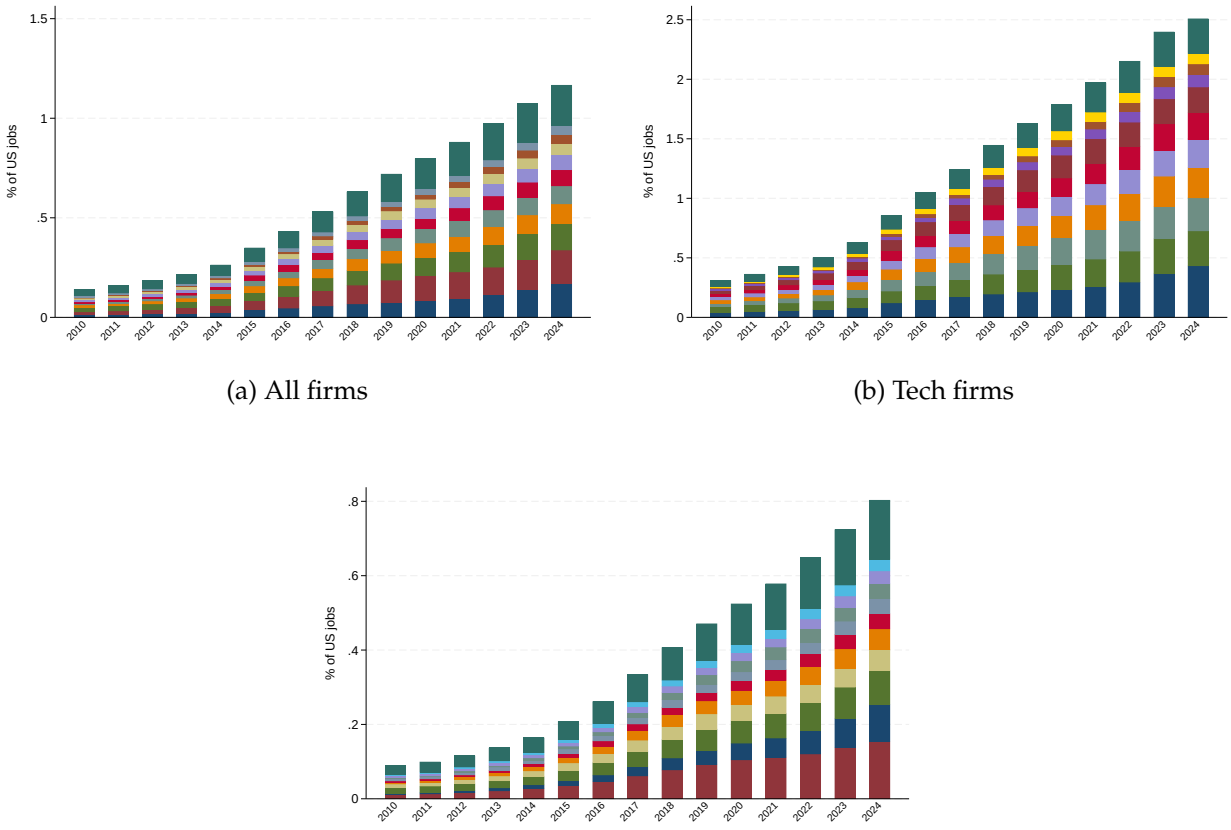
This figure plots the cross-sectional association between AI adoption intensity and productivity growth at different time points. The independent variable is the standardized change in a firm's share of AI jobs from 2018 to 2024, ΔShare $\text{AI}_i^{18\to24}$, which is the same independent variable in all regressions. Each point plots the coefficient $\hat{\beta}$ from a separate cross-sectional regression of $\Delta Y_i^{t\to t'}$ on ΔShare AI $\text{Jobs}_i^{18\to24}$, controlling for two-digit NAICS sector fixed effects and weighted by firm-level U.S. employees. Panel (a) reports results for log sales per worker, and Panel (b) is for total factor productivity. To the left of the dashed line, the dependent variable is $\Delta Y_i^{T\to2018}$ for $T = 2010, \ldots, 2017$. For example, the very first dot measures how changes in firm AI investments from 2018 to 2024 relate to ex-ante change in productivity from 2010 to 2018. The next dot shortens the period of when ex-ante productivity is measured by one year (from 2011 to 2018), with the last dot on the left using changes in productivity from 2017 to 2018. To the right of the dashed line, the dependent variable is $\Delta Y_i^{2018\to T}$ for $T = 2019, \ldots, 2024$: the very first dot after dashed line measures how changes in firm AI investments from 2018 to 2024 relate to change in productivity from 2018 to 2019, with the last dot using changes in productivity from 2018 to 2024. The standard errors are clustered at the five-digit NAICS level.

Figure 7. Share of Organization-Capital-Related Jobs Among AI and Non-AI Jobs, 2018–2024



This figure plots the share of AI jobs classified as building organization capital (organization-capital score ≥ 0.6), 2010–2024 in our matched public firm sample. Panel (a) plots the share of AI jobs that are organization-capital related. Panel (b) plots the share of AI jobs that are organization-capital related separately for tech firms (NAICS 51, 54, and semiconductor manufacturing) and all other firms. Panel (c) plots the share of non-AI jobs that are classified as building organization capital. Panel (d) plots the share of non-AI jobs that are organization-capital-building separately for tech and non-tech firms.

Figure 8. Organization-Capital Reason Categories for AI Jobs, 2010–2024



(a) All firms

(b) Tech firms

(c) Non-tech firms

This figure plots the share of organization-capital-related AI jobs in each reason category listed in Table 2 from 2010 to 2024 in our matched public firm sample. Panel (a) plots the shares for all firms. Panel (b) plots the shares for tech firms (NAICS 51, 54, and semiconductor manufacturing). Panel (c) plots the shares for non-tech firms. In each panel, we show the largest ten categories for that sample and group all remaining categories into "other". Colors denote the same category in every panel.

Table 1. Summary Statistics of Share of AI Jobs

		Mean (%)			SD (%)	% >5 AI	Max (%)	Percentiles of All AI Share (%)						
Year	N	All AI	Wave I	WaveII	All AI	Jobs	All AI	P1	P5	P25	P50	P75	P90	P99
2010	4,046	0.20	0.12	0.00	1.37	9.1	50.00	0.00	0.00	0.00	0.00	0.05	0.35	3.30
2011	4,038	0.21	0.13	0.01	1.06	10.6	33.33	0.00	0.00	0.00	0.00	0.07	0.39	3.78
2012	4,085	0.26	0.18	0.01	1.62	12.1	50.00	0.00	0.00	0.00	0.00	0.09	0.46	4.35
2013	4,271	0.31	0.22	0.01	2.07	13.8	100.00	0.00	0.00	0.00	0.00	0.12	0.55	5.00
2014	4,335	0.34	0.22	0.01	1.63	16.2	50.00	0.00	0.00	0.00	0.00	0.16	0.67	4.97
2015	4,254	0.40	0.27	0.01	2.16	18.2	100.00	0.00	0.00	0.00	0.00	0.22	0.85	5.93
2016	4,208	0.48	0.31	0.01	2.71	20.9	100.00	0.00	0.00	0.00	0.00	0.28	0.96	5.88
2017	4,153	0.59	0.39	0.02	2.93	24.2	100.00	0.00	0.00	0.00	0.00	0.39	1.23	7.51
2018	4,156	0.77	0.51	0.03	3.96	27.1	100.00	0.00	0.00	0.00	0.05	0.49	1.50	9.09
2019	4,100	0.81	0.51	0.04	3.29	29.3	100.00	0.00	0.00	0.00	0.10	0.60	1.75	9.70
2020	4,364	0.96	0.56	0.07	3.22	30.2	60.09	0.00	0.00	0.00	0.11	0.69	2.13	13.79
2021	4,589	1.09	0.58	0.09	3.30	31.6	54.79	0.00	0.00	0.00	0.15	0.81	2.60	15.85
2022	4,487	1.18	0.55	0.12	3.58	34.1	100.00	0.00	0.00	0.00	0.20	0.94	2.81	15.58
2023	4,263	1.31	0.55	0.19	3.75	36.1	100.00	0.00	0.00	0.00	0.27	1.08	3.15	16.79
2024	4,022	1.46	0.57	0.25	3.71	38.6	71.69	0.00	0.00	0.00	0.33	1.22	3.64	18.18
All	84,373	0.56	0.31	0.04	2.57	19.3	100.00	0.00	0.00	0.00	0.00	0.27	1.14	9.00

This table reports annual cross-sectional summary statistics for firm-level AI-related U.S. job shares from 2010 to 2024 for our matched public-firm sample. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm and is expressed in percent. AI-related jobs are identified using keyword matches in job titles or job descriptions. Wave I AI includes machine learning, natural language processing, and computer vision keywords; Wave II AI includes generative AI and agentic AI keywords. Wave I and Wave II categories are not mutually exclusive. $\% > 5$ AI jobs reports the share of firms with more than five AI-related U.S. jobs. Percentiles are computed from the cross-sectional distribution of All AI share. The All row pools all firm-years. Firm-years are restricted to those with positive U.S. employment in Revelio and a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio jobs divided by Compustat employment for each firm-year.

Table 2. Reason categories: distribution of AI versus non-AI jobs

Panel A: Builds organization capital

Reason category	AI jobs %	Non-AI jobs %
Reusable data platforms & pipelines	9.1	2.2
Reusable AI/ML platforms & MLOps	4.8	0.1
Firm-specific predictive & decision models	2.3	0.2
Proprietary-knowledge GenAI/NLP systems	6.9	0.4
Reusable software platforms & workflows	2.2	1.8
Cross-functional operating routines	6.3	4.8
Reusable cloud, infrastructure & reliability operations	3.4	2.0
Security, compliance & risk systems	1.9	1.4
Marketing & revenue-operations systems	4.1	1.9
Customer-service & support systems	2.0	0.8
Product & design systems	4.7	2.0
Manufacturing & process systems	1.6	3.5
Supply-chain & logistics systems	1.5	2.1
R&D & scientific platforms	1.4	0.5
Financial & risk decision systems	3.0	1.3
Hiring, training & talent development	1.4	5.8
Alliance & partnership systems	0.5	0.2
Knowledge organization & ontologies	0.4	0.0
Subtotal, builds org capital	57.5	31.0

Panel B: Does not build organization capital

Reason category	AI jobs %	Non-AI jobs %
Bespoke client consulting/implementation	7.8	8.8
Portable analytics & one-off models	12.6	4.1
Broadly applicable research & algorithms	8.9	1.3
Customer-facing products & IP	3.8	3.9
Standardized/portable software work	3.3	2.4
Transactional sales, marketing & account management	2.9	7.6
Individual recruiting & hiring transactions	0.9	0.7
Standardized routine operations, support, maintenance	0.7	37.9
Project/program management (episodic)	0.3	0.5
Training & general education	0.5	0.7
Portable legal & contracting	0.8	1.1
Subtotal, does not build	42.5	69.0

This table reports the distribution of organization-capital reason categories for AI-related U.S. jobs and non-AI U.S. jobs in the public-firm sample. AI-related jobs are identified using keyword matches in job titles or job descriptions. Each job is scored once and assigned to one standardized reason category. Panel A reports categories classified as building organization capital, corresponding to scores of at least 0.6. Panel B reports categories classified as not building organization capital. The non-AI sample consists of 1,000 randomly drawn non-AI U.S. jobs per year and is scored using the same taxonomy as the AI-related jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table 3. AI Investments and Productivity

Panel A: 2010–2018

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	-0.004	0.002	0.014	0.012
	(0.014)	(0.015)	(0.012)	(0.012)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.089	0.171	0.047	0.150
Observations	1,222	1,222	1,091	1,091

Panel B: 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.048**	0.062**	0.061***	0.075***
	(0.021)	(0.026)	(0.020)	(0.024)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.178	0.226	0.071	0.153
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2010 to 2018 in Panel A and from 2018 to 2024 in Panel B. Regressions are weighted by the number of workers in Revelio at the start of the period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period (2010 in Panel A and 2018 in Panel B): firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 4. Investments in Different Waves of AI Technologies and Productivity, 2018–2024

	Δ Log Sales per Worker						Δ TFP					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Δ Share Wave I Only AI Jobs	0.057**	0.065**					0.062**	0.069**				
	(0.024)	(0.028)					(0.024)	(0.027)				
Δ Share Wave II Only AI Jobs			0.034*	0.040**					0.058***	0.067***		
			(0.017)	(0.018)					(0.014)	(0.018)		
Δ Share Wave I & II AI Jobs					0.040**	0.056**					0.052***	0.070***
					(0.018)	(0.024)					(0.018)	(0.024)
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.184	0.230	0.161	0.203	0.181	0.230	0.068	0.145	0.051	0.127	0.078	0.165
Observations	1,321	1,321	1,321	1,321	1,321	1,321	1,173	1,173	1,173	1,173	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in share of AI jobs related to different waves of AI technologies. In columns 1, 2, 7, and 8, the independent variable is the change in the share of workers with only Wave I AI skills (Machine Learning, Natural Language Processing, and Computer Vision). In columns 3, 4, 9, and 10, the independent variable is the change in the share of workers with only Wave II AI skills (Generative AI and Agentic AI). In columns 5, 6, 11, and 12, the independent variable is the change in the share of workers with both waves of AI skills. In all columns, the independent variable is standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for detailed industry fixed effects and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5. AI Investments and Productivity: IV Estimates

	Δ Log Sales per Worker			Δ TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	0.160	0.158**	0.144***	0.135*	0.110**	0.114***
	(0.098)	(0.062)	(0.048)	(0.076)	(0.048)	(0.041)
Industry FE	Y	Y	Y	Y	Y	Y
CS Control	Y	Y	Y	Y	Y	Y
Baseline Controls	Y	Y	Y	Y	Y	Y
State FE	N	Y	Y	N	Y	Y
Control Pre-trend	N	N	Y	N	N	Y
F Statistic	10.1	17.9	26.2	10.5	18.8	26.0
Observations	633	627	531	576	570	528

This table estimates the relationship between AI investments and firm productivity growth from 2018 to 2024 for U.S. public firms (in non-tech sectors), where firm AI investments are instrumented with ex-ante firm-level exposure to AI-skilled graduates from AI-strong universities. The independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs) from 2018 to 2024, standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in 2018 (see the definition of the instrument in Section 3.4 and the details of instrument construction in Appendix C). We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). All specifications control for industry sector fixed effects and ex-ante exposure to universities that are strong in computer science research as well as top 20 universities. All columns also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 2, 3, 5, and 6 add state fixed effects, and columns 3 and 6 additionally control for firm-level changes in log sales, log employment, and TFP from 2002 to 2010. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. The first-stage F-statistics of the instrument are reported for all specifications. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6. AI Investments and Organization Capital

Panel A: 2010–2018

	Δ Log(Organizational Capital/PP&E)		Δ Log(Organizational Capital/Total Assets)	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	-0.053	-0.042	0.034	-0.002
	(0.037)	(0.046)	(0.022)	(0.034)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.104	0.209	0.171	0.262
Observations	851	851	1,028	1,028

Panel B: 2018–2024

	Δ Log(Organizational Capital/PP&E)		Δ Log(Organizational Capital/Total Assets)	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.104***	0.089***	0.119***	0.110***
	(0.028)	(0.027)	(0.043)	(0.036)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.095	0.124	0.092	0.126
Observations	886	886	1,105	1,105

This table reports the coefficients from long-differences regressions of organization capital (capitalized SG&A) on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 and 2 and changes in the natural log of organization capital normalized by total assets in columns 3 and 4. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2010 to 2018 in Panel A and from 2018 to 2024 in Panel B. Regressions are weighted by the number of workers in Revelio at the start of the period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period (2010 in Panel A and 2018 in Panel B): firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7. Organization-Capital-Related AI Jobs and Productivity, 2018–2024

	Δ Log Sales per Worker			Δ TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.069***	0.085***	0.084***	0.081***	0.100***	0.094***
	(0.025)	(0.028)	(0.026)	(0.025)	(0.027)	(0.025)
Δ Share of Non-Organization-Related AI Jobs			0.002			0.014
			(0.021)			(0.022)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.188	0.238	0.237	0.078	0.167	0.167
Observations	1,306	1,306	1,306	1,163	1,163	1,163

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in organization capital and non-organization-capital-related AI jobs among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). The independent variables are growth in the share of organization-capital-related AI jobs (normalized by firms' total U.S.-based jobs) in columns 1, 2, 4 and 5, and the share of AI jobs not related to organization capital (normalized by firms' total U.S.-based jobs) in columns 3 and 6, which are both standardized to mean zero and standard deviation of one. A job is organization-capital-related if the organization-capital score (based on job description in the Revelio resume data) is at least 0.6. All changes in productivity and the shares of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2, 3, 5, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 8. Heterogeneous Relationship between AI Investments and Productivity by Initial Organization Capital, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Workers*Below-Median Organization Capital	0.070*	0.092**	0.088**	0.105***
	(0.043)	(0.037)	(0.039)	(0.036)
Δ Share AI Workers*Above-Median Organization Capital	0.034	0.050	0.027	0.052
	(0.038)	(0.044)	(0.033)	(0.039)
NAICS2*Size tercile*Above-median org capital FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.200	0.259	0.184	0.245
Observations	875	875	836	836

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), separately for firms with above-median and below-median initial level of organization capital. Firms within each 2-digit NAICS industry sector are split at the median level of organizational capital normalized by PP&E in 2018. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the AI employment share, standardized to mean zero and unit standard deviation, interacted with an indicator for above- versus below-median initial organization capital. All changes in productivity and share of AI jobs are measured from 2010 to 2018. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix

A Data and Measurement **66**

A.1 Linking Revelio Labs Firms to Compustat 66

A.2 Construction of firm-level AI measures 74

A.3 Construction of AI job categories 92

A.4 Organization Capital: Definitions, Comparison to Neighboring Concepts, and Classification 97

A.5 Measures of Productivity: Definitions and Construction 138

B Robustness and Additional Results **205**

C Instrument Construction and Validation **250**

D A Model of AI and Organization Capital **262**

D.1 Model Setup 263

D.2 AI as a Codification Technology 270

D.3 Event-Study Dynamics 272

D.4 Endogenous Adoption over the Life Cycle 275

D.5 Taking the Model to Data 277

D.6 Extension: Organization Capital as a Substitute for Measured Inputs 279

A. Data and Measurement

A.1. Linking Revelio Labs Firms to Compustat

This section describes how we link each firm in Revelio Labs to Compustat over time. The resulting mapping is many-to-one in each year, since multiple Revelio firm identifiers (`rcid`) can map to the same Compustat firm (`gvkey`).[37] We construct our own mapping rather than use Revelio's direct linkage so that it can vary over time and accommodate mergers, acquisitions, and changes in corporate structure that alter firm identifiers.[38]

The procedure has five stages: (i) four direct matching methods that exploit the firm identifiers Revelio reports for each `rcid`; (ii) extension of these matches from parent firms to subsidiaries using Revelio's dynamic parent–subsidiary panel; (iii) consolidation of all candidate matches; (iv) a three-step false-positive screen that combines algorithmic similarity scoring, a large-language-model audit, and human review; and (v) production of the final panel.

A.1.1. Data inputs

(i) **Compustat firm–year master:** For each Compustat firm-year, we extract the firm identifier (`gvkey`), name, headquarters address (ZIP, city, state), NAICS code, IPO date, and other identifiers (`cusip`, `cik`, and `ein`). We drop firm-year observations with missing or zero employment.

(ii) **Revelio firm mapping file:** For each firm in Revelio Labs, we observe Revelio's firm identifier (`rcid`), firm name, ultimate-parent firm identifier (`ultimate_parent_rcid`) and name, headquarters location, founding year, NAICS code, and firm identifiers linked by Revelio, including `gvkey`, `cusip`, `cik`, and `ein`.

(iii) **Revelio employment:** Annual count of U.S. and global employed workers associated with each firm in Revelio Labs.

(iv) **Subsidiary structure data:** Revelio's panel of parent-subsidiary relationships with effective start and end dates, allowing us to identify which `rcid` is owned by which ultimate parent in any given year.

(v) **Historical identifier crosswalks from WRDS:** We use three WRDS link tables, each indexed by effective start and end dates, to build the historical mappings we need.

[37] Multiple Revelio firms sharing one `gvkey` are subsidiaries of a corporate parent.

[38] Examples include spin-offs, divestitures, reorganization, and reassignments of subsidiaries to a different parent.

- `ccmxpf_lnkhist` – the CRSP–Compustat Merged link history, which gives the historical mapping between `gvkey` and CRSP `permno` with link start (`linkdt`) and end (`linkenddt`) dates.
- `msenames` – the CRSP stock event names file, which gives the historical mapping between CRSP `permno` and `cusip` with effective name start (`namedt`) and end (`nameendt`) dates.
- `wciklink_gvkey` – the WRDS SEC `cik` link table, which gives the historical mapping between `cik` and `gvkey` with link start (`link_start_date`) and end (`link_end_date`) dates.

A.1.2. Four direct matching methods

For each `rcid`, Revelio reports other commonly used firm identifiers, such as `gvkey`, `cusip`, `cik`, and `ein`. These identifiers are also available for firms in Compustat. For each `rcid` in Revelio, we attempt four matches in parallel to firms in Compustat using these identifiers from both datasets, each producing a candidate `gvkey`:

(i) **Direct** `gvkey` **match:** Revelio assigns a `gvkey` to a subset of `rcids` based on its own internal matching. We take this assignment as the first candidate.

(ii) **Historical CUSIP:** Using the Compustat-CRSP link, we construct a year-indexed panel of all `gvkey-cusip` pairs that are linked in each year. In the year where the matched cusip for a gvkey switches, we use the new cusip for that year. For each Revelio firm-year, we match the `rcid`'s `cusip` to this panel. If the same `cusip` is linked to a `gvkey` in year `t`, that `gvkey` becomes a candidate.

(iii) **Historical CIK:** The same logic as above is applied to the `gvkey-cik` panel from `wciklink_gvkey`.

(iv) **EIN:** We match the `rcid`'s `ein` directly to the Compustat firm's `ein`. Unlike CUSIP and CIK, there is no widely available historical EIN-to-gvkey crosswalk, so we use Compustat's current EIN. This is the least reliable of the four methods for that reason.

Each method labels its successful matches with a tag (`match_on` $\in$ {`gvkey`, `cusip`, `cik`, `ein`}). A single `rcid` may be matched by more than one method, and to a `gvkey` different from another method's candidate in this step. We reconcile these conflicts in Section A.1.4.

A.1.3. Matching subsidiaries through their parent

Many Revelio firms are corporate subsidiaries that do not carry the parent's identifiers, so Section A.1.2 misses them. For example, the top-level Alphabet `rcid` carries Alphabet's

cusip and cik, so Section A.1.2 matches it directly to Alphabet's gvkey. But some of Alphabet's subsidiary rcids, such as YouTube, DeepMind, and so on, carry no Compustat identifier of their own in Revelio's company mapping file. In addition, the company mapping file is static and does not provide information on the timing of the acquisitions. To pick up the subsidiaries dynamically, we use Revelio's parent–subsidiary panel together with Section A.1.2.

For every (rcid, year) cell, we look up the ultimate parent rcid active in that year. If the parent was matched in Section A.1 and the ultimate parent's rcid is different from the firm's own rcid, we assign the parent's gvkey to the subsidiary for the same year and flag the match as one made through the parent (subsidiary $= 1$).

A.1.4. Consolidation

We append all direct matches from the above steps. When a single (rcid, year) cell has multiple direct matches that disagree (this is less than 1% of cases), we keep the most reliable one in the order gvkey $>$ cusip $>$ cik $>$ ein.

We then match each rcid-year observation to the gvkey of the ultimate parent. We replace the gvkey that is matched directly with the gvkey of the ultimate parent after acquisition. We also replace the gvkey as missing if the rcid is matched to the ultimate parent's gvkey via direct matching before the acquisition.

The output of this stage is a panel of candidate matches at the (rcid, year) level: each row identifies a candidate gvkey, a match_on tag indicating which method produced the match, and a subsidiary indicator for whether the match flows through the parent company.

A.1.5. False-positive screen

Identifier-based matches can still be wrong. Two firms may share a historical cusip or cik after a corporate reorganization, or Revelio's ein for a subsidiary may belong to a holding company that is itself listed under a different name. To remove these false-positives, we apply three screens in sequence.

Step 1: Similarity scoring. For each candidate (rcid, gvkey) match, we compute several similarity indicators between the Revelio firm and the Compustat firm:

- *Industry agreement* (NAICS match): equal to one if Revelio's 4-digit NAICS code and Compustat's 4-digit NAICS code match.
- *Location agreement* (state, city, ZIP): three separate binary indicators that equal one if the Revelio firm's headquarters state matches Compustat's, if the headquarters city matches, and if the headquarters ZIP code matches, respectively.

- *Name similarity*: a normalized string-similarity score between Revelio's company name and Compustat's company name.
- *Size ratio*: Revelio's annual employment count divided by Compustat's annual employment headcount.

A non-subsidiary candidate match is flagged as a suspected false positive if (i) the 4-digit NAICS codes disagree, the state, city, and ZIP code of the headquarters all disagree, and the name-similarity score is below 0.9; or (ii) Revelio's annual employment count for the `rcid` is more than twice Compustat's reported headcount for the `gvkey`; or (iii) the founded year in Revelio is after the IPO date in Compustat. We then review all potential false positives with over 50 employees in both Compustat and Revelio. For subsidiaries, we review all firms that have over 50 employees in both Compustat and Revelio and the subsidiary's employment in Revelio is over 5% of Compustat employment.

Step 2: Large language model audit. For each match flagged as suspect in Step 1, we ask the LLM (`Gemini-2.5 pro`) whether the Revelio firm and the Compustat firm represent the same business entity in the relevant year.[39] The prompt supplies the model with both firms' names, headquarters, industries, employment levels, IPO dates, and a brief description, and asks for a yes-or-no judgement together with a short reason. We run this audit separately for non-subsidiary candidates where Revelio matches a public firm directly, and for subsidiary candidates where Revelio matches through an ultimate parent. Prompt is as follows:

```
You are an expert financial analyst. Your task is to determine whether two
company names refer to the SAME consolidated legal entity in a GIVEN year.

Your response MUST be a JSON object with two keys:

* "FP_indicator": A boolean value. Set to false if they are the same firm,
and true if they are different.

* "reason": A concise, one-sentence justification for your decision,
referencing the specific year if applicable.

Criteria for SAME firm (FP_indicator: false):

1. One firm acquired, merged with, or officially changed its name to the
other in or before the given year.
```

[39] We use Gemini-2.5 Pro because of its ability to draw on Google Search when assessing company relationships.

```
2. One firm is the legal parent company and the other is its legal subsidiary
in the given year.

Criteria for DIFFERENT firms (FP_indicator: true):

1. The firms are unrelated.

2. One name is merely a product or brand name, such as "Windows," and not a
distinct legal company.

3. The merger, acquisition, name change, or parent-subsidiary relationship
was established after the given year.
```

Step 3: Human review of negative LLM judgment. For every match that the language model labels "no", we conduct a manual review using public sources such as firm websites, SEC filings, and press coverage. This step guards against false negatives from the model. If we confirm the model's rejection, we drop the match. If we overrule it, we retain the match.

Final decision rule. After the three screening steps, we keep a candidate match if either (i) it was never flagged as a suspect in Step 1, or (ii) it was flagged but cleared by the LLM in Step 2 and was not subsequently discarded by the human reviewer in Step 3. We drop the match if the human reviewer in Step 3 confirms it is a false-positive.

A.1.6. Final panel

The output of the procedure is an annual panel with one row per (`rcid`, `year`) cell, containing the matched Compustat `gvkey`, a `match_on` tag indicating which method produced the match (`gvkey`, `cusip`, `cik`, or `ein`), a `subsidiary` indicator equal to one if the match was made through the ultimate parent and zero if direct, and Revelio's annual employment count for the `rcid`.

A.1.7. Resume data employment coverage

Building on the identifier-level mapping panel above, we examine how firms' job counts in Revelio compare to the employment firms report in Compustat from 2010 through 2024. It is worth noting that the employment levels in the two datasets do not have to be equal: Compustat firms typically do not include temp workers, vendors, or contractors as part of their employee counts, while these workers might associate their employer with the firm that contracts with them directly or their direct employer. For example, Google is well known to have a two-tier worker system.[40] Representative surveys show that among

[40]Read about this in New York Times (2019) and Guardian (2018).

workers responding that they are employees, 5 to 8% are actually independent contractors (Abraham et al. 2024, 2026b,a). Hence, for some firms Revelio job counts can, in principle, exceed those in Compustat.

Since Compustat only provides global firm-level employment (U.S.- and non-U.S.-based workers), in Revelio data we similarly aggregate employees located in all countries to firm-year level. Figure A.9 sums employment each year across all matched firms: blue line reports total jobs in Revelio (left y-axis) and red line in Compustat (right y-axis). Two observations emerge. First, the two lines follow each other pretty closely over time. Second, global employment coverage improves over time: from 31% in 2010 to 53% in 2024.

The global employment coverage is lower than what Revelio reports for the U.S.-based workforce (about 2/3 of all U.S. workers are in Revelio data according to Revelio Labs): this is potentially due to lower usage of LinkedIn website—the core source of online resumes—outside of the United States (e.g., few workers in China use LinkedIn). To investigate this hypothesis, we examine the employment coverage ratios by two firm characteristics: firms that are headquartered in the U.S. (those likely with larger U.S.-based employment) versus non-U.S. headquartered firms, and firms with the majority of workers being based in the U.S. in Revelio data (versus minority). Table A.2 shows that, on average, more U.S.-based firms have about 20 percentage points higher coverage rates by both characteristics.[41]

Table A.3 complements Figure A.9 by comparing the firm-level distribution of Compustat employment and Revelio job counts in 2010, 2018, and 2024. Revelio job counts are lower than Compustat employment throughout the distribution, as expected given incomplete online-resume coverage, but as Figure A.9 shows the gap narrows over time to about 45% in 2018 and 53% in 2024.[42] For example, median Revelio jobs increase from 353 in 2010 to 644 in 2024, while median Compustat employment is roughly stable at about 1,165. The improvement is also visible in the upper tail, where the 95th percentile of Revelio jobs rises from 16,647 in 2010 to 34,207 in 2024. These firm-level patterns confirm that the aggregate improvement in employment coverage in Revelio in Figure A.9 is not driven only by a small number of very large firms.

We next examine the distribution of employment coverage by industry. Firms' industry

[41]We manually researched a few companies with headquarters outside of the U.S. but with high shares of employees in the U.S. in resume data. They confirm that these types of firms tend to have low coverage outside of the U.S. A few examples follow. "MITSUBISHI UFJ FINANCIAL GRP" (headquartered in Japan, but has more than 90% of employees in the U.S. according to Revelio data) reports in 2024 about 120,000 employees globally with the US division having 13,000 employees, suggesting its share of U.S. employment is only 11% (Mitsubishi UFJ Financial Group 2024). Similar discrepancies appear for Roche (headquartered in Switzerland): Revelio shows its U.S. share of employees is 97%, whereas public disclosures imply the U.S. employee share is no more than 26% (Roche Holding AG 2024).

[42]We define the coverage in year t as total global Revelio jobs summed across all matched U.S. public firms in that year divided by total Compustat employment summed across the same firms.

codes are from Compustat. First, we look at the employment-weighted coverage: Figure A.10 reports statistics on the ratio of the sum of job counts across all firms in Revelio to the sum of employment across all firms in Compustat within each 2-digit NAICS industry-year cell.[43] The figure shows substantial industry variation in employment coverage, with relatively low coverage in sectors such as retail, wholesale, accommodation & food services, where online-resume data under-represent lower-paid or less professionalized workers. Sectors with more high-skilled jobs, such as tech sectors, finance, or real estate, have the highest coverage ratios.

Second, we look at the equal-weighted employment coverage by industry: Figure A.11 reports, for each industry, the mean and median of firm-year coverage ratio pooled over 2010–2024. Specifically, we first compute the coverage ratio for each firm-year as total global Revelio jobs divided by Compustat employment, winsorize this ratio within year at the 1st and 99th percentiles, and then compute the mean and median across all firm-year observations within each industry. The equal-weighted coverage by industry is overall higher than value-weighted across most industries. This is because employment is highly concentrated among large firms, while smaller firm-years tend to have relatively higher coverage ratios.[44]

We next examine the distribution of firm-level employment coverage ratios over time, defined as Revelio jobs divided by Compustat employment for each firm in each year. Figure A.12 shows the mean (in blue), the median (in red), and the interquartile range (gray shade) of the coverage ratio variable. Consistent with the improving aggregate coverage over time in Figure A.9, this figure shows that, on average, firm-level coverage goes up from 61% in 2010 to 94% in 2024. The figure also shows that there is a wide distribution of employment coverage across firms: at the end of the sample this ratio ranges from 33% at the 25th percentile to 135% at the 75th percentile. We dig into firms for which Revelio employment count exceeds Compustat's reported employee headcount. Oftentimes these are predominantly gig-economy firms (e.g. ride-share, food-delivery) or firms that tend to hire subcontractors (e.g., tech firms) for which Revelio captures independent contractors that are not typically included in Compustat employment.

Given our main empirical specification uses within-firm variation (we examine *changes* in firm outcomes), we next investigate whether firms with high or low employment coverage have issues tracking firm-level employment changes over time (i.e., firms for which changes in employment in Compustat and Revelio are not closely correlated). We partition

[43] Bars report the mean and the median of the annual industry-level coverage ratio over the sample period from 2010 through 2024.

[44] The largest 10% of firm-years account for about 77% of total employment in our Compustat matched sample.

the distribution of this employment coverage ratio into six intervals: below 0.05, 0.05–0.2, 0.2–1, 1–2, 2–3, and at least 3.[45] The two tail intervals identify observations for which the employment comparison is the least informative, since in these ranges Revelio coverage is either nearly zero or exceeds Compustat by more than a factor of three. The remaining intervals distinguish low but nontrivial coverage, partial coverage, approximate parity, and moderate over-coverage. We then compare Revelio and Compustat employment within each coverage interval. In each employment coverage interval, Table A.4 reports the corresponding correlations in levels and after absorbing firm fixed effects. For firm-years with coverage ratios between 0.2 and 3, the cross-sectional correlations are 0.90, 0.99, and 0.998 in the 0.2–1, 1–2, and 2–3 intervals, respectively, and the within-firm correlations are 0.88, 0.94, and 0.99. These patterns suggest that moderate under-coverage and moderate over-coverage mainly reflect level differences in measured employment rather than a failure of Revelio to track firm-level changes over time reported in Compustat.

In the lower tail where the coverage ratio is below 0.05, the within-firm correlation falls to 0.33 indicating that Revelio does not capture a significant share of the variation in changes in firm employment reported in Compustat. The 0.05–0.2 interval is borderline with a high cross-sectional correlation of 0.97 but a more modest within-firm correlation of 0.47. In the extreme upper tail (where coverage is at least 3) Revelio records substantially more workers than Compustat and the within-firm correlation is only 0.38. This analysis suggests that firms in the two extreme tails do not track changes well in firm-level employment in Compustat.

Further analysis shows that firms with low employee counts (either in Compustat or Revelio) also tend to have a low correlation in employment in the two datasets. Table A.5 shows that for firms with fewer than 50 jobs in Revelio, the cross-sectional correlation between employment in Compustat and job counts in Revelio is -1.6% and the within-firm correlation is 5.6%. For firms with fewer than 50 employees in Compustat, we similarly observe low cross-sectional correlation of 3.7%, while within-firm correlation rises to 50.5%.

Informed by these analyses, our baseline firm-level sample for our main regression analysis retains firm-years with the employment coverage ratios between 0.05 and 3,[46] requires at least 50 workers in both Revelio and Compustat, and includes firms headquartered in the U.S. and having at least 50% of their workers in Revelio located in the U.S. For this sample of firms, we focus on the U.S.-based jobs to define AI-related and

[45] The cutoff of 3 corresponds approximately to the 99th percentile of the coverage ratio distribution.

[46] We keep the 0.05–0.2 interval because it continues to align strongly with Compustat in the cross section, and the 1–3 range because moderate over-coverage reflects a systematic level difference rather than poor correlation of within-firm employment changes.

non-AI-related jobs. Table A.6 reports how the sample changes under each filter from 2010 to 2024. The matched Compustat–Revelio sample contains 4,783 firms in 2010 and 4,562 firms in 2024.[47] The employment coverage ratio filter leaves 4,099 firms in 2010 and 4,141 firms in 2024. Among the other one-at-a-time filters, the U.S.-headquarters, Revelio U.S. employment, U.S.-majority-worker, and Compustat employment filters leave 4,040, 3,309, 3,563, and 4,264 firms in 2010, respectively. In 2024, the corresponding counts are 3,487, 3,295, 2,997, and 3,991. Applying all these filters jointly leaves 2,482 firms in 2010 and 2,179 firms in 2024. It is worth noting that our descriptive evidence on AI-related jobs provided below focuses on a broader sample of U.S. public firms with the only restriction of using firm-years with the employment coverage ratios between 0.05 and 3. This allows us to examine statistics over a broader sample of jobs and re-examine some of the data issues we identified in this section.

A.2. Construction of firm-level AI measures

In this section, we describe our methodology for identifying AI-related jobs in resume data: how we construct our AI-related keyword list and use it to measure the share of AI workers in each firm in our sample. Specifically, we proceed in several steps. We first build a list of AI-relevant skills from job postings in Lightcast data. We then expand that AI-relevant skill list using rule-based methods and an LLM. We next screen candidate AI keywords for false-positives and remove terms that tend to pick up AI-irrelevant jobs. Finally, we search the terms from the validated AI keyword list in resume data from Revelio to identify AI-relevant jobs.

We use this methodology because it has several key advantages. First, this approach for identifying AI-related jobs is 100% reproducible, which is critical in academic research. An alternative way is to ask an LLM to identify AI-related jobs, but then one needs to use judgment on how to define AI-related jobs (which is not trivial in quickly evolving AI technologies) and rely on nondeterministic LLM output (there is still variance in LLM classifications even if the LLM parameter to reduce variance—the temperature—is set to zero). Second, this approach is very inexpensive: one can run the relevant code on a laptop in an afternoon. In contrast, using a frontier LLM to classify billions of job records as AI-related or not would cost tens of thousands of dollars. Third, this approach is data-driven: it allows us to uncover AI-relevant job terms in a quickly evolving and dynamic AI field, and as we show, is able to identify jobs across different waves of AI, such as natural language processing and computer vision jobs early on and generative and agentic AI jobs most recently. Most importantly, we show in this extensive appendix that this

[47]The matched sample in Table B1 contains fewer firms because the AI-share calculation requires firms to have at least one U.S. job.

methodology works well: it mostly picks up jobs involved in internal AI development with a false positive rate of 2% and a peripheral (but still related to) AI job rate of 7%.

A.2.1. AI-relevant skills from job postings data

To build the list of AI-relevant skills, we start by following Babina et al. (2024) to identify AI-relevant skills in the job postings data. We use job posting data from Lightcast (which bought Burning Glass—an early provider of high quality job postings data), which covers a near-universe of U.S. online job postings from 2010 through 2025. Each job posting includes a list of skills required by the employer, usually requesting 5 to 20 skills (the interquartile range is 5 to 18 skills, and the 5th and 95th percentiles are 1 and 34 skills, respectively).

There are about 32,000 required skills across all years with the skills standardized by Lightcast. Four of those standardized skills are direct and unambiguous markers of AI technology over the past decade or so: *Artificial Intelligence*, *Machine Learning*, *Computer Vision*, and *Natural Language Processing (NLP)*, which are often abbreviated as "AI", "ML", "CV", and "NLP", respectively. We refer to these as the "core AI" skills.

We then identify other skills closely related to AI by measuring how often each skill appears in the same job posting as the core AI skills across all years of job posting data (2010–2025). For each skill s in the Lightcast skill taxonomy, we define

$$w_s^{AI} = \frac{\text{\# of jobs requiring skill } s \text{ and (ML, NLP, CV, or AI in required skills or in job title)}}{\text{\# of jobs requiring skill } s} \tag{A.1}$$

where w_s^{AI} is the **AI-relatedness score** of skill s, which is the share of job postings requiring skill s that also require at least one core AI skill. We retain skills with high AI-relatedness score (with $w_s^{AI} \geq 0.6$) that also appear in at least 50 job postings (otherwise we pick up skills that only appear a few times: e.g., "Continuous-Time Markov Chain" appears in 1 job posting and "Wit.ai" in 17). The threshold of AI-relatedness of 0.6 balances the risks of excluding highly relevant skills and including less-AI-related skills. In practice, most AI-related jobs are identified by terms above the 0.9 level. This gives us a starting list of 335 AI-relevant skills. In Table A.7, we report the list of the top 50 skills with the highest AI-relatedness scores. Outside of the top four core AI skills, the three skills with 100% AI-relatedness scores are "Zero Shot Learning", "Rweka", and "AgentGPT". Of these, "Zero Shot Learning" is particularly relevant for foundation models and LLMs: it is an ML approach where an AI system performs a task without having seen labeled examples of that exact task during training. For example, models like ChatGPT can answer questions, summarize text, classify content, write code, translate, or reason about new prompts without being specifically trained on each individual task. "Rweka" is an R

package that connects R to Weka, a Java-based collection of machine learning algorithms for data mining. Its strongest connection is to classical AI/ML. Finally, "AgentGPT" is a browser-based platform for creating and running autonomous AI agents. It is most relevant for agentic AI. It is encouraging to see that the methodology introduced in Babina et al. (2024) is able to pick up terms relevant for the recent AI waves such as generative AI (e.g., "AI Generated Code") or agentic AI (e.g., "Agentic Systems").

To avoid missing earlier AI skills that may have become less common over time, we also add the skills from Babina et al. (2024) that are missing in the starting list above, which is based on an earlier vintage of the same underlying job postings data. There are 12 such skills: Libsvm, Sentiment Classification, Torch (machine learning), Jung Framework, Latent Semantic Analysis, Stochastic Gradient Descent, Maximum Entropy Classifier, Expectation-Maximization (EM) algorithm, Classification Algorithms, Bayesian Networks, AI chatbox, ND4J (software).

To assess whether w_s^{AI} simply captures skills that "sound like AI" rather than skills that employers actually pair with AI requirements, we compare our co-occurrence-based measure to a semantic-similarity measure built from sentence embeddings. For each of the 335 seed skills, we compute the cosine similarity between the skill name and four AI anchor phrases ("artificial intelligence," "machine learning," "natural language processing," "computer vision") using a pretrained sentence embedding model (`all-MiniLM-L6-v2`).[48] The Pearson correlation between w_s^{AI} and this embedding-based score is only 0.06 (Spearman: 0.03), and similarly weak (0.04/0.01) when restricted to the 320 seed skills with an exact match in our final keyword dictionary. This weak correlation indicates that the two measures capture different things: terms like "Hugging Face Transformers" and "Seq2Seq" have high co-occurrence scores despite low lexical similarity to core AI terms, while "Text Mining" and "Machine Learning Algorithms" score higher on embedding similarity but co-occur less frequently with core AI skills because they also appear in generic analytics roles. This supports our use of co-occurrence rather than semantic similarity as the primary criterion for keyword selection.

A.2.2. AI keyword expansion and matching method

Having selected the seed AI skills in Section A.2.1, we next convert them into search terms that can be matched against job titles and descriptions in unstructured resume data from Revelio. A single AI skill can be written in many different variants in the text of jobs reported in resumes. For example, "Natural language processing" shows up as "NLP", "Natural Language Processing", "natural language processing", or as part of a

[48]We take the maximum cosine similarity across the four anchors as each skill's embedding-based AI-similarity score.

longer phrase like "NLP engineer". "Large language model" shows up as "LLM", "LLMs" (plural), or as "large language model". To pick these up, we first expand each seed skill into the variants it plausibly takes in resume text. Specifically, we generate variants covering (i) the full skill name; (ii) standard abbreviations (when appropriate); (iii) common variants (singular/plural); (iv) words that are hyphenated (e.g., *multi-modal model*), separated (e.g., *multi modal model*) or merged together (e.g., *multimodal model*), and (v) shorter forms whose meaning is preserved (e.g. "*hugging face*" for "*Hugging Face Transformers*").[49]

Based on the structure of these keyword variants and our extensive experience with NLP and text matching, we define three matching methods that vary in how strictly each candidate keyword must appear in the text:

- **Method 1 (case-sensitive, word boundary).** The keyword must appear in the exact case shown, with a word boundary on each side. A word boundary means that the keyword is delimited by the beginning and the end of the text string, whitespace, or punctuations. For example, "AI" matches the all-caps "AI" (the abbreviation for artificial intelligence), but not the lowercase "ai" or longer words containing that letter combination (e.g., "n**ai**cs" or "N**AI**CS").

- **Method 2 (case-insensitive, word boundary).** The keyword must appear with a word boundary on each side, but case does not matter. For example: "transformer" matches "Transformer" or "transformer", but not "Transformers" or "transformers".

- **Method 3 (case-insensitive, no word boundary).** The keyword need only appear somewhere in the text, ignoring case and word boundary. For example, a keyword "language model" picks up "language models" or "Multimodal Language Model" without us having to list these variants separately.

In this step, we transform each of the highly AI-relevant skills from the job postings into search terms across these three search methods. For example, "Machine Learning" gets searched as "ML" in Method 1 (case-sensitive, word boundary); as "ml/ai" and "ai/ml" in Method 2 (case-insensitive, word boundary); and as "machine learning" in Method 3 (case-insensitive, no word boundary). "Microsoft Cognitive Toolkit (CNTK)" gets searched for "CNTK" in Method 1 and "microsoft cognitive toolkit" in Method 3. "BERT (NLP Model)" gets searched for "BERT" in Method 1; "bert nlp" in Method 2; and "bidirectional encoder representations from transformer" in Method 3. This step expands the AI-relevant search term list to 433 unique terms.

[49] In principle, we could have used an LLM to do "fuzzy matching" on these terms. We decided to go with a more traditional "bag of words" search approach because: 1) this method is very transparent and reproducible; and 2) the computation costs outweigh the benefits of a small increase in precision.

To complement this manually created rule-based expansion, we use a large language model (`GPT-5-nano`) to suggest additional spellings and synonyms we might have missed. We feed to the model the AI-skill spreadsheet developed to this point one row at a time. For each row, the model sees the seed AI skill and its existing list of variants, and is asked to propose more variants. The prompt we give the model is the following:

```
We are trying to identify AI-related jobs in LinkedIn data by searching for
AI-relevant terms in people's job titles and job descriptions. For the
provided target term, suggest alternative ways people may write the same
concept on LinkedIn profiles or job postings.

Requirements:

- Only include highly related variants that improve recall.

- Include spelling variants, formatting variants, standard abbreviations,
official shorthand, branded sub-names, and very close synonyms only when they
are genuinely used as alternative search surfaces for the same concept.

- Do not include unrelated tools, competing tools, adjacent concepts, broader
categories, narrower subfields, job titles, or role labels unless they are
truly used as alternative names for the same concept.

- Do not include terms already covered by the existing terms list.

- If the existing terms already seem sufficient to capture the concept well,
return no additional suggestions.

- Prefer returning no suggestions over returning a borderline term.

- Flag terms that may create many false-positives.

For each suggestion, return: the term itself, its frequency in real text
(Very High / High / Medium / Low), its relevance to the seed concept (High /
Medium / Low), and its false-positive risk (Low / Moderate / High), along
with a short note.
```

For example, for the original skill in the job postings *Hugging Face Transformers*, the LLM suggested to add *HF transformers* and *HF transformer* (case sensitive and with word boundary). Lastly, we manually verify each GPT-suggested search term before adding it to the keyword list. For each candidate term, we conduct a web search and inspect the top results to assess whether the term is unambiguously AI-related. Terms whose results are clearly related to AI are retained, while terms yielding unrelated or ambiguous results are

excluded. This step expands the AI-relevant search term list by an additional 89 unique terms (although by construction all of them are highly related to the already included terms).

A.2.3. False-positive screen of AI-relevant search terms

Next, we investigate whether we included some AI-relevant search terms that can also have a meaning not related to AI and, hence, identify some jobs as AI-related when, in fact, they are not (i.e., generating false-positive AI jobs). To do so, for every keyword, we compute in 2010–2024 using Revelio data (i) the number of U.S. positions that the keyword matches unconditionally, and (ii) the additional jobs identified with the keyword in addition to the jobs identified with the list of core AI-relevant terms such as AI, ML, NLP. The keywords with a similar number of jobs in groups (i) and (ii) are flagged as false-positive candidates. The idea is that if a term appears in job descriptions that do not also contain key AI terms, this potentially indicates that the term might be picking up jobs that are not AI relevant. For example, the term "RL" (the abbreviation for reinforcement learning) appears in about 10,700 jobs, and 8,870 of those do not contain any core AI terms (we later learn that "RL" in resumes usually refers to "Ralph Lauren", an apparel company). We also flagged a few terms as potentially high false-positive risk because they might be referring to something else like "feature engineer" (someone who works on features of a product, system, or platform not necessarily AI-related).

We identify 26 terms as potentially high false-positive risk. The flagged keywords are then run through two complementary checks: LLM classification at scale and human verification on a random subsample.

LLM classification at scale. For each suspect keyword, we collect **all** job positions from our public firm sample in Revelio whose title or description contains the keyword. We pass each such position, one at a time, to a large language model (GPT-5.4) with a structured prompt that includes the keyword that triggered the match, the position's title, and the position's description. The model is asked to judge whether the position should count as AI-related given the context. The instructions we give the model are the following:

```
You are reviewing LinkedIn profile records to determine whether each job
should count as AI-related under a broad research definition.

Count a job as AI-related (is_ai_related = 1) if the role develops, uses,
supports, manages, researches, deploys, sells, or materially relies on
AI-related skills or tools. The role does NOT need to be purely AI-focused.
If AI-related work is one meaningful part of the job, it counts.

Broad AI-related examples that should usually count:
```

```
- machine learning, deep learning, neural networks, LLMs, generative AI,
prompt engineering

- data science, applied science, quantitative modeling, predictive modeling,
statistical learning

- NLP, computer vision, speech, robotics, recommender systems,
search/ranking, AI optimization

- MLOps, AI infrastructure, model deployment, model monitoring, inference
systems

- AI product management, AI program management, AI solutions engineering, AI
consulting

- software, analytics, or data roles that explicitly use ML/AI models, AI
APIs, or related methods

Mark a job as NOT AI-related (is_ai_related = 0) only when it is clearly
unrelated to AI/data-science/ML work and the matched keyword is incidental,
ambiguous, or refers to something else.

Important rules:

- Use a broad definition: if the job uses or entails AI-related skills, count
it.

- If the description is missing, vague, or uninformative, infer from the
title and matched keyword and make your best guess.

- When evidence is mixed, lean toward 1 rather than 0.

- Do not require the role to build frontier models or have 'AI' in the
official title.

Return ONLY valid JSON with exactly these keys:

{"is_ai_related": 0 or 1,
 "reasoning": "short explanation, at most 15 words"}
```

We then aggregate the model's judgments across keyword-matched positions to estimate each flagged keyword's false-positive rate, defined as the share of positions matched by that keyword that are not actually AI-related. This gives us false-positive rates for all jobs identified as AI-related for the list of suspect search terms.

It is important to note that in the next step we validate the model's ability to identify

AI-relevant jobs against the human classification by the authors on a random sample and find similar classification results.

Human verification on a random subsample. To verify whether the LLM correctly identifies AI- versus non-AI-relevant jobs in the above step in a systematic way, we draw a random subsample of keyword-matched Revelio positions for each suspect keyword, with 20 unique positions per keyword. We first ensure that this random sample is similar to the population of jobs with high potential false-positive risk: we use the same LLM prompt and model as in Section A.2.3 and compare the false-positive rates in this random sample to the universe of AI-related jobs identified in the step above. We confirm that the false-positive rates are very similar (the typical difference is less than 5%).

Next, using this random sample, each co-author independently reviews the position title and description and codes whether the position is genuinely AI-related, without seeing the LLM's classification. Aggregating these human labels gives a second, independent estimate of each suspect keyword's false-positive rate.

We then combine the LLM and the human false-positive estimates to decide whether to keep or drop each keyword. First, the agreement between the human and the LLM on which jobs are not AI-related and, hence, on false-positive rates, is quite high. Across the entire random sample, the human and the LLM classifications agree 91% of the time. This high agreement rate is reassuring, suggesting we can use this prompt to classify AI-tagged jobs as truly AI-related or as false positives.

Next, for each high potential false-positive term, we decide whether to keep it or drop it. If both estimates (the LLM's and the human's) of the false-positive rates are high (60% or higher), indicating that the keyword mostly matches non-AI positions, we drop it. For example, the candidate keyword *RL* (intended to capture the abbreviation for reinforcement learning) had a 95% false-positive rate by the LLM in the screened Revelio positions because it overwhelmingly referred to *Ralph Lauren* rather than *reinforcement learning*, and was therefore dropped. Another example—with a high false-positive rate of 90% that was dropped—is *loss adjusting*: in ML context it refers to changing how a model's loss function is calculated or weighted during training, but in resume data it typically refers to insurance context of investigating and assessing an insurance claim.

Next, if both false-positive estimates are low (less than 20%), meaning that both screens suggest the keyword mostly captures genuine AI jobs, we keep the keyword. For example, *language model* had only a 4% false-positive rate and was kept. Similarly, we keep *feature engineer* with a 15% false-positive rate.

Finally, if the two estimates disagree or have moderate false-positive rates of under 40%, we manually review additional matches for that keyword case by case until there is enough evidence to decide. Overall, we keep 14 out of 26 suspect AI terms.

After these steps, our final AI-relevant keyword list contains 510 keywords (433 original terms from humans plus 89 from the LLM minus 12 terms with high false-positive rates).

A.2.4. AI subcategories

AI techniques have evolved across several waves. The big breakthrough occurred in the 2010s with the "Machine Learning Era" when models began learning directly from data rather than following predefined logic. During this time, models remained focused on narrow applications: they performed well in specific domains—such as playing games or detecting fraudulent transactions. By the 2020s, the technology had evolved into generative AI, enabling more general-purpose tasks such as drafting research reports. More recently, firms have begun experimenting with agentic AI: systems built on top of generative models that can plan, take actions across tools, and iteratively pursue multi-step goals with limited human prompting.

To capture different waves of AI, we further classify our AI-related keyword list into AI subcategories: early waves of AI labeled as "Wave I AI": machine learning, natural language processing, computer vision, and more recent ones labeled as "Wave II AI" (generative AI, agentic AI). We also construct a sixth category, general AI, that is about more general AI terms that are not specific to any other AI category (e.g., *artificial intelligence*, *AI ethics*). We do so by using LLM-assisted classification, followed by author review.

We pass the keywords to Claude 3 Opus and ask it to assign each keyword a probability of belonging to each of six categories. These are the five AI subcategories and a residual "general AI" category. The six probabilities sum to one, and the model also provides a short explanation for each classification. The full prompt is reproduced below:

```
Can you classify the attached AI technologies into one of five categories:
machine learning, natural language processing, computer vision, generative
AI, agentic AI? The categories are mutually exclusive. Create one column for
each category with the probability that technology belongs to the category.
All five probabilities should sum up to one. Please also add a column with
your explanation for the classification. If a technology is general and does
not belong to a specific category, label the technology as "general AI" in a
new column. For each row, highlight the column with the highest probability.

Here are some examples in each category:

Machine learning: machine learning, deep learning

Natural language processing: NLP, chatbot

Computer vision: computer vision, object detection
```

```
Generative AI: GenAI, LLM, prompt engineering

Agentic AI: agentic system, CrewAI
```

Two co-authors independently review the output of the LLM classification and research ambiguous terms. We then assign each keyword to the AI category with the highest probability among the six if that maximum probability is equal to or above 0.5. Keywords for which the maximum probability is below 0.5 are placed in an "other" category. This produces a list of (mutually exclusive) terms specific to each of the six categories plus "other" category. The examples for each category include: *ethical ai* for general AI, *pytorch* for ML, *doc2vec* for NLP, *rekognition* for CV, *ai generated code* for generative AI, and *crewai* for agentic AI. The "other" terms tend to be concepts and tools that are important across different AI categories: for example, "attention mechanism" has been important in early work with text in NLP applications, but is also used in generative AI models. Of course, many of our AI search terms span across interrelated AI waves as AI progress tends to be cumulative with newer AI waves utilizing new and existing AI-related technologies, tools, and concepts. For example, while *deep learning* (classified as ML category with 85% probability and, hence, belongs to the ML subcategory) has emerged as an important methodology within a broader machine learning family in the 2010s, it is also an "engine" underlying many generative AI applications. So this subclassification should be thought of as approximate.

A.2.5. Identifying AI jobs in worker resume data

Using our comprehensive AI-related keyword set described above, we apply it to the worker resume data from Revelio Labs to search for AI-related jobs in our U.S. public firm sample. Revelio provides standardized data on each job a worker reports in his online resume with the job's starting and ending dates, and, critically, the job title and its description (we use both of them to search for our AI-related terms). While Revelio offers international coverage of worker resumes, we focus on U.S.-based positions for identifying AI-related jobs since, as we report in Part A.1.7, the U.S. coverage of employees of U.S. public firms tends to be more complete than its international coverage and since our keyword list is in English (creating and validating similar lists of AI-relevant search terms for all other languages using the steps above would be a major undertaking, which we leave for future research).

We first construct our main firm-level AI measure that spans all AI-related keyword terms. For each full-time U.S. job, we search its job title and job description against our comprehensive AI keyword list, applying the keyword-specific matching rule defined in

Section A.2.2.[50] A job is classified as AI-related if at least one AI keyword is matched in either the job title or the job description. We also construct firm-level AI measures for the AI subcategories described above: general AI, machine learning, natural language processing, computer vision, generative AI, agentic AI, and "other". We repeat the search for AI-relevant jobs but now use the AI-related keyword list specific to each AI subcategory. In addition, we construct firm-level AI measures that intend to capture early (2010s) versus more recent (2020s) waves of AI. We define "Wave I AI" jobs as jobs matching at least one of the ML, NLP, or CV search terms, and "Wave II AI" jobs as those that match at least one of the Gen AI or Agentic AI search terms. A job can be in several subcategories if it contains keywords from different AI subcategory lists.

To construct firm-level measures of AI-skilled jobs, we aggregate the firm-worker-job-year panel to the firm–year level by calculating the share of a firm's jobs that are AI-relevant in each year. Specifically, let $Job_{f,t,\text{US}}$ denote the set of full-time U.S. positions at firm f in year t, and let $AI_{f,t,\text{US}} \subseteq Job_{f,t,\text{US}}$ denote the subset classified as AI jobs. The firm's AI job share is defined as

$$\text{Share AI Jobs}_{f,t} \;=\; \frac{AI_{f,t,\text{US}}}{Job_{f,t,\text{US}}}.$$

We repeat this exercise by using the counts of AI-relevant jobs specific to each AI subcategory, and calculate the shares of AI workers for each subcategory (the number of workers in an AI subcategory divided by the total U.S. employment). We also calculate firm-level AI shares for "Wave I AI" jobs (NLP, ML, or CV) and for "Wave II AI" jobs (Gen AI or Agentic AI). Note that the sum of AI shares across subcategories does not have to add to the main AI share (generated across all search terms) since a job can be classified into several AI subcategories if it has search terms from several AI subcategories. To avoid this double counting and create mutually exclusive job categories, we also generate three additional versions of AI jobs: those *only* containing search terms from "Wave I Only AI Jobs" search term list, those *only* containing search terms from "Wave II Only AI Jobs" search term list, and those that contain search terms from *both* lists, "Wave I & II AI Jobs".

A.2.6. AI-job false-positive rate

We next offer several analyses of our new data on AI-related jobs. All of these statistics are based on the sample of U.S.-based workers of U.S. public firms matched to Revelio that have employment coverage between 0.05 and 3. This broad sample allows us to examine statistics over a comprehensive sample of U.S. jobs and re-examine some of the data issues

[50] En-dash and em-dash characters in job titles and descriptions are normalized to hyphens, and consecutive whitespace collapsed to a single space. We also exclude positions with titles containing terms that indicate internships, part-time work, student roles, postdoctoral roles, or other temporary positions.

we identified in Part A.1.7.

We start by examining the classification error of AI-related jobs: jobs that our methodology tags as AI-related but might not actually be AI-related, either because a search term might be used in non-AI contexts or because people put AI-related terms in their resumes without doing any significant development or implementation of AI in their jobs.[51] We also examine this error across all firms and by firm size (Figure B4 shows that the share of AI-related jobs is disproportionately concentrated in larger firms and the time-series of these shares are more volatile for smaller firms).

We start with Table A.8 which reports false-positive rates by firm-size terciles. Using the same LLM prompt and model as in Section A.2.3, we audit *all* 5,084 AI-classified job-year observations in the small-firm tercile, and equally sized random samples from the middle- and large-firm terciles for symmetric comparison (we draw random samples due to large number of jobs in these larger-firm terciles). The false-positive rate declines monotonically with firm size, from 5.47% among small firms to 3.38% among middle-tercile firms and 2.03% among large firms. Keyword-based AI classifications are therefore somewhat noisier among smaller firms, but the estimated false-positive rate remains modest across all three groups. These false-positive rates translate into an overall moderate false-positive rate of 2.13% across all firm sizes (by weighting the false-positive rates for each firm size tercile by the share of jobs in large, mid, and small firms, specifically 2.13% =0.55%*5.47%+3.38%*5.79%+93.66%*2.03%).

We next repeat the false-positive analysis for each AI subcategory by randomly drawing 5,000 AI-related tagged jobs in each AI subcategory. Table A.9 shows that the NLP category has the highest false-positive rate of 2.7%: it turns out that the NLP-specific term "BERT" (Google's early AI model) can also stand for "Bit Error Rate Tester", which is a test instrument used to measure how accurately a digital communication link transmits data and is not related to AI, pushing the false-positive rate up for the NLP subcategory. The general AI category has the second highest false-positive rate of 1.8%. Other subcategories have the following false-positive rates: 0.7% for agentic AI, 1.2% for CV, 0.5% for gen AI, 1.6% for ML, and 0.4% for "other" category. Overall, all these false-positive rates are small.

A.2.7. Summary statistics on AI-related jobs

We next provide summary statistics on the AI search terms, AI-related jobs, and firms with highest AI shares and job counts. As in the previous Part A.2.6, all statistics below are based on the sample of U.S.-based workers of U.S. public firms that have the employment

[51]The procedure on the removal of AI terms that lead to high false-positive AI job rates described in Part A.2.3 addresses whether individual keywords tend to match genuine AI-related positions, but likely does not eliminate all false-positive AI matches.

coverage between 0.05 and 3 (unless explicitly stated otherwise).

We start by plotting the number and the share of jobs that are AI-related in Figure 1 panels (a) and (b), respectively. We see a strong upward trend in both figures, suggesting there are more AI-related jobs over time overall and as a share of all jobs. By the end of the sample in 2024, there are more than 300,000 AI-related jobs which corresponds to about 2.1% of all jobs. These figures represent aggregate trends (the number of AI jobs is the sum over all AI jobs each year, and the share of AI jobs is the fraction of all AI jobs to total jobs). To examine firm-level AI workforce dynamics, we recalculate AI worker shares as the average of AI shares across all firms in a year and AI worker numbers as the average number of AI workers per firm. We plot these averages in Figure A.13. At the end of the sample, on average, a U.S. public firm has 78 AI-related jobs (panel (a)) and 1.5% share of AI-related jobs (panel (b)). Figure A.14 shows that the distribution of these AI-related job numbers and shares is quite wide across firms: this figure re-plots the means from Figure A.13 and also adds colored bands for different parts of the distribution: 90–95th percentile in the darkest blue, 75–90th percentile in the medium-light blue, and 50–75th percentile in the lightest blue. While the median firm has very few AI jobs, Figure A.14 shows that a firm at the 95th percentile has 213 AI jobs and 6.7% AI share.

Figure 5 shows aggregate AI shares by NAICS-2 industry sector at the beginning (2010), the middle (2018), and the end of our sample (2024). It shows that the tech sectors have the highest shares of AI workers, although the growth in AI-related workers over time occurs across all sectors.

We also examine aggregate AI shares over time by AI subcategories. Figure 3 shows that the growth in AI job shares in "Wave I" AI technologies (ML, NLP, CV) has either slowed down (ML) or started to decelerate (CV, NLP) at the end of the sample (although the levels are still high, especially for ML that identifies about 0.9% of jobs as AI-related in 2024). Instead, more recent waves of AI (generative AI and agentic AI) accelerate very rapidly at the end of the sample, with a 0.45% share of generative AI jobs in 2024. Table A.10 tabulates the numbers underlying these plots from Figure 3. Figure 4 shows these trends even more clearly where we aggregate ML, NLP, and CV into a single category "Wave I" AI, and generative AI with agentic AI into "Wave II" AI. In relative terms, "Wave I" AI jobs are about 1% and "Wave II" AI are about 0.5% of jobs in 2024. Note that the sum of these two shares is lower than the aggregate share 2.1% of AI jobs because these AI categories do not rely on search terms from "General AI" category (e.g., "AI") or "Other" category (e.g., "operationalizing ai").

We continue by examining which AI-relevant keywords identify most AI jobs overall and by six AI subcategories in 2024 (the last year of our resume sample). Table A.11 reports the top 50 keywords with the most AI jobs. The top 5 terms are *AI, machine learning,*

ML, generative ai, ai/ml. Table A.12 shows the top 10 search terms for each specific AI subcategory and all search terms make sense to us. Besides the core AI terms, the top three terms by job counts are as follows: for general AI, it is *ai research, ai innovation, ai security*; for ML it is *deep learning, numpy, tensorflow*; for NLP—*conversational ai, sentiment analysis, BERT*; for CV—*object detection, image recognition, opencv*; for generative AI—*prompt engineering, chatgpt, retrieval-augmented generation*; and for agentic AI—*langchain, ai agent, langgraph*. We also examine which words most often appear in job titles and descriptions for each AI subcategory: Figure 2 shows that while each AI subcategory has its own specific AI words, one common and important word is *data*: this makes sense since any type of AI needs data for AI to be useful.

We provide several examples of AI-related jobs for each AI subcategory in Table A.13. In the general AI category, we tend to observe leadership jobs, jobs that leverage AI but where AI is not the main focus, and cross-functional roles that help build AI into products and processes. In more specific AI subcategories (ML, NLP, CV, generative AI, agentic AI), AI-related roles tend to be more technical and data heavy. Agentic AI-related jobs tend to contain Gen AI related terms: this is not surprising since agentic AI relies heavily on GenAI capabilities.

We next examine more systematically job titles of AI-related jobs. We use job titles self-reported by workers, which tend to be less standardized than, say, an occupation: for a functionally similar job, different firms might use different job titles.[52] Table A.14 shows 30 job titles with the most AI-related jobs in 2024. Intuitively, it shows that AI-related jobs tend to be in software, engineering, scientific, and data-related positions. Table A.15 shows the top 10 job titles (by the number of AI jobs) for each AI subcategory. We observe a remarkable overlap of job titles across all AI subcategories.

Using jobs' occupation data standardized by Revelio, Table A.16 shows the top 30 occupations with the highest *shares* of AI workers in 2024. The top three occupations are: 1) "Computer and Information Research Scientists" with a whopping 0.66 AI share, 2) "Data Scientists" (0.43), and 3) "Mathematical Science Occupations, All Other" (0.39). Table A.17 reports top 30 occupations with the highest *number* of AI jobs in 2024: the top three are 1) "Data Scientists" with 50,403 AI jobs (with 0.43 AI share); 2) "Software Developers" with 40,418 jobs (0.06 AI share), and 3) "Computer and Information Research Scientists" with 23,985 jobs (0.66 AI share). Upon inspection, these top AI-related occupations intuitively cover many engineering, data, and scientific occupations, but also domain area experts (e.g., marketing managers, designers, financial quantitative analysts) whose knowledge is

[52]We lightly standardize job titles by aggregating across closely related title variants: seniority and level modifiers, such as "Senior," "Staff," "Principal," "Lead," "II," and "III," are combined when the underlying role is the same.

needed to apply AI effectively to solve business problems. These tables also highlight that using occupations to identify AI-related workers is not a good strategy since even the very top ones have high shares of non-AI-related workers.

We next examine education, compensation, and job seniority of AI- and non-AI-related jobs (we do not examine these characteristics at the worker-level since many workers start as non-AI-related and then transition into AI-related roles). We find that AI-related workers are more educated and more highly compensated, and tend to be in more senior roles. Among workers who report college education, Figure A.15 shows the shares of workers with the highest educational attainment being a Bachelor, a Master's, or a Doctoral degree for AI-related jobs (left bar) and non-AI-related jobs (right bar). Perhaps unsurprisingly, AI-related workers tend to be more educated: 11% have a Doctoral degree (versus 4% for non-AI workers), 52% have a Master's degree (34% for non-AI workers), and 37% have an undergraduate degree (62% for non-AI workers). We also examine the mean differences in the total compensation of AI and non-AI-related jobs (the compensation data is imputed by Revelio by using industry, location, and occupation of each job, so it is inherently noisy). Figure A.17 shows that, on average, AI-related workers have 62% higher compensation than non-AI-related workers: the average total compensation is about $224 thousand for AI-related, and $138 thousand for non-AI related jobs.

Lastly, we show that AI-related workers tend to be in more senior roles than non-AI-related workers in Figure A.16.[53] AI-related jobs are much less concentrated in junior roles than non-AI jobs: 20% of AI-related positions are junior, compared with 53% of non-AI positions. Instead, AI-related positions are more concentrated in mid-level and senior roles. Mid-level positions account for 42% of AI jobs, compared with 29% of non-AI jobs, while senior positions account for 38% of AI jobs, compared with 18% of non-AI jobs.

We switch gears to examine which specific firms have large AI-related workforces. It is worth a reminder that these firm-specific AI shares of workers are approximate due to the reasons described in Part A.1.7, and should be taken as a proxy for how AI-focused firms are. For firm-specific statistics, we keep firms with at least 50 jobs in Revelio and 50 employees in Compustat to avoid reporting on very small firms. We start by plotting the shares of AI workers over time for each firm with the highest AI share in 2024 within each NAICS-2 industry sector. Firms' industry codes are from Compustat. To observe time-series trends, we require a firm to have at least 5 years of data in Figure A.19 and 10 years in Figure A.20. While most firms across the two figures show an

[53] Revelio assigns jobs to seven seniority levels: 1 corresponds to entry-level roles, 2 to junior roles, 3 to associate roles, 4 to manager roles, 5 to director roles, 6 to executive roles, and 7 to senior executive roles. We aggregate these seven levels into three broader groups: junior roles include levels 1–2, mid-level roles include levels 3–4, and senior roles include levels 5–7.

upward trend in their AI shares, we keep both figures because they show different sample composition patterns. Starting with Figure A.20 that includes firms with longer time-series, we see that some of the well-known AI-intensive incumbents include "Expedia", "Visa", and "Nvidia".[54] Turning to Figure A.19 that includes firms with shorter time series, we see that many firms with the highest AI shares in their respective industry are newly public firms that are often "AI native". A few examples follow. "Corteva" in Agriculture focuses on seed, crop protection, and digital agriculture: AI helps predict pest-protein structures, identify crop-protection molecules, model thousands of molecule candidates, accelerate biological discovery, optimize fermentation, guide field decisions such as fungicide timing, and support regulatory dossier preparation. "DraftKings" in Arts/Entertainment (digital sports entertainment and gaming firm) uses AI in live-betting pricing and technology platform. "Upstart" in Finance and Insurance reports that its AI models are central to its lending marketplace and are used for credit-risk assessment, income and identity fraud, borrower acquisition, loan stacking, default and prepayment prediction, and servicing. "Ontrak" (in Health Care) uses AI-enabled predictive algorithms and claims analytics to identify members with behavioral-health needs and medical-cost risk. "C3.ai" (Information) is an enterprise AI software company selling AI platforms, applications, and generative-AI products.

Next we report top 3 firms with the highest *shares* and *numbers* of AI workers in 2024 within each NAICS-2 industry sector in Tables A.18 and A.19, respectively.[55] The two tables highlight a distinction between AI intensity and AI scale: the AI share ranking tends to identify firms for which AI-skilled labor is central to the organization, while the AI job count ranking captures firms with the largest absolute deployment of AI workers.

In Table A.18, the intensity leaders are typically AI-native, data-platform, or highly specialized firms. "Innodata" has the highest AI share in Professional and Business Services, with 71.7% of its U.S. workforce in AI-related roles. "C3.AI" leads Information with a 52% AI share, "Coupang" leads Retail Trade with 21%, "Nerdy" leads Education Services with 18%, "Tempus AI" leads Health Care with 11%, and "Upstart" leads Finance and Insurance with 10%. Many of these firms have modest AI headcounts, but AI accounts for a large fraction of their workforce. They also tend to be relatively small firms by overall headcount.

The scale leaders in Table A.19, by contrast, are large incumbents with many U.S.

[54] In agriculture, Figure A.20 shows zero AI shares for "Alico" (a US-based agribusiness and land management firm headquartered in Fort Myers, Florida): none of the firms in this sector with at least 10 years of data have positive AI workforce in 2024.

[55] We do not include firms in agriculture since there is only one firm that has non-zero AI workers and we already reported on it in Figure A.19.

workers. Unsurprisingly, Information sector has the firms with the highest numbers of AI jobs: "Microsoft", "Alphabet", and "Meta" lead Information sector with 26,169, 21,385, and 15,075 jobs, respectively. "Amazon" and "Walmart" lead Retail Trade with 14,668 and 1,876 AI jobs respectively. In Finance and Insurance, "JPMorgan Chase", "Capital One", and "Wells Fargo" lead with 4,527, 3,284, and 2,699 AI jobs, respectively. In Health Care, "CVS Health" leads with 2,173 AI jobs. In most sectors, the AI job share and the AI job count leaders differ, underscoring that AI intensity and AI scale capture distinct subpopulations of firms.

Because technology firms are natural AI adopters and dominate many aggregate AI rankings, we next separate technology and non-technology firms to examine top 30 firms with the highest numbers and shares of AI workers across these two groups of firms. Within the tech sectors (Information; Professional and Business Services), Tables A.20 and A.21 show that AI jobs are concentrated among large incumbent platforms and enterprise-technology firms, while AI intensity is concentrated among specialized AI and data firms. "Microsoft", "Alphabet", "Meta", and "IBM" together employ about 72,000 AI jobs in 2024. The next tier includes "Accenture", "Salesforce", "Oracle", "Adobe", "Verizon", "AT&T", "Booz Allen Hamilton", "ServiceNow", and "Intuit", each with more than 1,900 AI jobs. The AI share ranking, however, is led by smaller and more specialized AI and data-intensive firms. "Innodata", "C3.AI", "Coveo Solutions", "Baidu", "Cerence", "Veritone", "SoundHound AI", "LivePerson", and "Rekor Systems" report AI shares ranging from about 20% to 72%.[56] Thus, within technology, large platforms dominate the absolute stock of AI workers, while specialized AI and data-intensive firms are the most AI-intensive.

Outside the tech sectors, Tables A.22 and A.23 suggest a similar distinction across a broader set of industries. The non-tech AI share ranking is led by firms in semiconductors, AI hardware, biotechnology, fintech, education technology, and data-intensive platforms. Given the nature of the firm, "Nvidia" has one of the highest non-tech AI shares and one of the largest non-tech AI workforces, with an AI share of 30% and 8,965 AI jobs. Other high-share non-tech firms include "Xpeng" at 36%, "Absci" at 23%, "Coupang" at 21%, "Nerdy" at 18%, "Ambarella" at 18%, "Weride" at 18%, and "Ceva" at 16%. These firms illustrate that high AI intensity outside technology is concentrated among smaller firms whose products, production processes, or platforms are more closely tied to AI. The non-tech count ranking instead is led by large incumbents. "Amazon" has 14,668 AI jobs, followed by "Nvidia" with 8,965, "Apple" with 7,900, "Intel" with 5,430, "JPMorgan Chase" with 4,527, "Cisco" with 3,972, "Capital One" with 3,284, "Dell" with 3,163, "Wells Fargo" with 2,699, and "Tesla" with 2,545.

[56] Baidu has only a limited presence in the U.S.

Taken together, these tables show that AI-skilled labor has spread well beyond a narrow set of technology firms, but along two different margins. The firms with the highest AI shares are often smaller AI-native or data-heavy companies, where AI is central to the business model. The firms with the largest AI workforces, by contrast, are mostly large public firms in technology, manufacturing, finance, retail, and health care.

It is interesting to contrast these data points (showing that the most AI-intensive firms tend to be smaller firms) with the AI intensity across all firms by firm size. Figure B4 shows that by firm size terciles (small, medium, big), the shares of AI workers are largest for big firms, followed by medium-sized firms, and smallest for small firms. This is due to lower average AI adoption rates by small and medium-sized firms than by large firms: in 2024, the share of firms with at least one AI job was 28% in small-firm tercile, 73% in medium, and 94% in big firm tercile.

A.2.8. Comparison to firm-level data from Cognism

We next compare our new firm-level data built from resume data from Revelio Labs to the similar data we have built from resume data from Cognism (which we introduced in Babina et al. (2024)). Since the underlying source of resume data has changed and since we updated the methodology for identifying AI-related jobs, we want to investigate how employment count as well as AI measures compare across the two resume datasets.

Table A.24 reports the correlations between Revelio and Cognism AI measures in each year from 2010 to 2021 (last year when Cognism data were available) for the regression sample. The correlations reported in the table are generally high, indicating that the Revelio and Cognism datasets paint a broadly consistent picture of the firms' workforce while not being identical. Agreement is strongest for total job counts, where correlations are uniformly high—ranging from 87% in 2010 to 94% in 2021—and rise steadily over the period, suggesting the two sources track overall employment levels closely. The datasets are also positively correlated on the AI share, with agreement generally strengthening over time: going from 45% in 2010 and increasing to 72% in 2021. The consistent positive correlation indicates that the two sources move together in capturing the AI intensity of hiring.

A.2.9. Comparison to firm-level data from Lightcast

Because some resumes are updated retrospectively, workers might relabel past positions using AI terminology that became fashionable only later (Bloom et al., 2026b). Job postings data provides a natural check for this issue. They are contemporaneous employer records, so if resume-reported AI hiring reflects genuine firm demand rather than relabeling, the years in which a firm's workers report starting AI positions should coincide with the years in which the firm advertises AI-related vacancies. For each firm-year over 2010–2024, we

compare the number of new AI hires in the Revelio resume data with the number of the firm's Lightcast job postings that list the same AI-related skills from our keyword list. We also restrict the measure to skills that (i) are counted only from their first appearance in the postings data, so late-arriving vocabulary cannot backdate resume positions, and (ii) appear in both the firm's job postings and its new hires' resumes within the firm-year.

Appendix Figure A.21 shows a binned scatter plot of new AI resume jobs against AI-related postings for each firm-year, residualized on year fixed effects, total postings, and total U.S. resume jobs. The partial correlation is 0.89, and 0.74 when excluding firm-years with no AI activity in either source.

A.3. Construction of AI job categories

In our next step, we investigate what kinds of jobs employees do in our AI-related job sample. For this, we draw a random sample of 300 AI jobs each year from 2010 through 2024, and categorize this 4,500 sample of jobs into several functional roles. As a first step, we read through a random sample of thousands of job descriptions to build intuition about the most common types of AI-related jobs.

Two-step classification design. We implement the classification in two stages. In *Step 1*, the LLM constructs a mutually exclusive and collectively exhaustive taxonomy of 8–12 functional categories from the job sample. In *Step 2*, we freeze this taxonomy and assign each job to exactly one category based on the worker's job title and job description. For the primary results reported in Tables A.25 and A.26, we use *GPT-5.4* with a "low" reasoning effort to run Step 1 and Step 2 on all 4,500 jobs, yielding the 11-category taxonomy in Table A.25 and the category counts in Table A.26.

Initial pipeline and alternative approaches. We also evaluated an embedding-based clustering approach as a lower-cost alternative to LLM classification. The resulting clusters overlapped only partially with the LLM taxonomy, and boundary cases proved difficult to interpret consistently, so we did not adopt this method. Turning to the stability of the LLM approach itself, we find that independent Step 1 runs with a fixed prompt and fixed model settings yield taxonomies with the same number of categories and broadly similar structure. We therefore fix the primary taxonomy from the initial classification run and proceed to validating Step 2 reproducibility and classification accuracy.

Step 2 reproducibility. The following checks are separate from the primary classification run above. We assess Step 2 stability using two frozen taxonomies: the primary taxonomy (11 categories in Table A.25) and an alternative 12-category taxonomy from a separate Step 1 run with *GPT-5.6* (high reasoning). For each taxonomy, we run Step 2 twice independently on the same 4,500 jobs, using *GPT-5.4* with temperature set to zero and reasoning effort set

to none. The two passes agree on 4,326 of 4,500 jobs (96.1%) under the primary taxonomy and on 4,335 of 4,500 jobs (96.3%) under the *GPT-5.6* taxonomy. Step 2 is therefore highly reproducible under both codebooks.

Human audit. We further audit classification quality on a random subsample of 200 jobs from the same 4,500 jobs. Blind to the LLM labels, we manually assign each job to one category, once using the primary codebook and once using the *GPT-5.6* codebook. The manual labels agree with the first Step 2 validation pass on 176 of 200 jobs (88.0%) under the primary taxonomy and on 178 of 200 jobs (89.0%) under the *GPT-5.6* taxonomy. These checks indicate that LLM classifications align closely with human judgment under both codebooks.

```
AI Job Classification Task

You are an expert labor economist classifying LinkedIn resume entries. Each
row in the attached CSV or Excel file represents one job listed on a worker's
LinkedIn profile. title_raw is the worker's self-entered job title.
description is the worker's own description of their responsibilities in that
position. It is not an employer job advertisement. These jobs were selected
because they contain AI-related keywords in either the title or description.
However, not all of them represent workers whose primary responsibilities
involve AI development or implementation. Many only mention AI tools or
AI-related terminology. Your task is to distinguish these different types of
AI involvement.

Classification Objective

Classify each job according to the worker's primary role in the AI
production, deployment, governance, or use process. The goal is not to
determine whether a worker "works in AI," but rather how they contribute to
AI. The categories should distinguish between workers who: build AI systems,
support AI development, deploy or integrate AI, manage AI initiatives or
coordinate with AI workforce on AI implementation, govern AI, claim
implementation of AI tools in work processes but work in non-AI related
roles, use AI tools in non-AI work, or merely mention AI without meaningful
involvement.

Step 1: Construct a Taxonomy
```

Before classifying the jobs, develop a taxonomy consisting of 8-12 mutually exclusive and collectively exhaustive (MECE) categories. The taxonomy should be: comprehensive enough to cover all jobs, internally consistent, applicable across all uploaded files, based on the actual distribution of jobs in the data, stable so that similar jobs receive the same category. The categories should reflect the worker's function within the AI ecosystem, rather than industry, employer, seniority, or job title.

The following categories are intended as starting points and may be refined or reorganized if doing so improves consistency:

Direct AI model development, training, or engineering

Data engineering, data science, or annotation supporting AI model development

AI implementation, deployment, integration, or MLOps

AI product, project, or cross-functional program management and coordination with AI workforce on AI implementation

AI governance, ethics, compliance, documentation, education, or internal enablement

Business leaders or managers who oversee AI development and implementation without directly building models and who work in AI related roles

Business leaders or managers who oversee AI adoption without directly building models and work in non-AI related roles

Workers who use AI tools in their work but are not responsible for AI development or deployment

Incidental AI mention (AI keywords appear, but there is insufficient evidence that AI is a meaningful part of the job)

Step 2: Classify Every Job

Assign exactly one category to every row. The category should reflect the worker's primary responsibility, not every activity they mention. If multiple AI-related activities appear, classify according to the activity that appears most central to the role.

Evidence Hierarchy

Use the following evidence in order of priority:

Job description (highest priority)

Job title

AI-related keywords

If the description conflicts with the title, trust the description. Do not infer responsibilities that are not supported by the text.

Classification Principles

Judge the work actually performed.

Do not combine jobs with incidental AI mentions and those with insufficient evidence into one category. If there is insufficient evidence, try your best to assign it to a category, if you really cannot make informed decision, then put it in its own ‘‘insufficient evidence’’ category.

Ignore seniority when assigning categories.

Ignore company reputation.

Ignore formatting artifacts, HTML, encoding issues, or mojibake.

Do not assume technical AI expertise simply because AI tools or AI terminology are mentioned.

Mentioning ChatGPT, Copilot, LLMs, or similar tools alone does not imply that the worker develops or deploys AI systems.

Distinguish between:

creating AI systems,

supporting AI development,

deploying AI,

governing AI,

managing AI initiatives,

adopting AI tools within organization,

using AI as a productivity tool,

and merely mentioning AI.

```
If evidence is insufficient to distinguish between two categories, choose the
broader, less specific category rather than inferring technical AI
responsibilities.

If job description is missing or short, rely on job titles in the following
way: if job description or title has highly AI specific terms (e.g., ‘‘ML
engineer’’, ‘‘gen AI developer’’, ‘‘Deep learning researcher’’) then classify them
into one of the highly AI-relevant categories that is closest by nature of
the job title. Do not put these jobs into ‘‘insufficient evidence’’ jobs.
```

Table A.25 lists the functional categories and their definitions. The taxonomy covers a broad set of AI functions within a firm, while separating substantive roles in AI development, deployment, management, and governance from jobs with incidental AI mentions. Table A.26 reports the distribution of the 4,500 AI jobs across these categories. Only 336 jobs (7.5%) are classified as incidental or unsubstantiated AI mentions: after examining these jobs, they often contain legitimate mentions of AI-related keywords and some AI-related functions but are not focused exclusively on AI implementation within firms, or do not provide enough detail in job descriptions to classify them into one of the categories with high-enough confidence.

Most of the AI-related jobs are heavily involved in internal AI development, management, and support of firm AI efforts. About 45% of jobs are in core AI development: 650 jobs (14.4%) are classified as AI research and model development; 469 jobs (10.4%) as AI application and systems engineering; 273 jobs (6.1%) as AI platform, infrastructure, and MLOps; 531 jobs (11.8%) as domain analytics and predictive modeling; 51 jobs (1.1%) as AI data, annotation and evaluation support, and 1.4% of jobs provide governance of AI development. 22.7% of jobs are managing AI efforts: AI business leadership and functional oversight jobs (10.9% of AI jobs) provide executive, business-unit, or functional leadership over AI-related teams, portfolios, or investments; and AI product, program and delivery management jobs (11.8%) lead AI product strategy, program execution, or cross-functional delivery for AI products and features. A significant share of jobs help involve AI go-to-market, solutions and advisory (7.4%) that help sell and support AI products and services.

Finally, 17% of the jobs are general technical workers who do not seem to exclusively work on AI development and implementation but claim using AI technology in their jobs (e.g., “Won company-wide engineering contest classifying websites using *machine learning*”; “I am a Software architect with expertise in data science, cloud, *machine learning*, and cybersecurity leveraging languages C, C++, Python, Java”; “ Develop new business ideas in the Augmented Reality, Virtual Reality and *Artificial Intelligence* spaces”).

These patterns suggest that our resume-based AI measure primarily captures jobs

that contribute to firms' internal AI development and deployment as well as AI product management, support and monetization, rather than merely workers using external AI tools. When we force a model to have a separate classification category for *"Non-AI role using AI tools as productivity aid"* category, the model classifies only 0.9% of jobs into this category.

We also examine how these AI job categories are distributed within the tech and the non-tech sector firms. In Table A.27, we show that the share of AI jobs in AI management roles ("Management") is about the same in the two sub-samples. However, non-tech sector firms have significantly higher share of AI jobs in "Core AI development" (48% vs. 43% in tech firms), and also higher shares in "Other AI" jobs (27% vs. 22% in tech), while having much smaller share of AI workers in selling and supporting AI products and services (3% vs. 12% in tech). This is consistent with tech firms being more likely to use AI to build products and services to sell to others while non-tech firms investing in AI to build in house AI capabilities.

A.4. Organization Capital: Definitions, Comparison to Neighboring Concepts, and Classification

A.4.1. Definition and key characteristics

Organization Capital (OC) is durable, organization-specific intangible capability that improves how a firm combines inputs (people, tasks, routines, technologies, management systems, and physical assets), coordinates and governs activity, develops products, and brings output to market. It accumulates through learning (from experience, training, and management) and may be embodied in the organization itself or key employees. As firms learn how to do things better, they slowly accumulate organization capital and become more productive. Once built, the asset is non-rival within the firm: a documented procedure or a structured record serves many employees, units, and tasks at once without depletion, so organization capital scales internally at little additional cost. However, it travels poorly outside the firm: routines and systems tuned to one firm lose much of their value elsewhere, which is what makes the asset hard to separate, transfer, imitate quickly, or value from accounting data.

These five interrelated diagnostics differentiate organization capital from other concepts:

- **Built (vs purchased).** Organization capital assets are built up gradually from firms' learning-by-doing rather than bought on the market.
- **Durable (vs ephemeral).** OC represents repeatable useful knowledge rather than one-off outputs; it affects current and future production.

- **Organization-specific (vs general).** Its usefulness is highest in the firm in which it was accumulated: organization capital is less *portable* and *marketable* as some other intangible assets, such as patents. Organization specificity also tends to make organization capital relatively excludable.

- **Embedded (vs separable).** It resides in platforms, systems, routines, communication channels, culture, processes, workflows, and—to the extent it has not been written down—in workers themselves. This characteristic also differentiates organization capital from inputs into intangible capital building (e.g., R&D expense) and its outputs (e.g., patents).

- **Valuable (vs not).** Its use raises the productivity of inputs and market-facing activities by improving innovation, coordination, execution, production, and service. Knowledge that ceases to be useful—because the environment it describes changed—depreciates out of the stock of organization capital.

The existing literature often lists the following recurring families of OC examples, although they are often vague due to the difficulty of identifying and measuring OC in practice:

- **Information, learning, and matching:** personnel information, worker–task and worker–worker matches, task information, built-up plant expertise, and knowledge of how to use technology.

- **Human capital and key talent:** firm-specific skills, on-the-job training, informal learning, managers, engineers, product managers, software engineers, and other essential personnel.

- **Management and governance:** management practices, control systems, culture, relational contracts, leadership, capabilities, and management structures.

- **Processes and systems:** routines, internal communication, distribution, supply chains, systems (e.g., Enterprise Resource Planning, HR, Customer Relationship Management), platforms (e.g., cloud, enterprise application, developer and software platforms), knowledge sharing, and operational excellence.

A.4.2. Comparison to neighboring concepts

In terms of how to conceptually separate organization capital from related intangible investments and capital concepts, classify the durable OC stock separately from *inputs* (e.g., SG&A expenditures) or *efforts* to change something that creates OC and from *outputs* (product, patent, customer relationship) that OC enables. R&D spending, IT purchases,

sales and marketing, training, process redesign, and co-invention are expenditures, activities, and inputs. They may create several different outputs. Organization capital is the durable, firm-specific capability that persists after those activities are complete.

R&D and knowledge capital. R&D may denote an expenditure flow, a portfolio of projects, or the scientific and technical knowledge produced by those projects. In standard intangible-capital taxonomies, patents, formulas, designs, algorithms, experimental findings, and technical know-how are knowledge capital, whereas OC is the firm-specific system that selects, combines, coordinates, and deploys that knowledge (Corrado et al., 2009; Peters and Taylor, 2017; Eisfeldt et al., 2025a; Charoenwong et al., 2026). The two stocks are nevertheless complementary: both are cumulative, path-dependent, hard to observe, and often embodied in skilled labor. R&D can also build an organizational learning capability thus creating OC (Cohen and Levinthal 1989). The practical boundary is the locus and portability of value. A patent, molecule, design, or licensable code base is normally knowledge capital that is portable across firms; research-portfolio governance, cross-functional product-development routines, scientific communication networks, and architectural competence are OC when their value depends on the firm's configuration. Note also that OC need not be novel, whereas knowledge capital typically is. R&D and OC are therefore adjacent and complementary, not synonymous.

Information technology. The same distinction applies to IT. Hardware, software, databases, networks, and cloud infrastructure are technical or digital capital; OC is the complementary organizational asset that determines how effectively they are used. Productivity gains arise when technology is combined with decentralization, skilled labor, training, redesigned work practices, and complementary organizational investment (Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002; Brynjolfsson et al., 2002a). A packaged enterprise system is not OC merely because it is intangible: it is bought rather than accumulated, and the same system runs at any firm that licenses it, so it fails both the accrued and the organization-specific diagnostics. Custom code and proprietary data are the harder case, and the storage-medium criterion resolves it: a licensable code base is portable digital capital, whereas a data asset assembled from the firm's own operations, or an architecture that encodes its particular workflows, has little value elsewhere and is OC. Unambiguously OC, finally, are the authority structure, role definitions, exception-handling routines, training, and cross-unit coordination that make any of these systems productive. The fact that high-OC firms spend more on IT therefore indicates complementarity rather than identity (Eisfeldt and Papanikolaou, 2013a). Even when the two assets are co-specialized, empirical work should measure them separately.

Absorptive capacity. Where IT is an adjacent technical asset, absorptive capacity is closer to an internal organizational capability. Cohen and Levinthal (1990) define it as the ability to recognize the value of external information, assimilate it, and apply it commercially. It is path-dependent, rests on prior related knowledge, and can be built through R&D, production experience, technical training, and organizational interfaces. In pharmaceuticals, effective absorption also depends on incentives to publish and on connections to the wider scientific community, not only on the volume of in-house research (Cockburn and Henderson, 1998). These features make absorptive capacity a durable, firm-specific capability embodied in scientists, communication channels, incentives, shared language, and the architecture linking internal and external knowledge. Its scope is nonetheless narrower than OC: OC also includes worker matching, production routines, culture, supply-chain coordination, governance, and other capabilities unrelated to external learning. Absorptive capacity is therefore best treated as a subtype or dimension of OC.

Co-invention. This stock-versus-process distinction is sharper for co-invention. Bresnahan and Greenstein (1996) use the term for the user-side experimentation, discovery, and complementary redesign needed to make a general technology economically valuable. Adopters alter applications, workflows, skills, and organizational arrangements, so co-invention often builds routines, process knowledge, communication patterns, and governance that are slow to accumulate and hard to imitate. Yet co-invention is not itself a capital category. A project may yield custom software or data, a redesigned service, a customer interface, and persistent coordination routines; of these, the persistent coordination routines are unambiguously OC, and the data may be as well, depending on whether it was assembled from and used in the firm's own operations. Co-invention is therefore to OC as investment is to capital: an important formation mechanism, not the stock itself.

Process and product innovation. The same logic organizes process and product innovation. A business process innovation is a new or significantly improved process that has been brought into use, such as a production method, logistics system, administrative process, or information flow (OECD and Eurostat, 2018). Process innovation is therefore an implemented event, while OC is a durable capability. The overlap is substantial because new work design, cross-functional production, internal communication, supply-chain management, and process improvement can become OC once they are learned, stabilized, and embedded in routines, roles, and tacit coordination. Process innovations are also often difficult to sell separately and generate returns tied to the firm's own scale (Cohen and Klepper, 1996). Still, a process change may reside mainly in new equipment or in a purchased software package. The converse also holds: organization capital need not be innovative. An established routine that no competitor would find novel can be highly

valuable to the firm that operates it, and organization capital can arise from the codification of practice a firm already followed. Novelty and durable firm-specific capability are separate properties; OC requires the second, not the first. OC is thus the persistent organization-specific stock that some (but not all) process innovations create—and that also accumulates without any innovation event at all.

A product innovation is a new or significantly improved good or service introduced to the market (OECD and Eurostat, 2018). Its immediate object is the offering: the underlying design, patent, formula, or software is normally knowledge or digital capital, and market introduction is an innovation event. OC nevertheless enables product innovation through repeatable product-development systems, cross-functional coordination, launch routines, quality systems, platform governance, and the ability to scale reliable delivery. The clean boundary separates the artifact from the capability: the product or patent is not normally OC, but the system for repeatedly generating, integrating, launching, and scaling products can be.

Thus, OC can support both process and product innovation even though its own locus is closer to the persistent organizational capability than to the innovation event itself.

Architectural competence and architectural innovation. Finally, architectural competence and architectural innovation show why similarly named concepts can occupy different categories. Henderson and Cockburn (1994) distinguish local component competence from architectural competence in pharmaceutical research. Architectural competence integrates disciplinary knowledge, coordinates across programs, filters information, and develops new competencies; it is embodied in communication channels, managerial systems, incentives, norms, and the organization of research. It is collective, firm-specific, durable, and productive through integration rather than through a separable piece of technical knowledge. The best classification is therefore an R&D-specific subtype of OC, with a narrower domain than the firm-wide concept.

Architectural innovation is a different object. It is a product-system innovation that changes linkages among components while leaving core component concepts largely intact (Henderson and Clark, 1990). Established firms can fail because their communication channels and information filters embody the old architecture. The innovation event is not OC; rather, it reveals the value and rigidity of the architectural knowledge already embedded in the organization. It devalues existing OC and requires the accumulation of new OC.

Customer capital. A different boundary separates internal supply-side capabilities from demand-side assets. Customer capital is a durable intangible asset built through sales and marketing, customer service, brand-building, customer information, acquisition, retention,

and related strategies (He et al., 2024). Like OC, it is firm-specific, built gradually through people and routines, frequently expensed rather than capitalized, important for scale, and non-rival in the sense that it can serve many transactions at once. Broad definitions of OC may therefore include customer-facing capabilities that are tightly integrated with the business model. Under a narrower taxonomy, however, OC concerns how the firm coordinates, innovates, and produces, whereas customer capital concerns who buys, why customers return, and the value of relationships, reputation, and brand. A customer-service routine may be OC, while the relationships and reputation those systems help build is customer capital. This internal-versus-external distinction is also central to efforts to separate OC from brand and customer assets within SG&A (Chawla et al., 2025; Eisfeldt et al., 2025a; Charoenwong et al., 2026).

Summary classification. Table A1 applies the same rule across all neighboring concepts. The decisive questions are: What is the durable stock? Where does its value reside? How portable is it outside the firm? Expenditures, redesign activities, and innovation events should be classified separately from the persistent capabilities they create.

Table A1. Organization capital and neighboring concepts

Concept	**Relation to OC**	**Key boundary**
R&D / knowledge capital	*Adjacent stock*	Scientific and technical knowledge is knowledge capital. The organization of R&D—interfaces, incentives, cross-functional routines, and architectural competence—is OC when its value depends on the firm's configuration.
IT / digital capital	*Adjacent stock*	Hardware, packaged software, and networks are technical assets, as is portable code. Data and architecture built from the firm's own operations, and the work practices, authority, training, and process redesign that make IT productive, are OC.
Absorptive capacity	*OC subtype*	A path-dependent capability for identifying, assimilating, and using external knowledge. It fits within OC but is narrower than the firm-wide concept.
Co-invention	*Formation process*	User-side experimentation and redesign may create OC, digital capital, customer capital, or product outputs. Co-invention is a formation activity, not the resulting stock.
Process innovation	*Event → possible OC*	A new process is an innovation event. When it becomes a persistent, firm-specific routine or coordination system, the resulting stock is OC.

Continued on next page

Table A1 (continued)

Concept	Relation to OC	Key boundary
Product innovation	*Output enabled by OC*	The product, patent, or design is not normally OC. Repeatable product-development, launch, quality, and scaling capabilities can be OC.
Architectural competence	*OC subtype*	Integrative research competence is a close, domain-specific example of OC.
Architectural innovation	*Event*	Reconfiguring product linkages is an innovation event. It can devalue old OC and require new OC, but it is not the stock itself.
Customer capital	*Adjacent stock*	Customer relationships, brand, service, and retention assets are demand-side capital. The internal systems used to build them may be OC.

A.4.3. Classification of organization-capital-related jobs

We propose a new measure of organization capital using workers' job descriptions based on whether an individual job builds organization capital. The key advantage of this novel job-description measure is that it observes worker activities that build organization capital at the worker-job level—for example, codifying firm-specific knowledge into reusable systems, routines, and processes—rather than inferring organization capital from aggregate expenditures. This makes it substantially more informative than capitalized-SG&A measures, which capture an input into organization capital but cannot distinguish spending that actually builds durable organizational capabilities from unrelated SG&A spending. Relative to measures based on employee reviews of their employers—which captures employees' perceptions of the realized stock of culture, coordination, and management quality—the job-level measure identifies where and how organization capital is produced, allowing researchers to attribute its accumulation to particular workers, tasks, and technologies such as AI.

To develop our methodology for identifying organization-capital-building jobs, each coauthor independently built their own prompt and used different LLMs to classify jobs using the same 2,000 random sample of jobs. After comparing the outputs and the manual review of classifications, we converged on a prompt (copied below) that offered the highest accuracy and generalization in classification.

Our prompting approach has the following desirable features. First, it is agnostic on whether a job is AI-related or not: the prompt does not contain any mentions of AI or how AI jobs might be more or less likely to be building organization capital. This helps to *not* bias the results on classifying jobs into organization-capital-building based on how related they are to AI. Second, our definition of organization capital in the prompt comes from the

core seminal papers (Prescott and Visscher 1980; Atkeson and Kehoe 2005; Eisfeldt and Papanikolaou 2013a): this allows for a high-level conceptual but precise-enough definition of key features of organization capital outlined in Part A.4.1 above.We find this approach attractive since it does not rely on rough proxies (e.g., occupations, seniority, job titles, industry) for what types of jobs should be related to organization capital: a manager can oversee a few workers without building any organization capital or, instead, can lead and coordinate across teams to build systems to improve how firms operate—a good example of building organization capital. Instead, the model reads each job's description and title and identifies whether the job has the key features of building organization capital. Finally, we test our final prompt across LLMs and find very consistent outputs across multiple runs (on the order of around 95% agreements across different models).

To classify jobs as organization-capital-building or not, we first build two samples of jobs from the Revelio Labs resume data. The AI sample includes *all* 593,332 AI-related jobs in our public-firm sample with non-empty job descriptions (about 91% of all AI-related jobs). For cost-efficiency, the non-AI sample consists of a representative random sample of 15,000 non-AI jobs with non-empty job descriptions, drawn from the same firm-years as the AI jobs (1,000 jobs per year during 2010–2024). Table A.28 compares the 15,000-job random sample to the full population of non-AI jobs in terms of industry (Panel A) and firm size (Panel B) distribution. In Panel A, manufacturing accounts for 27.02% of all non-AI jobs (and 27.83% in the random sample), finance and insurance accounts for 17.67% of the population (and 16.14% of the sample), and retail trade for 14.64% (14.60%). Overall, across industries, the population and the random sample employment shares line up closely. In Panel B, the median firm has 23,457 employees in the population and 24,207 in the random sample, and the respective means are 52,280 and 53,728. In the non-AI job sample, firm size distribution matches almost exactly in the random sample and the whole population.

To reduce variability in the LLM's output and improve reproducibility, we set the temperature to zero. In both samples (AI and non-AI jobs), we score each job with *GPT-5.4* that reads each job title and description and returns an organization-capital score between 0 and 1. We classify a job as related to building organization capital ("org-related") if its score is at least 0.6, where 0.6 is the boundary between mixed evidence (0.4–0.6) and substantial evidence (0.6–0.85) on building OC. The full prompt is as below:

```
DEFINITION - ORGANIZATION CAPITAL
```

Organization capital is a durable, intangible, and firm-specific asset. Organization capital is knowledge specific to a firm that helps it convert inputs (e.g., labor, physical capital, data) into output more efficiently. Firms accumulate this knowledge through experience and learning-by-doing from their operations, and as they do, this knowledge helps them to operate more efficiently and become more productive. The faster the learning process, the more quickly firms accumulate organization capital and realize productivity gains. Organization capital is partly embodied in the knowledge of key employees and built by them. It is distinct from physical capital and from general, portable human capital (skills a worker takes to any employer). In sum, organization capital is firm-specific knowledge that is embodied in the organization itself (in its routines, practices, systems, processes, employees, databases), it is acquired through learning-by-doing and partially built by its employees.

STAGE 1 - SCORE INDIVIDUAL JOBS

You are an expert labor economist classifying LinkedIn resume entries. Each input is ONE position a worker listed on their LinkedIn profile: title_raw is the worker's raw, self-entered job title, and description is the worker's own free-text account of what they did in that role (not an employer ad).

TASK

Decide whether the work contributes to building the employer's ORGANIZATION CAPITAL and rate how strongly on a 0-1 scale, using the definition above.

HOW TO APPLY THE DEFINITION

1. Map the described activities onto the definition; a job scores high only to the extent its work builds firm-specific organization capital (accumulating operational/production knowledge through experience, codifying it in routines/systems/databases, or developing key employees' firm-specific knowledge).

2. General or portable skills, one-off individual output (a single analysis, model, design, or code change), or routine operation of existing assets without building firm-specific knowledge score low.

3. Rely on the definition rather than your own broader intuitions.

SCORING RUBRIC (0-1)

0.00-0.15 Unrelated.

0.16-0.40 Incidental.

0.41-0.60 Mixed.

0.61-0.85 Substantial (building firm-specific organizational capability is central).

0.86-1.00 Core (role exists primarily to build firm-specific organization capital).

OUTPUT

Return ONLY JSON with org_score (number 0.00-1.00). Judge the WORK described, not the title alone. Ignore mojibake/formatting artifacts. Do not invent activities not supported by the text.

STAGE 2A - NAME MECHANISMS FOR JOBS THAT BUILD ORGANIZATION CAPITAL

Below are JOB DESCRIPTIONS grouped by clustering. They share a common mechanism for HOW the work BUILDS the employer's ORGANIZATION CAPITAL.

Using the definition above, write ONE sentence naming the shared ECONOMIC MECHANISM by which these jobs build firm-specific organization capital. Examples include codifying firm-specific knowledge into systems/routines, developing key employees' firm-specific knowledge, or building reusable firm data assets.

Describe the reason, NOT the occupation or industry.

Return ONLY JSON: {"mechanism":<one sentence>}.

STAGE 2B - NAME MECHANISMS FOR JOBS THAT DO NOT BUILD ORGANIZATION CAPITAL

Below are JOB DESCRIPTIONS grouped by clustering. They share a common reason WHY the work does NOT build the employer's ORGANIZATION CAPITAL.

Using the definition above, write ONE sentence naming the shared ECONOMIC REASON these jobs do NOT build firm-specific organization capital. Examples include general/portable skills, one-off individual output, or routine operation of existing assets.

Describe the reason, NOT the occupation or industry.

Return ONLY JSON: {"mechanism":<one sentence>}.

STAGE 3 - STANDARDIZE MECHANISM TAXONOMY

You are standardizing a taxonomy of economic mechanisms for ORGANIZATION CAPITAL.

```
Below are candidate mechanism statements, each summarizing a cluster of real
jobs with job counts. Merge them into between 20 and 40 DISTINCT standardized
categories. Each category is ONE full sentence stating the ECONOMIC REASON
the work [builds / does not build] firm-specific organization capital.

RULES

Merge statements that express the SAME economic reason even if they come from
different industries or occupations.

Do NOT base categories on occupation or industry.

Do NOT create a catch-all "other" or "miscellaneous" category.

If one reason would dominate the data, SPLIT it into distinct sub-mechanisms.

Ground every category in the definition of organization capital above.

Return ONLY JSON: {"categories":[{"category_id":"[prefix]1","label":<one
sentence>}, ...]}.
```

One potential concern is that our organization-capital classification simply pick up workers who write higher-quality job descriptions, but we do not find that to be the case. To assess this, we construct some simple writing-quality measures. Organization-capital jobs are about twice as long, but they use a more repetitive vocabulary (78 distinct words among the first 100, versus 84), cite fewer specific numbers and tools (7.9 versus 9.3 mentions per 100 words), and contain spelling errors at the same rate as other AI job descriptions (7.1 versus 6.5 errors per 100 words). When comparing descriptions of similar length, the two groups are essentially indistinguishable across all quality measures.

We find that AI jobs are more likely to build organization capital than non-AI jobs. Table A.29 shows that 58% of AI jobs are org-related versus 31% of non-AI jobs, and AI jobs have a higher mean organization capital score of 0.60 versus 0.42 for non-AI jobs. This pattern holds across the whole score distribution—most of the AI jobs fall in the "substantial" band and most of the non-AI jobs in the "incidental" band.

We next look at how jobs that build organization capital differ from those that do not across several dimensions. First, we examine word clouds for job descriptions of jobs categorized as organization-capital-building in non-AI-related jobs and in AI-related jobs in Figure A.22. Among AI-related jobs, intuitively both types of jobs have *data* and *ai* as key words, but organization-capital-related positions emphasize *team*, *lead*, *product*, *business*, *design* and *process*, combining language about leading teams and building busi-

ness products and processes, whereas non-organization-capital-related jobs are led by *model*, *learning*, and *develop*, emphasizing technical model development and deployment. Among non-AI-related jobs, organization-capital-related positions are led by *team*, *process*, *lead*, *management*, and *business*, highlighting operational and managerial routines, whereas non-organization-capital-related jobs are led by *customer*, *service*, *sale*, *provide*, *team*, *product*, *ensure*, emphasizing customer-facing service, sales, and support work. Comparing organization-capital-building jobs across the AI and non-AI divide, both groups emphasize *team*,*lead*, and *business* as well as *product* and *process*, but AI-related positions also emphasize *data*, *ai*, and *model*, whereas non-AI-related positions stress *management* and *project*.

Second, we find that non-tech firms have significantly higher shares of AI jobs that build organization capital: 65.1% versus 50.6% in tech firms (defined as firms in NAICS 2-digit sector 51 or 54, or in semiconductor manufacturing). These differences are consistent with the evidence in Appendix Table A.27 that, compared to tech firms, non-tech firms are more likely to employ AI-skilled workers in core AI development and less likely to have AI jobs selling and consulting on AI products and services: tech firms tend to deploy AI talent to build products for others, while non-tech firms use it to build capabilities in house.

Third, consistent with prior work that argues that workers in management roles are likely to build organization capital (Eisfeldt and Papanikolaou, 2013a; Dessein and Prat, 2022), we find that 79% of AI-related workers in management roles are classified into organization-capital-building roles for all firms (Panel A of Table A.31), and this ratio is higher for non-tech firms (86%, Panel B) than for tech firms (73%, Panel C). Consistent with the idea that customer service roles are more likely to serve customers and are less likely to build durable organizational capabilities, only 22.6% of AI workers in "Support for AI products/services" build organization capital. Interestingly, among the group of workers who are more heavily involved in core AI development, about half are tagged as organization-capital-building (50% for non-tech and 43% in tech firms), consistent with the discussion above that tech-facing roles do not automatically build organization capital.

We next examine the time-series dynamics. In Figure 7, we find that the majority of AI jobs build organization capital, and this share rises from about one half in 2010 to around two thirds in 2024 (panel (a)), both within the tech sector and outside of it (panel (b)). In contrast, only about one third of non-AI jobs build organization capital (panel (c)), and this share is flat over the entire period in both sub-sectors (panel (d)). Appendix Figure A.23 shows that the same pattern holds using the average continuous organization-capital score, confirming that the trend is not an artifact of the classification threshold. The stability of the time series of non-AI organization-capital-building jobs indicates that the rise of

org-related share of AI jobs reflects a genuine change in the nature of AI work rather than a more general trend of jobs becoming more organization-capital-building or a drift in job descriptions of AI-related jobs. Consistent with the changing nature of AI work over time, we find that 57% of AI jobs in the earlier Wave I categories (ML, NLP, and CV) build organization capital, compared with 66% of jobs in the more recent Wave II categories (generative and agentic AI). More recent vintages of AI technologies appear increasingly oriented toward building organization capital, with the latest vintage—agentic AI—having the highest share of organization-capital-building jobs (71%) among all AI vintages.

We then move beyond whether a job builds organization capital to *how* it does so: two jobs with identical organization-capital scores may build very different kinds of firm-specific assets—one codifying tacit knowledge into systems and routines, another accumulating reusable data assets, a third building expertise in product innovation. Characterizing these mechanisms lets us ask not only whether AI work is more organization-capital-intensive than other work, but whether it builds organization capital through different channels. We therefore assign each job a standardized reason category alongside its organization-capital score. For org-related jobs, the category records the mechanism through which the work builds firm-specific organization capital; for non-org-related jobs, it records why the work does not — because the skills are general and portable, because the output is one-off and individual, or because the role operates existing assets rather than creating new ones. The LLM first induces a standardized set of mechanisms from the data and then assigns each job to its nearest category based on the job description. Appendix Table A.30 provides detailed definitions of each reason category.

Table 2 shows how AI and non-AI jobs are distributed across the reason categories within each group—organization-capital-related and non-organization-capital-related. Panel A shows that all organization-capital-related job categories involve building systems, platforms, workflows, routines, and firm-specific knowledge and models. These categories also provide a sense of the types of assets created through organization-capital-building jobs. Among these categories, AI jobs are concentrated in codifying data, models, and proprietary knowledge into reusable technical systems: data platforms (9.1% of all AI jobs), proprietary knowledge generative-AI/NLP (6.9%), reusable AI/ML platforms (4.8%), and firm-specific predictive and decision models (2.3%). In contrast, non-AI-related jobs are more operational in nature, concentrating in hiring and talent development (5.8%) and manufacturing & process systems (3.5%). Both AI and non-AI jobs rank cross-functional operating routines among their top five categories (6.3% and 4.8%, respectively), highlighting the importance of within-firm coordination as a source of organization capital accumulation.

It is worth emphasizing the diversity of the types of organization capital built by both

AI and non-AI workers. For example, AI workers build product and design systems (4.7%), marketing and revenue systems (4.1%), cloud, infrastructure and reliability operations (3.4%), financial and risk decision systems (3%), customer support systems (2%), security, compliance and risk systems (1.9%), R&D and scientific platforms (1.4%), and other operating processes. These activities represent the creation of durable organizational capabilities rather than one-off analytics, portable algorithms, or standardized routine work. The diversity of organization capital assets built by workers also underscores the difficulty of measuring organization capital using coarse signals like occupations, job seniority, or departments, since workers across each of these categories can contribute to its accumulation.

In Panel B, which covers jobs that do not build organization capital, AI jobs are concentrated in portable analytics and one-off models (12.6% of all AI jobs), broadly applicable research and algorithms (8.9%), and bespoke client consulting and implementation (7.8%). By contrast, non-AI non-org-jobs are overwhelmingly concentrated in standardized operations and support (37.9%), bespoke client consulting and implementation (8.8%), and transactional sales, marketing and account management (7.6%).

Appendix Table A.32 provides examples of AI-related job descriptions in our sample that do and do not build organization capital, together with the reason for each classification. Organization-capital-related AI jobs build systems the firm continue to use, such as data platforms, evaluation frameworks, and copilots built on internal workflows. Other AI jobs apply similar technical skills to one-time analyses, general research, client work, and products and sales. These examples illustrate that the LLM classifies an AI-related job as organization-capital-building only when there is *sufficient evidence* that the work results in *durable* organizational capabilities—it does not assume that AI work is organization-capital-building by default. For example, the following jobs are—correctly—not classified as organization-capital-building. First, a job titled "Research Scientist" is "responsible for developing advanced and reliable industrial *software* for automation applications such as smart manufacturing, industrial robots and cyber physical systems." Although the job involves building software, as we discuss above in Part A.4.2, software itself does not automatically qualify as organization capital. Second, a "Software Engineer" uses computer vision and "designed and implemented a *method* of hiding and retrieving imperceptible patterns in images displayed on computer monitors and projection displays." Here, the method is an output of the research process and may or may not ultimately be implemented in the firm's systems or workflows.

Appendix Table A.33 provides analogous examples for non-AI-related jobs. Organization-capital-related non-AI jobs build the systems and practices through which the firm operates, including standardizing procedures, codifying safety and risk programs, training

employees, and coordinating and collaborating across teams. In contrast, other non-AI jobs perform day-to-day tasks within existing systems or use skills that workers can readily transfer to other employers. The examples also highlight differences in the types of organization capital built across the two samples. Organization-capital-related AI jobs primarily codify the firm's data, models, and proprietary knowledge into technical systems, whereas organization-capital-related non-AI jobs more often build operational practices such as procedures, quality programs, and training, consistent with the distribution of reason categories in Appendix Table A.27.

A.4.4. Case studies of AI building organization capital

Organization capital is difficult to observe directly, which is why it is usually inferred from spending or from residuals; the case studies below take the opposite approach, using public accounts and detailed job descriptions of workers to show what building it actually looks like inside particular companies. The eight firms below span an unusually wide range of industries—investment banking and wealth management, consumer packaged goods, biotech, home improvement retail, integrated healthcare and pharmacy, telecommunications, athletic footwear and apparel, and online travel—and an equally wide range of what AI is actually doing inside them. Morgan Stanley codifies advisor and underwriting judgment out of research documents and client meetings; Procter & Gamble builds simulatable models of its plants from process telemetry; Moderna converts a short experimental history into predictive models while its organizational routines are still forming; Home Depot puts the expertise of store associates into a tool available in every aisle; CVS Health makes usable a claims and clinical record it already possessed; Verizon compresses the gap between its best and newest care representatives and field technicians; Nike turns costing and design judgment into tools and extends the record across a network of factories; and Expedia codifies how its own engineers diagnose incidents and ship software. The knowledge being converted is variously tacit clinical, commercial, mechanical, and engineering judgment, and the systems built to hold it range from digital twins and dispatch platforms to retrieval systems and standard operating procedures.

What these cases share is more striking than what separates them. In every instance the raw material is the firm's own operational record—the documents, transcripts, telemetry, claims, tickets, and transactions that accumulate as a by-product of doing business—rather than anything purchased or generally available; the same models trained on public data would build nothing firm-specific. In every instance the conversion is performed by the firm's own employees: the internal AI roles our measure captures are the labor that turns that record into a durable asset, and their descriptions show the work being done in the categories our classification identifies. And in every instance the resulting asset is non-rival within the firm and of limited use outside it—once the Assistant, the twin, or

the knowledge base exists, the cost of the next employee using it is close to zero, while a competitor could not use it at all. The industries and applications differ, but the underlying mechanism does not. Table A.1 summarizes which organization-capital categories each case illustrates, and the coverage is wide: the eight firms together span seventeen distinct categories, nearly the full classification. Reusable data platforms and pipelines is the only category present at all eight—every one of them, whatever its industry, is converting its own operational record into a shared and reusable asset—and cross-functional operating routines is next at six.

Table A.1. Organization-Capital Categories Represented in Each Case Study

	Category	Morgan Stanley	P&G	Moderna	Home Depot	CVS Health	Verizon	Nike	Expedia
O1	Reusable data platforms & pipelines	•	•	•	•	•	•	•	•
O2	Reusable AI/ML platforms & MLOps							•	•
O3	Firm-specific predictive & decision models				•				•
O4	Proprietary-knowledge GenAI/NLP systems	•	•	•		•			
O5	Reusable software platforms & workflows	•			•	•			
O6	Cross-functional operating routines	•		•	•	•	•	•	
O7	Reusable cloud, infrastructure & reliability operations						•		•
O8	Security, compliance & risk systems					•			
O9	Marketing & revenue-operations systems		•		•				•
O10	Customer-service & support systems						•		•
O11	Product & design systems							•	
O12	Manufacturing & process systems		•	•					
O13	Supply-chain & logistics systems		•					•	
O14	R&D & scientific platforms			•					
O15	Financial & risk decision systems	•							
O16	Hiring, training & talent development			•					
O18	Knowledge organization & ontologies						•		

Notes: A filled circle indicates that the category is illustrated in that firm's case study. The categories shown are, with few exceptions, among the five largest for the firm. Category codes and definitions follow Table 2.

Morgan Stanley. Morgan Stanley offers an unusually well-documented illustration of how AI investment builds organization capital. The premise is that AI lets a firm convert the operational knowledge generated by its daily work—what its people know how to do, recorded in documents, transcripts, code, and client interactions—into a codified stock the firm owns and can redeploy. Morgan Stanley's head of firmwide AI, Jeff McMillan,

has described five basic enterprise use cases for generative AI: the software development lifecycle, document management, translation, document creation, and marketing. Four of the five are exactly this conversion. Each takes a record of work the firm was already producing—code and design discussions, documents in digital and paper form, presentations and client materials, marketing content—and turns it into a reusable asset.

The advisor tools make the conversion concrete. The firm's AI Assistant, built on a corpus of roughly 100,000 proprietary research documents, reached over 98% adoption among financial-advisor teams, with document retrieval efficiency rising from 20% to 80% (OpenAI, 2024). The Debrief tool extends the same logic to a record that previously left no trace at all: with client consent it takes notes during meetings, summarizes key points, and writes a note into the firm's CRM. Advisor knowledge that had lived in individual memory and ad hoc notes becomes a firm asset—the transformation Cullen et al. (2026) describe. It is also why the firm's own framing of its competitive position has shifted toward proprietary data and business logic as the foundation of what it sells (Son, 2026).

The AI jobs we observe at Morgan Stanley are the labor that performs this conversion. Our classification groups them not by application, as McMillan does, but by the kind of asset the work produces, so the four conversion-oriented use cases above reappear below distributed across those categories. The descriptions cluster into the categories listed in Table 2 (and described in detail in Appendix Table A.30 that our classification identifies as most common at the firm:

- *Proprietary-knowledge GenAI/NLP systems—codifying documents and judgment rules into retrieval and generative systems.* The most visible body of work codifies the firm's documents and judgment rules into retrieval and generative systems. A GenAI engineer describes delivering "an enterprise GenAI Research & Policy Copilot (RAG + agent workflows) used by 600+ analysts/associates, handling 12K+ queries/week, reducing document search time by 35-45% across Wealth/Risk/Compliance teams." An executive director leads "the firmwide modernization of enterprise search," integrating "key platforms (e.g., HR, Social, Confluence) into Fusion and GenAI to deliver a unified knowledge experience." A senior business analyst serves as product owner for "AI Agent Assist call note summarization"—the Debrief capability, staffed—and an analyst describes "conversationalizing and orchestrating tools used by ~16,000 Financial Advisors into the AI @ Morgan Stanley Assistant chatbot."

- *Reusable software platforms & workflows—standardizing execution and reducing reliance on individual employees.* The largest category turns tacit process and exception-handling knowledge into platforms other teams build on. A vice president describes an "AIMLOps Platform … redefining the firm's approach to AI/ML operationalization,

including integration of Generative AI (GAI) and Large Language Models," with features for "agility, automation, scalability, reproducibility, CI/CD, monitoring," reporting a reduction in time-to-market of more than 35%. Another describes an internal GenAI platform serving "80,000+ Morgan Stanley employees." A head of client service change describes "delivering a knowledge management platform leveraging AI and AEM, re-architecting service analytics capabilities and implementing a CRM Agent Desktop solution using Salesforce."

- *Reusable data platforms & pipelines—unifying the records the rest depends on.* Codified data definitions and integration logic are the substrate the rest depends on. A senior data engineer describes designing "a cloud-native AI-powered investment management and regulatory data platform" that "unified trading, market, and risk data from Bloomberg, Reuters, ESG feeds, and OLTP systems." The Research & Policy Copilot above rests on an "end-to-end RAG ingestion pipeline" processing "2.5M+ documents (PDFs, policies, research notes)," which raised top-three retrieval precision from 62% to 82%.

- *Financial & risk decision systems—turning underwriting and control knowledge into standing decision routines.* Underwriting, fraud, and control knowledge is embedded into standing decision routines. A vice president in mortgage operations describes architecting "a GenAI-powered agent to analyze call transcripts and flag unaddressed items for mortgage Operations" and deploying "an LLM-based mortgage underwriting assistant to standardize underwriting memos, projected to save over 10,000 hours annually"—judgment that had resided with individual staff becoming a firm-level standard.

- *Cross-functional operating routines—institutionalizing how complex work gets coordinated.* A smaller but persistent category. A vice president in product and technical program management describes leading "finance technology transformation, defining multi-year product vision, roadmap, and OKRs to modernize accounting platforms and streamline Wealth Management operations," delivering "an enterprise sub-ledger that reduced error rates by 40%."

Figure A.1 places these examples in the context of the firm's hiring. Panel (a) shows that both organization-capital-related AI roles and other AI roles have grown steadily as a share of Morgan Stanley's US positions, with the growth accelerating after 2021; throughout the period the organization-capital-related series runs roughly two to three times the size of the other, so the firm's AI hiring has been tilted toward work that converts proprietary operations into durable assets rather than toward generic technical capability,

and that tilt has widened rather than narrowed as the level has risen. Panel (b) breaks that series into the reason categories in Table 2, showing the firm's five largest categories and grouping all remaining categories into "other." Reusable software platforms and workflows, together with reusable data platforms and pipelines, are present throughout and account for much of the level—they are the infrastructure on which anything else is codified—while financial and risk decision systems form a steady third component, consistent with the underwriting and compliance examples above. The most striking feature is the category of systems built on proprietary knowledge using generative AI and natural language processing, essentially absent before the early 2020s and a substantial share by the end of the sample. That timing marks a change in what could be codified at all: retrieval over the firm's research library, summarization of advisor meetings, and standardization of written judgment were not feasible at scale under the earlier technology.

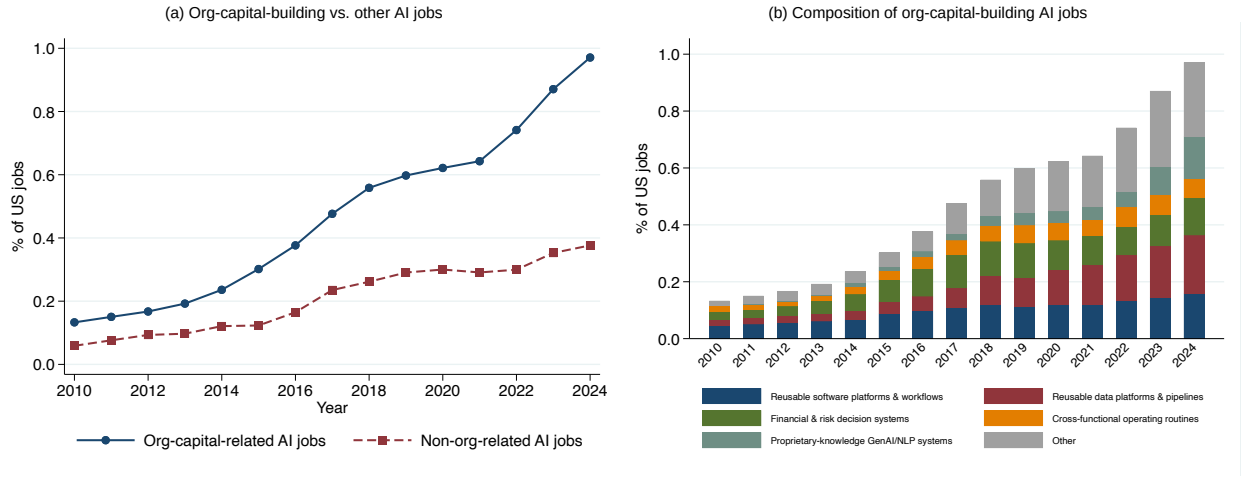


Figure A.1. AI jobs at Morgan Stanley as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification.

These uses of AI raise measured productivity through two channels. The first is that the same inputs yield more output: the research library, the advisors, and the client relationships are unchanged, but a far larger share of what the firm already knows is actually reachable at the moment it is needed—retrieval efficiency rising from 20% to 80%, or search time falling by 35 to 45%, is that change in mechanical form. Codification also raises the floor rather than the ceiling, since a standardized underwriting memo makes the median underwriter perform closer to the best one, which is where the projected saving of over 10,000 hours comes from. The second channel is that the resulting asset is non-rival within the firm: once the Assistant exists, the cost of the sixteen-thousandth advisor using it is essentially zero. That is why these investments show up as firm-level productivity

rather than as isolated wins, and it is also why the returns accrue to the firm rather than to the individuals whose expertise was codified.

Two features of the case are worth emphasizing. First, the codified stock is built out of operations rather than purchased: the assets being created are indexes of the firm's own research library, models trained on its own call transcripts, and workflow rules extracted from its own underwriting and origination processes. A firm cannot buy this stock; it can only operate and record. Second, not every AI use case builds it. Translation is one of McMillan's five but is generic—it converts nothing firm-specific into a durable asset—and consistent with this, a job in our data whose description centers on portable machine-learning methods applied to project-specific outputs, including machine translation, is classified as not organization-capital-related. The distinction between AI that raises a firm's capability in general and AI that codifies what a firm already does is one our classification can and does make.

Procter & Gamble. Procter & Gamble illustrates the same logic in a very different setting. A consumer goods manufacturer's accumulated advantage lies in how it runs plants, moves product, and learns what consumers want—knowledge that has historically lived in the experience of process engineers, planners, and brand managers. AI lets the firm convert that operational knowledge into a codified stock it can query, simulate against, and redeploy across sites.

The clearest instance is the digital twin. Rather than halting a production line to trial a process change, P&G runs the trial against a model of the site built from its own telemetry—line speeds, cycle times, quality data—joined to R&D and demand data. The World Economic Forum's Global Lighthouse Network reports that at P&G's Takasaki plant, which deployed data flow integration, digital twins, and machine learning across the value chain from R&D to customers, innovation lead time accelerated by 72% and shutdown days for trials fell by 21% (World Economic Forum, 2023); at its Lima site, digital twins and advanced analytics helped bring new products to market ten times faster (World Economic Forum, 2021). What makes this organization capital rather than a purchased tool is that the twin is only as good as the plant history it encodes. It is built out of operations and is worth little to anyone else.

The AI jobs we observe at P&G are the labor that performs this conversion, and they cluster into the categories our classification identifies as most common at the firm:

- *Manufacturing & process systems—turning plant experience into simulatable models.* The digital twin work sits here. A platform head describes having "[s]pearheaded the development of an AI engine to optimize production lines" and "[o]wned the development and deployment of digital twin models for a new manufacturing platform."

An IT director for industrial IoT describes "[d]riving P&G's companywide IIoT platform for managing near real-time ML manufacturing models." A modeling director describes a program "that aims at digitizing end to end the innovation process from an idea on the paper to the production line."

- *Supply-chain & logistics systems—embedding planning judgment into standing systems.* A supply chain executive describes deploying "AI-based route optimization tools (P&G ORION system)" and later leading "the development of a machine learning-based dynamic replenishment model," a framework subsequently adopted into a firmwide AI inventory platform. An engineering manager describes co-founding an "internal startup delivering AI/algorithmic warehouse optimization solutions across P&G's global network." A principal data scientist extends the twin idea from machines to processes, listing "[r]eal-time Digital Twins of business processes" including a "[c]ustomer order management digital twin to minimize on-shelf out-of-stocks."

- *Marketing & revenue-operations systems—consolidating consumer learning into reusable assets.* A product manager leads a Voice of Consumer suite described as "an enterprise ecosystem of AI, ML, and large language model (LLM) products that transform how 10 global businesses understand and act on consumer feedback." A brand director describes "[s]pearheading the use of generative AI in brand building, from developing proprietary knowledge to conducting pilot projects and scaling successful initiatives."

- *Proprietary-knowledge GenAI/NLP systems—codifying tacit expertise directly.* The same principal data scientist describes "Corporate Knowledge Management using Generative AI," comprising an "[i]nsight engine for explicit knowledge using large language models" alongside "[c]apture and insight building from tacit knowledge residing with Experts (SpeechBrain)"—an unusually explicit statement of the conversion this section is about, and one that names the instrument: a speech toolkit, applied to what experts say rather than to what they have already written down. The same description places this work directly alongside "[r]eal-time Digital Twins of business processes," so the codification of expert judgment and the codification of process behavior are treated as one program. In HR, a manager describes prompt-engineering a system to write job descriptions and co-developing a review-process assistant that became "the most used P&G-wide GPT assistant."

- *Reusable data platforms & pipelines—turning the firm's own data into shared, reusable assets.* A chief data and AI officer describes having "[b]uilt the enterprise AI Factory & GenAI platform, compressing time-to-value from months to weeks." A group scientist describes delivering "enterprise AI platforms integrating multi-agent ar-

chitectures, workflow orchestration, and retrieval-augmented generation (RAG) … powered by proprietary data."

Figure A.2 places these examples in the context of the firm's hiring. Panel (a) shows that organization-capital-related AI roles were flat as a share of P&G's US positions until about 2013 and then rose sharply, roughly a tenfold increase by the early 2020s before levelling off, while other AI roles grew far more modestly and plateaued after 2020. The gap between the two series widens substantially over the period, so the firm's AI hiring has been directed overwhelmingly toward work that converts proprietary operations into durable assets rather than toward generic technical capability. Panel (b) breaks that series into the kinds of work that qualify it. Reusable data platforms and pipelines account for much of the level throughout, and marketing and revenue-operations systems form a steady base—consistent with the consumer-knowledge and brand examples above—while supply-chain and logistics systems and cross-functional operating routines together reflect the planning and replenishment work described earlier. The timing of the takeoff is itself informative: the digital twin work predates the generative-AI wave by nearly a decade, since codifying plant telemetry into a simulatable model required only conventional machine learning applied to data the plants were already producing. That is why P&G's series accelerates from around 2014, whereas at Morgan Stanley it accelerates after 2021—document and knowledge work became codifiable at scale only once language models arrived. Consistent with this, the category of systems built on proprietary knowledge using generative AI and natural language processing is essentially absent before the early 2020s here too, arriving late but on top of a much larger pre-existing base.

The productivity channel here is unusually visible. A physical trial consumes capacity: the line stops, and output is forgone to learn something. Running the trial against a digital twin decouples learning from production, so the plant keeps producing while the experiment happens elsewhere—a 21% reduction in trial shutdown days is additional output from the identical stock of plant and equipment. The same mechanism compresses the innovation cycle, since more candidate formulations and process settings can be evaluated per unit of engineering and plant time, which is what a 72% reduction in innovation lead time and a tenfold acceleration in speed to market represent. On the supply side, replenishment and routing models deliver the same service levels with less inventory and fewer miles, lowering the working capital and logistics cost embedded in each unit sold. And because a twin built at one site can be adapted to the next, the fixed cost of codifying is spread across a network of plants and lines rather than borne once per site.

Two features stand out. First, in nearly every case the input is data the firm was already generating—process telemetry, order flows, consumer feedback, supplier prices—and the

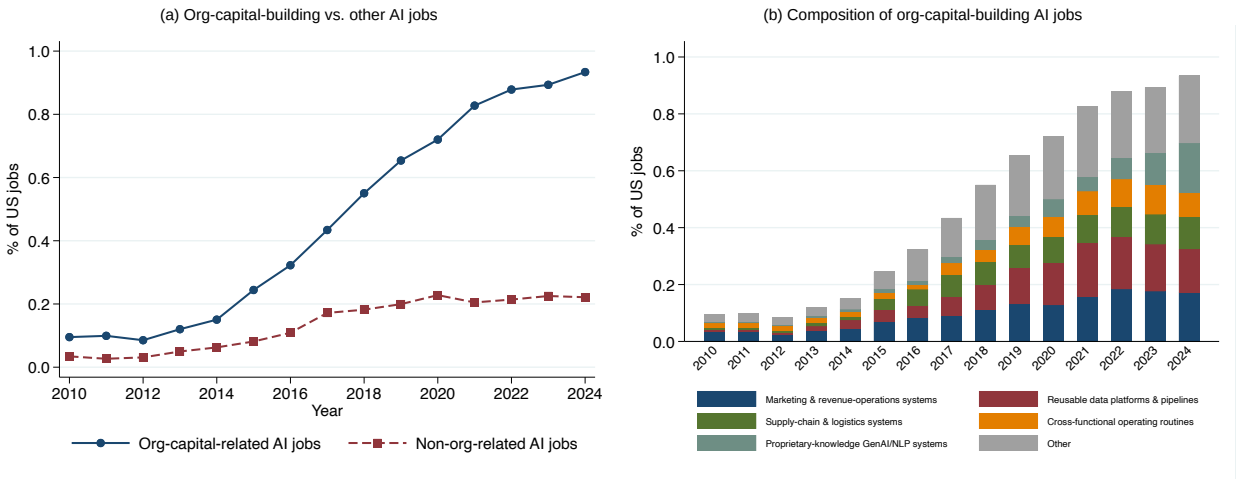


Figure A.2. AI jobs at Procter & Gamble as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification.

output is an asset that persists after the people who built it move on. P&G's own account of its AI strategy emphasizes exactly this, unlocking insight from data the company already holds (MIT Sloan Management Review). Second, a substantial share of this work is not any single application but infrastructure for producing applications: platforms, model governance, and internal accelerators that lower the cost of the next conversion. That is what makes the accumulation compound, and it is also why the payoff is unlikely to be immediate.

Moderna. Moderna presents a different starting point from the other two cases: a firm founded in 2010 and without a commercial product until 2020, which scaled from clinical-stage research to global manufacturing in a matter of years. Its job descriptions read accordingly: a computational sciences executive describes having "[f]ounded, built, and led the first ever ML/AI capability at Moderna, the first custom informatics platform"; a commercial data director describes leading "the establishment of data infrastructure from the ground up for a rapidly scaling pharmaceutical company, transforming it from a startup to a commercial enterprise." What Moderna lacked was not operational knowledge but codified organizational assets—the procedures, routines, and reusable systems that a firm normally accumulates over decades. On the contrary, the company describes a decade-long library of data as its single biggest asset, and being cloud-native from the outset as what let it avoid the data silos that slow codification elsewhere (Challamel, 2023). A firm generating knowledge far faster than it has codified it is exactly where the mechanism we describe binds hardest.

Moderna's response was to make codification something every employee does rather than a project run by a central team. Its internal chatbot, mChat, was built in about two weeks and launched in May 2023; within months roughly two-thirds of employees were active users, embedding it into their own workflows (Challamel, 2023). The company subsequently deployed ChatGPT Enterprise with a builder that lets employees create assistants trained on their own team's documents and terminology, and within roughly two months employees had built more than 750 such tools across legal, research, manufacturing, and commercial functions. The significance is not the count but the mechanism: the people who hold the operational knowledge are the ones converting it into durable form, without waiting for engineering capacity.

The AI jobs we observe at Moderna show what this looks like in practice, and cluster into the categories our classification identifies as most common at the firm:

- *R&D & scientific platforms—turning a short experimental history into predictive models.* The dominant category in Moderna's early years. An RNA process engineer describes having "founded and launched the RNA AI initiative by building predictive ML models from thousands of development experiment data points," work that "accelerated development time by several months, reduced unexpected manufacturing failures." A limited stock of experiments becomes a reusable asset that substitutes for experience the firm has not yet had time to accumulate. A quantitative biology executive describes integrating "robotic plate handling and AI/ML tools into NGS Prep" and improving "time to result from 4 hr/run to 20 min/run with improved accuracy."

- *Reusable data platforms & pipelines—building the firm's data foundation from nothing.* A VP of data and robotics platforms describes having "[d]esigned Moderna's data platform strategy" and "[d]eveloped and launched an AI platform integrating LLMs with proprietary data, both structured and unstructured," including an operational control tower for the supply chain. A VP of AI products describes introducing "AI as a utility to all knowledge workers through a company-wide transformation program."

- *Proprietary-knowledge GenAI/NLP systems—codifying process and documentation directly.* A manager in Japan describes developing "custom GPT tools, including 'Meeting Minutes GPT' and 'SOP Creation GPT,' significantly improving operational efficiency and reducing document preparation time from days to hours." An HR director describes deploying "custom GPT solutions that reduced manual inquiries and strengthened knowledge management." Standard operating procedures are organization capital in its most literal form, and here they are written with the assistance of

the tool.

- *Cross-functional operating routines—institutionalizing operating routines and coordination mechanisms.* A supply chain associate director describes championing "AI-driven process optimization using Mimica task mining to identify automation opportunities across planning and logistics workflows," and co-designing a "multi-agent strategy integrating AI, IoT, and digital twins to enable autonomous operational decision-making"— instrumenting operations to find what is worth codifying next.
- *Manufacturing & process systems—converting production and quality experience into reusable routines.* A quality specialist describes spearheading "AI-enabled, LEAN-based opportunity engines across Quality, CMC, and Manufacturing," extending codification into the regulated production environment the company had to build from nothing.
- *Hiring, training & talent development—institutionalizing AI skill-building into training curricula.* A talent development head lists the "launch of Moderna's AI Academy"; a learning director describes creating "Moderna's Digital College" with "comprehensive learning paths in AI, Cybersecurity, and Digital Tooling." This category is small in the aggregate but central here: distributing the capability to codify was the company's explicit strategy.

Figure A.3 places this in the context of the firm's hiring, and the levels are striking: organization-capital-related AI roles run at one to three percent of Moderna's US positions, several times the shares at Morgan Stanley or P&G. Panel (a) shows the series dipping through 2020—the pandemic years, when the company was hiring at scale for manufacturing and distribution—and then roughly tripling between 2020 and 2024, while other AI roles stay flat below half a percent throughout. The acceleration begins after the commercial scale-up rather than during it, consistent with a firm that first built the operation and then codified it. Panel (b) shows what changed. R&D and scientific platforms dominate the early years, when AI at Moderna was largely a research capability; from 2021 onward the growth comes from reusable data platforms and pipelines, cross-functional operating routines, and manufacturing and process systems—the codification of how the company runs rather than what it discovers. Systems built on proprietary knowledge using generative AI appear from 2022 and expand quickly, matching the mChat and custom-GPT timeline described above.

The productivity channel is one the company states directly: whereas conventional operating models require headcount to grow with the business, its aim is to "scale the value of their people"—which is output per worker (Challamel, 2023). The mechanism is

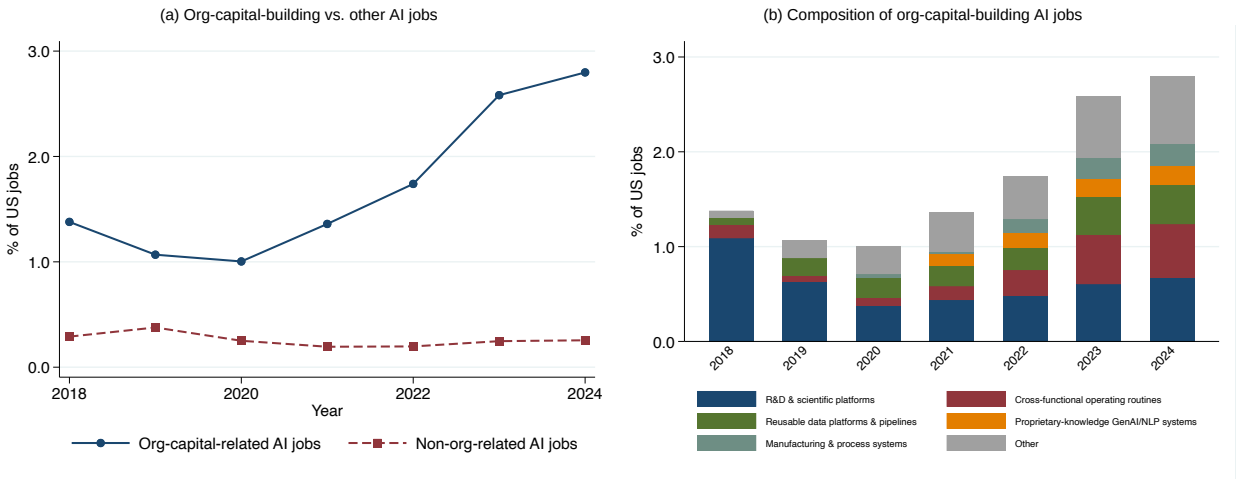


Figure A.3. AI jobs at Moderna as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs from the corresponding figures for Morgan Stanley and Procter & Gamble.

visible in the individual cases. Predictive models built from a few thousand development experiments let process engineers avoid failures they have not personally encountered, so output rises not because inputs increased but because the firm stopped rediscovering the same things. Drafting SOPs, meeting records, and regulatory documentation consumes scientific and professional hours that generate no output themselves; compressing that work reallocates skilled time toward activities that do. And because a GPT built by one team is available to the next at essentially no cost, each conversion raises the productivity of people who had no part in building it.

The broader point is the one our heterogeneity result makes. Moderna entered the AI era holding little organization capital as new firm, and so had the most to gain from a technology that speeds the accumulation of organization capital—which is exactly what followed, as the company built that stock far faster than a firm accumulating it through ordinary experience could have. That it chose to distribute the capability to every employee rather than centralize it follows from the same reasoning: when the codified stock is thin relative to the operation, the binding constraint is the rate of conversion, not the quality of any single converted asset.

Home Depot. Home Depot presents the mechanism in its most literal form. The firm's competitive asset has always been what it calls orange apron knowledge—the expertise of store associates, many of them former tradespeople, who can tell a customer which grout works with glass tile and where to find it. That knowledge is the definition of tacit: it lives in individuals, it is available only in one aisle of one store at a time, and it walks

out the door with the associate. In a retail workforce of over 470,000 associates across more than 2,300 stores, with the turnover rates the sector implies, the firm has historically had to rebuild it continuously. AI lets Home Depot convert that knowledge into a stock the company holds.

The clearest instance is Magic Apron, an assistant that began on product pages and has become a conversational tool spanning the firm's digital platforms, extending to real-time store inventory and aisle-level product locations. What it is built from is the point. The company describes it as powered by a proprietary home improvement knowledge base that combines large-scale datasets with its own expertise and product information, trained on Home Depot's proprietary project expertise and product catalog alongside information from the wider internet, and explicitly designed to leverage the company's own human expertise; the tool continues to learn from customer and associate feedback and serves as a foundation for store and contact center associates (The Home Depot, 2025a,b). The executive who leads the effort has framed the conversion in both directions—customers have always relied on the expertise of orange-aproned associates in the aisles, and the aim is to bring that same expertise to the digital world (The Home Depot, 2025a), "putting 'Orange Apron' expertise in the pocket of every customer" (The Home Depot, 2026b). Alongside it sit AI-powered materials lists, which let pros describe a job in plain language and receive a grouped list including items they had omitted, and blueprint takeoffs that produce quotes in days rather than weeks (The Home Depot, 2026a). Less visible but equally relevant are the operational systems: order intelligence that identifies delivery misses before they reach the customer, and route intelligence that layers customer-specific data with road and weather conditions to anticipate access problems at the jobsite (The Home Depot, 2026b).

The AI jobs we observe at Home Depot show the work of conversion, and cluster into the categories our classification identifies as most common at the firm:

- *Reusable data platforms & pipelines—assembling orange apron expertise into a queryable stock.* The largest category. The knowledge base behind Magic Apron is built from the company's accumulated project expertise and refined continuously by associate feedback (The Home Depot, 2025a), and the job descriptions show the internal counterpart: a senior director of data science describes leading enterprise generative AI foundation work to "build the federated home improvement knowledge brain and AI common services" for the company's applications.
- *Firm-specific predictive & decision models—embedding store-level experience into forecasting and decision models.* A principal data and analytics lead describes engineering "a novel On-Shelf Availability (OSA) metric by integrating inventory, bay sequencing, and

pack-down data"—converting what experienced staff knew informally about stock position into a predictive measure usable everywhere. A senior director of operations describes adopting "AI and machine learning to improve labor forecasting, task prioritization, and operational visibility."

- *Marketing & revenue-operations systems—turning customer interactions into structured project knowledge.* A senior data science manager describes having "[u]sed deep learning & knowledge graph to recognize customer's DIY project intents in real time, providing holistic product, how-to guide, rental and home services recommendations"—customer interactions turned into structured project knowledge.
- *Reusable software platforms & workflows—standardizing execution across a distributed store network.* A product leader in operations describes designing "platform capabilities that strengthened visibility, standardization, and performance consistency across a distributed store network," with tools that "empowered frontline associates with real-time insights." A manager in inventory transformation describes launching "a reset platform, anticipated to eliminate 450 hours per week of non-value-added work."
- *Cross-functional operating routines—codifying frontline procedure.* An analyst in process and knowledge management describes "simplifying process and instruction documentation using Lean Six Sigma (LSS) and Knowledge-Centered Service (KCS) methodologies," work that improved guided workflow integration while "reducing onboarding costs by $3k per hire"—documentation as the mechanism by which a new associate reaches proficiency faster. A conversational UX designer describes "[e]stablishing Conversation Design practices at the Home Depot" and creating a guild to teach others where conversational AI applies.

Figure A.4 places these examples in the context of the firm's hiring. Panel (a) shows organization-capital-related AI roles rising almost continuously as a share of Home Depot's US positions, with the increase steepening after 2014 and continuing through 2024, while other AI roles grow more slowly and flatten after 2018. The levels are worth noting: they are well below those at the technology-intensive firms in our set, which is what one expects of a retailer whose workforce is overwhelmingly in stores rather than at headquarters. Panel (b) decomposes that series. Reusable data platforms and pipelines and marketing and revenue-operations systems form much of the base, consistent with the consumer-knowledge and merchandising work described earlier, while firm-specific predictive and decision models together with supply-chain and logistics systems reflect the on-shelf availability, labor forecasting, and delivery examples. The composition is notably

stable over time relative to the other firms we examine: Home Depot was codifying store and supply-chain operations with conventional machine learning long before generative models arrived, and the recent Magic Apron work sits on top of that accumulated base rather than initiating it.

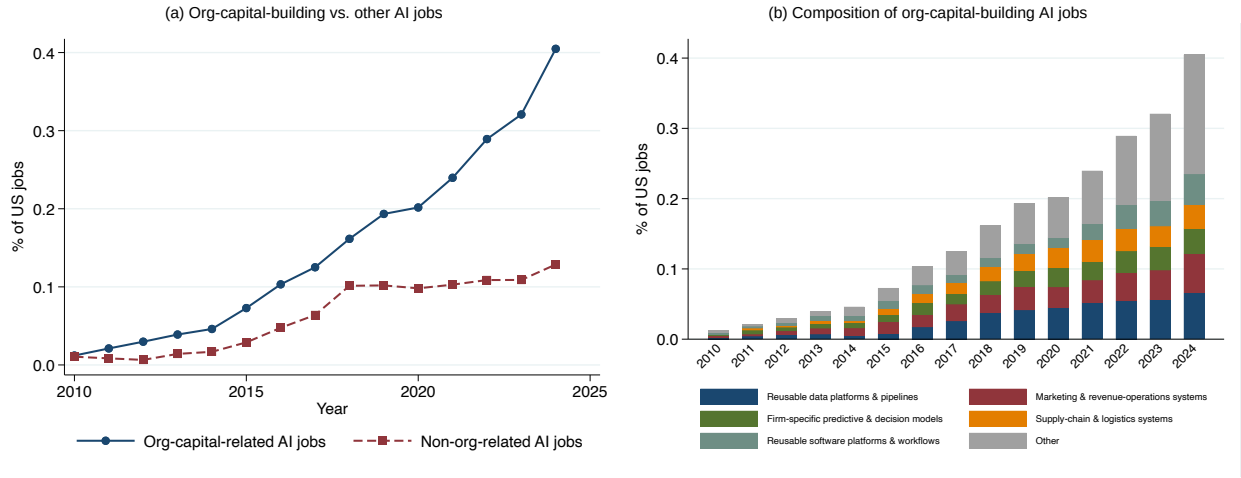


Figure A.4. AI jobs at Home Depot as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs from the corresponding figures for the other firms.

The productivity channel is the sharpest of our cases because the multiplication factor is largest. Expertise codified once is available in every aisle of every store simultaneously, at no additional cost per store; the same asset that helps one associate helps all of them. This matters most precisely where embodied knowledge is least durable. In a workforce with high turnover, an associate's departure has historically meant losing what they knew, and the firm paid to rebuild it with each hire—which is why the onboarding-cost reduction above is not a minor administrative saving but the mechanism working. On the customer-facing side, materials lists and blueprint takeoffs compress estimating work that generates no output itself, and order and route intelligence prevent failed deliveries, which are pure waste of labor, fuel, and truck capacity.

CVS Health. CVS Health presents the mechanism in an unusually integrated setting. The company spans roughly 9,000 retail pharmacy locations, more than a thousand walk-in and primary care clinics, a pharmacy benefits manager with about 88 million plan members, and health insurance serving more than 37 million people through Aetna (CVS Health, 2026a). Operational knowledge accumulates across all of it automatically, as a by-product of filling prescriptions, paying claims, and managing care. The constraint

has never been generating that knowledge but rendering it usable: much of it arrives as unstructured clinical documentation—as Aetna's chief digital and technology officer describes the records, (Olavsrud, 2026) "medical charts that have white space filled with handwritten notes." The firm possessed this knowledge all along without being able to act on it at scale.

The clearest instance is AI Medical Chart Review at Aetna. Under HEDIS quality measures, large managed care providers review more than ten million medical records annually to identify gaps in care—missed screenings, overdue tests, lapsed chronic disease management. Industry benchmarks put the annual review at roughly 50,000 work weeks, nearly a thousand dedicated full-time employees; a team of fifty working manually would need more than two decades to complete a single year. The platform extracts clinically relevant data and ranks records by likelihood of measure closure and evidence strength, passing a prioritized set to human reviewers. It has cut the need for manual review by 65% and processed fourteen million documents, with capacity formerly spent on review shifting toward quality control (Olavsrud, 2026). The same logic runs through claims and authorization: the second-generation Claims Assist Manager cuts processing time by more than 20% on complex claims requiring manual review (CVS Health, 2026a), and prior authorization reform has standardized data requirements for 88% of volume, moved 83% of authorizations to real-time processing, and eliminated more than a million provider calls (CVS Health, 2026b). On the pharmacy side, member-facing and colleague-facing tools draw on the same accumulating record, and CVS Health has since announced Health100, a subsidiary intended to deliver an integrated engagement platform across these surfaces (CVS Health, 2026b).

The AI jobs we observe show the work of conversion, span the company's businesses, and cluster into the categories our classification identifies as most common at the firm:

- *Reusable data platforms & pipelines—making a record the firm already held usable.* By far the largest category, and the binding constraint for a firm whose problem is not generating knowledge but rendering it usable. A lead director of machine learning engineering describes "[b]uilding the world's largest RAG system for knowledge management for CVS Health's 300,000+ employees that will scale to 1M+ documents." A senior manager in data governance describes pioneering "creation of a business glossary, standardizing terminology" across the firm's data.
- *Proprietary-knowledge GenAI/NLP systems—extracting clinical knowledge from unstructured documents.* An NLP technical lead describes responsibility for "the creation and maintenance of the enterprise NLP platform," used across domains for "bot creation, customer IVR systems, financial contract extraction, and legal text analysis." An AI

research architect describes having "mapped over 1,600 distinct healthcare outcomes and their associated friction points," work that "collapsed delivery timelines from six weeks to one," and a research operations manager describes an "AI-enhanced research repository with automated tagging and intelligent retrieval."

- *Cross-functional operating routines—codifying procedure as the precondition for AI.* An executive director of employee experience describes having "[d]igitized and standardized HR knowledge, policy, and procedure content and implemented a governance framework as the foundation for AI"—an explicit statement that documentation comes first and the model second. A head of enterprise AI governance describes establishing the company's Responsible AI operating model and policy framework across all lines of business, necessary infrastructure in a regulated setting and codified organizational practice in its own right.

- *Reusable software platforms & workflows—translating pharmacy practice into standardized workflows.* A director of product management describes leading development of "an AI-powered platform that helped reduce medication errors across 9,000+ pharmacies." A platform head describes having "[b]uilt and scaled a Virtual Pharmacist to answer drug and side-effect questions," and a retail IT director describes building "StoreOS, a unified IoT/AI-driven platform redefining colleague workflows" including pharmacy pill verification AI.

- *Security, compliance & risk systems—turning audit and control experience into full-coverage review.* A director of business innovation describes having "[r]eplaced manual sampling with algorithm-driven 100% transaction review, reducing compliance risk," alongside unlocking eighteen thousand workforce hours through automation. In the PBM, a lead director describes delivering a "GenAI-powered Audit Advisor that reduced contract risk and improved rebate accuracy."

Figure A.5 places these examples in the context of the firm's hiring. Panel (a) shows organization-capital-related AI roles growing slowly through the mid-2010s, accelerating sharply after 2016—around the Aetna acquisition, which combined the pharmacy, PBM, and insurance records into a single firm—and continuing to climb steeply through 2024, reaching a level roughly three times that of other AI roles. The two series diverge progressively rather than moving together, so the firm's AI hiring has been directed increasingly toward work that converts proprietary operations into durable assets. Panel (b) shows what that work consists of, and the composition is distinctive among our cases: reusable data platforms and pipelines dominate throughout and account for close to half the total by the end of the sample. That is what one expects of a firm whose codification problem

is less about capturing knowledge that was never recorded than about making usable an operational record it already possessed—the pipelines are the binding constraint. Cross-functional operating routines and marketing and revenue-operations systems form steady secondary components, and systems built on proprietary knowledge using generative AI and natural language processing appear from the early 2020s and expand quickly, matching the chart review, claims, and conversational tools described above.

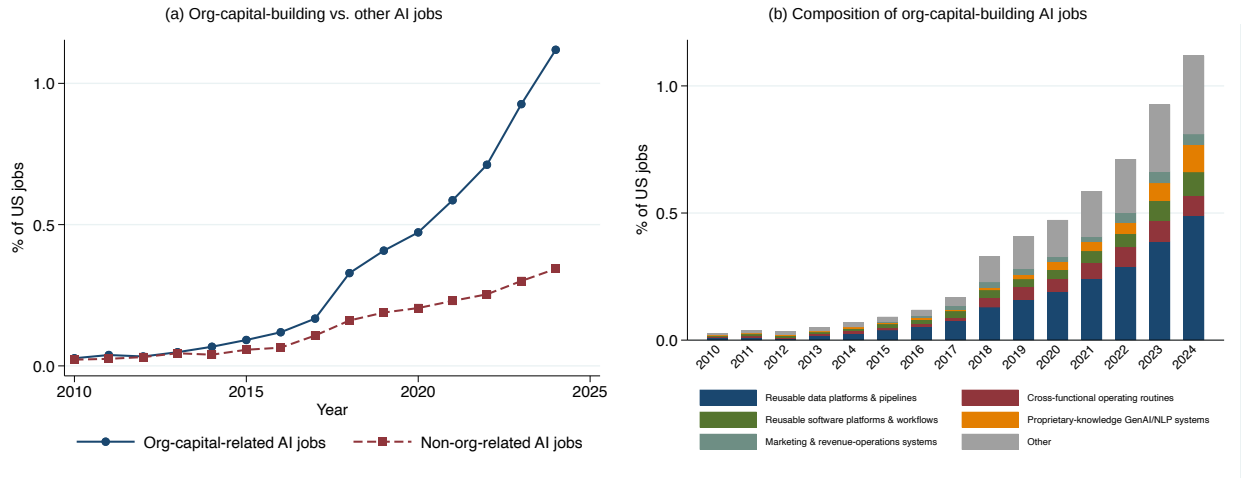


Figure A.5. AI jobs at CVS Health as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs across the corresponding figures for other firms.

Two productivity channels stand out, and the first is distinctive to this case. In our other examples AI speeds up work the firm was already doing; here it makes feasible a task that full coverage had made impossible. Reviewing every chart, or every transaction rather than a sample, was never a matter of working faster—it exceeded what any realistic workforce could do, which is why the industry relied on partial review. Codification converts sampling into census, and the firm learns things about its own operations that were previously unknowable at any cost. The measured gains follow: the same coders process a far larger share of the record, with freed capacity redirected to quality control rather than eliminated, so output per worker rises without headcount falling; higher gap closure feeds Star Ratings and reimbursement, making this a revenue channel and not only a cost one; and eliminating a million provider calls removes administrative work from both sides of a transaction. The second channel is integration. Because the same person may fill a prescription at a CVS pharmacy, see a clinician at MinuteClinic, and file a claim through Aetna, knowledge codified in one business is usable in the others—an asset built once and drawn on repeatedly across surfaces that a monoline insurer or a standalone

pharmacy could not connect.

Verizon. Verizon runs two large workforces whose knowledge has historically been hard to hold: tens of thousands of field technicians who install and repair the network, and tens of thousands of care representatives who resolve customer problems. In both cases the expertise is procedural, learned on the job, unevenly distributed, and lost when someone leaves. The firm also generates an enormous operational record—dispatch logs, trouble tickets, network telemetry, support transcripts—as a by-product of running the network. AI lets it convert that record into an asset the company holds.

The clearest instance is on the care side, and its construction is explicit about what is being converted. Verizon's internal assistant was built by feeding roughly fifteen thousand internal documents into Google's Gemini model (Cai, 2025). Deployed across 28,000 care representatives and retail locations, the Personal Research Assistant surfaces context-aware answers and has been extended with automated conversation summarization and follow-up reminders; a companion Problem Solver agent has been reported to benefit newer representatives in particular, raising their confidence on complex queries (Google Cloud, 2025). The reported effect was not substitution but reallocation: shorter calls freed representatives to sell, with sales through the service team up sharply, which Verizon's consumer chief described as reskilling care agents into selling agents (Cai, 2025). On the network side, the company has moved toward agentic automation running on its own cloud platform (Benton Institute for Broadband and Society, 2026).

The AI jobs we observe show the work of conversion, and cluster into the categories our classification identifies as most common at the firm:

- *Customer-service & support systems—putting the best representative's answer on every screen.* The largest category at Verizon by a wide margin, and the one no other firm in our set shares. A head of data and analytics describes "AI-enabled productivity platforms" transforming "workforce operations across 65K+ retail and call center employees." The internal assistant built on the firm's own documents belongs here, and its effect was reported as reallocated time rather than reduced headcount.

- *Reusable cloud, infrastructure & reliability operations—codifying network and field operating procedures into automation.* A senior director of software development describes driving "Zero-Touch Operations (ZTO) across Verizon's wireless and wireline network operations" and pioneering a "Dispatch and Field Operations platform, achieving 85% digitization of manual service transactions via Machine Learning." A network operations technology head describes systems enabling a "30k+ workforce to be efficiently dispatched," delivering "mobile tools for our frontline 30k+ customer service technicians and their supervisors." An associate fellow describes an "OpenSON

ecosystem," a "real-time, event driven automation ecosystem employing AI/ML pipelines" across radio access network domains.

- *Reusable data platforms & pipelines—unifying the record so it can be codified.* A head of data and analytics describes designing an organization "integrating data engineering, science, and governance under a unified model," the precondition for codifying anything at scale across a firm of this size.
- *Cross-functional operating routines—institutionalizing coaching as a standing management routine.* A Lean Six Sigma transformation manager describes leading development of "Smart AI Coach (SAIC), an early AI application that provided technicians and their managers with real-time comparative results and coaching information," and a Data Driven Process Analysis program deployed after acquiring "data feeds from over a dozen separate systems" and training "over 1600 Local Managers." A vice president for data and AI describes establishing an enterprise AI council and Responsible AI practice across care, marketing, operations, and network.
- *Knowledge organization & ontologies—organizing specialized knowledge into ontologies and semantic layers.* Small in the aggregate but unusually explicit here. A principal engineer describes having "[b]uilt the Knowledge Engineering discipline from scratch, authoring lifecycle standards and governance frameworks adopted across Data & Analytics," and designing "a Unified Semantic Knowledge Layer with 1800+ classes supporting data literacy through business-friendly querying," including a "Network Operations ontology enabling semantic search." Few descriptions in our data state the mechanism as directly.

Figure A.6 places these examples in the context of the firm's hiring. Panel (a) shows organization-capital-related AI roles rising steadily from 2010 and accelerating from about 2015, reaching a level roughly three times that of other AI roles by 2024, with the gap widening throughout rather than converging. Panel (b) explains why this case looks different from the others in our set: customer-service and support systems are the single largest category in every year and account for close to a third of the total by the end of the sample. No other firm we examine has its codification effort so concentrated on a single function. That concentration matches the case: the binding constraint is the dispersion of procedural knowledge across tens of thousands of care representatives and field technicians, so the asset worth building is the one that makes the best representative's answer available to all of them. Reusable data platforms and pipelines and cloud and reliability operations form the infrastructure beneath it, cross-functional operating routines appear steadily from the mid-2010s, and systems built on proprietary knowledge using

generative AI and natural language processing emerge from the early 2020s—the timing of the assistant built on the firm's own internal documents.

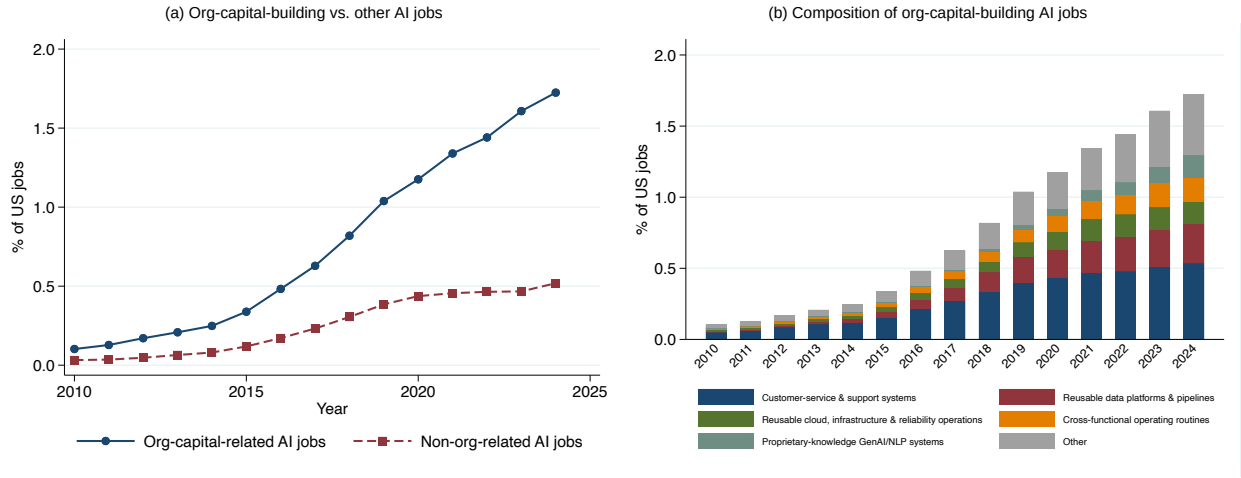


Figure A.6. AI jobs at Verizon as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs across the corresponding figures for other firms.

Two features of the productivity channel stand out. First, the gains land disproportionately on the less experienced. The care assistant compresses the gap between a new representative and a seasoned one by making the seasoned one's answers available on screen, which is what codifying dispersed expertise into a common asset should do, and it is why the reported effect appears as reallocated time rather than reduced headcount—shorter calls converted into selling time. Second, the field organization illustrates the same logic in physical work. Digitizing manual service transactions and coaching technicians against comparative results turns the tacit skill of the best installers into a standard the rest can follow; the output is more jobs completed correctly per technician-day, with fewer repeat visits. In both cases the asset is non-rival within the firm: an answer codified once serves all 28,000 representatives, and a dispatch improvement serves the whole footprint.

Nike. Nike presents a variant of the mechanism our other cases do not. Much of the operational knowledge the firm depends on is generated outside its own walls: Nike designs and sells footwear and apparel but does not own the factories that make it. The know-how of how a shoe is actually built—tooling, materials behavior, line balance, quality yield, cost drivers—accumulates in a network of contract manufacturers and in the experience of the Nike staff who work with them. That knowledge has historically been partly tacit and partly stranded at the partner, which makes codification both more valuable and harder than in a firm that owns its plants. Nike acquired several analytics

and data-integration firms between 2018 and 2021 (Kitishian, 2025; Sohel, 2026), but the acquisitions supplied components rather than the asset itself; what the internal roles below describe is the work of attaching those components to the firm's own processes and records, which is where the stock is actually built.

The clearest internal instance is the manufacturing backbone. A senior director describes being "[p]romoted to jump-start and lead the first digital transformation of Nike's manufacturing landscape of 500+ factories, through modern platforms, software and connectivity," producing "a connected and resilient ecosystem of multiple Nike functions and their manufacturing partners, delivering 1.4B+ units of Nike product on time." The instrument is described as "the industry leading data-streaming and partner-collaboration Nike Manufacturing Backbone platform," digitizing "critical supply chain collaboration, visibility, resilience, and sustainability business processes." The object being built is a shared record across a boundary the firm does not control.

The AI roles we observe show the work of conversion, and cluster into the categories our classification identifies as most common at the firm:

- *Reusable data platforms & pipelines—merging decades of incompatible records into one source of truth.* The largest category, and the step at which dispersed operational records become a single addressable stock. A director of enterprise data and analytics describes building a cloud data lake as "one 'source of truth' for data and insights," ingesting "transactional ERP, point of sale, digital/consumer, and product creation" data. The records being merged are the firm's own, accumulated over decades in incompatible systems.

- *Product & design systems—codifying product-creation expertise into tools.* A director of footwear cost engineering describes shifting the function "from 18 Individual Contributors to a 4-person Center of Excellence (COE) that provides 'smart' tools to Product Creation team, enabling quick decisions," including "an AI/ML-based auto-budgeting tool that reduced creation time from 3 hours to 2 minutes"—costing judgment that had lived in eighteen experienced people becoming a tool available to the whole creation organization. The publicly documented design-side analogue is the same substitution at larger scale: Nike's Athlete Imagined Revolution project applies generative models to footwear concept generation with elite athletes, compressing concept-to-prototype from weeks or months to hours (NIKE, Inc., 2024; Team DigitalDefynd, 2026), with the underlying asset described as a proprietary model trained on decades of Nike's own athlete performance data rather than a general corpus (Kitishian, 2025).

- *Supply-chain & logistics systems—codifying supplier and fulfillment learning across factories*

Nike does not own. A senior director describes owning product and platform strategy for "Manufacturing, Digital Consumer Order Management, Wholesale Customer Order Management, Inventory Management, Global Fulfillment, Global Distribution, Global Customs & Trade, & Global Transportation," directing hundreds of engineers across a portfolio intended to integrate these processes end to end. The manufacturing backbone described above is the flagship instance of this category.

- *Cross-functional operating routines—agreeing what the numbers mean across functions.* The same enterprise data director describes pairing the data lake with "governance, master data management," and "KPI/Metric alignment across stakeholder teams"—agreeing what a metric means across functions is itself codified practice, and the precondition for the marketing and revenue-operations work that forms the demand-side counterpart to the supply-chain systems above.
- *Reusable AI/ML platforms & MLOps—turning model-development experience into tooling many teams can use.* A vice president for global architecture describes leading the transformation of "a $47B company with over 79,000 employees, and over 50 years of technology accumulation, through a platform transformation." A senior director of technical product management describes an ecosystem of "standardized, fully managed and composable control planes" spanning developer, data, and AI capabilities, exposed to agents through standard interfaces.

Figure A.7 places these examples in the context of the firm's hiring. Panel (a) shows organization-capital-related AI roles essentially flat through 2013 and then rising sharply, with a particularly steep step around 2018 and continued growth to roughly five times the level of other AI roles by 2024. The timing is consistent with the case: the manufacturing and supply-chain digitization described above dates from the late 2010s, and the acceleration follows it rather than preceding it. Panel (b) decomposes the series. Reusable data platforms and pipelines form the base throughout, with marketing and revenue-operations systems and supply-chain and logistics systems as substantial and roughly comparable components—a split that matches a firm codifying on two fronts at once, consumer demand knowledge on one side and manufacturing and distribution knowledge on the other. Cross-functional operating routines and reusable AI/ML platform work grow steadily from the late 2010s, reflecting the platform and center-of-excellence roles described above. Notably, no generative-AI category appears in Nike's composition, in contrast to Morgan Stanley, CVS Health, and Verizon. Nike's codification effort runs through structured operational and commercial data rather than through documents and language, which is what one would expect when the knowledge being converted concerns how a product is costed, made, and moved rather than what an expert would say.

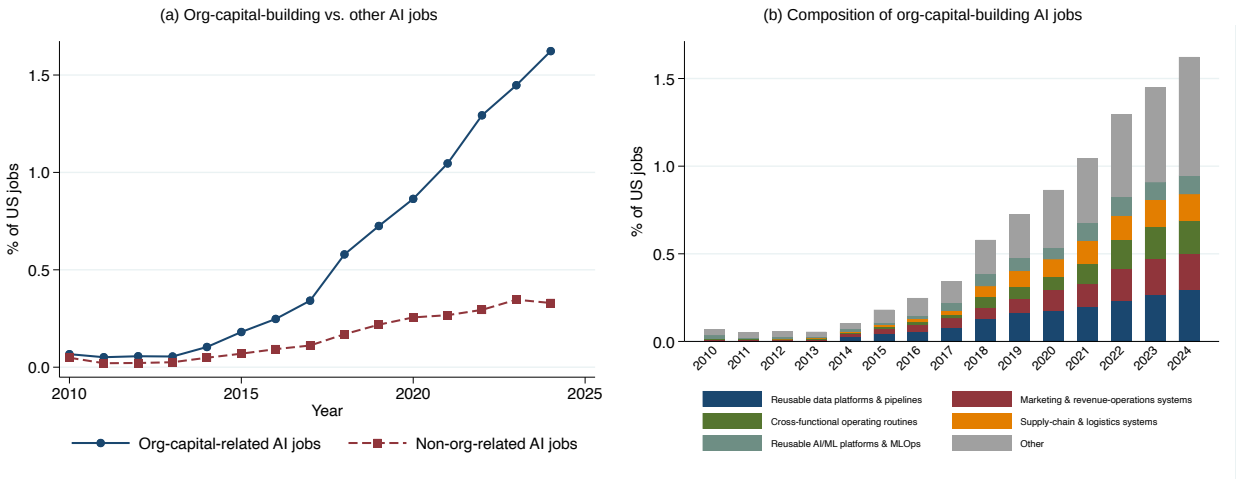


Figure A.7. AI jobs at Nike as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs across the corresponding figures for other firms.

The productivity channel runs through decision speed and through yield. Product creation at Nike is a repeated design problem under cost and performance constraints, and the binding input is the time of experienced designers, developers, and engineers. Converting costing and specification judgment into tools does not make those people better; it means the median decision is made with the knowledge of the best decision-maker, in minutes rather than hours, so more design iterations occur per unit of skilled time. On the manufacturing side, a shared digital record across hundreds of factories reduces the coordination failures—late changes, misread specifications, quality escapes—that consume material and rework capacity at the partner and delivery reliability at Nike. On the inventory side, the return is the difference between a demand forecast made once per season and one made continuously, which shows up as avoided markdowns rather than as higher output.

One qualification follows. Knowledge codified into a partner-facing platform is not entirely firm-specific: contract manufacturers serve competing brands, so practice Nike documents and teaches may diffuse. Our assumption that the codified stock stays with the firm is reasonable for internal documentation and poor for a shared platform spanning independent suppliers, which makes Nike the case where appropriability is most clearly at issue, and where the return to codification depends on how much of what is written down the firm can keep.

Expedia Group. Expedia Group presents the mechanism in a firm whose product is software and whose operational knowledge is therefore largely engineering knowledge: how its

services fail, how incidents are diagnosed, how a feature moves from specification to production, how the marketplace behaves under stress. That knowledge accumulates in incident records, code repositories, runbooks, and the practice of individual engineers, and it has historically been unevenly distributed—the best engineer on a team knows things the others do not, and the firm has no way to hold what they know apart from their continued presence. What makes this case distinctive is that Expedia has treated the diffusion of that expertise as an explicit strategy rather than a by-product.

The company's approach has two parts. The first is infrastructure: a shared internal stack built before any particular application, including agent testing and evaluation infrastructure and an internal environment giving employees access to many large language models to build their own assistants. Since early 2025 employees have built more than 1,500 agents in that environment (Fortune, 2026). The second is the diffusion strategy, and it states the mechanism of this paper almost exactly. After a nine-month study of how its roughly 5,000 engineers used AI coding assistants, the company found that its strongest performers gained far more than the average engineer, and set out to identify those engineers, capture their methods, and spread them across the organization (Skift, 2026). One engineer is described as having assembled a personal workbench of roughly ten agents supported by fifty to sixty specialized skills, changing his work from writing code to orchestrating agents (High, 2026). Codifying that workbench so others can use it is precisely the conversion of individual expertise into an asset the firm holds.

The AI jobs we observe show the work of conversion, and cluster into the categories our classification identifies as most common at the firm:

- *Reusable AI/ML platforms & MLOps—codifying how software gets built.* The shared stack the company built before pursuing particular applications. A director of software engineering describes leading development of a business-availability AIOps solution and having "[a]rchitected Agentic AI integrated with observability platforms, assisting developers with SRE principles and reducing support tickets through internal knowledge base ingestion," alongside "automated single-click workflows via Slack/UI, improving developer productivity by 30%." A director of product management for developer experience describes driving "AI-native development and tooling across the SDLC adopted by 80% of engineers, including an MCP registry, agentic workflows, and spec-driven development templates," reporting a 20%+ improvement in developer productivity. Templates and registries are codified practice: they make one team's way of working the default for the next.

- *Reusable data platforms & pipelines—building the data substrate codification requires.* A senior director of machine learning engineering describes delivering "structured

clickstream eventing, visitor context sharing, segmentation, fenced pricing, bot detection, identity classification, telemetry, knowledge graph, data discovery and federated governance," and unifying traveler identity across brands. A data stewardship director describes establishing "a data management program from the ground up" with a "hybrid" model pairing a central team with domain stewards across the organization.

- *Marketing & revenue-operations systems—making experimentation cheap enough to learn from.* The largest category in the firm's composition. A vice president for data platform and products describes "the consolidation of three brand-specific A/B testing platforms into one, resulting in a 5x increase in experimentation velocity"—infrastructure that converts commercial guesses into recorded results and improves future targeting and launches.

- *Firm-specific predictive & decision models—embedding marketplace learning into decision systems.* A technology vice president describes delivering "a world-class fraud detection decision engine" processing petabytes of data and "computing 100K+ complex features in milli-seconds, training and executing hundreds of models, and thousands of rules to make split-second risk decisions."

- *Reusable cloud, infrastructure & reliability operations—turning incident knowledge into an assistant.* A lead site reliability engineer describes leading "war rooms" for high-priority incidents and mentoring engineers on troubleshooting, then building an in-house capability that "replaced third-party limitations with a proprietary AIOps system capable of automated incident correlation and real-time revenue impact analysis," alongside a global "Gap Process" for around-the-clock operations pods.

- *Customer-service & support systems—codifying service knowledge.* A product vice president describes influencing the "transformation of the business unit from a traditional call center operation to an AI-powered conversational and self-service platform," delivering a multi-channel conversation platform.

Figure A.8 places these examples in the context of the firm's hiring, and the levels are the highest of any firm we examine—organization-capital-related AI roles reach over four percent of Expedia's US positions by 2024. That is what one expects of a firm whose product is software and whose workforce is concentrated in engineering rather than in stores, plants, or branches. Panel (a) shows the series rising almost continuously from 2010 with no discernible inflection at the arrival of generative AI, reaching roughly two and a half times the level of other AI roles. Expedia has been codifying operational knowledge for the whole period; the recent agentic work extends a long trajectory rather than starting one. Panel (b)

is consistent with the case. Reusable data platforms and pipelines and reusable AI/ML platform and MLOps work together account for a large share throughout—the shared stack that the company built before pursuing particular applications—alongside firm-specific predictive and decision models and marketing and revenue-operations systems, reflecting the marketplace pricing, fraud, and traffic-acquisition work described above. Customer-service and support systems form a smaller but steady component, consistent with the conversation platform. As at Nike, no generative-AI category appears: Expedia's codification runs through platforms, models, and structured operational data rather than through documents and language.

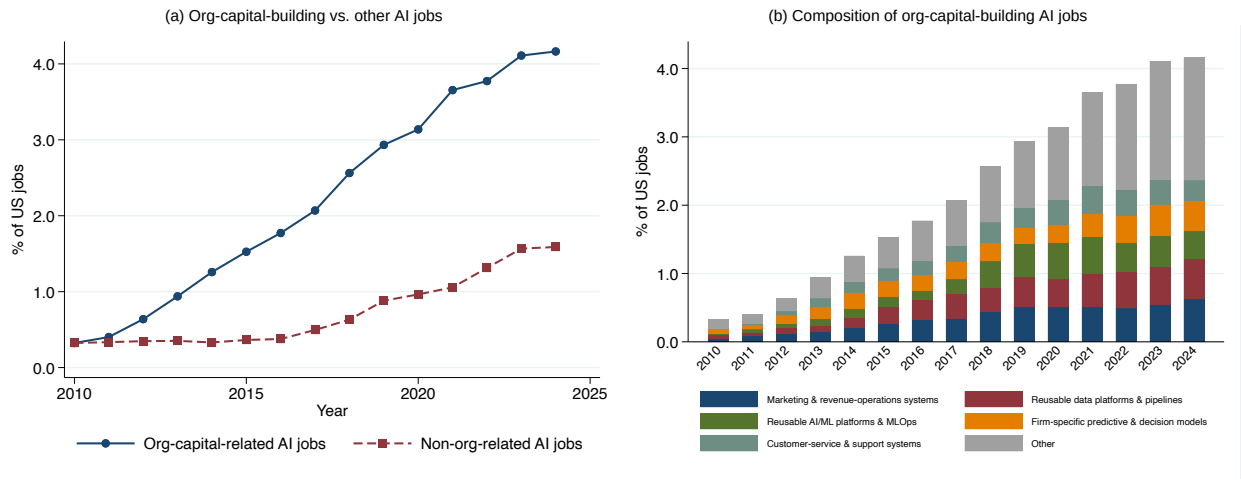


Figure A.8. AI jobs at Expedia Group as a share of the firm's US positions. Panel (a) compares organization-capital-related AI jobs with non-organization-capital-related AI jobs. Panel (b) decomposes the organization-capital-related jobs by reason for classification. Note the vertical scale differs across the corresponding figures for other firms.

The productivity channel is unusually well documented at Expedia, and it has a distinctive shape. Reported adoption of coding assistants rose from roughly two-thirds of developers with an estimated 20% productivity lift in early 2026 (Fortune, 2026) to 92% of engineers a few months later, with more than a tenth reported to be achieving two- to fivefold improvements (High, 2026). The dispersion is the interesting part. Where codified expertise compresses the gap between novice and expert in the cases we examine elsewhere—care representatives, store associates, underwriters—here the reported gains are concentrated among the already-strong, which is why the firm's response has been to codify what those engineers do rather than simply to buy more licenses. Both patterns are the same mechanism seen from opposite ends: the value comes from making one person's method available to everyone, and whether that raises the floor or the ceiling depends on whose method is being captured. The company also frames the resulting capacity as

substitutable for headcount at the margin, with its CTO suggesting that token spending may become a standard workforce budget line (Skift, 2026)—an unusually direct statement that codified capability and labor are alternatives—a possibility we take up in an extension of the model in Section D.6.

A.5. Measures of Productivity: Definitions and Construction

This Appendix provides definitions of our two firm-level TFP measures. Our baseline measure is revenue TFP (Appendix A.5.1), and an alternative measure based on the procedure of Imrohoroğlu and Tüzel (2014) is used in robustness checks (Appendix A.5.2.

A.5.1. Revenue TFP

Our revenue TFP is the residual from an industry-level production function estimated by OLS. For firm i in industry s and year t, we estimate the following specification separately for each industry,

$$\log y_{it} = \mu_i + \mu_t + \alpha_s^l \log(l_{it}) + \alpha_s^k \log(k_{it-1}) + \varepsilon_{it}, \tag{A.1}$$

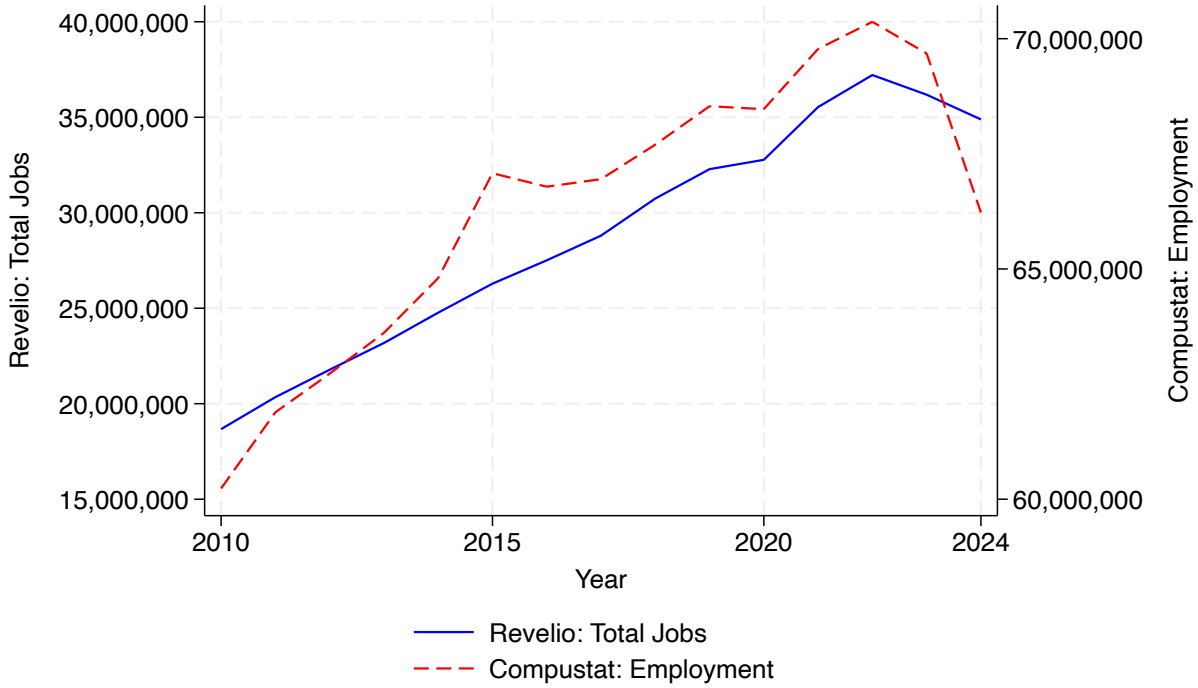
where y_{it} is real sales, l_{it} is employment, k_{it-1} is the beginning-of-period capital stock, and μ_i and μ_t are firm and year fixed effects. The capital stock is constructed using the perpetual inventory method. Revenue TFP is the estimated residual $\hat{\varepsilon}_{it}$. This measure is specific to Cobb–Douglas production functions.

A.5.2. TFP based on the control-function approach

We follow the implementation of Imrohoroğlu and Tüzel (2014) using the control-function approach by Olley and Pakes (1996) to construct an alternative measure of firm-level TFP. This approach addresses the correlation between firm's input choices and unobserved productivity. First, we measure each firm's real value added (sales less materials), labor input (number of employees), and capital stock from Compustat, with output and inputs deflated by industry and aggregate price indices. The capital stock is measured from gross property, plant, and equipment and deflated using investment prices from the years in which the capital was purchased. Second, we estimate a value-added production function in labor and capital separately for each industry, using the procedure of Olley and Pakes (1996), which uses firms' investment controls for the correlation between unobserved productivity and input choices. Third, firm-year TFP is the residual from this production function, equal to log value added minus the estimated contributions of labor and capital. To expand the sample size, we lower the minimum employment requirement from 100 to 50 employees, matching our baseline sample; we relax the requirement that a firm have at least two years of Compustat history; and rather than drop industry with

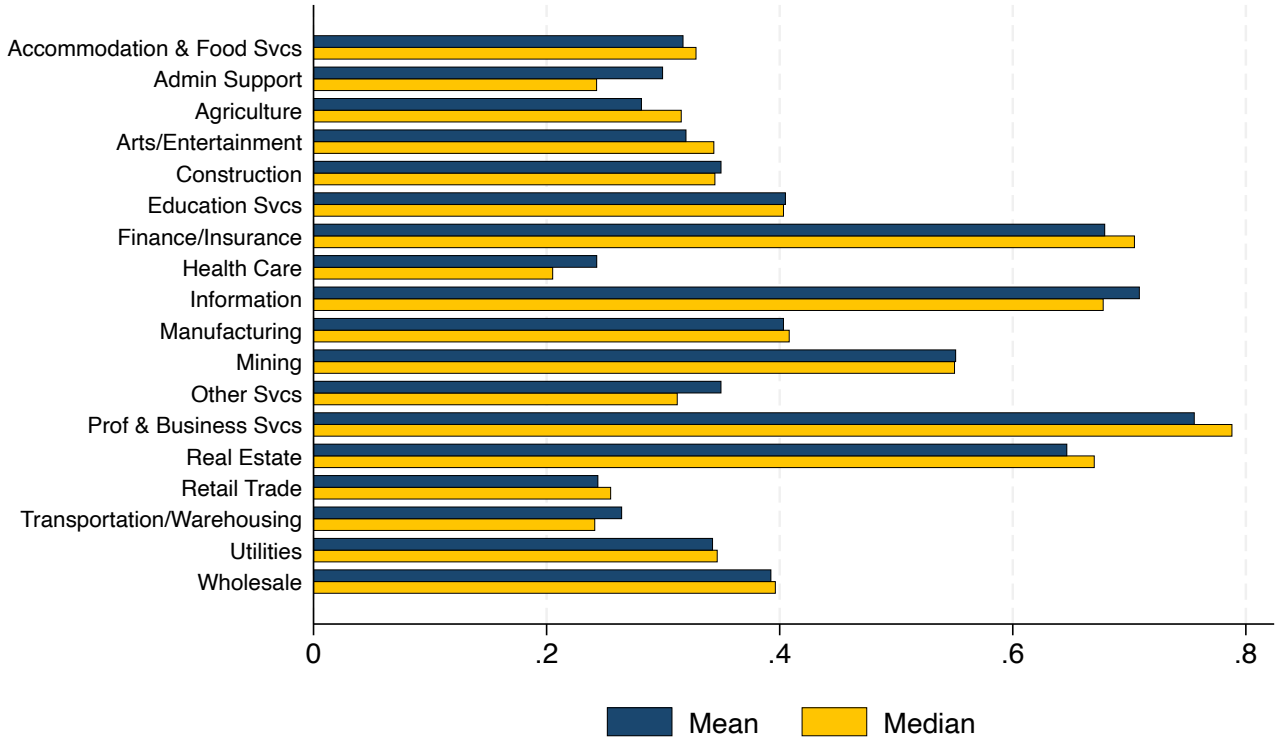
fewer than five firms, we pool them into their two-digit industry when the pooled industry exceeds five firms. Because the procedure imposes many sample restrictions and excludes industries where physical inputs are not well defined, the measure is available for a smaller subset of firms. For instance, it drops firms in the finance and utilities sectors, firms with missing investment, and firms with low sales or material costs. Our resulting productivity measure has a correlation of 0.986 with the productivity data in Imrohoroğlu and Tüzel (2014). We use it only in robustness checks, as a productivity measure that is methodologically independent of our revenue-TFP residual.

Figure A.9. Revelio and Compustat Employment for Matched Public Firms



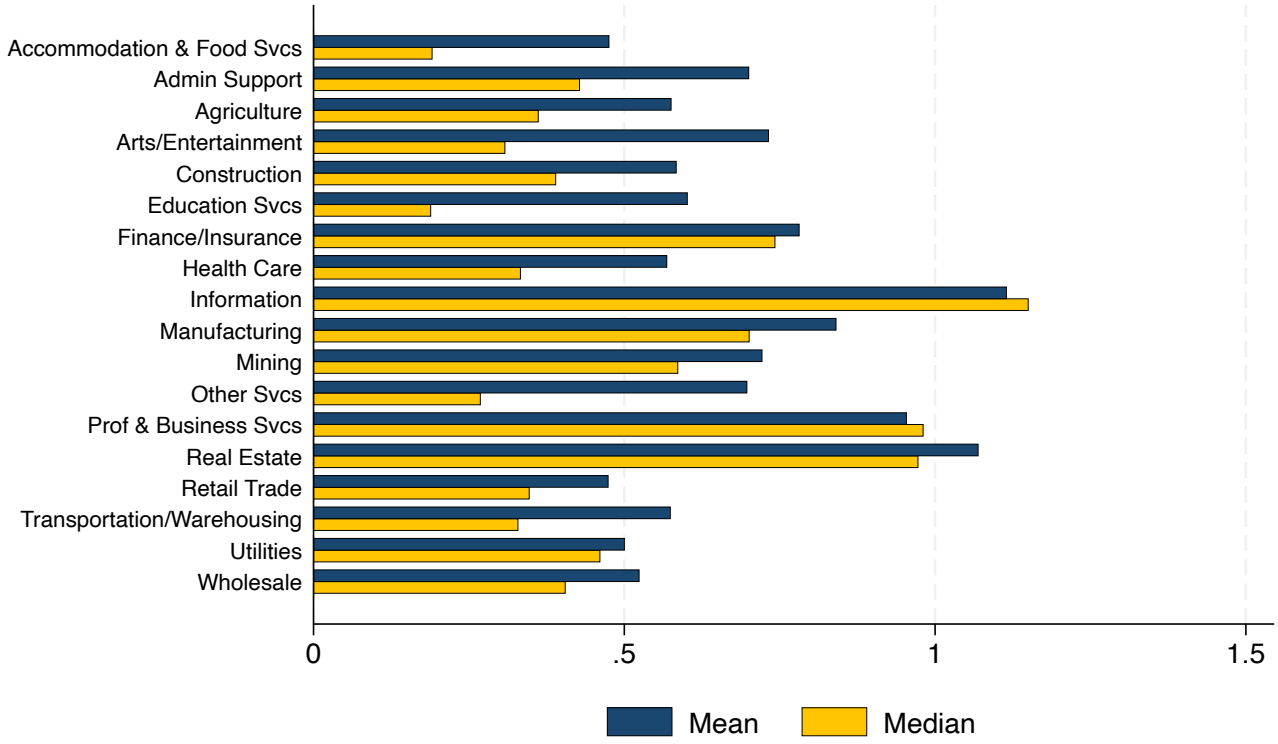
This figure compares aggregate employment in Revelio and Compustat for our matched U.S. public firms from 2010 to 2024. The solid blue line reports the total global Revelio jobs each year, and the dashed red line reports aggregate Compustat employment for the same matched U.S. public-firm sample.

Figure A.10. Industry-level Employment Coverage Ratios (Employment-weighted)



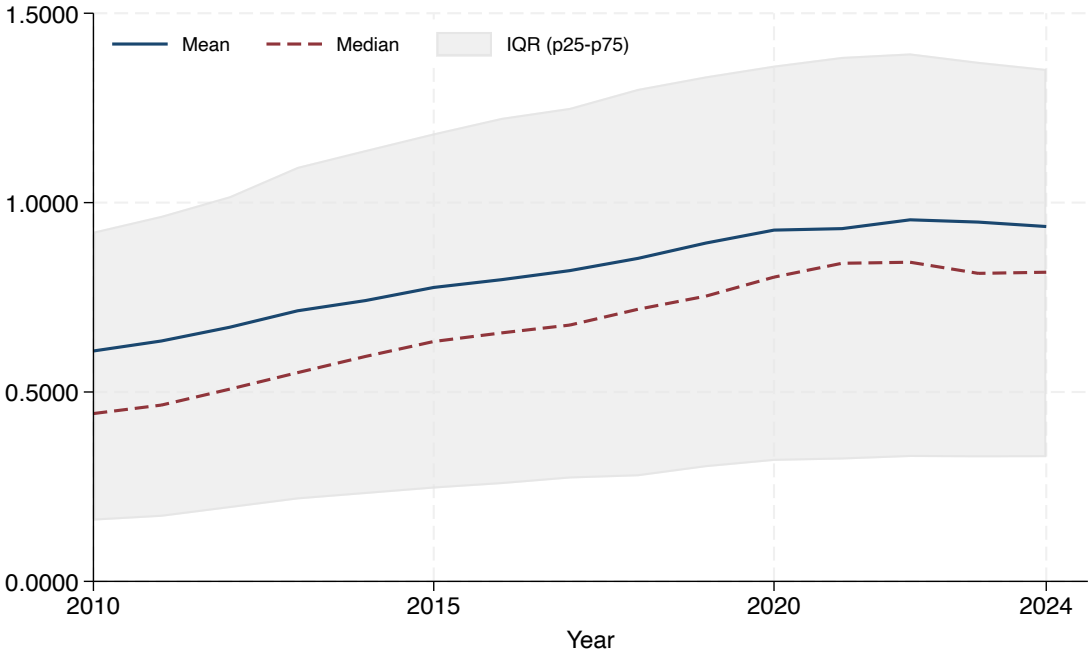
This figure reports industry-level employment coverage ratios for all matched U.S. public firms from 2010 to 2024. For each two-digit NAICS industry-year, the coverage ratio is defined as the sum of all global Revelio jobs divided by the sum of Compustat employment. Bars report the mean and the median of the annual industry-level coverage ratio over the sample period.

Figure A.11. Industry-level Employment Coverage Ratios (Equal-weighted)



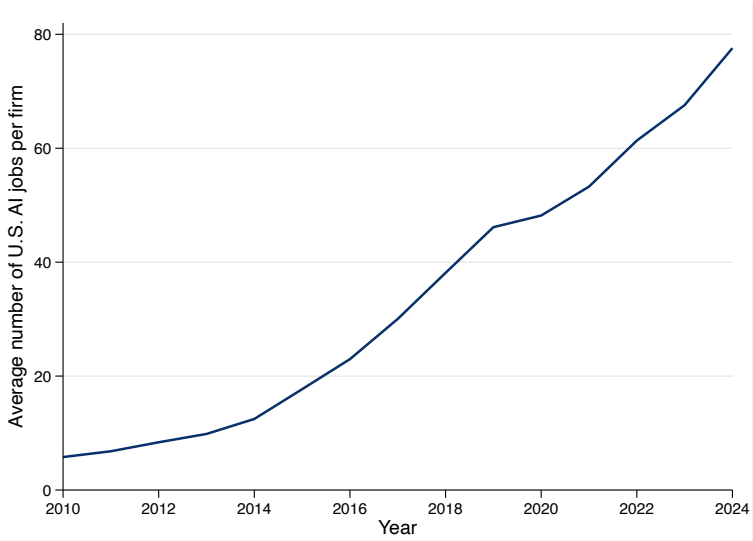
This figure reports equal-weighted industry-level statistics for the firm-level employment coverage ratio for all matched U.S. public firm-years from 2010 to 2024. The coverage ratio is defined as total global Revelio jobs divided by Compustat employment and is computed at the firm-year level. The ratio is winsorized within year at the 1st and 99th percentiles before computing industry statistics. Bars report the mean and median coverage ratio across all firm-year observations within each two-digit NAICS industry over the sample period.

Figure A.12. Firm-level Employment Coverage Ratio by Year

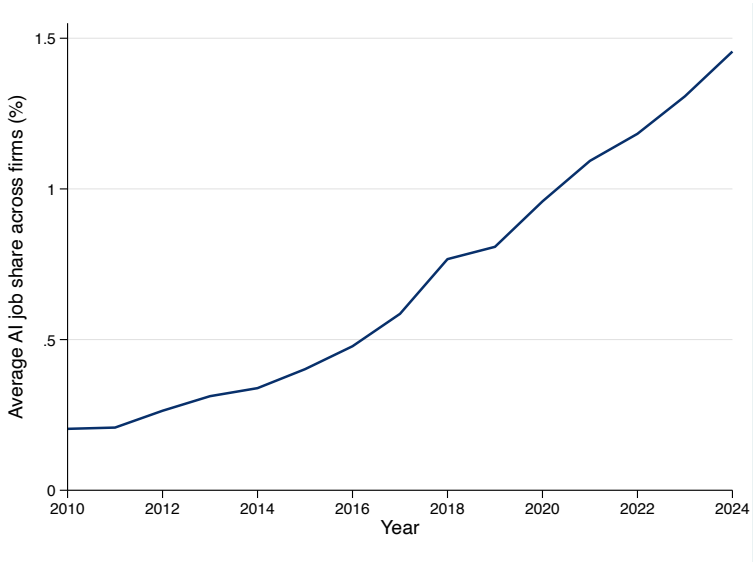


This figure reports annual distributional statistics for the firm-level employment coverage ratio for all matched U.S. public firm-years from 2010 to 2024. The coverage ratio is defined as total global Revelio jobs divided by Compustat employment and is computed at the firm-year level. The ratio is winsorized within year at the 1st and 99th percentiles before computing annual statistics. The solid blue line reports the cross-sectional mean by year, the dashed red line reports the cross-sectional median by year, and the gray shaded area denotes the interquartile range, defined by the 25th and 75th percentiles.

Figure A.13. Average Firm-Level AI-Related Jobs and AI Job Shares



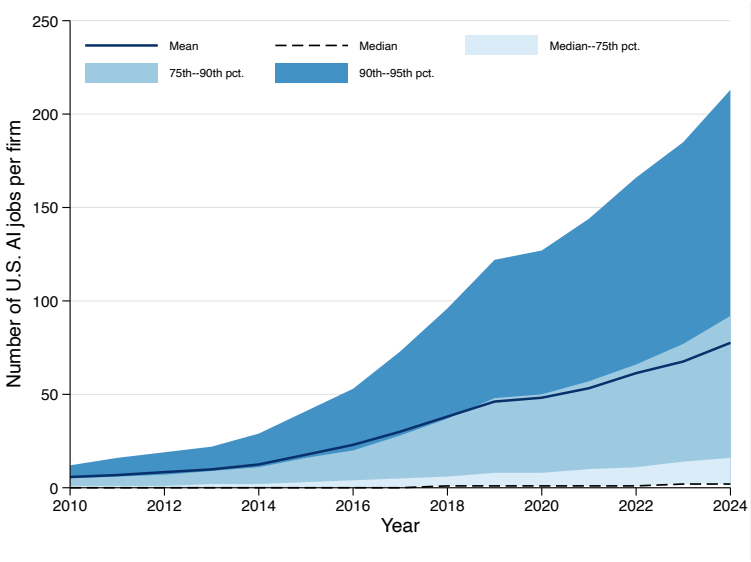
(a) AI jobs



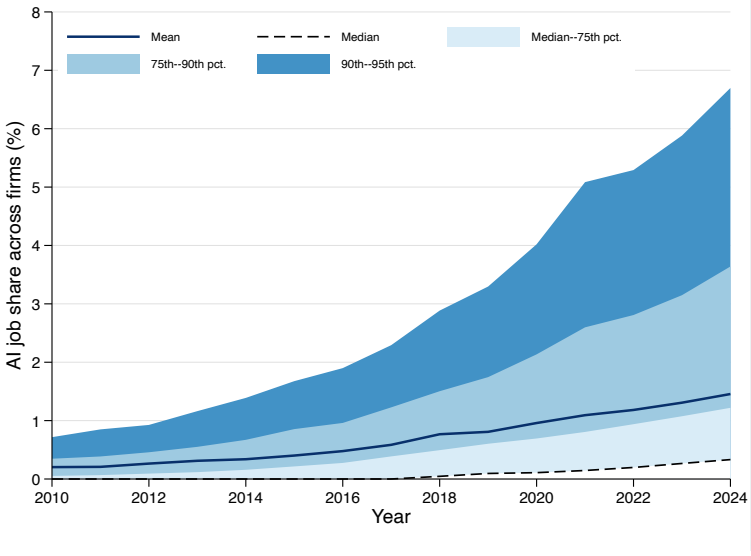
(b) AI shares

This figure reports average firm-level AI-related employment from 2010 to 2024 for our matched public-firm sample. Panel (a) plots the average number of AI-related U.S. jobs per firm. Panel (b) plots the average firm-level AI job share, defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. Averages are computed across firms within each year. AI-related jobs are identified using keyword matches in job titles or job descriptions. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.14. Distribution of Firm-Level AI-Related Jobs and AI Job Shares



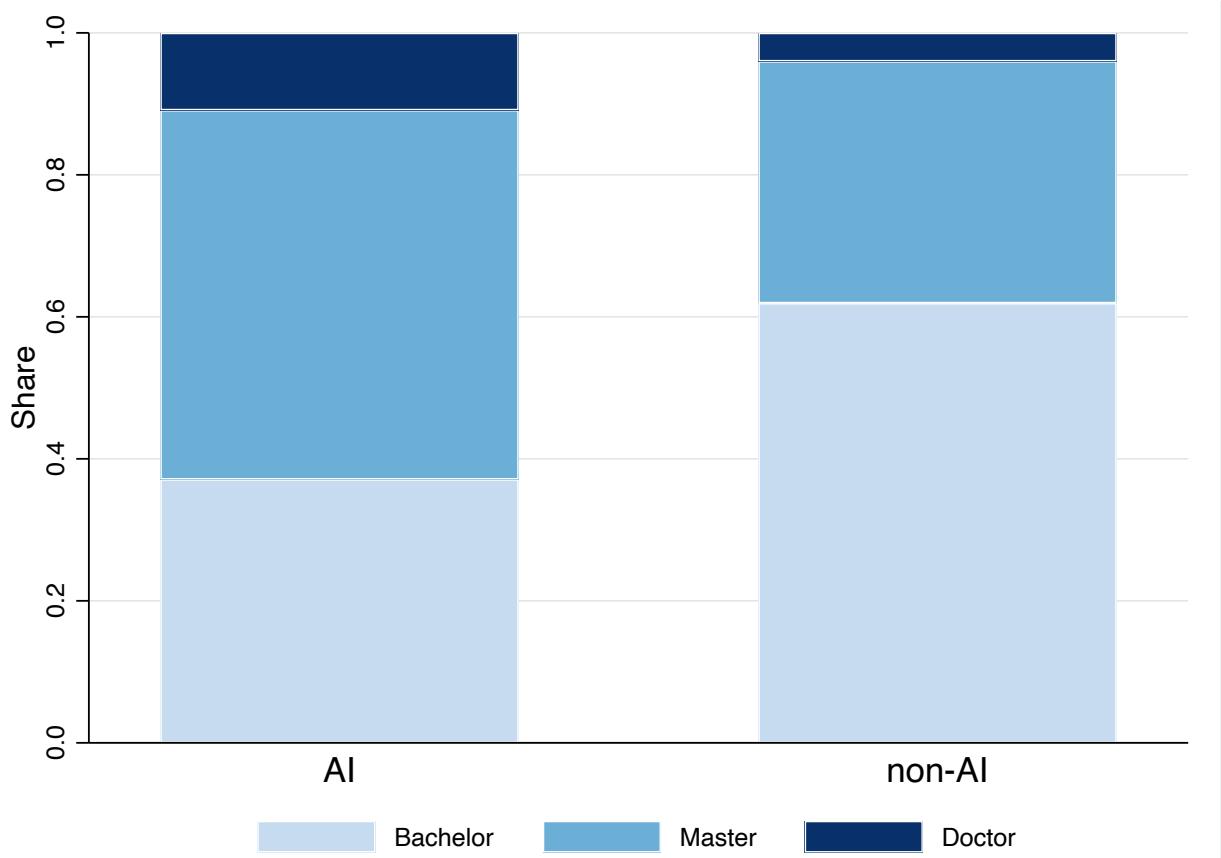
(a) AI jobs



(b) AI shares

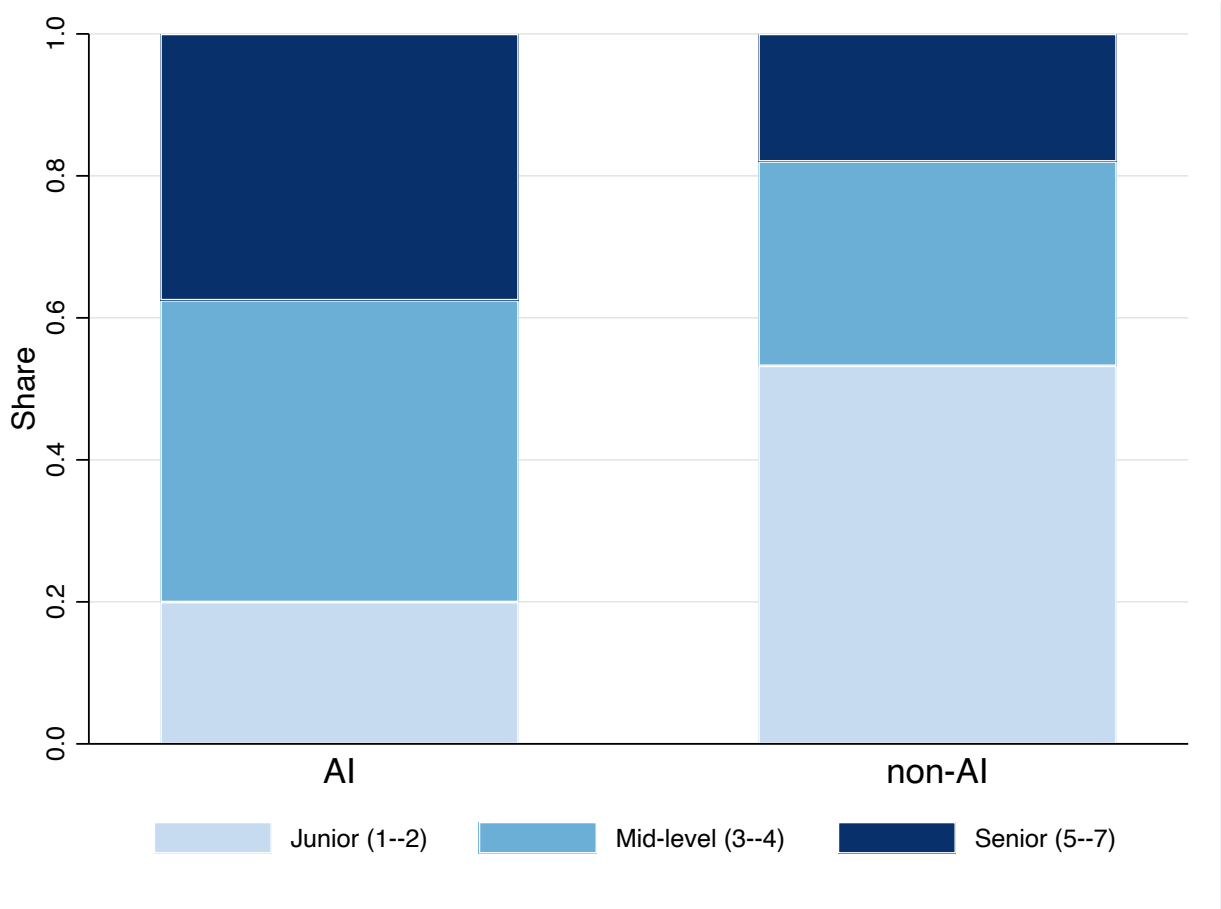
This figure reports the distribution of firm-level AI-related employment from 2010 to 2024 for our matched public-firm sample. Panel (a) plots the number of AI-related U.S. jobs per firm. Panel (b) plots the firm-level AI job share, defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. Statistics are computed across firms within each year. AI-related jobs are identified using keyword matches in job titles or job descriptions. The shaded areas report the median–75th percentile, 75th–90th percentile, and 90th–95th percentile ranges. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.15. Educational Attainment of AI-Related versus Non-AI-Related Workers



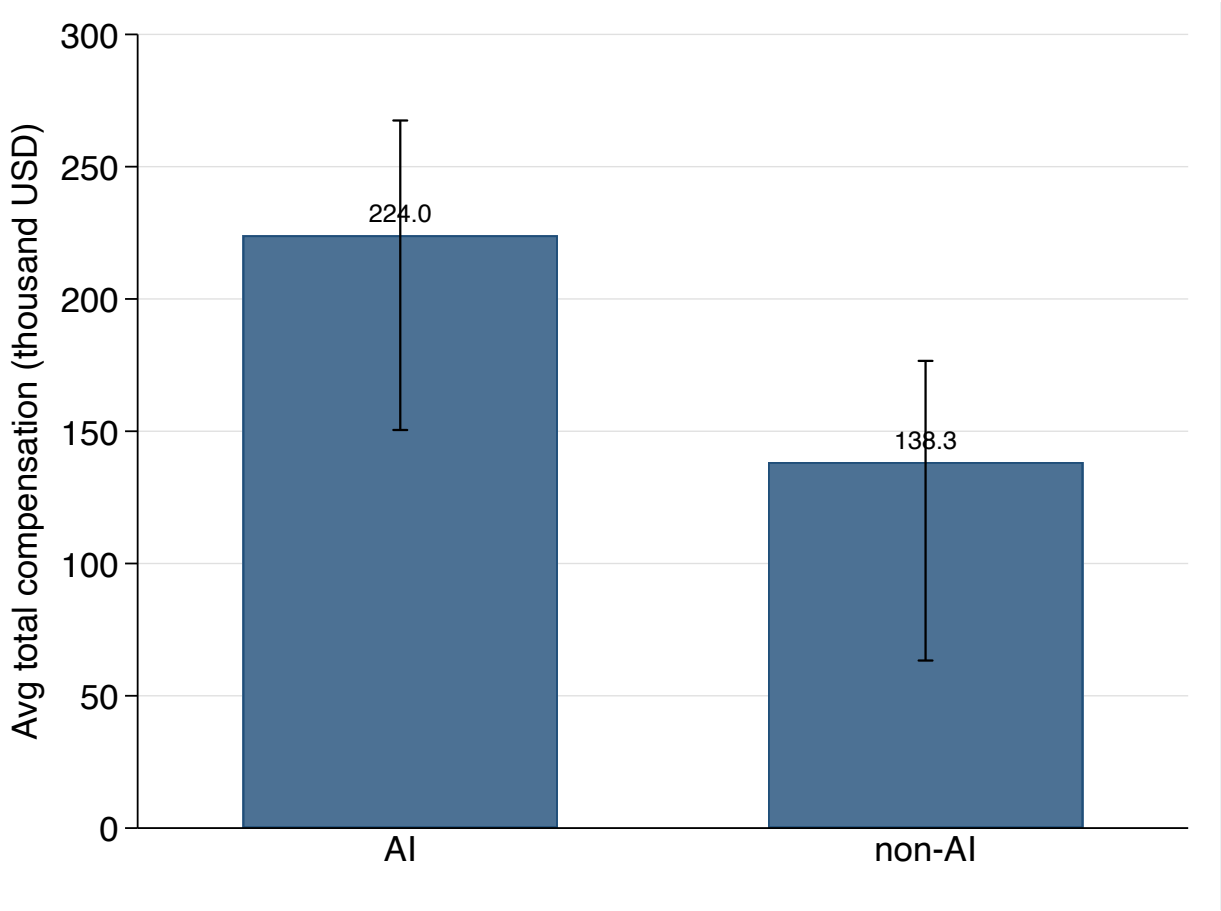
This figure plots the highest educational attainment of AI-related and non-AI-related workers among U.S. jobs at the job level. Each bar reports the share of U.S. positions within AI-related or non-AI-related jobs that falls into the corresponding degree category in 2024 for our matched U.S. public-firm sample. Bachelor includes bachelor's and associate degrees, Master includes master's degrees and MBAs, and Doctor includes doctoral degrees. Positions with missing or unrecognized education and high-school degrees are excluded. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.16. Seniority Distribution of AI-Related versus Non-AI-Related Workers



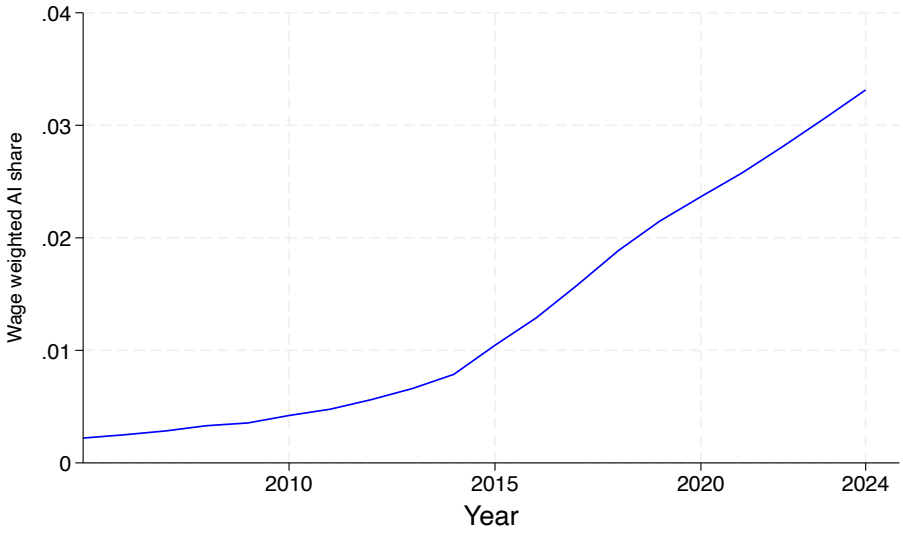
This figure plots the seniority distribution of AI-related and non-AI-related workers among U.S. jobs at the job level. Each bar reports the share of U.S. positions within AI-related or non-AI-related jobs that falls into the corresponding seniority category in 2024 for our matched U.S. public-firm sample. Seniority is based on Revelio's seven-level seniority measure: 1 is Entry Level, 2 is Junior Level, 3 is Associate Level, 4 is Manager Level, 5 is Director Level, 6 is Executive Level, and 7 is Senior Executive Level. We group levels 1–2 as Junior, levels 3–4 as Mid-level, and levels 5–7 as Senior. Positions with missing or unrecognized seniority are excluded. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.17. Average Total Compensation of AI-Related versus Non-AI-Related U.S. Jobs



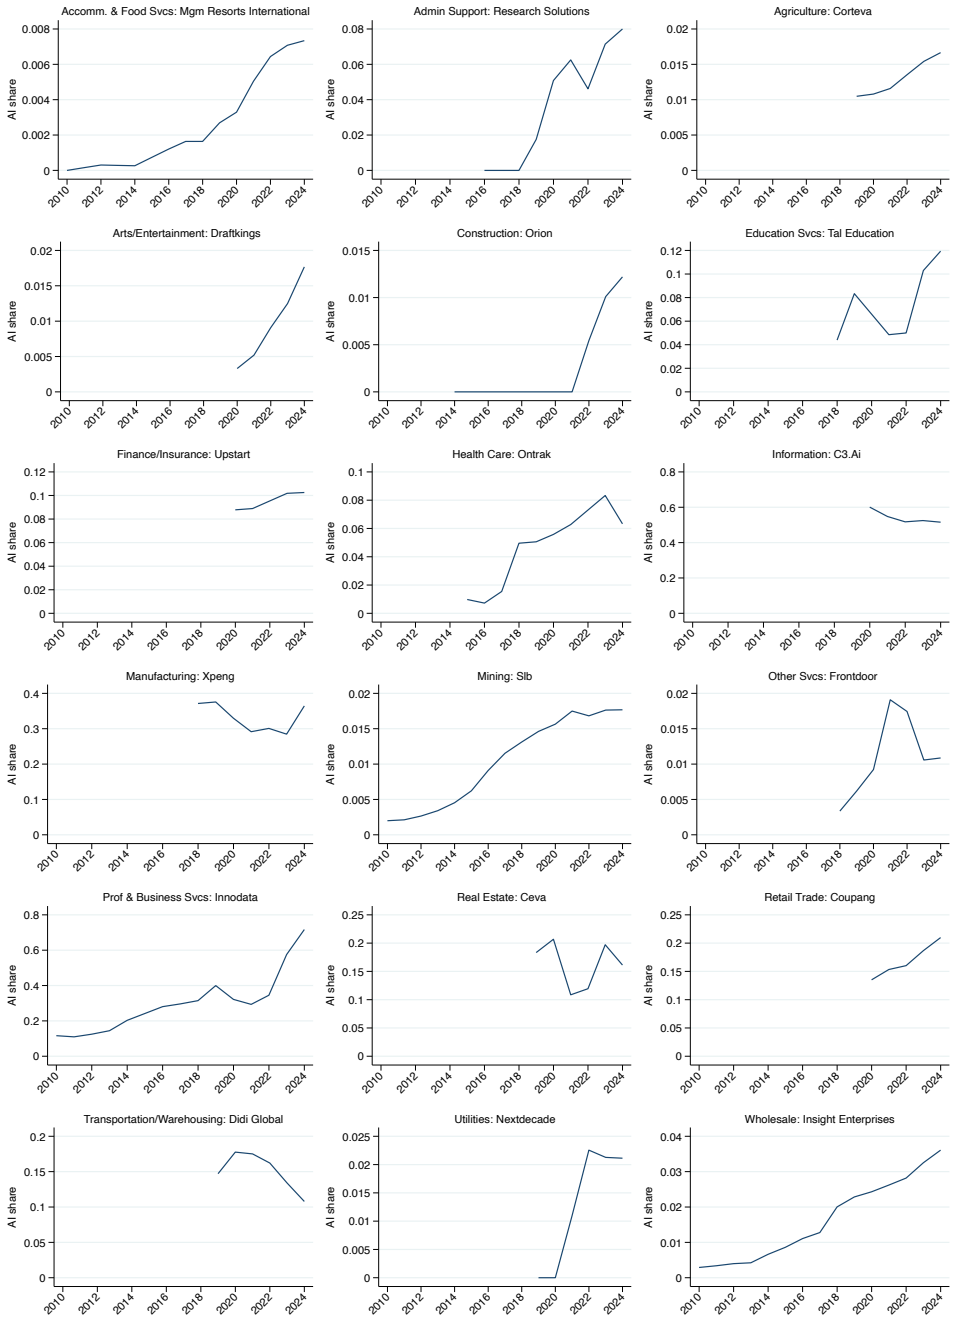
This figure plots average total compensation for AI-related and non-AI-related U.S. jobs in 2024 for our matched U.S. public-firm sample. Total compensation for each position is reported by Revelio. Bars report aggregate total compensation divided by the number of jobs, computed separately for AI-related and non-AI-related U.S. jobs, and expressed in thousands of U.S. dollars. Vertical lines show the 25th and 75th percentiles of job-level total compensation within AI-related and non-AI-related U.S. jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total Revelio global workers divided by Compustat employment for each firm-year.

Figure A.18. Wage-Weighted Share of AI Jobs



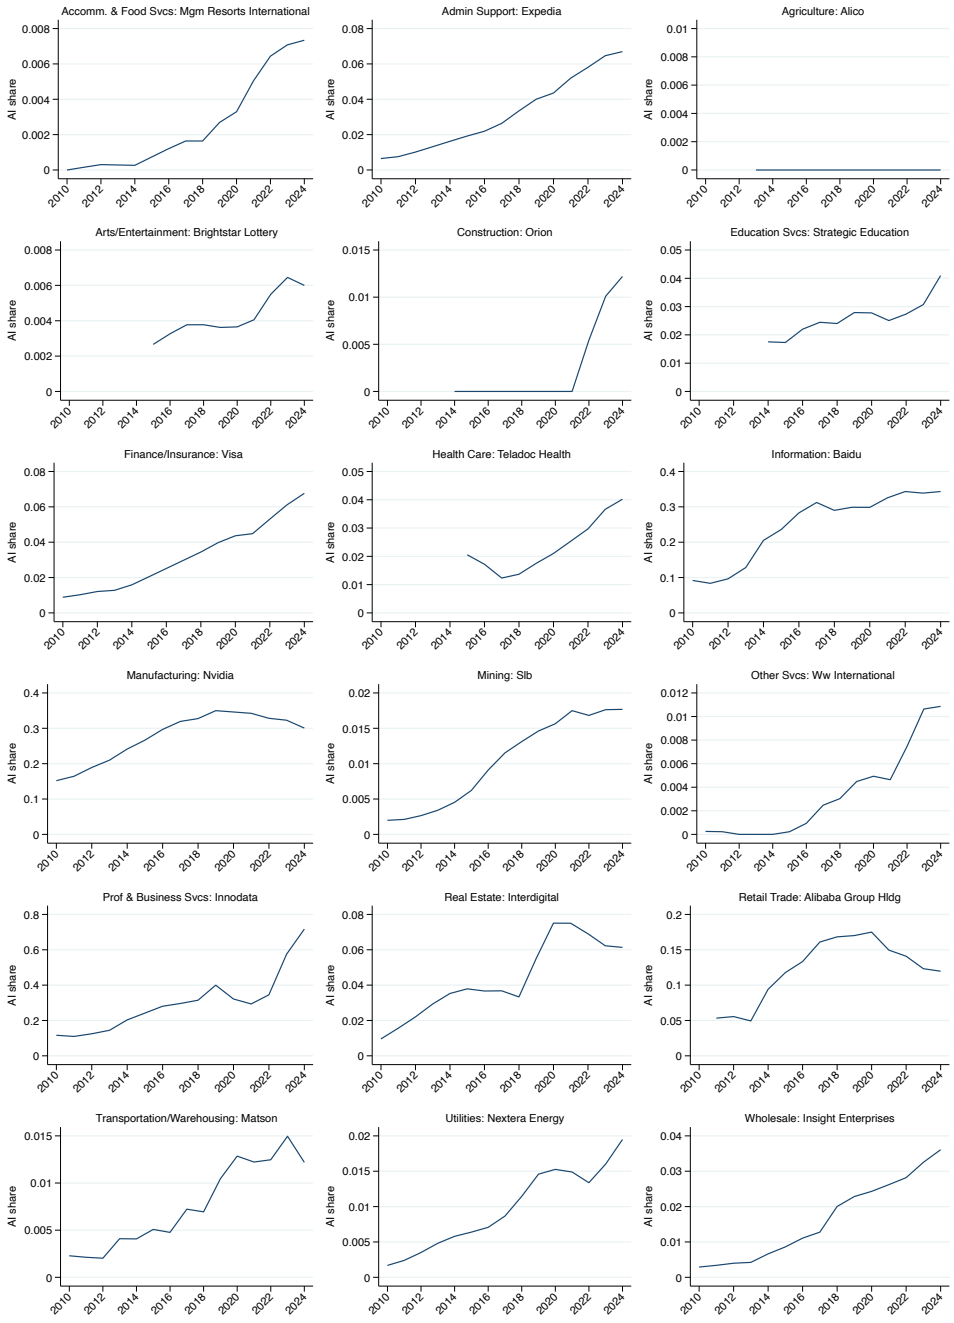
This figure plots the wage-weighted share of AI-related jobs among U.S. jobs in our matched U.S. public-firm sample from 2010 to 2024. We weight jobs by their annual compensation reported by Revelio. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where coverage is total Revelio global workers divided by Compustat employment.

Figure A.19. AI Share of 2024 Industry Leaders with Five-Year Continuous Histories



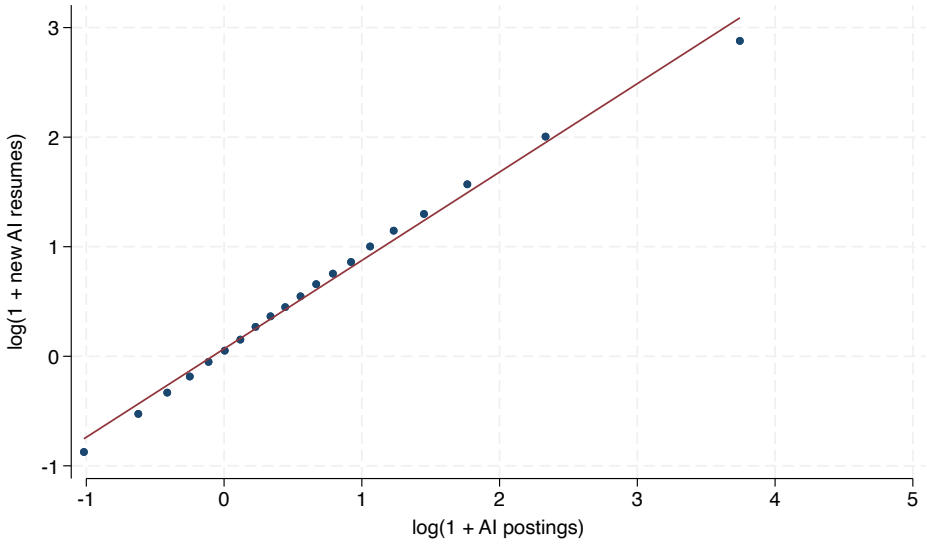
This figure plots the time series of AI share for the top firm in each two-digit NAICS industry from 2010 to 2024 for our matched U.S. public-firm sample. The top firm in each industry is defined as the public firm with the highest AI share in 2024 among firms observed continuously over the five-year period ending in 2024 and with at least 50 U.S. workers in Revelio and 50 employees in Compustat. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.20. AI Share of 2024 Industry Leaders with Ten-Year Continuous Histories



This figure plots the time series of AI share for the top firm in each two-digit NAICS industry from 2010 to 2024 for our matched U.S. public-firm sample. The top firm in each industry is defined as the public firm with the highest AI share in 2024 among firms observed continuously over the ten-year period ending in 2024 and with at least 50 U.S. workers in Revelio and 50 employees in Compustat. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

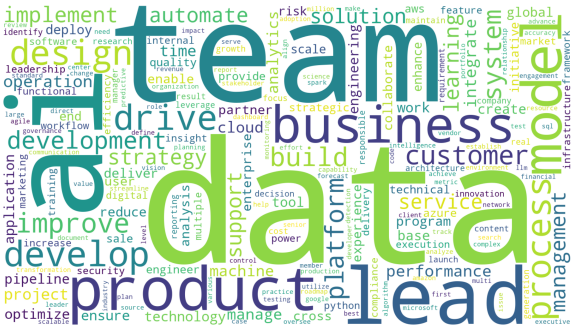
Figure A.21. AI-related Job Postings versus New AI Resumes



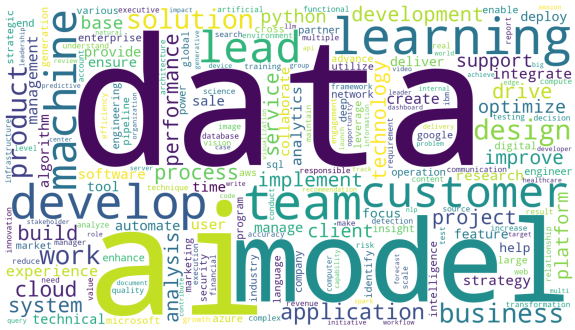
This figure shows a residualized binscatter of log(1 + new AI resume jobs) against log(1 + AI-related job postings) at the firm-year level, 2010–2024. A new AI resume job is a Revelio position that begins at firm f in year t and whose text matches an AI skill term; an AI-related posting is a Lightcast job ad posted by firm f in year t that lists one of the same AI skills. Both sides are identified by a common skill list: a skill enters the list in the first year it appears in the postings data and is applied to both sources from that year onward, and within each firm-year only skills observed in both the firm's postings and its new resumes are used, so the two measures reflect an identical, contemporaneous definition of AI. Both series are residualized on year fixed effects, the log of the firm's total job postings, and the log of the firm's total U.S. resume jobs.

Figure A.22. Word Clouds for Organization-Capital-Related and Non-Organization-Capital-Related Jobs

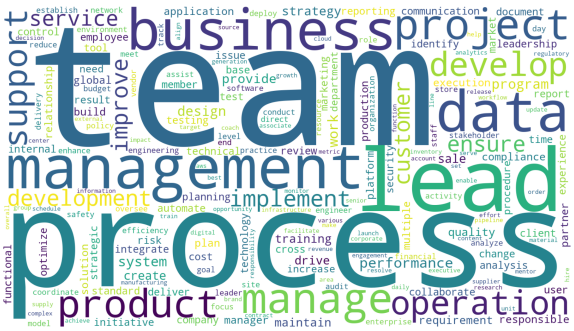
(a) AI, org-capital-related



(b) AI, not org-capital-related



(c) Non-AI, org-capital-related

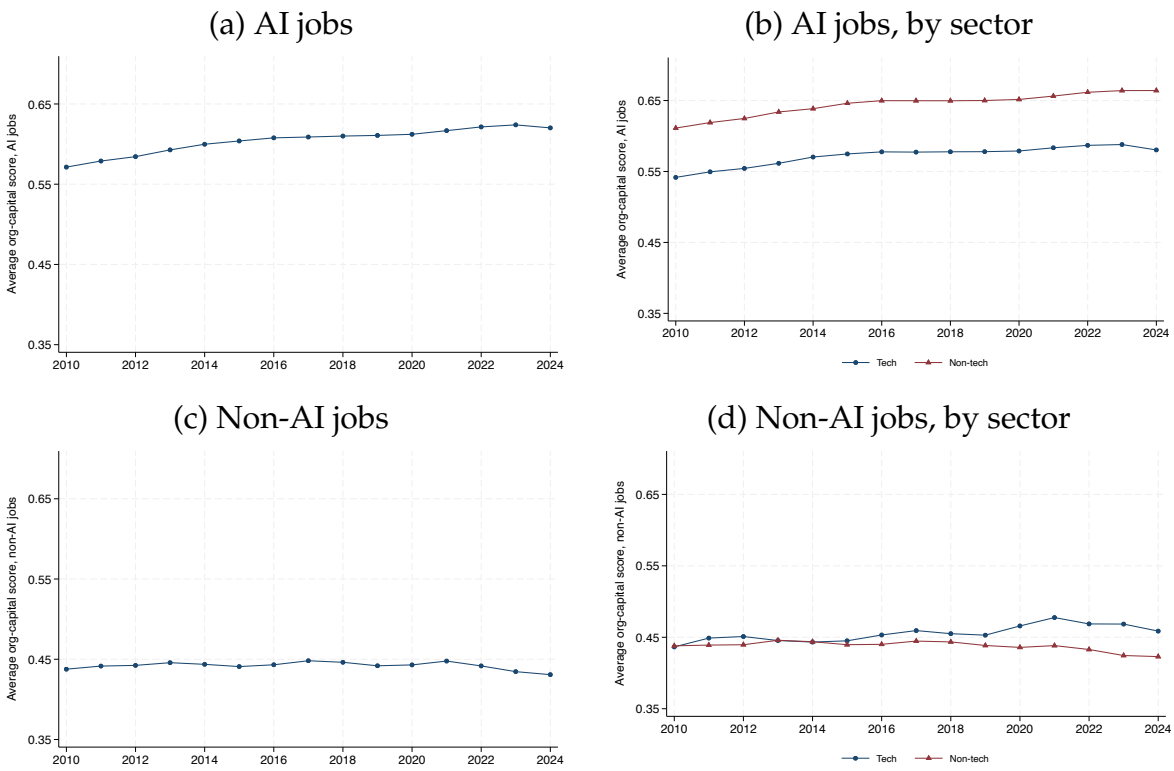


(d) Non-AI, not org-capital-related



This figure plots word clouds for four job groups defined by the cross of our AI-related job indicator and our organization-capital classification: (a) AI-related, organization-capital-related jobs; (b) AI-related jobs not classified as organization-capital-related; (c) non-AI, organization-capital-related jobs; and (d) non-AI jobs not classified as organization-capital-related. The word clouds are constructed from the job descriptions of U.S. jobs in our matched U.S. public-firm sample. For each group, words are extracted from positions assigned to that group based on our AI keyword matches (for AI-related jobs) and our LLM-based organization-capital classification. Words are lemmatized and related derivational forms are merged. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure A.23. Average Organization-Capital Score: AI versus Non-AI Jobs



This figure plots the continuous organization-capital score, which ranges from 0 to 1 and measures the extent to which a job builds organization capital. Panel (a) plots the average organization-capital score of all AI jobs. Panel (b) plots the same average score separately for tech firms (NAICS 51, 54 and semiconductor manufacturing) and all other firms. Panel (c) plots the average organization-capital score of non-AI jobs, based on the random sample of 1,000 non-AI jobs per year. Panel (d) plots that average score separately for tech and non-tech firms.

Table A.2. Employment Coverage Rates by Sample Characteristics

Sample	N	Mean	Std. dev.	p25	p50	p75
U.S. headquartered	58,260	86.0	72.3	30.0	71.9	128.8
Non-U.S. headquartered	14,329	63.1	66.6	10.2	43.2	100.5
U.S.-majority workers	51,411	86.5	70.7	31.1	73.4	129.4
Non-U.S.-majority workers	21,178	69.3	72.9	13.3	49.3	107.4

This table reports firm-year employment coverage rates for our matched U.S. public firms from 2010 to 2024. The coverage rate is defined as total global Revelio jobs divided by Compustat employment for the same firm-year and is expressed in percentage points. The rate is winsorized within year at the 1st and 99th percentiles before computing statistics. All reported moments are computed across firm-years. Rows report separate sample splits and are not cumulative. U.S.-headquartered firms are firms classified as headquartered in the United States. U.S.-majority worker firms are firm-years for which more than half of workers observed in Revelio are located in the United States.

Table A.3. Summary Statistics for Compustat Employment and Revelio Jobs

	Compustat Employment			Revelio Jobs		
Statistic	2010	2018	2024	2010	2018	2024
p1	5	4	5	1	1	2
p5	19	13	16	4	5	8
p25	193	149	151	60	81	98
p50	1,164	1,173	1,166	353	554	644
p75	6,200	7,400	7,500	1,780	3,186	3,803
p95	58,300	69,000	68,852	16,647	28,652	34,207
Mean	12,593	14,355	14,517	3,903	6,519	7,648
St.d	49,620	55,945	58,819	15,808	24,688	30,869
Skewness	19	17	17	13	10	13
Kurtosis	682	528	476	259	156	277
N	4,783	4,716	4,562	4,783	4,716	4,562

This table reports summary statistics for Compustat employment and total global Revelio jobs for our matched U.S. public firms in 2010, 2018, and 2024. Revelio jobs are the annual count of all jobs observed in Revelio for each firm-year. Statistics are computed across firms within each year and are unweighted.

Table A.4. Correlation Between Revelio and Compustat Employment by Employment Coverage Ratio

Coverage bin	N	Number of Firms	Total	Firm FE
All	72,589	9,086	0.610	0.618
< 0.05	6,894	1,278	0.674	0.328
0.05–0.2	9,137	1,898	0.972	0.469
0.2–1	33,602	5,321	0.904	0.873
1–2	20,224	4,011	0.987	0.937
2–3	1,672	795	0.998	0.991
≥ 3	1,060	361	0.609	0.376

This table reports Pearson correlations between total global Revelio jobs and Compustat employment for all matched U.S. public firm-years from 2010 to 2024. The employment coverage ratio is defined as total global Revelio workers divided by Compustat employment for the same firm-year. Coverage bins are assigned at the firm-year level. "N" reports the number of firm-year observations, and "Number of Firms" reports the number of unique firms in each bin. "Total" stands for the cross-sectional correlation that is computed in employment levels. "Firm FE" stands for the correlation that is computed after demeaning both employment measures by firm within each coverage-bin sample.

Table A.5. Correlation Between Revelio and Compustat Employment by Firm Size

Sample	N	Number of Firms	Total	Firm FE
All	72,589	9,086	0.610	0.618
Revelio Worker ≥ 50	57,404	6,983	0.607	0.623
Revelio Worker < 50	15,185	3,157	-0.016	0.056
Compustat employment ≥ 50	63,279	7,842	0.607	0.618
Compustat employment < 50	9,310	2,269	0.037	0.505

This table reports Pearson correlations between total global Revelio jobs and Compustat employment for all matched U.S. public firm-years from 2010 to 2024. Rows split the sample by firm size, using either total Revelio workers or Compustat employment for the same firm-year. "N" reports the number of firm-year observations, and "Number of Firms" reports the number of unique firms in each row. The cross-sectional correlation is computed in levels. The firm FE correlation is computed after demeaning total global Revelio jobs and Compustat employment by firm within each row's sample.

Table A.6. Firm Counts by Sample Filter

Year	Matched	Coverage $(0.05, 3)$	U.S. headquartered	Revelio U.S. workers ≥ 50	U.S.-majority	Compustat employment ≥ 50	All filters
2010	4,783	4,099	4,040	3,309	3,563	4,264	2,482
2011	4,747	4,102	3,972	3,308	3,514	4,231	2,471
2012	4,761	4,160	3,947	3,290	3,489	4,167	2,424
2013	4,928	4,350	4,078	3,391	3,617	4,244	2,491
2014	4,971	4,415	4,095	3,453	3,649	4,255	2,526
2015	4,866	4,331	3,985	3,430	3,548	4,184	2,483
2016	4,790	4,282	3,875	3,376	3,463	4,138	2,445
2017	4,727	4,235	3,807	3,349	3,387	4,102	2,386
2018	4,716	4,228	3,774	3,373	3,339	4,099	2,391
2019	4,650	4,168	3,685	3,371	3,267	4,078	2,362
2020	4,953	4,458	3,857	3,517	3,401	4,283	2,434
2021	5,196	4,687	4,032	3,690	3,535	4,541	2,550
2022	5,083	4,597	3,911	3,654	3,424	4,450	2,479
2023	4,856	4,382	3,715	3,473	3,218	4,252	2,335
2024	4,562	4,141	3,487	3,295	2,997	3,991	2,179

This table reports the number of matched U.S. public firms by year under alternative sample filters. "Matched" firms are firms in the matched Revelio–Compustat sample. The "Coverage" filter keeps firm-years with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year. "U.S.-headquartered" firms are firms classified as headquartered in the United States. "Revelio U.S. workers ≥ 50" keeps firm-years with at least 50 U.S.-based workers observed in Revelio. "U.S.-majority" firms are firm-years for which more than half of workers observed in Revelio are located in the United States. "Compustat employment ≥ 50" keeps firm-years with at least 50 employees in Compustat. Each filter column applies the listed condition separately. The final column applies all listed filters jointly.

Table A.7. Skills with Highest AI-Relatedness Measures in Lightcast Job Postings

#	Skills	AI-relatedness Score
1	Machine Learning	1.000
2	Artificial Intelligence	1.000
3	Natural Language Processing (NLP)	1.000
4	Computer Vision	1.000
5	Zero Shot Learning	1.000
6	Rweka	1.000
7	AgentGPT	1.000
8	Doc2vec	0.995
9	Few Shot Learning	0.991
10	fastText	0.989
11	Hugging Face Transformers	0.989
12	Confusion Matrix	0.986
13	Hugging Face (NLP Framework)	0.985
14	CrewAI	0.985
15	Seq2Seq	0.983
16	AI-Generated Code	0.983
17	Hyperparameter Optimization	0.981
18	Agentic Systems	0.980
19	Standard ML Library	0.977
20	PyTorch Lightning	0.977
21	Amazon Rekognition	0.977
22	Neural Architecture Compression	0.976
23	AWS Inferentia	0.976
24	Microsoft Cognitive Toolkit (CNTK)	0.976
25	Dbscan	0.976
26	Word2Vec Models	0.975
27	CatBoost (Machine Learning Library)	0.972
28	Caffe2	0.972
29	Apache MXNet	0.971
30	Feature Learning	0.969
31	Reinforcement Learning from Human Feedback (RLHF)	0.969
32	OpenAI Gym	0.966
33	BERT (NLP Model)	0.966
34	PyTorch (Machine Learning Library)	0.966
35	Adversarial Machine Learning	0.964
36	Vowpal Wabbit	0.963
37	Attention Mechanisms	0.962
38	Retrieval Augmented Generation	0.962
39	Open Neural Network Exchange (ONNX)	0.961
40	Long Short-Term Memory (LSTM)	0.961
41	Weaviate	0.960
42	Small Language Model	0.960
43	Text Summarization	0.957
44	Part-of-Speech Tagging	0.957
45	DeepSpeech	0.957
46	Autoencoders	0.956
47	SpaCy (NLP Software)	0.956
48	Word Embedding	0.956
49	Neural Architecture Search (NAS)	0.956
50	Keras (Neural Network Library)	0.955

This table lists the top skills in the Lightcast data ranked by the skill-level AI measure w_s^{AI}. For each skill, we report the fraction of jobs requiring that skill that also require one of the four core AI skills—artificial intelligence, machine learning, computer vision, and natural language processing. Only skills that appear in at least 50 job postings are included.

Table A.8. False-Positive Rates for AI-Related Jobs by Firm Size

Tercile	N	FPs	FP rate
Small	5,084	278	5.47%
Mid	5,084	172	3.38%
Large	5,084	103	2.03%

This table reports false-positive rates for AI-related U.S. jobs by firm-size tercile. The sample covers 2010–2024, excludes internships and part-time positions, and is restricted to matched U.S. public firm-years with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year. Firm-size terciles are assigned within year using total global Revelio job counts. All AI-related jobs at small-tercile firms are included. The mid- and large-firm samples are year-matched random draws from their respective tercile job-year pools, with draws matched to the small-firm sample by year. "N" reports the number of sampled position-years, "FPs" reports the number classified as false positives, and "FP rate" shows the percentage of false positives jobs.

Table A.9. False-Positive Rates for AI-Related Jobs by AI Subcategory

Subcategory	*N*	FPs	FP rate
AI	5,000	88	1.76%
Agentic AI	5,000	35	0.70%
CV	5,000	62	1.24%
Gen AI	5,000	24	0.48%
ML	5,000	80	1.60%
NLP	5,000	134	2.68%
Other	5,000	18	0.36%

This table reports false-positive rates from random samples of AI-related U.S. jobs by AI subcategory. For each subcategory, the sample includes 5,000 jobs from 2010 to 2024 in the matched U.S. public-firm sample, excluding internships and part-time positions. Subcategories are assigned using keyword matches in job titles or job descriptions. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year. "*N*" reports the number of sampled jobs, "FPs" reports the number classified as false positives, and "FP rate" shows the percentage of false positives jobs.

Table A.10. Annual AI-Related Job Shares by Technology Category

Year	(1) Share AI all (%)	(2) Share AI (%)	(3) Share ML (%)	(4) Share NLP (%)	(5) Share CV (%)	(6) Share Gen AI (%)	(7) Share Agentic AI (%)	(8) Share AI Other (%)
2010	0.2508	0.1167	0.1347	0.0338	0.0232	0.0083	0.0015	0.0140
2011	0.2753	0.1292	0.1507	0.0360	0.0249	0.0094	0.0017	0.0155
2012	0.3221	0.1517	0.1799	0.0418	0.0295	0.0110	0.0018	0.0191
2013	0.3733	0.1771	0.2141	0.0473	0.0320	0.0125	0.0021	0.0228
2014	0.4491	0.2152	0.2643	0.0563	0.0361	0.0156	0.0025	0.0284
2015	0.5968	0.2908	0.3594	0.0722	0.0460	0.0225	0.0035	0.0411
2016	0.7390	0.3769	0.4470	0.0851	0.0508	0.0294	0.0045	0.0539
2017	0.9195	0.4835	0.5591	0.1039	0.0594	0.0395	0.0061	0.0694
2018	1.1027	0.6010	0.6646	0.1232	0.0670	0.0544	0.0081	0.0893
2019	1.2697	0.7162	0.7527	0.1385	0.0724	0.0771	0.0115	0.1115
2020	1.4037	0.8252	0.8079	0.1464	0.0762	0.1075	0.0162	0.1321
2021	1.5426	0.9513	0.8431	0.1462	0.0758	0.1591	0.0244	0.1558
2022	1.7029	1.1125	0.8556	0.1403	0.0735	0.2450	0.0385	0.1758
2023	1.8845	1.2958	0.8610	0.1390	0.0694	0.3527	0.0596	0.1936
2024	2.0990	1.5008	0.8710	0.1367	0.0640	0.4514	0.0921	0.2106

This table reports annual AI-related job shares from 2010 to 2024 for our matched U.S. public-firm sample. All shares are computed using U.S. jobs only. Share is defined as the fraction of U.S. jobs whose title or description contains at least one AI keyword. "Share AI all (%)" uses the full AI keyword dictionary. The remaining columns report shares for six AI technology categories: general AI, machine learning (ML), natural language processing (NLP), computer vision (CV), generative AI (Gen AI), and agentic AI. A job is counted at most once within a category even if it matches multiple keywords in that category, but categories are not mutually exclusive across columns. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.11. Top 50 Keywords with the Most AI Jobs in 2024

Rank	Keyword	Job Counts	Rank	Keyword	Job Counts
1	AI	219,613	26	feature engineer	3,176
2	machine learning	66,932	27	ai agents	3,121
3	ML	59,366	28	prompt engineering	2,974
4	generative ai	23,930	29	xgboost	2,765
5	ai/ml	23,672	30	matplotlib	2,729
6	LLM	18,986	31	chatgpt	2,650
7	artificial intelligence	18,549	32	conversational ai	2,618
8	genai	17,995	33	mlflow	2,431
9	NLP	10,046	34	retrieval-augmented generation	2,358
10	deep learning	9,959	35	random forest	2,337
11	LLMs	9,664	36	vertex ai	2,195
12	computer vision	7,313	37	sentiment analysis	2,084
13	language model	6,717	38	aws sagemaker	1,982
14	large language model	5,508	39	reinforcement learning	1,910
15	numpy	5,483	40	microsoft copilot	1,897
16	tensorflow	4,982	41	BERT	1,871
17	mlops	4,891	42	machine learning technique	1,842
18	agentic ai	4,715	43	ai agent	1,826
19	pytorch	4,477	44	github copilot	1,786
20	natural language processing	4,088	45	azure openai	1,783
21	scikit-learn	4,020	46	ml/ai	1,741
22	langchain	3,823	47	ai innovation	1,646
23	machine learning algorithm	3,488	48	gpt-4	1,628
24	neural network	3,331	49	scipy	1,544
25	ai research	3,313	50	hugging face	1,458

This table reports the 50 AI keywords with the largest number of job-level matches among U.S. jobs in 2024 for our matched U.S. public-firm sample. A job is counted once for a given keyword if the keyword appears in either the job title or the job description. A job can match multiple keywords, so counts are not mutually exclusive. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.12. Top 10 Keywords with the Most AI Jobs by AI Subcategory in 2024

AI		**ML**	
Keyword	Job Counts	Keyword	Job Counts
AI	219,613	machine learning	66,932
artificial intelligence	18,549	ML	59,366
ai research	3,313	deep learning	9,959
ai innovation	1,646	numpy	5,483
ai security	951	tensorflow	4,982
ai development	933	mlops	4,891
ethical ai	573	pytorch	4,477
artificial general intelligence	451	scikit-learn	4,020
explainable ai	403	machine learning algorithm	3,488
ai safety	356	neural network	3,331
NLP		**CV**	
Keyword	Job Counts	Keyword	Job Counts
NLP	10,046	computer vision	7,313
natural language processing	4,088	object detection	1,115
conversational ai	2,618	image recognition	625
sentiment analysis	2,084	opencv	624
BERT	1,871	CNNs	386
AIML	1,257	image segmentation	207
spacy	772	3d reconstruction	189
nltk	752	scene understanding	143
NLTK	703	pose estimation	111
natural language understanding	668	object recognition	101
Gen AI		**Agentic AI**	
Keyword	Job Counts	Keyword	Job Counts
generative ai	23,930	agentic ai	4,715
LLM	18,986	langchain	3,823
genai	17,995	ai agents	3,121
LLMs	9,664	ai agent	1,826
language model	6,717	langgraph	1,053
large language model	5,508	copilot studio	941
prompt engineering	2,974	agentic system	327
chatgpt	2,650	autogen	276
retrieval-augmented generation	2,358	intelligent agent	255
microsoft copilot	1,897	crewai	249

This table reports the 10 AI keywords with the largest number of job-level matches in 2024 within each of six AI subcategories. The sample includes U.S. jobs in our matched U.S. public-firm sample. A job is counted once for a given keyword if the keyword appears in either the job title or the job description. A job can match multiple keywords, so counts are not mutually exclusive. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.13. Examples of AI-Related Jobs by AI Subcategory

Example	Job title	Job description
Panel A: General AI		
1	Director, Medical Affairs Oncology **AI Innovation**	Driving **AI Innovation**: Leading the integration of **artificial intelligence** into Medical Affairs to optimize workflows, enhance decision-making, and deliver impactful insights. Empowering Oncology Teams: Providing **AI**-driven tools and solutions that enable the oncology team to address complex challenges and improve patient outcomes. Strategic Leadership: Developing a comprehensive **AI** roadmap that aligns with business goals and ensures scalable and sustainable solutions. Evidence Generation: Leveraging **AI** for real-world evidence (RWE) analysis and predictive analytics to support clinical and therapeutic advancements. Collaboration Across Functions: Partnering with cross-functional teams, including R&D, regulatory, and commercial, to align **AI** initiatives with broader organizational objectives. Stakeholder Engagement: Building relationships with internal and external stakeholders to drive the adoption and trust of **AI**-driven solutions in oncology. **Ethical AI** Implementation: Ensuring all **AI** solutions are compliant, unbiased, and focused on improving patient care.
2	Information Architect, UI/UX Design Lead	Transforming user complexity into intuitive digital experiences that deliver measurable results. Selected to transition from 3M to Solventum, helping build the new company's digital product foundation. • Leading Vision-to-Execution Design: Owning design strategy from concept to launch, aligning the team on visually compelling interfaces for the best user experience. • Bridging Design & Development: Facilitating collaboration across teams to ensure the design vision is translated flawlessly into shippable products. • Designing Intuitive User Journeys: Architecting the underlying structure of digital products, championing user needs to boost engagement and retention. • Integrating **AI** in the Design Workflow: Leveraging **AI** to automate tasks, empowering the team to accelerate prototyping and focus on innovation.
3	Senior Manager, Field & Account Based Marketing	Oversees Zillow's largest execution-focused B2B marketing team, leading Field and Account-Based Marketing strategy and delivery for top partner segments. Owns strategy and leadership for Field and Account-Based Marketing programs that drive product adoption across Zillow's mortgage offerings and real estate CRM platform, contributing to record-setting month-over-month transaction growth. Hosted Zillow's 2024 National Home Buyer Webinar Series, leveraging real estate expertise and on-camera experience to strengthen Zillow's consumer education presence. Embeds **AI** strategy into team operations (e.g., **AI** voiceovers, market updates, and enablement content) to accelerate execution and improve team efficiency. Restructured team model to scale with business needs, improved cross-functional alignment, and built internal leadership capacity. Is recognized across Marketing and Sales leadership as a performance-focused, high-impact people leader.
4	Senior Director, Finance Data Strategy Leader	Lead the design and implementation of the global Financial Data Strategy to modernize finance through data integration, automation, and advanced analytics. Build and scale a financial data platform using cutting age technologies, enabling seamless access to ERP and non-ERP data for real-time reporting, deep analytical caoabiities, and decision support. Drive Digital Transformation initiatives focused on process optimization, **AI** integration, and automation of financial data pipelines. Establish Data Governance frameworks to ensure data quality, security, and compliance across the finance organization. Develop and lead a Finance Upskilling program to enhance capabilities in data literacy, analytics, and technology adoption. Define KPIs, success metrics, and governance models to measure impact, improve operational efficiency, and support a culture of continuous innovation.
5	VP, Retail and Digital Technology	Leading Product and Engineering functions across our Digital and Shop technologies as a proven technology and product leader who led a 30+ person team to deliver digital solutions that fueled 33% revenue growth, scaling the business to $1.28B in just 15 months. Drove mobile platform adoption with 75% loyalty usage and 15% channel growth, while aligning closely with C-level stakeholders to unlock $120M in new revenue through mobile ordering. Modernized product delivery by introducing agile frameworks, dual-track workflows, and CI/CD pipelines-accelerating time to market and improving delivery confidence. Applied **AI** and automation to cut support call volume by 20% and reduce handle time by 30%. Built high-performing cross-functional teams, overhauled legacy systems to reduce operational costs by 30%, and implemented a hybrid talent model that lowered development costs by 20%. Known for delivering scalable strategy, strong execution, and measurable business value.
6	Senior Research Scientist	Research on Digital Humans, 3D Generative Models, and Embodied **AI**.
7	Vice President, Product Manager	Wealth Plan Driving innovation in digital wealth management by building and launching Wealth Plan, a personalized financial planning experience used by millions of clients. Partnered with engineering, design, and data science to deliver seamless goal-setting, **AI**-driven insights, and intuitive user experiences that empower clients to plan, track, and achieve their long-term financial goals. Skilled at transforming complex financial products into simple, client-first digital solutions at scale.
8	SVP, Chief Data Officer & Enterprise Innovation	Driving the business process optimization governance frameworks that guide and monitor the relationship between Xerox and Business Process Outsourcing provider(s), and designing and developing a holistic data and digital strategy across the organization. • Leading **AI** & Data Governance Center of excellence • Developing data and digital strategy, including master data management and governance processes to execute across operational units. • Creating a roadmap for an advanced data infrastructure, including a data warehouse, data lake, ingestion pipeline, and accessible APIs, to drive efficiencies across the organization. • Driving Business Process Outsourcing (BPO) strategy and governance,

Continued on next page

Table A.13 continued

Example	Job title	Job description
9	Innovation Manager	Develop virtual reality + **artificial intelligence** solutions to improve the interactive experience of intelligent VR scenes.
10	Lead Product Designer	• Lead Product Designer: • Contextual **AI** recommendations • Compute Usage Optimizer • Customer Support Homepage widget • **AI**-powered Logs Analysis sidekick • Materialized Views
Panel B: Machine Learning		
1	Measurement and Insights	Analyzed qualitative and quantitative data in Databricks and Tableau to identify the impact of sponsorship at Live Nation festivals and events. Distilled and organized information for decks from raw data for Fortune 500 partners and event deliverables. Built a **random forest** feature importance model, using Python to identify factors highly correlated with high fan satisfaction, which informed future resource allocation and brand marketing strategy
2	NLP Data Scientist & **ML** Engineer	• Achieved a 15x reduction in data processing costs by building RAG architecture with embedding and re-ranking models, accelerating drug discovery insights. • Enabled data extraction from scientific papers PDFs by implementing BERT-based document chunking. • Translated business goals into technical steps by coordinating with stakeholders on 5 targeted projects.
3	Data Scientist	At UBS Wealth Management, I build and deploy AI/**ML** models to strengthen portfolio risk management and optimize asset allocation. Using **XGBoost**, **LightGBM**, and statistical techniques, I developed predictive models and risk scoring algorithms that improved returns and reduced high-risk exposure. I collaborate with portfolio managers and analysts to translate complex data into actionable investment insights, while ensuring compliance with GDPR and SOX. I also design dashboards in Power BI and Tableau and maintain ETL pipelines in Python and SQL, enabling real-time decision-making supported by accurate, secure, and scalable data pipelines.
4	Vice President - AI/**ML** Product Manager - Portfolio Analysis Group	• Leading AI/**ML** product strategy for PAG, partnering with PAG analysts, front office, and risk teams to modernize exception management and performance attribution workflows. • Defined enterprise AI/**ML** roadmap for exception management and attribution workflows, aligning initiatives with Asset Management strategic priorities and presented results to senior PAG leadership. • Productionized a **random forest**-based after-tax outlier review solution, reducing workload by 2 FTEs. • Develop anomaly detection models for AM WRAP Composites, with the intention on reducing false positives. • Building NLP-based agent interactions to enable Front Office teams to query and derive insights from attribution datasets. • Oversee proof-of-concept initiatives for document extraction and composite description evaluation, scaling across PAG teams. • Conduct demos and training sessions for Front Office teams on new capabilities as well as gathering feedback.
5	Sr. Principal Systems Engineer	Working on advanced **neural network**s for radar and edge AI sensor processing. Working with external SMEs and Raytheon Fellows from across the country. Company AI/**ML** instructor. Codeveloped a tutorial on event-based sensing and **spiking neural network**s (SNNs).
6	Staff **Machine Learning** Engineer	Eater Incentive & Pricing team. I built the **ML**-based solutions to grant the right promotions/fees to the right users/cohorts. Applied the latest **ML** techniques such as sequential modeling, causal modeling et.al. to maximize the uplift of revenue and minimize the cost of promos, by giving personalized promotions to the consumers.
7	**Machine Learning** Engineer	Engineered an end-to-end ETL pipeline to process 52,000+ songs using the Spotify API, automating extraction and storage in PostgreSQL. Performed extensive **feature engineering** and data normalization for downstream **ML** applications Conducted a comprehensive **ML** study leveraging to disentangle and quantify the effects of audio attributes (e.g., tempo, danceability, acousticness) and metadata (e.g., genre, release year) on Spotify song popularity. **Pytorch** was used for building a ANN to model the relationship between 17 features. Incorporated model diagnostics, variance inflation factor analysis, and out-of-sample validation to ensure robustness; insights contributed to understanding content personalization and music recommendation strategies
8	**Machine Learning** Engineer	• Worked on team of 3 engineers to develop automated data collection pipelines for 14+ data sources using Python and SQL with Databricks platform; pipelines served as foundation for **ML** models to provide key marketing insights to downstream teams. • Spearheaded agile scrum meetings and serve as liason for team by successfully managing stakeholder communications, cross-functional collaboration, internal documentation, and vendor relations.

Continued on next page

Table A.13 continued

Example	Job title	Job description
9	Product Manager	• Led end-to-end product development of an AI-driven embedded finance platform, enabling partner applications to deliver real-time credit decisions via APIs, reducing loan approval time from days to under 5 seconds. • Defined product vision, strategy, and roadmap aligned with business objectives; collaborated with engineering, data science, and compliance teams to ensure innovation within FCRA, ECOA, and GDPR frameworks. • Partnered with data scientists to design and deploy an explainable **ML**-based credit-scoring model, achieving 94% predictive accuracy and transparent outputs using SHAP-based interpretability. • Managed a backlog of 40+ features, prioritized initiatives using the RICE framework, and delivered the MVP 20% ahead of schedule while maintaining enterprise-grade reliability and uptime. • Launched a self-service partner dashboard enabling instant API integration, credit analytics, and performance monitoring - improving partner onboarding efficiency by 30%. • Established measurable success metrics and continuous feedback loops, tracking KPIs such as approval-rate lift (+15%), model drift, and partner retention to drive iterative product optimization.
10	Data Scientist	• Developed AI-driven image processing models for OLED TVs using **TensorFlow** and **PyTorch**, achieving a 25% improvement in image clarity (PSNR metric) and enabling real-time 4K upscaling for the Z95A series launch in September 2024. • Designed predictive maintenance models with **XGBoost** and Apache Spark for manufacturing equipment, reducing downtime by 20% (MTTR metric) and supporting sustainability goals showcased at CES 2025. • Built customer segmentation models using BERT and Snowflake for market re-entry, increasing targeted campaign ROI by 15% through enhanced audience precision and data-driven insights. • Implemented **MLOps** pipelines with **AWS SageMaker** and **MLflow**, automating model deployment for the "Panasonic Go" AI initiative, achieving 30% faster deployment cycles and 99.9% model uptime. • Created interactive Tableau dashboards for sustainability metrics, integrating Google BigQuery data, enabling 10% faster decision-making on carbon emission reductions, aligning with Panasonic's net-zero 2030 target.
Panel C: Natural Language Processing		
1	Data Engineer	Developed a predictive model using CatBoost to forecast credit risk, achieving a 15% improvement in prediction accuracy (Mean Absolute Error) over prior models. This refinement has significantly strengthened our risk management capabilities. Spearheaded research using Bertopic and LDAvis for **NLP topic modeling**, effectively pinpointing key themes in customer feedback across various channels. This initiative has directly influenced product improvements and customer service strategies. Created and fine-tuned an ARIMAX model to accurately predict monthly loan demand, directly contributing to $590,000 in annual savings through optimized resource allocation and enhanced financial planning.
2	Business Data Analyst	Built and automated 35+ dashboards using Tableau, Qlik, and Streamlit, reducing manual reporting time by 60% across departments. Analyzed structured & semi-structured data with Python and SQL, boosting campaign ROI and cutting churn by 25%. Optimized ETL pipelines with SSIS and Qlik Data Integration, improving refresh speed by 70% and reducing latency. Designed scalable data models and star schemas across Snowflake, MySQL, and Qlik, enabling real-time analytics for 5+ teams. Cleansed and transformed datasets with 10M+ rows, ensuring accurate reporting for finance, sales, and operations. Analyzed 100K+ customer reviews, increasing **sentiment classification** precision by 18% and informing product strategy. Created Excel models with pivot tables, macros, and lookups, cutting spreadsheet errors by 35% and enabling faster financial scenarios. Developed interactive Qlik Sense dashboards with drill-downs, improving decision-making speed by 40%. Automated data ingestion from 10+ sources using Qlik pipelines, reducing manual prep by 50%. Enhanced dashboard performance with advanced Qlik scripting (SET analysis, variables), cutting load times by 30% on 20M+ row datasets
3	Software Engineer	• Developed an AI Assist Intelligent Bot using Oracle cloud, Python, and **NLP**, and trained models for context-sensitive user assistance to enable real-time information retrieval and interactive support, significantly enhancing user engagement and satisfaction. • Developed an interactive dashboard for business users to provide better insights into metrics, KPIs, and trends on a SaaS application. Utilized Python, HTML, CSS, and JavaScript, following an API-first approach microservice-based architecture, and Integrated it with the AI Help Tool. • Served as Scrum Master, effectively managing sprint planning and execution using Jira. Created and prioritized user stories, logged work, and facilitated agile ceremonies to improve team productivity and Project outcomes. • Enhanced my knowledge in Oracle SaaS and AI, earning key certifications and applying these skills to drive innovative solutions and business success.
4	**NLP** Algorithm Engineer	Worked on text processing projects, focusing on **Named Entity Recognition** (NER) related to e-commerce, leveraging **Natural Language Processing** (**NLP**) techniques to enhance the recommendation system's accuracy and relevance. Collaborated closely with cross-functional teams to identify and resolve complex challenges, ensuring seamless project execution.

Continued on next page

Table A.13 continued

Example	Job title	Job description
5	Staff Software Engineer (Apple Intelligence)	Scalable Knowledge Graphs for On-Device Intelligence • Designed and implemented scalable knowledge graph architectures optimized for inference, enabling real-time contextual intelligence in resource-constrained environments. • Developed graph ingestion pipelines to process large-scale heterogeneous data sources, ensuring efficient entity resolution, **relationship extraction**, and semantic enrichment. • Integrated lightweight graph query engines and custom embeddings to support intelligent decision-making on edge devices. • Focused on optimization for latency, memory footprint, and incremental updates to enable robust offline performance and personalization at scale. • On-device embeddings computation and store it in Vector DB
6	Data Scientist	Developed and trained a **natural language processing** (**NLP**) model to automate invoice processing, boosting the rate of fully automated invoices from 13% to 60% and significantly reducing manual effort and processing costs.
7	Senior Consultant	Specialized in forensic analytics with hands-on experience supporting Anti-Money Laundering (AML) and Bank Secrecy Act (BSA) compliance investigations. Designed and implemented ETL pipelines, developed automation solutions, and built predictive models leveraging advanced **natural language processing** (**NLP**) techniques.
8	AI Software Developer	**NLP**-driven land price forecasting.
9	Data Scientist	• Developed and deployed predictive models, enhancing customer engagement and informing data-driven strategies. • Engineered ETL pipelines and clustering solutions for effective customer segmentation, optimizing marketing efforts. • Created executive Tableau dashboards and utilized **NLP** for **sentiment analysis**, providing actionable insights for leadership.
10	Data Analyst	Developed scalable data workflows using Python and SQL to streamline data collection, wrangling, and preprocessing tasks, enabling timely delivery of analytical insights to business stakeholders. Applied statistical modeling and predictive analytics techniques to support 20+ business initiatives, improving operational forecasting accuracy by 30% and enabling proactive resource planning. Designed and launched 15+ interactive dashboards using Tableau and Excel, increasing executive reporting engagement by 40% through effective data storytelling and visualization strategies. Developed and maintained custom analytics models in Python (Pandas, Scikit-learn), identifying key behavioral patterns that led to a 25% boost in customer retention and a 20% increase in sales conversion. Built and automated cloud-based data pipelines across AWS and GCP environments, reducing manual intervention by 60% and achieving 99.9% uptime for real-time analytics processing. Constructed and validated supervised and unsupervised ML models, including classification, regression, clustering, and **NLP** models, leading to a 40% increase in predictive accuracy for business-critical use cases. Partnered with business and technical teams to document data requirements, deploy version-controlled code via GitHub, and ensure collaboration throughout the project lifecycle.
Panel D: Computer Vision		
1	Data Scientist	Developed and deployed classical **computer vision** algorithms (e.g., image filtering, morphological operations, feature extraction, **template matching**, object segmentation) to automate manufacturing processes in aluminum recycling, increasing system efficiency and reducing manual intervention Engineered performance metrics and designed key performance indicators (KPIs) to assess and enhance the accuracy of feature detection, resulting in a 10% improvement in detection precision Implemented a multi-threaded image processing pipeline, optimizing real-time feature detection performance, reducing processing latency by 60% Calibrated machine vision cameras to ensure precise pixel-to-world distance measurements and consistent object localization Collaborated with the **computer vision** team to integrate new algorithms into the existing pipeline, utilizing continuous integration workflows via GitHub Projects to streamline development and ensure seamless version control and testing
2	AI Research Scientist	Development of a Generalized Pipeline for Carotid Vessel Segmentation Across Diverse MRI Sequences: Advanced Research in Data Processing, Normalization, and **Computer Vision** Models Optimization Leveraging Supercomputing.
3	Senior Data Analyst	• Analyze large volumes of supply chain data to identify bottlenecks, track KPIs, and support strategic decision-making in reverse logistics. • Collaborate with cross-functional teams to integrate AI tools into operational systems, achieving a 20% improvement in processing efficiency. • Led AI-driven process improvements in Verizon's reverse logistics operations, using **computer vision** (YOLO, CNN) to optimize grading machine performance and increase accuracy of returned product assessments. • Apply machine learning models to automate and enhance refund approval workflows, reducing processing time and human error.
4	Senior **Computer Vision** Scientist	Working on the Research and Development of framework for AI assisted eye surgery
5	Software Engineer, Fundamental AI Research (FAIR)	Core contributor to the Segment Anything 2 (SAM 2) research demo, which served the SAM 2 model and all its capabilities to x00,000 global users. Building new research playgrounds, human annotation tasks, and data visualization tools for **computer vision** research (object segmentation in video, semantic grounding in images and videos) FAIR (Fundamental AI Research) - Perception

Continued on next page

Table A.13 continued

Example	Job title	Job description
6	Senior Software Engineer	The company developed concentrated solar power (CSP) technology including novel heliostats and multithreaded C# control software • Created the heliostat battery management system's central controller, which used MQTT protocol; to prevent total battery discharge, functions included a low power mode, a nightly power off, and a failsafe heartbeat; enabled a 21% increase in test site uptime • Improved a **computer vision** gradient anomaly mask for cloud detection and created a test visualization tool to enable future work • Performed a root cause analysis to successfully identify temporary server overloads as the cause of communication losses • Collaborated on robustness improvements to the 4th generation heliostats' TCP/IP-to-serial network bridge and protocol handler
7	Data Scientist	**Computer Vision** (AI driven content for the website and identifying products in images)
8	Principal Solutions Architect	As Principal Solutions Architect at Berkshire Grey, I lead the end-to-end design and deployment of AI-powered robotic automation systems that optimize large-scale supply chain operations. My role blends deep technical architecture with strategic leadership, guiding cross-functional teams to deliver complex solutions that integrate robotics, AI/ML, **computer vision**, and orchestration software. I work directly with enterprise clients to translate operational challenges into scalable automation strategies, while also influencing product direction through field-driven insights. Focused on execution and impact, I've delivered high-throughput systems that drive measurable efficiency gains across retail, eCommerce, and logistics environments.
9	Director of Engineering	I lead and manage multiple **computer vision**/AI/machine learning engineering teams, developing organizational-level technology strategies and roadmaps to ensure engineering deliverables align with Intel's overall business and product strategy.
10	AP Engineer	. Develop and optimize AI models for autonomous driving, robotics, and energy systems, focusing on scalability and real-world application.. Design deep learning algorithms for **computer vision**, sensor fusion, and predictive analytics to enhance Tesla's Full Self-Driving (FSD) capabilities. • Process and analyze large-scale datasets from Tesla's fleet to identify and address edge cases.. Collaborate with cross-functional teams to integrate AI solutions into production and proprietary hardware.. Optimize machine learning models for deployment on Tesla's FSD hardware, ensuring minimal latency and high performance.. Conduct research on emerging AI technologies to drive innovation aligned with Tesla's mission of accelerating sustainable energy adoption.
Panel E: Generative AI		
1	Data Science / AI Engineer	• Implementing best-practice **LLM** prompting (prompt chaining and modularization) for concurrent taxonomy generation and content tagging to improve B.O. prediction • Utilizing agent-based prompting (task agent and verification agent) to generate in-depth feature explanations for multidimensional film similarity searches, ensuring **LLM** answers are grounded in real tabular data. • Proposing and testing various clustering, representation similarity, and A/B testing strategies to validate attributes generated by **LLMs**
2	Applied Scientist	**Large language models** for Amazon shopping and customer understanding. Pre-training (long context and hybrid attention architectures). Pre-train evaluations.
3	Full-stack Developer	Developed a **GenAI**-powered customer insights engine integrating GPT APIs & LangChain with Spring Boot microservices. Built RAG pipelines (**Pinecone**, FAISS) enabling semantic search across financial documents & compliance records. Designed AI-powered chatbots & knowledge assistants (React.js + Spring Boot), improving support accuracy by 40%. Integrated Kafka event streaming + **LLM** APIs for real-time fraud detection & anomaly insights. Deployed **LLMs** on AWS (Bedrock, SageMaker) for scalable, hybrid rule-based + AI decision systems. Implemented secure **prompt engineering**, API throttling, and logging for AI observability (Splunk, Grafana). Enhanced backend for transaction monitoring & fraud detection while embedding **GenAI** features into workflows.
4	**GenAI** Engineer	I design and deploy **LLM**-powered **Generative AI** applications using LangChain, Hugging Face, and OpenAI APIs, integrating RAG pipelines for enterprise-scale knowledge retrieval. I have fine-tuned multilingual models like T5 and GPT-Neo with DeepSpeed and Accelerate, boosting domain-specific accuracy by up to 20%. My work includes building document Q&A systems, AI chatbots, and multi-agent frameworks with FAISS and **Pinecone**, and implementing scalable inference via Triton. I also design benchmarking dashboards to track latency, hallucination, and token usage, and lead A/B testing across **GPT-4**, Claude, and Cohere, optimizing performance and reliability.
5	**Generative AI** Engineer	Designing and deploying enterprise-grade **Generative AI** solutions with a focus on **LLMs**, RAG pipelines, and agentic AI systems. Built autonomous AI workflows (AutoGen, CrewAI) to automate reasoning and multi-step tasks, integrated with Azure + **vector databases** (**Pinecone**). Deployed scalable models via Docker, Kubernetes, REST APIs, and implemented monitoring with Grafana & Prometheus. Applied evaluation frameworks (RAGAS, G-Eval) to ensure accuracy, safety, and business relevance.

Continued on next page

Table A.13 continued

Example	Job title	Job description
6	Senior Dot Net Full Stack Developer	Architecting operational data mesh on ADLS Gen2, Delta Lake, Unity Catalog Implementing PySpark ETL frameworks with schema governance, incremental upserts Integrating **Azure OpenAI** embeddings & RAG pipelines into operational datasets Designing CDC ingestion using Event Hubs + Delta checkpoints Building FastAPI microservices exposing reliability signals Creating ADF pipelines with lineage and dependency graphs Developing reliability semantic layers in Looker for SLO/SLA reporting Instrumenting ETL & **LLM** workflows using Azure Monitor, App Insights Enforcing cost governance & secure operational boundaries Automating CI/CD using GitHub Actions for notebooks and Delta tables
7	Senior Software Engineer	Designed and deployed end-to-end Voice AI Agents using LiveKit, Retell AI, Vapi, and Twilio, integrating real-time STT, LLM_x0002_based dialogue management, and natural TTS responses to handle dynamic phone conversations, automate customer service workflows, and connect seamlessly with CRM and backend APIs through n8n and webhooks for data synchronization and task automation.
8	AI Enablement & Community Adoption	At AT&T, I've had the privilege of leading and contributing to several transformative AI initiatives-most notably the design, development, and deployment of advanced AI assistants powered by **Retrieval-Augmented Generation** (RAG) and fine-tuned **large language models** (**LLMs**). One of the most exciting projects I spearheaded was the creation of a multi-purpose RAG-based AI assistant that supports customer service, sales, and technical troubleshooting teams. This assistant draws from AT&T's vast ecosystem of structured and unstructured data-ranging from product catalogs and billing documentation to network troubleshooting manuals and customer interaction histories-enabling it to deliver highly accurate, context-aware, and grounded responses in real-time. To build this, I integrated **LLMs** like OpenAI's GPT, Anthropic's Claude, Cohere, and Mistral with LangChain for orchestration and **Pinecone**/FAISS for semantic retrieval. By combining the power of foundation models with precise information retrieval systems, we created a tool that drastically reduced time-to-resolution for customers and improved operational efficiency for internal teams. But generic **LLMs** don't understand domain-specific language or nuances-so I fine-tuned models like Mistral and LLaMA on AT&T-specific datasets including customer service transcripts, FAQs, and internal technical documentation. This improved accuracy, increased trust in AI responses, and allowed us to maintain relevance even in highly specialized telecom conversations. I also developed optimized **GenAI** pipelines involving embedding refinement, dynamic chunking strategies, and multi-step prompt chaining. These pipelines ensured our assistants could not only retrieve the right knowledge but synthesize it effectively-even when customers asked layered, ambiguous, or multi-intent queries. The result was a robust conversational system capable of delivering accurate, contextual, and natural-sounding answers that mirrored human-level support.
9	Associate Engineer, AI	- AI & Advanced Analytics- Internal **Generative AI** use-case/experimentation enablement
10	Analyst	Product Lead on SnackGPT, a **GenAI** platform that analyzed complex sets of business data using Gemini and YAML query engineering. Responsibilities involved leading daily stand-up meetings with business and tech employees, testing the app, and resolving front- and back-end bugs when they arose. Front-end developer: Designed, built, tested and deployed several websites and dashboards that were used in day-to-day business activities. Products included the NA Digital Portal, a unified landing page built using JavaScript integrated with Metric Insights API, and several Power BI dashboards. Business products involved querying and transforming large sets of data from sources like Google Cloud Platform and designing a convenient UX, as well as leading training sessions with 200+ business users to introduce the product.
Panel F: Agentic AI		
1	Software Engineer	Implemented a RAG chatbot using Hugging Face, **LangChain**, and Oracle Autonomous Database to enable semantic search across internal docs, adopted by internal engineering teams.
2	Head of Product, Agentforce Guardrails	Building AgentForce - the world's first and most innovative platform for Trusted **AI Agents** for enterprise. I lead the product vertical for trusted, safe and secure **AI Agents**. Learn more:
3	Sr Manager, Product Management - **Agentic AI**	• Defined and led WEX's enterprise **Agentic AI** strategy ($2B+ portfolio) in partnership with CDO, CTO, and business leaders across Mobility, Corporate Payments, and Benefits. • Built consensus among 15+ stakeholders across Technology, Operations, and Customer Experience to define enterprise Virtual Assistant strategy, architecting unified multi-agent ecosystem (Benefits Assist, Claims Processing AI, Agent Assist) with human-centered design principles and API-first microservices architecture deployed across voice, web, messaging, and mobile channels. • Scaled Virtual Assistant platform to deliver measurable business impact powering 5M+ annual interactions, improving self-service containment, reducing average handle time, and generating $2M+ in operational savings while maintaining 4.4/5 customer satisfaction score. • Spearheaded Finn, an enterprise agentic marketplace unlocking $2M(25K hours) in productivity savings. • Managed PMs (ICs and Managers) and a team of 80+ engineers, shipping embedded GenAI products (AutoIVR, Claims Automation, Fleet Intelligence) generating $10M+ in new revenue. • Built AI partner ecosystem through strategic partnerships with UiPath (RPA), ThoughtSpot, Moveworks, Decagon, AWS, GCP and OpenAI, accelerating intelligent automation, and enabling scalable GenAI enterprise adoption. • Oversaw 40+ AI experiments across diverse customer segments to validate pain points and de-risk product ideas.

Continued on next page

Table A.13 continued

Example	Job title	Job description
4	Gen AI Developer	Engineered the AI Lending Insight Platform leveraging LLaMA-based generative AI for contextual credit and payment behavior analysis, generating narrative explanations for loan approval, BNPL eligibility, and transaction patterns, reducing manual review effort by ~12%. • Integrated generative AI into Python-based backend microservices (FastAPI/Flask), embedding LLaMA inference to generate real-time consumer credit summaries and payment recommendations, accelerating decision-making turnaround by nearly 15% for lending and risk teams. • Orchestrated **AI agents** to correlate transactional behavior, repayment history, and alternative financial signals while designing and optimizing PostgreSQL and Aurora schemas, applying indexing, partitioning, and consistency strategies to support high-volume consumer lending workloads. • Architected a BNPL risk and payment decision microservices using Python, Kafka, and gRPC, using OAuth authentication, AWS Cognito, end-to-end encryption, secrets management, JWT validation, audit logging, and RBAC, ensuring platform security and regulatory compliance. • Developed real-time merchant and consumer dashboards using React, TypeScript, Redux Toolkit, WebSockets, and D3.js to visualize payment trends, credit utilization, and risk indicators. Applied unit and integration testing with Pytest, Selenium, and automated CI/CD pipelines, improving release reliability by 13%. • Collaborated with DevOps teams to deploy containerized Python services on AWS EKS using Docker and Kubernetes, integrating OpenTelemetry, Prometheus, Grafana, and ELK for enhanced observability, autoscaling, and production stability, improving system uptime and monitoring efficiency.
5	Software Engineer II	Developing tools to objectively measure and refine the performance of **AI agents**. Ensuring every AI-driven assistant stays aligned with your goals and keeps improving through guided feedback.
6	Backend & Gen AI Engineer	• Designed, built, and scaled a production-grade **Agentic AI** platform delivering conversational assistants, task-based copilots, and autonomous **AI agents** for enterprise analytics and decision support. • Evaluated and integrated LLMs and multimodal foundation models from multiple providers (OpenAI, Google, Anthropic, open-source via Hugging Face/Ollama), selecting models dynamically based on latency, cost, and task complexity. • Architected agentic workflows using **LangChain** and OpenAI Agents SDK, enabling tool-calling, multi-step reasoning, memory, and multi-agent orchestration to autonomously plan and execute complex tasks. • Built advanced RAG pipelines with: • Vector embeddings and databases (Pinecone / PGVector / Vertex AI Search) • Multi-source retrieval across documents, tickets, logs, and structured data • Reranking strategies and semantic caching to improve relevance and response quality • Implemented prompt engineering best practices, including system instructions, few-shot prompting, JSON/JSON Schema structured outputs, and function/tool calling to ensure deterministic and auditable AI behavior. • Integrated safety and guardrail layers (content moderation APIs, custom policies, schema validation, hallucination detection) to minimize unsafe outputs and keep conversations focused and compliant. • Developed full-stack AI experiences, integrating **AI agents** into React-based front-end applications with real-time interactions via WebSockets and secure backend orchestration. • Built scalable backend services using Python, exposing AI orchestration via REST, gRPC, and WebSockets, and deployed them in containerized, cloud-native environments using Docker and Kubernetes. • Designed and maintained LLM evaluation and monitoring pipelines, tracking offline and online metrics such as accuracy, relevance, latency, and cost; used LangSmith and OpenTelemetry for tracing and observability.
7	Senior Engineering Manager - Virtual Expert Platform	• Innovate in the assisted services domain using **agentic ai** to build done for you experiences. • Lead engineers on a few core engineering teams to enhance the Virtual Expert Platform, driving customer-expert engagements. • Implement AI-driven workflows that streamline interactions across Intuit products, creating seamless user experiences. • Foster collaboration among diverse personas to develop innovative agentic solutions, significantly improving customer satisfaction.
8	Python Developer	Designed and implemented FastAPI-based microservices for real-time trade risk analytics, reducing API latency by 42%. • Developed automated ETL pipelines using Python, Pandas, and SQLAlchemy to sync trade data between internal systems and third party platforms (e.g., Bloomberg). • Integrated JWT/OAuth2-secured APIs for account and portfolio services, ensuring robust authorization and compliance with FINRA and SOX guidelines. • Containerized applications using Docker and orchestrated deployments via Kubernetes (EKS) for high availability and scalability. • Optimized PostgreSQL queries and database indexing strategies, resulting in a 35% performance gain for analytical dashboards. • Created internal Python libraries for financial data transformation, utilized across 5+ teams for standardization. • Integrated LLM-based summarization models (using OpenAI APIs and **LangChain**) for generating investment insights and market summaries. • Collaborated in Agile teams with product managers, DevOps, and data engineers using Jira, Confluence, and GitHub Actions

Continued on next page

Table A.13 continued

Example	Job title	Job description
9	Program Manager	• Spearheaded GenAI business review cadences (5 VP-level MBRs), aligning Directors across 25+ initiatives and directly reprioritizing $15M+ investments in **Agentic AI** and Foundational Tech programs • Orchestrated 7 cross-functional (Product, Engineering and Legal) partnerships to streamline GenAI experimentation pipelines, accelerating time-to-feedback by 45% and enabling 3 high-impact pilots to scale across customer touchpoints • Drove 50% increase in GenAI program adoption by launching an impact-tracking system that influenced $10M+ in funding decisions and helped sunset 3 low-ROI initiatives unlocking $5M+ in variable cost efficiency • Orchestrated cross-portfolio revenue forecasting models through BIE partnerships spanning 25+ initiatives, enabling VP-level strategic pivots that redirected $15M in GenAI investments toward highest-ROI opportunities • Deployed scalable automation dashboards covering 120+ tech teams, reducing manual reporting overhead by 70% and informing 3 engineering practice pivots based on surfaced telemetry patterns • Directed automated pipeline attribution system development across 120+ product teams, surfacing performance correlation patterns that informed 3 strategic operational pivots and improved delivery timeline accuracy by 35%
10	Lead Generative AI Engineer	Enterprise AI-Powered Clinical & Claims Automation Platform • Architected enterprise-grade LLM-powered Clinical Decision Support System using GPT-4, **LangChain**, and RAG 2.0 • Designed scalable **Agentic AI** frameworks with multi-agent orchestration and tool calling • Built AI microservices using Python (FastAPI) and Java (Spring Boot) • Implemented HIPAA-compliant AWS deployments (EKS, Lambda, S3, IAM) • Integrated AI modules with legacy J2EE healthcare systems and HL7/FHIR APIs • Developed RAG pipelines using LlamaIndex, Pinecone, FAISS • Implemented Kafka-driven real-time claims processing • Led LLM evaluation frameworks with A/B testing and human-in-the-loop validation • Ensured enterprise security (OAuth2, JWT, RBAC, AES-256 encryption)

This table reports 10 examples of AI-related U.S. jobs by AI subcategory for our matched U.S. public-firm sample from 2010 to 2024. Jobs are assigned to subcategories based on keyword matches in either the job title or the job description. The job descriptions are shown to illustrate the type of text captured by the keyword-based classification, and bold text highlights AI-related terms or phrases under each subcategory.

Table A.14. Top 30 Job Titles with the Most AI Jobs in 2024

Rank	Job-title	AI positions	Rank	Job-title	AI positions
1	Software Engineer	19,515	16	Director of Product Management	725
2	Machine Learning Engineer	9,792	17	Machine Learning Scientist / Researcher	665
3	Data Scientist	8,170	18	Product Designer	637
4	Product Manager	4,695	19	Account Executive	630
5	Data Engineer	4,523	20	Business Analyst	548
6	Data Analyst	4,120	21	Vice President	413
7	Engineering Manager	3,974	22	Manager	376
8	AI Engineer	3,778	23	AI Software Engineer	312
9	Program / Project Manager	2,555	24	Data Science Manager	307
10	Research Scientist	1,900	25	Solutions Architect	269
11	Software Development Engineer	1,411	26	Technical Recruiter	242
12	Software Developer	1,406	27	Computer Vision Engineer	203
13	Director	1,373	28	Generative AI Associate	183
14	Applied Scientist	1,127	29	Artificial Intelligence Partner	179
15	Technical Lead / Principal Engineer	1,106	30	Product Marketing Manager	168

This table reports the 30 job titles with the largest number of AI-related U.S. jobs in 2024 for our matched U.S. public firm sample. Job titles are ranked by the total number of AI-related jobs. Titles are based on raw LinkedIn job-title strings observed in Revelio. To generate this table, we aggregate closely related title variants ourselves: seniority and level modifiers, such as "Senior," "Staff," "Principal," "Lead," "II," and "III," are combined when the underlying role is the same. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.15. Top 10 Job Titles with the Most AI Jobs by AI Subcategory in 2024

Rank	AI	ML	NLP	CV	Gen AI	Agentic AI
1	Software Engineer	Machine Learning Engineer	Data Scientist	Software Engineer	Software Engineer	Software Engineer
2	Product Manager	Software Engineer	Software Engineer	Data Scientist	Data Scientist	Data Scientist
3	Data Scientist	Data Scientist	Machine Learning Engineer	Computer Vision Engineer	Machine Learning Engineer	AI/ML Engineer
4	AI Engineer	Data Engineer	AI/ML Engineer	Machine Learning Engineer	Product Manager	Machine Learning Engineer
5	Machine Learning Engineer	Data Analyst	Data Engineer	AI Engineer	Applied Scientist	Product Manager
6	AI/ML Engineer	AI/ML Engineer	Data Analyst	AI/ML Engineer	Data Engineer	AI Engineer
7	Data Engineer	Product Manager	AI Engineer	Artificial Intelligence Engineer	Software Development Engineer	Data Engineer
8	Technical Program Manager	Machine Learning Software Engineer	Applied Scientist	Applied Scientist	AI/ML Engineer	Generative AI Engineer
9	Engineering Manager	Python Developer	Product Manager	Research Scientist	Research Scientist	Software Development Engineer
10	Program Manager	Engineering Manager	Artificial Intelligence Engineer	Autonomous Vehicles Software Engineering Manager	AI Engineer	Gen AI Engineer

This table reports the 10 job titles with the largest number of AI-related U.S. jobs in 2024 within each AI subcategory. Job titles are ranked by the number of AI-related jobs within each subcategory. Titles are based on raw LinkedIn job-title strings observed in Revelio. To generate this table, we aggregate closely related title variants ourselves: seniority and level modifiers, such as "Senior," "Staff," "Principal," "Lead," "II," and "III," are combined when the underlying role is the same. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.16. Top 30 Occupations with the Highest AI Shares in 2024

Rank	O*NET code	Occupation	AI share	AI jobs	Total jobs
1	15-1221.00	Computer and Information Research Scientists	0.66	23,985	36,568
2	15-2051.00	Data Scientists	0.43	50,403	117,157
3	15-2099.00	Mathematical Science Occupations, All Other	0.39	335	867
4	17-2199.08	Robotics Engineers	0.24	1,515	6,222
5	25-1022.00	Mathematical Science Teachers, Postsecondary	0.22	68	305
6	19-2099.01	Remote Sensing Scientists and Technologists	0.17	46	265
7	19-1029.01	Bioinformatics Scientists	0.17	785	4,611
8	15-1255.01	Video Game Designers	0.14	532	3,846
9	19-3022.00	Survey Researchers	0.14	999	7,363
10	15-1299.07	Blockchain Engineers	0.13	1,605	12,313
11	15-1243.00	Database Architects	0.11	12,253	107,464
12	29-2036.00	Medical Dosimetrists	0.10	69	676
13	15-1255.00	Web and Digital Interface Designers	0.10	6,110	61,608
14	27-3091.00	Interpreters and Translators	0.10	227	2,353
15	25-1021.00	Computer Science Teachers, Postsecondary	0.09	180	1,971
16	17-2031.00	Bioengineers and Biomedical Engineers	0.08	1,364	16,404
17	13-2099.01	Financial Quantitative Analysts	0.08	1,103	13,713
18	15-2021.00	Mathematicians	0.08	42	526
19	19-3041.00	Sociologists	0.08	25	314
20	13-1161.01	Search Marketing Strategists	0.06	2,026	32,151
21	19-2012.00	Physicists	0.06	36	581
22	17-2072.01	Radio Frequency Identification Device Specialists	0.06	1,959	31,756
23	25-1042.00	Biological Science Teachers, Postsecondary	0.06	26	432
24	17-2199.06	Microsystems Engineers	0.06	8,247	139,429
25	19-1099.00	Life Scientists, All Other	0.06	425	7,369
26	27-1029.00	Designers, All Other	0.06	2,365	42,656
27	15-1252.00	Software Developers	0.06	40,418	729,588
28	19-3094.00	Political Scientists	0.06	62	1,125
29	19-3011.00	Economists	0.05	118	2,347
30	11-2021.00	Marketing Managers	0.05	17,849	358,050

This table reports the 30 occupations with the highest AI-related job shares among U.S. jobs in 2024 for our matched U.S. public-firm sample. Occupations are based on O*NET occupation codes. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs in the same occupation and is reported as a fraction. AI-related jobs are identified using keyword matches in job titles or job descriptions. The sample is restricted to firm-years with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to occupations with at least 100 total U.S. jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.17. Top 30 Occupations with the Most AI Jobs in 2024

Rank	O*NET code	Occupation	AI jobs	AI share	Total jobs
1	15-2051.00	Data Scientists	50,403	0.4302	117,157
2	15-1252.00	Software Developers	40,418	0.0554	729,588
3	15-1221.00	Computer and Information Research Scientists	23,985	0.6559	36,568
4	11-2021.00	Marketing Managers	17,849	0.0499	358,050
5	15-1299.09	Information Technology Project Managers	16,234	0.0365	444,979
6	15-2051.01	Business Intelligence Analysts	14,719	0.0456	323,089
7	15-1299.08	Computer Systems Engineers/Architects	12,667	0.0403	314,006
8	15-1243.00	Database Architects	12,253	0.114	107,464
9	17-2199.06	Microsystems Engineers	8,247	0.0591	139,429
10	15-1255.00	Web and Digital Interface Designers	6,110	0.0992	61,608
11	41-9031.00	Sales Engineers	5,126	0.0432	118,686
12	13-1071.00	Human Resources Specialists	4,196	0.027	155,460
13	41-3091.00	Sales Representatives of Services, Except Advertising, Insurance, Financial Services, and Travel	3,413	0.0152	224,262
14	11-1011.00	Chief Executives	3,395	0.0187	181,871
15	17-2199.00	Engineers, All Other	3,092	0.0171	180,902
16	15-1299.05	Information Security Engineers	2,713	0.0371	73,218
17	11-2022.00	Sales Managers	2,618	0.0081	324,345
18	19-1029.00	Biological Scientists, All Other	2,586	0.0294	87,953
19	11-3021.00	Computer and Information Systems Managers	2,398	0.0315	76,132
20	27-1029.00	Designers, All Other	2,365	0.0554	42,656
21	23-1011.00	Lawyers	2,336	0.0326	71,718
22	41-4012.00	Sales Representatives, Wholesale and Manufacturing, Except Technical and Scientific Products	2,214	0.006	367,741
23	15-1254.00	Web Developers	2,213	0.0285	77,710
24	13-1161.00	Market Research Analysts and Marketing Specialists	2,171	0.0493	44,005
25	11-3031.01	Treasurers and Controllers	2,070	0.0074	279,856
26	13-1199.00	Business Operations Specialists, All Other	2,039	0.0168	121,426
27	13-1161.01	Search Marketing Strategists	2,026	0.063	32,151
28	17-2072.01	Radio Frequency Identification Device Specialists	1,959	0.0617	31,756
29	15-1253.00	Software Quality Assurance Analysts and Testers	1,949	0.0175	111,525
30	15-1243.01	Data Warehousing Specialists	1,912	0.0319	59,908

This table reports the 30 occupations with the largest number of AI-related U.S. jobs in 2024 for our matched U.S. public-firm sample. Occupations are based on O*NET occupation codes. AI-related jobs are identified using keyword matches in job titles or job descriptions. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs in the same occupation and is reported as a fraction. The sample is restricted to firm-years with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to occupations with at least 100 total U.S. jobs. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.18. Top 3 Firms with Highest AI Shares by Industry in 2024

Industry	Rank	Firm	AI share	AI jobs	Total jobs
Accommodation & Food Svcs	1	SONDER HOLDINGS INC	0.0118	12	1,015
	2	MGM RESORTS INTERNATIONAL	0.0073	72	9,809
	3	BOYD GAMING CORP	0.0070	15	2,145
Admin Support	1	STEM INC	0.1308	62	474
	2	KNIGHTSCOPE INC	0.1061	14	132
	3	RESEARCH SOLUTIONS INC	0.0800	6	75
Arts/Entertainment	1	DRAFTKINGS INC	0.0177	75	4,243
	2	FLUTTER ENTERTAINMENT PLC	0.0172	1	58
	3	ENDEAVOR GROUP HOLDINGS	0.0091	22	2,429
Construction	1	ORION GROUP HOLDINGS INC	0.0122	3	246
	2	BEAZER HOMES USA INC	0.0077	12	1,556
	3	FERROVIAL SE	0.0068	5	731
Education Svcs	1	NERDY INC	0.1835	29	158
	2	TAL EDUCATION GROUP	0.1194	8	67
	3	DUOLINGO INC	0.0749	100	1,335
Finance/Insurance	1	UPSTART HOLDINGS INC	0.1025	199	1,941
	2	DAVE INC	0.0687	29	422
	3	VISA INC	0.0677	1,316	19,450
Health Care	1	TEMPUS AI INC	0.1134	403	3,555
	2	DARIOHEALTH CORP	0.0793	18	227
	3	ONTRAK INC	0.0633	15	237
Information	1	C3.AI INC	0.5156	960	1,862
	2	COVEO SOLUTIONS INC	0.3553	27	76
	3	BAIDU INC	0.3434	114	332
Manufacturing	1	XPENG INC	0.3645	109	299
	2	NVIDIA CORP	0.3007	8,965	29,818
	3	ABSCI CORP	0.2301	55	239
Mining	1	ENCORE ENERGY CORP	0.0192	1	52
	2	SLB LTD	0.0177	302	17,080
	3	VITAL ENGY INC	0.0176	9	511
Other Svcs	1	FRONTDOOR INC	0.0109	16	1,472
	2	WW INTERNATIONAL INC	0.0109	39	3,590
	3	CHARGEPOINT HOLDINGS INC	0.0094	11	1,166
Prof & Business Svcs	1	INNODATA INC	0.7169	595	830
	2	STAGWELL INC	0.1193	109	914
	3	FLUENT INC	0.1066	37	347
Real Estate	1	CEVA INC	0.1613	10	62
	2	INTERDIGITAL INC	0.0613	20	326
	3	DOLBY LABORATORIES INC	0.0466	77	1,651
Retail Trade	1	COUPANG INC	0.2098	214	1,020
	2	ALIBABA GROUP HLDG	0.1197	116	969
	3	XOMETRY INC	0.0702	61	869
Transportation/Warehousing	1	DIDI GLOBAL INC	0.1079	26	241
	2	SERVE ROBOTICS INC	0.1061	21	198
	3	GRAB HOLDINGS LIMITED	0.0394	5	127
Utilities	1	NEXTDECADE CORP	0.0211	6	284
	2	NEXTERA ENERGY INC	0.0195	191	9,801
	3	AES CORP (THE)	0.0185	70	3,791
Wholesale	1	INSIGHT ENTERPRISES INC	0.0361	318	8,812
	2	EPLUS INC	0.0285	58	2,037
	3	MCKESSON CORP	0.0266	574	21,541

This table reports the three firms with the highest AI-related job shares within each two-digit NAICS industry in 2024 for our matched U.S. public-firm sample. Rank is assigned separately within each industry. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. AI-related jobs are identified using keyword matches in job titles or job descriptions. The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.19. Top 3 Firms with Most AI Jobs by Industry in 2024

Industry	Rank	Firm	AI jobs	AI share	Total jobs
Accommodation & Food Svcs	1	STARBUCKS CORP	253	0.0021	119,293
	2	MARRIOTT INTL INC	159	0.0036	44,168
	3	MCDONALD'S CORP	146	0.0023	62,615
Admin Support	1	EXPEDIA GROUP INC	801	0.0670	11,960
	2	S&P GLOBAL INC	457	0.0312	14,637
	3	EQUIFAX INC	381	0.0401	9,504
Arts/Entertainment	1	DRAFTKINGS INC	75	0.0177	4,243
	2	LIVE NATION ENTERTAINMENT	73	0.0047	15,632
	3	BRIGHTSTAR LOTTERY PLC	31	0.0060	5,166
Construction	1	LENNAR CORP	55	0.0058	9,483
	2	FLUOR CORP	19	0.0018	10,346
	3	BEAZER HOMES USA INC	12	0.0077	1,556
Education Svcs	1	PEARSON PLC	312	0.0220	14,212
	2	DUOLINGO INC	100	0.0749	1,335
	3	SKILLSOFT CORP	96	0.0620	1,548
Finance/Insurance	1	JPMORGAN CHASE & CO	4,527	0.0236	191,930
	2	CAPITAL ONE FINANCIAL CORP	3,284	0.0458	71,757
	3	WELLS FARGO & CO	2,699	0.0130	206,854
Health Care	1	CVS HEALTH CORP	2,173	0.0165	131,826
	2	HUMANA INC	910	0.0228	39,905
	3	HCA HEALTHCARE INC	417	0.0089	46,674
Information	1	MICROSOFT CORP	26,169	0.1355	193,200
	2	ALPHABET INC	21,385	0.1194	179,131
	3	META PLATFORMS INC	15,075	0.1781	84,644
Manufacturing	1	NVIDIA CORP	8,965	0.3007	29,818
	2	APPLE INC	7,900	0.0539	146,636
	3	INTEL CORP	5,430	0.0791	68,679
Mining	1	SLB LTD	302	0.0177	17,080
	2	BAKER HUGHES CO	133	0.0094	14,119
	3	CONOCOPHILLIPS	117	0.0099	11,861
Other Svcs	1	WW INTERNATIONAL INC	39	0.0109	3,590
	2	CINTAS CORP	23	0.0010	23,462
	3	FRONTDOOR INC	16	0.0109	1,472
Prof & Business Svcs	1	ACCENTURE PLC	4,571	0.0586	77,983
	2	BOOZ ALLEN HAMILTON HLDG CP	2,034	0.0481	42,244
	3	COGNIZANT TECH SOLUTIONS	1,732	0.0443	39,070
Real Estate	1	CBRE GROUP INC	229	0.0060	37,878
	2	EQUINIX INC	229	0.0412	5,565
	3	JONES LANG LASALLE INC	159	0.0046	34,665
Retail Trade	1	AMAZON.COM INC	14,668	0.0410	358,106
	2	WALMART INC	1,876	0.0075	250,219
	3	EBAY INC	818	0.0505	16,211
Transportation/Warehousing	1	UNITED PARCEL SERVICE INC	284	0.0032	89,810
	2	DELTA AIR LINES INC	280	0.0058	48,629
	3	UNITED AIRLINES HOLDINGS INC	269	0.0064	41,854
Utilities	1	NEXTERA ENERGY INC	191	0.0195	9,801
	2	NRG ENERGY INC	160	0.0125	12,814
	3	DUKE ENERGY CORP	132	0.0070	18,778
Wholesale	1	MCKESSON CORP	574	0.0266	21,541
	2	CARDINAL HEALTH INC	467	0.0214	21,859
	3	INSIGHT ENTERPRISES INC	318	0.0361	8,812

This table reports the three firms with the largest number of AI-related U.S. jobs within each two-digit NAICS industry in 2024 for our matched U.S. public-firm sample. Rank is assigned separately within each industry. AI-related jobs are identified using keyword matches in job titles or job descriptions. The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.20. Top 30 Tech-Sector Firms by AI Job Share in 2024

Rank	Firm	AI share	Industry	AI jobs	Total jobs
1	INNODATA INC	0.7169	Prof & Business Svcs	595	830
2	C3.AI INC	0.5156	Information	960	1,862
3	COVEO SOLUTIONS INC	0.3553	Information	27	76
4	BAIDU INC	0.3434	Information	114	332
5	CERENCE INC	0.3067	Information	69	225
6	VERITONE INC	0.3037	Information	147	484
7	SOUNDHOUND AI INC	0.2965	Information	110	371
8	LIVEPERSON INC	0.2137	Information	178	833
9	REKOR SYSTEMS INC	0.2044	Information	37	181
10	UIPATH INC	0.1869	Information	386	2,065
11	META PLATFORMS INC	0.1781	Information	15,075	84,644
12	GENIUS SPORTS LTD	0.1745	Information	26	149
13	STRATEGY INC	0.1711	Information	253	1,479
14	MOBILEYE GLOBAL INC	0.1695	Information	20	118
15	EGAIN CORP	0.1548	Information	61	394
16	PONY AI INC	0.1542	Information	31	201
17	CYNGN INC	0.1429	Information	16	112
18	APPLOVIN CORP	0.1360	Information	144	1,059
19	MICROSOFT CORP	0.1355	Information	26,169	193,200
20	PINTEREST INC	0.1341	Information	810	6,042
21	APPLIED DIGITAL CORP	0.1329	Information	23	173
22	INUVO INC	0.1319	Information	19	144
23	SNOWFLAKE INC	0.1257	Information	954	7,591
24	FISCALNOTE HOLDINGS INC	0.1194	Information	64	536
25	ALPHABET INC	0.1194	Information	21,385	179,131
26	STAGWELL INC	0.1193	Prof & Business Svcs	109	914
27	CS DISCO INC	0.1178	Information	76	645
28	SNAP INC	0.1160	Information	746	6,431
29	COREWEAVE INC	0.1142	Information	87	762
30	BIGBEAR.AI HOLDINGS INC	0.1140	Information	75	658

This table reports the 30 tech-sector firms with the highest AI-related job shares in 2024 for our matched U.S. public-firm sample. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. AI-related jobs are identified using keyword matches in job titles or job descriptions. The tech-sector sample includes firms in information (NAICS 51) and professional and business services (NAICS 54). The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.21. Top 30 Tech-Sector Firms by Number of AI Jobs in 2024

Rank	Firm	AI jobs	AI share	Industry	Total jobs
1	MICROSOFT CORP	26,169	0.1355	Information	193,200
2	ALPHABET INC	21,385	0.1194	Information	179,131
3	META PLATFORMS INC	15,075	0.1781	Information	84,644
4	INTL BUSINESS MACHINES CORP	9,440	0.0900	Information	104,923
5	ACCENTURE PLC	4,571	0.0586	Prof & Business Svcs	77,983
6	SALESFORCE INC	4,433	0.0802	Information	55,245
7	ORACLE CORP	4,037	0.0453	Information	89,111
8	ADOBE INC	2,775	0.0996	Information	27,873
9	VERIZON COMMUNICATIONS INC	2,465	0.0253	Information	97,536
10	AT&T INC	2,259	0.0158	Information	142,681
11	BOOZ ALLEN HAMILTON HLDG CP	2,034	0.0481	Prof & Business Svcs	42,244
12	SERVICENOW INC	1,904	0.0999	Information	19,067
13	INTUIT INC	1,902	0.0789	Information	24,109
14	COGNIZANT TECH SOLUTIONS	1,732	0.0443	Prof & Business Svcs	39,070
15	COMCAST CORP	1,457	0.0179	Information	81,392
16	SAP SE	1,300	0.0543	Information	23,939
17	DEUTSCHE TELEKOM	1,243	0.0167	Information	74,584
18	WORKDAY INC	980	0.0577	Information	16,982
19	C3.AI INC	960	0.5156	Information	1,862
20	SNOWFLAKE INC	954	0.1257	Information	7,591
21	NETFLIX INC	931	0.0520	Information	17,897
22	PALO ALTO NETWORKS INC	875	0.0705	Information	12,416
23	INFOSYS LTD	856	0.0300	Prof & Business Svcs	28,504
24	PINTEREST INC	810	0.1341	Information	6,042
25	SNAP INC	746	0.1160	Information	6,431
26	THOMSON-REUTERS CORP	658	0.0478	Information	13,767
27	AUTODESK INC	643	0.0706	Information	9,111
28	ATLASSIAN CORP	613	0.0864	Information	7,091
29	GARTNER INC	612	0.0371	Prof & Business Svcs	16,507
30	INNODATA INC	595	0.7169	Prof & Business Svcs	830

This table reports the 30 tech-sector firms with the largest number of AI-related U.S. jobs in 2024 for our matched U.S. public-firm sample. AI-related jobs are U.S. jobs identified using keyword matches in job titles or job descriptions. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. The tech-sector sample includes firms in information (NAICS 51) and professional and business services (NAICS 54). The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.22. Top 30 Non-Tech Firms by AI Job Share in 2024

Rank	Firm	AI share	Industry	AI jobs	Total jobs
1	XPENG INC	0.3645	Manufacturing	109	299
2	NVIDIA CORP	0.3007	Manufacturing	8,965	29,818
3	ABSCI CORP	0.2301	Manufacturing	55	239
4	COUPANG INC	0.2098	Retail Trade	214	1,020
5	NERDY INC	0.1835	Education Svcs	29	158
6	AMBARELLA INC	0.1833	Manufacturing	68	371
7	WERIDE INC	0.1765	Manufacturing	33	187
8	CEVA INC	0.1613	Real Estate	10	62
9	ZEPP HEALTH CORP	0.1594	Manufacturing	11	69
10	MERCEDES BENZ GROUP AG	0.1424	Manufacturing	198	1,390
11	PENGUIN SOLUTIONS INC	0.1362	Manufacturing	88	646
12	SES AI CORP	0.1321	Manufacturing	21	159
13	STEM INC	0.1308	Admin Support	62	474
14	ALIBABA GROUP HLDG	0.1197	Retail Trade	116	969
15	TAL EDUCATION GROUP	0.1194	Education Svcs	8	67
16	ONE STOP SYSTEMS INC	0.1149	Manufacturing	10	87
17	SUPER MICRO COMPUTER INC	0.1138	Manufacturing	279	2,452
18	NIO INC	0.1134	Manufacturing	39	344
19	TEMPUS AI INC	0.1134	Health Care	403	3,555
20	RECURSION PHARMA INC	0.1101	Manufacturing	97	881
21	ASTERA LABS INC	0.1082	Manufacturing	42	388
22	ADVANCED MICRO DEVICES	0.1081	Manufacturing	1,670	15,444
23	DIDI GLOBAL INC	0.1079	Transportation/Warehousing	26	241
24	ALLURION TECHNOLOGIES INC	0.1064	Manufacturing	15	141
25	SERVE ROBOTICS INC	0.1061	Transportation/Warehousing	21	198
26	KNIGHTSCOPE INC	0.1061	Admin Support	14	132
27	UPSTART HOLDINGS INC	0.1025	Finance/Insurance	199	1,941
28	CERIBELL INC	0.1000	Manufacturing	42	420
29	ALTO NEUROSCIENCE INC	0.0877	Manufacturing	10	114
30	BUTTERFLY NETWORK INC	0.0816	Manufacturing	32	392

This table reports the 30 non-tech firms with the highest AI-related job shares in 2024 for our matched U.S. public-firm sample. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. AI-related jobs are U.S. jobs identified using keyword matches in job titles or job descriptions. The non-tech sample excludes firms in information (NAICS 51) and professional and business services (NAICS 54). The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.23. Top 30 Non-Tech Firms by Number of AI Jobs in 2024

Rank	Firm	AI jobs	AI share	Industry	Total jobs
1	AMAZON.COM INC	14,668	0.0410	Retail Trade	358,106
2	NVIDIA CORP	8,965	0.3007	Manufacturing	29,818
3	APPLE INC	7,900	0.0539	Manufacturing	146,636
4	INTEL CORP	5,430	0.0791	Manufacturing	68,679
5	JPMORGAN CHASE & CO	4,527	0.0236	Finance/Insurance	191,930
6	CISCO SYSTEMS INC	3,972	0.0582	Manufacturing	68,223
7	CAPITAL ONE FINANCIAL CORP	3,284	0.0458	Finance/Insurance	71,757
8	DELL TECHNOLOGIES INC	3,163	0.0521	Manufacturing	60,732
9	WELLS FARGO & CO	2,699	0.0130	Finance/Insurance	206,854
10	TESLA INC	2,545	0.0424	Manufacturing	60,006
11	CVS HEALTH CORP	2,173	0.0165	Health Care	131,826
12	CITIGROUP INC	1,877	0.0207	Finance/Insurance	90,850
13	WALMART INC	1,876	0.0075	Retail Trade	250,219
14	BANK OF AMERICA CORP	1,790	0.0096	Finance/Insurance	185,569
15	ADVANCED MICRO DEVICES	1,670	0.1081	Manufacturing	15,444
16	QUALCOMM INC	1,546	0.0777	Manufacturing	19,907
17	UNITEDHEALTH GROUP INC	1,392	0.0189	Finance/Insurance	73,651
18	HP INC	1,333	0.0652	Manufacturing	20,436
19	LOCKHEED MARTIN CORP	1,323	0.0133	Manufacturing	99,685
20	VISA INC	1,316	0.0677	Finance/Insurance	19,450
21	GENERAL MOTORS CO	1,269	0.0178	Manufacturing	71,209
22	JOHNSON & JOHNSON	1,251	0.0268	Manufacturing	46,766
23	ELEVANCE HEALTH INC	1,177	0.0223	Finance/Insurance	52,772
24	AMERICAN EXPRESS CO	1,138	0.0297	Finance/Insurance	38,326
25	FORD MOTOR CO	1,082	0.0217	Manufacturing	49,831
26	CIGNA GROUP (THE)	1,065	0.0228	Finance/Insurance	46,702
27	MORGAN STANLEY	1,053	0.0151	Finance/Insurance	69,844
28	PFIZER INC	1,043	0.0214	Manufacturing	48,678
29	PAYPAL HOLDINGS INC	1,018	0.0647	Finance/Insurance	15,746
30	BROADCOM INC	992	0.0322	Manufacturing	30,813

This table reports the 30 non-tech firms with the largest number of AI-related U.S. jobs in 2024 for our matched U.S. public-firm sample. AI-related jobs are U.S. jobs identified using keyword matches in job titles or job descriptions. AI share is defined as AI-related U.S. jobs divided by total U.S. jobs at the firm. The non-tech sample excludes firms in information (NAICS 51) and professional and business services (NAICS 54). The 2024 firm-year sample is restricted to firms with at least 50 U.S. workers observed in Revelio, at least 50 employees in Compustat, and to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.24. Correlations between Cognism and Revelio AI Measures

Year	Number of Firms	Corr. AI share	Corr. AI job count	Corr. total job count
2010	668	0.450	0.771	0.873
2011	710	0.517	0.837	0.876
2012	753	0.357	0.811	0.886
2013	832	0.351	0.839	0.888
2014	832	0.427	0.871	0.904
2015	856	0.511	0.873	0.910
2016	877	0.629	0.854	0.912
2017	929	0.515	0.918	0.920
2018	940	0.717	0.929	0.926
2019	922	0.477	0.936	0.927
2020	848	0.784	0.931	0.932
2021	598	0.723	0.929	0.937

The table reports year-by-year correlations between two resume data sources, Revelio and Cognism, across three different variables, covering 2010 to 2021. Each row corresponds to a calendar year. The "Number of Firms" column gives the number of underlying firms on which each year's correlations are computed. The remaining three columns report Pearson correlation coefficients between the Revelio and Cognism measures of, respectively: the AI share (the share of jobs classified as AI-related), the AI count (the absolute number of AI-related jobs), and the total job count (the overall number of jobs).

Table A.25. AI Job Functional Taxonomy

Category	Definition
AI Research & Model Development	Builds or directly advances AI models or core algorithms, including machine learning, deep learning, NLP, computer vision, reinforcement learning, and applied AI research.
AI Data, Annotation & Evaluation Support	Creates, curates, labels, validates, or analyzes data and evaluation assets primarily to support AI model training, testing, or improvement.
AI Platform, Infrastructure & MLOps	Builds or operates the technical platforms, pipelines, infrastructure, and deployment systems used to train, serve, monitor, or scale AI workloads.
AI Application & Systems Engineering	Integrates AI models or methods into software, devices, products, or operational systems where the core job is engineering AI-enabled functionality rather than creating the underlying model.
AI Product, Program & Delivery Management	Leads product strategy, program execution, or cross-functional delivery for AI products, AI features, or AI implementation without being the primary hands-on model builder.
AI Governance, Risk & Responsible AI	Sets, enforces, audits, or advises on AI-related governance, compliance, ethics, privacy, security, model risk, or responsible AI practices.
AI Business Leadership & Functional Oversight	Provides executive, business-unit, or functional leadership over AI-related teams, portfolios, or investments where the person oversees AI direction but is not primarily doing hands-on AI work.
Domain Analytics & Predictive Modeling	Uses statistical analysis, forecasting, optimization, or machine learning mainly to generate business, scientific, clinical, financial, or operational insights rather than to build AI products or platforms.
AI Go-to-Market, Solutions & Advisory	Sells, solutions, markets, evangelizes, or advises on AI-related products or transformations for customers or internal stakeholders without being the primary builder.
General Digital/Technical Role with AI Tooling or Claims	Works in a broader software, operations, consulting, or business role where AI is presented as one tool, initiative, or supporting claim but not the central function of the job.
Incidental or Unsubstantiated AI Mention	Mentions AI-related terms but provides insufficient evidence that AI work, AI deployment, or meaningful AI use was a substantive part of the role.

This table defines the functional categories used to describe what employees do within the AI-related job sample. The categories group sampled AI-related jobs by the worker's primary role in the AI ecosystem. The taxonomy is constructed from a random sample of 4,500 AI-related jobs, with 300 jobs per year from 2010 through 2024. Each category is mutually exclusive and is based on the worker's job title and self-reported job description rather than industry, employer, or seniority. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio jobs divided by Compustat employment for each firm-year.

Table A.26. Distribution of AI Jobs Across Functional Categories

Category	N	% jobs
AI Research & Model Development	650	14.4
AI Data, Annotation & Evaluation Support	51	1.1
AI Platform, Infrastructure & MLOps	273	6.1
AI Application & Systems Engineering	469	10.4
AI Product, Program & Delivery Management	529	11.8
AI Governance, Risk & Responsible AI	61	1.4
AI Business Leadership & Functional Oversight	491	10.9
Domain Analytics & Predictive Modeling	531	11.8
AI Go-to-Market, Solutions & Advisory	334	7.4
General Digital/Technical Role with AI Tooling or Claims	775	17.2
Incidental or Unsubstantiated AI Mention	336	7.5

This table reports the distribution of the 4,500 sampled AI-related jobs across the taxonomy of functional AI job categories. Each job is assigned to exactly one category using the worker's job title and job description. "N" reports the number of jobs assigned to each category, and "% jobs" reports the corresponding share of the classified sample. Categories are mutually exclusive. Table A.25 provides the definitions for functional categories. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio jobs divided by Compustat employment for each firm-year.

Table A.27. Share of AI Jobs by Aggregated AI Functional Category

	All	Tech	Non-tech
Core AI development	45.2%	42.9%	47.9%
Management	22.7%	23.1%	22.2%
Support for AI products/services	7.4%	11.6%	2.5%
Other	24.7%	22.4%	27.4%
Jobs	4,500	2,420	2,080

This table reports the distribution of 4,500 sampled AI-related jobs across aggregated taxonomy categories. Each job is assigned to exactly one category using the worker's job title and self-reported job description, which we then aggregate into four groups. "Core AI development" combines "AI Research & Model Development", "AI Data, Annotation & Evaluation Support", "AI Platform, Infrastructure & MLOps", "AI Application & Systems Engineering", "AI Governance, Risk & Responsible AI", and "Domain Analytics & Predictive Modeling". "Management" combines "AI Product, Program & Delivery Management" and "AI Business Leadership & Functional Oversight". "Support for AI products/services" consists of "AI Go-to-Market, Solutions & Advisory". "Other" combines "General Digital/Technical Role with AI Tooling or Claims" and "Incidental or Unsubstantiated AI Mention". Table A.25 provides the definitions for functional categories. Column 1 includes all sampled jobs, and columns 2 and 3 split firms into the tech sector—NAICS 51, 54, and semiconductor manufacturing (NAICS 334413)—and all remaining (non-tech) sectors. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio jobs divided by Compustat employment for each firm-year.

Table A.28. Representativeness of the non-AI sample

Panel A: Distribution across industries (% of non-AI job-years)

Industry (2-digit NAICS)	Population %	Sample %
Manufacturing	27.02	27.83
Finance/Insurance	17.67	16.14
Retail Trade	14.64	14.60
Transportation/Warehousing	4.02	3.19
Accommodation & Food Svcs	3.67	3.87
Health Care	2.68	2.49
Wholesale	2.03	1.78
Admin Support	1.65	1.86
Utilities	1.71	1.39
Real Estate	1.42	1.11
Mining	1.25	1.19
Education Svcs	0.46	0.47
Arts/Entertainment	0.44	0.60
Construction	0.37	0.28
Other Svcs	0.26	0.21
Agriculture	0.07	0.07

Panel B: Firm size (US jobs per firm-year)

U.S. jobs per firm-year	Population	Sample
10th percentile	2,923	2,959
25th percentile	8,044	8,291
Median	23,457	24,207
75th percentile	70,311	73,298
90th percentile	153,629	154,255
Mean	52,280	53,728

This table compares the randomly drawn non-AI job sample with the population of non-AI U.S. jobs used for sampling. The population includes all non-AI U.S. jobs in firm-years with at least one AI-related U.S. job. The sample consists of 15,000 randomly drawn non-AI U.S. jobs from this population each year from 2010 through 2024. Panel A reports the industry distribution for the two-digit NAICS industries, excluding IT sectors. Panel B reports the distribution of firm size, measured as the total number of U.S. jobs at the firm-year level. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.29. Organization-capital AI versus non-AI jobs

	AI jobs	Non-AI jobs
N	593,332	15,000
Mean organization-capital score	0.601	0.424
Share org-related (score ≥ 0.6)	57.5%	31.0%
Distribution across score bands		
Unrelated (0.00–0.15)	2.6%	17.8%
Incidental (0.16–0.40)	21.3%	35.3%
Mixed (0.41–0.60)	18.6%	15.9%
Substantial (0.61–0.85)	48.7%	28.7%
Core (0.86–1.00)	8.7%	2.3%

This table compares organization capital scores for AI-related U.S. jobs and non-AI U.S. jobs. AI-related jobs are identified using keyword matches in job titles or job descriptions. The non-AI sample consists of 15,000 randomly drawn non-AI U.S. jobs each year from 2010 through 2024. Each job is scored once. "Org-related" denotes a score of at least 0.6. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Table A.30. Definition of Organization Capital Reason Category

Panel A: Builds organization capital

Reason category	Definition
Reusable data platforms & pipelines	They build firm-specific organization capital by codifying the firm's data definitions, integration logic, governance rules, and reporting needs into reusable data platforms, pipelines, warehouses, and shared datasets that preserve operational learning and make future analytics and coordination more efficient.
Reusable AI/ML platforms & MLOps	They build firm-specific organization capital by converting the firm's accumulated model-development and deployment experience into reusable AI/ML platforms, MLOps tooling, evaluation standards, and monitoring routines that let many teams build, launch, and improve models more efficiently over time.
Firm-specific predictive & decision models	They build firm-specific organization capital by embedding firm-specific patterns, decision rules, and labeled experience into reusable predictive, ranking, recommendation, search, pricing, forecasting, and optimization models that improve future decisions as the firm learns from its own data.
Proprietary-knowledge GenAI/NLP systems	They build firm-specific organization capital by codifying proprietary documents, workflow knowledge, and judgment rules into reusable generative-AI, conversational-AI, retrieval, and NLP systems that automate recurring work and make the firm's internal knowledge easier to reuse.
Reusable software platforms & workflows	They build firm-specific organization capital by translating tacit business processes, integrations, and exception-handling know-how into reusable software platforms, APIs, enterprise applications, and automated workflows that standardize execution and reduce reliance on individual employees.
Cross-functional operating routines	They build firm-specific organization capital by institutionalizing cross-functional operating knowledge into reusable governance routines, planning cadences, delivery processes, roadmaps, documentation standards, and coordination mechanisms that make complex execution more repeatable at scale.
Reusable cloud, infrastructure & reliability operations	They build firm-specific organization capital by codifying the firm's infrastructure and reliability know-how into reusable cloud architectures, deployment pipelines, automation, monitoring, network configurations, and operating procedures that make future technology operations more scalable and dependable.
Security, compliance & risk systems	They build firm-specific organization capital by translating accumulated security, privacy, legal, compliance, and risk experience into reusable policies, controls, detection rules, audit systems, governance frameworks, and response routines that let the firm manage risk and regulated activity more efficiently over time.
Marketing & revenue-operations systems	They build firm-specific organization capital by codifying customer, market, channel, and monetization learning into reusable marketing systems, advertising models, experimentation frameworks, revenue operations tools, and commercialization playbooks that improve future targeting, launches, and sales productivity.
Customer-service & support systems	They build firm-specific organization capital by embedding customer-service and support experience into reusable workflows, self-service tools, conversational systems, escalation routines, performance metrics, and service platforms that make future service delivery more consistent and efficient.
Product & design systems	They build firm-specific organization capital by codifying product, user, and market learning into reusable product strategies, roadmaps, design systems, UX standards, lifecycle processes, and cross-functional development routines that improve the firm's future ability to create and scale offerings.
Manufacturing & process systems	They build firm-specific organization capital by converting production, quality, and plant operating experience into reusable manufacturing processes, control systems, test frameworks, automation, and cross-site routines that raise future yield, reliability, and scale efficiency.
Supply-chain & logistics systems	They build firm-specific organization capital by codifying the firm's logistics, procurement, inventory, supplier, and fulfillment learning into reusable forecasting, routing, planning, optimization, and supply-chain management systems that improve future resource allocation and operational performance.
R&D & scientific platforms	They build firm-specific organization capital by embedding scientific, clinical, experimental, and regulatory learning into reusable R&D platforms, bioinformatics pipelines, validated methods, data assets, and cross-functional development routines that make future discovery, trials, and product approvals faster and more effective.
Financial & risk decision systems	They build firm-specific organization capital by codifying financial, underwriting, fraud, portfolio, and planning knowledge into reusable data assets, models, reporting systems, and decision routines that improve future capital allocation, risk decisions, and control processes.
Hiring, training & talent development	They build firm-specific organization capital by institutionalizing the firm's people-management and skill-building know-how into reusable hiring systems, assessment methods, onboarding, training curricula, leadership practices, and talent-development routines that spread firm-specific knowledge across employees.
Alliance & partnership systems	They build firm-specific organization capital by codifying alliance, partner, acquisition, and ecosystem-management experience into reusable integration models, co-sell processes, enablement systems, governance structures, and relationship routines that improve the firm's ability to repeatedly scale through external relationships.
Knowledge organization & ontologies	They build firm-specific organization capital by organizing the firm's specialized knowledge into reusable ontologies, schemas, knowledge graphs, semantic indexes, documentation systems, and content standards that make internal information easier to integrate, search, interpret, and apply in future work.

Panel B: Does not build organization capital

Reason category	Definition
Bespoke client consulting/implementation	This work mainly delivers bespoke consulting, implementation, solution-design, or advisory services to external clients, so learning is transferred to customer-specific engagements rather than accumulating as durable firm-specific routines inside the employer.
Portable analytics & one-off models	This work mainly applies general, portable analytics, statistics, or machine-learning methods to produce project-specific models, forecasts, reports, or recommendations, so value is embodied in one-off analytical outputs and worker human capital rather than in enduring firm-specific operating knowledge.
Broadly applicable research & algorithms	This work mainly creates broadly applicable research, algorithms, prototypes, or scientific and engineering know-how that are codified and portable across employers, so the gains accrue to general technical knowledge rather than to firm-specific organizational routines.
Customer-facing products & IP	This work mainly creates customer-facing products, features, models, designs, or intellectual property whose value is embodied in marketable technology output, not in durable internal routines or process knowledge that make the employer organization itself more efficient.
Standardized/portable software work	This work mainly uses standardized software, data, or engineering tools and frameworks to produce code, systems, or technical deliverables, so the knowledge created remains largely portable across firms instead of becoming employer-specific organization capital.
Transactional sales, marketing & account management	This work mainly generates revenue through sales, marketing, partnerships, account management, or client relationship activities, so value comes from transactions and external commercial outcomes rather than from building durable internal routines or production know-how.
Individual recruiting & hiring transactions	This work mainly applies portable recruiting and labor-market matching skills to fill individual vacancies, so it produces transactional hiring outcomes rather than durable firm-specific systems, routines, or knowledge.
Standardized routine operations, support, maintenance	This work mainly executes standardized operations, support, maintenance, administrative, service, or compliance tasks within existing procedures and systems, so it sustains current output rather than creating new firm-specific knowledge that compounds into organization capital.
Project/program management (episodic)	This work mainly applies portable project, product, or program management and coordination skills to deliver discrete launches or projects, so the output is episodic execution rather than repeatable firm-specific routines that improve future operations.
Training & general education	This work mainly provides classroom-style training, rotational experience, or standalone educational projects that build general human capital, rather than firm-specific knowledge developed through learning-by-doing in the employer's own ongoing operations.
Portable legal & contracting	This work mainly applies portable legal, contracting, negotiation, or dispute-management expertise to case-specific matters, so value rests in one-off professional judgment rather than in durable firm-specific operating routines or systems.

This table reports the definition of the reason categories tabulated in Table 2. Panel A lists the categories of jobs classified as building organization capital (score ≥ 0.6). Panel B lists the categories of the remaining jobs.

Table A.31. Organization Capital and AI Job Categories

Panel A: All firms

	Org-related	Not org-related	Jobs	Share of jobs	Org-related share
Core AI development	856	986	1,842	44.2%	46.5%
Management	741	196	937	22.5%	79.1%
Support for AI products/services	67	229	296	7.1%	22.6%
Other	662	434	1,096	26.3%	60.4%
Total	2,326	1,845	4,171	100.0%	55.8%

Panel B: Non-tech firms

	Org-related	Not org-related	Jobs	Share of jobs	Org-related share
Core AI development	447	452	899	46.7%	49.7%
Management	362	57	419	21.8%	86.4%
Support for AI products/services	12	34	46	2.4%	26.1%
Other	382	179	561	29.1%	68.1%
Total	1,203	722	1,925	100.0%	62.5%

Panel C: Tech firms

	Org-related	Not org-related	Jobs	Share of jobs	Org-related share
Core AI development	409	534	943	42.0%	43.4%
Management	379	139	518	23.1%	73.2%
Support for AI products/services	55	195	250	11.1%	22.0%
Other	280	255	535	23.8%	52.3%
Total	1,123	1,123	2,246	100.0%	50.0%

This table reports the distribution of organization-capital-related and other AI jobs across aggregated AI job categories, for all firms (Panel A), non-tech firms (Panel B), and tech firms (NAICS 51, 54, and semiconductor manufacturing; Panel C). Jobs come from the classified random sample of 4,500 AI jobs, and we exclude the 329 jobs without the job description. The aggregated categories follow Table A.27. "Org-related share" is the share of the category's jobs that are classified as building organization capital.

Table A.32. Examples of AI Jobs by Organization-Capital Classification

Year	Job title	Job description	Reason category
Panel A: Organization-capital-related AI jobs			
2010	Director of Trading Surveillance and Data Strategy	Designed and built a scalable statistical surveillance platform with profiling, Complex Event Processing and real time capabilities to address surveillance challenges that are non-deterministic in nature. Conceptualized and initiated a project to use machine learning and Big Data to connect data points across trading and E-Communications. Appointed as the Compliance Data Officer for Institutional Compliance globally. Responsible for all data strategy including data acquisition and on-going data quality monitoring.	Security, compliance & risk systems
2014	Lead Associate	Led advanced AI initiatives in defense and intelligence, overseeing data scientists across various customers, focused on advanced analytic development in Machine Learning and Generative AI applications. Spearheaded proposal development for new business capture, as well as corporate innovation in several areas, including expanding the firm's reach into new industries, developing new ways to productize the firm's data science services, and internal talent development initiatives. Managed combined team of 20 data scientists developing some early generative AI applications. Prepared and rolled out technology proposals ranging from $3M to $30M, securing over $20M in awards. Captured and led work with pro sports organizations & personally landed Superbowl-winning head coach. Fundraised nearly $1M of internal investment to build a data science as a service marketplace. Created joint Incubator + Co-op for firm's internship program, responsible for fostering problem-solving and entrepreneurship, and awarding the best projects Co-op's.	Proprietary-knowledge GenAI/NLP systems

Year	Job title	Job description	Reason category
2017	Program Director, Honeywell Competency Management	Selected by Honeywell executive management to lead the development and commercialization of Honeywell's Competency Management practice focused on building and maintaining employee skills while transitioning educational methods to meet the needs of the next generation of workers. Leveraged IoT technologies including AR/VR, AI, wearable and other smart technologies to capture and transfer knowledge and skills. Managed the development of advanced analytics designed to assess employee skills as they perform their job on an ongoing basis. Created the contracts, business processes, KPI Scorecards and metrics required to support the delivery of a formal competency program.	Hiring, training & talent development
2018	Computer Vision AI/ML Engineer	As a core member of the global Automatic Visual Inspection (AVI) project, provide image analysis solutions, including deep learning and rule-based algorithms. Lead and standardize the MLOps ecosystem, tools and processes to support the development and deployment of AI/ML algorithms for AVI solutions. Build communication platforms and lead forums to promote knowledge sharing and increase the visibility of the AVI project, while achieving technology promotion and application in more than 80 manufacturing bases around the world.	Firm-specific predictive & decision models
2023	Senior Gen AI Engineer	Led the architecture and production deployment of an enterprise-scale FraudGuard Copilot platform supporting 5M+ monthly transactions across fraud detection and risk analytics systems. Designed real-time streaming pipelines (AWS MSK/Kafka), scalable ML fraud scoring services (Python, XGBoost, EKS), and a policy-aware RAG-based GenAI architecture leveraging OpenSearch Vector Engine, LangChain, and AWS Bedrock.	Financial & risk decision systems

Year	Job title	Job description	Reason category
		Delivered citation-backed AI insights, reduced investigation response times by 30%, improved latency by 35%, and implemented secure MLOps, audit traceability, and governance-aligned AI controls within regulated financial environments.	
2023	UX Director, AI	Lead AI design team + strategy across HubSpot's product suite. Define principles, frameworks, and interaction models for integrating AI into CRM and workflow tools. Grew AI UX discipline 0->20 people in 2 years. Shaped the design of Breeze, an AI-powered product line, including agents and assistants for integrated SMB AI use-cases Developed approach to conversational interfaces, AI collaboration and AI UX patterns Helped evolve HubSpot's design system to support AI-native features Collaborated across product, engineering, and marketing to align AI capabilities with strategic goals Driver of company wide adoption for AI product design practices, principles and tool adoption	Product & design systems
2024	Operation & Technology Data Office - Director of Data Governance & AI Accelerator	Lead enterprise-wide Data Governance and AI transformation initiatives that drive operational excellence and measurable business value. Enterprise Data Governance: Enforce comprehensive data governance frameworks, including risk and control strategies and organizational taxonomies. Drive accountability through consistent policy adherence and sustainable data management practices aligned with business objectives. Data Quality Leadership: Build and manage teams responsible for enterprise data quality standards and controls. Establish risk assessment protocols and implementation processes that reduce data-related risks to acceptable levels, preventing operational disruptions and compliance violations. AI Transformation & Adoption: Serve as AI Accelerator, promoting an AI-first organizational culture.	Reusable data platforms & pipelines

Year	Job title	Job description	Reason category
		Lead deployment of enterprise AI tools and process re-engineering initiatives that deliver measurable gains in operational efficiency and workflow simplicity. AI-Powered Productivity: Leverage advanced prompting techniques to streamline document review and analysis workflows. Utilize AI to summarize complex control documentation, simplify technical content for stakeholder consumption, and accelerate quality review processes-significantly improving personal and team productivity. Strategic Analytics: Direct analytics teams in transforming complex data into actionable executive insights. Oversee the complete analytics lifecycle-from data collection and validation through advanced analysis and visualization-enabling data-driven decision-making across all business units. Executive Reporting: Develop and present program performance reports to C-suite executives, board committees, and governance bodies. Deliver strategic recommendations, risk assessments, and business impact analyses that inform critical leadership decisions.	
2024	Product Management, AV MLOps	I lead the Product Management team within MLOps for Autonomous Vehicles within NVIDIA. We're enabling data flywheels for machine learning engineers across perception, planning/control, GT generation, and more!	Reusable AI/ML platforms & MLOps
2024	Lead Software Engineer - D365 Sales AI	I design and ship agentic AI and the evaluation framework that keeps it trustworthy at scale. Evaluations: Architected a generic, extensible framework (connectors, LLM APIs, evaluators, reporting) that gives teams consistent rubrics, automated scoring, and decision-ready dashboards. It's now used by 30+ teams to gate releases and track improvements.	Proprietary-knowledge GenAI/NLP systems

Year	Job title	Job description	Reason category
		Leading team of 5 engineers to deliver AI quality across CXA Org. Led statistical and competitive benchmarking of multiple LLMs, applying rigorous quantitative analysis to evaluate and compare AI agent performance across models. Agents: Delivered a production competitor-research AI agent for sales workflows; led end-to-end quality, safety checks, and unblock strategies across partners. Enablement: Mentor engineers, host office hours, and publish playbooks so teams can self-serve evaluations. Now : Expanding adoption beyond my org and partnering with our experimentation platform to make evals first-class in the build-measure-learn loop.	
2024	Software Engineer - Speech AI, Evals	Led evaluation and fine-tuning of transformer based in-house multilingual Automatic Speech Recognition (ASR) models, including Whisper and Wav2Vec, enabling them to process over 34 million+ daily requests with high accuracy. Prototyped a large-scale annotation, ranking, and auto-evaluation pipeline using multiple ASR engines + LLM-as-judge, enabling scalable and reliable ground-truth generation. Designed a synthetic query generation system (LLM + RAG) to produce context-aware, policy-violating test cases, stress-testing conversational AI guardrails and validating model robustness. Also introduced a probability drift metric for early detection of ASR performance degradation, reducing customer-impacting issues. Engineered "Sherlock," an in-house model evaluation framework built with Java 17, significantly enhancing ASR model assessments. My contributions also include deploying ASR models in production and QA environments that mirror production, reducing pre-deployment issues by 40%, and building an automated testing pipeline that processes 10,000+ production audio files daily, ensuring system reliability under production-level loads.	Knowledge organization & ontologies

Year	Job title	Job description	Reason category
		By designing ELK and Grafana dashboards, I proactively monitored microservices, reducing downtime by 25%, while collaborating with cross-functional teams to maintain a 99.9% uptime SLA. My work leverages technologies like Python, Java, PySpark, Docker, ELK, Grafana, Spring Boot, Databricks, and AWS to drive innovation and reliability in SpeechAI systems.	
Panel B: Non-organization-capital-related AI jobs			
2010	Managing Consultant	Management Consultant leveraging Big Data and AI to solve business problems at clients like PepsiCo, Amgen, Genentech and SuperValu	Bespoke client consulting/implementation
2010	Equity Research Analyst (buy-side)	Built and maintained models on all portfolio holdings and companies on watchlist (∼200 companies) Invested in Industrial, Energy, and Consumer sectors with a focus on DCF, share buybacks, and operating margin mean reversion Incorporated technical and market sentiment analysis to identify optimal trading entry and exit points Ran a mock L/S equity portfolio to track my stock recommendations which outperformed its index by 34.7% over 5 months	Portable analytics & one-off models
2010	Data Scientist	Build and optimize machine learning models for analyzing massive amounts of data and providing business insights. Perform data processing, feature engineering, and model evaluation using Python and R. Collaborate on the development of end-to-end data science solutions to provide personalized recommendations to clients.	Portable analytics & one-off models
2010	Software Engineer	During my time at Interphase, I worked on a new product (Penveu) based around the use of computer vision and sensor fusion in a handheld form factor. Designed and implemented a method of hiding and retrieving imperceptible patterns in images displayed on computer monitors and projection displays Studied and provided design parameters for component specifications, including human design considerations	Customer-facing products & IP

Year	Job title	Job description	Reason category
2010	Software Engineer	Participated in the development of Oracle cloud computing platform, focusing on the implementation of big data analysis and machine learning functions. Developed distributed data processing systems using Java and Python, optimized customers' business processes and increased data processing speed. Designed and deployed multiple data analysis tools to help corporate customers better understand and analyze their business data. Participated in cross-departmental projects to improve product scalability and support the global deployment of cloud computing platforms.	Standardized/portable software work
2010	Senior Staff Software Engineer	Design, implementation and optimization of hardware VP9 encoders / decoders. Chief in inventing and engineering world class VP9 4K60-capable (single-core, 600MHz) hardware encoder, some of the fancier features being • practically unlimited motion search window over 3 ref frames. • RD-optimal mode selection. • trellis quantization. • temporal denoising.	Customer-facing products & IP
2018	Distinguished Research Scientist	Advanced robotics research in machine learning, motion generation, control, and integrated systems for real-world dextrous manipulation. At NVIDIA's robotics lab in Seattle we're driving core algorithmic technologies toward the democratization of advanced robotics systems to unlock new applications across manufacturing, logistics, healthcare, service, and homes. With some of the world's foremost experts in bleeding edge algorithms for robotics, we're driving toward a unique and deep understanding of robot intelligence and pushing this technology toward unprecedented performance.	Broadly applicable research & algorithms
2021	Research Scientist	Responsible for developing advanced and reliable industrial software for automation applications such as smart manufacturing, industrial robots and cyber physical systems. Technical expertise includes automated reasoning, automatic control and artificial intelligence.	Broadly applicable research & algorithms

Year	Job title	Job description	Reason category
2019	AI Engineer	Responsible for designing and developing highly scalable artificial intelligence and machine learning applications for industrial, energy and financial applications. Utilize technologies such as deep learning, natural language processing and image recognition to solve complex business problems faced by customers.	Portable analytics & one-off models
2023	Account Executive, Small Business	Sold Salesforce's multi cloud and AI products to new and existing customer accounts 262% FY25 ACV attainment ($1.1 M closed ACV), 130% FY25 Multiyear attainment, 92% FY25 New Logo attainment Recognized as South MVP for FY25 for exceptional performance; ranked Top 5 in Small Business (Americas) and earned the South 6 Achievement Award Managed accounts across the U.S. in key sectors including professional services, technology, manufacturing, engineering, and energy	Transactional sales, marketing & account management

This table shows examples of AI jobs classified as organization-capital-related (Panel A) and non-organization-capital-related (Panel B), together with each worker's job title, job description, and reason category. We score each job with *GPT-5.4* that reads each job title and description and returns an organization-capital score between 0 and 1, and classify a job as related to building organization capital if its score is at least 0.6. The reason category captures the mechanism behind the classification. See Appendix Table A.30 for definitions of each category.

Table A.33. Examples of Non-AI Jobs by Organization-Capital Classification

Year	Job title	Job description	Reason category
Panel A: Organization-capital-related non-AI jobs			
2010	Restaurant Guest Service Manager of Disney's Flying Fish Cafe	As RGSM at the Flying Fish, I led and supported a cast of approximately 50 servers, seaters and busers to ensure that Disney standards and basic was followed at all times, I educated cast members on food safety and health procedures and familiarized them with how best to ensure these standards are met in our facility. I dedicated a significant amount of time and energy to being visible on the floor, interacting with guests and demonstrating our dedication to providing them with an unforgettable dining experience. Accomplishments include reorganizing our wine storeroom in compliance with the updated receiving standards. In this position, I worked to further the development of our seasoned cast members through pre-meal coaching and one-to-one meetings.	Hiring, training & talent development
2015	EVP/Head of the Corporate Risk Program Office	Head of the Corporate Risk Program Office responsible for management strategy and execution across Wells Fargo's various functional areas, platforms, and risk disciplines. She works to strengthen the risk management function across the company and ensure that the risk perspective is integrated with major initiatives, including those driving improvements in the customer experience.	Security, compliance & risk systems
2016	Senior Software Engineer	I began working in software configuration management and have moved to performance simulation where I have been working in build management. I am the current build management team lead, leading a team of 6 of us. I am responsible for managing our regression testing and test scripts, planning out the direction of our testing and capabilities for multiple programs, as well as defining the overall process and development life-cycle for the program.	Product & design systems

Year	Job title	Job description	Reason category
2017	Talent Portfolio Delivery Planner	Designing and planning delivery for enterprise talent experiences Create and maintaining global delivery standards Fostering delivery innovation by sensing and incorporating industry best practices. Partnering with regional teams to plan and execute enterprise talent experiences Being the single partner contact for the Talent Management Organization talent segment and design partners.	Hiring, training & talent development
2019	Site Quality Manager	Created a centralized Quality System that allowed continued operation of ISO 17025, NADCAP, ISO 9001, and IATF Certifications under a 75% headcount reduction as the only person with a quality background on site. Developed and sourced material coupon manufacturing and inspection from a 3rd party facility in 5 weeks with no on-site support under the immediate loss of the current in house facility.	Manufacturing & process systems
2019	Engineering Manager, Studio Apps	My senior engineering team builds workflow and productivity applications that power Netflix's content business and legal teams: content and music rights management, contracts, content protection. On the Content Intelligence side, my team's apps support submissions (i.e. pitches) and script management.	Firm-specific predictive & decision models
2020	Operations Manager II	Controls specific day-to-day operations which include training, recognition, coaching, performance evaluations, safety training, and staff productivity Manage a team of 50+ Manage 120 truck and delivery routes Administer and adhere to company safety strategies and programs Develop and support new operational strategies to improve overall success and safety	Hiring, training & talent development
2021	Fulfillment Supervisor	Led outbound sortation and shipping in a 1M+ sq. ft. Amazon FC with 250+ associates and 4 direct shift leads. Improved productivity and reduced overtime by 18% through labor optimization and real- time monitoring. Implemented workflow changes based on A/B testing and Six Sigma methodology. Built high-performing teams by creating mentorship and peer-coaching programs.	Supply-chain & logistics systems

Year	Job title	Job description	Reason category
		Partnered with HR and Loss Prevention to address performance, attendance, and integrity cases. Oversaw inventory reconciliation efforts and resolved discrepancies in coordination with ICQA team.	
2022	Sales Facilitator	New Hire facilitation for both Residential Voice and Residential Chat. Facilitate learning to develop participants' skill sets. This includes classroom instruction, team meetings, and online Learning debriefs. Perform all administrative duties including class and lab preparation, administration, and participant record maintenance. Coach learners on specific skills within a role, process or system. Develop new hire curriculum for effectiveness and appropriate content.	Hiring, training & talent development
2023	Customer Experience Manager	• Collaborate with Cross-functional teams to align customer experience with business objectives. • Ensured comprehensive trainings are delivered to internal teams for newly launching processes. • Collaborate with Care and Sales to streamline processes and optimize workflow within Government Education Programs. • Conduct Steering Meetings to collaborate and align on program goals. • Conducted a thorough analysis of existing processes and workflows to identify obstructions and areas of optimization. • Work as a SME for Government processes to launch new initiatives.	Customer-service & support systems
Panel B: Non-organization-capital-related non-AI jobs			
2010	National Account Executive - Field Sales	Online advertising sales. Prospect for and grow national advertising partners. Assist in driving sales from the largest real estate audience in the US (37 million unique users/mo). My specific areas of focus have been homes builders, real estate brokerages and auction houses.	Transactional sales, marketing & account management

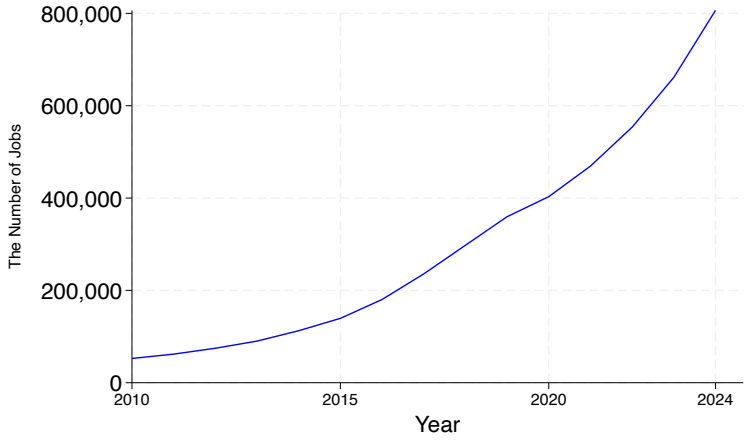
Year	Job title	Job description	Reason category
2012	Copy and Print Associate	Worked with customers to assess the needs of their particular project. Experience operating photocopiers, large format printers, book binders, and cutting machines. Performed quality checks on customer orders and worked with others to complete orders. Used programs such as Adobe Acrobat, Word, Publisher, Excel, Cute PDF, and Photoshop. Maintained a work log that helped to prioritize workflow.	Standardized routine operations, support, maintenance
2013	Mortgage Resolution Analyst	Analyze attorney firms billing invoices for accuracy and confirm invoices are received within the allowable timeframe. Research denied invoices for payment resolution. Review and verify accuracy of invoices and supporting documentation prior to processing. Submit invoices to Account Payables for payment. Ability to process a high volume of invoices per month.	Standardized routine operations, support, maintenance
2014	Licensed Health Insurance Broker and District Manager	Licensed Health Insurance Broker (June 2014-Present) Delivered high impact educational presentations regarding law changes, The Affordable Care Act, product overviews, and how clients can streamline back office administration Performed quote reviews with client to choose appropriate plan options to satisfy and retain employees Responsible for leading and coordinating all Health Insurance enrollments for new clients	Transactional sales, marketing & account management
2016	Trailer Mechanic	Attend to all aspects of trailer maintenance and repair according to company procedure. Restore, lubricate, and adjust hydraulic drum & shoe brake systems. Diagnose and repair electrical lighting issues. Mend damages in trailer body and frame using an angle grinder and MIG welder. Safely use a variety of hand, power, and pneumatic tools while utilizing proper PPE.	Standardized routine operations, support, maintenance
2017	Web Content Writer	I create compelling content and communications strategy for technology, science, retail, health and legal companies ranging from startups to Fortune 500 companies.	Customer-facing products & IP

Year	Job title	Job description	Reason category
		I write and edit blog posts, white papers, web copy, thought leadership pieces, trend articles, press releases, brand copy, social media content, newsletters, internal publications, online questionnaires and video scripts.	
2018	Internet/VoIP Support Technician	Resolved customer connectivity issues related to internet and wireless network configurations, ensuring stable and reliable service. Assisted in porting phone numbers from other carriers, streamlining transitions for VoIP users. Collaborated with internal teams to deliver seamless customer experiences and address complex system issues.	Standardized routine operations, support, maintenance
2019	Customer Service Representative	During my time at Sykes, I worked for the PNC Customer Care Center. I would service an average of 30-40 customers with a range of issues that they had with their bank account. Some of the concerns include: balance discrepancies, online banking troubleshooting, identity theft prevention, dispute management & other responsibilities.	Standardized routine operations, support, maintenance
2020	Cashier	Receiving payments and issuing receipts Keeping track of all cash and credit transactions Communicating effectively with customers Using point-of-sale technologies, such as cash register, computer, barcode scanners, etc.	Standardized routine operations, support, maintenance
2022	Regional Sales Director, Southwest	Capital Rep for Swoop Point of Care MRI system throughout the Southwest. Territory includes Arizona, Southern California, Hawaii & Las Vegas.	Bespoke client consulting/implementation

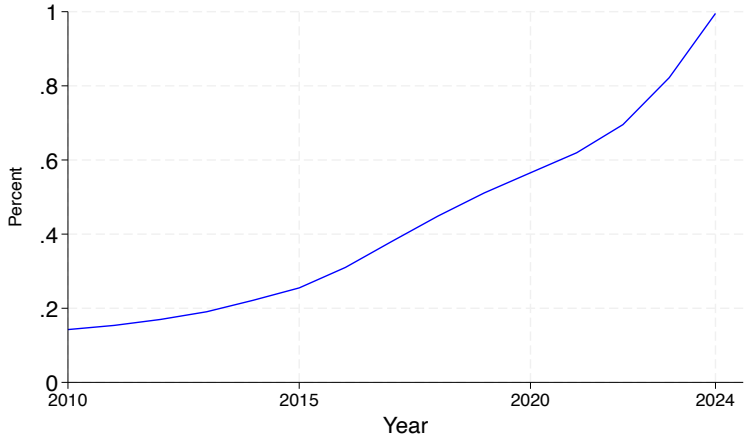
This table shows examples of non-AI jobs classified as organization-capital-related (Panel A) and non-organization-capital-related (Panel B), together with each worker's job title, job description, and reason category. The non-AI sample is a representative random sample of 15,000 non-AI jobs drawn from the same firm-years as the AI jobs. We score each job with *GPT-5.4* that reads each job title and description and returns an organization-capital score between 0 and 1, and classify a job as related to building organization capital if its score is at least 0.6. The reason category captures the mechanism behind the classification. See Table A.30 for definitions of each category.

B. Robustness and Additional Results

Figure B1. Time Series of AI Jobs (Non-Public Firms)



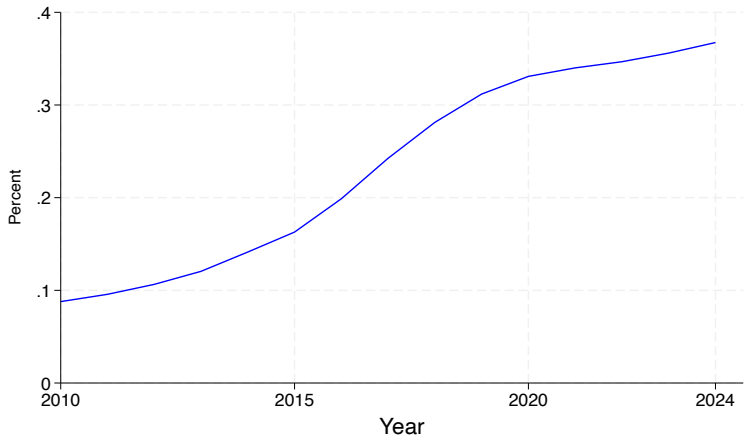
(a) Number of AI Jobs



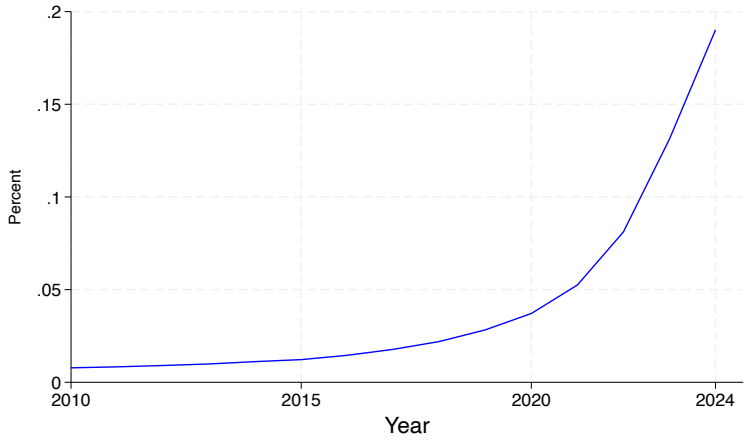
(b) Share of AI Jobs

This figure plots annual AI-related U.S.-based jobs among employers that are not U.S. public firms in Revelio from 2010 to 2024. These employers are mostly privately held firms, government, and non-profit sector entities. Panel (a) reports the total number of AI-related U.S. jobs. Panel (b) reports the aggregate AI job share, defined as AI-related U.S. jobs divided by total U.S. jobs in the same year. AI-related jobs are identified using keyword matches in job titles or job descriptions.

Figure B2. Share of AI Jobs: Wave I vs. Wave II Technologies (Non-Public Firms)



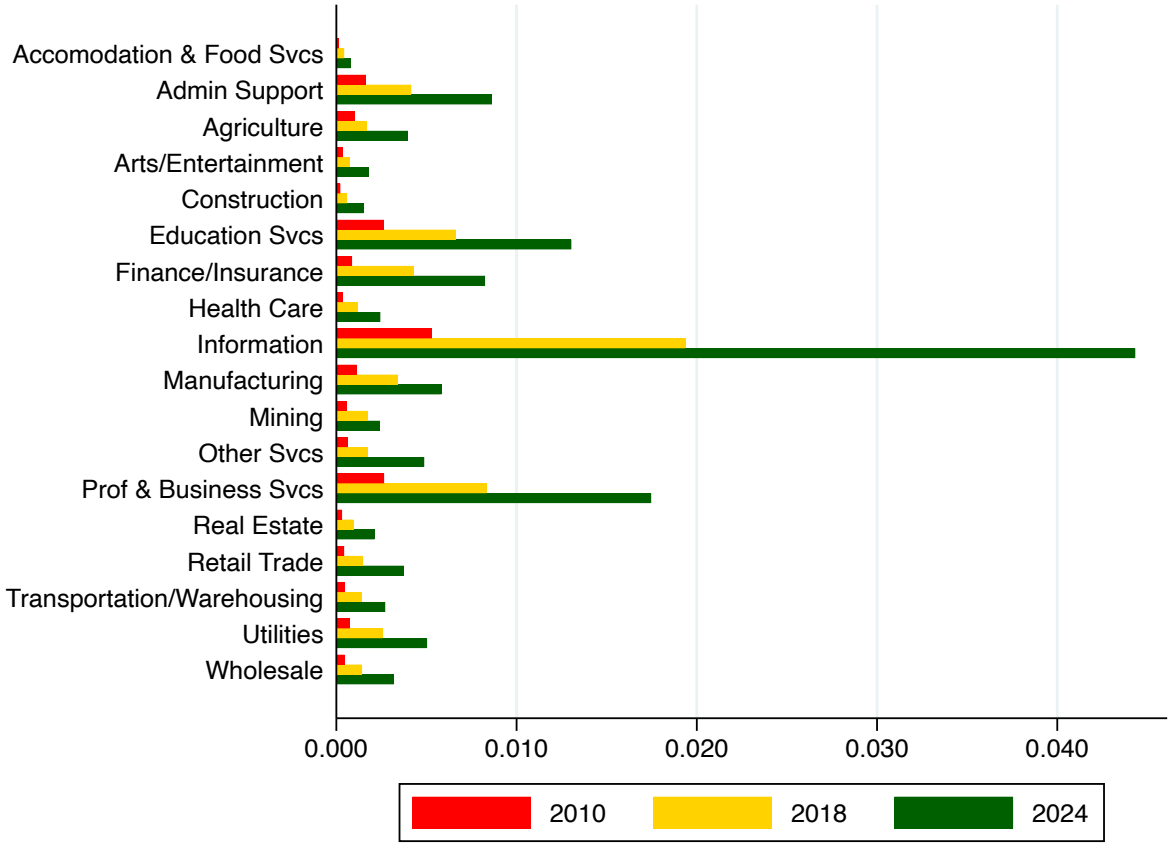
(a) Wave I



(b) Wave II

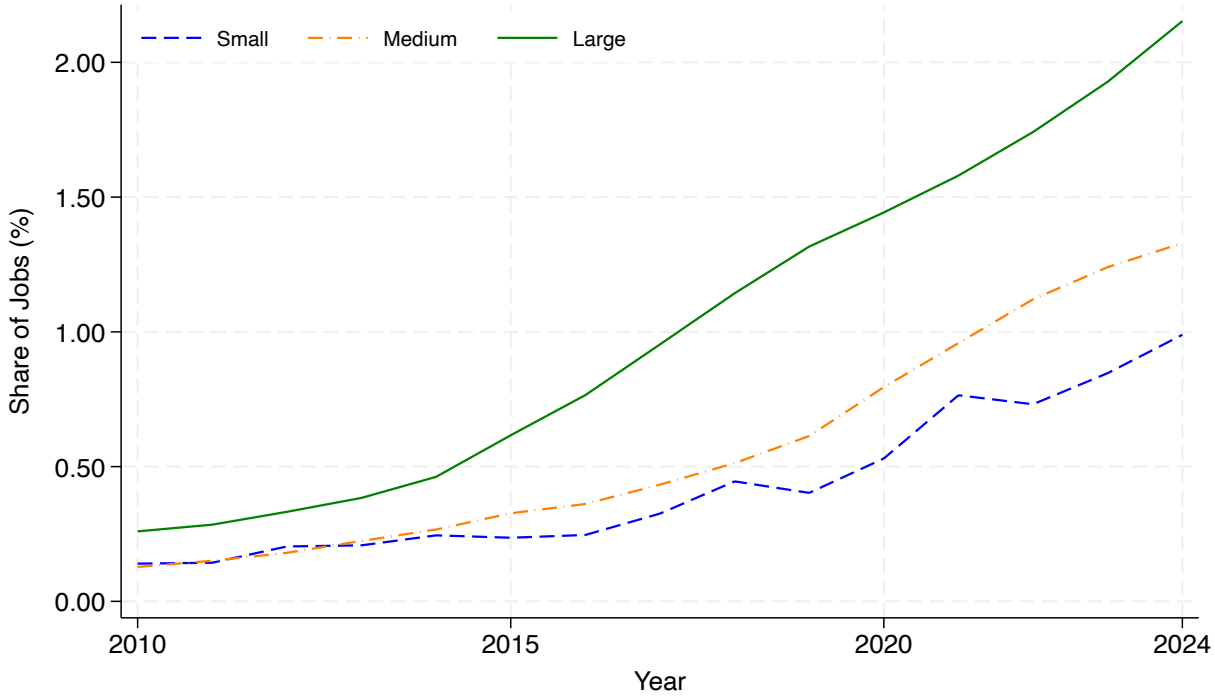
This figure plots annual AI-related U.S.-based jobs among employers that are not U.S. public firms in Revelio from 2010 to 2024, separately for Wave I and Wave II AI technologies. These employers are mostly privately held firms, government and non-profit sector entities. Panel (a) reports the share of "Wave I" AI jobs, defined as jobs containing at least one of the NLP, ML, or CV keywords. Panel (b) reports the share of "Wave II" AI jobs, defined as jobs containing at least one of the generative AI or agentic AI keywords.

Figure B3. Share of AI Jobs by Industry Sector (Non-Public Firms)



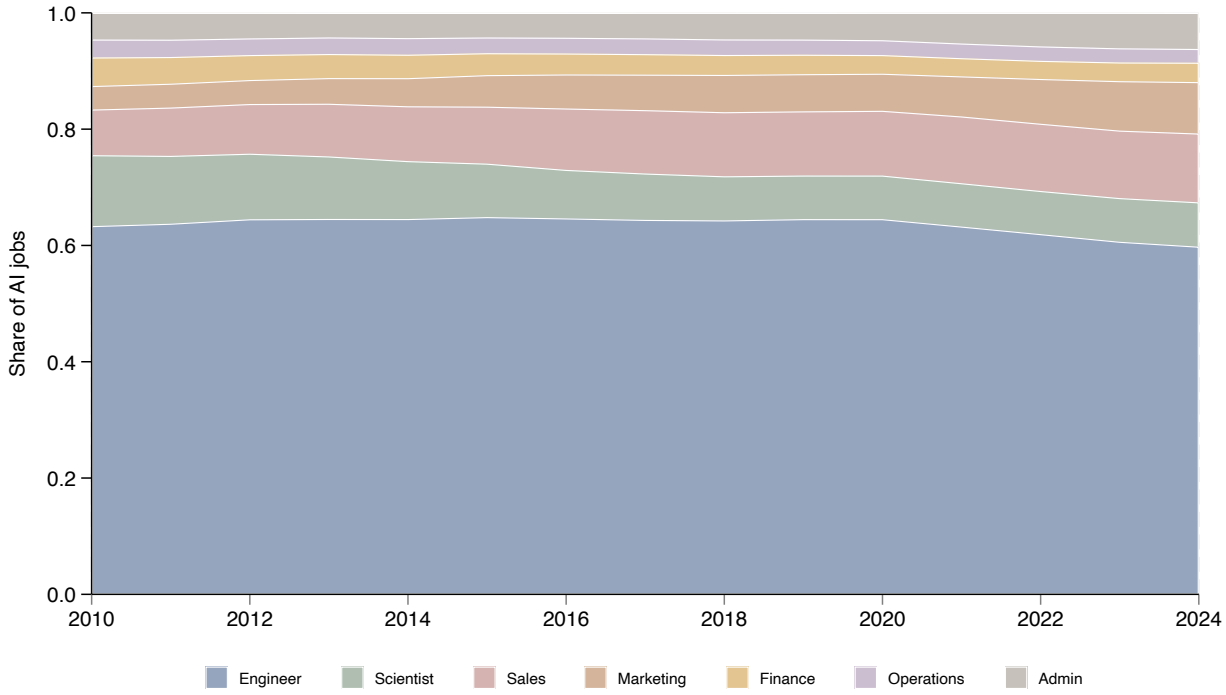
This figure presents the share of U.S. AI jobs at the industry level, based on the sample of jobs among employers that are not U.S. public firms in Revelio. These employers are mostly privately held firms, government entities, and non-profit sector entities. For each sector (defined by two-digit NAICS industry codes), we compute the fraction of AI-related jobs in the resume data. The statistics are computed across all non-public employers in each sector in three years: 2010, 2018, and 2024. In each year, we sum AI jobs within an industry and divide by total jobs in that industry.

Figure B4. Share of AI Jobs by Firm Size Tercile



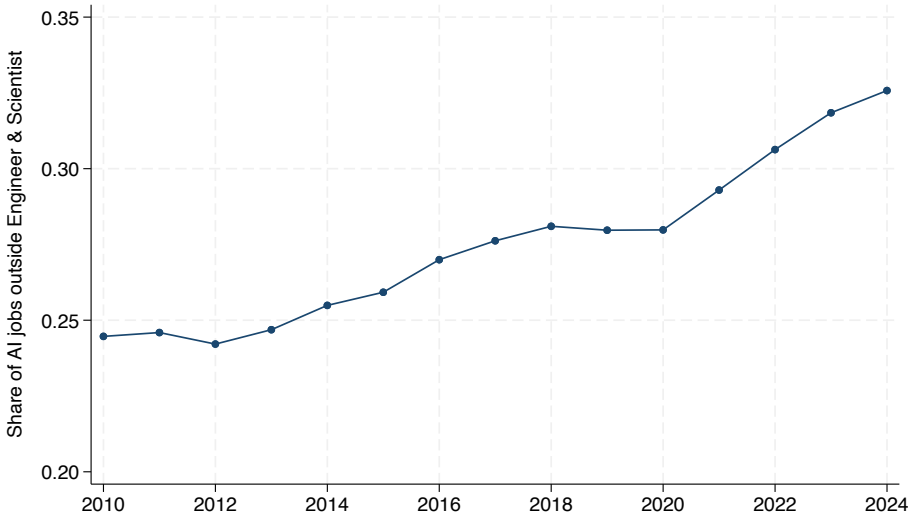
This figure plots the aggregate share of AI-related U.S. jobs by firm-size tercile from 2010 to 2024 for our matched U.S. public-firm sample. Each year, firms are sorted into terciles based on total global Revelio job counts. Within each tercile-year cell, AI share is defined as AI-related U.S. jobs divided by total U.S. jobs across all firms in that cell. AI-related jobs are identified using keyword matches in job titles or job descriptions. Firm-years are restricted to those with a coverage ratio between 0.05 and 3, where the coverage ratio is defined as total global Revelio workers divided by Compustat employment for each firm-year.

Figure B5. AI Jobs by Business Function



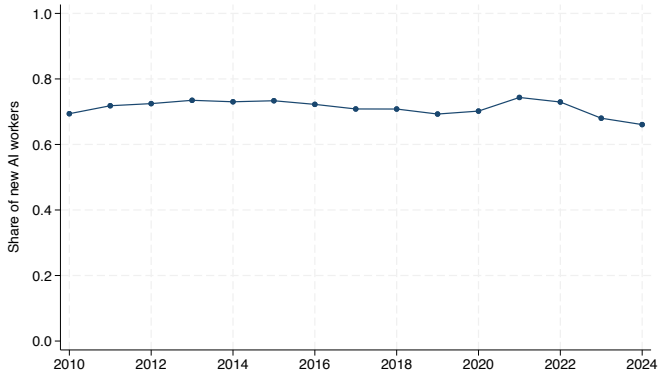
Notes. This figure plots the composition of AI-related U.S. jobs at U.S. public firms across seven broad job functions defined by Revelio from 2010 to 2024. Job functions are as reported by Revelio, which classifies each position into one of seven broad business functions—Engineer, Scientist, Sales, Marketing, Finance, Operations, and Admin—based on its job role.

Figure B6. Share of AI Jobs Outside Engineering and Science Functions

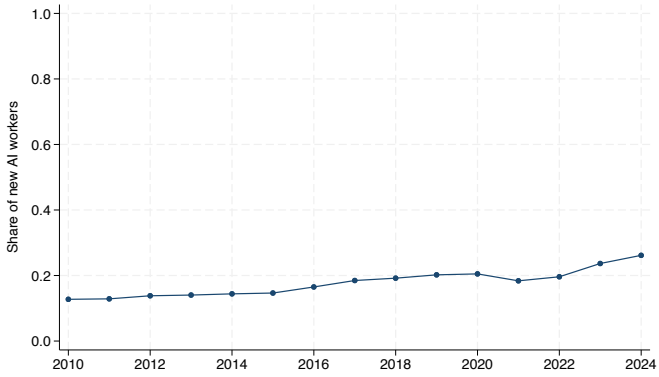


Notes. This figure plots the share of AI-related U.S. jobs from 2010 to 2024 at our matched public-firm sample whose job function is outside the Engineer and Scientist functions. Revelio defines seven categories for each job: Engineer, Scientist, Sales, Marketing, Finance, Operations, and Admin.

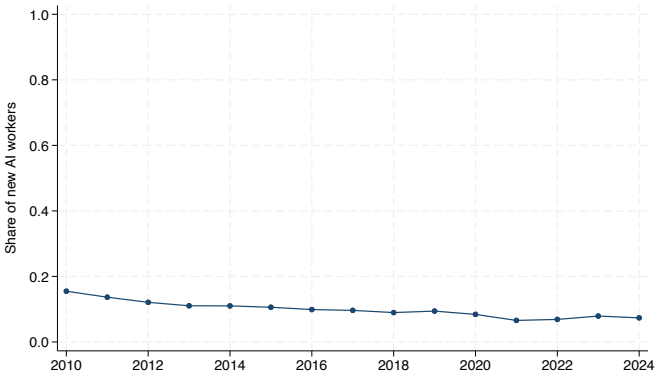
Figure B7. Origins of New AI Workers



(a) Retrained from U.S. Jobs



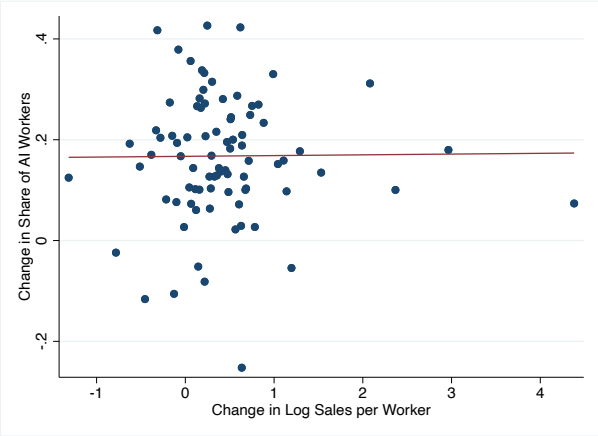
(b) Recent U.S. Graduates



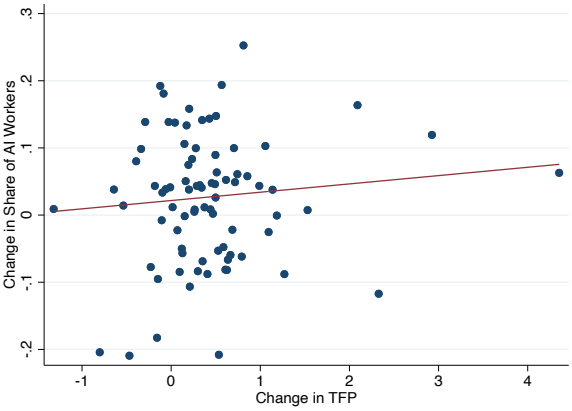
(c) Other

This figure classifies workers who begin their first AI-related job at a U.S. public firm in each year 2010–2024 by their history prior to that job, using their full Revelio employment records (including positions at non-public employers) and education histories. Each panel plots one origin group's share of that year's new AI workers. Panel (a): workers with prior U.S. non-AI work experience. Panel (b): workers whose most recent degree, from a U.S. university, was completed within one year of starting the AI job. Panel (c): workers whose most recent degree, from a non-U.S. university, was completed within one year of starting the AI job, and workers with neither prior U.S. employment nor a U.S. degree before the AI job. The residual group of workers who do not belong to any of the above three categories (e.g. workers with a U.S. degree completed more than one year before the AI job but no recorded U.S. employment prior to it) is about 1% throughout the year 2010–2024.

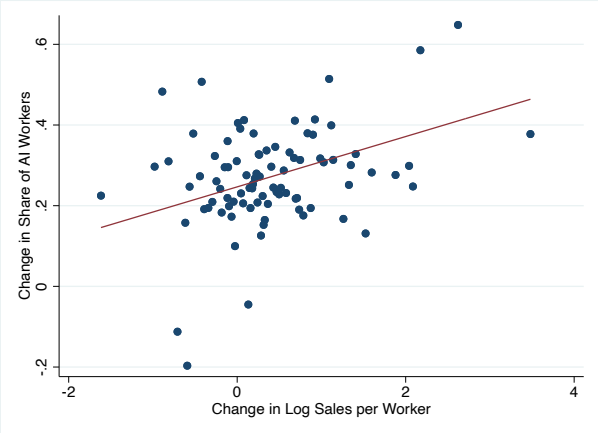
Figure B8. Relationship Between AI Investments and Productivity Growth



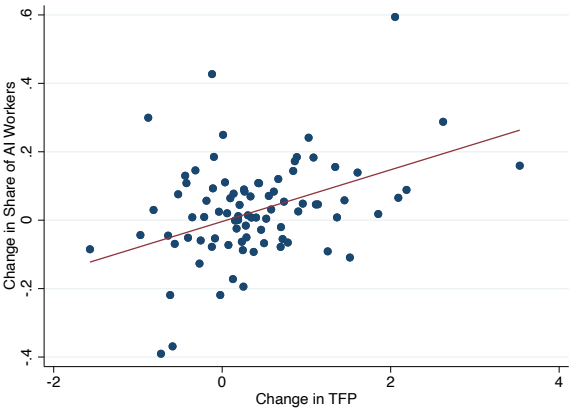
(a) Log sales per worker, 2010–2018



(b) TFP, 2010–2018



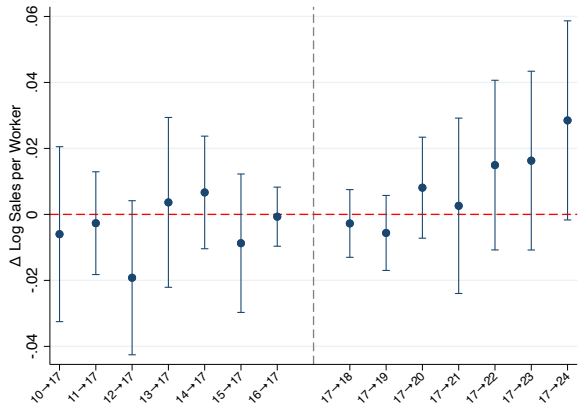
(c) Log sales per worker, 2018–2024



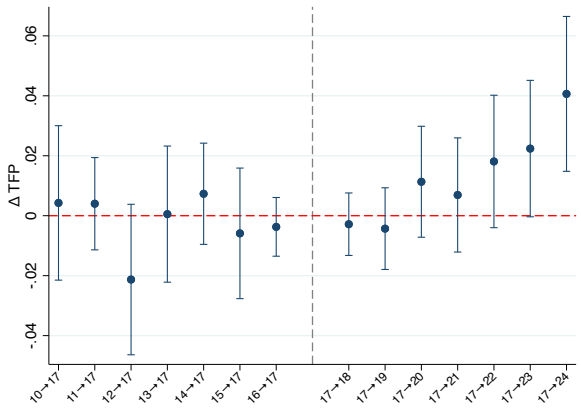
(d) TFP, 2018–2024

This figure is a binned scatterplot of the change in log sales per worker or TFP against the change in the share of AI jobs. All changes are measured from 2010 to 2018 in Panels (a) and (b) and from 2018 to 2024 in Panels (c) and (d). Each dot represents roughly the same number of firms, and the solid line is the fitted regression line. Observations are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects and baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers).

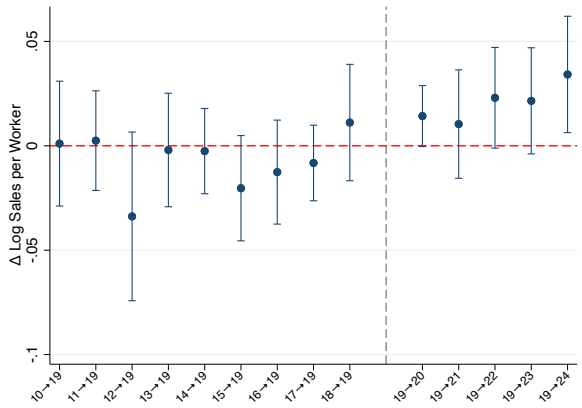
Figure B9. AI Investments and Productivity Changes Over Time: Alternative Base Years



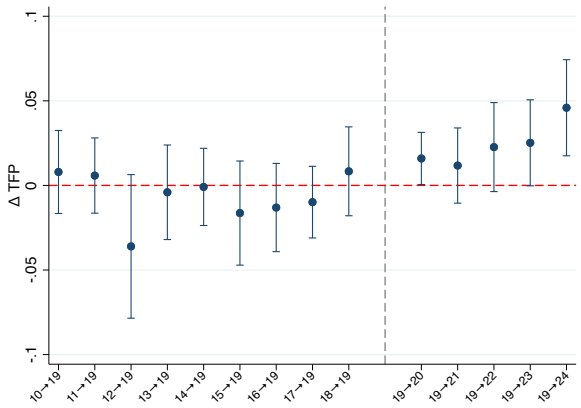
(a) Log Sales per Worker, Base Year 2017



(b) TFP, Base Year 2017



(c) Log Sales per Worker, Base Year 2019



(d) TFP, Base Year 2019

This figure plots the cross-sectional association between AI adoption intensity and productivity growth at different time points, using two alternative base years. In Panels (a) and (b), the independent variable is the standardized change in a firm's share of AI jobs from 2017 to 2024, ΔShare $\text{AI}_i^{17\to24}$; in Panels (c) and (d), it is the corresponding change from 2019 to 2024, ΔShare $\text{AI}_i^{19\to24}$. Within each pair of panels, the independent variable is the same across all regressions. Each point plots the coefficient $\hat{\beta}$ from a separate cross-sectional regression of $\Delta Y_i^{t\to t'}$ on the change in the AI job share, controlling for two-digit NAICS sector fixed effects and weighted by firm-level U.S. employees. Panels (a) and (c) report results for log sales per worker; Panels (b) and (d) report results for total factor productivity. Let B denote the base year (2017 in the top row, 2019 in the bottom row). To the left of the dashed line, the dependent variable is $\Delta Y_i^{T\to B}$ for $T = 2010, \ldots, B-1$. The first dot measures how changes in firm AI investments from B to 2024 relate to the ex-ante change in productivity from 2010 to B; each successive dot shortens the window over which ex-ante productivity is measured by one year, with the last dot on the left using the change from $B-1$ to B. To the right of the dashed line, the dependent variable is $\Delta Y_i^{B\to T}$ for $T = B+1, \ldots, 2024$: the first dot after the dashed line uses the change in productivity from B to $B+1$, and the last dot uses the change from B to 2024. Standard errors are clustered at the five-digit NAICS level.

Table B1. Sample Construction: Effect of Filters on AI Share in 2024

Step	Filter	Firms	Mean (Share of AI Jobs)
0	Compustat–Revelio match + coverage ratio $\in (0.05, 3)$	4,022	1.46%
1	+ U.S. headquartered	3,270	1.39%
2	+ Revelio U.S. workers ≥ 50	2,748	1.42%
3	+ Majority U.S. workers (>50%)	2,262	1.20%
4	+ Compustat employment ≥ 50	2,215	1.18%
5	+ Exclude IT sectors (NAICS 51, 54)	1,923	0.77%
6	+ Drop NAICS 334413 (semiconductor)	1,894	0.73%
7	+ Winsorize Share AI at 1/99%	1,894	0.64%
8	+ Long-difference panel (2018 & 2024)	1,344	0.56%
9	+ Regression sample (controls non-missing)	1,321	0.55%

This table reports the effect of our sample filters on the number of firms and the mean firm-level share of U.S. AI jobs (AI jobs / total U.S. jobs) in 2024. Steps 0–5 apply the basic sample restrictions (coverage, U.S. headquartered, U.S. employment, exclude Information / Professional sectors). Step 6 drops semiconductor manufacturers (NAICS 334413), Step 7 winsorizes the AI share at 1/99%, Step 8 keeps firms observed in both 2018 and 2024 (the long-difference panel), and Step 9 restricts to the exact regression sample (firms with all controls non-missing).

Table B2. Summary Statistics

Panel A: 2010–2018

Variable	N	mean	sd	p1	p5	p10	p25	p50	p75	p90	> p95	p99
Change in share of AI jobs, 2010–2018	1222	.0021	.004	-.0031	-.00019	0	0	.00075	.0028	.006	.0091	.024
Change in log sales, 2010–2018	1222	.48	.67	-1.4	-.53	-.24	.12	.44	.8	1.3	1.7	2.7
Change in log employment, 2010–2018	1222	.29	.61	-1.7	-.69	-.37	-.022	.27	.6	1	1.3	2.2
Change in log sales per worker, 2010–2018	1222	.19	.43	-.74	-.37	-.22	-.0083	.16	.33	.61	.89	1.5
Change in revenue TFP, 2010–2018	1091	.039	.34	-.98	-.48	-.32	-.12	.018	.19	.39	.58	1.4
Share of AI jobs, 2010	1224	.00088	.0028	0	0	0	0	0	.00058	.0022	.0044	.013
Have at least 5 AI jobs, 2010	1222	.1	.31	0	0	0	0	0	0	1	1	1
Share of AI jobs, 2018	1222	.003	.0055	0	0	0	0	.0011	.0034	.0078	.012	.033
Have at least 5 AI jobs, 2018	1222	.35	.48	0	0	0	0	0	1	1	1	1

Panel B: 2018–2024

Variable	N	mean	sd	p1	p5	p10	p25	p50	p75	p90	> p95	p99
Change in share of AI jobs, 2018–2024	1321	.0028	.0047	-.0099	-.0019	-.00017	0	.0015	.0044	.0087	.013	.02
Change in log sales, 2018–2024	1321	.44	.63	-.98	-.41	-.17	.1	.36	.67	1.1	1.5	3.1
Change in log employment, 2018–2024	1321	.16	.48	-1.2	-.57	-.36	-.09	.11	.37	.73	1	1.8
Change in log sales per worker, 2018–2024	1321	.29	.45	-.96	-.27	-.11	.083	.25	.45	.71	.99	2
Change in revenue TFP, 2018–2024	1173	.024	.39	-1.4	-.51	-.32	-.14	.0036	.17	.38	.59	1.7
Share of AI jobs, 2018	1321	.0029	.0054	0	0	0	0	.00094	.0034	.0076	.012	.031
Have at least 5 AI jobs, 2018	1321	.31	.46	0	0	0	0	0	1	1	1	1
Share of AI jobs, 2024	1321	.0055	.0073	0	0	0	.00057	.0032	.0071	.014	.02	.038
Have at least 5 AI jobs, 2024	1321	.48	.5	0	0	0	0	0	1	1	1	1

This table reports summary statistics for the sample of firms in our baseline regressions (including 1222 firms in the regressions from 2010 to 2018 and 1321 firms in the regressions from 2018 to 2024). For each variable, we report the number of observations, the mean, the standard deviation, the median, and 1st, 5th, 10th, 25th, 75th, 90th, 95th, and 99th percentiles.

Table B3. AI Investments and Firm Growth

Panel A: 2010–2018

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.092**	0.075**	0.093***	0.074*
	(0.038)	(0.038)	(0.029)	(0.039)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.127	0.264	0.114	0.212
Observations	1,222	1,222	1,222	1,222

Panel B: 2018–2024

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.068***	0.067**	0.020	0.005
	(0.026)	(0.034)	(0.018)	(0.026)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.162	0.211	0.055	0.109
Observations	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of firm growth on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider three measures of firm growth: changes in log sales (columns 1 and 2) and changes in log employment (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2010 to 2018 in Panel A and from 2018 to 2024 in Panel B. Regressions are weighted by the number of workers in Revelio at the start of the period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period (2010 in Panel A and 2018 in Panel B): firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B4. AI Investments and Firm Growth Based on the Cognism Sample

Panel A: 2010–2018

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.145***	0.140***	0.139***	0.129***
	(0.042)	(0.047)	(0.033)	(0.049)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.146	0.298	0.111	0.238
Observations	627	627	627	627

Panel B: 2018–2024

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.082**	0.077	0.014	-0.007
	(0.038)	(0.049)	(0.020)	(0.029)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.217	0.284	0.102	0.157
Observations	493	493	493	493

This table reports the coefficients from long-differences regressions of firm growth on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors) that are covered in the Cognism dataset. We consider three measures of firm growth: changes in log sales (columns 1 and 2) and changes in log employment (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2010 to 2018 in Panel A and from 2018 to 2024 in Panel B. Regressions are weighted by the number of workers in Revelio at the start of the period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period (2010 in Panel A and 2018 in Panel B): firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B5. AI Investments and Employment in Revelio

	Δ Log Number of Workers in Revelio		Δ Log Number of Non-AI Workers in Revelio	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.011	0.004	0.006	-0.002
	(0.011)	(0.014)	(0.011)	(0.014)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.135	0.210	0.130	0.205
Observations	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of the number of workers in Revelio on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). The dependent variables are log change in the number of workers in Revelio in columns 1 and 2 and log change in the number of non-AI workers in Revelio in columns 3 and 4. The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B6. AI Investments and Productivity across Sectors, 2018–2024

	Manufacturing		Wholesale & Retail		Finance		Other Non-tech		Tech	
	Sales per Worker	TFP	Sales per Worker	TFP	Sales per Worker	TFP	Sales per Worker	TFP	Sales per Worker	TFP
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Share AI Jobs	0.046**	0.059**	0.135**	0.123**	0.068	0.102	0.098***	0.087***	0.029	0.034*
	(0.019)	(0.024)	(0.065)	(0.055)	(0.069)	(0.072)	(0.033)	(0.031)	(0.019)	(0.019)
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.239	0.218	0.243	0.216	0.132	0.217	0.276	0.188	0.246	0.197
Observations	505	474	146	141	273	227	397	331	176	161

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments, separately for each sector: Manufacturing (columns 1 and 2), Wholesale & Retail (columns 3 and 4), Finance (columns 5 and 6), Other Non-tech (columns 7 and 8), and Tech (Information and Professional Services, columns 9 and 10). For each sector, we report results for changes in log sales per worker and changes in revenue TFP. The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for detailed industry fixed effects and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B7. AI Investment and Subsequent Productivity Growth

	Δ Log Sales per Worker, t to $t+4$			Δ TFP, t to $t+4$		
	t =2018	t =2019	t =2020	t =2018	t =2019	t =2020
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	-0.005	0.012	0.017*	-0.006	0.018	0.030***
	(0.013)	(0.013)	(0.011)	(0.017)	(0.013)	(0.011)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.166	0.089	0.249	0.070	0.093	0.140
Observations	1,215	1,205	1,179	1,127	1,120	1,104

This table reports the coefficients from OLS regressions of changes in firm productivity from year t to $t+4$ on the firm-level changes in AI investments among U.S. public firms (in non-tech sectors) from year $t-4$ to t, where year t is 2018, 2019, or 2020. We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in year t. All specifications control for 2-digit NAICS industry sector fixed effects and include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B8. AI Investment and Previous Productivity Growth

	Δ Log Sales per Worker, $t-4$ to t			Δ TFP, $t-4$ to t		
	t=2014	t=2015	t=2016	t=2014	t=2015	t=2016
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers, t to $t+4$	-0.015	0.006	-0.001	0.004	0.015	0.002
	(0.019)	(0.013)	(0.015)	(0.014)	(0.010)	(0.018)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.169	0.066	0.054	0.082	0.040	0.029
Observations	1,208	1,188	1,187	1,086	1,067	1,063

This table reports the coefficients from OLS regressions of changes in firm productivity from year $t-4$ to t on the firm-level changes in AI investments among U.S. public firms (in non-tech sectors) from year t to $t+4$, where year t is 2014, 2015, or 2016. We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in year t. All specifications control for 2-digit NAICS industry sector fixed effects and include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B9. AI Investments and Productivity Using Alternative TFP Measure, 2018–2024

	Δ TFP (İmrohoroğlu and Tüzel)		Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	0.063*	0.066**	0.053**	0.077***	0.056***	0.078***
	(0.033)	(0.032)	(0.025)	(0.025)	(0.019)	(0.021)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.078	0.179	0.150	0.238	0.114	0.208
Observations	709	709	709	709	709	709

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider three measures of productivity: changes in the TFP measure from Imrohoroğlu and Tüzel (2014) (columns 1 and 2), changes in log sales per worker (columns 3 and 4) and changes in the baseline revenue TFP measure (columns 5 and 6). In all columns, the sample consists of firms with a non-missing 2018–2024 change in the İmrohoroğlu and Tüzel TFP measure. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2, 4, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B10. AI Investments and Productivity Excluding Entrants After 2010, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.039**	0.049**	0.053***	0.064***
	(0.017)	(0.019)	(0.013)	(0.015)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.237	0.324	0.105	0.216
Observations	881	881	818	818

This table reports the coefficients from long-differences regressions of firm productivity on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), excluding firms that entered after 2010. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B11. AI Investments and Productivity Using Share of Wage Bill of AI Workers, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Share of Wage Bill of AI Workers	0.043**	0.053**	0.053***	0.063***
	(0.019)	(0.021)	(0.019)	(0.021)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.170	0.216	0.055	0.132
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is the growth in the share of the wage bill going to AI-related jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of wage bill of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B12. AI Investments and Productivity Using Global Share of AI Workers, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.044**	0.058**	0.057***	0.070***
	(0.021)	(0.026)	(0.021)	(0.026)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.173	0.221	0.063	0.143
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in global AI investments among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the global share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B13. AI Investments and Productivity Using 2025 Revelio Labs Vintage, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.051**	0.063***	0.065***	0.079***
	(0.022)	(0.024)	(0.020)	(0.023)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.177	0.222	0.070	0.148
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the global share of AI jobs, based on the Revelio Labs vintage delivered in October 2025 and standardized to mean zero and standard deviation one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B14. AI Investments and Productivity Using Year-Specific AI Skills

Panel A: 2010–2018

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs (year-specific)	-0.004	-0.000	0.014	0.011
	(0.014)	(0.015)	(0.012)	(0.011)
Industry FE	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes
Adj R-Squared	0.089	0.171	0.047	0.150
Observations	1,222	1,222	1,091	1,091

Panel B: 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs (year-specific)	0.049**	0.064**	0.062***	0.078***
	(0.022)	(0.026)	(0.021)	(0.025)
Industry FE	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes
Adj R-Squared	0.179	0.229	0.075	0.158
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity on changes in AI investments among U.S. public firms in non-tech sectors. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is the change in the share of AI jobs using the year-specific measure constructed from Revelio Labs resume data, standardized to mean zero and standard deviation of one. To construct the exact-year measure, we build a separate AI keyword list for each year: AI-related skills are identified from Lightcast job postings using the same method as in Appendix A.2.1, and a skill's keywords enter year t's list only if the skill appears in postings in year t. A job is then classified as an AI job in year t only if its title or description contains a keyword from year t's list. All changes in productivity and the share of AI jobs are measured from 2010 to 2018 in Panel A and from 2018 to 2024 in Panel B. Regressions are weighted by the number of workers in Revelio at the start of the period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include baseline controls measured at the start of each period (2010 in Panel A and 2018 in Panel B): firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B15. AI Investments and Productivity Excluding Subsidiaries' AI Investments, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.049**	0.062**	0.058***	0.071***
	(0.021)	(0.024)	(0.020)	(0.024)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.179	0.228	0.068	0.151
Observations	1,309	1,309	1,162	1,162

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs, excluding jobs in subsidiary firms, standardized to mean zero and standard deviation one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B16. AI Investments and Balance-Sheet Intangibles, 2018–2024

	Δ Balance-Sheet Intangibles/Total Assets		Δ Acquired Intangibles/Total Assets	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.006	0.002	-0.003	-0.005
	(0.004)	(0.004)	(0.003)	(0.003)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.099	0.149	0.102	0.156
Observations	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of firms' intangibles on the balance sheet on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). In columns 1 and 2, the dependent variable is the change in balance-sheet intangibles (*intan* from Compustat) normalized by total assets; in columns 3 and 4, the dependent variable is the change in acquired intangibles (*acqintan* from Compustat) normalized by total assets. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B17. AI Investments and Productivity Excluding Magnificent 7 Stocks, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.061**	0.073**	0.070***	0.087***
	(0.026)	(0.029)	(0.026)	(0.027)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.193	0.238	0.083	0.160
Observations	1,319	1,319	1,171	1,171

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors) and excluding Magnificent 7 stocks (Apple, Microsoft, Amazon, Alphabet, Meta, Tesla, and Nvidia). We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B18. AI Investments and Productivity Controlling for State Fixed Effects and Initial Productivity, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.056***	0.052***	0.066***	0.054***
	(0.020)	(0.018)	(0.018)	(0.018)
TFP in 2010		-0.795***		-0.821***
		(0.133)		(0.117)
Log Sales per Worker in 2010		-0.068**		-0.019
		(0.031)		(0.031)
Industry FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Adj R-Squared	0.295	0.453	0.240	0.415
Observations	1,320	1,179	1,172	1,172

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), controlling for state fixed effects and initial productivity levels. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. Columns 2 and 4 additionally control for initial TFP in 2010 and log sales per worker in 2010. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for industry fixed effects, state fixed effects, and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B19. AI Investments and Productivity Controlling for Pre-trends, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.059**	0.067**	0.077***	0.081***
	(0.029)	(0.029)	(0.027)	(0.027)
Industry FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Pre-trend in sales and employment	Y	Y	Y	Y
Pre-trend in TFP	N	Y	N	Y
Observations	1,007	879	928	876

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), controlling for pre-trends in firm outcomes. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All specifications control for pre-trends in sales and employment from 2010 to 2018. Columns 2 and 4 additionally control for pre-trends in TFP. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for industry fixed effects and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B20. AI Investments and Productivity Controlling for Detailed Industry Fixed Effects, 2018–2024

	Δ Log Sales per Worker				Δ TFP			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Jobs	0.062**	0.070**	0.078*	0.086*	0.075***	0.092***	0.087**	0.095**
	(0.026)	(0.030)	(0.041)	(0.046)	(0.024)	(0.028)	(0.041)	(0.044)
Industry FE	Y	N	N	N	Y	N	N	N
NAICS3 FE	N	Y	N	N	N	Y	N	N
NAICS4 FE	N	N	Y	N	N	N	Y	N
NAICS5 FE	N	N	N	Y	N	N	N	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.226	0.273	0.321	0.336	0.153	0.213	0.262	0.255
Observations	1,321	1,311	1,266	1,171	1,173	1,166	1,119	1,031

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), controlling for increasingly detailed industry fixed effects. We consider two measures of productivity: changes in log sales per worker (columns 1–4) and changes in revenue TFP (columns 5–8). Each column uses a different level of industry classification: 2-digit NAICS (columns 1 and 5), 3-digit NAICS (columns 2 and 6), 4-digit NAICS (columns 3 and 7), and 5-digit NAICS (columns 4 and 8). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B21. AI Investments and Productivity Controlling for Covid Exposure and Remote Work Suitability, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.065**	0.062**	0.078***	0.074***
	(0.028)	(0.027)	(0.027)	(0.026)
Covid Exposure	-0.029*		-0.032**	
	(0.015)		(0.015)	
Remote Work Suitability Share		-0.012		0.033
		(0.122)		(0.138)
Industry FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Adj R-Squared	0.228	0.225	0.163	0.152
Observations	1,138	1,321	1,021	1,173

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), controlling for covid exposure and remote work suitability. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for industry fixed effects, and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 1 and 3 additionally control for firms' Covid exposure between 2020 Q2 and Q4 in Hassan et al. (2023), and columns 2 and 4 additionally control for the share of remote-work-suitable jobs (from Dingel and Neiman (2020)) in 2018. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B22. AI Investments and Productivity Controlling for Non-AI IT Investments, 2018–2024

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.047**	0.060**	0.061***	0.074***
	(0.022)	(0.026)	(0.020)	(0.025)
Δ Share Non-AI IT Jobs	0.030*	0.016	0.003	0.014
	(0.018)	(0.018)	(0.021)	(0.016)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.184	0.227	0.070	0.154
Observations	1,321	1,321	1,173	1,173

This table reports the coefficients from long-differences regressions of firm productivity from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), controlling for contemporaneous firm-level changes in non-AI IT investments. We consider two measures of productivity: changes in log sales per worker (columns 1 and 2) and changes in revenue TFP (columns 3 and 4). The main independent variable is growth in the share of AI jobs, standardized to mean zero and standard deviation one. The control is growth in the share of non-AI IT jobs, defined as non-AI jobs in SOC 15-1200 computer and IT occupations, also standardized to mean zero and standard deviation one. Both measures are constructed from the resume data. All changes in productivity and share of AI and IT jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for industry fixed effects, and include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B23. AI Investments and Firm Tangible and Intangible Capital

	Δ Log SG&A		Δ Log R&D		Δ Log PP&E	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	0.051**	0.052**	0.054**	0.025*	-0.012	-0.007
	(0.022)	(0.023)	(0.025)	(0.014)	(0.025)	(0.024)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.103	0.169	0.195	0.285	0.068	0.193
Observations	1,066	1,066	1,321	1,321	1,238	1,238

This table reports the coefficients from long-differences regressions of firm measures of intangible and tangible capital from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider three dependent variables: changes in log SG&A (columns 1 and 2), changes in log R&D (columns 3 and 4), and changes in log PP&E (columns 5 and 6). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2, 4, and 6 also include the baseline controls measured at the start of the period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B24. AI Investments and Employment Shares by Education and Seniority Levels, 2018–2024

Panel A: Education

	No College Degree		College Degrees		Master Degrees		Doctor Degrees	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Jobs	-0.686***	-0.454***	0.156	-0.045	0.462***	0.438***	0.069	0.062**
	(0.171)	(0.159)	(0.136)	(0.138)	(0.146)	(0.090)	(0.046)	(0.024)
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.224	0.301	0.218	0.323	0.128	0.273	0.105	0.224
Observations	1,321	1,321	1,321	1,321	1,321	1,321	1,321	1,321

Panel B: Seniority

	Low Seniority		Medium Seniority		High Seniority	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	-1.134***	-1.168***	0.213	0.270	0.922***	0.899***
	(0.160)	(0.236)	(0.154)	(0.197)	(0.077)	(0.078)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.259	0.316	0.146	0.185	0.280	0.311
Observations	1,321	1,321	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of changes in employment shares by education and seniority levels from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). The dependent variables are changes in the share of employment (in percentages) of workers in each education level in Panel A and in each seniority level in Panel B. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Even-numbered columns also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B25. AI Investments and Log Employment by Education and Seniority Levels

Panel A: Education

	No College Degree		College Degrees		Master Degrees		Doctor Degrees	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Jobs	-0.014	-0.030*	0.015	-0.000	0.052***	0.041**	0.057***	0.050***
	(0.017)	(0.018)	(0.014)	(0.018)	(0.016)	(0.016)	(0.021)	(0.018)
Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.069	0.128	0.173	0.253	0.279	0.355	0.382	0.482
Observations	1,321	1,321	1,321	1,321	1,321	1,321	1,321	1,321

Panel B: Seniority

	Low Seniority		Medium Seniority		High Seniority	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	-0.017	-0.023	0.005	0.003	0.029**	0.021
	(0.011)	(0.017)	(0.014)	(0.014)	(0.013)	(0.016)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.078	0.164	0.091	0.142	0.112	0.149
Observations	1,321	1,321	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of changes in log employment by education and seniority levels from 2018 to 2024 on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). The dependent variables are changes in log employment (in percentages) of workers in each education level in Panel A and in each seniority level in Panel B. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Even-numbered columns also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B26. AI Investments and Firm Wages

	Δ Log Average Wage		Δ Log Average Wage of Non-AI Workers	
	(1)	(2)	(3)	(4)
Δ Share AI Jobs	0.013***	0.017***	0.011***	0.016***
	(0.004)	(0.004)	(0.004)	(0.004)
Industry FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.193	0.254	0.182	0.242
Observations	1,321	1,321	1,321	1,321

This table reports the coefficients from long-differences regressions of firm wages on contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider two measures of wages: changes in log average wage (columns 1 and 2), and changes in log average wage of non-AI workers (columns 3 and 4). The main independent variable is growth in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in firm outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of each period (2010 in Panel A and 2018 in Panel B). All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured at the start of each period: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B27. AI Investment and Subsequent Organization Capital Growth

	Δ Log(Organizational Capital/PP&E), t to $t+4$			Δ Log(Organizational Capital/Total Assets), t to $t+4$		
	t=2018	t=2019	t=2020	t=2018	t=2019	t=2020
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.052**	0.050**	0.046*	0.083**	0.068**	0.048**
	(0.026)	(0.020)	(0.024)	(0.039)	(0.033)	(0.023)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.056	0.103	0.098	0.063	0.104	0.093
Observations	841	826	793	1,025	1,018	990

This table reports the coefficients from OLS regressions of changes in organization capital (capitalized SG&A) from year t to $t+4$ on the firm-level changes in AI investments among U.S. public firms (in non-tech sectors) from year $t-4$ to t, where year t is 2018, 2019, or 2020. The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 to 3 and changes in the natural log of organization capital normalized by total assets in columns 4 to 6. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in year t. All specifications control for 2-digit NAICS industry sector fixed effects and include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B28. AI Investment and Previous Organization Capital Growth

	Δ Log(Organizational Capital/PP&E), $t-4$ to t			Δ Log(Organizational Capital/Total Assets), $t-4$ to t		
	t=2014	t=2015	t=2016	t=2014	t=2015	t=2016
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	-0.062	-0.024	-0.040	-0.013	0.022	-0.006
	(0.040)	(0.033)	(0.039)	(0.012)	(0.016)	(0.043)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.077	0.087	0.114	0.109	0.070	0.065
Observations	848	824	812	1,019	999	992

This table reports the coefficients from OLS regressions of changes in organization capital (capitalized SG&A) from year $t-4$ to t on the firm-level changes in AI investments among U.S. public firms (in non-tech sectors) from year t to $t+4$, where year t is 2014, 2015, or 2016. The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 to 3 and changes in the natural log of organization capital normalized by total assets in columns 4 to 6. The main independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs), standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in year t. All specifications control for 2-digit NAICS industry sector fixed effects and include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B29. AI Investments and Organization Capital: IV Estimates

	Δ Log(Organizational Capital/PP&E)			Δ Log(Organizational Capital/Total Assets)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Jobs	0.318	0.279*	0.303**	0.253	0.189	0.239**
	(0.244)	(0.146)	(0.138)	(0.259)	(0.129)	(0.109)
Industry FE	Y	Y	Y	Y	Y	Y
CS Control	Y	Y	Y	Y	Y	Y
Baseline Controls	Y	Y	Y	Y	Y	Y
State FE	N	Y	Y	N	Y	Y
Control Pre-trend	N	N	Y	N	N	Y
F Statistic	6.3	16.6	19.2	6.3	16.5	22.8
Observations	419	408	383	523	514	439

This table estimates the relationship between AI investments and growth in organization capital from 2018 to 2024 for U.S. public firms (in non-tech sectors), where firm AI investments are instrumented with ex-ante firm-level exposure to AI-skilled graduates from AI-strong universities. The independent variable is growth in the share of AI jobs (based on the resume data from Revelio Labs) from 2018 to 2024, standardized to mean zero and standard deviation of one. Regressions are weighted by the number of workers in Revelio in 2018 (see the definition of the instrument in Section 3.4 and the details of instrument construction in Appendix C). The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 to 3 and changes in the natural log of organization capital normalized by total assets in columns 4 to 6. All specifications control for industry sector fixed effects and ex-ante exposure to universities that are strong in computer science research as well as top 20 universities. All columns also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 2, 3, 5, and 6 add state fixed effects, and columns 3 and 6 additionally control for firm-level changes in log sales, log employment, and TFP from 2002 to 2010. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. The first-stage F-statistics of the instrument are reported for all specifications. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B30. Organization-Capital-Related AI Jobs and Productivity with Alternative Thresholds, 2018–2024

Panel A: Using 0.5 Threshold for Organization-Capital-Related Jobs

	Δ Log Sales per Worker			Δ TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.068***	0.083***	0.089***	0.080***	0.099***	0.100***
	(0.024)	(0.028)	(0.028)	(0.024)	(0.026)	(0.026)
Δ Share of Non-Organization-Related AI Jobs			-0.014			-0.004
			(0.019)			(0.019)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.191	0.241	0.241	0.084	0.172	0.172
Observations	1,306	1,306	1,306	1,163	1,163	1,163

Panel B: Using 0.7 Threshold for Organization-Capital-Related Jobs

	Δ Log Sales per Worker			Δ TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.064***	0.079***	0.054***	0.076***	0.095***	0.061***
	(0.024)	(0.026)	(0.019)	(0.025)	(0.026)	(0.019)
Δ Share of Non-Organization-Related AI Jobs			0.040			0.052*
			(0.029)			(0.028)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.178	0.226	0.233	0.062	0.149	0.164
Observations	1,306	1,306	1,306	1,163	1,163	1,163

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in organization capital and non-organization-capital-related AI jobs among U.S. public firms (in non-tech sectors) under alternative thresholds for organization-capital-related jobs. We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). The independent variables are growth in the share of organization-capital-related AI jobs (normalized by firms' total U.S.-based jobs) in columns 1, 2, 4 and 5, and the share of AI jobs not related to organization capital (normalized by firms' total U.S.-based jobs) in columns 3 and 6, which are both standardized to mean zero and standard deviation of one. In Panel A, a job is organization-capital-related if the organization-capital score (based on job description in the Revelio resume data) is at least 0.5. In Panel B, a job is organization-capital-related if the organization-capital score is at least 0.7. All changes in productivity and the shares of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2, 3, 5, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B31. Organization-Capital-Related AI Jobs and Organization Capital, 2018–2024

	Δ Log(Organizational Capital/PP&E)			Δ Log(Organizational Capital/Total Assets)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.145***	0.131***	0.115***	0.159***	0.134***	0.103**
	(0.019)	(0.024)	(0.028)	(0.048)	(0.041)	(0.043)
Δ Share of Non-Organization-Related AI Jobs			0.038			0.065*
			(0.033)			(0.036)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.120	0.146	0.149	0.105	0.129	0.137
Observations	877	877	877	1,092	1,092	1,092

This table reports the coefficients from long-differences regressions of organization capital on the contemporaneous firm-level changes in organization capital and non-organization-capital-related AI jobs among U.S. public firms (in non-tech sectors). The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 to 3 and changes in the natural log of organization capital normalized by total assets in columns 4 to 6. The independent variables are growth in the share of organization-capital-related AI jobs (normalized by firms' total U.S.-based jobs) in columns 1, 2, 4 and 5, and the share of AI jobs not related to organization capital (normalized by firms' total U.S.-based jobs) in columns 3 and 6, which are both standardized to mean zero and standard deviation of one. A job is organization-capital-related if the organization-capital score (based on job description in the Revelio resume data) is at least 0.6. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2, 3 5, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B32. Organization-Capital-Related AI Jobs and Productivity Controlling for Initial Organization Capital, 2018–2024

	Δ Log Sales per Worker			Δ TFP		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.087***	0.110***	0.107***	0.090***	0.115***	0.102***
	(0.025)	(0.028)	(0.027)	(0.023)	(0.025)	(0.025)
Δ Share of Non-Organization-Related AI Jobs			0.006			0.027
			(0.027)			(0.026)
Log(Organizational Capital/PP&E), 2018	-0.010	-0.005	-0.005	-0.027**	-0.020**	-0.020**
	(0.012)	(0.010)	(0.010)	(0.011)	(0.009)	(0.009)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.164	0.232	0.231	0.142	0.217	0.220
Observations	881	881	881	842	842	842

This table reports the coefficients from long-differences regressions of firm productivity on the contemporaneous firm-level changes in organization capital and non-organization-capital-related AI jobs among U.S. public firms (in non-tech sectors). We consider two measures of productivity: changes in log sales per worker (columns 1 to 3) and changes in revenue TFP (columns 4 to 6). The independent variables are growth in the share of organization-capital-related AI jobs (normalized by firms' total U.S.-based jobs) in columns 1, 2, 4 and 5, and the share of AI jobs not related to organization capital (normalized by firms' total U.S.-based jobs) in columns 3 and 6, which are both standardized to mean zero and standard deviation of one. A job is organization-capital-related if the organization-capital score (based on job description in the Revelio resume data) is at least 0.6. All changes in productivity and the shares of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects and the natural log of organization capital normalized by PP&E in 2018. Columns 2, 3, 5, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B33. Organization-Capital-Related AI Jobs and Organization Capital Controlling for Initial Organization Capital, 2018–2024

	Δ Log(Organizational Capital/PP&E)			Δ Log(Organizational Capital/Total Assets)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of Organization-Related AI Jobs	0.129***	0.114***	0.109***	0.151***	0.123**	0.110**
	(0.018)	(0.020)	(0.023)	(0.051)	(0.048)	(0.049)
Δ Share of Non-Organization-Related AI Jobs			0.013			0.032
			(0.032)			(0.037)
Log(Organizational Capital/PP&E), 2018	0.077*	0.091**	0.090**	0.092**	0.108***	0.106***
	(0.045)	(0.044)	(0.045)	(0.038)	(0.035)	(0.037)
Industry FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	Y	N	Y	Y
Adj R-Squared	0.179	0.215	0.214	0.178	0.211	0.211
Observations	877	877	877	881	881	881

This table reports the coefficients from long-differences regressions of organization capital on the contemporaneous firm-level changes in organization capital and non-organization-capital-related AI jobs among U.S. public firms (in non-tech sectors). The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 to 3 and changes in the natural log of organization capital normalized by total assets in columns 4 to 6. The independent variables are growth in the share of organization-capital-related AI jobs (normalized by firms' total U.S.-based jobs) in columns 1, 2, 4 and 5, and the share of AI jobs not related to organization capital (normalized by firms' total U.S.-based jobs) in columns 3 and 6, which are both standardized to mean zero and standard deviation of one. A job is organization-capital-related if the organization-capital score (based on job description in the Revelio resume data) is at least 0.6. All changes in productivity and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the number of workers in Revelio at the start of the period. All specifications control for 2-digit NAICS industry sector fixed effects and the natural log of organization capital normalized by PP&E in 2018. Columns 2, 3 5, and 6 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B34. Heterogeneous Relationship between AI Investments and Organization Capital Growth by Initial Organization Capital, 2018–2024

	Δ Log(Organizational Capital/PP&E)		Δ Log(Organizational Capital/Total Assets)	
	(1)	(2)	(3)	(4)
Δ Share AI Workers*Below-Median Organization Capital	0.153***	0.141***	0.214***	0.193***
	(0.035)	(0.043)	(0.049)	(0.052)
Δ Share AI Workers*Above-Median Organization Capital	0.064***	0.065**	0.024	0.015
	(0.024)	(0.030)	(0.024)	(0.029)
NAICS2*Size tercile*Above-median org capital FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.162	0.188	0.187	0.214
Observations	871	871	875	875

This table reports the coefficients from long-differences regressions of organization capital (capitalized SG&A) on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors), separately for firms with above-median and below-median initial level of organization capital. Firms within each 2-digit NAICS industry sector are split at the median level of organizational capital normalized by PP&E in 2018. The dependent variables are changes in the natural log of organization capital normalized by PP&E in columns 1 and 2 and changes in the natural log of organization capital normalized by total assets in columns 3 and 4. The main independent variable is growth in the AI employment share, standardized to mean zero and unit standard deviation, interacted with an indicator for above- versus below-median initial organization capital. All changes in productivity and share of AI jobs are measured from 2010 to 2018. Regressions are weighted by the number of workers in Revelio in 2018. All specifications control for 2-digit NAICS industry sector fixed effects. Columns 2 and 4 also include the baseline controls measured in 2018: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B35. AI Investments and Industry Productivity and Prices, 2018–2024

	Δ Log Sales per Worker		Δ TFP		Δ Log PPI	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of AI Jobs	0.073***	0.078***	0.040***	0.033***	0.005	0.009
	(0.023)	(0.030)	(0.010)	(0.010)	(0.016)	(0.014)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.217	0.214	0.247	0.268	0.117	0.201
Observations	287	287	287	287	210	210

This table reports the coefficients from industry-level long-differences regressions of the changes in industry productivity and prices on contemporaneous changes in industry-level AI investments from 2018 to 2024. All industry-level variables are calculated for all firms in Compustat (regardless of whether they are in our main regression sample or not). Each observation is a 5-digit NAICS industry, and (as in our main analysis) we exclude tech sectors. The independent variable is the change in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. All changes in industry outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the total (industry-level) number of Revelio resumes in 2018. The dependent variables are the changes in log sales per worker (log of total industry sales divided by total industry employment) in columns 1 and 2, sales-weighted average TFP in columns 3 and 4, and industry-level producer price index (PPI) in columns 5 and 6. All specifications control for industry sector fixed effects. Regressions in columns 2, 4, and 6 also include industry-level controls for log total employment, log total sales, and log average wage in 2010. Standard errors are robust against heteroskedasticity and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table B36. AI Investments and Industry Prices and Employment Including Non-Public Firms, 2018–2024

	Δ Log PPI		Δ Log Emp (Revelio)		Δ Log Emp (QWI)	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share of AI Jobs	-0.018	-0.009	0.008	-0.011	-0.008	-0.019
	(0.014)	(0.013)	(0.015)	(0.018)	(0.017)	(0.019)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.110	0.389	0.322	0.374	0.183	0.246
Observations	210	181	287	248	284	246

This table reports the coefficients from industry-level long-differences regressions of the changes in industry prices and employment on contemporaneous changes in industry-level AI investments from 2018 to 2024. Each observation is a 5-digit NAICS industry, and (as in our main analysis) we exclude tech sectors. All changes in industry outcomes and share of AI jobs are measured from 2018 to 2024. Regressions are weighted by the total (industry-level) number of Revelio resumes in 2018. All industry-level variables are calculated for all public and private firms in Compustat (regardless of whether they are in our main regression sample or not). The independent variable is the change in the share of AI jobs (based on the resume data), standardized to mean zero and standard deviation of one. The dependent variables are industry-level producer price index (PPI) in columns 1 and 2, the changes in log sales per worker (log of total industry sales divided by total industry employment) in columns 3 and 4, and sales-weighted average TFP in columns 5 and 6. All specifications control for industry sector fixed effects. Regressions in columns 2, 4, and 6 also include industry-level controls for log total employment, log total sales, and log average wage in 2010. Standard errors are robust against heteroskedasticity and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

C. Instrument Construction and Validation

Our instrumental variable strategy builds on Babina et al. (2024) and exploits differences in firms' pre-existing ties to universities with historically strong AI research traditions to generate variation in firm-level AI investment. Building the instrument requires two distinct data inputs. The first is an indicator of each university's pre-period strength in AI research. The second is a measure of the hiring linkages between firms and specific universities prior to the treatment period. For the first component, we categorize universities as either ex-ante strong in AI or not, on the basis of the count of faculty members publishing AI-related work at each institution before 2010. For the second component—firm-university hiring networks—we draw on the Revelio resume data to identify which universities the existing employees of each firm in 2010 had attended.

Data Construction.

To identify universities with pre-existing strength in AI research, we rely on the definition of AI-strong universities from Babina et al. (2024), which provides a detailed description of the methodology. In a nutshell, we use data from Open Academic Graph (OAGv2) to quantify AI-related publications linked to each institution. We then link 659 research institutions—representing 99% of entries—in the Higher Education Research and Development Survey (HERDS) to faculty records in OAGv2. The HERDS data, collected by the National Science Foundation, cover all U.S. universities with at least $150,000 in annual R&D expenditures. For each university matched to the HERDS roster, we collect all OAGv2 publications in each year for which at least one co-author lists an affiliation with that institution. While we are interested specifically in AI-related researchers to build our instrument, we also identify non-AI CS-related researchers to use as a control variable since AI- and CS-research strengths might be highly correlated. We classify researchers as AI and CS researchers based on the share of their publications in AI- or CS-related venues, and calculate the annual proportion of researchers classified as AI researchers and the proportion classified as CS researchers at the university level. A university is classified as AI-strong if, in at least one year between 2005 and 2009, it satisfies either of the following conditions: (i) its count of AI researchers ranks in the top 5% across all universities in that year; (ii) its count of AI researchers ranks in the top 10% and the *share* of AI researchers relative to all other researchers in the OAGv2 data ranks in the top 5% across all universities in that year.

To build the second component of our instrument—firm-level exposure to historically AI-strong universities through pre-existing hiring networks—we turn to our Revelio resume data. These records list the degree-granting institutions associated with every credential on workers' resumes. First, we disambiguate university names in the Revelio data and match them to the HERDS roster. We then define individual i as a graduate of

university u if the most recent degree (graduate or undergraduate) reported on individual i's resume is from university u. We classify individual i as a *fresh graduate* of university u in year t if i entered employment at a firm in year t and the job is the first job after graduating from university u. The data offer extensive university coverage: in 2010, for example, 643 out of 659 universities in the HERDS dataset have at least one fresh graduate represented in our resume records. Because firms' recruitment patterns for STEM versus non-STEM workers may differ—a firm might draw economics talent from one university's policy program and management talent from another's business school—we also construct firm-university hiring networks restricted to STEM employees, which may be more relevant for AI recruitment. We define STEM workers as those holding at least one degree in engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., mathematics, physics, chemistry, computer science, statistics), or biological sciences (e.g., biology, pharmacology). In our analysis, we exclude firms for which the fraction of STEM workers with missing university information exceeds 50 percent.

We benchmark the coverage of our university graduate data against official statistics and find that our resume records capture a substantial share of U.S. university graduates. Aggregating to the university-year level, we compute the total number of fresh graduates from each institution in each year and compare these figures to the total degrees (bachelors, masters, and PhDs) conferred annually, as reported in the Integrated Postsecondary Education Data System (IPEDS), which covers all post-secondary institutions in the United States. As of 2012 (the most recent year in the IPEDS data), our resume data account for an average of 46% of all fresh graduates at each university. The cross-sectional correlation between our fresh graduate counts and the IPEDS figures is also high, at 0.67.

Lastly, we use the Revelio resume data to compute the annual share of fresh university graduates from each institution who enter AI-skilled positions, covering the period from 2006 to 2024. Appendix Figure C.1 shows that the share of AI graduates increased from less than 0.5% in the early 2010s to 4% in 2024, suggesting that universities scale up the production of AI-skilled graduates once commercial demand for AI applications expanded after the mid-2010s.

Instrument Validation.

We begin by testing several core assumptions that underpin the logic of our instrument. Consistent with our central argument, we demonstrate that the growth in AI-trained graduates during the 2010s was far more pronounced at AI-strong universities than at other institutions. Appendix Figure C.2 displays the share of fresh graduates entering AI-skilled jobs, plotted separately for AI-strong and non-AI-strong universities over the 2006–2024 period. In 2006, the proportion of AI graduates was negligible across the board, remaining well below 1% for both groups. In 2018, the share grew to 4% at AI-strong

universities, but remained at around 1% in other universities. Between 2018 and 2024, AI-strong universities saw this share further double to about 8%, while non-AI-strong universities grew the share slowly to 3%.

We next examine which universities produce more AI graduates. Appendix Table C.1 and Appendix Table C.2 lists the top 10 universities with the highest share and number of AI graduates in 2018 and 2024. When measured by the share of graduates trained in AI, elite technical institutions dominate: Caltech, Carnegie Mellon, Stanford, and MIT consistently rank at the top in both 2018 and 2024, with shares rising markedly over time (e.g., Carnegie Mellon from 12% to 17%). When measured by absolute numbers, large research universities lead: Georgia Tech tops both years (421 in 2018, 545 in 2024), followed by Berkeley, Carnegie Mellon, and Columbia, with most schools seeing substantial growth in AI graduate counts.

We next assess whether firm-university hiring networks generate the variation needed for our identification strategy. Our instrument relies on cross-firm differences in exposure to AI-strong universities, which requires that firms draw their talent from distinct sets of institutions rather than hiring uniformly from the same pool. Empirically, hiring is highly concentrated: the average firm sources 19% of its fresh graduates from a single dominant university, 47% from its top five universities, and 61% from its top ten. By comparison, the single largest university accounts for only 1.6% of the nationwide pool of fresh graduates, the five largest for 7.1%, and the ten largest for 12.9%. Geographic proximity also plays a meaningful role: on average, firms recruit 36% of all fresh graduates, 34% of STEM fresh graduates, and 32% of AI fresh graduates from universities located in the same state as their headquarters. In addition, for the instrument to function, pre-existing firm-university networks must be persistent enough to predict post-period hiring of AI-skilled workers. In column 1 of Table C.3, we regress the share of fresh graduates recruited from each university after 2018 on the corresponding pre-2010 share, finding a strong positive association that confirms the temporal stability of these networks. Column 2 uses the share of all workers employed at the firm between 2005 and 2010 who graduated from each university as the predictor, with similarly strong results. This persistence extends to AI hiring specifically: columns 3 and 4 show that the universities supplying a firm's overall workforce before 2010 are also strong predictors of the universities from which that firm later recruits AI-skilled employees. Finally, columns 5 and 6 demonstrate that pre-2010 firm-university hiring networks constructed solely from STEM workers also have considerable predictive power for post-2018 AI recruitment.

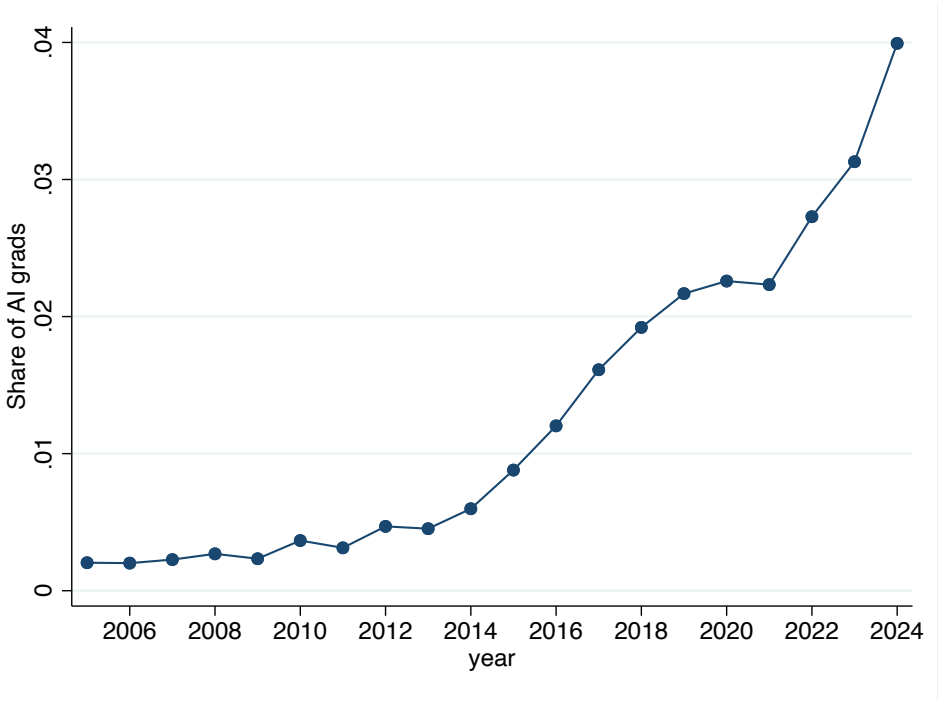
Robustness.

An important identification concern is that if firms anticipated the coming wave of AI demand, they might have proactively strengthened their recruiting ties to AI-strong

universities before 2010, rendering the 2010 hiring networks endogenous to future AI recruitment capacity. This scenario appears unlikely, given the absence of both significant corporate interest in AI and meaningful production of AI-skilled graduates by universities prior to 2010. We also provide direct empirical support: firms connected to AI-strong universities did not disproportionately increase their hiring from those institutions between 2005 and 2010. Appendix Table C.5 confirms that there is no significant association between changes in the share of fresh graduates from AI-strong universities during the pre-period (2005–2010) and our instrument.

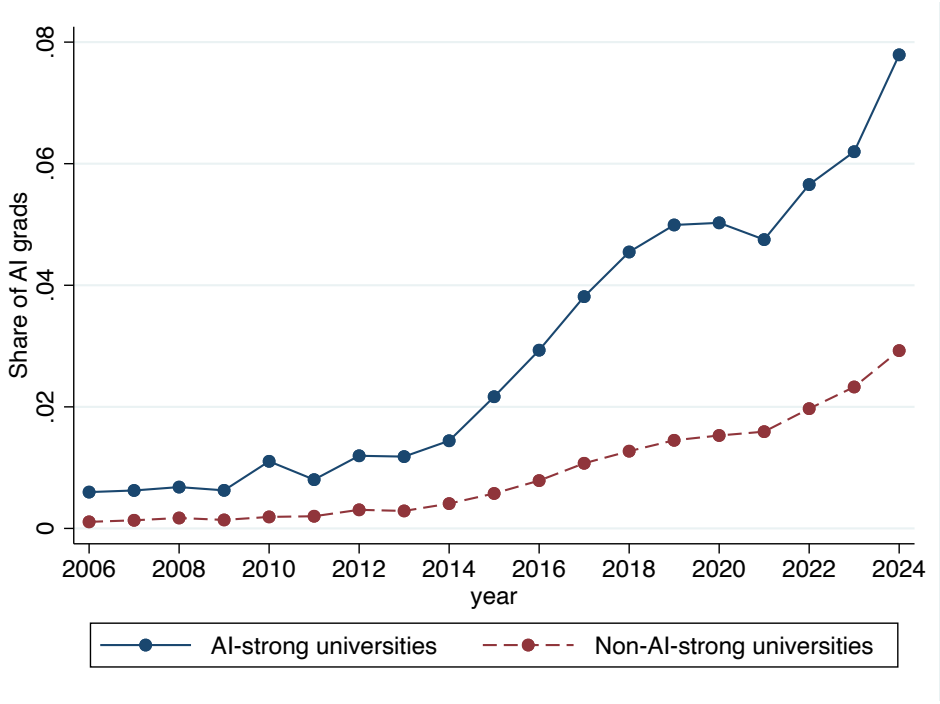
Appendix Table C.6 demonstrates that firms with greater exposure to AI-strong universities did not experience faster growth in productivity prior to 2010, lending support to the exclusion restriction that pre-existing ties to AI-strong universities affect firm performance only through AI investments made after 2010. Panel B further shows that exposure to AI-strong universities is not associated with higher productivity growth over 2010–2018, consistent with the lagged effects of AI investment on productivity.

Figure C.1. Share of AI-trained Fresh Graduates Over Time



This figure plots the share of AI-trained fresh graduates out of all fresh graduates from 2005 to 2024. We define an individual i as a fresh graduate from university u in year t if individual i joined a firm in year t and got the most recent degree from university u and the job is the first job after graduation. An individual is considered an AI-trained fresh graduate in year t if the individual is a fresh graduate in year t and that individual's first job after graduation is an AI-skilled job. AI-skilled jobs are defined based on the methodology described in Section 2.1 and used throughout the paper.

Figure C.2. Time Series of the Share of AI-trained Fresh Graduates from Ex-ante AI-strong Universities and Other Universities



This figure plots the share of AI-trained fresh graduates out of all fresh graduates from 2005 to 2024, separately for ex-ante AI-strong universities and non-AI-strong universities. We define a university as an AI-strong university if it satisfies one of the following two criteria in at least one year between 2005 and 2009: (i) the number of AI researchers is in the top 5% of the distribution across all universities in a given year; (ii) the number of AI researchers is in the top 10% of the distribution, and the share of AI researchers is in the top 5% of the distribution across all universities in a given year. We define an individual i as a fresh graduate from university u in year t if individual i joined a firm in year t and got the most recent degree from university u and the job is the first job after graduation. An individual is considered an AI-trained fresh graduate in year t if the individual is a fresh graduate in year t and that individual's first job after graduation is an AI-skilled job. AI-skilled jobs are defined based on the methodology described in Section 2.1 and used throughout the paper.

Table C.1. Top 10 Universities by Share of AI Graduates

	2018		2024	
Rank	**University**	**Share**	**University**	**Share**
1	California Institute of Technology	0.119	CUNY Graduate Center	0.176
2	Carnegie Mellon University	0.117	Carnegie Mellon University	0.170
3	Stanford University	0.097	Rockefeller University	0.167
4	Massachusetts Institute of Technology	0.097	Stanford University	0.143
5	Harvey Mudd College	0.094	Claremont Graduate University	0.133
6	Georgia Institute of Technology	0.068	Massachusetts Institute of Technology	0.132
7	CUNY Graduate Center	0.067	California Institute of Technology	0.121
8	Princeton University	0.056	Illinois Institute of Technology	0.108
9	Rice University	0.055	Georgia Institute of Technology	0.102
10	Williams College	0.050	Columbia University	0.097

This table lists the top 10 universities with the highest share of AI-trained fresh graduates in 2018 and 2024. An individual is considered an AI-trained fresh graduate in year t if the individual is a fresh graduate in year t and that individual's first job after graduation is an AI-skilled job. AI-skilled jobs are defined based on the methodology described in Section 2.1 and used throughout the paper.

Table C.2. Top 10 Universities by Number of AI Graduates

	2018		2024	
Rank	**University**	**Count**	**University**	**Count**
1	Georgia Institute of Technology	421	Georgia Institute of Technology	545
2	Carnegie Mellon University	353	University of California, Berkeley	468
3	University of California, Berkeley	306	Carnegie Mellon University	430
4	Stanford University	292	Northeastern University	412
5	University of Southern California	264	Columbia University	389
6	Columbia University	250	University of Southern California	370
7	New York University	224	New York University	352
8	University of Illinois at Urbana-Champaign	214	University of Illinois at Urbana-Champaign	349
9	University of California, San Diego	177	University of California, San Diego	344
10	Massachusetts Institute of Technology	167	Stanford University	290

This table lists the top 10 universities with the largest number of AI-trained fresh graduates in 2018 and 2024. An individual is considered an AI-trained fresh graduate in year t if the individual is a fresh graduate in year t and that individual's first job after graduation is an AI-skilled job. AI-skilled jobs are defined based on the methodology described in Section 2.1 and used throughout the paper.

Table C.3. Persistence of Firm-University Hiring Networks

	Share of Post-2018 Hires		Share of Post-2018 AI Hires			
	(1)	(2)	(3)	(4)	(5)	(6)
Share of Pre-2010 Hires	0.583***		0.542***			
	(0.021)		(0.040)			
Share of Pre-2010 Workers		0.636***		0.586***		
		(0.023)		(0.050)		
Share of Pre-2010 STEM Hires					0.172***	
					(0.020)	
Share of Pre-2010 STEM Workers						0.292***
						(0.032)
Firm FE	Y	Y	Y	Y	Y	Y
University FE	Y	Y	Y	Y	Y	Y
Observations	118,897	175,419	108,763	157,584	108,651	157,524

This table reports the coefficients from regressing the share of each firm's fresh graduates hired from each university after 2018 on the pre-2010 firm-university network. Each observation is a firm-university pair. The dependent variable, constructed using the Revelio resume data, is the share of all fresh graduates hired from each university after 2018 in columns 1 and 2 and the share of AI-trained fresh graduates hired from each university after 2018 in columns 3–6. In columns 1 and 3, the independent variable is the share of all fresh graduates hired from each university between 2005 and 2010. We define an individual i as a fresh graduate from university u in year t if individual i joined a firm in year t and got the most recent degree from university u and the job is the first job after graduation. In columns 2 and 4, the independent variable is the share of all workers in the firm between 2005 and 2010 who graduated from each university. In column 5, the independent variable is the share of all STEM fresh graduates hired from each university between 2005 and 2010. We define STEM workers as employees who have at least one degree with a major in either engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., math, physics, chemistry, computer science, statistics), or biological sciences (e.g., biology, pharmacology). In column 6, the independent variable is the share of STEM workers in the firm between 2005 and 2010 who graduated from each university. All columns control for firm fixed effects and university fixed effects. Standard errors are clustered at the university level and reported in parentheses. *, **, and *** denote statistical significance at the 10

Table C.4. First Stage of the Instrument

	Δ Share of AI Jobs		
	(1)	(2)	(3)
Instrument	0.362***	0.498***	0.687***
	(0.114)	(0.118)	(0.134)
Industry FE	Y	Y	Y
CS Control	Y	Y	Y
Baseline Controls	Y	Y	Y
State FE	N	Y	Y
Control Pre-trend	N	N	Y
F Statistic	10.1	17.9	26.2
Observations	633	627	531

This table reports the first stage of the instrument, where we regress our key independent variable—firm-level changes in the share of AI-skilled workers from 2018 to 2024—on the instrument, which measures ex-ante firm-level exposure to the supply of AI-trained university graduates from AI-strong universities. The independent variable is the share of STEM workers in the firm between 2005 and 2010 who graduate from ex-ante AI-strong research universities. The dependent variable is the resume-based measure of the growth in the share of AI jobs from 2018 to 2024. Both the independent variable and the dependent variable are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Revelio jobs in 2018. All specifications control for industry sector fixed effects, ex-ante firm-level exposure to universities that are historically strong in CS research as well as top 20 universities, and the baseline controls measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 2 and 3 add state fixed effects. Column 3 additionally controls for firm-level pre-trends: changes in log sales, log employment, and TFP from 2002 to 2010. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. The first-stage F-statistics of the instrument are reported in both tables for all specifications. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table C.5. Changes in Hiring from Ex-ante AI-strong Universities during the Pre-period (2005-2010)

	Δ Share of Fresh Graduates Hired from AI Hubs 2005-2010		
	(1)	(2)	(3)
Instrument	0.020	0.035	0.016
	(0.029)	(0.038)	(0.036)
Industry FE	Y	Y	Y
CS Control	Y	Y	Y
Baseline Controls	Y	Y	Y
State FE	N	Y	Y
Control Pre-trend	N	N	Y
Adj R-Squared	0.047	0.070	0.164
Observations	470	463	378

This table reports the coefficients from regressing the change in the share of fresh graduates from AI-strong universities from 2005 to 2010 on the instrument (the share of STEM workers in the firm between 2005 and 2010 who graduated from AI-strong universities). The independent variable is standardized to mean zero and standard deviation of one. All columns control for ex-ante exposure to universities that are strong in CS research and top 20 universities, industry sector fixed effects, and the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 2 and 3 add state fixed effects. Column 3 additionally controls for firm-level pre-trends: changes in log sales, log employment, and TFP from 2002 to 2010. Regressions are weighted by the number of Revelio resumes in 2010. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table C.6. Firm Connections to AI-strong Universities and Productivity Growth in Earlier Periods

Panel A: 2002–2010

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Instrument	0.021	-0.016	0.003	-0.021
	(0.032)	(0.030)	(0.034)	(0.036)
Industry FE	Y	Y	Y	Y
CS Control	Y	Y	Y	Y
Baseline Controls	Y	Y	Y	Y
State FE	N	Y	N	Y
Adj R-Squared	0.339	0.409	0.289	0.341
Observations	643	639	550	545

Panel B: 2010–2018

	Δ Log Sales per Worker		Δ TFP	
	(1)	(2)	(3)	(4)
Instrument	-0.011	0.012	-0.034	-0.019
	(0.032)	(0.028)	(0.032)	(0.031)
Industry FE	Y	Y	Y	Y
CS Control	Y	Y	Y	Y
Baseline Controls	Y	Y	Y	Y
State FE	N	Y	N	Y
Adj R-Squared	0.278	0.292	0.243	0.271
Observations	808	806	714	711

This table reports the coefficients from regressions of firm growth in earlier periods on the instrument – ex-ante firm-level exposure to AI-skilled graduates from AI-strong universities. Regressions are weighted by the number of Revelio resumes in 2010. The instrument are standardized to mean zero and standard deviation of one. The dependent variables are changes in firm productivity from 2002 to 2010 in Panel A and from 2010 to 2018 in Panel B. In each panel, we consider changes in log sales per worker in columns 1 and 2, and changes in TFP in columns 3 and 4. All specifications control for industry sector fixed effects and ex-ante exposure to universities that are strong in computer science research as well as top 20 universities. All columns also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, markup, and log number of workers in Revelio), log industry wage, and characteristics of the commuting zones where the firms are located (log average wage, log population, exposure to Chinese import competition, the share of foreign-born workers, the share of routine workers, the share of workers in manufacturing, finance, and IT-related occupations, the share of college-educated workers, and the share of female workers). Columns 2 and 4 additionally control for state fixed effects. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

D. A Model of AI and Organization Capital

This appendix sets out a way of thinking about how AI affects organization capital and productivity growth. We extend the firm life-cycle model of Atkeson and Kehoe (2005) to separate two state variables inside a firm: *idiosyncratic productivity* A, a reduced-form advantage following their stochastic life-cycle process with declining drift, and *organization capital* B, the durable documented asset a firm builds out of running its operations. Measured productivity is $\tilde{A} = AB^{\psi}$. AI is modeled as a codification technology: it raises the rate χ at which operations are converted into organization capital, leaving the idiosyncratic process untouched. Three conclusions follow.

1. **A permanent but gradual level effect.** Adoption raises measured productivity permanently by $\psi \log(\chi'/\chi)$ and permanently shifts the composition of firm assets toward the durable stock. The effect is a level shift, not a change in growth: firms do not grow faster forever, they converge to a higher path.

2. **Slow adjustment, with no frictions required.** The post-adoption gap in organization capital and productivity follows an exact first-order recursion whose convergence rate is the depreciation rate δ_B plus the organization-capital growth of untreated firms. Plausible depreciation rates imply half-lives of five years or more. Responses that take many years to build are therefore not evidence of adjustment costs or implementation lags—they are what a stock–flow system delivers when stocks are large relative to flows.

3. **Heterogeneity runs through the organization capital ratio.** Both the size of the impact response and the speed of convergence depend on the firm's state only through its pre-adoption organization capital ratio $r = B/A$. Firms with a low level of codified knowledge stock respond most and fastest, which we test empirically.

Introduction

Atkeson and Kehoe (2005) propose that the value of a firm resides in its *organization capital*: the built-up, firm-specific expertise that makes an established producer more productive than a new one. In their model this expertise is a pair (A, s)—a specific productivity and an age—evolving stochastically with an age-dependent drift and lost permanently when the firm closes.

Their framework treats the firm's accumulated advantage as a single object. We split it in two. *Idiosyncratic productivity* is the reduced-form state governing how large a firm is relative to its competitors. It is deliberately agnostic about mechanism: it may reflect the tacit knowledge its workforce holds but has not written down, its technical efficiency,

product appeal, or an accumulated customer base—Foster et al. (2016) show that the slow growth of young producers is driven by demand-side accumulation rather than productivity gaps. What these have in common is their dynamics: they build up gradually over the life cycle, are buffeted by idiosyncratic shocks, and erode when not renewed (Benkard, 2000; Argote et al., 1990). *Organization capital* is the durable, documented asset a firm builds out of running its operations: procedures, structured data, trained models, executable workflow, customer records. It is bookable, it is immune to the shocks that buffet period-to-period performance, and it goes stale only as the environment moves on.

Artificial intelligence, on this view, is a technology for converting operations into durable assets. It transcribes, structures, and operationalizes what a firm does, turning activity that would otherwise leave no persistent trace into a stock. We model it as an increase in the conversion rate between the two, and ask what that hypothesis implies. Beraja and Talamàs (2026) share this conceptualization, treating AI as an organizational learning technology whose aggregate value is governed by how much learning firms have left to do. We work at a finer resolution, splitting the accumulated advantage into a tacit component and a codified one and letting AI act only on the conversion rate between them.

The exercise has three products. The first is a set of life-cycle predictions for firm size and asset composition. The second is a characterization of post-adoption dynamics, stated exactly in Proposition 3, which is what lets the model speak to our empirical finding that responses build over many years. The third is a cross-sectional test built on the organization capital ratio, developed in Section D.3 and shown in Section D.4 to govern the private value of adopting as well.

Terminology. Atkeson and Kehoe (2005) use "organization capital" for the pair (A, s), which in our notation is idiosyncratic productivity. We reserve the term for the codified stock B, in line with the measurement literature (Prescott and Visscher, 1980; Eisfeldt and Papanikolaou, 2013a; Corrado et al., 2009), where organization capital denotes the durable, bookable organizational asset rather than idiosyncratic productivity. This is the object empirical proxies target. The relabeling is terminological; no equation of Atkeson and Kehoe (2005) is altered by it.

D.1. Model Setup

D.1.1. Households

Time is discrete, $t = 0, 1, 2, \ldots$. There is a unit continuum of households with preferences $\sum_t \beta^t \log c_t$. Each household contains one worker and one manager, each supplying one unit of labor inelastically. Households own the initial capital stock and the firms existing at date 0, face sequences of worker wages w_t, managerial wages w_{mt}, and intertemporal

prices p_t, and choose consumption subject to the usual present-value budget constraint. This side of the economy is exactly as in Atkeson and Kehoe (2005).

D.1.2. Firms

A firm is characterized by (A, B, s): idiosyncratic productivity, organization capital, and age. Define *effective specific productivity*

$$\tilde{A} = A\, B^{\psi}, \qquad \psi \geq 0, \tag{D.1}$$

and let a firm operating with one manager, capital k, and labor l produce

$$y = z_t\, \tilde{A}^{\,1-n} F(k, l)^n, \qquad n \in (0, 1), \tag{D.2}$$

with F homogeneous of degree one, z_t economywide productivity, and n the span-of-control parameter. Setting $\psi = 0$ recovers Atkeson and Kehoe (2005).

Idiosyncratic productivity. Conditional on operating, $A_{s+1} = A_s e_s$ with $\log e_s \sim N(\mu_s, \sigma^2)$ and the Atkeson-Kehoe declining drift

$$\mu_s = \kappa_1 + \kappa_2 \left[\tfrac{S-s}{S}\right]^2 \mathbf{1}\{s \leq S\}. \tag{D.3}$$

We take A as a reduced form and stay agnostic about its content: it may be technical efficiency, product appeal, an accumulated customer base (Foster et al., 2016), the firm's position relative to rivals, or the tacit, non-codified knowledge held by its workforce—the routines and judgment that have not been written down and that leave with the people who hold them. What (D.3) imposes is only the shape of the resulting life cycle—rapid gains when young, a peak, and erosion at rate $|\kappa_1|$ once mature.

Organization capital. Conditional on operating,

$$B_{s+1} = (1 - \delta_B) B_s + \chi A_s, \qquad B_0 = b_0 \tau_t, \tag{D.4}$$

where $\chi > 0$ is the rate at which operations are converted into codified form, $b_0 > 0$ is the codified content embodied in the entrant's blueprint (process templates, purchased systems), and δ_B is the depreciation rate of organization capital, reflecting obsolescence as documented processes, systems, and records go stale. A firm codifies out of what it does, so the flow into B is proportional to the scale of operations. Once codified, the asset is not exposed to the innovation shocks: a bad quarter does not delete the process documentation, and a departing employee does not take the customer database.

The separation of the two stocks is motivated by a long tradition treating the binding constraint on organizational growth as the difficulty of transferring knowledge that resides in individuals and resists articulation (Polanyi, 1966; Nonaka, 1994). The constraint binds hardest early in the life cycle, before routines have stabilized and while learning depends on sparse and noisy feedback rather than on accumulated experience and documented practice (March, 1991; Grant, 1996). Knowledge held this way is also fragile: it erodes when not exercised (Argote et al., 1990) and degrades as practice is transmitted informally rather than written down. This is one reason organization capital accumulates as slowly as it does (Atkeson and Kehoe, 2005), and it is what χ represents in (D.4): the rate at which a firm's operations leave a durable trace rather than remaining with the people who generated them.

Relation to other forms of intangible capital. Organization capital as modeled here shares the defining features of intangible capital more generally. Crouzet et al. (2022) emphasize that intangible capital must be stored on some medium, and that where it is stored governs who can use it and who can appropriate its returns. That observation is the organizing principle of our two-stock structure: idiosyncratic productivity is stored in people, arrangements, and relationships and is lost when they change, while organization capital is stored on a durable medium and persists independently of them. The parameter χ is the rate of transfer between the two media, and Assumption 1 is the claim that AI acts on that rate. Second, like all capital, the stock accumulates and depreciates, though δ_B here represents obsolescence of content rather than physical decay—documentation and records go stale because the environment they describe has moved on. Third, and more specifically, organization capital in (D.4) is a *by-product of production*: the flow χA_s is proportional to the scale of operations rather than freely chosen. This is the property that data has in Veldkamp and Chung (2024), where data is generated by economic activity rather than produced for its own sake, and it distinguishes organization capital sharply from R&D capital, which is a deliberate investment flow. A firm cannot buy its way to a large codified stock; it can only operate and record. Fourth, the stock is non-rival within the firm in the sense of Jones and Tonetti (2020): a documented procedure or a structured record can be used by every employee at once without depletion.[57]

Two features set organization capital apart, and both follow from firm specificity. First, it is not shared across firms. Data can be sold and patents licensed or disclosed, and both raise questions about the efficient allocation of a non-rival good across users; organization

[57]Non-rivalry generates increasing returns: scaling (A, B, k, l) by λ scales output by $\lambda^{1+\psi(1-n)}$. This is where (D.2) departs from the workhorse formulations in Crouzet and Eberly (2019) and Crouzet and Eberly (2021), in which intangibles enter as an additional rival input alongside physical capital; here B shifts $\tilde{A}$ outside $F(k, l)$, which is what converts non-rivalry into the scale economies above.

capital is non-rival within the firm but has little value outside it. Second, firm specificity also tends to make organization capital relatively excludable, which is worth noting since limited excludability is often treated as characteristic of intangibles (Crouzet et al., 2022). Patents rely on legal protection and trade secrets on secrecy, whereas a firm's documented routines are of limited use elsewhere, so specificity does some of the work that legal or contractual protection does for other intangibles. Two further consequences follow. Because it is neither traded nor separable from the firm, organization capital has no market price, which is why it must be inferred from firm behavior rather than read off transactions. And because it cannot be reallocated, it is destroyed at exit rather than sold, unlike physical capital or patents, which survive their owners.

Firm specificity. Both stocks are firm-specific. Nothing spills over to entrants or rivals, and exit destroys (A, B) jointly, preserving the Atkeson-Kehoe assumption that the accumulated advantage of a closing producer "is lost permanently."

Define the *organization capital ratio* $r_s \equiv B_s / A_s$, with $r_{s+1} = [(1 - \delta_B) r_s + \chi]/e_s$. Growth in effective productivity decomposes as

$$\Delta \log \tilde{A}_s = \underbrace{\Delta \log A_s}_{\text{idiosyncratic, } \mathbb{E}=\mu_s} + \psi \underbrace{\log\big[(1 - \delta_B) + \chi / r_s\big]}_{\text{organization capital, } \equiv \gamma_s}. \tag{D.5}$$

The term γ_s refers to the growth rate of organization capital.

Capital and labor are freely mobile each period, so the input choice is static:

$$\max_{k,l} z_t \tilde{A}^{1-n} F(k, l)^n - r_t k - w_t l - w_{mt}.$$

Because F is homogeneous of degree one and there is a single fixed factor, all operating firms use inputs in the same proportions.

D.1.3. Entry, exit, and firm value

The decision to operate is dynamic. Firm value satisfies

$$V_t(A, B, s) = \max\Big\{0,\ d_t(\tilde{A}) - w_{mt} + \tfrac{p_{t+1}}{p_t} \mathbb{E}_s V_{t+1}\big(Ae, (1 - \delta_B)B + \chi A, s + 1\big)\Big\}. \tag{D.6}$$

The maximization in (D.6) is over the firm's only dynamic choice: the binary decision to operate this period or to exit. The first argument is the value of exit—the firm produces nothing and both stocks are destroyed permanently. The second is the value of operating: current variable profits $d_t(\tilde{A})$, which equal revenue net of payments to capital and labor at the static optimum, minus the fixed cost of hiring the manager, plus the discounted expected continuation value, where the expectation is over the innovation e to idiosyncratic

productivity drawn from its age-s distribution.

Entry. Entry works as in Atkeson and Kehoe (2005). Before a new firm can enter in period t, a manager must spend period $t-1$ preparing and adopting a blueprint for constructing the firm, and the blueprint determines the firm's initial state. Blueprints adopted in period $t-1$ embody the frontier of knowledge regarding the design of firms at that point in time; this frontier τ_t evolves exogenously, growing at rate g_τ. Thus a firm built in $t-1$ starts period t with idiosyncratic productivity τ_t and organization capital $b_0\tau_t$, so its state is $(A, B, s) = (\tau_t,\, b_0\tau_t,\, 0)$: the blueprint embodies not only an initial market position but also a starter stock of codified material—process templates, purchased software and systems—scaled by the same frontier. Free entry into the activity of preparing blueprints drives the net value of entry to zero, which pins down the managerial wage:

$$w_{mt} = \tfrac{1}{1+i}\, V_{t+1}\big(\tau_{t+1},\, b_0\tau_{t+1},\, 0\big), \tag{D.7}$$

where i is the equilibrium interest rate.

Managers. Managers are the economy's fixed factor and are used in two activities. Let f_t denote the measure of managers spending period t preparing blueprints for firms that will enter in $t+1$, and let M_t denote the measure of firms operating in t, each of which employs exactly one manager as its fixed factor. Since each household supplies one manager, the resource constraint for managers is $f_t + M_t = 1$.

As in Atkeson and Kehoe (2005), two features of (D.6) matter below. V is convex in A—the flow is linear in $\tilde{A}$ and the continuation is truncated at zero—and young firms typically have $d < w_m$: organization rents are back-loaded and young firms operate on option value.

D.1.4. Balanced growth and stationary composition

Following Atkeson and Kehoe (2005), we define a balanced growth path as an equilibrium in which the quality of the best available blueprint τ_t and economywide productivity z_t grow at constant rates g_τ and g_z; aggregate output, consumption, capital, and the two wages grow at a common constant rate; the interest rate, the entry measure f_t, and the measure of operating firms M_t are constant; and the distribution of firm states is stationary once each cohort is detrended by its blueprint—$(A_s/\tau, B_s/\tau)$ has the same law for every cohort, so each generation of firms traverses the same life cycle, scaled up by the frontier. Along such a path effective productivity across cohorts grows at rate $(1+g_\tau)^{1+\psi} - 1$, since both stocks scale with the cohort blueprint, and the aggregate growth rate follows from the Atkeson-Kehoe formula with this embodied rate.

Stationarity of the firm's asset *composition* requires a restriction linking the two decay rates.

Proposition 1 (Stationary composition)**.** *Consider a mature firm drawing from the terminal distribution with drift κ_1. Along the median path the organization capital ratio converges to*

$$r^* = \frac{\chi}{e^{\kappa_1} - (1 - \delta_B)}$$

if and only if $\delta_B > 1 - e^{\kappa_1}$, and the fixed point is globally stable on $\mathbb{R}_+$. If $\delta_B < 1 - e^{\kappa_1}$, then $r_s \to \infty$ from any initial condition.

Proof. On the median path $e_s = e^{\kappa_1}$, so the organization capital ratio follows the affine map

$$r_{s+1} = \phi(r_s) \equiv a\, r_s + c, \qquad a = (1 - \delta_B)e^{-\kappa_1}, \quad c = \chi e^{-\kappa_1} > 0.$$

Iterating, $r_s = a^s r_0 + c\,(1 - a^s)/(1 - a)$ for $a \neq 1$. If $a < 1$—that is, $1 - \delta_B < e^{\kappa_1}$, or $\delta_B > 1 - e^{\kappa_1}$—then $a^s \to 0$ and $r_s \to c/(1 - a) = \chi/(e^{\kappa_1} - (1 - \delta_B)) = r^* > 0$ monotonically from any $r_0 \geq 0$, and since $|\phi'| = a < 1$ the fixed point is a global contraction. If $a \geq 1$, then $a^s r_0$ is nondecreasing and the geometric sum $c(1 + a + \cdots + a^{s-1})$ diverges, so $r_s \to \infty$. □

The economics of the condition is that organization capital must go stale at least as fast as the firm's idiosyncratic productivity advantage erodes; otherwise the documentation outlives the business it documents. At the Atkeson-Kehoe drift the requirement is $\delta_B > 1 - e^{\kappa_1} = 10.8\%$ per year. Direct evidence suggests this is comfortably satisfied: Benkard (2000) finds that production experience in aircraft manufacturing depreciates substantially over time, Argote et al. (1990) estimate very rapid depreciation of the accumulated experience stock in shipbuilding, and Foster et al. (2016) estimate an annual depreciation rate of roughly 36% for the accumulated demand stock. Abis and Veldkamp (2024) estimate a monthly depreciation rate of 3% for data, implying an annual rate of 31%. Accumulated firm advantages go stale quickly; the stationarity condition is therefore a description of the data, not a strained requirement.

D.1.5. Calibration and the baseline life cycle

Macro parameters are from Atkeson and Kehoe (2005): $\beta = .98$, $n = .85$, physical capital share 19.9%, $i = 4.1\%$, $g = 2.02\%$. The parameters of (D.3) are their estimates from the 1988 LRD panel (Davis et al., 1996). The depreciation rate δ_B is set to 0.15, the value conventionally used in the organization capital literature (Eisfeldt and Papanikolaou, 2013b). The codification-side parameters are inferred based on our estimates and discussed in Section D.5.

Table D.1. Parameters used in the numerical illustrations

	Parameter	Value	Source
Idiosyncratic productivity	κ_1	-0.1139	Atkeson and Kehoe (2005)
	κ_2	0.1741	Atkeson and Kehoe (2005)
	σ	0.1945	Atkeson and Kehoe (2005)
	S	100	Atkeson and Kehoe (2005)
Organization capital	ψ	0.50	based on reduced-form estimates (Section D.5)
	χ	0.15	illustrative
	δ_B	0.15	$> 1 - e^{\kappa_1} = 0.108$; cf. depreciation estimates
	b_0	0.5	entrants start half-codified
AI treatment intensity	χ'/χ	2.0	based on reduced-form estimates (Section D.5)
Adoption values	ρ	0.85	discount $\times$ survival

All figures plot median paths (e_s at its median), isolating life-cycle profiles from the dispersion generated by σ. Figure D.1 shows the baseline. Idiosyncratic productivity peaks at 0.59 log points around age 20 as the drift crosses zero; effective productivity—firm size—peaks at 1.18 log points at age 22, growing at 13.0% per year at entry (6.0 points from idiosyncratic productivity plus $\psi\gamma_0 = 7.0$ points from codification). Organization capital is a lagged, smoothed integral of idiosyncratic productivity: its growth rate γ_s starts at 14.0% per year, declines monotonically, and turns negative at age 26—exactly the age at which the organization capital ratio crosses χ/δ_B—before converging to κ_1 as the composition stabilizes at $r^* = 3.54$.

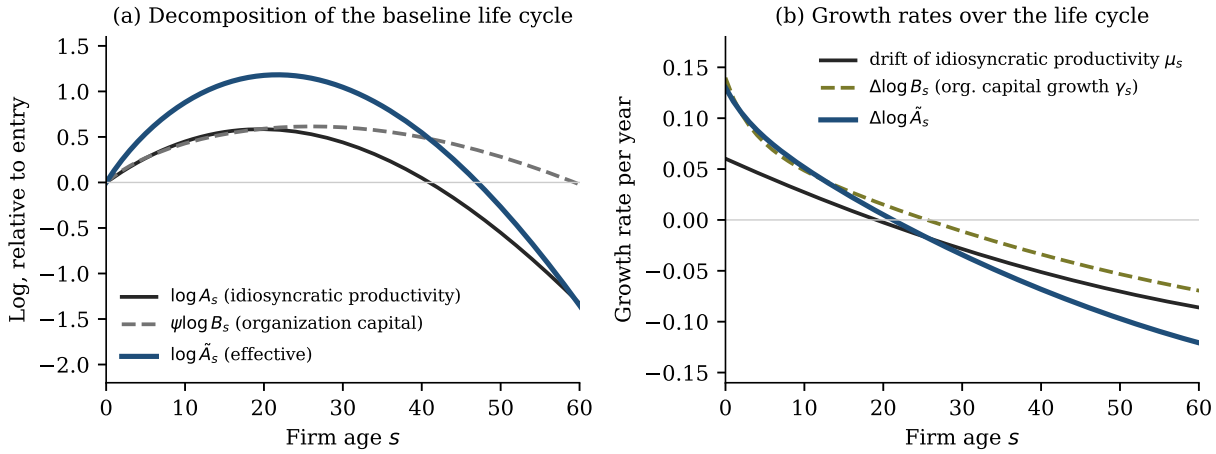


Figure D.1. The baseline (pre-AI) firm life cycle. Panel (a): decomposition of effective productivity into idiosyncratic productivity and the contribution of organization capital. Panel (b): the drift of idiosyncratic productivity μ_s, organization-capital growth γ_s, and effective productivity growth. Median paths at the parameters of Table D.1.

D.2. AI as a Codification Technology

Assumption 1 (Codification)**.** *AI raises the conversion rate: $\chi \mapsto \chi' > \chi$ from the date of adoption onward. The idiosyncratic productivity process $\{\mu_s, \sigma\}$, the depreciation rate δ_B, and the blueprint content b_0 are unchanged.*

Modeling AI as an increase in χ reflects what these systems do when trained on a firm's own records of work. The raw material is the record the firm's operations generate: the accumulated output of individual employees—code repositories, service transcripts, clinical documentation—but also the machine-generated trace of the operation itself, such as line speeds, cycle times, and quality readings, all of which previously served only to document what had been done. Training a model on it converts that record into an operational asset, and does so far faster than the alternative: what would otherwise diffuse slowly through mentoring, documentation, and turnover is extracted in a single pass, so the problem-solving approaches and conventions embedded in the work of particular employees become available to everyone else in the organization almost at once (Brynjolfsson et al., 2025b; Singhal et al., 2023).[58]

This is precisely the transformation (D.4) describes, and Appendix A.4.4 documents it across eight firms: the flow into organization capital is proportional to the scale of operations because operations are what the model is trained on, and χ governs how much of that activity is converted into a form the firm retains. Cullen et al. (2026) make the point explicit—expertise that had been an attribute of individual workers becomes, once embedded in a model trained on their output, an asset the firm owns and can deploy independently of them. Where the record is machine-generated rather than human, the conversion is the same but the constraint it relaxes is different: a firm cannot run experiments on a production line without stopping it, and a model of the line, built from the line's own history, removes that limit. Another complementary channel is that AI extracts more from a given operational record than earlier methods could (Agrawal et al., 2022), so a firm need not wait to accumulate decades of history before its records support a usable asset. Moderna, discussed in Appendix A.4.4, illustrates the point: a firm without a commercial product until 2020, it built predictive models from thousands of development experiment data points that let its process engineers avoid failures they had not yet encountered.

Proposition 2 (Properties of codification)**.** *Under Assumption 1: (i)* Permanent level effect. *B_s is affine in χ with a fixed anchor $b_0\tau$, and effective productivity is permanently higher by*

[58]For example, Brynjolfsson et al. (2025b) show that generative AI systems capture precisely the behaviors hardest to write into explicit procedure—the judgment and communication patterns of a firm's strongest performers—and diffuse them through real-time guidance. That is a change in how much of a firm's existing capability becomes durable and shareable, not a change in the capability itself.

$\psi \log(\chi'/\chi)$. *(ii)* Composition shift. *The stationary organization capital ratio rises to* $(\chi'/\chi)\, r^*$*:* $\log r_s$ *is permanently higher by* $\log(\chi'/\chi)$. *(iii)* Idiosyncratic productivity process unchanged. *The drift profile and innovation variance of A are unaffected.*

Proof. Solving (D.4) forward,

$$B_s = \underbrace{(1-\delta_B)^s\, b_0\tau}_{\text{depreciated endowment}} + \chi \underbrace{\sum_{j=0}^{s-1}(1-\delta_B)^{s-1-j}A_j}_{\text{depreciated sum of past operations}}, \tag{D.8}$$

so organization capital is the sum of two pieces. The first is what remains of the codified endowment the firm entered with; it does not involve χ at all. The second is the depreciated sum of everything the firm has converted since birth; it is exactly proportional to χ. "Affine with a fixed anchor" means just this: doubling the conversion rate doubles the accumulated piece while leaving the endowment piece untouched. Early in life the endowment still matters, so the effect of a higher χ is partial; as the endowment depreciates away, organization capital becomes proportional to χ.

(i) Equation (D.8) holds for any shock path. An adopter and its non-adopting counterfactual share the same path of A (by part (iii)) and the same anchor term, so

$$\frac{B_s^1}{B_s^0} = \frac{(1-\delta_B)^s b_0\tau + \chi'\Sigma_s}{(1-\delta_B)^s b_0\tau + \chi\,\Sigma_s} \longrightarrow \frac{\chi'}{\chi} \quad \text{as } s \to \infty,$$

where Σ_s is the depreciated sum in (D.8), since the anchor terms vanish relative to Σ_s. The log gap in organization capital converges to $\log(\chi'/\chi)$, and by (D.1) the log gap in effective productivity converges to $\psi \log(\chi'/\chi)$. (ii) Immediate from Proposition 1: r^* is linear in χ. (iii) By inspection: the law of motion (D.3) of A involves neither χ nor B. □

Figure D.2 illustrates. Adoption at entry raises growth at birth from 13.0% to 24.6% per year, raises the peak from 1.18 to 1.53 log points at the same age, and shifts the entire mature path up by the permanent $\psi \log(\chi'/\chi) = 0.347$: the gap reaches 0.325 by age 10 and is essentially complete by age 20. Under an illustrative exit threshold the lifespan extends from 53 to 56 years; more generally, because B is immune to the innovation shocks, adopters carry a larger buffer against bad draws, mature exit falls, and exits follow prolonged decline rather than single shocks. Panel (b) shows the composition ratio converging to a permanently higher plateau, $\log r^*$ shifting up by exactly $\log(\chi'/\chi)$.

Two further predictions follow directly from (D.5). *Persistence:* post-adoption, a larger share of effective productivity is the smoothed integral of past idiosyncratic productivity, so a shock to A passes through with long-run elasticity $1+\psi$ at a rate governed by δ_B, and the autocorrelation of size growth rises. *Selection into adoption:* conditional on age and size,

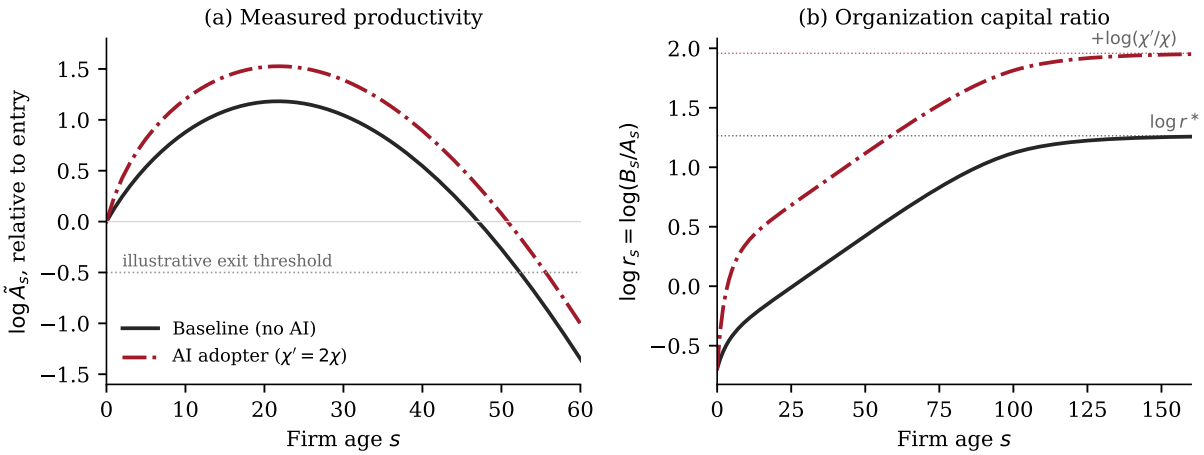


Figure D.2. Life-cycle paths under no AI and under adoption at entry. Panel (a): measured productivity $\log \tilde{A}$; the dotted line is an illustrative exit threshold, crossed at ages 53 and 56. Panel (b): the organization capital ratio $\log r_s$, converging to $\log r^*$ in the baseline and to $\log r^* + \log(\chi'/\chi)$ under AI.

the firms with the most to gain are those with low r—large operations that have codified little of what they do—a point developed in Sections D.3 and D.4.

D.3. Event-Study Dynamics

The bridge from the model to panel evidence is that the post-adoption gap in organization capital obeys an exact first-order recursion, with coefficients built from observables.

Proposition 3 (Event-study dynamics)**.** *Consider a firm adopting at age s_0 and its never-adopting counterfactual with the same shock path. Let B^1, B^0 denote the two organization-capital paths, $x_h \equiv (B^1_{s_0+h} - B^0_{s_0+h})/B^0_{s_0+h}$ the proportional gap, and $G_h \equiv B^0_{s_0+h+1}/B^0_{s_0+h} = e^{\gamma_{s_0+h}}$ the counterfactual growth factor of organization capital. Then for all $h \geq 0$:*

(i) Exact AR(1). $x_0 = 0$ *and*

$$x_{h+1} = \rho_h\, x_h + (1-\rho_h)\,\frac{\chi' - \chi}{\chi}, \qquad \rho_h = \frac{1-\delta_B}{G_h} = (1-\delta_B)e^{-\gamma_{s_0+h}}.$$

(ii) Common plateau. $x_h \to (\chi' - \chi)/\chi$ *regardless of s_0 and of the realized path of A: the log gap in organization capital converges to $\log(\chi'/\chi)$ for every adoption age.*

(iii) Proportionality at every horizon. *The measured-productivity gap is exactly ψ times the organization-capital gap: $\log \tilde{A}^1_{s_0+h} - \log \tilde{A}^0_{s_0+h} = \psi \log(1 + x_h)$ for all h.*

(iv) Impact response. *Setting $h = 0$, the first-period response of organization-capital growth is*

$$\Delta\gamma(s_0) = \log \frac{(1-\delta_B) + \chi'/r_{s_0}}{(1-\delta_B) + \chi/r_{s_0}},$$

strictly decreasing in the pre-adoption organization capital ratio r_{s_0}.

Proof. Let $D_h = B^1 - B^0$ at horizon h. Subtracting the two accumulation equations, $D_{h+1} = (1-\delta_B)D_h + (\chi' - \chi)A_{s_0+h}$, with $D_0 = 0$. Dividing by $B^0_{s_0+h+1} = G_h B^0_{s_0+h} = (1-\delta_B)B^0_{s_0+h} + \chi A_{s_0+h}$ and using $\chi A_{s_0+h}/B^0_{s_0+h} = G_h - (1-\delta_B)$ gives (i). Since $\rho_h \in (0,1)$, x_h converges to the fixed point $(\chi' - \chi)/\chi$, which does not involve the A path, giving (ii). Part (iii) is immediate from $\log \tilde{A} = \log A + \psi \log B$ and identical A paths under Assumption 1(iii). For (iv), $x_1 = (1-\rho_0)(\chi' - \chi)/\chi$ with $\rho_0 = (1-\delta_B)/[(1-\delta_B) + \chi/r_{s_0}]$; the first-period growth-rate difference is $\log(B^1_{s_0+1}/B^1_{s_0}) - \log(B^0_{s_0+1}/B^0_{s_0}) = \log(1+x_1)$, and substituting yields $1 + x_1 = [(1-\delta_B) + \chi'/r_{s_0}]/[(1-\delta_B) + \chi/r_{s_0}]$. □

Figure D.3 displays the gap dynamics for adoption at ages 5, 30, and 60: a common plateau reached at different speeds, with the productivity panel exactly ψ times the organization-capital panel at every horizon. Idiosyncratic productivity A is *identical* for adopters and counterfactuals by Assumption 1, so all the divergence in both panels is due to organization capital.

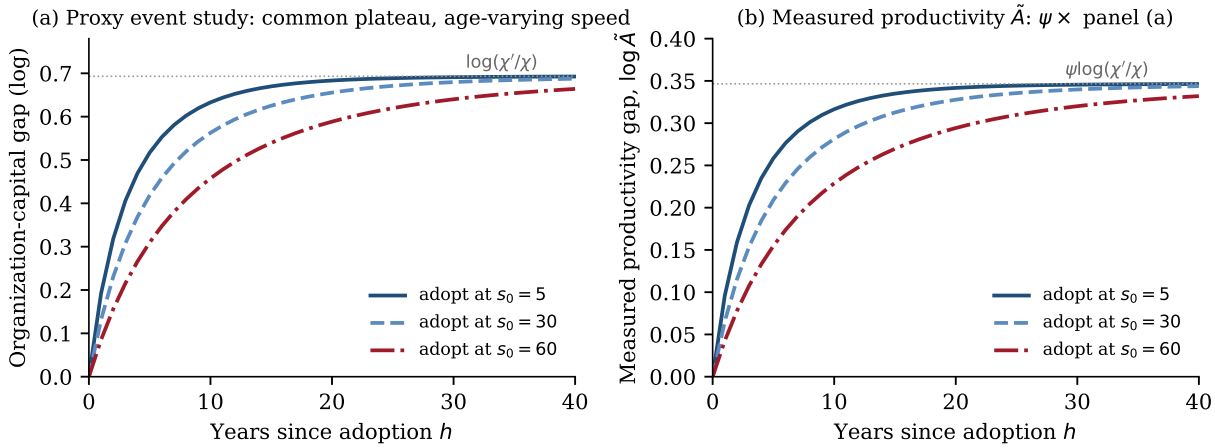


Figure D.3. Event studies around AI adoption at ages 5, 30, and 60, relative to never-adopting firms of the same age. Panel (a): the organization-capital (proxy) gap converges to the common plateau $\log(\chi'/\chi)$ at speeds $\delta_B + \gamma_s$ that fall with the pre-adoption organization capital ratio. Panel (b): the measured-productivity gap $\log \tilde{A}$ is exactly ψ times panel (a) at every horizon.

D.3.1. Why adjustment is slow

The convergence speed in Proposition 3(i) is $1 - \rho_h = 1 - (1 - \delta_B)e^{-\gamma_s}$, approximately $\delta_B + \gamma_s$: the depreciation rate of organization capital *plus the growth rate of the untreated stock*. The intuition is a savings account. Each year the firm deposits χA and the balance depreciates at δ_B; AI raises the deposit rate. How fast the balance reprices to the new deposit rate depends on how large deposits are relative to the accumulated balance—and that ratio is $\gamma_s + \delta_B$. A young firm's stock is mostly recent deposits, so it reprices in a few years. A mature firm's stock is mostly legacy accumulation, and for such a firm γ_s is near zero or negative, so repricing is slow: at the calibration, half-lives are 3, 4, and 6 years for adoption at ages 5, 30, and 60, and the mature limit ($\gamma_s \to \kappa_1$) is 14.6 years.

The magnitudes matter for how we read the evidence. In our empirical work, organization capital and measured productivity take many years to respond fully to AI investment. In this model that is not evidence of adjustment costs, implementation lags, or slow organizational learning—none of which are present. It is what the accounting of a stock–flow system delivers for mature firms: for any firm whose organization capital is stable or declining ($\gamma_s \leq 0$), every depreciation rate below roughly 13–22% implies a half-life above five years, and the depreciation estimates cited in Section D.1 sit squarely in that range.

D.3.2. Heterogeneity test

Proposition 3 has a sharp cross-sectional implication. Both the impact response (part iv) and the convergence speed (part i, through γ_s) depend on the adopting firm's state only through its pre-adoption organization capital ratio r_{s_0}. Two firms with the same r respond identically *regardless of age*; the plateau, meanwhile, is common to all. It follows that at any finite horizon the treatment effect—on organization capital and, by the corollary, on measured productivity—is *decreasing in the initial organization capital ratio*: low-r firms respond more on impact and converge faster toward the same eventual gain. The reason r_s is a sufficient statistic for the response is visible in the definition of γ_s itself: from (D.5) the untreated growth rate of organization capital is $\gamma_s = \log\big[(1 - \delta_B) + \chi / r_s\big]$, so both objects that govern the response—the impact effect $\Delta\gamma(s) = \log \frac{(1-\delta_B)+\chi'/r_s}{(1-\delta_B)+\chi/r_s}$ and the convergence factor $\rho_h = (1 - \delta_B)e^{-\gamma_s} = (1 - \delta_B)\big/\big[(1 - \delta_B) + \chi/r_s\big]$—depend on the firm's state through r_s and nothing else.

This is the heterogeneity test we take to the data: interact AI investment with the pre-adoption ratio of the organization-capital proxy to firm size, and predict larger medium-run responses for firms with low initial ratios, with the heterogeneity shrinking at long horizons as plateaus equalize.

Why not simply test for larger effects among young firms, since r_s rises with age in

the model? Three reasons, one practical and two conceptual. The practical one: firm age is essentially unobserved in Compustat—incorporation and listing dates measure age at listing, not the age of the underlying business, and are especially unreliable for firms formed through mergers and spinoffs.The first conceptual reason is sufficiency: age matters only through the codification ratio, and idiosyncratic shocks disperse that ratio within a cohort, so age is a noisy proxy for the state variable that actually governs the response—while the ratio itself is directly observable as organization capital relative to size. The second conceptual reason is confounding through the treatment intensity. Outside the strict model, χ'/χ need not be a constant: the productivity of AI plausibly *rises* with a firm's accumulated data and organizational experience, since models trained on richer operational histories are better models. If so, age-based heterogeneity confounds two opposing forces—older firms have lower capture speed (smaller effect) but possibly higher treatment intensity (larger effect)—with an ambiguous net sign, and conditioning on r isolates the capture mechanism.[59]

D.4. Endogenous Adoption over the Life Cycle

So far adoption has been exogenous. This section makes it a choice. A firm of age s_0 may pay a one-time irreversible cost κ, in units of entry-size output, to raise its conversion rate permanently from χ to χ'. Writing $\tilde{A}^1$ and $\tilde{A}^0$ for the adopting and non-adopting paths from s_0 onward and ρ for the product of the discount factor and the survival rate, the gross gain from adopting at s_0 is

$$G(s_0) \;=\; \sum_{h\geq 0} \rho^h \big[\tilde{A}^1_{s_0+h} - \tilde{A}^0_{s_0+h}\big], \tag{D.9}$$

which is proportional to the discounted gain in output and variable profits. A firm facing a one-time cost κ adopts at the age maximizing $\rho^{s_0}[G(s_0) - \kappa]$; we focus here on the shape of G, which is what governs both the adoption margin and the response heterogeneity.

Proposition 4 (Value of adopting)**.** *To first order in the treatment intensity,*

$$G(s_0) \;=\; \underbrace{\psi}_{\textit{elasticity}} \times \underbrace{\bar{x}}_{\textit{intensity}} \times \underbrace{W(s_0)}_{\textit{remaining scale}} \times \underbrace{\Lambda(s_0)}_{\textit{capture fraction}} ,$$

[59]This confound signs the bias in our favor: if intensity rises with the accumulated stock, then low-r firms receive *weaker* treatments, which works against finding larger responses for them. The organization capital-ratio test is therefore conservative with respect to the leading violation of the constant-intensity assumption.

where $\bar{x} = (\chi' - \chi)/\chi$, $W(s_0) = \sum_{h \geq 0} \rho^h \tilde{A}^0_{s_0+h}$ *is the firm's discounted remaining size, and*

$$\Lambda(s_0) = \frac{\sum_h \rho^h \tilde{A}^0_{s_0+h} \, x_h(s_0)}{\bar{x} \sum_h \rho^h \tilde{A}^0_{s_0+h}} \in (0,1)$$

is the size-weighted fraction of the eventual plateau that the firm captures. Λ *is decreasing in the pre-adoption organization capital ratio* r_{s_0} *and hence, along the life cycle, decreasing in age;* W *is hump-shaped. The product is hump-shaped in* s_0*.*

The decomposition separates two economically distinct forces, plotted in Figure D.4. *Remaining scale* rises through the first two decades, because a young firm's discounted future size is dominated by a size peak it has not yet reached, and falls thereafter; at the calibration it peaks at age 17. *Capture fraction* falls monotonically—from 0.59 at entry to 0.45 at age 12 and 0.27 at age 40—because a firm that has already codified much of its operation converges to the plateau slowly, by the same ρ_h of Proposition 3 that governs Section D.3.1: what makes mature adopters' event studies slow makes their adoption less valuable. The gain peaks at age 12, before both the scale peak at 17 and the firm's size peak at 22. The hump is, if anything, conservative about how late in life adoption is worth most.

Note that the state variable doing the work is again the organization capital ratio, not age: at age 15, a firm holding half the modal organization capital ($r = 0.41$) has a gain of 4.65 entry-size units against 4.30 for the modal firm and 3.85 for one holding twice the modal stock—a 21% spread across a factor of four in organization capital intensity, entirely through Λ. The adoption margin and the treatment-response heterogeneity of Section D.3.2 are the same economics: AI is worth most, and works fastest, in firms that are large relative to what they have codified. One might have expected the return to codification to be largest for the largest firms, since the flow gain scales with size; it is not, because the largest firms are the oldest and have already converted the bulk of what they do into durable form.

This age profile is visible in the data. Figure D.5 plots the share of a firm's organization-capital-related AI jobs against firm age, averaged within age bins across the 2024 cross-section. The pattern is hump-shaped in the way the model implies. Very young firms invest relatively little—about 0.5% of the workforce at ages zero to four—despite having the most left to codify. Investment then rises steeply, roughly two and a half times, to peak near 1.2% between ages five and thirteen, and declines steadily thereafter, falling to around 0.3% among firms over seventy. Two forces of Proposition 4 map onto the two arms of the hump: the rise reflects growth in remaining scale, as young firms have not yet reached the size over which the return to codification is earned, while the decline reflects the falling capture fraction, as older firms have already converted much of what they do

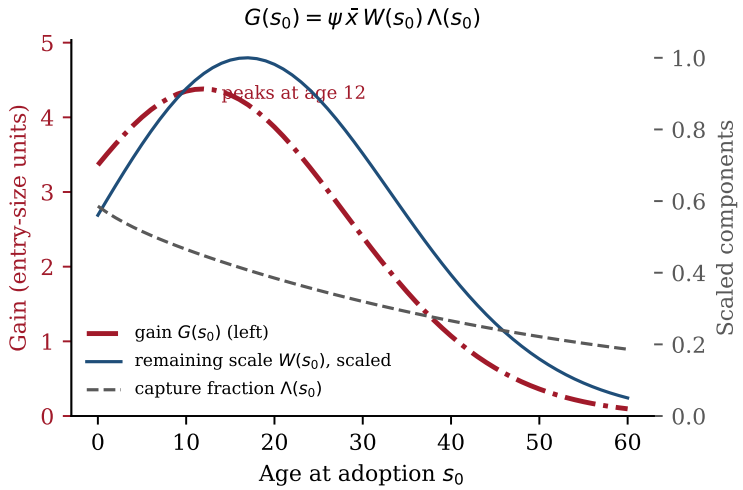


Figure D.4. The gross gain from adopting at age s_0 (left axis) and its two components from Proposition 4: discounted remaining scale $W(s_0)$ and capture fraction $\Lambda(s_0)$ (right axis, scaled).

into durable form and their existing stock dilutes what they add.

D.5. Taking the Model to Data

We conclude with the intuition for how the model organizes panel evidence on AI investment, organization capital, and productivity. A full estimation is beyond this paper; here we record only how the reduced-form objects map into the model's parameters.

Empirically, we measure organization capital by capitalizing SG&A, and read the resulting stock as a constant multiple of B; the unknown scale never matters, because every relevant statistic is a growth rate, a gap, or a ratio in which it differences out. The obvious concern with this proxy is that SG&A bundles spending that builds the codified stock (systems, documentation, data infrastructure) with spending that builds idiosyncratic productivity (tacit knowledge, customer acquisition), so the constructed measure is better read as B contaminated by some component of A. This matters less than it might appear, because our estimates come from *changes* around adoption rather than from levels. Under Assumption 1 AI leaves A untouched, so the contaminating component does not respond to adoption: it shifts the level of the proxy without contributing to the event-study response, and the response still traces the organization-capital gap g_h that Proposition 3 describes.[60]

[60] Because the contamination inflates the measured stock without inflating its response, it attenuates the estimated growth rates and the implied plateau, so $\widehat{\chi'/\chi}$ is biased toward zero—a conservative direction.

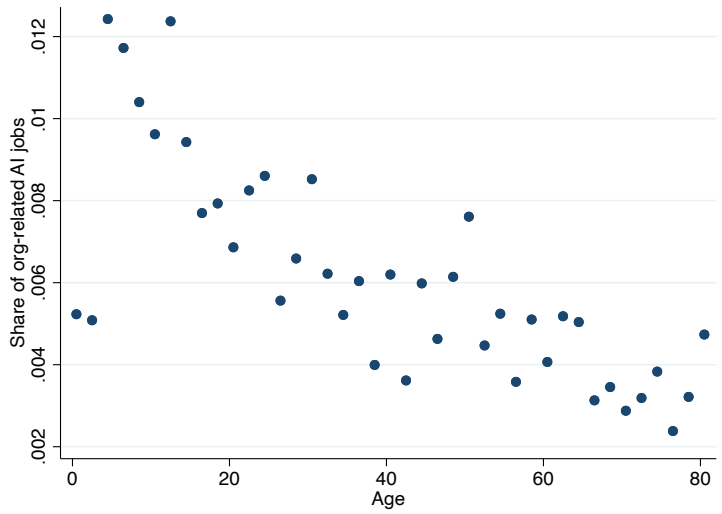


Figure D.5. Share of organization-capital-related AI jobs, by firm age. Firm age is measured in 2024 using founding year from Revelio. Each point is the average share across firms within a two-year age bin, with firm age top-coded at 80.

Recovering ψ and χ'/χ

Our specification is not an event study but a long difference: we regress the change in proxy for log organization capital, and separately the change in measured productivity, on the change in AI investment over the sample window. Proposition 3 maps directly into this design. Let H denote the length of the window and $g_H \equiv \log(1 + x_H)$ the log organization-capital gap that adoption has opened up by the end of it. The first regression recovers g_H and, by Proposition 3(iii), the second recovers ψ times the same object:

$$\beta_B = g_H, \qquad \beta_{\tilde{A}} = \psi\, g_H. \tag{D.10}$$

The elasticity is therefore the ratio of the two coefficients, and the treatment intensity is obtained by scaling the first up by however far convergence has run:

$$\psi = \frac{\beta_{\tilde{A}}}{\beta_B}, \qquad \log\frac{\chi'}{\chi} = \frac{g_H}{\varphi_H}, \qquad \varphi_H \equiv \frac{x_H}{\bar{x}} = 1 - \prod_{h<H} \rho_h. \tag{D.11}$$

The two objects are identified on very different terms, and it is worth being explicit about which is robust. The ratio in (D.11) requires no assumption about the horizon: by Proposition 3(iii) the productivity gap is ψ times the organization-capital gap at *every* h, so a partially converged long difference recovers ψ as cleanly as a fully converged one. It is also invariant to how the AI investment measure is scaled, since the normalization enters both coefficients and cancels, and it is robust to anything that moves B for reasons

outside the model, since such forces move measured productivity by ψ times as much. Recovering χ'/χ, by contrast, requires taking a stand on φ_H and on what a unit of the treatment represents.

In our data the instrumented long-difference coefficient is 0.24–0.30 for organization capital (columns 3 and 6 of Table B29) and 0.11–0.14 for firm productivity (columns 3 and 6 of Table 5), so (D.11) gives $\psi \approx 0.5$. Taking the view that our sample window captures roughly half of the eventual response—plausible given the half-lives implied in Section D.3.1—sets $\varphi_H \approx 0.5$ and hence $\log(\chi'/\chi) \approx 0.6$, or $\chi'/\chi \approx 2$: AI roughly doubles the rate at which operations are converted into organization capital. These are the values reported in Table D.1.

Both numbers should be read as illustrative magnitudes rather than precise estimates, because capitalized SG&A is a noisy measure of B and its errors do not obviously wash out. Two work in opposite directions. To the extent that the proxy contains a component of idiosyncratic productivity that does not respond to adoption, the measured stock is inflated without its response being inflated, attenuating β_B and thus overstating ψ while understating the plateau. To the extent that AI expenditure is itself partially booked in SG&A, part of β_B is accounting mechanics rather than organization-capital accumulation, which pushes ψ the other way. We report $\psi \approx 0.5$ and $\chi'/\chi \approx 2$ as the central case without claiming the two biases cancel.

D.6. Extension: Organization Capital as a Substitute for Measured Inputs

The environment above treats organization capital as a productivity shifter: it raises the marginal product of capital and labor, so a firm that accumulates B optimally expands its purchased inputs alongside its output. That is not what we find: sales respond to AI investment while employment and PP&E do not, so the coefficients on sales and on sales per worker are close to equal. This subsection sketches an extension in which that is the case. Organization capital is intangible: capitalized SG&A is by construction not part of PP&E, and the employees who build it are already counted in headcount. If the codified stock substitutes for purchased inputs rather than merely raising their productivity—a digital twin replacing physical trial capacity, a knowledge base replacing staffed support, route optimization replacing trucks and miles—then the firm's effective input bundle can expand while its measured inputs do not. Write

$$y = z\, A^{1-n}\, B^{\beta}\, \big[\, \underbrace{X + \lambda B}_{\tilde{X}} \,\big]^{n}, \qquad X = k^{\alpha} l^{1-\alpha}, \tag{D.12}$$

where X is purchased at market prices and λB is supplied by the firm's own codified stock at zero marginal cost. Organization capital enters twice, and both terms are necessary: β

governs whether output rises, and λ whether purchased inputs rise or fall. With $\beta = 0$, the firm holds output fixed and simply swaps purchased inputs for its own stock, so codification would be a pure cost saving, predicting flat sales and *falling* capital and employment, which is inconsistent with our empirical findings. Setting $\lambda = 0$ recovers the baseline environment exactly, so (D.12) adds one parameter rather than replacing the model.

Solving the static problem gives $\tilde{X} \propto A\, B^{\beta/(1-n)}$ and hence

$$y \;\propto\; A\, B^{\beta/(1-n)}, \tag{D.13}$$

so the extension is *observationally equivalent to the baseline for output*: the composite $\beta/(1-n)$ plays exactly the role ψ plays there, and output data cannot separate the two. We therefore write the extension in terms of the same ψ, with $\beta = \psi(1-n) = 0.075$ its structural counterpart given n. The life-cycle profiles in Figures D.1 and D.2 are unchanged, as are the event-study dynamics of Proposition 3, which concern the law of motion (D.4) rather than the production function. What the extension changes is the input side. Letting $s \equiv \lambda B/\tilde{X}$ denote the share of the effective bundle supplied by organization capital,

$$\frac{d\log X}{d\log B} = \frac{\psi - s}{1-s}, \qquad \frac{d\log(y/l)}{d\log B} = \psi - \frac{\psi - s}{1-s}. \tag{D.14}$$

Figure D.6 plots these. The baseline is the case $s = 0$: purchased inputs move with output and sales per worker does not respond. As s rises, a growing part of the expansion is met from the firm's own stock, purchased inputs respond less, and sales per worker responds more. Flat measured inputs correspond to the knife-edge $s = \psi$, at which purchased inputs are exactly unchanged and sales per worker moves one for one with sales—the pattern in our estimates.[61] Panel (b) shows that $s \propto B/y$, which is the observable proxy-to-size ratio of Section D.3.2; it is roughly flat over the first three decades of a firm's life and rises thereafter, and it shifts up permanently for an adopter.

To summarize, modeling AI as a codification technology—a change in the rate at which a firm's operations become durable, documented assets—yields a permanent level effect on firm size that arrives gradually, a permanent shift in the composition of firm assets, and post-adoption dynamics governed by a single exact recursion. Two implications organize the empirical work. Adjustment is slow for mature firms not because of frictions but because their stocks are large relative to their flows: plausible depreciation rates deliver

[61] In a more general specification that replaces $X + \lambda B$ with a CES aggregator of X and λB—which nests (D.12) as the limiting case of an infinite elasticity of substitution—flat inputs can arise for a range of substitution elasticities and input shares, rather than at a single value of s.

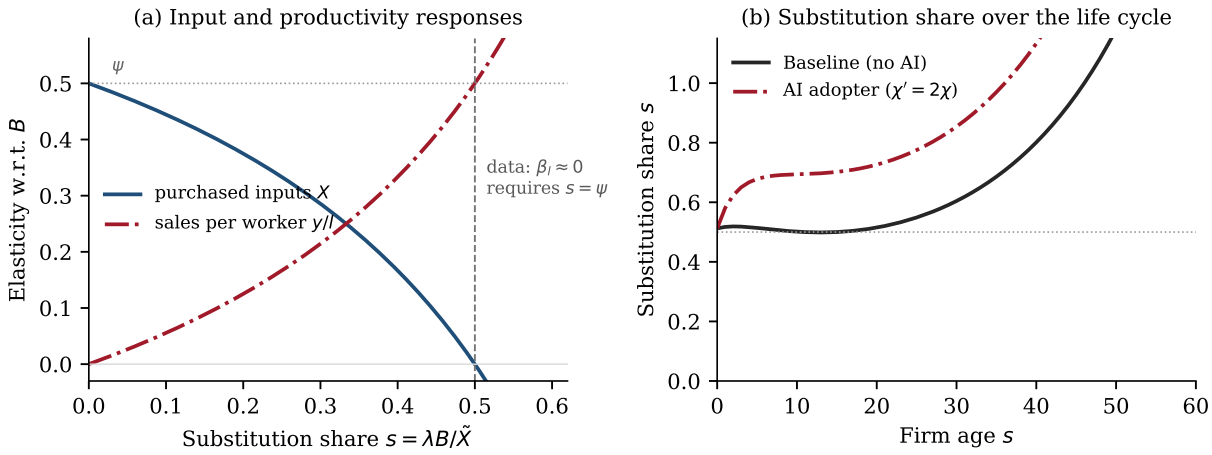


Figure D.6. Organization capital as a substitute for measured inputs. Panel (a): elasticities of purchased inputs and of sales per worker with respect to B, as functions of the substitution share s, from (D.14). Panel (b): the substitution share over the firm life cycle, normalized so that $s = \psi$ in mid-life.

half-lives of five years or more, matching the sluggish responses we estimate. And the response is governed by the organization capital ratio: firms that are large relative to what they have codified respond more, respond faster, and gain most from adopting—which is why the heterogeneity test we take to the data conditions on the pre-adoption level of organization capital to size.