Artificial Intelligence, Firm Growth, and Industry Concentration*

Tania Babina† Anastassia Fedyk‡

Alex He§ James Hodson¶

May 2021

Abstract

We study the use and economic impact of AI technologies among US firms. We propose a new human-capital-based measure of AI investments, using a unique combination of job postings and worker resume datasets. Our measure reveals a stark increase in AI investments across most sectors over the past decade. We document that firms that invest more in AI experience faster growth in sales and employment. While AI-investing firms do not see increased productivity, they grow through increased product offerings and geographic expansion. AI growth effects concentrate among the ex-ante largest firms, leading to higher industry concentration and reinforcing winner-take-most dynamics.

Keywords: artificial intelligence, technological change, technology adoption, economic growth, human capital, superstar firms, industry concentration

*First draft: March 16, 2020. The authors are grateful to Daron Acemoglu, Philippe Aghion, David Autor, Ann Bartel, Jan Bena, Francesco D'Acunto, Nicolas Crouzet, Xavier Giroud, Maarten de Ridder, Larry Katz, Anton Korinek, David Levine, Max Maksimovic, Gustavo Manso, Filippo Mezzanotti, Thomas Philippon, David Sraer, Tano Santos, Bledi Taska, Laura Veldkamp, Shang-Jin Wei, Michal Zator, and seminar participants at AFFECT, Babson College, Columbia University, FOM Conference, Labor and Finance Group Conference, Michigan State University, NBER Corporate Finance, NBER SI Macroeconomics and Productivity, NBER productivity seminar, NYU WARPFIN, Tilburg University, Triangle Macro-Finance Workshop, Queen Mary University of London, Sao Paulo School of Economics, UC Berkeley, UNC Junior Roundtable, University of Illinois Chicago, University of Maryland, University of Oklahoma, University of Washington, and Vienna Graduate School of Finance for helpful suggestions. The authors thank Cognism Ltd. for providing the employment data, Burning Glass Technologies for providing the job postings data, and Bernhard Ganglmair, W. Keith Robinson, and Michael Seeligson for sharing their data on patent (product vs. process) classification. Nikita Chuvakhin, Joanna Harris, Derek Ho, Gary Nguyen, and Hieu Nguyen provided excellent research assistance. Fedyk gratefully acknowledges financial support from the Clausen Center for International Business and Policy at UC Berkeley.

†Columbia University. Email: tania.babina@gsb.columbia.edu.

‡University of California, Berkeley. Email: fedyk@berkeley.edu.

§University of Maryland. Email: axhe@umd.edu.

¶AI for Good Foundation. Email: hodson@ai4good.org.

Technological change is a key driver of economic growth. The past decade has seen a new technological shift: advances in computing power and data availability led to substantial developments in artificial intelligence (AI) technologies and their wide-spread commercial applications.[1] In surveys of executives at companies investing in AI, 70% anticipate that AI investments will fundamentally transform their companies and industries within the next five years.[2] Yet it remains an open question whether artificial intelligence can transform economies and spur economic growth. On the one hand, AI is considered a general purpose technology that can generate growth through increased productivity and product innovation (Aghion et al., 2017; Cockburn et al., 2018). On the other hand, lackluster aggregate productivity growth over the past decade has led to concerns that the benefits of AI may be over-hyped or take a much longer time to materialize (Mihet and Philippon, 2019). To date, the lack of comprehensive data on firms' use of AI technologies has posed the key challenge to understanding the prevalence and the economic impact of AI investments (Seamans and Raj, 2018).

In this paper, we propose a new measure of investments in artificial intelligence technologies based on firms' AI-skilled human capital. The heavy reliance of AI on human expertise makes the human-capital-based approach particularly well-suited in our setting. Our measure takes advantage of a unique combination of datasets that capture both the *stock* and the *demand* for AI-related employees among U.S. firms: job postings data from Burning Glass, which include 180 million job vacancies, and resume data from Cognism Inc, which offer job histories for 235 million individual employees. This new measure allows us to document the patterns of AI adoption and its impact on the adopting firms and industries. We find a stark increase in AI investments across all major sectors in the last decade, with firms investing more in AI seeing higher growth through increased product offerings and geographic expansion. At the aggregate level, AI investments are associated with increased industry concentration as larger firms benefit most from AI investments.

Our work offers several innovations over the existing literature. First, while prior research has made progress in studying the impact of AI in various applications and settings (e.g., DAcunto et al., 2019; Grennan and Michaely, 2019), we are able to measure AI investments and their potential effects for a large sample of firms across a wide range of industries. Second, our paper is the first to cross-validate AI labor demand identified from job postings with the stock of AI workers identified from resumes, which captures actual inflows and outflows of AI workers, including those on-boarded via acquisitions. Third, our rich data on firms' employees and their jobs allow us to measure and control for confounding factors such as the use of non-AI information tech-

[1] Bughin et al. (2018) and Furman and Seamans (2019) provide an overview of AI investments in the private sector. When it comes to public investments, the U.S. government is looking to double its non-defense research and development (R&D) in AI (Executive Office of the President, 2019); the European Union has called for $24 billion investments in AI research by 2020 (European Commission, 2020); and China aims for $150 billion invested in the domestic AI market by 2030 (Mou, 2019).

[2] See here for a survey by Deloitte in 2018.

nologies. Finally, our comprehensive cross-industry data provide a first glimpse at the aggregate industry-level implications of AI investments.

Even with our detailed data, identifying firms' investments in AI is difficult, due to the multi-faceted nature of AI applications.[3] We circumvent this challenge by proposing a new data-driven approach to identify AI-related jobs, which does not depend on pre-specified lists of keywords. Instead, our algorithm learns the AI-relatedness of each job posting empirically. First, our approach measures the AI-relatedness of each skill in the job postings data, based on that skill's co-occurrence with the core AI skills—machine learning, computer vision, and natural language processing.[4] Second, we obtain a measure of AI-relatedness of each job posting by averaging the AI-relatedness of all skills required in that job posting. Finally, we leverage the most AI-related skills in the job postings data to classify AI workers in the less structured resume data. For each employee, we consider whether skills with the highest AI-relatedness (e.g. "deep learning") appear either in the job title, in the job description, or in any publications, patents, or awards received during that job. This gives us a classification of each employee of each firm at each point in time. We aggregate both job postings data and resume data to the firm level and match to public firms in the Compustat database.

Encouragingly, the measures of AI investments, although based on two independent datasets, are highly correlated and yield consistent results throughout the paper. We also confirm that our human-capital-based measures of AI investments display intuitive properties. First, we manually inspect large samples of AI-classified jobs and confirm that indeed our classification picks up highly AI-skilled positions. Second, given that we do not use job titles in the job postings data to filter AI-skilled jobs, we validate our measure by confirming that the job postings with the highest AI-relatedness measures skew heavily towards highly AI-specific job titles. Third, we confirm that our AI measure does not pick up general data-related skills, only those that are highly associated with AI implementation. Fourth, we dive into the specific applications of AI within several firms and provide detailed case studies. Fifth, since AI investments involve a great deal of experimentation with applying the new technology, we expect AI-investing firms to increase research and development (R&D) expenditures, which is indeed what we find. Finally, we enrich our baseline measure by incorporating the use of external AI solutions and software (e.g. IBM Watson, IPSoft Amelia), and find similar results.

[3]Some firms have dedicated AI research groups, while others host AI workers in their Technology division, and others still have distributed AI human capital across divisions. Within a single firm, e.g. Caterpillar Inc., AI can have use cases ranging from demand forecasting to improved machinery via computer vision.

[4]The three key AI sub-fields—machine learning, computer vision, and natural language processing—capture the bulk of private sector investments in AI (see a survey by Deloitte here). Our method is able to effectively separate specifically AI-related skills from skills that are related to more general topics (such as data analysis, statistical methods, econometrics) with a high degree of precision. For example, the AI-relatedness measure for "deep learning" is 0.86 (on a continuous scale from 0 to 1, meaning that 86% of all jobs that require deep learning also require one of the four core AI skills), while the measure is 0.37 for "information retrieval", 0.09 for "regression analysis", 0.04 for "data modeling", 0.03 for "quantitative analysis", and 0.003 for "communication".

We begin our analysis by describing key patterns in AI investments. In both datasets, the fraction of AI jobs has increased dramatically over time, growing more than five-fold from 2010 to 2018. The share of AI jobs is highest in the technology sector, but the rate of increase in AI investments over time is similar across all sectors. At the firm level, AI investments are more pronounced among larger firms and firms with ex ante higher cash holdings, mark-ups, and R&D intensity. Looking at the local labor market conditions, we observe that higher-wage and more educated areas experience faster growth in AI-skilled hiring.

We next address the central question of whether AI investments are associated with higher firm growth. Given that AI investments are gradual over time, and that the effect of AI may not be immediate, our primary specification is a long-differences regression of changes in firm outcomes from 2010 to 2018 on contemporaneous changes in the share of AI workers. We also control for industry fixed effects and firm-, industry-, and commuting-zone-level characteristics in 2010. We document a strong and consistent pattern: firms that invest in AI grow more. For example, a one standard deviation increase in the resume-based measure of AI investments over the 8-year period corresponds to a 15.6% increase in sales, a 15.2% increase in employment, and a 1.4 percentage point increase in market share within the 5-digit NAICS industry. Moreover, the results are ubiquitous across major industry sectors (e.g., manufacturing, finance, and retail), supporting the idea that AI is a general purpose technology.

Our main empirical design effectively compares firm-level growth across firms with larger versus smaller changes in AI investments. This specification controls for time-invariant firm characteristics but does not rule out concerns that unobserved time-varying shocks or reverse causality can drive the results. We perform several tests to address these concerns and buttress a causal interpretation of our results. First, to address concerns about reverse causality, we conduct an event study and examine the dynamics of firm growth in years around AI investments. We find no differential trends in firm growth prior to AI investments and a positive effect on sales and employment after a lag of two years following AI investments. Second, we confirm that our results reflect specifically investments in AI, rather than other technologies. We show that the relationship between AI investments and firm-level growth is robust to controlling for contemporaneous firm-level investments in robotics, non-AI information technologies, and non-AI data analytics. Third, we use two instrumental variable strategies to address unobserved firm- and industry-level shocks that could be correlated with AI investments. The first one builds on prior work of Acemoglu and Restrepo (2020), where US technology adoption is instrumented with the European industry-level adoption rates—in our case for AI investments. The second is a shift-share instrument that relies on firms' exposure to the local supply of AI labor based on pre-existing firm locations and local industry composition.

The benefits from AI investments are unevenly distributed across firms. First, following the hypothesis that AI can favor large firms with more data and increase inequality among firms (Mi-

het and Philippon, 2019; Farboodi et al., 2019), we estimate the effect of AI investments within groups of firms by initial size. The positive relationship between AI investments and growth is much stronger among the ex ante largest firms. Second, we test whether more or less productive firms experience higher growth following AI investments. Recent literature argues that more productive firms are able to leverage new technologies better, leading to a phenomenon of "superstar" firms (Autor et al., 2020), but some of the previous technological innovations saw laggard firms benefiting most (Cockburn et al., 2000). In the case of AI, we find that the benefits are concentrated among ex ante more productive firms. These cross-sectional results appear most consistent with increasing firm inequality, where AI facilitates scale advantages for the ex ante largest and most productive firms.

To understand the mechanisms through which AI could induce firm growth, we examine two key non-mutually-exclusive channels: (i) productivity enhancements and cost cutting; and (ii) product innovations and expansion. Empirically, we do not find support for the productivity channel, at least in the short run, with AI investments having zero effect on both sales per worker and total factor productivity. The second channel reflects the idea that big data and AI enable granular product tailoring and unleash new modes of product innovations. We do find empirical support for the second channel, documenting that firms with larger AI investments see increased product creation and expand across geographic markets. Specifically, AI investments are related to increased patenting activity, driven exclusively by product (rather than process) patents, increased trademarks, increased hiring of product managers, and an increased number of commuting zones spanned by the firm.

Our final set of results focuses on potential aggregate effects of AI on growth and industry dynamics. We first test whether firm-level growth translates into industry-level growth in sales and employment. It is possible that the positive effects on firms investing in AI are offset or even dominated by negative spillovers to competitors within the industry as output and labor are reallocated from other firms to the AI-investing firms. Nevertheless, we find that industries that invest more in AI experience an overall increase in sales and employment within the sample of Compustat firms. The increase in employment is especially surprising, given the labor-replacing effect of other digitization and automation technologies (Zator, 2019) and wide-spread concerns regarding the potential for AI to replace labor (Frank et al., 2019).[5] Second, AI investments are associated with increased industry concentration, as measured by the Herfindahl-Hirschman Index (HHI) or the fraction of total industry sales accruing to the single largest firm. We consider the implications of increased concentration by examining whether AI-investing firms increase their market power, which speaks to the potential anti-competitive effects of AI. However, we document zero relationship between firms' AI investments and changes in markups, suggesting that current uses

[5]For theoretical treatments of the impact of artificial intelligence technologies on labor displacement, see Korinek and Stiglitz (2017), Acemoglu and Restrepo (2019), and Agrawal et al. (2019).

of AI, while facilitating the expansion of large firms and increasing industry concentration, do not predict measurable increases in firms' market power.

Our paper contributes to the growing literature in corporate finance where recent studies make progress in examining the impact of AI technologies in various settings: robo-advising (DAcunto et al., 2019), financial analysts (Grennan and Michaely, 2019; Coleman et al., 2020), fintech innovation (Davis et al., 2019), loan underwriting (Jansen et al., 2020), investment management (Abis and Veldkamp, 2020), mortgages (Fuster et al., 2020), auction pricing (Aubry et al., 2020), and corporate boards (Erel et al., Forthcoming). Our key contribution to this nascent literature is to offer our comprehensive data to measure AI investments in a large sample of firms and industries. This allows us to examine the incidence of AI investments and their potential effects on firm growth and industry dynamics, as well as the economic mechanisms driving these effects.[6]

Methodologically, this paper offers a new measure of firms' intangible capital based on human capital, with a specific application to capturing investments in AI. Despite ongoing efforts to incorporate more comprehensive measures of intangibles at the national level (Corrado et al. 2016), most firm-level measures of intangible capital use cost items such as R&D and SG&A (e.g., Eisfeldt and Papanikolaou, 2013; Peters and Taylor, 2017; Crouzet and Eberly, 2018). Our methodology complements these measures by offering a new measure of intangibles that is consistent across firms and sectors. Importantly, our methodology can be applied to measure other forms of intangible assets, especially those based on human expertise. For example, while our focus is on AI investments, we are also able to measure (and control for) firm investments in robotics, non-AI information technology, and non-AI data analytics. Our methodology of using firm workers and their skills also complements the corporate finance literature that uses firms' financial filings to measure intangible assets (e.g., Hoberg and Phillips, 2016; Hoberg and Maksimovic, 2019).

Finally, we contribute to the active debate on the causes and consequences of rising industry concentration (e.g., Gutiérrez and Philippon, 2017; Grullon et al., 2019). One proposed channel is that intangible assets propel growth of the largest firms, contributing to increased industry concentration (e.g., Crouzet and Eberly, 2019; Ayyagari et al., 2019).[7] Our results support this hypothesis and suggest that technologies like AI can contribute to increased concentration by enabling large firms to grow even larger and reinforcing winner-takes-most dynamics. Importantly, AI investments are not associated with higher market power. Instead, AI-investing firms expand their product offerings and geographic footprint. As such, our evidence supports the recently proposed theoretical mechanism that new technologies can reduce the overhead costs of spanning multiple markets and enable the most efficient firms to scale more easily (Aghion et al., 2019).

[6]To foster further research on the economic impacts of AI, all code to generate our AI investment measures and our firm-level AI data will be publicly available on the authors' websites.

[7]Other proposed explanations for growing concentration include increasing barriers to entry and lax antitrust enforcement (e.g., Grullon et al., 2019; Covarrubias et al., 2019), low interest rates (Liu et al., 2020), and globalization (e.g., Elsby et al., 2013).

The paper proceeds as follows. Section 1 provides background on artificial intelligence and its potential impact on firms. We introduce our data in Section 2 and detail our methodology for measuring AI investments in Section 3. Section 4 addresses the question of which firms choose to invest in AI, while Section 5 presents our main results on the relationship between AI investments and firm growth, including the mechanisms behind the main effects. Section 6 considers potential aggregate industry-level effects, Section 7 addresses alternative explanations, and Section 8 concludes.

1 Artificial Intelligence: Background and Mechanisms

According to Organisation for Economic Co-operation and Development (2019), an AI system is defined as a *"Machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments."* We provide a brief overview of the current state-of-the-art and key features of AI, followed by a discussion on economic mechanisms through which AI investments might benefit a broad range of firms.

1.1 Artificial Intelligence: A Brief Overview

Commercial applications and investments in AI have increased exponentially over the past decade. While there are no systematic data on AI investments by firms, recent estimates hover around $70 billion globally per year (see Stanford's AI Index Report 2021 here). Historically, the U.S. is considered to be the leader in both academic research and private investments in AI, but other regions including China and the E.U. have recently been spearheading their investments (Knight, 2017). There has also been an expansion of AI investments across industry sectors. While the tech sector was an early adopter of AI, surveys of executives indicate widespread adoption of AI technologies by firms in all industries, with respondents in the automobile and healthcare industries most likely to report increased investments in AI in 2020 (see here for a survey by McKinsey).

Academic research in AI has flourished for decades since John McCarthy coined the term in 1956 (McCarthy et al., 2018). The recent explosion of interest in and commercial applications of AI by the private sector are driven by a rapid accumulation of data, decreasing costs of computation, and advances in methodologies such as deep learning. In terms of commercial applications, three key areas of artificial intelligence have captured the bulk of private sector investments: machine learning, natural language processing, and computer vision (see here for a survey by Deloitte in 2018).[8] These core techniques are united by their ability to perform high-skilled, non-routine tasks such as prediction, detection, and classification. Their main distinction from traditional meth-

[8]While our focus is on artificial intelligence technologies rather than more general automation technologies like ATMs and industrial robots, our measure does incorporate specific recent robotics technologies (e.g., autonomous vehicles, vision-guided robots) that are highly related to computer vision and machine learning technologies.

ods of data analysis consists of these techniques' ability to learn to optimize objectives from vast quantities of high-dimensional data (including text, speech, and image data) and significantly improve the accuracy of predictions. For example, the ImageNet challenge of 2012[9] led to an almost halving of image recognition error rates (relative to traditional methods), which launched large corporate interest in the computer vision space. The ability of AI to aid in the decision-making process and in solving complex cognition problems has led to concerns that AI can disrupt many high-skill and high-wage occupations, in contrast to previous waves of technology adoption. For example, the literature has shown that the widespread adoption of IT in the 1990s and early 2000s mainly aided and replaced medium-skilled routine and codifiable tasks, whereas industrial robots performed mainly low-skilled manual and repetitive tasks.

AI has several important economic properties. First, economists have argued that AI may be a general purpose technology (GPT): AI can be leveraged across different business segments and sectors to solve a wide range of business problems. Well-known examples of GPT include the steam engine, electricity, and internal combustion. Second, investments in AI center around human expertise, with complementary investments in computing technology and data infrastructure. As such, AI is an intangible asset, reflecting the broader shift towards intangible capital (Mihet and Philippon, 2019). The third key feature of AI technologies is that they are information goods with non-rival uses: new algorithms are usually published openly and can be used simultaneously by many firms. However, the extent to which AI can be beneficial to firms depends to a large extent on who owns big data—the key input to AI technologies (Fedyk, 2016; Jones and Tonetti, 2020).

1.2 Artificial Intelligence and Firm Growth: Mechanisms

It is an open question whether and how investments in AI technologies benefit firms. On the one hand, as a potential general purpose technology, AI might raise productivity and facilitate the creation of new products or the entry into new markets. On the other hand, current attention to AI may be over-hyped (Mihet and Philippon, 2019), or AI may still be too early in the adoption cycle to have a meaningful impact on firm growth (Brynjolfsson et al., 2020). Below, we discuss two non-mutually-exclusive channels through which investments in AI can benefit firms: improvements in productivity and expansion into new products and markets.

AI as a Driver of Productivity Growth. Technological innovations often aim at streamlining operations and improving productivity. For example, information technologies during the 1990s brought about a number of productivity improvements (Bartel et al., 2007), and analogous effects have been observed for robotics (Acemoglu et al., 2020; Graetz and Michaels, 2018). When it comes to AI, the technology can increase productivity in at least two ways. First, AI can po-

[9] See ImageNet Large Scale Visual Recognition Challenge 2012 (ILSVRC2012)

tentially replace human labor for some tasks (Webb, 2020; Agrawal et al., 2019), cutting per-unit labor costs. Second, AI can increase efficiency through better forecasting and decision-making (Mihet and Philippon, 2019; Agrawal et al., 2020). This aspect is explored in depth by Tanaka et al. (2019), who present a model of firm input choice under uncertainty and costly adjustment, where forecast errors result in under- or over-investment.[10] The forecasting applications of artificial intelligence are indeed prevalent in our data. In Internet Appendix A.1, we conduct case studies of the applications of AI in four firms across four industries (UnitedHealth Group, JPMorgan Chase, Caterpillar, and Qualcomm) based on our detailed resume data. The case studies highlight how AI-enabled forecasting is implemented across a variety of industries: for example, AI workers at JPMorgan Chase model default of non-performing loans; Caterpillar leverages AI to for unstable demand anticipation and inventory management; and UnitedHealth uses AI to support efficient medical billing.

These potential improvements can manifest through higher firm-level productivity. Empirically, AI-induced productivity gains should translate into increases in measured total factor productivity (TFP) in a Cobb-Douglas framework and higher measured output per worker under more general production functions.

AI as a Driver of Expansion into New Products and Markets. AI has the potential to fuel product innovations by improving forecasting and reducing the uncertainty of exploration (Cockburn et al., 2018). According to surveys of executives, a top use of AI is the enhancement of existing products and services and the creation of new ones (see a 2020 survey from Deloitte here).[11] Industry reports also highlight that spending on AI is allocated "roughly equally among radical innovation; major new products, services, and internal process changes; and improvement of existing products, services, and processes" (Ringel et al., 2019). The radical innovation portion includes the ability to expand into products and services that were previously either infeasible or not cost-effective: for example, the introduction of AI-driven trading platform DeepX at JPMorgan and the "smart" machinery at Caterpillar. Moreover, AI can also allow firms to tailor product and service offerings to customers' tastes and needs. For example, Amazon may use machine learning to build more specialized products tailored to certain customers.

The ability to make predictions with AI may increase not only the variety of products offered by AI-investing firms but also the set of customers targeted. For example, firms such as JPMorgan Chase highlight the importance of AI technology in finding new clients (i.e., "prospecting").[12] Similarly, Repsol, an energy and utility company, uses AI to prepare personalized offers for 8

[10] Their empirical evidence supports the prediction that forecast accuracy is associated with higher profitability among Japanese firms. Chase (2013) draws on executive interviews to point out that big data can streamline demand forecasting, leading to more efficient inventory management.

[11] This is further underscored by university curricula; for example, Duke University offers a master of engineering degree specifically in Artificial Intelligence for Product Innovation.

[12] The letter to the shareholders from the 2020 Annual Report by JPMorgan states: "We already extensively use AI, quite successfully, in fraud and risk, marketing, prospecting, idea generation, operations, trading and in other areas".

million customers, significantly expanding the geographic reach of its client base. This indicates that AI could enable firms to scale more easily across different segments of consumers.

These applications of AI are consistent with recent theories of technological change, which model technological innovations as lowering the overhead costs of spanning multiple markets (Aghion et al., 2019; Hsieh and Rossi-Hansberg, 2019). Empirically, if AI enables product innovation and expansion, we would observe that AI investments are associated with increased product innovation and variety as well as expansion into new markets, e.g. geographically.

2 Data

We provide a uniquely comprehensive perspective on firm-level AI investments by simultaneously measuring firms' *demand* for AI workers through job postings and the *stock* of AI workers through employment profiles. We detail each dataset and describe our sample construction.

2.1 Job Postings from Burning Glass

The first dataset we use is a proprietary dataset covering over 180 million electronic job postings in the United States in 2007 and 2010–2018. The dataset is provided by Burning Glass Technologies (BG in short) and draws from a rich set of sources. BG examines more than 40,000 online job boards and company websites, aggregates the job postings data, parses them into a systematic, machine-readable form, and creates labor market analytic products. The company employs a sophisticated deduplication algorithm to avoid double counting vacancies that post on multiple job boards. BG data contain detailed information for each job posting, including job title, job location, occupation, and employer name. Importantly, the job postings are tagged with thousands of specific skills standardized from the open text in each job opening. The main advantages of the BG dataset are the breadth of its coverage and the rich detail of the individual job postings. The dataset captures a near-universe of jobs that were posted online and covers approximately 60–70% of all vacancies posted in the U.S., either online or offline (Carnevale et al., 2014). Hershbein and Kahn (2018) provide a detailed description of the BG data and show that their representativeness is stable over time at the occupation level.

We focus on jobs with non-missing employer names and at least one required skill. About 65% of the job postings have employer information and 93% of the job postings require at least one skill.[13] We also drop job postings that are internships. We then match the employer firms in the remaining job postings to Compustat firms. This step is necessary to aggregate job postings to firm level and merge with other firm-level variables. We perform a fuzzy matching between firm names in BG and Compustat after stripping out common endings such as "Inc" and "L.P.". For

[13]The job postings with missing employer names are primarily those listed on recruiting websites that mask the employers' identities.

observations that do not match exactly on firm name, we manually assess the top ten potential fuzzy matches by looking at the firm name, industry, and location. Out of 112 million job postings with non-missing employer names and skills, 42 million (38%) are matched to Compustat firms. This slightly overrepresents employees of publicly listed firms, which constitute more than one fourth of U.S. employment in the non-farm business sector (Davis et al., 2006).

2.2 Employment Profiles from Cognism

While job postings provide a useful measure of the firms' labor *demand,* we also use a second dataset providing comprehensive information on the *stock* of employees at each firm in each year. In particular, we leverage a novel dataset of approximately 235 million individual profiles provided by Cognism, an aggregator of employment profiles for lead generation and client relationship management services. Cognism obtains the resumes from a variety of sources, including publicly available online profiles, collaborations with recruiting agencies, third party resume aggregators, and human resources databases of partner organizations, as well as direct user contributed data.[14] These data are introduced and described in detail in Fedyk and Hodson (2019). While the data slightly over-represent high-skilled employees, they cover approximately 43% of the entire U.S. workforce as of 2019 and offer a representative breakdown across 2-digit NAICS industries. For each employment record listed by the individual, we see the start and end dates, the job title, the company name, and the job description. Individuals may also volunteer self-identified skills and list their courses attended, certifications, patents, awards, and publications.

We match employer names in the profile data to the names of publicly traded firms using a similar approach to matching employers in BG data to Compustat firms. Fedyk and Hodson (2019) provide further details on the procedure as applied to resume data. Of the 462 million US-based person-firm-year employment records between 2007 and 2018, 86 million (19%) are matched to U.S. public firms that are headquartered in the U.S. This is consistent with approximately 26% of overall U.S. employment being accounted for by publicly listed firms (Davis et al., 2006), only our sample additionally excludes cross-listed firms or firms with headquarters outside of the U.S. The sample of 86 million person-firm-years matched to U.S. public firms is comprised of 14 million distinct individual employees.

2.3 Additional Data Sources

We merge the Burning Glass job postings data and the Cognism resume data to several additional data sources. We use commuting-zone-level wage and education information from the Census American Community Surveys (ACS), price data from the Bureau of Labor Statistics (BLS), and industry-level wages and employment from the Census Quarterly Workforce Indicators (QWI).

[14] The processing of all profiles is compliant with the applicable GDPR and CCPA regulations.

Firm-level operational variables (including sales, employment, assets, net income, cash, costs of goods sold, operating expenses, R&D expenditures, and SG&A expenses) come from Compustat.

3 Methodology: Measuring AI Investments by Firms

3.1 AI Investments from Job Postings (Burning Glass)

We take advantage of the detailed information on required skills in the job postings data to propose a new data-driven methodology for identifying AI-related jobs. Other work classifies job postings based on the presence of key terms from a pre-specified list,[15] which is likely to suffer from both Type I (incorrectly labeling tangentially-related employees as AI-related) and Type II (missing real AI skills that did not make the initial dictionary) errors due to the arbitrariness of the list of keywords. This is especially relevant in a quickly-evolving domain such as AI, where new emerging skills can easily be missed. Our methodology circumvents these challenges by learning the AI-relatedness of each of approximately 15,000 unique skills directly from the job postings data, based on their empirical co-occurrence (within the required list of skills for each job posting) with unambiguous core AI skills. We then aggregate the skill-level measure to the job level by generating a continuous measure of AI-relatedness for each job posting, from which we can classify employees into AI-workers and non-AI-workers.

To measure the AI-relatedness of each skill, we calculate the skill's co-occurrence with Artificial Intelligence (AI) and its three main sub-fields: machine learning (ML), natural language processing (NLP), and computer vision (CV) (Deloitte, 2018):

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

Intuitively, this measure captures how correlated each skill s is to the core AI skills. For example, the skill "Tensorflow" has a value of 0.9, which means that 90% of job postings with Tensorflow as a required skill also require one of the core AI skills or contain one of the core AI skills in the job title. Hence, a "Tensorflow" requirement in a job posting is highly indicative of that job being AI-related. On the other hand, the skill "Communication" is required in a large share of jobs across the board, and its AI-relatedness measure is only 0.003. We list the skills with highest AI-relatedness measures in Table A.1.

One potential concern is that this measure might be picking up general data analytics skills not specific to AI. To address this concern, we further refine our measure and screen out non-

[15] For example, Hershbein and Kahn (2018) identify jobs requiring cognitive abilities if any listed skills include at least one of the following terms: "research," "analy-," "decision," "solving," "math," "statistic," or "thinking." Similar bag-of-words approaches with pre-specified search terms are used to identify AI-related employees (e.g., Alekseeva et al., 2020; Acemoglu et al., 2021).

AI-specific skills: we manually categorize around 700 skills that have an AI-relatedness measure above 0.05 and are required in at least 50 job postings into "narrow" and "broad" AI skills. The threshold of 0.05 is set sufficiently low to ensure that we do not miss any AI skills. We manually screen out skills that represent general programming languages (e.g., Python) or statistics (e.g., linear regression) and leave only those skills that relate specifically to AI, including AI methodology or algorithms (e.g., supervised learning) and AI software (e.g., Tensorflow).

We define the job-level AI-relatedness measure (ω_j^{AllAI}) for a given job posting j as the mean skill-level measure (w_s^{AllAI}) across all required skills s listed in job posting j. Throughout the main analysis we concentrate on the refined measure that leaves out the non-AI-specific skills, denoted by $\omega_j^{NarrowAI}$ for job posting j. We transform this continuous measure into a binary indicator by defining each job posting j as AI-related if the measure $\omega_j^{NarrowAI}$ is above 0.1.[16] The firm-level measure $Share_{f,t}^{NarrowAI}$ is then defined as the fraction of job postings by firm f in year t that are AI-related (i.e. $\omega_j^{NarrowAI} > 0.1$).[17] We use a discrete classification for ease of interpretability and consistency with the resume-based measure in 3.2, but we show in Section 5.1 that the results are robust to (i) varying the threshold of the discrete measure, and (ii) using the continuous measures ($\omega_j^{NarrowAI}$ or ω_j^{AllAI}) aggregated to the firm level.

Table A.2 presents some examples of AI and non-AI job postings. For each job, the narrow AI measure is the average AI-relatedness of skills that are classified as narrow AI skills. These skills are bolded in the table with their AI-relatedness score in parentheses.[18] Our measure enables us to capture a wide range of AI-related jobs, from data scientists to speech recognition scientists to autonomous vehicle engineers. While many AI-related jobs are data scientists and similar data-analysis-related jobs, our measure differentiates data-analysis jobs specifically related to AI (job postings numbered 6-10) from data-analysis jobs that are not specific to AI and focus on more traditional statistical methods (job postings numbered 11-15).

3.2 AI Investments from Resumes (Cognism)

We complement the job postings with employee resumes, which allow us to measure the actual stock of AI workers at each firm and help address potential concerns around job postings: for example, if a firm is not able to hire despite active job postings, if a firm posts numerous job openings for one job, or if a firm on-boards AI talent through acquisitions rather than direct hiring.

[16] This cut-off was selected based on manual inspection of the data to capture the full range of AI-related technical jobs while minimizing any false positives. We show that our results are robust to alternative cut-offs, e.g. 0.05 and 0.15.

[17] Throughout our empirical analyses, we focus on jobs that are matched to Compustat firms. Figure A.5 plots the share of all job postings and the share of AI-related job postings that are matched to Compustat in each year. Although publicly listed firms constitute 38% of all job postings, they account for about half of all AI-related job postings. This suggests that, on average, publicly-listed firms hire more AI workers than private firms.

[18] Narrow AI skills include AI-specific tools and applications as well as programming and data skills highly related to AI (e.g., Apache Hadoop, Boosting) but exclude programming and data skills that are not highly specific to AI (e.g., SAS, Web Analytics, Exploratory Analysis).

In the Cognism resume data, we identify AI-related employees as those whose current positions directly involve AI. We begin with the set of 67 keywords in Table A.1, which have the highest skill-level AI-relatedness measures. We then search for these terms in every employment record of each individual in the resume data. Specifically, for each job record, we consider four aspects: (i) whether that job (role and description) directly includes any of the identified AI terms; (ii) whether any patents obtained during the year of interest or the two following years (to account for the time lag between the work and the patent grant) include these AI terms; (iii) whether any publications during the year of interest or the following year include the AI terms; and (iv) whether any of the identified AI terms appear in awards received during the year of interest or the following year. If any of these conditions are met, then that person at that firm at that time is classified as an AI-related employee. For example, jobs with titles such as "senior **machine learning** developer" or descriptions such as "develop chatbots using Python with **Tensorflow** and **deep learning** models" are identified as AI jobs.

After classifying each individual at each point in time, we use the number of AI-related employees and the number of total employees at each firm in each year to compute the percentage of employees of that firm in that year who are classified as AI-related. Our firm-level measure focuses on the employees of each firm within the same country as the firm's listed exchange: i.e., for firms listed on U.S. exchanges, we consider employees currently based in the U.S. For all analysis involving the resume data, we exclude firm-years with fewer than 50 employees to reduce noise due to low coverage.

3.3 Summary Statistics and Validation

We examine both of our constructed measures of AI investments, confirm that they display intuitive properties, and show how potential limitations with measuring AI investments through job postings are addressed with the resume data.

First, we document that both measures—based on job postings and resumes—display a natural rise over time, increasing more than fivefold from 2010 to 2018. Panel (a) of Figure 1 focuses on the job-postings-based measure and shows that the average continuous job-level measure $\omega_j^{NarrowAI}$ starts out at 0.02% at the beginning of the sample in 2007. The measure rises monotonically over time, with the increase speeding up from 2014 to 2018, and peaks at 0.19% at the end of the sample in 2018.[19] Panel (b) shows analogous patterns in the resume data, plotting the fraction of all employees in each year who are classified as AI-related. The fraction of AI employees starts at 0.03% in 2007 and reaches 0.24% in 2018.

[19] In unreported results, we confirm that the rise in $\omega_j^{NarrowAI}$ accounts for almost all of the rise in w_j^{AllAI} over time, which means that the increase in AI-relatedness of jobs is entirely driven by the increase in the frequency of required skills that are highly AI-specific and form the basis of our refined "Narrow AI" measure instead of other less AI-specific skills (e.g. statistics or programming languages).

Second, we document that the increase in AI jobs displays an intuitive distribution across industries. Panel (a) of Figure 2 plots the average AI-relatedness measure of job postings in each of the 2-digit NAICS sectors, separately for the years 2007–2014 and 2015–2018. Panel (b) repeats the same analysis for the average AI-relatedness of employees. The figure highlights that AI investments are highest in the "Information" sector, growing from 0.19% (0.09%) in the early years of 2007–2014 to 0.49% (0.52%) in the later period of 2015–2018 based on job postings (resumes). However, all economic sectors see a meaningful increase in both measures from earlier (2007–2014) to later (2015–2018) in the sample period, supporting the notion that AI is a general purpose technology (Goldfarb et al., 2019). The ability of our measures to pick up AI investments in a broad cross-section of economic sectors highlights a key advantage of our human-capital-based approach. Within each industry, we also document that AI investments correlate positively with increased R&D expenditures. For example, changes in the resume-based share of AI workers from 2010 to 2018 display a correlation of 0.30 with changes in log R&D expenditures over the same time period, controlling for industry fixed effects. The pattern of AI-investing firms increasing research and development (R&D) expenditures supports the notion that AI investments involve a great deal of experimentation with applying the new technology.

Third, digging deeper into the skills and jobs that receive the highest AI-relatedness measures according to our methodology, we observe that our measure is indeed capturing the essence of AI investments by firms. As illustrated in Table A.2, our measure picks up jobs that are specific to AI, rather than broader data analysis jobs. The skills with highest AI-relatedness measures, presented in Table A.1, capture highly AI-specific skills, such as "Tensorflow" and "random forests". On the other hand, general data-analytics-related skills have low AI-relatedness measures: for example, the measure is equal to 0.04 for "data modeling" and 0.03 for "quantitative analysis". Similarly, Table A.3 shows that the job titles associated with the highest job-level measure of AI-relatedness, $\omega_j^{NarrowAI}$, are all very relevant postings such as "Artificial Intelligence Engineer" (average AI-relatedness measure of 0.476), "Senior Data Scientist - Machine Learning Engineer" (0.367), "AI Consultant" (0.365), and "AI Senior Analyst" (0.354).

As a further validation, it is worth noting the geographic locations of the identified AI jobs. We aggregate firm-level AI investments for all Compustat firms to the commuting zone level and link the commuting-zone-level changes in the share of AI workers from 2010 to 2018 to 2010 commuting zone characteristics from the Census American Community Survey. Figure A.6 displays a heat map of the growth in the average job-level AI relatedness measure from 2010 to 2018 and shows that there is significant variation in AI investments across commuting zones. Figure A.7 examines the sources of this variation. Panel (a) presents a binned scatter plot of the change in the share of AI workers from 2010 to 2018 against the average commuting-zone-level log wage in 2010, with the fitted regression line in red. We find a strong positive relationship between the ex ante average wage of a region and the local future growth in AI jobs—the R-squared of the regression is 0.43.

Figure A.7 (b) demonstrates that the growth in AI workers is also concentrated in commuting zones with a large fraction of college-educated workers. These patterns are intuitive, given that AI employees tend to be high-skilled technologically-oriented workers, and contrast with investments in robotics, which concentrate in areas with larger shares of manufacturing employment (Acemoglu et al., 2020).

Lastly, we observe a high correlation between our two measures of AI investments, which helps address the concern that firms' job postings may not be sufficient to capture firms' actual hiring of AI talent. For example, if a firm is unable to fill AI-related vacancies, the job postings measure will overstate the firm's investments in AI. In practice, this does not appear to be a main driver of the job-postings-based measure, because the two measures of AI investments—using job postings and resumes—are highly correlated and yield similar results throughout the remainder of the paper. Table 1 displays the cross-sectional correlations between the two measures at the firm level for each year when job postings data are available (2007 and 2010–2018). We report correlations for three variable pairs: (i) the absolute number of AI job postings in Burning Glass against the absolute number of AI employees in Cognism; (ii) the fractions of AI employees in the two datasets; and (iii) the fraction of AI employees in Cognism against the average continuous measure of AI-relatedness from BG, $\omega_j^{NarrowAI}$. Panel 1 presents Pearson correlations, while Panel 2 tabulates Spearman (rank) correlations. The correlations are quite high, especially in the recent years: e.g., in 2018, we see a Pearson correlation of 0.82 for the absolute numbers of AI jobs, 0.55 for the fractions, and 0.57 for the Cognism fraction against the continuous Burning Glass measure.

High correlations and consistency across our two measures also address the potential concern that job postings do not reflect firms investing in AI by acquiring other firms (e.g., AI startups). Human capital on-boarded through acquisitions *is* captured by the resume data, where employees of acquisition targets are counted as employees of the acquirer subsequent to the acquisition. The consistency of our results across the two measures (based on job postings and resumes) indicates that on-boarding of AI talent through acquisitions is likely not a main driver of firms' AI investments. This notion is also supported by the smooth increase (without jumps) of most firms' AI workers (representative time series for a few large firms are presented in Internet Appendix A.1) and by industry reports estimating that 90% of firms' investments in AI are internal and 10% is spent on acquisitions (Bughin et al., 2017).

4 Firm-level Determinants of AI Investments

We consider the determinants of investments in AI technologies, and document that larger firms—with more employees, higher sales, and larger market share—see greater levels of AI investments between 2010 and 2018. In addition, firms with higher markups tend to invest in AI more aggres-

sively, as are firms with higher cash reserves and more sizable R&D activities.

In Table 2, we examine which ex ante firm characteristics predict future growth in firm-level AI investments. Since our focus is on understanding the *use* of AI technologies by firms, we exclude firms in the tech sector from our main empirical analyses in this and the following sections. Firms in the tech sector are likely to be inventors of new AI technologies or suppliers of AI solutions, in which case their hiring of AI employees could lead to different dynamics than in other sectors.[20] For each measure of AI investment, we estimate the following specification:

$$\Delta ShareAIWorkers_{i,[2010,2018]} = \beta FirmVariable_{i,2010} + IndustryFE + \epsilon_i, \quad (1)$$

where $\Delta ShareAIWorkers_{i,[2010,2018]}$ denotes the change in the share of firm i's AI-related employees from 2010 to 2018 in the regressions in Panel 1, and the change in the share of firm i's AI-related job postings in Panel 2. All regressions include 2-digit NAICS industry fixed effects. Here and throughout all subsequent analyses, the $\Delta ShareAIWorkers_{i,[2010,2018]}$ variables are standardized to mean zero and standard deviation one to aid in economic interpretation. $FirmVariable_{i,2010}$ represents one of the ex ante firm characteristics of interest measured as of 2010: log firm employment in column 1, firm market share in the corresponding 5-digit NAICS industry in column 2, log sales in column 3, the ratio of cash to assets (Cash/Assets) in column 4, the ratio of R&D expenditures to sales (R&D/Sales) in column 5, return on sales (ROS) measured as the ratio of net income plus interest expense to sales following Fracassi and Tate (2012) in column 6, log markup measured as the log of the ratio of sales to cost of goods sold following De Loecker et al. (2020) in column 7, and log markup measured as the log of the ratio of sales to operating expenses following Traina (2018) in column 8.[21] Column 9 includes all variables in a multivariate specification, except for employment and market share, which are highly correlated with sales. To account for differences in precision in the measurement of AI investments due to the number of observations available to calculate the measure for each firm, the estimating equation is weighted by each firm's number of

[20]The tech sector includes NAICS 2-digit industry codes 51 (Information) and 54 (Professional, Scientific, and Technical Services). In later analyses, we confirm that the main effects of AI spurring firm-level growth are also present and even stronger in these industries, and a complementary treatment of the impact of AI on specifically AI-inventing firms is provided by Alderucci et al. (2020). Our main regression sample is comprised as follows. Among the 2969 non-IT-sector firms in Compustat with U.S. headquarters and positive sales and employment in 2010, 2074 firms are matched to Cognism in both 2010 and 2018, and 1664 firms are matched to Burning Glass in both 2010 and 2018. For the Cognism sample, we further restrict to firms with at least 50 U.S. jobs in both 2010 and 2018 to ensure good coverage of the firm's workforce, which leaves us with approximately 1400 firms. For the Burning Glass sample, we further restrict to firms that are also matched to Cognism so that we can cross-validate with the actual hiring, leaving us with approximately 1200 firms.

[21]We examine log markups, because our subsequent analyses of the impact of AI on firms focus on changes in firm outcomes, and the change in log markups captures the percent change in markups. Log markups are also examined by De Loecker and Warzynski (2012). In De Loecker et al. (2020), firm-level markups are equal to $\mu_{it} = \theta_{st} \frac{Revenue_{it}}{Variable\,Cost_{it}}$, where θ_{st} is the degree of returns to scale in industry s in year t. When taking the logs, $\log \mu_{it} = \log(\theta_{st}) + \log(Revenue_{it}/VariableCost_{it})$, the change in term $\log(\theta_{st})$ is absorbed by industry fixed effects, and therefore we can focus our empirical tests on the change in $\log(Revenue_{it}/VariableCost_{it})$.

resumes (job postings) in 2010.[22]

The results reported in Table 2 highlight that ex ante larger firms experience higher growth in AI investments. For example, using the Cognism-based measure in Panel 1, a one-standard-deviation increase in log sales in 2010 (which equals 2.0) corresponds to the share of AI workers increasing by 26% of the standard deviation from 2010 to 2018, significant at the 1% level. In addition, firms with higher starting Cash/Assets and higher R&D/Sales also see greater investment in AI, consistent with contemporaneous work of Alekseeva et al. (2020). While overall return on sales is not robustly predictive of future AI investments, operating expenses-based markups positively predict future AI investments. Importantly, the results for firm-level demand for AI talent measured with Burning Glass data are very consistent with the results using Cognism data, reinforcing the high correlations documented in Table 1. This consistency suggests that, in the absence of matched employer-employee data, our methodology for identifying AI investments from the Burning Glass data can be a good proxy for firms' actual AI hiring.

5 AI Investment and Firm-level Results

We present our main results on the relationship between AI investments and firm-level outcomes. We document that firms investing in AI technologies grow faster, and the positive effect of AI investments on growth is concentrated among ex ante larger and more productive firms. While AI-investing firms do not see changes in their productivity, they tend to grow by expanding their product offerings and geographic reach.

5.1 Long-differences Results

We begin the analysis by examining whether firms that invest in AI see faster growth from 2010 to 2018. Given that we do not expect to observe the effects of AI investments on firms immediately, we use a long-differences specification. This specification ensures that no time-invariant firm characteristics can drive the results. In Table 3, we report the estimates from the following regression:[23]

$$\Delta FirmVariable_{i,[2010,2018]} = \beta \Delta ShareAIWorkers_{i,[2010,2018]} + Controls'_{i,2010}\gamma + IndustryFE + \epsilon_i, \quad (2)$$

where the main independent variable, $\Delta ShareAIWorkers_{i,[2010,2018]}$, captures the change in the

[22] Since the numbers of worker resumes and job postings are highly correlated with the size of the firm, this weighting scheme also roughly weights firms in accordance to their contribution to the economy. The results are robust to weighting each firm by its sales or employment in 2010.

[23] In the baseline regressions, we include firms that are observed in the Compustat sample both in 2010 and 2018. We consider entry and exit in industry-level results in Section 6.

share of AI workers in firm i from 2010 to 2018, standardized to mean 0 and standard deviation 1. As in Section 4, this analysis excludes firms in the tech sector. $IndustryFE$ are 2-digit NAICS fixed effects. In Panel 1, we report the coefficients for the resume-based measure of AI investments, while Panel 2 considers the job-postings-based measure. In columns 1, 3, and 5 we include only industry fixed effects to examine the unconditional relationship between AI investments and firm growth. In columns 2, 4, and 6, we include a rich set of controls that are all measured at the start of the sample period in 2010: (a) initial firm-level characteristics that predict investment in AI, including the log of the total number of jobs (or job postings), cash/assets, log sales, R&D/Sales, and log markups (both COGS-based and operating-expense-based); and (b) controls for the characteristics of commuting zones where the firms are located (log average wage, the share of college-educated workers, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), as well as the log industry average wage.[24]

In columns 1 and 2 of Table 3, the dependent variable is the firm-level change in log sales from 2010 to 2018.[25] Both measures of AI investment are associated with a significant and economically meaningful increase in sales growth: a one-standard-deviation increase in the share of AI workers over an 8-year period corresponds to an additional 12% to 16% growth in sales, depending on the specification. In columns 3 and 4, we find a positive effect on employment of similar magnitude to the effect on sales: a one-standard-deviation increase in the share of AI workers over an 8-year period predicts an increase of 11% to 16% in firm-level employment. This suggests that AI is not displacing firms' workforces, at least on net and in the short-run, although we do not rule out reallocation of labor across different job functions or tasks. Columns 5 and 6 show that firms investing in AI grow more than their industry peers: a one-standard-deviation increase in the share of AI workers is associated with a 0.9–1.5 percentage point increase in a firm's market share in its 5-digit NAICS industry, although the effect is not always statistically significant. In untabulated analyses, we confirm that the results are similar (with positive and significant effects on the order of 10-15%) when using different measures of firm growth: market value rather than sales and the number of workers in each firm in the Cognism resume data rather than Compustat employment.

[24]When firms span multiple commuting zones, we calculate these variables as the weighted average, using numbers of BG job postings in each commuting zone as weights, which restricts the sample in Cognism regression analysis to firms that are also matched to the Burning Glass data. The results are similar in magnitude and economic significance if we only include firm-level controls enumerated in list (a). We also confirm, in Table A.4, that the results are robust to the addition of (i) state fixed effects to control for all omitted variables at the state level, and (ii) initial Tobin's Q as of 2010, which proxies for the firm's growth opportunities.

[25]Out of the approximately 1400 (1200) firms in the Cognism (Burning Glass) sample in Table 2, 1088 (934) firms have positive sales and employment in 2018, which are used to calculate the dependent variables. We further restrict the sample to firms with non-missing control variables throughout, to keep the sample composition stable. This results in a sample of 766 firms in Cognism and 849 firms in Burning Glass. The results of the regressions without controls are similar when estimated on the entire available sample.

The positive relationship between AI investments and firm growth is ubiquitous across different sectors of the economy, reinforcing the notion that AI is a general purpose technology. Table A.5 shows the results from regressing changes in log sales and log employment on the change in the share of AI workers, separately for the largest 2-digit NAICS sectors: (i) Manufacturing, (ii) Wholesale and Retail Trade, (iii) Finance, and (iv) the remaining non-tech sectors. We also estimate the regression for the two tech sectors—Information and Professional and Business Services—and find that AI has an even stronger positive relationship with growth in tech firms, which is consistent with the tech firms being more advanced in their AI efforts than firms in other sectors (see Table A.6). Overall, we observe that investments in AI are associated with economically significant increases in firms' operations, and these effects are meaningful across key economic sectors.

We also confirm that our results are not sensitive to different approaches of constructing the AI measures. In Table A.7, we show that the results are robust to using firm-level average continuous AI-relatedness measures of job postings, which are defined in Section 3.1. In addition, we also find similar results when using higher or lower cutoffs for classifying AI jobs (Table A.8).

Finally, while our measure is centered on internal AI investments, we also consider whether firms' use of external AI solutions might affect the interpretation of our results. To this point, it is worth mentioning that even external AI software requires internal data management and implementation guidance by AI-skilled workers to be effective (Fedyk, 2016), and industry reports underscore that AI-skilled labor is the most critical input to successful deployment of AI programs (Bughin et al., 2018). Nevertheless, we leverage the rich detail of the Cognism resume data to confirm that our approach of focusing on AI workers to identify AI investments is indeed a suitable one. First, we undertake a deep dive into case studies of individual firms (see Internet Appendix A.1 for examples) and observe that (i) our measure captures internal AI investments well, and (ii) the use of external AI software solutions (e.g., IBM Watson, IPSoft Amelia) tends to be complementary to internal AI hiring, consistent with patterns highlighted by surveys of executives. Second, we process the individual job descriptions and job titles for any mention of external AI software (including IBM Watson Studio, Symphony, AyasdiAI, Salesforce Einstein, and about a hundred other key AI-powered solutions) to construct a proxy for firms' reliance on external AI solutions.[26] In Table A.9, we confirm that our results are robust to directly including this proxy in our overall measure of AI investments.

[26]The Cognism resume data are especially well-suited to capture the use of external technological solutions, given Cognism's emphasis on developing "technographic data" (defined by Cognism as "the technologies that the employee or company is using"). Cognism advertises these data for two purposes: (i) enhancing technology-providers' capabilities in targeted marketing of their products and (ii) improving individual firms' understanding of which technologies are used by their competitors.

5.2 Dynamic Effects

In this section, we use firm-level panel data to examine firm growth dynamically around the year of AI investment. In addition to elucidating the lag between AI investments and their realized effects, this analysis helps to address reverse causality and concerns about AI-investing firms being on differential growth trajectories.

Firms tend to invest in AI on a continuous basis, rather than make a lumpy investment in a single year, which precludes us from examining dynamic effects in a standard event-study framework.[27] Instead, we estimate the dynamic effects of AI investments using a distributed lead-lag model, which allows for continuous variation in the independent variable (Aghion et al., 2020; Stock and Watson, 2015). In particular, the model is specified as:

$$Y_{it} = \sum_{k=-2}^{5} \delta_k \Delta ShareAIWorkers_{i,t-k} + \mu_i + \lambda_{st} + \epsilon_{it} \tag{3}$$

where $\Delta ShareAIWorkers_{i,t-k}$ is the annual change in the share of AI workers from year $t-k-1$ to year $t-k$, normalized to have mean zero and standard deviation of one, and Y_{it} is the outcome variable (either log sales or log employment) in year t. We include firm fixed effects μ_i to absorb firm-specific time-invariant factors and industry-year fixed effects λ_{st} to control for industry-specific trends. Each lead-lag coefficient δ_k captures the cumulative response of the outcome variable in year t (Y_{it}) to AI investments in year $t-k$, holding fixed the path of AI investments in all other years. As such, specification (3) incorporates both immediate and delayed responses of sales or employment to the firms' AI investments. For each firm-year sales and employment observation between 2010 and 2016, we consider five lags and two leads so that we estimate the cumulative impact of AI investments on firm growth from two years before the investments to five years after the investments.[28] The estimated coefficients for the leads can be used as a pre-trend test: if firms investing in AI are on similar growth trends as other firms prior to AI investments, δ_k with $k < 0$ should be statistically indistinguishable from zero.

Figure 3 reports the coefficients from the lead-lag regressions. The top panel shows that sales increase following AI investments. The effect is not immediate, and it takes two to three years for firms to realize the benefits from AI investments. The cumulative effect of a one-standard-deviation increase in annual AI investments on log sales is between 1% and 2% and remains steady five years after the investments. This is consistent with the long-differences estimates in Section

[27] The percentage of AI-investing firms that only invest in a single year is 29.5%, compared to 70.6% for robots (Humlum, 2019).

[28] Since the data on AI investments end in 2018, we include only two leads to keep all firm-year observations up to 2016. We obtain similar results when including only one lead or no leads at all. Furthermore, analysis of dynamic effects focuses on the Cognism resume data because these data offer full coverage of AI investments from 2005 to 2018. By contrast, Burning Glass job postings data have a more limited time series, where including 2 leads and 5 lags would restrict the sample to only firm-year observations in 2015 and 2016.

5.1, with a one-standard-deviation increase in AI investments associated with a 12-16% increase in sales over 8 years. The bottom panel shows that AI investments have a similar positive effect on firms' employment.

Importantly, there is no evidence of pre-trends, which helps to assuage concerns about reverse causality and AI-investing firms being on differential growth trends: conditional on the controls we include, firms that invest more in AI in any given year show comparable sales and employment paths in prior years and start diverging only afterwards. This suggests that our results are not capturing the reverse effect of firm growth on AI investments and restricts potential confounders to the set of shocks that occur simultaneously with AI investments. We further consider and rule out potential alternative explanations for our results in Section 7, where we show that contemporaneous investments in other non-AI technologies do not drive the results, and that the results remain robust to using two different instruments for AI investments.

5.3 Firm-level Heterogeneity

We next examine how the documented effects vary across firms' ex ante characteristics. To understand how AI might change inequality among firms (Mihet and Philippon, 2019), we first examine whether larger or smaller firms benefit more from their AI efforts. On the one hand, AI can reduce the barriers to growth among smaller firms. Unlike proprietary IT systems that require large upfront capital investments, AI implementation is largely dependent on human expertise, with auxiliary costs centering on data storage and computing, which can be purchased from specialized providers on a per use basis and do not require lumpy upfront investments (Organisation for Economic Co-operation and Development, 2015; CorrelationOne, 2019). On the other hand, big data and AI technologies can have scale effects that favor large firms with large amounts of data (Farboodi et al., 2019). Table 4 shows the relationship between AI investments and firm growth, across terciles of firms by employment in 2010 (within the firm's 2-digit NAICS sector), controlling for initial size. For employment, sales, and market share, the effect of AI investments is monotonically increasing in the firm's initial size. For example, using the resume-based measure, the results in column 2 indicate that a one-standard-deviation increase in the share of AI workers is associated with a 17% increase in sales for firms in the top size tercile, a 4% increase for firms in the middle tercile, and no increase for firms in the bottom tercile. The difference in the coefficients between the top and bottom size terciles is statistically significant at the 5% level. Hence, the positive relationship between AI investments and growth concentrates among ex ante larger firms.

We also test whether more or less productive firms experience higher growth associated with their AI investments. Prior literature on technology adoption finds mixed results. Less productive firms in least favorable positions have benefitted most from innovations such as science-driven drug discovery (Cockburn et al., 2000). Yet, recent literature points out that more productive firms

may be able to leverage new technologies better, leading to a phenomenon of "superstar" firms (Autor et al., 2020; Aghion et al., 2019). When we group firms in each 2-digit NAICS industry into terciles based on initial productivity (measured as of 2010), we see that the growth fueled by AI is mostly concentrated among ex ante most productive firms. This analysis is presented in the Internet Appendix, with Table A.10 slicing firms into terciles based on revenue TFP as of 2010, and in unreported results we find a similar pattern based on labor productivity (sales per worker) as of 2010. Overall, these cross-sectional results appear most consistent with increasing firm inequality where AI facilitates scale advantages for the ex ante largest and most productive firms.

5.4 Mechanisms

Finally, we examine the drivers of the AI-fueled growth by considering the two non-mutually-exclusive mechanisms detailed in Section 1. We find no evidence that firm AI investments are associated with increased productivity but do observe that AI-investing firms are able to significantly expand their reach, both in terms of product offerings and geographic scope.

5.4.1 AI as a Driver of Productivity Growth

We explore whether the increase in firm growth from AI investments could be driven by AI technologies making firms more productive. We consider two measures of productivity: sales per worker (or labor productivity) and revenue Total Factor Productivity (TFP).[29] Table 5 shows that the relationship between AI investments and both productivity measures is a precisely-estimated zero. The lack of growth in labor productivity is consistent with the results in Section 5 that AI investments predict similar increases in sales and employment, challenging the view that the current primary use of AI is to replace human tasks and cut down labor costs. Similarly, we also do not observe any evidence of a decrease in costs: the growth in costs of goods sold (COGS) and operating expenses associated with AI investments is similar in magnitude to the growth in firm sales. For example, a one-standard-deviation increase in the share of AI workers measured using the resume data corresponds to a 15% increase in operating expenses when all controls are included (column 8 of Panel 1).

Overall, we do not find evidence that investments in AI make firms more productive, at least in the short term. This speaks to the broader debate on the timing of productivity gains from general purpose technologies: while it has long been argued that the adoption of general purpose technologies leads to delayed productivity benefits (David, 1990; Brynjolfsson et al., 2020), recent

[29] We use standard methodology to calculate the revenue TFP as the residual from regressing log real sales on log employment and log capital controlling for firm fixed effects and year fixed effects: $\log y_{it} = \mu_i + \mu_t + \alpha_s^l \log(l_{it}) + \alpha_s^k \log(k_{it-1}) + \varepsilon_{it}$. The regression is estimated using OLS separately for each industry. The capital stock is constructed using the perpetual inventory method. The TFP measure is specific to Cobb-Douglas production functions, while sales per worker measures labor productivity for more general production functions.

evidence in Fizsbein et al. (2020) shows rapid productivity growth following electrification. In Table A.11, we examine the effect of AI investments during the first half of the period (2010–2014) on productivity growth through 2018 and do not find any significant positive effect. Hence, even with a lag of a few years, AI investments are not yet associated with productivity gains.

5.4.2 AI as a Driver of Expansion into New Products and Markets

Next, we explore the possibility that investments in artificial intelligence allow firms to expand into new products and markets. First, we show that AI-investing firms experience robust increases in product innovations, which is visible in trademarks, patents related to new products, and the hiring of product managers. Trademarks are a good proxy for the creation of new products and services, because trademarks are registered whenever new products or services are ready for commercialization (Hsu et al., 2021). The first two columns in Table 6 present the results from the long-differences regressions for the number of USPTO trademarks.[30] When all controls are included, a one-standard-deviation increase in the share of AI workers in the resume (job postings) data over an 8-year period translates into a 9.1% (17.6%) increase in the firm's number of trademarks. Columns 3 and 4 reveal a similar relationship between AI investments and the number of product patents, which are patents specifically focusing on product innovations.[31] A one-standard-deviation increase over 8 years in the share of AI workers based on resumes (job postings) corresponds to a 27.2% (23.1%) increase in the number of product patents. Importantly, we do not observe any analogous increase in process patents (columns 5 and 6), which reflect process innovations and potential improvements in efficiency; as such, the null result on process patents reinforces the null results on productivity documented in Table 5. In the last two columns of Table 6, we show that AI investments are associated not only with increased product output (trademarks and product patents) but also with greater emphasis on product development within the firm: AI investments are linked to increased job postings for product managers, who align the firm's product capabilities to specific market segments and are responsible for product innovations and servicing.

Second, we examine whether AI investments are associated with the expansion of firm boundaries along two dimensions: (i) the geographic reach of the firm, measured as the number of counties with at least 1% of the firm's job postings in any given year; and (ii) the number of industries at the most granular (6-digit NAICS) level with at least 1% of the firm's job postings. Geographic reach is a proxy for firms' ability to extend their operations into new geographic markets or service

[30] The dependent variable is the change in log(1+trademarks) from 2010 to 2018 so that the regression takes into account firms with zero trademarks in either 2010 or 2018, and the results are also robust to using the inverse hyperbolic sine transformation (i.e., $\ln(x + \sqrt{(1 + x^2)})$). The regression sample is smaller than our baseline sample because not all public firms file trademarks (we include firms with at least one trademark in 2009-2018).

[31] See Ganglmair et al. (2021) for the methodology to distinguish between product patents and process patents. The regression sample is smaller than our baseline sample because not all public firms file patents.

new clients. The results, reported in the first two columns of Table 7 show that a one-standard-deviation increase in the share of AI workers is accompanied by a 7–9% increase in the number of counties that the firm operates in. The test on the number of industries examines whether the product innovations we document in Table 6 reflect AI-investing firms' foray into new industry segments outside of these firms' existing operations. We observe that the change in the number of different industries spanned is mild, at 0 to 2%, and not statistically significant, suggesting that AI investments are less strongly associated with expansion into new industries and that, so far, AI investments are not a strong predictor of firm diversification activities.

Overall, the juxtaposition of positive results on firms' product innovations and geographic expansion with null results on productivity points towards firms utilizing AI to expand product variety and customer reach, without necessarily increasing efficiency or cutting costs in their existing operations.

6 AI Investments and Industry-level Outcomes

To shed light on the potential aggregate effects of AI, we examine the relationship between industry-level variation in AI investments and: (i) industry growth; (ii) industry concentration.

While AI-investing firms grow faster, the gain in industry sales and employment may be zero-sum if the rivalrous nature of some data and the use of AI technologies creates a business-stealing effect on the competitors (Bloom et al., 2013). For example, the negative spillovers have been shown to dominate the positive effect on firms investing in new technologies in the case of robotics, leading to an overall negative effect on aggregate employment (Acemoglu et al., 2020). Hence, signing the aggregate effect is an empirical question. We estimate the following long-differences regression at the industry level:

$$\Delta \ln y_{j,[2010,2018]} = \gamma \Delta ShareAIWorkers_{j,[2010,2018]} + IndustryFE + \epsilon_j \tag{4}$$

where $\Delta \ln y_{j,[2010,2018]}$ is the change in total sales or employment for all Compustat firms (including those that entered the sample after 2010 or exited before 2018) in 5-digit NAICS industry j, and $\Delta ShareAIWorkers_{j,[2010,2018]}$ is the change in the share of AI workers among Compustat firms in industry j from 2010 to 2018. Analogously to the firm-level tests, the regressions are weighted by the total number of resumes (or job postings) in each industry in 2010.

Columns 1–4 of Table 8 show that AI investments are associated with a robust increase in employment and sales at the industry level. In both panels (Panel 1: resume-based AI measure, and Panel 2: job-postings-based AI measure), odd columns estimate the unconditional relationship (with 2-digit NAICS fixed effects only), and even columns add controls for log employment, log sales, and log average wages in 2010. A one-standard-deviation increase in the industry-level

share of AI workers in the resume data is associated with a 15% increase in sales and a 17% increase in employment (columns 2 and 4). We also evaluate potential spillovers on firms outside of the Compustat sample. In Table A.12 we estimate the impact of AI investments by Compustat firms on the aggregate industry-level employment of all (public and private) firms from the Census Quarterly Workforce Indicators (which does not have sales data). The net effect is approximately zero, indicating that although there may be some negative spillovers from public to private firms, these negative spillovers do not dominate and, contrary to robots, AI is not associated with reduced industry employment.[32]

We next examine whether the higher AI-fueled growth of the largest firms is substantial enough to translate into increased industry concentration. This test would speak to whether, as an intangible asset, AI investments could be a mechanism contributing to the increased industry concentration. We link industry-level growth in AI investments to contemporaneous changes in industry concentration from 2010 to 2018. Following Autor et al. (2020), we use the Herfindahl-Hirschman Index (HHI) to measure industry concentration. To examine winner-take-most dynamics, we also consider the fraction of sales accruing to the largest firm in each 5-digit NAICS industry among the Compustat firms. Columns 5–8 of Table 8 show a positive relationship between changes in industry concentration and industry-level AI investments. For example, in columns 6 and 8, which include all controls, a one-standard-deviation change in the industry share of AI workers measured using the resume data corresponds to a 2.2 percentage point increase in the HHI and a 1.7 percentage point increase in the market share accruing to the top firm.

To enhance the economic interpretation of these industry-level results, we examine whether AI-investing firms increase their product market power, which would speak to potential anti-competitive effects of AI. To test this hypothesis, we look at the effect of AI investments on firm-level markups. We consider three measures of markups: the first two measures are based on De Loecker et al. (2020) and Traina (2018), with log markups defined as the log of revenues divided by Cost of Goods Sold (COGS) and Operating Expenses, respectively. For robustness, our third measure of markups is the operating profit rate, termed the "Lerner Index" (Gutiérrez and Philippon, 2017), which is defined as operating income before depreciation and amortization (OIBDA) minus depreciation, scaled by sales.[33] Internet Appendix Table A.13 shows that for all

[32] A caveat with these results is that the Compustat sample assigns each firm to a single main industry, even for firms that might have operations in several industries. This caveat is unlikely to affect the interpretation of our results. First, using US Census microdata, prior research shows that for a typical US public firm the large majority of its operations fall within one main industry (Babina, 2020). Second, our results in Section 5.4 indicate that firms with larger AI investments do not experience increased industrial diversification.

[33] Our measures of markups are robust to a wide range of production functions with constant returns to scale. However, the first two markup measures may not reflect market power accurately if AI investments change firms' production functions. Empirically, the possibility that AI changes the production function, for example through shifts from variable to fixed costs (De Ridder, 2019), would bias both markup measures in favor of detecting increased markups and against the null results we document. Regardless of the potential changes in the production function, the Lerner Index offers an accounting-based measure of monopoly power enjoyed by firms.

three measures, the relationship between AI investments and markups is statistically insignificantly different from zero. This is in line with the significant increase in variable costs (COGS or Operating Expenses) reported in Table 5. Overall, the results on markups are inconsistent with the explanation that AI increases the market power of firms, at least in the short run.

Overall, our results are consistent with Crouzet and Eberly (2019) and Ayyagari et al. (2019), who argue that investments in intangible assets are responsible for the rise in industry concentration observed in the U.S. data. Our results suggest that, as a general purpose technology that can be applied across many industries, AI has the potential to further increase concentration across a broad range of industries. However, instead of increasing market power, the observable effects of AI investments to date have been on expansion into new products and geographic markets for the largest and most efficient firms.

7 Robustness

In this section, we consider potential alternative explanations for our main results. First, we show that the effect of AI investments is not driven by concurrent investments in related technologies.Second, an important concern is that firms investing in AI are experiencing other positive shocks at the same time. We use two instrumental variable strategies and show that unobserved firm and industry shocks are unlikely to drive our results.

7.1 Investments in Related Technologies

In Section 3.1, we confirm that our methodology for identifying AI human capital captures employees whose work is specifically related to AI. However, it is possible that hiring of AI human capital happens as part of broader expansion in firms' technological infrastructure, so that our measure is a proxy for general tech investments rather than AI specifically. Our rich data allow us to investigate this possibility: we develop measures of non-AI IT, robot, and non-AI data analytics investments from the job postings data which parallel the measure of AI investments: for each firm, we measure the percentage of job postings in each year requiring IT-, robotics- or data-related skills that are not related to AI.

First, we note that AI investments are not strongly correlated with concurrent non-AI IT investments. For example, in the Cognism data, the correlation between the change in the share of AI workers and the change in the share of IT workers between 2010-2018 is 0.02, and the correlation between the change in the share of AI workers and the change in the share of robotics workers is 0.10. This is perhaps not surprising: IT and robotics investments have been ongoing for several decades, while most AI investments are made within the last decade. Furthermore, the effects that we document for AI investments are distinct from those previously shown for general IT: previ-

ous IT investments are usually associated with productivity gains and benefit smaller firms more (Hobijn and Jovanovic, 2001; Cardona et al., 2013), whereas we find that AI investments are not associated with significant productivity changes and benefit larger firms more.

Second, in Table A.14, we explicitly test whether our measure of AI investments captures investments in other technologies by controlling for other investments. In the first four columns, we estimate the relationship between AI investments and firm-level growth, controlling for (i) investments in robots and (ii) investments in other information technologies (IT). The estimated relationship between AI investments and firm growth remain similar with the addition of these controls.

A related alternative story is that AI investments reflect more general investments in data or predictive technologies. To investigate this hypothesis, we use the Burning Glass job postings data to identify data-related skills (such as "Data Cleaning") and analysis-related skills (such as "SAS") and the share of jobs that require these skills but are not classified as AI-specific. In columns 5 to 8 of Table A.14, we rerun our main firm-level analysis with an additional control for firms' contemporaneous investments in non-AI data or data analysis technologies. The effect of AI investments remains unchanged with the inclusion of these controls. In unreported analyses, we also confirm that the results are similar when controlling for the initial share of robot-, IT-, or data-analysis-related workers in 2010. Our findings are consistent with Abis and Veldkamp (2020), who show that AI and machine learning are more scalable than "old technologies" and can fully unleash the potential from data. Overall, these results confirm that the positive effect on firm growth is specifically driven by AI rather than other technologies.

7.2 Firm-specific Shocks to Demand and Growth Opportunities

Although our long-differences specification controls for time-invariant firm-specific characteristics, and we find no pre-trends in the analysis of dynamic effects, our estimates of the impact of AI may still be biased if there are unobserved shocks that occur simultaneously with AI investments and correlate with firm growth. Specifically, firms with positive demand shocks may invest more in AI, leading to a spurious positive relationship between AI investments and firm growth.

To address this concern, we instrument for a firm's change in the share of AI workers from 2010 to 2018 using the contemporaneous change in the share of AI workers of European firms in the same 5-digit NAICS industry. This instrument is based on prior work by Autor et al. (2013) and Acemoglu and Restrepo (2020). Since the instrument is at the industry level, the IV estimates are not confounded by firm-specific shocks to demand or growth opportunities. The identifying assumption is that the common within-industry component of the rising share of AI workers is due to supply-side related differences in data availability and technological feasibility across industries, rather than demand shocks. This is supported by industry reports, such as Bughin et

al. (2017), which argue that AI adoption across industries is largely driven by the availability of data and technical capabilities, which are crucial inputs for implementing AI investments.

We construct the instrument using all firms that are listed on stock exchanges in Europe.[34] We focus on European firms for two reasons. First, similar to the U.S., Europe has experienced a surge in AI investment in recent years. Figure A.8 shows that Europe has the largest increase in the share of AI workers between 2010 and 2018 outside of the U.S. Second, while it is generally difficult to obtain consistent firm-level measures of AI investments in the U.S. and abroad, Cognism has good coverage of employee resumes in Europe.

The results are reported in Table A.15. As in long-differences OLS regressions, we control for 2-digit NAICS industry fixed effects in all columns, and even columns include the set of full controls. Columns 1 and 2 show that the instrument has a strong first stage for both AI measures, with F-statistics ranging from 19 to 32. We report the instrumented effect of AI investments on the change in log sales in columns 3 and 4 and on the change in log employment in columns 5 and 6. We do not look at the changes in market share in this set of IV results, since the instrument is based on industry-level variation. In all specifications and for both measures of AI investments, we find positive and significant effects of AI investments on firm growth, which are larger than, although not statistically significantly different, from the long-differences estimates.

The main threat to the validity of our instrument is that industries with higher rates of AI investment (in the United States and Europe) experience a coincidental upward trend because of, for example, positive demand shocks. Figure A.9 tests whether AI investments are correlated with common demand shocks by plotting the growth in industry-level AI investments in Europe against U.S. industry-level price changes, which should be correlated with industry-level demand shocks. We document a weak negative correlation, pointing against common demand shocks in AI-investing industries driving the IV estimates. We also observe no pre-trends in firm growth in Table A.16, where we regress firm-level changes in sales and employment from 1999 to 2007 on subsequent changes in AI investments from 2010 to 2018.[35]

7.3 Confounding Industry Shocks

A remaining concern is that our results reflect *industry*-level shocks that happen contemporaneously with AI investments in both US and Europe, leading to a spurious correlation between AI investments and firm growth even within the foreign industry IV specification. This alternative explanation relies on strong assumptions: industry-level shocks would have to occur at the same time as AI investments given the lack of observable pre-trends and would need to correlate with

[34] To ensure that there is no overlap between the endogenous variable and the IV, we exclude European workers of U.S. multinational firms that are listed in the U.S. and the non-European workers of European firms.

[35] We look at 1999-2007 to avoid the confounding effect of the financial crisis between 2008 and 2010. We find similar results when using 2000-2008.

firm growth opportunities in the U.S. *without* correlating with price changes. Nonetheless, we address this alternative explicitly by using a Bartik-style instrumental variables strategy, which does not rely on industry-level variation, and instead leverages pre-existing geographic variation in co-located businesses to capture firms' exposure to the local supply of AI labor.

For each firm, we compute a weighted average of AI investments across industries that operate at the pre-period locations of the firm's operations, with weights proportional to the local industry employment shares. Specifically, let s_{ic} denote the share of job postings of firm i that falls within commuting zone c. Let θ_{cj} be the share of commuting zone c's employment that is mapped to (5-digit NAICS) industry j in 2010, obtained from the County Business Patterns data. Our instrument for firm i's change in the share of AI workers from 2010 to 2018 ($\Delta ShareAIWorkers_{i,[2010,2018]}$) is:

$$\sum_c s_{ic} \left(\sum_j \theta_{cj} \Delta AI_{j,[2010,2018]} \right),$$

where $\Delta AI_{j,[2010,2018]}$ is industry j's change in the share of AI workers from 2010 to 2018.

This instrument is conceptually similar to a Bartik-style shift-share instrument (Bartik, 1991), with one important difference. Whereas a canonical Bartik instrument is based on industries in which a firm operates, our instrument is based on the industries geographically close to the firm. For example, if the financial industry adopts AI technologies heavily and attracts AI skilled labor into New York City, then a healthcare company in New York may be more likely to adopt AI than a healthcare company in DC because of proximity to the financial industry. The IV compares firms in the same industry across different locations, and is therefore not confounded by industry-wide shocks. The identifying assumption is that conditional on local characteristics (such as share of college-education workers and average wage), a firm's geographic proximity to industries that subsequently invest in AI is pre-determined and not correlated with changes in firm outcomes in ways other than AI investments. Goldsmith-Pinkham et al. (2020) suggest that one way to test the plausibility of this assumption is to check whether there are pre-trends before the shocks. We test for pre-trends in Table A.18, and find that there is no relationship between future shocks to instrumented AI investments and past firm growth.

Table A.17 reports the results from the shift-share IV for both of our measures of AI investments. Columns 1 and 2 display the results from the first stage regressions, while columns 3 through 8 show the second stage results for the changes in firm size. The results show a robust and significant effect of AI investments on sales and a positive effect on employment that is significant in 3 out of 4 specifications.

Overall, the results from the shift-share IV are consistent with the OLS results and the foreign industry IV results: all three analyses show that AI investments predict robust and economically meaningful growth at the firm-level. Moreover, the magnitudes of the effects across the two IV

strategies are similar, which are based on different variation and identifying assumptions. This combination suggests that alternative explanations based on firm- or industry-level shocks are unlikely to be driving our results.

8 Conclusion

In this paper, we study how firms invest in and benefit from one of the most important new technologies of the last decade–artificial intelligence. We introduce a novel measure of investments in AI technologies at the firm level using two detailed datasets on human capital: job postings from Burning Glass Technologies, which indicate each firm's demand for particular skills, and resume data from Cognism, which reveal the actual composition of a firm's workforce. Our unique measure allows us to examine both the determinants and consequences of AI investments by firms across a wide range of sectors. We find a positive feedback loop between AI investments and firm size: AI investments concentrate among the largest firms, and as firms invest in AI, they grow larger, gaining sales, employment, and market share. This AI-fueled growth does not appear to stem from increased productivity; instead, AI-investing firms expand by increasing their product offerings and geographic reach.

Our findings highlight important differences between the adoption of AI technologies and the adoption of information technology (IT) in the 1980s and 1990s.[36] Much of the previous literature finds that IT investments were associated with economically large productivity increases but mixed results on firm performance measures such as market share. By contrast, we find that AI-investing firms see gains in market share but not in productivity. Our results also show higher AI adoption and larger gains from AI investments for larger firms, which contrasts with prior work on diffusion patterns for IT (Hobijn and Jovanovic, 2001). These differences underscore the distinctive features of AI relative to previous waves of IT: as a general purpose technology with non-rival uses across different operations of the firm, AI facilitates expansion into new markets, and the reliance of AI on data confers a special advantage to larger firms.

Our industry-level results point to new technologies, such as AI, being an important factor contributing to the increase in industry concentration and the rise of "superstar" firms documented in recent work (Gutiérrez and Philippon, 2017; Autor et al., 2020). We find little evidence of higher market power of AI-investing firms ("bad" concentration) and instead see that AI allows the most efficient firms to overcome supply-side constraints and expand their scale. Further understanding how AI affects production processes and corporate strategies and business models of firms and assessing the distribution of gains from investing in AI technologies across firms and workers are fruitful avenues for future research.

[36] See Dedrick et al. (2003) and Cardona et al. (2013) for a review of that literature.

References

Abis, Simona and Laura Veldkamp, "The changing economics of knowledge production," *Working Paper*, 2020.

Acemoglu, Daron and Pascual Restrepo, "Automation and new tasks: How technology displaces and reinstates labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and jobs: Evidence from U.S. labor markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **, Claire LeLarge, and Pascual Restrepo**, "Competing with robots: Firm-level evidence from France," *AEA Papers and Proceedings*, 2020, *110*, 383–388.

_ **, David Autor, Jonathon Hazell, and Pascual Restrepo**, "AI and Jobs: Evidence from Online Vacancies," 2021.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li, "A theory of falling growth and rising rents," *NBER Working Paper w26448*, 2019.

_ **, Benjamin F. Jones, and Charles I. Jones**, "Artificial intelligence and economic growth," *NBER Working paper w23928*, 2017.

_ **, Céline Antonin, Simon Bunel, and Xavier Jaravel**, "What are the labor and product market effects of automation? New evidence from France," *CEPR Discussion Paper No. DP14443*, 2020.

Agrawal, Ajay, John McHale, and Alex Oettl, "Finding needles in haystacks: Artificial intelligence and recombinant growth," *NBER Working Paper w24541*, 2020.

_ **, Joshua S. Gans, and Avi Goldfarb**, "Artificial intelligence: the ambiguous labor market impact of automating prediction.," *Journal of Economic Perspectives*, 2019, *33*, 31–50.

Alderucci, Dean, Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas, "Quantifying the impact of AI on productivity and labor demand: Evidence from U.S. Census microdata," *Working paper*, 2020.

Alekseeva, Liudmila, José Azar, Mireia Gine, Sampsa Samila, and Bledi Taska, "The demand for AI skills in the labor market," *CEPR Discussion Paper No. DP14320*, 2020.

Aubry, M, R Kräussl, G Manso, and C Spaenjers, "Biased Auctioneers," Technical Report, Working Paper 2020.

Autor, David, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

_ , _ **, Lawrence F Katz, Christina Patterson, and John Van Reenen**, "The fall of the labor share and the rise of superstar firms," *Quarterly Journal of Economics*, 2020, *Forthcoming*.

Ayyagari, Meghana, Asli Demirguc-Kunt, and Vojislav Maksimovic, "The rise of star firms: Intangible capital and competition," *Working Paper*, 2019.

Babina, Tania, "Destructive creation at work: How financial distress spurs entrepreneurship," *Review of Financial Studies*, 2020, *33* (9), 4061–4101.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw, "How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills," *Quarterly Journal of Economics*, 2007, *122* (4), 1721–1758.

Bartik, Timothy, "Who benefits from state and local economic development policies?," *W. E. Up-*

john Institute For Employment Research, 1991.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying technology spillovers and product market rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The productivity J-curve: How intangibles complement general purpose technologies," *NBER Working Paper w25148*, 2020.

Bughin, Jacques, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlström, Nicolaus Henke, and Monica Trench, "Artificial intelligence the next digital frontier?," *McKinsey and Company Global Institute*, 2017.

_ , Jeongmin Seong, James Manyika, Michael Chui, and Raoul Joshi, "Notes from the AI frontier: Modeling the impact of AI on the world economy," *McKinsey and Company Global Institute*, 2018.

Cardona, Melisande, Tobias Kretschmer, and Thomas Strobel, "ICT and productivity: conclusions from the empirical literature," *Information Economics and Policy*, 2013, *25* (3), 109–125.

Carnevale, Anthony P., Tamara Jayasundera, and Dmitri Repnikov, "Understanding online job ads data," *Georgetown University, Center on Education and the Workforce, Technical Report (April)*, 2014.

Chase, Charles W., "Using big data to enhance demand-driven forecasting and planning," *Journal of Business Forecasting: Methods & Systems*, 2013, *32*, 27–34.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern, "The impact of artificial intelligence on innovation," *NBER Working Paper w24449*, 2018.

Cockburn, Iain M, Rebecca M Henderson, and Scott Stern, "Untangling the origins of competitive advantage," *Strategic management journal*, 2000, *21* (10-11), 1123–1145.

Coleman, Braiden, Kenneth J Merkley, and Joseph Pacelli, "Man Versus Machine: A Comparison of Robo-Analyst and Traditional Research Analyst Investment Recommendations," *Available at SSRN 3514879*, 2020.

Corrado, Carol, Jonathan Haskel, Cecilia Jona-Lasinio, and Massimiliano Iommi, "Intangible investment in the E.U. and U.S. before and since the great recession and its contribution to productivity growth," *EIB Working Papers No. 2016/08*, 2016.

CorrelationOne, *Future of Data Talent, 2019 Annual Report* 2019.

Covarrubias, Matias, Germán Guriérrez, and Thomas Philippon, "From good to bad concentration? U.S. industries over the past 30 years," *NBER Working paper w25983*, 2019.

Crouzet, Nicolas and Janice C Eberly, "Intangibles, Investment, and Efficiency," *AEA Papers and Proceedings*, 2018, *108*, 426–31.

_ and _ , "Understanding weak capital investment: The role of market concentration and intangibles," *NBER Working Paper w25869*, 2019.

DAcunto, Francesco, Nagpurnanand Prabhala, and Alberto G Rossi, "The promises and pitfalls of robo-advising," *Review of Financial Studies*, 2019, *32* (5), 1983–2020.

David, Paul, "The dynamo and the computer: an historical perspective on the modern productivity paradox," *AEA Papers and Proceedings*, 1990, *80* (2), 355–61.

Davis, Steven J., John Haltiwanger, Ron Jarmin, Javier Miranda, Christopher Foote, and Eva Nagypal, "Volatility and dispersion in business growth rates: Publicly traded versus privately held firms," *NBER Macroeconomics Annual*, 2006, *21*, 107–179.

_ , _ , _ , _ , _ , **and** _ , "How valuable Is FinTech innovation?," *Review of Financial Studies*, 2019, *32*, 2062–2106.

Dedrick, Jason, Vijay Gurbaxani, and Kenneth L Kraemer, "Information technology and economic performance: A critical review of the empirical evidence," *ACM Computing Surveys (CSUR)*, 2003, *35* (1), 1–28.

Eisfeldt, Andrea L. and Dimitris Papanikolaou, "Organization capital and the cross-section of expected returns," *Journal of Finance*, 2013, *68*, 1365–1406.

Elsby, Michael WL, Bart Hobijn, and Ayşegül Şahin, "The decline of the U.S. labor share," *Brookings Papers on Economic Activity*, 2013, *2013* (2), 1–63.

Erel, Isil, Léa Stern, Chenhao Tan, and Michael S Weisbach, "Selecting Directors Using Machine Learning," *Review of Financial Studies*, Forthcoming.

European Commission, "White paper on Artificial intelligence: A European approach to excellence and trust," 2020.

Executive Office of the President, "Maintaining American leadership in artificial intelligence," *Executive Order 13859*, 2019.

Farboodi, Maryam, Roxana Mihet, Thomas Philippon, and Laura Veldkamp, "Big data and firm dynamics.," *AEA Papers and Proceedings*, 2019, *109*, 38–42.

Fedyk, Anastassia, "How to tell if machine learning can solve your business problem," *Harvard Business Review*, 2016.

_ **and James Hodson**, "Trading on talent: Human capital and firm performance," *Working paper*, 2019.

Fizsbein, Martin, Jeanne Lafortune, Ethan Lewis, and Jose Tessada, "Electrifying? How new technologies impact productivity and jobs," *Working Paper*, 2020.

Fracassi, Cesare and Geoffrey Tate, "External networking and internal firm governance," *Journal of Finance*, 2012, *67* (1), 153–194.

Frank, Morgan R., David Autor, James E. Bessen, Erik Brynjolfsson, Manuel Cebrian, David J. Deming, Maryann Feldman, Matthew Groh, José Lobo, Esteban Moro, Dashun Wang, Hyejin Youn, and Iyad Rahwan, "Toward understanding the impact of artificial intelligence on labor," *Proceedings of the National Academy of Sciences*, 2019, *116* (12), 6531–6539.

Furman, Jason and Robert Seamans, "AI and the economy," *Innovation Policy and the Economy*, 2019, *19*, 161–191.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, "Predictably unequal? the effects of machine learning on credit markets," *The Effects of Machine Learning on Credit Markets (October 1, 2020)*, 2020.

Ganglmair, Bernhard, W. Keith Robinson, and Michael Seeligson, "The Rise of Process Claims: Evidence from a Century of U.S. Patents," *Working paper, University of Mannheim*, 2021.

Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis, "Could machine learning be a general-purpose technology? Evidence from online job postings," *Working Paper*, 2019.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110*, 2586–2624.

Graetz, Georg and Guy Michaels, "Robots at work," *Review of Economics and Statistics*, 2018, *100*, 753–768.

Grennan, Jillian and Roni Michaely, "Artificial intelligence and the future of work: Evidence from analysts," *Working Paper*, 2019.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely, "Are US industries becoming more concentrated?," *Review of Finance*, 2019, *23* (4), 697–743.

Gutiérrez, Germán and Thomas Philippon, "Declining competition and investment in the U.S.," *NBER Working Paper w23583*, 2017.

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

Hoberg, Gerard and Gordon Phillips, "Text-based network industries and endogenous product differentiation," *Journal of Political Economy*, 2016, *124* (5), 1423–1465.

_ **and Vojislav Maksimovic**, "Product life cycles in corporate finance," *Available at SSRN 3182158*, 2019.

Hobijn, Bart and Boyan Jovanovic, "The information-technology revolution and the stock market: Evidence," *American Economic Review*, 2001, *91.*

Hsieh, Chang-Tai and Esteban Rossi-Hansberg, "The industrial revolution in services.," *NBER Working Paper w25968*, 2019.

Hsu, Po-Hsuan, Dongmei Li, Qin Li, Siew Hong Teoh, and Kevin Tseng, "Valuation of new trademarks," *Management Science*, 2021.

Humlum, Anders, "Robot adoption and labor market dynamics," *Working Paper*, 2019.

Jansen, Mark, Hieu Nguyen, and Amin Shams, "Human vs. Machine: Underwriting Decisions in Finance," *Fisher College of Business Working Paper*, 2020, (2020-03), 019.

Jones, Charles I. and Christopher Tonetti, "Nonrivalry and the economics of data," *American Economic Review*, September 2020, *110* (9), 2819–58.

Knight, Will, "The U.S. Leads in Artificial Intelligence, but for How Long?," *MIT Technology Review*, 2017.

Korinek, Anton and Joseph E. Stiglitz, "Artificial intelligence and its implications for income distribution and unemployment," *NBER Working paper w24174*, 2017.

Liu, Ernest, Atif Mian, and Amir Sufi, "Low interest rates, market power, and productivity growth," 2020.

Loecker, Jan De and Frederic Warzynski, "Markups and firm-Level export status," *American Economic Review*, 2012, *102* (6), 2437–71.

_ **, Jan Eeckhout, and Gabriel Unger**, "The rise of market power and the macroeconomic implications," *Quarterly Journal of Economics*, 2020, *Forthcoming.*

McCarthy, John, Marvin L Minsky, Nathaniel Rochester, and Claude E Shannon, "A proposal for the Dartmouth summer research project on artificial intelligence (1955)," *Reprinted online at http://www-formal. stanford. edu/jmc/history/dartmouth/dartmouth. html*, 2018.

Mihet, Roxana and Thomas. Philippon, "The economics of big data and artificial intelligence," *Disruptive Innovation in Business and Finance in the Digital World (International Finance Review)*, 2019, *20*, 29–43.

Mou, Xiaomin, "Artificial intelligence: Investment trends and selected industry uses," *International Finance Corporation*, 2019.

Organisation for Economic Co-operation and Development, *OECD digital economy outlook 2015*

2015.

_ , "Artificial intelligence in society," *OECD Publishing: Paris*, 2019.

Peters, Ryan and Lucien A. Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, *123*, 251–272.

Ridder, Maarten De, "Market power and innovation in the intangible economy," *Working paper*, 2019.

Ringel, Michael, Florian Grassl, Ramón Baeza, Derek Kennedy, Michael Spira, and Justin Manly, "The Rise of AI, Platforms, and Ecosystems," *BCG Most Innovative Companies Report*, 2019.

Seamans, Robert and Manav Raj, "AI, labor, productivity and the need for firm-level data," *NBER Working paper w24239*, 2018.

Stock, James H and Mark W Watson, "Introduction to econometrics (3rd updated edition)," *Age (X3)*, 2015, *3* (0.22).

Tanaka, Mari, Nicholas Bloom Bloom, Joel M. David, and Maiko Koga, "Firm performance and macro forecast accuracy.," *Journal of Monetary Economics*, 2019, *forthcoming*.

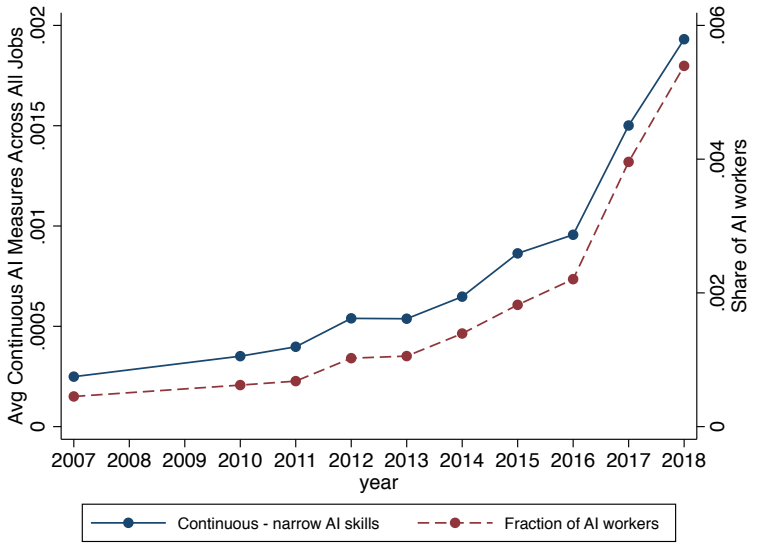
Traina, James, "Is Aggregate market power increasing? Production trends using financial statements," *Working Paper*, 2018.

Webb, Michael, "The impact of artificial intelligence on the labor market," *Working Paper*, 2020.

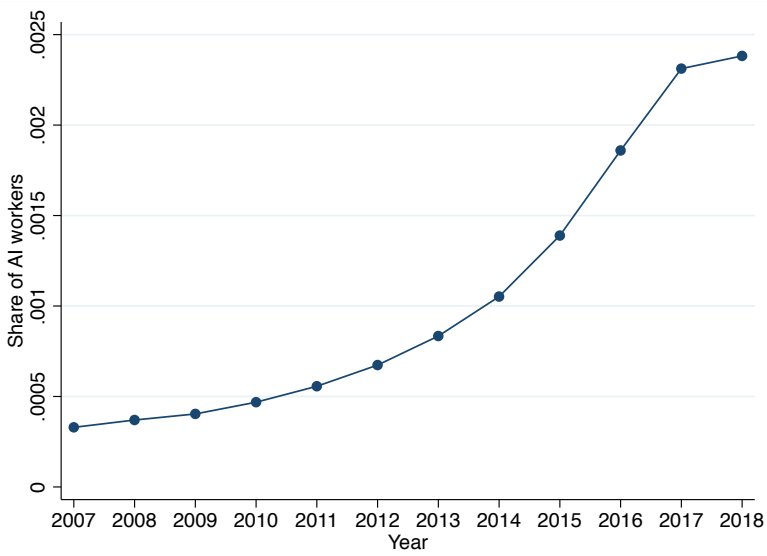
Zator, Michal, "Digitization and automation: Firm investment and labor outcomes," *Working Paper*, 2019.

Figure 1. Time Series of the AI Investments Measure

This figure shows the time series of the two measures of AI investments. Panel (a) reports averages for 2007 and 2010-2018, based on the job postings in Burning Glass with employer firms matched to Compustat. The solid line shows the average job-level continuous measure based on narrow AI skills ($\omega_j^{NarrowAI}$) (left y-axis), and the dashed line tracks the fraction of jobs with continous narrow AI measure above 0.1 (right y-axis). Panel (b) shows the fraction of all employees (across all firms) in a given year who are classified as holding directly AI-related positions in the Cognism resume data from 2007 to 2018.



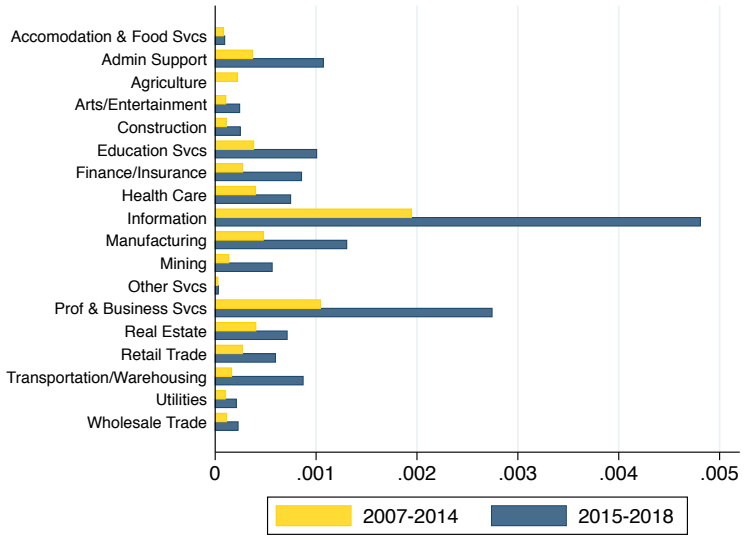
(a) Job posting data (Burning Glass)



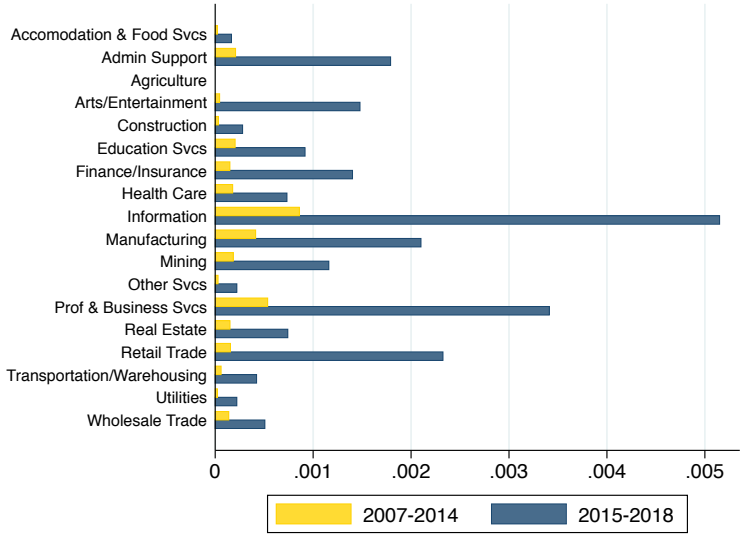
(b) Resume data (Cognism)

Figure 2. AI Investments by Industry Sector

This figure presents the two measures of AI investments at the industry level (based on the sample of public firms). For each sector (based on NAICS–2 digit industry codes), we compute the average job-level continuous narrow AI measure across all job postings (in Panel (a)) and the fraction of AI employees (in Panel (b)) for all firms in that sector across two sub-periods: 2007–2014 and 2015–2018.



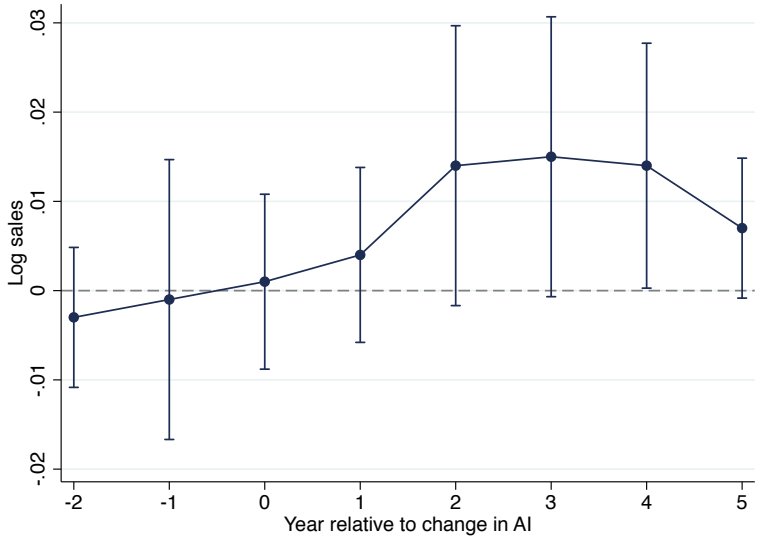
(a) Job posting data (Burning Glass)



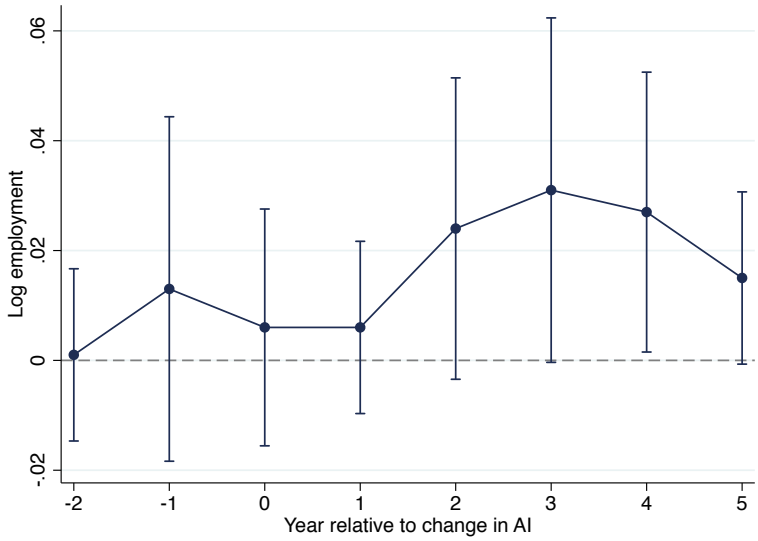
(b) Resume data (Cognism)

Figure 3. AI Investments and Firm Growth Over Time

This figure plots the coefficients from the distributed lead-lag model. The dependent variable is annual log sales in Panel (a) and log employment in Panel (b). The independent variable is the annual change in the share of AI workers in Cognism resume data, standardized to have mean zero and standard deviation of one. Regressions include firm-level sales (or employment) observations between 2010 and 2016 and control for firm fixed effects and 2-digit NAICS industry-by-year fixed effects. Regressions are weighted by the number of workers in Cognism resume data. Standard errors are clustered at the 5-digit NAICS level.



(a) Sales



(b) Employment

Table 1. Correlations between Job-posting-based and Resume-based AI Measures

This table reports correlations between three pairs of firm-level variables: (i) the absolute number of AI job postings in Burning Glass against the absolute number of AI employees in resumes from Cognism; (ii) the fraction of employees classified as AI-related in the two datasets; and (iii) the fraction of AI employees in Cognism against the average continuous narrow-AI measure in Burning Glass. Panel 1 shows the Pearson correlation, and Panel 2 displays the Spearman rank correlation, with both correlations computed over the cross-section of firms with at least 50 total employees in the Cognism resume data in each year of the sample.

Panel 1: Pearson correlation

	Correlations between:		
Year	Numbers of AI jobs	Fractions of AI Jobs	Cognism fraction & BG $share^{NarrowAI}$
2007	0.651	0.127	0.181
2010	0.888	0.192	0.292
2011	0.783	0.113	0.242
2012	0.830	0.190	0.440
2013	0.808	0.325	0.405
2014	0.777	0.102	0.232
2015	0.803	0.422	0.502
2016	0.694	0.460	0.522
2017	0.715	0.426	0.560
2018	0.818	0.547	0.567

Panel 2: Spearman correlation

	Correlations between:		
Year	Numbers of AI jobs	Fractions of AI Jobs	Cognism fraction & BG $share^{NarrowAI}$
2007	0.381	0.342	0.288
2010	0.400	0.370	0.358
2011	0.410	0.376	0.336
2012	0.357	0.322	0.326
2013	0.458	0.409	0.383
2014	0.504	0.459	0.418
2015	0.514	0.449	0.445
2016	0.554	0.489	0.457
2017	0.608	0.525	0.510
2018	0.599	0.521	0.513

Table 2. Firm-level Determinants of AI Investments

This table reports the coefficients from regressions of cross-sectional changes in AI investments by U.S. public firms (in non-tech sectors) from 2010 to 2018 on the following ex-ante firm characteristics measured in 2010: log firm employment in column 1, market share within the 5-digit NAICS industry in column 2, log sales in column 3, Cash/Assets in column 4, R&D/Sales in column 5, return on sales in column 6, log markup measured following De Loecker et al. (2020) in column 7, and log markup measured following Traina (2018) in column 8. In Panel 1, the dependent variable is the growth in the share of AI workers from 2010 to 2018 using the resume data from Cognism. In Panel 2, the dependent variable is the growth in the share of AI workers from 2010 to 2018 using the job posting data from Burning Glass. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. The dependent variable is normalized to have a mean of zero and a standard deviation of one. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Share of AI Workers, 2010–2018								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Employment 2010	0.071*								
	(0.043)								
Market Share 2010		0.443*							
		(0.234)							
Log Sales 2010			0.131***						0.137***
			(0.045)						(0.031)
Cash/Assets 2010				2.932***					3.017***
				(0.729)					(0.675)
R&D/Sales 2010					2.845**				2.060*
					(1.206)				(1.053)
ROS 2010						1.171**			0.090
						(0.525)			(0.319)
Log Markup (COGS) 2010							0.429*		-0.387**
							(0.220)		(0.165)
Log Markup (Total Exp) 2010								1.388**	1.833**
								(0.635)	(0.775)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.126	0.124	0.166	0.236	0.149	0.140	0.147	0.171	0.372
Observations	1,361	1,412	1,410	1,412	1,410	1,377	1,409	1,410	1,376

Panel 2: AI measure from job postings data

	Δ Share of AI Workers, 2010–2018								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Employment 2010	0.032								
	(0.053)								
Market Share 2010		0.495*							
		(0.264)							
Log Sales 2010			0.146***						0.168***
			(0.042)						(0.041)
Cash/Assets 2010				3.227***					3.006**
				(1.235)					(1.295)
R&D/Sales 2010					0.715*				2.852***
					(0.382)				(1.034)
ROS 2010						0.748*			1.125*
						(0.433)			(0.657)
Log Markup (COGS) 2010							0.340		-0.215
							(0.220)		(0.312)
Log Markup (Total Exp) 2010								1.220***	1.398**
								(0.466)	(0.677)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.124	0.133	0.150	0.216	0.128	0.140	0.134	0.154	0.317
Observations	1,171	1,202	1,200	1,202	1,200	1,181	1,200	1,200	1,181

Table 3. AI Investments and Firm Growth

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 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), changes in log employment (columns 3 and 4), and changes in market share within the NAICS 5-digit industry (columns 5 and 6). The dependent variables are measured as growth from 2010 to 2018. The main independent variable is the growth in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.129*	0.156***	0.134	0.152***	0.015	0.014*
	(0.073)	(0.055)	(0.082)	(0.059)	(0.012)	(0.007)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.163	0.336	0.111	0.261	0.229	0.292
Observations	766	766	766	766	766	766

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.138***	0.121***	0.155**	0.109**	0.012*	0.009
	(0.052)	(0.041)	(0.069)	(0.054)	(0.007)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.208	0.398	0.330	0.421	0.140	0.259
Observations	849	849	849	849	849	849

Table 4. Heterogeneous Relationship between AI Investments and Firm Growth by Initial Firm Size

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on contemporaneous changes in AI investments among US public firms (in non-tech sectors), separately for each tercile of initial firm size. Firms in each 2-digit NAICS sector are divided into terciles based on employment in 2010. We consider three measures of firm-level growth for the dependent variable: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the 5-digit NAICS industry (columns 5 and 6). The main independent variable is the growth in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects and initial firm size tercile fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers*Size Tercile 1	0.016	-0.000	0.037	0.008	-0.007	-0.007
	(0.037)	(0.040)	(0.036)	(0.047)	(0.004)	(0.005)
Δ Share AI Workers*Size Tercile 2	0.090	0.043	0.139**	0.116**	0.009	0.008
	(0.060)	(0.067)	(0.054)	(0.058)	(0.008)	(0.009)
Δ Share AI Workers*Size Tercile 3	0.149*	0.174***	0.153	0.164**	0.017	0.016*
	(0.083)	(0.066)	(0.093)	(0.070)	(0.014)	(0.009)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Size tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.193	0.339	0.141	0.262	0.228	0.290
Observations	766	766	766	766	766	766
T-test statistic	2.3	5.5	1.4	3.0	2.7	6.3
T-test p value	0.130	0.019	0.245	0.083	0.103	0.012

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers*Size Tercile 1	-0.057	-0.149***	-0.118	-0.207**	-0.008	-0.012
	(0.079)	(0.054)	(0.099)	(0.090)	(0.007)	(0.009)
Δ Share AI Workers*Size Tercile 2	0.092*	-0.000	0.064	0.018	0.004	-0.001
	(0.049)	(0.047)	(0.061)	(0.054)	(0.011)	(0.012)
Δ Share AI Workers*Size Tercile 3	0.141***	0.126***	0.162**	0.114**	0.013*	0.009
	(0.053)	(0.042)	(0.071)	(0.056)	(0.007)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Size tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.210	0.399	0.339	0.423	0.137	0.256
Observations	849	849	849	849	849	849
T-test statistic	4.6	18.0	5.2	11.0	4.5	6.5
T-test p value	0.033	0.000	0.024	0.001	0.034	0.011

Table 5. AI Investments and Firm Productivity

This table reports the coefficients from long-differences regressions of changes in firm productivity from 2010 to 2018 on contemporaneous changes in AI investments by US public firms (in non-tech sectors). We consider two measures of productivity: log sales per worker (columns 1–2) and revenue TFP (columns 3–4). Revenue TFP is the residual from regressing log revenue on log employment and log capital (constructed using the perpetual inventory method), with separate regressions for each industry sector. We also look at two measures of costs: log COGS in columns 5 and 6 and log operating expenses in columns 7 and 8. The main independent variable is the growth in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2, 4, 6, and 8 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales per Worker		Δ Revenue TFP		Δ Log COGS		Δ Log Operating Expense	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	-0.015	-0.009	-0.006	-0.003	0.142**	0.142***	0.133*	0.154***
	(0.030)	(0.026)	(0.032)	(0.032)	(0.065)	(0.048)	(0.069)	(0.049)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.209	0.264	0.158	0.222	0.157	0.303	0.173	0.354
Observations	766	766	720	720	766	766	766	766

Panel 2: AI measure from job postings data

	Δ Log Sales per Worker		Δ Revenue TFP		Δ Log COGS		Δ Log Operating Expense	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	-0.045	-0.020	-0.017	-0.019	0.159***	0.124***	0.145***	0.121***
	(0.043)	(0.037)	(0.037)	(0.034)	(0.038)	(0.028)	(0.046)	(0.033)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.471	0.545	0.236	0.301	0.211	0.329	0.209	0.374
Observations	849	849	788	788	849	849	849	849

Table 6. AI Investments and Product Innovation

This table reports the coefficients from long-differences regressions of the changes in measures of product innovation from 2010 to 2018 on the contemporaneous changes in AI investments by U.S. public firms (in non-tech sectors). The dependent variable is the change in log(1+number of trademarks) in columns 1 and 2; the change in log(1+number of product patents) in columns 3 and 4; the change in log(1+number of process patents) in columns 5 and 6; andthe change in (1+log number of product manager job postings) in columns 7 and 8. Product (or process) patents are patents with over 50% of the claims being product (or process) claims, following the categorization in Ganglmair et al. (2021). For the main independent variable, Panel 1 considers the resume-based measure of the growth in the share of AI workers from 2010 to 2018, while Panel 2 looks at the job-posting-based measure. Both measures are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2, 4, 6, and 8 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Number of Trademarks		Δ Log Number of Product Patents		Δ Log Number of Process Patents		Δ Log Number of Product Managers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.148**	0.091	0.275***	0.272***	0.002	-0.009	0.209***	0.137
	(0.066)	(0.087)	(0.032)	(0.043)	(0.064)	(0.075)	(0.059)	(0.083)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.194	0.227	0.639	0.680	0.526	0.599	0.063	0.108
Observations	459	459	447	447	447	447	756	756

Panel 2: AI measure from job postings data

	Δ Log Number of Trademarks		Δ Log Number of Product Patents		Δ Log Number of Process Patents		Δ Log Number of Product Managers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.203***	0.176***	0.229***	0.231***	0.127	0.122	0.193***	0.149**
	(0.059)	(0.058)	(0.056)	(0.034)	(0.092)	(0.074)	(0.043)	(0.075)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.403	0.466	0.600	0.688	0.646	0.700	0.119	0.190
Observations	497	497	495	495	495	495	849	849

Table 7. AI Investments and Expansion into New Markets

This table reports the coefficients from long-differences regressions of the changes in the number of markets spanned by U.S. public firms (in non-tech sectors) from 2010 to 2018 on the contemporaneous changes in AI investments. In columns 1 and 2, the dependent variable is the change in the log number of counties with at least 1% of the firm's job postings in Burning Glass; in columns 3 and 4, the dependent variable is the change in the log number of NAICS 6-digit industries associated with at least 1% of the firm's job postings in Burning Glass. For the main independent variable, Panel 1 considers the resume-based measure of the growth in the share of AI workers from 2010 to 2018, while Panel 2 looks at the job-posting-based measure. Both measures are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2 and 4 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Number of Counties		Δ Log Number of Industries	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.081***	0.085***	-0.001	0.023
	(0.020)	(0.028)	(0.015)	(0.021)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.254	0.328	0.130	0.167
Observations	757	757	759	759

Panel 2: AI measure from job postings data

	Δ Log Number of Counties		Δ Log Number of Industries	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.038*	0.071***	0.003	0.017
	(0.023)	(0.023)	(0.009)	(0.015)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.441	0.610	0.076	0.106
Observations	847	847	849	849

Table 8. AI Investments and Changes in Industry Growth and Concentration

This table reports the coefficients from industry-level long-differences regressions of the changes in industry concentration on contemporaneous changes in AI investments. Each observation is a 5-digit NAICS industry, and we exclude tech sectors. The independent variable is the change in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the total industry number of Cognism resumes in 2010 in Panel 1 and the total industry number of Burning Glass job postings in 2010 in Panel 2. The dependent variables are the changes, from 2010 to 2018, in log total sales (including entrants and exits between 2010 and 2018) in columns 1 and 2, in log total employment in columns 3 and 4, in the Herfindahl-Hirschman Index (HHI) in columns 5 and 6 and in the market share of the top firm in an industry in columns 7 and 8. Both measures are calculated using all Compustat firms. All specifications control for industry sector fixed effects. Regressions in columns 2, 4, 6, and 8 also control for log total employment, log total sales, and log average wage in 2010. Standard errors are robust against heteroskedasticity. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Log Sales		Log Employment		HHI		Top Firm Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.134***	0.150***	0.144**	0.167***	0.023***	0.022***	0.016***	0.017***
	(0.046)	(0.035)	(0.056)	(0.047)	(0.005)	(0.006)	(0.005)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Observations	233	233	233	233	233	233	233	233

Panel 2: AI measure from job postings data

	Log Sales		Log Employment		HHI		Top Firm Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.150***	0.142***	0.154*	0.138*	0.023***	0.026***	0.013**	0.013*
	(0.055)	(0.055)	(0.086)	(0.082)	(0.006)	(0.007)	(0.006)	(0.008)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Observations	243	243	243	243	243	243	243	243

A Internet Appendix

A.1 Case Studies on Firms' AI Investments

In order to illustrate the wide range of applications of AI technologies by individual firms, we provide detailed summaries of the investment patterns and uses of AI technologies within four firms in four different industries.

A.1.1 UnitedHealth Group

UnitedHealth Group (UNH) is a large managed healthcare company based in Minnetonka, Minnesota. The group includes a healthcare arm (UnitedHealthcare) established in 1977 and a new technology arm founded in 2011 (Optum). While the UnitedHealthcare arm makes use of AI techniques to optimize operations ranging from cost projections to fraud detection in medical claims, the launch of Optum highlights the way in which firms such as UNH can leverage AI technologies to expand operations by creating new products and entering new market segments. UNH is one of very few companies with access to detailed patient, patient-physician, and drug-patient interaction data for large portions of the U.S. and many additional global locations, making it perfectly placed to harness AI in its operations.

AI use cases and product impact. Most of the AI investments and impact at UNH center around its Optum arm. The traditional UnitedHealthcare part of UNH uses AI in a limited capacity for predictive analytics that inform business decisions and safeguards for vulnerabilities such as fraud detection. The launch of Optum in 2011 has enabled UNH to leverage AI technologies to deliver new products across several healthcare markets. At its core, Optum is a vast data store of proprietary and 3rd party datasets linked together to enable machine-learning-based analysis. Specifically, the AI-powered Optum products include: (i) statistics on drugs and potential alternatives through the pharmaceutical platform Optum Rx; (ii) analysis of electronic medical records through the Optum One platform for physicians; and (iii) the Optum Population Health Management platform for larger institutions (including employers and federal and state agencies) to optimize costs and accessibility to care. The AI-powered OptumIQ system, which is leveraged throughout the Optum solutions, also targets machine-learning-based prediction and diagnostics for diseases such as atrial fibrillation.

Timeline of AI investments at UNH. The use of AI technologies at UNH traces further back than at most firms. As early as the 1990s, UNH piloted AdjudiPro, an AI-powered platform for processing claims from physicians. However, the presence of AI-skilled labor at UNH remained low throughout the 1990s and 2000s, noticeably picking up in 2011 with the launch of the Optum platform. Thereafter, UNH's investment in AI human capital rose steadily throughout the 2010s. The Optum arm of the firm released the Optum360 and Impact Pro products in 2013 and the

Optum One Analytics Platform in 2014, prompting a further acceleration in the rise of UNH's AI human capital in the second half of the decade. The timeline of AI investments at UNH is displayed in Figure A.1.

Internal structure of UNH's AI workforce. UNH has a centralized approach to AI integration, with strategic decisions primarily coming from the headquarters in Minnesota and regional offices handling specific applications. Correspondingly, the majority of UNH's AI workforce concentrates in Minneapolis and Minnetonka, including senior personnel heading AI and machine learning efforts, automation/deployment, consumer analytics, and Optum enterprise analytics. Locations outside of the headquarters tend to employ predominantly engineering and general IT personnel to support the AI efforts.

A.1.2 JPMorgan Chase & Co

JPMorgan Chase & Co (JPM) is the largest bank in the U.S., based in New York City, NY, with consumer banking that has relationships with more than half of U.S. households, a commercial banking arm, a large investment banking business, and a sizable asset management arm. The bank stores hundreds of petabytes of data ranging from credit card transactions and loan applications, to financial news and market data, to alternative data sources.

AI use cases and product impact. The main use cases for AI at JPM fall into the following categories: (i) risk modeling and management ranging from internal cybersecurity to fraud detection in consumer banking and assessment of geo-political risks; (ii) quantitative analysis and algorithmic investment products, including the Algo Central, LOXM, and DeepX programs aimed at executing trades at both maximal speed and optimal prices; (iii) general analytics for Big Data use in broad internal applications including recruiting; and (iv) product development, including enhancements to mobile apps and customer support through AI-powered virtual assistants. In addition, JPM also employs AI in more peripheral applications, for example, with methods for processing of alternative data such as satellite images and mapping contingency plans for AI-driven workforce disruptions. The use of AI at JPM is aimed at both cutting costs (e.g., through risk assessment) and creating new products (e.g., machine-learning-powered trading platforms such as DeepX).

Timeline of AI investments at JPM. As highlighted by Figure A.2, investments in AI at JPM began at the turn of the century, with a steady increase through the first decade turning into an exponential growth in the second decade. The explosion in AI investments at JPM during the 2010s is marked by the acquisition of the multimedia recommendations patent in 2011; an underscoring of the risks associated with data security following a data leak in 2016; and finally the establishment of a dedicated AI research initiative (Machine Learning Center for Excellence) spearheaded by Dr. Manuela Veloso (previously the Chair of the Machine Learning Department

at Carnegie Mellon University) in 2018.

Internal structure of JP Morgan's AI workforce. AI efforts at JPM are centered in the New York location, with peripheral AI expertise throughout the U.S., in London, and in India. JPM has taken a top-down approach to AI investments, with involvement from the highest levels of management and the establishment of a dedicated AI research team in 2018. At the same time, JPM's investments in AI have seen not only the formation of dedicated AI hubs, but also a different approach to corporate structuring. Specifically, the firm's approach relies heavily on small skilled and responsive AI "task-forces" specializing in different sectors (quantitative analysis, user experience, etc.), which can alternatively work on experimental projects (e.g., satellite imagery analysis) or coordinate together to work on core applications (trading algorithms, firm-wide cybersecurity).

A.1.3 Caterpillar Inc.

Caterpillar Inc. is a large construction manufacturing firm headquartered in Deerfield, IL, with a variety of additional business activities including financial products and insurance. The firm has correspondingly varied applications for AI, ranging from inventory management to part recognition, to credit scoring for machinery financing.

AI use cases and product impact. AI investments at Caterpillar are organized along several key verticals. First, the Data Innovation Lab at UIUC conducts core projects in demand forecasting (unstable demand anticipation) and inventory management, in part identification (using techniques from image recognition), and in tracking and tracing technology for fleet management. Second, Caterpillar's asset intelligence efforts include a product line of Internet of Things (IoT) style analytics for managers and machine operators, which facilitates data collection, interpretation, predictive maintenance, and integration. Lastly, smaller targeted efforts at Caterpillar also employ AI techniques in other parts of the business, including leveraging sensor-based data for equipment management and using drone data to optimize job site organization. Caterpillar's uses of AI serve to modernize the firm's machinery, streamline operations and reduce waste through better forecasting and inventory management, and expand the product offerings with the IoT product line and efficient long-term service contracts.

Timeline of AI investments at Caterpillar. Caterpillar began employing workers with AI expertise at the turn of the century, but the growth in the firm's AI workforce went hand-in-hand with the growth in the firm's overall workforce throughout the 2000s (with a dip during the financial crisis). The share of AI employees at Caterpillar noticeably picked up only in mid-2010s, with the CEO Douglas Obenhelmer underscoring the importance of capitalizing on the firm's vast available data resources. Since 2014, Caterpillar has aggressively pursued the development of "smart" machinery, connecting it to predictive IoT-style networks and developing better models for demand prediction. In 2015, Caterpillar established the Analytics and Innovation Division headed

by Greg Foley, and in 2016, the firm hired Morgan Vawters as the Chief of Analytics. The timeline of Caterpillar's investments in AI human capital is presented in Figure A.3.

Internal structure of Caterpillar's AI workforce. The majority of the AI employees at Caterpillar are in the firm's Technology division, with notable presence also in Business and Production departments. The major locations setting the trend for Caterpillar's AI adoption are the company centers in Chicago and Peoria, Illinios, with projects percolating through the dedicated research centers such as the Champaign Innovation Center and production centers such as the manufacturing plant in Aurora, Illinois.

A.1.4 Qualcomm Inc.

Qualcomm Inc. is a wireless telecommunications firm headquartered in San Diego, CA. The firm produces a number of products including semiconductors, hardware, software, and other services related to wireless technology. Device manufacturers such as Apple are Qualcomm's primary clients.

AI use cases and product impact. The principal use of AI at Qualcomm over the past decade and a half has been the improvement of its core products. This includes optimization of chips within mobile devices, improvements to the camera using techniques from computer vision for face recognition and auto-adjustments, audio and video processing, physical sensitivity, power use, and location tracking capabilities. More recently, Qualcomm made a large investment in the development of the Snapdragon Neural Processing Engine (SNPE) platform, which offers a combination of hardware and software on android devices that allows developers to more easily create AI-powered or assisted applications. With the exception of a few stand-alone projects for internal data processing efficiency (e.g., improving internal servers), Qualcomm does not appear to be heavily invested in applying AI for applications such as sales or supply chain optimization, unlike Caterpillar Inc. described above.

Outside of its core businesses, Qualcomm has invested in a number of side products at more exploratory or proof-of-concept stages, such as general work on autonomous vehicles, or enterprise partnerships, for example with Accenture and Kellogg on virtual reality tracking of customers for marketing purposes. This highlights the broad scope of AI technologies that facilitate firms entering new markets: for example, the autonomous vehicle work at Qualcomm makes use of the efforts aimed at enhancing smartphone components, only applied to a different domain.

Timeline of AI investments at Qualcomm. As can be seen from the timeline in Figure A.4, the presence of AI employees at Qualcomm began earlier than in the other firms, and by 2007 the firm initiated dedicated AI research projects in its research arm. The ramp up continued through 2013, marked by collaborations with outside partners such as Brain Corp and internal projects on problems such as face detection. After 2013, Qualcomm saw notable consequences of the earlier

investments, including the first release of SNPE and the formation of an organizationally separate AI research group, but the share of Qualcomm's overall workforce that is skilled in AI remained approximately flat from 2013 to 2018.

Internal structure of Qualcomm's AI workforce. Between 2000 and 2018, the majority of Qualcomm's AI employees have been engineers focused on the improvement of the core product being developed at each point in time, supported by an auxiliary staff of patent counsels and data scientists. In 2018, Qualcomm established a separate AI research group, which is bringing about increased centralization of its AI workforce. Specifically, AI efforts at Qualcomm are organized around the San Diego headquarters, with leadership on overall AI strategy, the newly formed AI research group, and teams spanning nearly every project from computer vision R&D to GPU architecture. Smaller AI offices, scattered mostly throughout the U.S. and Canada, tend to focus on single elements of Qualcomm's AI initiative (for example, SNPE in Toronto and positioning sensors in Santa Clara).

Figure A.1. Timeline of AI investments by UnitedHealth Group

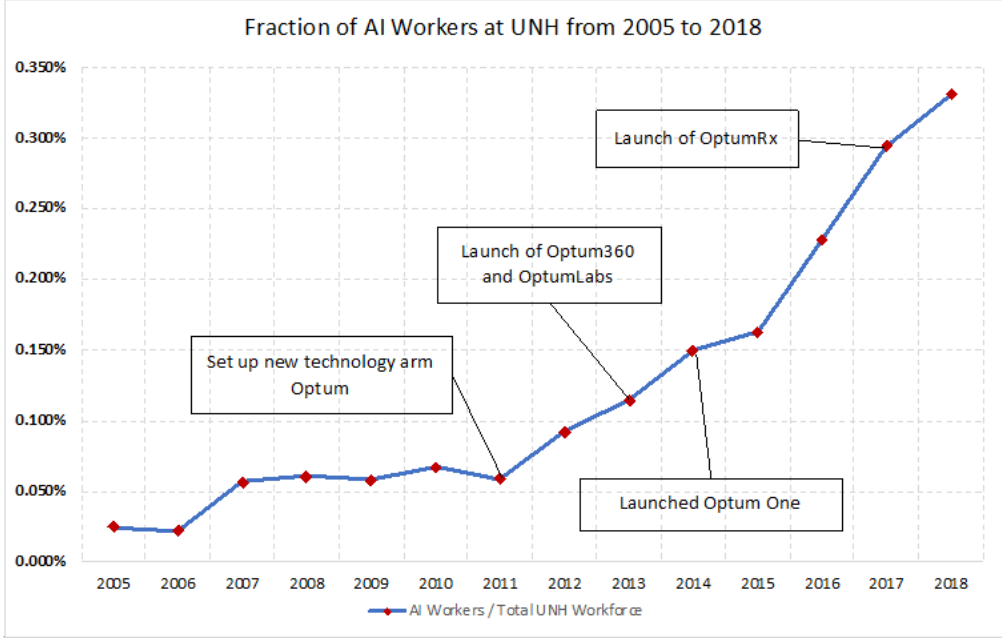


Figure A.2. Timeline of AI investments by JPMorgan Chase & Co

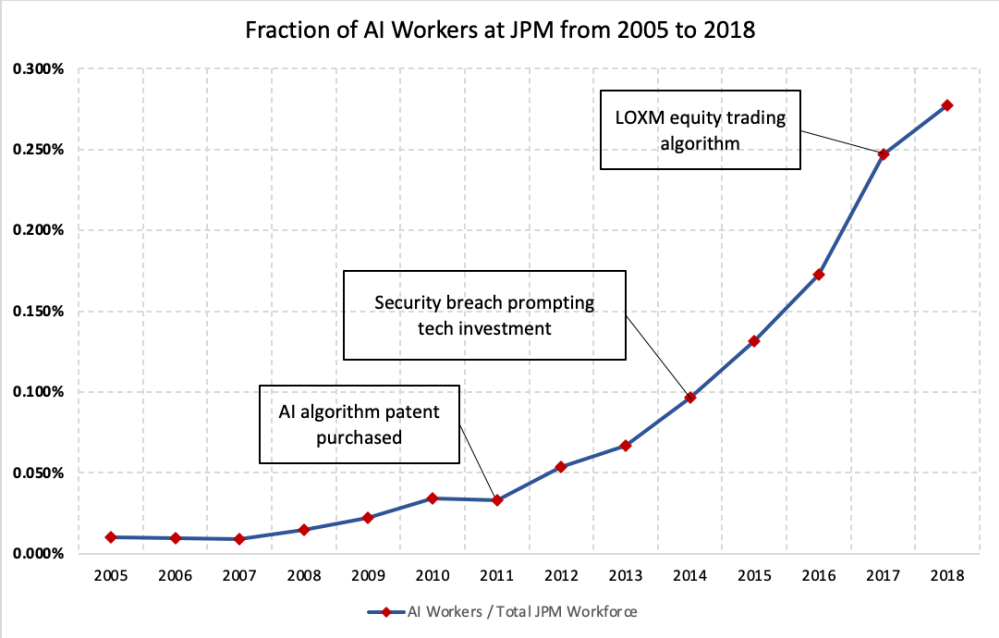


Figure A.3. Timeline of AI investments by Caterpillar Inc

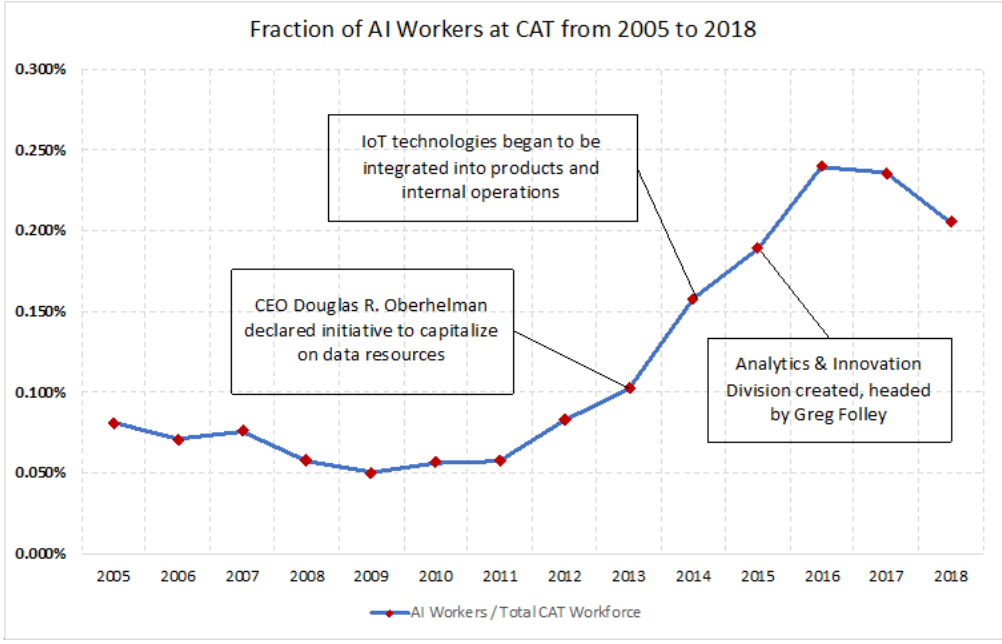
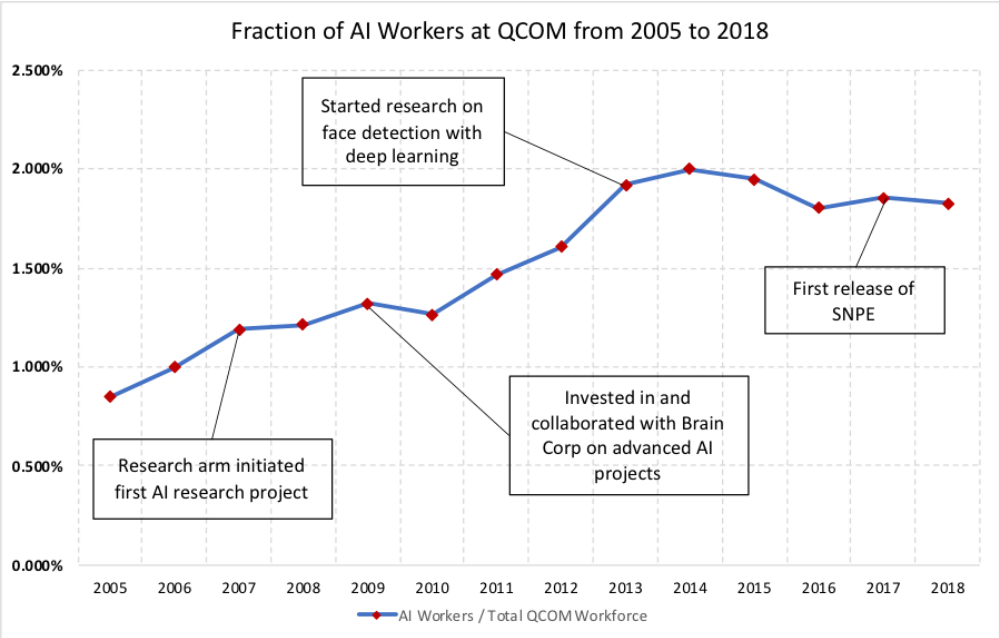


Figure A.4. Timeline of AI investments by Qualcomm Inc



A.2 Additional Figures and Tables

Figure A.5. Matching Rate to Compustat in Job Posting Data

This figure shows the time series of the share of all job postings and the share of AI job postings (job postings with continuous measure $\omega^{NarrowAI}$ above 0.1) that are matched to Compustat firms in the Burning Glass data in 2007 and in each year from 2010 to 2018.

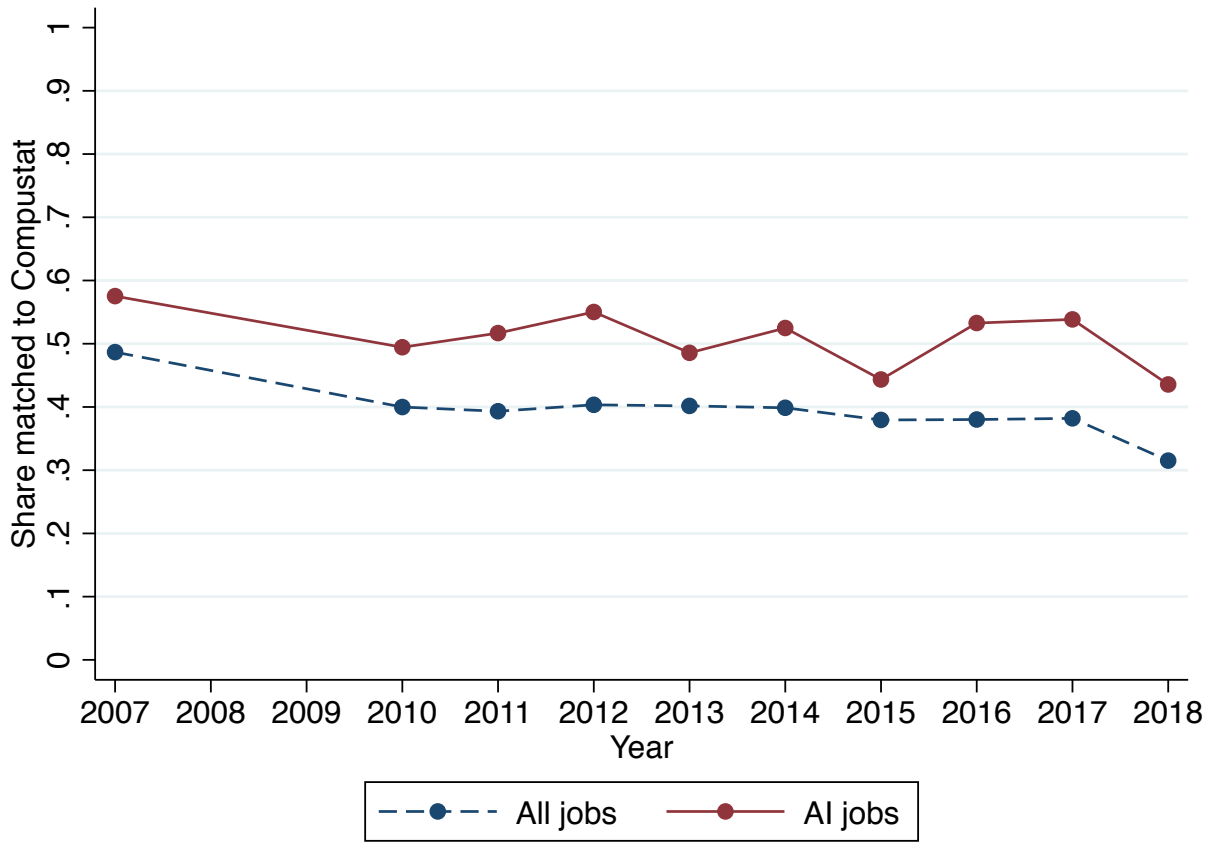


Figure A.6. Distribution of AI Investments across US Geographies

This figure plots the heat map of changes in the job-posting-based measure of AI investments across geographies in the U.S. The figure plots the change in the average AI-relatedness measure ($w_j^{NarrowAI}$) of job postings of public firms in each commuting zone from 2010 to 2018.

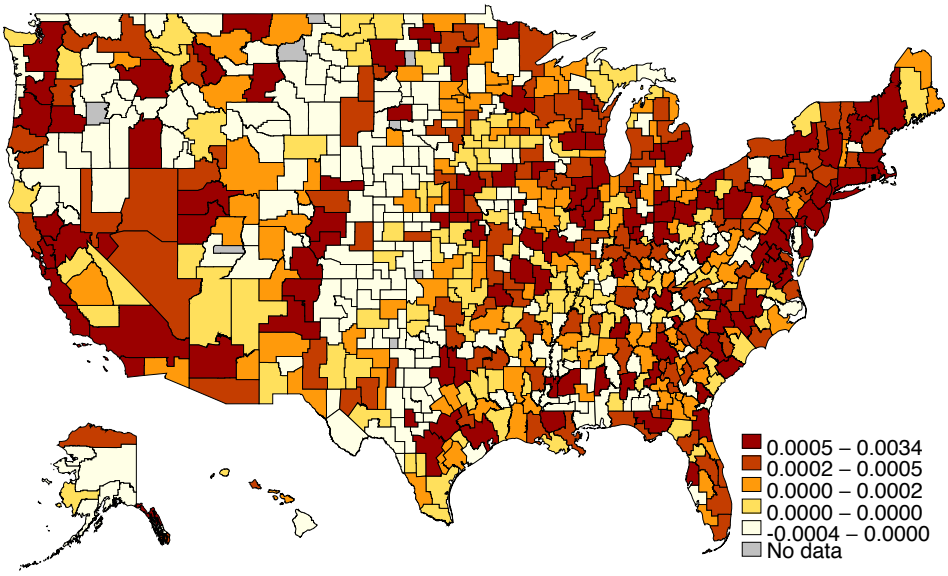
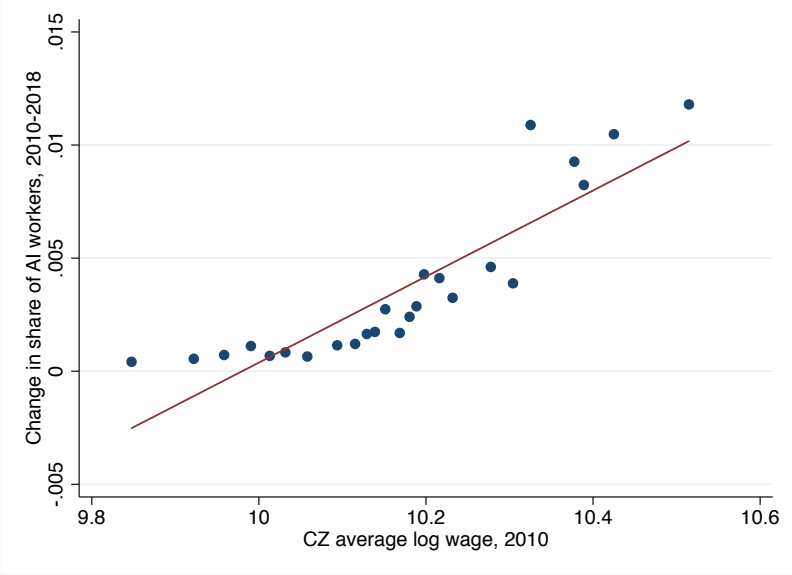
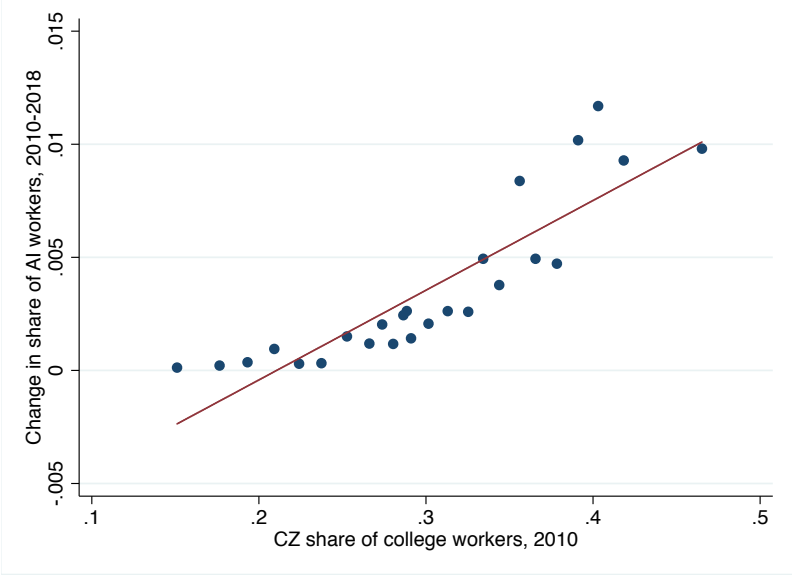


Figure A.7. AI Investments and Local Conditions

This figure is a binned scatter plot of commuting-zone-level AI investments against local conditions. The solid line is the fitted regression line, where regressions are weighted by commuting zones' population in 2010. The y-axis is the change in AI investments (measured as the share of AI workers) from 2010 to 2018, using the Burning Glass data (based on the sample of public firms). The x-axis in Panel (a) is the average log wage of a commuting zone in 2010. The x-axis in Panel (b) is the share of college educated workers in a commuting zone in 2010. The log wage and the share of college-educated workers are from the Census American Community Survey. The t-statistic on the regression slope is 23.6 in the top figure and 23.9 in the bottom figure.



(a) AI Investments and Local Average Wage



(b) AI Investments and Local Share of College-educated Workers

Figure A.8. Time Series of the Share of AI Workers in Foreign Countries from Resume Data

This figure plots the share of AI workers identified in the Cognism data for public firms in non-U.S. countries in each year. All lines exclude firms listed in the U.S. and workers of multinational firms working in the U.S. The group "EU" includes all firms listed on Euronext Paris, Frankfurt Stock Exchange, Borsa Italiana (Milan), SIX Swiss Exchange, NASDAQ Stockholm, NASDAQ Copenhagen, Oslo Stock Exchange, Warsaw Stock Exchange, Vienna Stock Exchange, and Madrid Stock Exchange. The group "UK" includes all firms listed on the London Stock Exchange. The group "India" includes all firms listed on the National Stock Exchange of India and Bombay Stock Exchange. Firms listed on other non-U.S. stock exchanges are included in the group "Rest of the World".

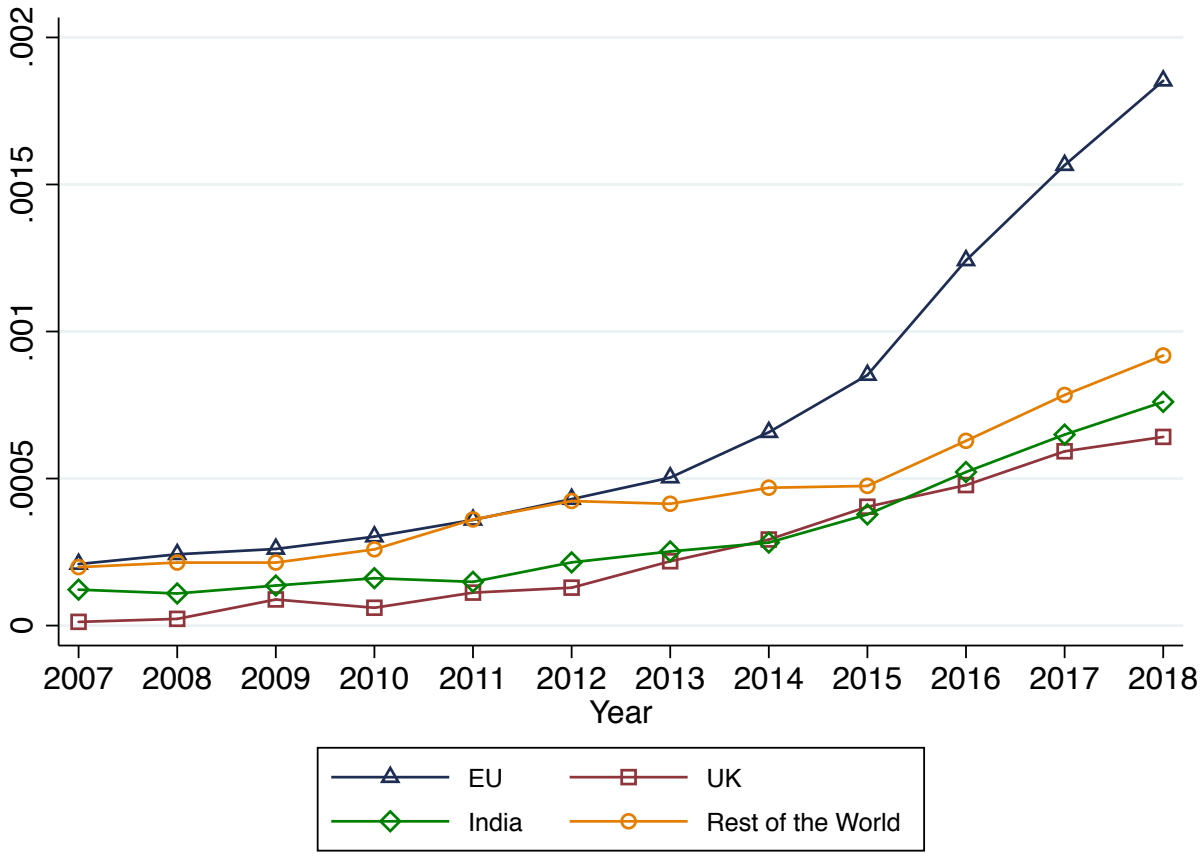


Figure A.9. Correlation Between Changes in the Producer Price Index (PPI) and Changes in the Share of AI Workers in Europe at the Industry Level

This figure is a binned scatter plot of industry-level AI investments against industry-level inflation (measured as PPI). Each dot represents the same number of NAICS 5 digit industries. The solid line is the fitted regression line. The y-axis is the change in AI investments by European firms (measured as the share of AI workers in Cognism data) from 2010 to 2018, which is our first instrument. The x-axis is the change in log PPI in the U.S. from 2010 to 2018. Both variables are measured at the 5-digit NAICS level. The PPI data are from the Bureau of Labor Statistics. The correlation between the two variables is -0.06.

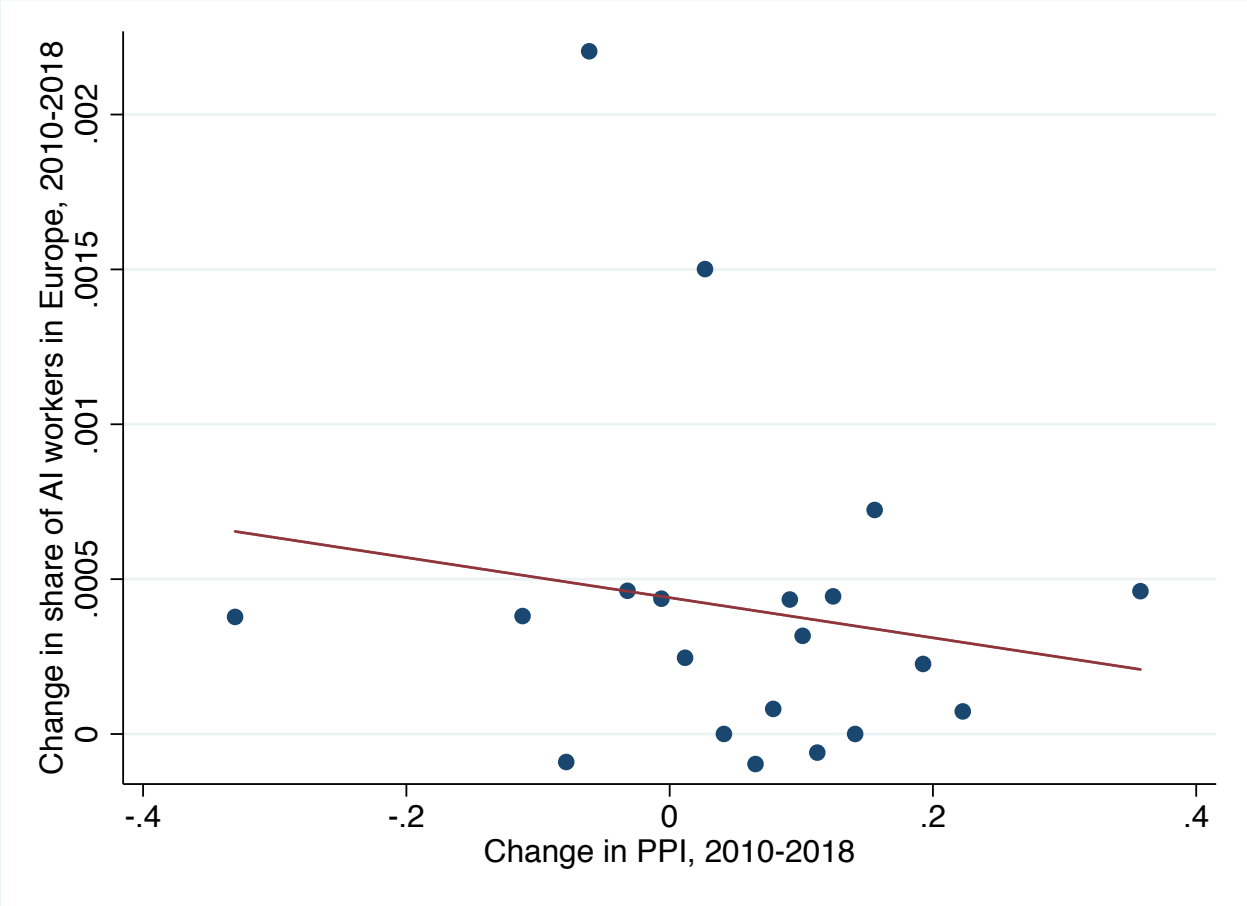


Table A.1. Skills with Highest AI Measures in Burning Glass Job Postings

This table lists the top skills in the Burning Glass data ranked by their narrow-AI measures $\omega_s^{NarrowAI}$. Only skills that are required in at least 50 job postings are included.

#	Skills	AI-relatedness Score
1	Natural Language Processing	1.000
2	Computer Vision	1.000
3	Artificial Intelligence	1.000
4	Machine Learning	1.000
5	ND4J (software)	0.980
6	Kernel Methods	0.979
7	Instance-Based Learning	0.975
8	Microsoft Cognitive Toolkit	0.975
9	Xgboost	0.972
10	Sentiment Classification	0.971
11	Long Short-Term Memory (LSTM)	0.971
12	Libsvm	0.968
13	Semi-Supervised Learning	0.968
14	Recurrent Neural Network (RNN)	0.965
15	Word2Vec	0.956
16	MXNet	0.953
17	Caffe Deep Learning Framework	0.950
18	Autoencoders	0.949
19	MLPACK (C++ library)	0.942
20	Keras	0.941
21	Theano	0.938
22	Torch (Machine Learning)	0.932
23	Wabbit	0.929
24	Boosting (Machine Learning)	0.905
25	TensorFlow	0.904
26	Vowpal	0.903
27	Convolutional Neural Network (CNN)	0.897
28	Jung Framework	0.894
29	OpenNLP	0.894
30	Natural Language Toolkit (NLTK)	0.892
31	Unsupervised Learning	0.891
32	Dlib	0.891
33	Scikit-learn	0.889
34	Latent Semantic Analysis	0.889
35	Latent Dirichlet Allocation	0.889
36	Stochastic Gradient Descent (SGD)	0.881
37	Gradient boosting	0.872
38	Dimensionality Reduction	0.861
39	Deep Learning	0.859
40	DBSCAN (Density-Based Spatial Clustering of Applications with Noise)	0.855
41	AI ChatBot	0.844
42	Recommender Systems	0.842
43	Random Forests	0.840
44	Deeplearning4j	0.839
45	AdaBoost algorithm	0.818
46	Support Vector Machines (SVM)	0.817
47	Unstructured Information Management Architecture	0.806
48	Apache UIMA	0.805
49	Maximum Entropy Classifier	0.799
50	Pybrain	0.786
51	Computational Linguistics	0.780
52	Naive Bayes	0.768
53	H2O (software)	0.763
54	WEKA	0.761
55	Clustering Algorithms	0.740
56	Matrix Factorization	0.739
57	Object Recognition	0.727
58	Classification Algorithms	0.721
59	Information Extraction	0.709
60	Image Recognition	0.706
61	Bayesian Networks	0.705
62	Supervised Learning (Machine Learning)	0.695
63	OpenCV	0.688
64	K-Means	0.683
65	Sentiment Analysis / Opinion Mining	0.679
66	Machine Translation (MT)	0.655
67	Neural Networks	0.640

Table A.2. Examples of AI and Non-AI Job Postings in Burning Glass

This table displays examples of job postings and their narrow-AI measure $\omega_j^{NarrowAI}$. Skills that are classfied as narrow AI skills are highlighted, with the AI relatedness score of skills in parentheses.

#	Job Title	Employer	Skills	Score
	AI jobs			
1	Research Engineer - Natural Language Processing	InterActiveCorp	**Machine Learning(1), Natural Language Processing(1), Natural Language Toolkit(0.895), Computational Linguistics(0.777), WEKA(0.760), Information Extraction(0.709), Mahout(0.593), Information Retrieval(0.360), Apache Hadoop(0.204), Lucene(0.188), SOLR(0.142),** C++, Software Engineering, Python, Lexical Semantics, Ontologies, Java, PERL Scripting Language, Relational Databases, SQL, Search Analytics, Shell Scripting, Web Analytics, Research, Online Research, Extensible Markup Language	0.248
2	Computer Vision & Image Processing Researcher	Rambus Incorporated	**Computer Vision(1), Object Recognition(0.725), OpenCV(0.689), Pattern Recognition(0.442), CUDA(0.362), Image Processing(0.179),** Troubleshooting Technical Issues, C++, Communication Skills, MATLAB, Self-Motivation, Optical System Design and Analysis, Research, Writing, OpenGL, Prototyping, Very Large Scale Integration, Creativity	0.189
3	Algorithm Developer	IBM	**Natural Language Processing(1), Machine Learning(1), IBM Watson(0.125),** Java, Software Development, Candidate Generation, Creativity, Troubleshooting, English	0.236
4	Localization Software Engineer	Nvidia Corporation	**Computer Vision(1), Deep Learning(0.859),** Linear Algebra, OpenGL, C++, Software Engineering, Geometry, Motor Vehicle Operation, Teamwork / Collaboration, Calibration	0.186
5	Speech Recognition Scientist	Vocera Communications	**Computational Linguistics(0.780), Automatic Speech Recognition(0.457), Speech Recognition(0.215),** Experiments, Performance tuning, Research, Written Communication	0.207
6	Data Scientist	Zappos	**Machine Learning(1), Natural Language Processing(1), Boosting (Machine Learning)(0.902), Support Vector Machines (0.816), Naive Bayes(0.759), Matrix Factorization(0.738), Classification Algorithms(0.718), Data Science(0.379), Data Mining(0.159), NoSQL(0.119), Clustering(0.103), Data Structures(0.069),** Relational DataBase Management System, SQL, Attribution Modeling, Detail-Oriented, Revenue Projections, Traffic Maintenance	0.377
7	Data Mining Engineer	Apple Inc.	**Artificial Intelligence(1), Natural Language Processing(1), Machine Learning(1), Unsupervised Learning(0.891), Supervised Learning(0.696), Mahout(0.593), Pattern Recognition(0.442), Apache Hadoop(0.204), Image Processing(0.179), Data Mining(0.159), NoSQL(0.119),** Data Collection, Communication Skills, Java, Detail-Oriented, MATLAB, SQL, Network Engineering, Research, Python, Meeting Deadlines, R, Predictive Models	0.273
8	Big Data Engineer	Socialwire	**Machine Learning(1), Recommender Systems(0.843), MapReduce(0.285), Apache Hadoop(0.204), Big Data(0.196),** Facebook, R, Pinterest, Writing, MATLAB	0.253
9	Big Data Senior Data Scientist	AT&T	**Machine Learning(1), WEKA(0.760), Clustering Algorithms(0.738), Mahout(0.593), Data Science(0.379), Big Data(0.196), Data Mining(0.159), Clustering(0.103),** Simulation, Experimental Testing, R, SPSS, Creativity, SAS, Information Systems, Experiments, Presentation Skills, Research, Data Quality	0.207
10	Data Scientist	Warby Parker	**Natural Language Processing(1), Natural Language Toolkit(0.895), Random Forests(0.839), Pandas(0.498), Data Science(0.379), PIG(0.290), Apache Hadoop(0.204), Data Mining(0.159), Data Visualization(0.136), Tableau(0.074), Pentaho(.058),** NumPy, SQL, Python, Java, DevOps, Agile Development, Creativity, Django, Apache Webserver, Predictive Models, Relational Databases, Data Modeling	0.200
	Non-AI jobs			
11	Data Analyst-Predicitive Modeler	Insys Corporation	**Data Mining(0.159),** Predictive Analytics, SPSS Modeler, SAS, Enterprise Miner, Data Mining, SPSS, SAS, Data Analysis, SQL, UNIX, Communication Skills	0.016
12	Director Of Business Intelligence	Odesus Incorporated	**Data Science(0.379), Data Transformation(0.060),** SQL, Communication Skills, SQL Server Reporting Services, SQL Server, SQL Server Analysis Services, Budgeting, Microsoft Sharepoint, Data Warehousing, MySQL, Key Performance Indicators, Problem Solving, Web Analytics, Market Research, Data Modeling, Business Intelligence, Creativity	0.026
13	Data Scientist	IBM	**Data Science(0.379), Data Mining(0.159),** Software Development, Generalized Linear Model, Communication Skills, Java, Simulation, Research, Product Sales, SAS, Exploratory Analysis, English, Product Management, S-Plus, Thought Leadership	0.036
14	Director, Data & Analytics	Decision Resources	**Big Data(0.196),** Business Intelligence, Business Intelligence Industry Knowledge, Teamwork / Collaboration, Biopharmaceutical Industry Knowledge, Communication Skills	0.033
15	Senior Healthcare Economics Data Analyst	UnitedHealth Group	**Tableau(0.076),** Advanced Statistics, SAS, Data Analysis, SQL, Economics, Database Design, Clinical Data Analysis, Clinical Data Review, Business Process	0.008
16	Delivery Driver And Technician	Rotech Healthcare	Physical Abilities, Lifting Ability, Caregiving, Patient Contact, Patient Transportation and Transfer, HAZMAT, Hazardous Materials Endorsement	0
17	Vice President Underwriting	Morgan Stanley	Workflow Management, Written Communication, Detail-Oriented, Financial Analysis, Mortgage Underwriting, Staff Management	0
18	Quality Assurance Engineer	Amazon	Computer Engineering, Software Development, User Interface (UI) Design, Software Quality Assurance, Black-box testing, Quality Assurance and Control, Consumer Electronics	0
19	Sales Associate	GNC	Sales, Retail Industry Knowledge, Retail Sales, Basic Mathematics	0
20	Dog And Cat Department Manager	Petco	Creativity, Leadership, Budgeting, Sales Goals, Retail Industry Knowledge, Physical Abilities, Inventory Management	0

Table A.3. Job Titles with the Highest Average AI-relatedness Measures

This table reports the job titles in Burning Glass with the highest average AI measure $\omega_j^{NarrowAI}$. We only include job titles that have at least 50 job postings and are matched to Compustat firms.

	cleantitle	share narrowai
1	Artificial Intelligence Engineer	0.476
2	Senior Data Scientist - Machine Learning Engineer	0.367
3	Ai Consultant	0.365
4	Ai Senior Analyst	0.354
5	Lead Machine Learning Scientist - Enterprise Products	0.313
6	Technician Architecture Delivery Senior Analyst Ai	0.294
7	Artificial Intelligence Analyst	0.291
8	Artificial Intelligence Architect	0.288
9	Software Engineer, Machine Learning	0.283
10	Machine Learning Engineer	0.283
11	Computer Vision Engineer	0.272
12	Machine Learning Researcher	0.261
13	Senior Machine Learning Engineer	0.260
14	Senior Software Engineer - Machine Learning	0.252
15	Senior Machine Learning Scientist	0.250
16	Artificial Intelligence Consultant	0.245
17	Computer Vision Scientist	0.233
18	Senior Ai Engineer	0.229
19	Senior Engineer II - Data Scientist	0.227
20	Senior Machine Learning Researcher	0.226
21	Artificial Intelligence Manager	0.216
22	Senior Applied Scientist	0.215
23	Lead Machine Learning Researcher	0.215
24	Vice President- Data Analytics	0.202
25	Big Data Hadoop Consultant	0.198
26	Machine Learning Scientist	0.195
27	Software Engineer - Data Mining/Data Analysis/Machine Learning	0.193
28	Applied Scientist	0.189
29	Senior Engineer - Machine Learning	0.187
30	Senior Associate, Data Scientist	0.183
31	Data Scientist - Engineer	0.181
32	Architect - Relevance Infrastructure	0.177
33	Data Scientist Specialist	0.175
34	Director, Data Scientist	0.175
35	Senior Staff Data Scientist	0.171
36	Data Scientist, Junior	0.169
37	Engineering Manager - Feed Personalization Platform	0.167
38	Manager, Data Scientist	0.162
39	Principal Data Scientist	0.161
40	Director, Data Science	0.161
41	Senior Risk Modeler	0.160
42	Data Science Specialist	0.160
43	Manager/Senior Manager Small Business Open Digital Acquisition	0.159
44	Chief Data Scientist	0.159
45	Manager, Data Science	0.157
46	Research And Development Engineer - Data Mining/Data Analysis/Machine Learning	0.157
47	Senior Manager, Data Science	0.157
48	Big Data Scientist	0.153
49	Data Scientist II	0.152
50	Senior Data Science Engineer	0.151

Table A.4. AI Investments and Firm Growth Controlling for State FE and Tobin's q

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). We consider three measures of growth: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the NAICS 5-digit industry (columns 5 and 6). The main independent variable is the growth in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects and state fixed effects. All columns also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Columns 2, 4, and 6 also control for Tobin's Q in 2010, defined as market value of assets divided by book value of assets. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.154***	0.127***	0.148***	0.126***	0.021***	0.019***
	(0.047)	(0.032)	(0.048)	(0.035)	(0.008)	(0.006)
Tobin's Q 2010		0.225***		0.196***		0.022**
		(0.055)		(0.047)		(0.010)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.405	0.481	0.422	0.470	0.331	0.371
Observations	746	733	746	733	746	733

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.112***	0.079***	0.110***	0.072***	0.014**	0.013**
	(0.034)	(0.022)	(0.038)	(0.027)	(0.007)	(0.006)
Tobin's Q 2010		0.222***		0.248***		0.006
		(0.040)		(0.045)		(0.008)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.464	0.526	0.671	0.707	0.423	0.430
Observations	832	814	832	814	832	814

Table A.5. AI Investments and Firm Growth across Non-Tech Sectors

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous changes in AI investments among U.S. public firms (in non-tech sectors), separately by broad industry sector. Columns 1 and 2 consider firms in the manufacturing sector (2-digit NAICS = 31, 32, 33), columns 3 and 4 consider firms in the wholesale and retail trade sectors (2-digit NAICS = 42, 44, 45), columns 5 and 6 look at firms in the finance sector (2-digit NAICS = 52), and columns 7 and 8 include firms in the other non-tech sectors (all 2-digit NAICS sectors, except those listed above and 51 and 54). Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. The changes in the AI measures are standardized to mean zero and standard deviation of one. We consider two measures of firm growth: changes in log sales in odd columns and changes in log employment in even columns. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All regressions control for industry sector fixed effects, log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Manufacturing		Wholesale & Retail		Finance		Other	
	Sales	Employment	Sales	Employment	Sales	Employment	Sales	Employment
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.145**	0.137**	0.222***	0.253**	0.212*	0.307***	0.125	0.096
	(0.059)	(0.062)	(0.076)	(0.102)	(0.101)	(0.097)	(0.092)	(0.113)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.319	0.281	0.649	0.667	0.478	0.462	0.511	0.255
Observations	382	382	98	98	122	122	164	164

Panel 2: AI measure from job postings data

	Manufacturing		Wholesale & Retail		Finance		Other	
	Sales	Employment	Sales	Employment	Sales	Employment	Sales	Employment
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.136***	0.086**	0.276***	0.325***	0.055**	0.069**	0.163*	0.018
	(0.048)	(0.038)	(0.032)	(0.032)	(0.019)	(0.029)	(0.093)	(0.111)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.319	0.251	0.749	0.791	0.502	0.524	0.380	0.702
Observations	431	431	100	100	130	130	188	188

Table A.6. AI Investments and Firm Growth in Tech Sectors

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous changes in AI investments among U.S. public firms in the tech sectors. Columns 1 and 2 consider firms in the information sector (2-digit NAICS = 51), and columns 3 and 4 consider firms in the professional and business services sector (2-digit NAICS = 54). Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. The changes in the AI measures are standardized to mean zero and standard deviation of one. We consider two measures of firm growth: changes in log sales in odd columns and changes in log employment in even columns. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All regressions control for industry sector fixed effects, log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Information		Prof. & Business Svcs	
	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.226**	0.249**	0.113	0.120
	(0.084)	(0.081)	(0.114)	(0.089)
NAICS2 FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Adj R-Squared	0.442	0.408	0.686	0.692
Obs	100	100	45	45

Panel 2: AI measure from job postings data

	Information		Prof. & Business Svcs	
	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment
	(1)	(2)	(3)	(4)
Δ Share AI Workers	-0.073	0.046	0.132	-0.002
	(0.129)	(0.115)	(0.141)	(0.120)
NAICS2 FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Adj R-Squared	0.317	0.379	0.501	0.495
Obs	114	114	49	49

Table A.7. Continuous AI Measures

This table reports the coefficients from long-differences regressions of growth of U.S. public firms (in non-tech sectors) from 2010 to 2018 on the contemporaneous changes in the average job-level continuous AI measure across all Burning Glass job postings of a firm. Panel 1 presents results from the continuous measure based on narrow AI skills ($\omega_j^{NarrowAI}$), and Panel 2 shows results from the continuous measure based on all skills (ω_j^{AllAI}). The independent variable is standardized to mean zero and standard deviation of one. We consider three measures of firm growth: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the 5-digit NAICS industry (columns 5 and 6). Regressions are weighted by the number of Burning Glass job postings in 2010. All specifications control for industry sector fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: Continuous Narrow AI Measure

	Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.140***	0.128***	0.157**	0.112**	0.012*	0.010
	(0.053)	(0.042)	(0.070)	(0.057)	(0.007)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.208	0.388	0.330	0.423	0.140	0.260
Observations	849	849	849	849	849	849

Panel 2: Continuous All-skill AI measure

	Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.136**	0.123**	0.147*	0.092	0.014*	0.014*
	(0.063)	(0.049)	(0.083)	(0.066)	(0.008)	(0.008)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.180	0.369	0.310	0.412	0.139	0.263
Observations	849	849	849	849	849	849

Table A.8. Alternative Cutoffs for AI Jobs in Burning Glass

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous changes in the share of AI job postings of U.S. public firms (in non-tech sectors). AI job postings are defined as job postings with continuous job-level $\omega_j^{NarrowAI}$ above 0.05 in Panel 1, and job postings with continuous job-level $\omega_j^{NarrowAI}$ above 0.15 in Panel 2. The independent variable is standardized to mean zero and standard deviation of one. We consider three measures of firm growth: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the 5-digit NAICS industry (columns 5 and 6). Regressions are weighted by the number of Burning Glass job postings in 2010. All regressions control for industry sector fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: Cutoff = 0.05

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.139***	0.128***	0.162**	0.123**	0.013*	0.010*
	(0.050)	(0.041)	(0.066)	(0.054)	(0.007)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.211	0.402	0.335	0.426	0.143	0.260
Observations	849	849	849	849	849	849

Panel 2: Cutoff = 0.15

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.127**	0.103**	0.137*	0.088*	0.010	0.007
	(0.056)	(0.042)	(0.071)	(0.052)	(0.007)	(0.005)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.194	0.387	0.318	0.416	0.136	0.257
Observations	849	849	849	849	849	849

Table A.9. AI Investments (Including External AI Software) on Firm Growth

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous firm-level changes in AI investments among U.S. public firms (in non-tech sectors). The main independent variable is the growth in the share of AI-utilizing workers (including both workers with AI skills and workers with jobs referencing external AI software, such as IBM Watson) from 2010 to 2018, standardized to mean zero and standard deviation of one. We consider three measures of firm growth: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the NAICS 5-digit industry (columns 5 and 6). Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.170**	0.146**	0.176**	0.145***	0.027*	0.020*
	(0.076)	(0.067)	(0.075)	(0.056)	(0.015)	(0.011)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.139	0.307	0.107	0.256	0.174	0.247
Observations	864	864	864	864	864	864

Table A.10. Heterogeneous Relationship between AI Investments and Firm Growth by Initial Firm Productivity (TFP)

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous changes in AI investments among US public firms (in non-tech sectors), separately for each tercile of initial firm productivity. Firms in each 2-digit NAICS sector are divided into terciles based on revenue TFP in 2010. We consider three measures of firm growth: changes in log sales (columns 1 and 2), changes in log employment (columns 3 and 4), and changes in market share within the 5-digit NAICS industry (columns 5 and 6). The independent variables are changes in the share of AI workers from 2010 to 2018 interacted with indicator variables for the productivity tercile in 2010. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Both measures are standardized to mean zero and standardized deviation of one. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects and productivity tercile fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers*TFP Tercile 1	0.062	0.069	0.149**	0.152**	0.010	0.005
	(0.070)	(0.073)	(0.066)	(0.066)	(0.012)	(0.010)
Δ Share AI Workers*TFP Tercile 2	0.038	0.095***	0.021	0.076*	0.002	0.010*
	(0.029)	(0.033)	(0.035)	(0.038)	(0.004)	(0.005)
Δ Share AI Workers*TFP Tercile 3	0.350***	0.346***	0.315***	0.294***	0.043**	0.032**
	(0.072)	(0.073)	(0.082)	(0.090)	(0.020)	(0.015)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
TFP tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.257	0.368	0.229	0.319	0.274	0.319
Observations	724	724	724	724	724	724
T-test statistic	8.6	7.3	2.8	1.7	1.8	2.0
T-test p value	0.004	0.008	0.099	0.199	0.177	0.161

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers*TFP Tercile 1	0.035	0.023	0.169**	0.140***	-0.004	-0.007
	(0.049)	(0.055)	(0.073)	(0.052)	(0.011)	(0.011)
Δ Share AI Workers*TFP Tercile 2	0.051**	0.045**	0.045*	0.024	-0.002	0.001
	(0.022)	(0.022)	(0.026)	(0.029)	(0.005)	(0.005)
Δ Share AI Workers*TFP Tercile 3	0.252***	0.236***	0.234***	0.181***	0.032***	0.025**
	(0.046)	(0.034)	(0.080)	(0.062)	(0.010)	(0.011)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
TFP tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.279	0.441	0.419	0.522	0.210	0.310
Observations	793	793	793	793	793	793
T-test statistic	12.2	14.1	0.3	0.2	4.2	4.3
T-test p value	0.001	0.000	0.555	0.628	0.042	0.040

Table A.11. AI Investments and Productivity of Early Adopters

This table reports the coefficients from long-differences regressions of changes in firm productivity from 2010 to 2018 on the changes in AI investments by U.S. public firms (in non-tech sectors) from 2010 to 2014. We consider two measures of productivity: log sales per worker and in columns 1 an 2, and revenue TFP in columns 3 and 4. Revenue TFP is the residual from regressing log revenue on log employment and log capital (constructed using the perpetual inventory method), with separate regressions for each industry sector. The main independent variable is growth in the share of AI workers from 2010 to 2014, calculated based on resumes in Panel 1 and based on job postings in Panel 2. All independent variables are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2 and 4 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level.

Panel 1: AI measure from resume data

	Δ Log Sales per Worker		Δ Revenue TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Workers 2010-2014	-0.039	-0.030	-0.021	-0.017
	(0.047)	(0.037)	(0.051)	(0.046)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.212	0.265	0.160	0.223
Observations	766	766	720	720

Panel 2: AI measure from job postings data

	Δ Log Sales per Worker		Δ Revenue TFP	
	(1)	(2)	(3)	(4)
Δ Share AI Workers 2010-2014	-0.039	-0.030	-0.021	-0.017
	(0.047)	(0.037)	(0.051)	(0.046)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.212	0.265	0.160	0.223
Observations	766	766	720	720

Table A.12. AI Investments and Industry-level Employment from QWI

This table reports the coefficients from industry-level regressions of the changes in total employment for all (public and nonpublic) US firms on the contemporaneous changes in AI investments of Compustat firms. The dependent variable is the change in log industry employment from 2010 to 2018 calculated from Census Quarterly Workforce Indicators (QWI) data. Each observation is a 4-digit NAICS industry, which is the finest industry level available in the QWI data. The independent variable is the change in the share of AI workers from 2010 to 2018, standardized to mean zero and standard deviation of one. Columns 1 and 2 consider the resume-based measure of AI investments, while columns 3 and 4 look at the job-posting-based measure. Regressions are weighted by the total industry number of Cognism resumes in 2010 in columns 1 and 2 and the total industry number of Burning Glass job postings in 2010 in columns 3 and 4. All specifications control for industry sector fixed effects. Regressions in columns 2 and 4 also control for log total employment, log total sales, and log average wage in 2010. Standard errors are robust against heteroskedasticity.

	Dependent variable: Δ Log Employment			
	Cognism		Burning Glass	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	-0.038	-0.031	-0.003	-0.020
	(0.038)	(0.039)	(0.036)	(0.039)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Observations	149	149	155	155

Table A.13. AI Investments and Markups

This table reports the coefficients from long-differences regressions of the changes in markups from 2010 to 2018 on the contemporaneous changes in AI investments by US public firms (in non-tech sectors). We consider three measures of mark-ups: sales divided by cost of goods sold (COGS) in columns 1 and 2, sales over total operating expenses in columns 3 and 4, and the Lerner Index (operating income before depreciation and amortization minus depreciation scaled by sales) in columns 5 and 6. For the main independent variable, Panel 1 considers the resume-based measure of the growth in the share of AI workers from 2010 to 2018, while Panel 2 looks at the job-posting-based measure. Both measures are standardized to mean zero and standardized deviation of one. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All specifications control for industry sector fixed effects. Columns 2, 4, 6, and 8 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Markup (COGS)		Δ Log Markup (Total Exp)		Δ Lerner Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	-0.010	0.013	-0.004	0.003	-0.007*	-0.002
	(0.016)	(0.022)	(0.006)	(0.009)	(0.004)	(0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.332	0.411	0.260	0.338	0.173	0.340
Observations	766	766	766	766	766	766

Panel 2: AI measure from job postings data

	Δ Log Markup (COGS)		Δ Log Markup (Total Exp)		Δ Lerner Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	-0.020	-0.004	-0.008	0.001	-0.005	-0.001
	(0.029)	(0.031)	(0.009)	(0.012)	(0.006)	(0.009)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.277	0.386	0.309	0.457	0.119	0.587
Observations	849	849	849	849	849	849

Table A.14. Controlling for Other Technologies

This table reports the coefficients from long-differences regressions of firm growth from 2010 to 2018 on the contemporaneous changes in each firm's share of AI workers among U.S. public firms (in non-tech sectors) , controlling for investments in other technologies measured using the Burning Glass data. Columns 1 and 2 control for the 2010-2018 changes in the firm's share of non-AI IT jobs, columns 3 and 4 control for changes in the share of robot-related jobs, columns 5 and 6 control for changes in the share of non-AI data-related jobs, and columns 7 and 8 control for changes in the share of non-AI data analysis jobs. An IT job is defined as a job for which at least 10% of the required skills are in the "Information Technology" skill cluster; a robot-related job is a job with a robot relatedness score (constructed with the same methodology as the AI-relatedness score but using the core skill of "Robotics") above 0.1; a data-related job is a job with at least 10% of required skills in data-related skill clusters; a data analysis job is a job with at least 10% of required skills in the "Analysis" skill cluster. All measures are standardized to mean zero and standard deviation of one. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. All regressions control for industry sector fixed effects, log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: AI measure from resume data

	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.159***	0.150**	0.155***	0.149**	0.157***	0.150**	0.157***	0.150**
	(0.053)	(0.062)	(0.052)	(0.060)	(0.052)	(0.060)	(0.053)	(0.061)
Δ Share Non-AI IT Workers	0.019	-0.003						
	(0.038)	(0.043)						
Δ Share Robot Workers			0.013	0.009				
			(0.033)	(0.037)				
Δ Share Non-AI Data Workers					0.056*	0.053		
					(0.033)	(0.040)		
Δ Share Non-AI Data Analysis Workers							0.048	0.038
							(0.029)	(0.035)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.339	0.255	0.339	0.255	0.344	0.259	0.344	0.258
Observations	756	756	756	756	756	756	756	756

Panel 2: AI measure from job postings data

	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.121***	0.105*	0.108***	0.097**	0.128***	0.111**	0.124***	0.108**
	(0.040)	(0.055)	(0.032)	(0.046)	(0.038)	(0.054)	(0.039)	(0.053)
Δ Share Non-AI IT Workers	0.038	0.049						
	(0.059)	(0.062)						
Δ Share Robot Workers			0.125**	0.088				
			(0.056)	(0.062)				
Δ Share Non-AI Data Workers					-0.036	-0.032		
					(0.064)	(0.077)		
Δ Share Non-AI Data Analysis Workers							0.007	0.007
							(0.057)	(0.068)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.388	0.423	0.404	0.426	0.388	0.422	0.387	0.421
Observations	849	849	849	849	849	849	849	849

Table A.15. AI Investments and Firm Growth: Foreign Industry IV

This table estimates the relationship between AI investments and firm growth from 2010 to 2018 for U.S. public firms (in non-tech sectors), using the change in the share of AI workers in European public firms in the same NAICS 5-digit industry as an instrument for the two firm-level measures of the growth in the share of AI workers from 2010 to 2018. Panel 1 presents the results for the resume-based measure of the share of AI workers, while Panel 2 focuses on the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. The first stage is tabulated in columns 1 and 2, and the second stage results are displayed in columns 3-6. The independent variable and the IV are standardized to mean zero and standard deviation of one. As the dependent variable, we consider changes in log sales in columns 3 and 4 and changes in log employment in columns 5 and 6. All specifications control for industry sector fixed effects. Columns 2, 4, and 6 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industries, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: IV results using AI measure from resume data

	First Stage		Second Stage			
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Instrument	0.398***	0.281***				
	(0.092)	(0.050)				
Δ Share AI Workers			0.359***	0.419***	0.313*	0.356*
			(0.100)	(0.103)	(0.160)	(0.185)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	18.6	31.6	18.6	31.6	18.6	31.6
Observations	643	643	643	643	643	643

Panel 2: IV results using AI measure from job postings data

	First Stage		Second Stage			
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Instrument	0.548***	0.495***				
	(0.151)	(0.090)				
Δ Share AI Workers			0.224***	0.190***	0.253**	0.232**
			(0.065)	(0.060)	(0.099)	(0.093)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	13.2	29.9	13.2	29.9	13.2	29.9
Observations	714	714	714	714	714	714

Table A.16. Pre-trend of Foreign Industry IV

This table investigates pre-trends in growth of U.S. public firms (in non-tech sectors), measured over 1999-2007, prior to AI investments in 2010-2018. We estimate an IV regression of each pre-trend variable against the two measures of growth in the share of AI workers from 2010 to 2018 (based on resume data in Panel 1 and based on job posting data in Panel 2), standardized to mean zero and standard deviation of one. The independent variable is instrumented using the change in the share of AI workers for European public firms in the NAICS 5-digit industry. Column 1 looks at pre-trends in log sales, and column 2 considers pre-trends in log employment. All specifications control for industry sector fixed effects and log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industry, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level.

Panel 1: Testing for pre-trends using AI measure from resume data

	Δ Log Sales, 1999–2007	Δ Log Employment, 1999–2007
	(1)	(2)
Δ Share AI Workers	0.011	-0.047
	(0.163)	(0.178)
NAICS2 FE	Y	Y
Controls	Y	Y
F Statistic	27.2	24.9
Observations	529	498

Panel 2: Testing for pre-trends using AI measure from job postings data

	Δ Log Sales, 1999–2007	Δ Log Employment, 1999–2007
	(1)	(2)
Δ Share AI Workers	0.082	0.040
	(0.065)	(0.077)
NAICS2 FE	Y	Y
Controls	Y	Y
F Statistic	27.9	24.8
Observations	577	541

Table A.17. Effect of AI Investments on Firm Growth: Bartik IV

This table estimates the relationship between AI investments and firm growth from 2010 to 2018 for U.S. public firms (in non-tech sectors), using a weighted average of changes in the share of AI workers across industries operating in the same location as a given firm as an instrument for the two firm-level measures of the growth in the share of AI workers from 2010 to 2018. Panel 1 presents the results for the resume-based measure of AI workers, while Panel 2 focuses on the job-posting-based measure. Regressions are weighted by the number of Cognism resumes in 2010 in Panel 1 and the number of Burning Glass job postings in 2010 in Panel 2. The first stage is provided in columns 1 and 2, and the second stage results are displayed in columns 3-8. The independent variable and the IV are standardized to mean zero and standard deviation of one. We consider changes in log sales in columns 3 and 4, log employment in columns 5 and 6, and market share within the 5-digit NAICS industry in columns 7 and 8. All specifications control for industry sector fixed effects. Columns 2, 4, 6, and 8 also control for log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industry, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel 1: IV results using AI measure from resume data

	First Stage		Second Stage					
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Instrument	0.642***	0.754***						
	(0.109)	(0.192)						
Δ Share AI Workers			0.343***	0.317***	0.297***	0.261***	0.036	0.024*
			(0.092)	(0.077)	(0.093)	(0.088)	(0.022)	(0.013)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	34.5	15.5	34.5	15.5	34.5	15.5	34.5	15.5
Observations	766	766	766	766	766	766	766	766

Panel 2: IV results using AI measure from job postings data

	First Stage		Second Stage					
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment		Δ Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Instrument	0.771***	0.489**						
	(0.209)	(0.198)						
Δ Share AI Workers			0.230***	0.360**	0.281**	0.237	0.018	0.053
			(0.053)	(0.159)	(0.123)	(0.254)	(0.017)	(0.040)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	13.6	6.1	13.6	6.1	13.6	6.1	13.6	6.1
Observations	849	849	849	849	849	849	849	849

Table A.18. Pre-trend of Bartik IV

This table investigates pre-trends in growth of U.S. public firms (in non-tech sectors), measured over 1999-2007, prior to AI investments in 2010-2018. We estimate an IV regression of each pre-trend variable against the two measures of the growth in the share of AI workers from 2010 to 2018 (based on resume data in Panel 1 and based on job posting data in Panel 2), standardized to mean zero and standard deviation of one. The independent variable is instrumented using a weighted average of changes in the share of AI workers across industries operating in the same location as a given firm. Column 1 looks at pre-trends in log sales; column 2 considers pre-trends in log employment, and column 3 reports pre-trends in market share within the 5-digit NAICS industry. All specifications control for industry sector fixed effects and log employment, cash/assets, log sales, log industry wages, R&D/Sales, and log markups, as well as characteristics of the commuting zones where the firms are located (average log wage, the share of college graduates, the share of routine workers, the share of workers in finance and manufacturing industry, the share of workers in IT-related occupations, and the share of female and foreign-born workers), all measured as of 2010. Standard errors are clustered at the 5–digit NAICS industry level.

Panel 1: Testing for pre-trends using AI measure from resume data

	Δ Log Sales, 1999–2007	Δ Log Employment, 1999–2007	Δ Market Share, 1999–2007
	(1)	(2)	(3)
Δ Share AI Workers	0.260	0.204	0.008
	(0.178)	(0.160)	(0.023)
NAICS2 FE	Y	Y	Y
Controls	Y	Y	Y
F Statistic	13.7	12.9	13.7
Observations	627	591	631

Panel 2: Testing for pre-trends using AI measure from job postings data

	Δ Log Sales, 1999–2007	Δ Log Employment, 1999–2007	Δ Market Share, 1999–2007
	(1)	(2)	(3)
Δ Share AI Workers	-0.142	-0.131	-0.019
	(0.270)	(0.298)	(0.043)
NAICS2 FE	Y	Y	Y
Controls	Y	Y	Y
F Statistic	5.5	4.8	5.4
Observations	687	647	690