

Artificial Intelligence, Firm Growth, and Industry Concentration*

Tania Babina[†] Anastassia Fedyk[‡]

Alex He[§] James Hodson[¶]

September 2020

Abstract

Which firms invest in artificial intelligence (AI) technologies, and how do these investments affect individual firms and industries? We provide a comprehensive picture of the use of AI technologies and their impact among US firms over the last decade, using a unique combination of job postings and individual-level employment profiles. We introduce a novel measure of investments in AI technologies based on human capital and document that larger firms with higher sales, markups, and cash holdings tend to invest more in AI. Firms that invest in AI experience faster growth in both sales and employment, which translates into analogous growth at the industry level. The positive effects are concentrated among the ex ante largest firms, leading to a positive correlation between AI investments and an increase in industry concentration. However, the increase in concentration is not accompanied by either increased markups or increased productivity. Instead, firms tend to expand into new product and geographic markets. Our results are robust to instrumenting firm-level AI investments with foreign industry-level AI investments and with local variation in industry-level AI investments, and to controlling for investments in general information technology and robotics. We also document consistent patterns across measures of AI using firms' demand for AI talent (job postings) and actual AI talent (resumes). Overall, our findings support the view that new technologies, such as AI, increase the scale of the most productive firms and contribute to the rise of superstar firms.

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, 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 Columbia University, Sao Paulo School of Economics, UC Berkeley, University of Maryland, AFFECT, NBER SI Macroeconomics and Productivity, UNC Junior Roundtable, and the Labor and Finance Group Conference for helpful suggestions. The authors thank Cognism Ltd. for providing the employment data and Burning Glass Technologies for providing the job postings data. 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 have led to substantial developments in artificial intelligence (AI) technologies,¹ enabling their commercial applications across a broad landscape of firms, industries, and countries.² Yet it remains an open question whether artificial intelligence can transform economies and spur economic growth. On the one hand, as a potential general purpose technology, AI can generate growth through product innovations and increased productivity (Aghion et al., 2017; Acemoglu and Restrepo 2018; Cockburn et al., 2018; Mihet and Philippon, 2019). On the other hand, aggregate productivity growth over the past decade has slowed significantly, leading to concerns that AI may not deliver economic growth or take a much longer time to reach its potential (Brynjolfsson et al., 2019). So far, there has been little systematic evidence on AI investments and their economic impact.

In this paper, we examine the adoption of artificial intelligence technologies (e.g., machine learning, natural language processing, and computer vision) and their impact on the growth of US firms and industries over the past decade. Artificial intelligence technologies can lead to firm growth through several non-mutually-exclusive channels with differing implications. First, AI technologies can streamline business processes involving prediction and anomaly detection, leading to improvements in forecast accuracy and resource allocation (Brynjolfsson et al., 2011; Agrawal et al., 2019; Mihet and Philippon, 2019; Tanaka et al., 2019). This type of automation can potentially stimulate growth by streamlining production and increasing productivity. Second, among the more visible applications of AI to date have been the tailoring of product offerings and the targeting of online advertisements, which can potentially enable firms to price discriminate and gain market power (Shiller, 2016; Mihet and Philippon, 2019). Third, AI can create scale advantages that benefit the ex ante largest and most productive firms. For example, AI can change cost structures by reducing the cost of spanning multiple markets (Aghion et al., 2019) or disproportionately benefit firms with extensive operations due to AI's unique reliance on big data (Farboodi et al., 2019).

A major challenge in studying the impact of AI is a lack of data on the use of AI technologies

¹According 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.”

²For example, a third of companies surveyed globally by Deloitte in 2018 have a comprehensive AI strategy; see [here](#). Bughin et al. (2018) and Furman and Seamans (2019) provide an overview of 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).

at the firm level. We overcome this challenge by developing a novel measure of AI investments based on detailed data on firms' human capital, motivated by the heavy reliance of AI on human expertise rather than physical capital. We use a unique combination of datasets to measure firms' AI-related human capital: job postings data from Burning Glass Technologies, which include the near-universe of online job vacancies, and resume data from Cognism Inc, which offer job histories for hundreds of millions of individual employees. These datasets enable us to construct firm-level measures of AI investments by observing both the stock and the hiring of AI-related employees.

We offer a new approach to identify AI-related jobs, which consists of several steps. First, we measure AI-relatedness of each skill using the job postings data. The measure is based on the intuition that if a given skill is related to AI, then jobs requiring that skill should also require some of the core AI skills. We define four core AI skills: "machine learning", "natural language processing", "computer vision", and "artificial intelligence". We measure the AI-relatedness of a given skill in the job postings data as the fraction of jobs requiring that skill that also require one of the four core AI skills.³ Second, we obtain a measure of AI-relatedness of each job posting by averaging the AI-relatedness measure across all required skills listed in that job posting. This gives a continuous measure of AI relatedness ranging from 0 to 1 for each job. We define AI jobs as jobs whose AI-relatedness measure exceeds a certain cutoff (e.g., 0.1). Third, we leverage the most AI-related skill terms to classify AI jobs in the less structured resume data. For each employee, at each point in time, we consider whether terms with the highest AI-relatedness (e.g. "deep learning") appear either in the contemporaneous job title, in the job description, in any patents or publications produced during the job, or in any awards received during the job. This gives us a classification of each employee of each firm at each point in time.

Our measure of AI jobs offers several advantages over previous studies. First, compared to previous studies using bag-of-words approaches with job postings data, our method does not require researchers to pre-specify a list of AI-related keywords and instead learns the most relevant terms from the data. Second, our measure is built on a continuous (rather than binary) classification, which allows us to capture a wide range of AI-related skills and differentiate more AI-related skills (e.g., deep learning) from less AI-related skills (e.g., information retrieval). Our paper is also the first one to cross-validate AI demand identified from job postings with AI hiring identified from resumes. Importantly, our methodology for identifying AI-skilled jobs can also be

³For example, the AI-relatedness measure for "deep learning" is 0.86, meaning that 86% of all jobs that require deep learning also require one of the four core AI skills. In contrast, the measure for "information retrieval" is 0.37, the measure for "regression analysis" is 0.09, and the measure for "communication skills" is 0.003.

applied to identify jobs related to a wide range of other technologies (e.g., blockchain, robotics, and digitization), especially those centering on human expertise.

We confirm that our human-capital-based measures of AI investments constructed from resumes and job postings display intuitive properties. In both datasets, the fraction of AI jobs has monotonically increased over time, growing more than five-fold from 2010 to 2018. Firms in all sectors increased their AI investments by similar relative magnitudes, consistent with AI being a general purpose technology. Encouragingly, the two measures of AI investments—based on a firm’s *current* employees (from resumes) and the demand for *additional* AI employees (from job postings)—are highly correlated, and all subsequent tests in the paper yield consistent results across these two independent datasets.

We begin our empirical analysis by examining which firms invest in AI. We aggregate both job postings data and resume data to the firm level and match to public firms in the Compustat database. We predict growth in the share of AI-related employees from 2010 to 2018 based on firm characteristics measured as of 2010. Our results indicate that larger firms, in terms of both sales and market share, are more likely to invest in AI, consistent with the evidence by [Alekseeva et al. \(2020\)](#). Furthermore, AI investments are stronger among firms with larger cash holdings, higher mark-ups, and higher 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 then address the fundamental question of whether artificial intelligence is able to stimulate growth for the firms investing in AI. Given that AI investments are gradual over time, and that we do not expect to observe their effects on firms immediately, our primary specification is a long-differences regression of changes in firm-level outcomes on contemporaneous changes in the share of AI workers from 2010 to 2018. We control for industry fixed effects and firm-, industry-, and commuting-zone-level characteristics from 2010 that predict AI investments. We document a strong and consistent pattern: firms that invest in AI grow more. Specifically, a one-standard-deviation increase in the share of AI workers based on the resume data 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. The results are similar using the job-postings-based measure of AI investments, where a one-standard-deviation change in the share of AI workers corresponds to a 12.1% increase in sales, a 10.9% increase in employment, and a 0.9 percentage point increase in market share. Moreover, the results are ubiquitous across major industry sectors, supporting the claim that artificial intelligence is a general purpose technology.

We perform two additional tests to rule out potential confounding factors. First, we conduct an event study analysis using a distributed lead-lag model (Stock and Watson, 2015; Aghion et al., 2020) to measure the dynamic effects of AI investments and address concerns about reverse causality. We find no increases in firm growth prior to AI investments and confirm that the effects are not immediate: the impact of AI on both sales and employment becomes discernible only after a lag of two years. Second, we confirm that our results reflect specifically investments in artificial intelligence, rather than other related technologies. We show that the relationship between AI investments and firm-level growth is robust to controlling for contemporaneous firm-level investments in robotics and other non-AI information technologies.

Although we control for factors that predict AI investments, the decision to invest in AI may be endogenous even conditional on controls, which can lead to biased estimates. The direction of this bias is ex ante ambiguous, with potential upward bias if firms investing in AI are on a faster growth trajectory or have better management, and potential downward bias due to measurement error or firms investing in AI in response to fewer growth opportunities or anticipated negative shocks. These concerns are partially assuaged by the event study, which does not show any pre-trends in either sales or employment. However, if AI investments coincide with shocks that are exactly contemporaneous and not pre-existing, those would not be visible in the pre-trends.

In order to further address the endogeneity concerns around firms' decisions to invest in AI, we take advantage of two different instruments for AI investments. First, we employ a strategy similar to Autor et al. (2013) and Acemoglu et al. (2020) and instrument the change in each firm's share of AI workers using the change in the share of AI workers among European firms in the same 5-digit NAICS industry. Our unique resume data, which include consistent coverage of European firms, enable the detailed measurement of firm- and industry-level AI investments in Europe. Second, we use a shift-share design where firm-level AI investments are instrumented with a weighted average of national industry-level AI investment rates, where the weights are given by the industry share at the location(s) where the firm operates. Both instruments have a strong first stage, with F-statistics above 10. Consistent with the argument that differences in AI investments across industries are largely driven by differences in technical feasibility and availability of data rather than different demand shocks or growth trajectories (Bughin et al., 2017), we show that firms that have higher predicted AI investments are not on different growth trajectories prior to 2010, and controlling for pre-2010 growth does not change the results. We also show that industry-level growth in the share of AI employees in Europe is not positively correlated with changes in

industry-level prices in the U.S. from 2010 to 2018, suggesting that the relevance of the European IV is not driven by correlated industry-level demand shocks. Both IV strategies yield results that are consistent with the estimates from the OLS regressions.

We next turn to whether firm-level growth aggregates into industry-level growth in sales and employment. It is possible that the positive effects on employment and sales of 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 (Acemoglu et al., 2020). Nevertheless, we find that industries that invest more in AI experience an overall increase in sales and employment within the sample of Compustat firms. We estimate industry-level growth at the 5-digit NAICS level, which is the same level of granularity as the instrument based on AI investments in Europe, and we confirm that the industry-level results are consistent across OLS and IV regressions. The positive effect on employment is especially surprising, given wide-spread concerns of the potential for AI to replace labor (Frank et al., 2019). This highlights the differences between AI and previous technologies, such as robots and automation: although larger firms are also more likely to adopt robots (Humlum, 2019; Acemoglu et al., 2020), most of the prior evidence finds that robot adoption leads to lower aggregate employment (Autor and Salomons, 2018; Acemoglu and Restrepo, 2019; Zator, 2019).⁴

AI investments not only spur industry growth, but also increase industry concentration. By interacting AI investments with firms' initial size, we show that the positive relationship between AI investments and growth concentrates among the ex ante largest firms. For example, a one-standard-deviation increase in the share of AI workers based on resume data increases sales by 17.3% in the top tercile of initial firm size, 4.3% in the middle tercile, and 0.0% in the bottom tercile. Furthermore, AI investments lead to higher industry concentration measured by both the Herfindahl-Hirschman Index (HHI) and the fraction of total industry sales accruing to the single largest firm.

We outline three non-mutually-exclusive channels through which investments in AI technologies can generate the empirical patterns we observe in terms of increased firm-level growth and higher industry concentration: (i) productivity improvements; (ii) market power; and (iii) scale advantages. For the first channel, we consider whether the observed effects stem from AI invest-

⁴More recent evidence in Graetz and Michaels (2018), Aghion et al. (2020), and Fujiwara and Zhu (2020) suggests that automation can have zero or even positive effects on aggregate employment. For theoretical treatments of the impact of automation on labor displacement, see Korinek and Stiglitz (2017), Acemoglu and Restrepo (2019), and Agrawal et al. (2019). See also Brynjolfsson et al. (2018) and Athey et al. (2020) for task allocation across human and artificial intelligence.

ments increasing firms' productivity. Empirically, we do not find much support for this hypothesis, at least in the short run. We document a weak, statistically insignificant, and slightly *negative* relationship between AI investments and sales per worker or revenue total factor productivity (TFP). The second channel reflects the idea that big data and AI enable granular product tailoring that can potentially facilitate price discrimination. We test this channel empirically by evaluating whether the AI-fueled growth is driven by increased market power captured by the AI-investing firms. We find statistically insignificant and economically small effects on firm-level markups (on the order of a +/-1% change in markups for a one-standard-deviation change in AI investments).

We find that our results appear most consistent with the third channel: that AI facilitates scale advantages for the ex ante largest and most productive firms (Lashkari et al., 2018; Ayyagari et al., 2019; Autor et al., 2020). Large firms are more likely to invest in AI, and these investments allow the firms to grow even larger, with the positive effects of AI concentrating among the ex ante largest firms. In addition to the results by initial firm size, we perform two further tests of the scale advantage mechanism. First, we document that the positive effects of AI on firm sales growth are concentrated in the most ex ante productive firms, with large positive effects for firms in the highest productivity tercile in 2010 and small and insignificant effects for firms in lower terciles. Second, we document that AI investments are related to firms' expansion across both product markets and geographic regions, which is consistent with the theoretical argument by Aghion et al. (2019) that new technologies lower the overhead costs of spanning multiple markets and allow the most productive firms to expand.

Our paper contributes to the literature on the adoption and economic impact of new technologies. Specifically, we contribute to the growing literature on the advent of artificial intelligence and its impact on the economy. While a number of theories have been proposed about how artificial intelligence and big data could affect the economy (e.g., Korinek and Stiglitz, 2017; Aghion et al., 2017; Brynjolfsson et al., n.d.; Brynjolfsson et al., 2019; Farboodi et al., 2019), there has been a dearth of empirical evidence until recently due to the lack of data (Brynjolfsson and Mitchell, 2017; Seamans and Raj, 2018). Most empirical evidence to date focuses on the effect of AI on the labor market (e.g., Erel et al., 2019; Felten and Seamans (2018); Grennan and Michaely, 2019; Cowgill, 2019; Webb, 2020). Some recent papers also document that AI increases the market value of firms (Rock, 2019) and changes the knowledge production functions (Abis and Veldkamp, 2020). In related work, Alderucci et al. (2020) explore the effect of AI-related patents on firm outcomes. Our measure of AI investments using a combination of job postings and resumes allows us to estimate

the effects of AI investments in a wide range of firms and industries, with a focus on firms using AI technologies, rather than firms inventing new AI technologies.

We also contribute to the literature on the causes and consequences of increasing industry concentration. A growing literature documents the rise of concentration and market power in the U.S. (Gutiérrez and Philippon, 2017; Grullon et al., 2019; Autor et al., 2020; Barkai, 2020; De Loecker et al., 2020).⁵ Our results suggest that, while AI technologies contribute to increased concentration, they do not tend to increase firms' market power, at least in the short-run. Instead, our results lend support to the hypothesis that new technologies, such as AI, can disproportionately favor large firms by enabling the most efficient firms to scale more easily, supporting the mechanism proposed by Aghion et al. (2019). While the AI-investing firms are able to enter new markets and expand their product offerings, they do not see immediate productivity gains, nor do they charge higher markups on their products. This potentially reflects the argument by Brynjolfsson et al. (2019) that the productivity benefits of AI investments can take a long time to materialize and is consistent with evidence on previous technologies during industrialization (Juhász et al., 2020; Braguinsky et al., 2020).

Methodologically, this paper offers a new measure of firms' intangible capital based on human capital. While we employ this methodology to look specifically at investments in AI, our methodology can be applied more generally, including to other skills and technologies. Despite ongoing efforts to incorporate more comprehensive measures of intangibles at the national level (e.g., Corrado et al. 2016), most firm-level measures of intangible capital use cost items such as R&D and SG&A (Eisfeldt and Papanikolaou, 2013; Peters and Taylor, 2017). Our methodology offers a measure of intangibles that is consistent across all firms and sectors, can be defined as broadly or narrowly as needed, and is not subject to different R&D or SG&A norms across firms and industries.

The remainder of the paper proceeds as follows. Section 1 develops our main hypotheses. We introduce the two primary datasets (job postings and resumes) in Section 2 and detail our methodology for constructing the measures of AI investments in Section 3. Section 4 addresses the question of which firms choose to invest in AI, while Section 5 considers the impact of AI investments on firm growth and industry concentration. Section 6 explores the mechanisms. Section 7 concludes.

⁵Other explanations for growing concentration include increasing barriers to entry and lax antitrust enforcement (e.g., Grullon et al., 2019 and Covarrubias et al., 2019), low interest rates (Liu et al., 2020), and globalization (e.g., Elsby et al., 2013).

1 Conceptual Framework

It is an open question whether AI investments fuel firms' growth. On the one hand, as a general purpose technology, AI can be widely applicable across firms' operations, empowering the creation of new products and services or facilitating entry into new markets.⁶ This would lead to the expansion of the AI-investing firms' operations and growth in sales. On the other hand, current attention to AI investments 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., n.d.). Moreover, even if the AI investments are already benefiting firms via faster sales growth, the impact of AI on employment, which is of first-order importance (Frank et al., 2019), is ambiguous. Since AI capabilities have the potential to displace a large share of occupations, employment may decline even as output increases.

At a broader level, we are also interested in how investments in AI technologies affect industry concentration. Following a well-documented increase in industry concentration over the past several decades (Furman and Orszag, 2015; Grullon et al., 2019), an active debate emerged about the causes and consequences of the increase in industry concentration (see Syverson, 2019 and Covarrubias et al., 2019 for a review). One proposed channel driving this and other important trends, including the decrease in the labor share, is improvements in information technology (IT) (e.g., Karabarbounis and Neiman, 2014; Crouzet and Eberly, 2019; Lashkari et al., 2018; Aghion et al., 2019). Empirically, Bessen (2017) finds a positive relationship between the rise in industry concentration and the use of proprietary IT systems in the U.S. He stresses that the scalability of IT is advantageous to firms that are already large. However, in the case of AI, the effect on industry concentration is ex ante ambiguous. On the one hand, AI can democratize adoption among smaller firms. Unlike proprietary IT systems that require large upfront investments, AI implementation is largely dependent on human expertise, with auxiliary investments centering on data storage and computing, which can be purchased from specialized providers on a per use basis and hence do not require lumpy upfront investments (Organisation for Economic Co-operation and Development, 2015; Van Ark, 2016). On the other hand, big data and AI technologies have scale effects that favor large firms and industry leaders with large amounts of data, which can contribute to the increase in industry concentration and winner-take-all dynamics (Farboodi et al.,

⁶For example, the JPMorgan Chase 2017 annual report highlights the broad applicability of AI: "Artificial intelligence, big data and machine learning are helping us reduce risk and fraud, upgrade service, improve underwriting and enhance marketing across the firm." A detailed case study of JPMorgan's use of AI technologies is provided in Appendix A.1.

2019).

Below, we discuss specific channels through which investments in artificial intelligence can lead to firm growth and higher industry concentration. First, as a technological advancement, AI can potentially stimulate growth by streamlining production processes and increasing productivity. Second, some of the specific applications of AI could be used to price discriminate and increase firms' market power. Third, AI has unique features, such as reliance on big data, that can favor ex ante larger, more productive firms and increase the scale of these firms. While all three channels can lead to higher growth of individual firms and increases in industry concentration, they have different implications for competitive industry dynamics and varying predictions for firm-level productivity and markups, leading to different economic implications associated with "good" vs. "bad" concentration (Covarrubias et al., 2019).

1.1 AI as a Driver of Productivity Growth

Technological innovations aim at streamlining operations and improving productivity. For example, previous waves of information technologies brought about significant productivity improvements through a number of channels (Bartel et al., 2007). When it comes to AI, the technology can increase productivity in at least two ways. First, AI can potentially replace or augment human labor for some tasks (Webb, 2020; Agrawal et al., 2019) and cut per-unit labor costs. Empirically, this appears to be the case for other recent technological innovations such as automation and robots (Acemoglu et al., 2020; Graetz and Michaels, 2018). Second, big data and AI can increase efficiency through better forecasting (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.⁷ The forecasting applications of artificial intelligence are indeed prevalent in the data. Chase (2013) draws on executive interviews to point out that big data can streamline demand forecasting, leading to more efficient inventory management. Our detailed resume data 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; ExxonMobil invests in AI for assessment and mitigation of risks in oil exploration; and General Electric uses AI in decision support for jet engine preventive maintenance. For further examples, detailed case studies of the applications of AI in

⁷Their empirical evidence supports the prediction that forecast accuracy is associated with higher profitability among Japanese firms. Relatedly, Brynjolfsson et al. (2011) use detailed survey data from 179 large public firms to document a positive link between the use of "data driven decision making" and firm-level output and productivity.

four firms across four different industries—UnitedHealth Group, JPMorgan Chase, Caterpillar, and Qualcomm – are provided in Appendix A.1.

This varied set of potential productivity improvements can manifest through greater productivity of the standard labor and capital inputs, as well as through changes in the production function that reflect data as an input. Empirically, in a Cobb-Douglas framework, AI-induced productivity gains should translate into increases in measured total factor productivity (TFP). Under more general production functions, higher productivity should also be reflected in higher measured output per worker. Specific improvements (e.g., automation of tasks previously performed by human labor) can be further reflected in decreased costs or employment.

Prediction 1 (Productivity Channel). *Productivity gains from artificial intelligence would be observed empirically as:*

1. *AI-investing firms see increases in productivity (TFP) and labor productivity (sales per worker).*
2. *AI-investing firms potentially see decreases in costs or employment.*

1.2 AI as a Driver of Price Discrimination and Market Power

Three of the most notable applications of artificial intelligence to date are: (i) targeting and pricing of online ads, (ii) tailoring product offerings to customers' tastes, and (iii) using consumer data to price products. For example, a consumer product company may use machine learning to build more specialized products tailored to certain customers, thus shielding themselves from competition and being able to charge higher prices. Mihet and Philippon (2019) highlight the uses of artificial intelligence by companies such as Amazon to improve matching of products with consumers and to deliver more tailored recommendations, which in turn can enable firms to set prices based on a large number of features including consumers' demographics (Varian, 2018). As a result, firms can price discriminate by offering contracts based on personalized customer information (Brunnermeier et al., 2020). For example, data on individual behaviors such as web browsing history can enable even better approximations of individual demand functions than pure demographics and can potentially lead to large heterogeneity in prices charged to different consumers (Shiller, 2016).⁸

⁸At the same time, the transparency of online pricing may limit firms' ability to successfully price discriminate (Cavallo, 2017; Ater and Rigby, 2018).

If this effect is present, then the ability to tailor products (thereby making them less substitutable) and price discriminate would grant greater market power to firms investing in AI, enabling them to extract more consumer surplus. Empirically, this greater market power would appear in the form of higher markups charged by the AI-investing firms (Syverson, 2019).

Prediction 2 (Market Power Channel). *Applications of artificial intelligence that enable price discrimination would generate the following effects:*

1. *Investments in AI lead to higher firm-level markups.*

1.3 AI as a Driver of Scale Advantages

As an information good, AI can also have scale effects that would facilitate higher growth of large firms and industry leaders even without increases in productivity or market power. If this is the case, then AI can contribute to the rise of superstar firms, as in (Autor et al., 2020).

Big data and artificial intelligence are intangible assets (Mihet and Philippon, 2019). Crouzet and Eberly (2019) highlight that intangible assets are more “scalable” than physical capital, and De Ridder (2019) conceptualizes intangible assets as a shift towards fixed costs and away from variable costs. In particular, successful implementation of AI technologies relies strongly on data availability (Fedyk, 2016), and Farboodi et al. (2019) point out that there is a positive feedback loop between firm size and the firm’s data assets, driven by the fact that data are a “by-product of economic activity”. This makes data and AI skilled labor complementary inputs, where breakthroughs in AI technologies can enable firms with extensive datasets to produce output that was not feasible previously. As a result, firms’ investments in AI technologies can help mitigate diminishing returns to data inputs (Abis and Veldkamp, 2020) or even increase the returns to scale, akin to what Lashkari et al. (2018) suggest for IT technologies more generally.

AI can also directly increase firm scale by allowing the most efficient firms to expand more easily across different markets. Aghion et al. (2019) argue that when there are lower overhead costs of spanning multiple markets due to new technologies like IT, high-productivity firms expand into new markets, leading to higher sales and market share for these firms. Relatedly, Hsieh and Rossi-Hansberg (2019) model the ongoing transformation of the services industry as technological innovation that increases fixed costs but reduces variable costs, allowing productive firms to extend their boundaries to new geographical markets. In our data, we see that similar methods and skillsets are often used to leverage AI in different business segments, and when firms begin invest-

ing in AI, they tend to do so concurrently on multiple fronts. For example, investments in AI by the construction manufacturing firm Caterpillar Inc. range from using techniques from computer vision for part recognition to credit scoring for machinery financing (see Appendix A.1). Some of these innovations include expansion 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 manufacturing of “smart” machinery at Caterpillar. This underscores the argument made by [Braguinsky et al. \(2020\)](#) that new technologies help firms overcome supply-side constraints and grow through product innovations.

Overall, AI can create scale advantages by changing the cost structure (e.g., reducing the costs of spanning multiple markets or moving towards a greater reliance on fixed costs) or by directly impacting the production function (e.g., increasing returns to scale or using data and AI labor as complementary inputs). Empirically, all of these channels predict that larger firms should both (a) be more willing to invest in AI and (b) benefit more from AI investments. At the industry level, firms’ ability to scale more easily can lead to greater market share accruing to the most ex ante productive firms and a winner-takes-most phenomenon. At the same time, scale advantages from AI would not necessarily be visible in empirical measures of productivity; for example, if AI helps reduce the overhead costs of spanning multiple markets or creating new products, firms will expand their operations but will not necessarily have higher productivity in any given market.

Prediction 3 (Scale Advantages Channel). *Scale advantages from AI generate the following predictions:*

1. *Larger firms are more likely to invest in AI.*
2. *AI investments lead to the largest and most productive firms growing more and accruing greater market share.*
3. *AI investments are associated with expansion into new markets.*

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 in turn 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. Burning Glass examines more than 40,000 online job boards and company websites to aggregate the job postings data, parse them into a systematic, machine-readable form, and create labor market analytic products. The company employs a sophisticated deduplication algorithm to avoid double counting vacancies that post on multiple job boards. [Hershbein and Kahn \(2018\)](#) provide a detailed description of the BG data.

The 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 posting.⁹ The main advantages of the BG dataset are the breadth of its coverage and the detail of the individual jobs in the sample. The dataset captures a near-universe of jobs that were posted online and covers approximately 60–70% of vacancies posted in the U.S., either online or offline ([Carnevale et al., 2014](#)). The broad coverage of the data presents a substantial advantage over datasets based on a single vacancy source, such as CareerBuilder.com. [Hershbein and Kahn \(2018\)](#) show that the representativeness of Burning Glass is stable over time at the occupation level. In other words, although BG slightly over-represents some occupations relative to the U.S. Census Current Population Survey, the degree to which these occupations are over-represented does not change over the sample period.

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 are linked to at least one skill.¹⁰ 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 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.¹¹ Out of 112

⁹For example, a job posting might ask for a worker who is bilingual or who can organize and manage a team. BG cleans and codes these and other skills into a taxonomy of thousands of unique but standardized requirements. Beginning with a set of predefined possible skills, BG does a fuzzy search of each job posting’s text for an indication that any given skill is required. For example, for team work, the BG algorithm searches for the key words “team work” but also looks for variations such as “ability to work in a team.”

¹⁰The job postings with missing employer names are primarily those listed on recruiting websites that mask the employers’ identities.

¹¹Observations without an exact match generally fall into the following categories: (i) the employer name in BG

million job postings with non-missing employer names and skills, 42 million (38%) are matched to Compustat firms. This is consistent with the fact that publicly listed firms constitute about one-third of U.S. employment in the non-farm business sector (Davis et al., 2006).

2.2 Employment Profiles from Cognism

While job postings data provide an important look at the firms' demand for certain types of employees, vacancies data represent just one aspect of a firm's adjustment of labor inputs: stated, but not necessarily realized, demand (Hershbein and Kahn, 2018), which is a less complete view of the firm's labor inputs than the actual employees. To address this concern, we complement our job postings data with comprehensive information on the actual individuals employed at each firm. To do so, 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.

For each individual in our sample, we have the following general information: a unique identifier, location (city and country), an approximate age derived from the individual's educational record, gender classified based on the first name, and a short bio sketch (where provided). 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 (where provided). Similarly, each education record includes start and end dates, the name of the institution, and the degree (major). In addition, individuals may volunteer self-identified skills and list their courses attended, certifications, patents, awards, and publications.

We take several steps to disambiguate self-reported employer names in the profile data to the names of publicly traded firms. First, we follow the procedure outlined in Fedyk and Hodson (2019): (i) begin with a comprehensive list of publicly traded companies from the exchanges (NASDAQ, NYSE) and common datasets (CRSP and Compustat), (ii) strip out common endings (e.g., "Inc" and "L.P."); (iii) run a fuzzy matching algorithm from the self-reported employer names to the official company names; and (iv) augment the algorithm by mapping the self-reported employer names to semantic entities in the WikiData project. In addition, for the set of Compustat firms that are not mapped to any companies in the employment data, we perform a manual attempt at finding matches. Similarly to our procedure matching BG job postings to Compustat matches to several firms in Compustat; (ii) the entries contain extra words (e.g., "X company" vs. "X company international").

firms, we manually check candidate non-exact matches. Of the 462 million U.S.-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 one third of overall U.S. employment being accounted for by publicly listed firms, with the additional exclusion of cross-listed firms with headquarters outside of the U.S. The sample of 86 million person-firm-years matched to U.S. public firms accounts for 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) and industry-level wage, employment, and price data from the Bureau of Labor Statistics (BLS) and the Census Quarterly Workforce Indicators (QWI). Operational variables such as sales, employment, assets, net income, cash, costs of goods sold, operating expenses, R&D expenditures, and SG&A expenses come from annual accounting data available through Compustat.

3 Methodology: AI Investments by Firms

In this section, we introduce a new methodology to proxy for firm-level AI investments with job postings data. We provide summary statistics of our AI investments measure and validate the (job-postings-based) measure using resume data. The two measures—based on job openings versus current employees—display analogous trends over time and across industries, and show high correlations with each other.

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 methodology for identifying AI-related jobs. Previous work classifies job postings based on the presence of key terms from a pre-specified list.¹² This approach presents significant measurement challenges, as word lists are highly subjective and are likely to suffer from both

¹²For 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](#)).

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. In a quickly-evolving domain such as AI, identifying an accurate and complete set of search terms is especially challenging, as newer emerging skills can easily be missed. Our methodology circumvents these challenges by not requiring researchers to impose any subjective assessments on which skills are AI-related ex ante, instead learning the AI-relatedness of each of approximately 15,000 unique skills directly from the job postings data, based on their empirical co-occurrence 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.

We consider four core AI skills: Artificial Intelligence (AI), Machine Learning (ML), Natural Language Processing (NLP), and Computer Vision (CV). For each skill s , we define a metric that captures the relatedness of that skill to these four core AI technologies (AI, ML, NLP, and CV):

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

Intuitively, this measure captures how related 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, requiring “Tensorflow” as a skill for a job is indicative of a job being AI-related. On the other hand, the skill “Communication” is required in a large share of jobs across the board, so its AI-relatedness measure is only 0.03%, and having the “Communication” skill does not indicate that the job is AI-related.

We define the job-level AI-relatedness measure for a given job posting as the mean (skill-level) measure across all required skills listed in that job posting. Specifically, letting N denote the number of required skills listed for job posting j , our job-level AI-relatedness measure is:

$$\omega_j^{AllAI} = \frac{1}{N} \sum_{s=1}^N w_s^{AllAI}$$

To further refine our measure and screen out general skills (e.g., the programming language “R”, which has a value of 0.25), we manually categorize all skills that have an AI-relatedness measure w_s^{AllAI} above 0.05 and that 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

ex-ante important AI-skills. There are about 700 skills with $w_s^{AllAI} > 0.05$ and at least 50 jobs. We group these skills into nine categories: computing (e.g. GPU), data (e.g. NoSQL), general programming (e.g. Python), AI software (e.g. Tensorflow), AI methodology or algorithm (e.g. supervised learning), AI application (e.g. Chatbot), AI core (AI, ML, NLP, CV), statistics (e.g. linear regression), and other. Of these, we consider “computing”, “data”, “AI software”, “AI methodology or algorithm”, “AI application”, and “AI core” as narrow AI skills, while “general programming”, “statistics”, and “other” represent more general skills.

We can decompose the job-level measure ω_j^{AllAI} into components corresponding to these categories:

$$\begin{aligned}\omega_j^{AllAI} &= \omega_j^{computing} + \omega_j^{data} + \omega_j^{AIsoftware} + \omega_j^{AImethodology} + \omega_j^{AIapplication} + \omega_j^{AIcore} \\ &\quad + \omega_j^{programming} + \omega_j^{other} + \omega_j^{statistics} + \omega_j^{w<0.05} \\ &= \omega_j^{NarrowAI} + \omega_j^{programming} + \omega_j^{statistics} + \omega_j^{other} + \omega_j^{w<0.05}\end{aligned}\tag{1}$$

For the remainder of the paper, we use $\omega_j^{NarrowAI}$ as the primary continuous measure of the AI-relatedness of jobs. We transform this continuous measure into a discrete indicator by defining each job posting j as AI-related if and only if the measure $\omega_j^{NarrowAI}$ is above 0.1, which upon manual examination captures the majority of AI-related job postings. The firm-level measure $Share_{f,t}^{NarrowAI}$ is the fraction of job postings by firm f in year t whose $\omega_j^{NarrowAI}$ measure exceeds the 0.1 threshold. The reason for using a discrete classification of each job as AI-related or non-AI-related is twofold. First, this increases the interpretability of the firm-level measure, which captures the share of the firm’s employees classified as AI workers. Second, we apply the same binary classification to the firm’s current employees when constructing the resume-based measure in Section 3.2, leading to consistent approaches and interpretations across the two datasets. In Section 5.1, we confirm that our key results are robust to averaging the continuous narrow-AI measure ($\omega_j^{NarrowAI}$) or the continuous measure using all skills (ω_j^{AllAI}) across all jobs in a given firm.

3.1.1 Summary Statistics

Our firm-level measure displays intuitive properties: the overall share of AI workers rises over time, is highest in the “Information” sector, and is driven by relevant job titles such as “Artificial Intelligence Researcher” and “Deep Learning Engineer”.

The trends over time are displayed in Figure 1. The average continuous job-level measure $\omega_j^{NarrowAI}$ starts out close to zero, at 0.02%, at the beginning of the sample in 2007. The average AI-relatedness of job postings rises monotonically over time, with the increase speeding up from 2014 to 2018. The measure peaks at 0.2% at the end of the sample in 2018.¹³ The increase in AI jobs is present across industries, as can be seen in Figure 2, in line with the notion that AI is a general purpose technology (Goldfarb et al., 2019). This figure plots the average AI-relatedness measure of job postings in each of the NAICS sectors, separately for the years 2007–2014 and 2015–2018. The figure highlights that AI investments are highest in the “Information” sector, growing from 0.19% in the early years of 2007–2014 to 0.49% in the later period of 2015–2018. AI investments in nearly all sectors show a substantial (two- to six-fold) increase from the earlier to the later period. The heterogeneity in AI investments across industries is consistent with supply-side arguments made in industry reports (Bughin et al., 2017): AI adoption across industries is largely driven by availability of data and technical capabilities, which are crucial inputs for the AI production function. At the same time, the ability of our measure to pick up AI investments in the broad cross-section of economic sectors highlights a key advantage of our human-capital-based measure, which does not rely on specialized outputs such as patents.

Additional checks on the data confirm that our measure is indeed capturing the essence of AI investments by firms. For example, Table A.1 shows that the job titles associated with the highest job-level measure of AI-relatedness, $\omega_j^{NarrowAI}$, are all very relevant postings: “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). Similarly, Table A.2 shows that the job titles that contribute the highest number of AI-related job postings (i.e., job postings with the $\omega_j^{NarrowAI}$ measure above 0.1) are relevant titles such as “Data Scientist”, “Senior Data Scientist”, “Software Engineer”, “Principal Data Scientist”, and “Data Engineer”. AI jobs are also concentrated in the BLS occupations “Computer and Information Research Scientists”, “Software Developers, Applications”, and “Computer Occupations” (Table A.3).

Throughout our empirical analyses, we consider only 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

¹³In 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 we manually categorize as “Narrow AI” skills instead of other less AI-specific skills (e.g. statistics or programming languages) or skills that have AI-relatedness measures below 0.05.

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.

3.2 AI Investments from Resumes (Cognism)

We validate our job-postings-based measure of firms' investments in AI against an analogous measure using profiles of all firm employees with available resume records. This helps address concerns that the job postings data are not fully representative of firm activities—for example, if a firm is not able to hire despite active job postings, or if a firm posts numerous job openings due to high employee turnover. This type of data issue does not appear to drive our job-postings-based measure. Instead, the measure using resumes displays very similar trends (e.g., across industries and over time) to the measure using job postings and the two measures are highly correlated, as we show in Section 3.3.

We use resume data to identify AI-related employees as those whose current positions directly involve AI through a holistic approach covering each person's entire profile. We begin with the set of keywords that are classified as having an AI-relatedness measure w_s^{AllAI} above 0.70 in the Burning Glass skills data. This includes 73 terms that are most relevant for AI-skilled jobs. We then search for these terms in every employment record of each individual in the resume data. Specifically, for each particular employment 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; (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 employee.

After classifying each individual at each point in time, we aggregate the number of AI and all other employees up to the firm level: for each firm in each year, we compute the percentage of employees of that firm in that year who are classified as AI-related. Our firm-level measures focus 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., and for firms listed on European exchanges, we look at the employees based in corresponding European countries.¹⁴

¹⁴For firms cross-listed on U.S. and European exchanges, we assign firms to a single region based on their headquarters. For example, Nokia, listed on both NASDAQ Helsinki and the New York Stock Exchange, is captured as a European firm headquartered in Finland, and we consider only its employees based in Finland. We exclude cross-listed

For all analysis involving the resume data, we exclude firms that have fewer than 50 employees in every year to reduce noise from small firms or firms with poor coverage.

3.2.1 Summary Statistics

The general patterns of AI investment are very similar using the resume-based measure and the job-postings-based measure. Figure 3 displays the accelerating time trends of the resume-based measure, plotting the fraction of all employees in each year who are classified as AI-related. Figure 4 shows the distribution of the fraction of AI employees across industries, separately for the 2007–2014 and the 2015–2018 sub-periods. The rise in the fraction of AI employees (Figure 3) is very similar to the contemporaneous rise in AI job postings (Figure 1), beginning very close to zero, at 0.03%, in 2007 and reaching 0.24% in 2018. Analogously to the job-postings-based measure, the resume-based measure increases from earlier (2007–2014) to later (2015–2018) in the sample period for all sectors.

3.3 Correlations between AI Investment Measures from Job Postings (Burning Glass) and Resumes (Cognism)

The two proxies of AI investments—using job postings and resumes—are highly but imperfectly correlated. Table 1 displays the cross-sectional correlations between the two measures at the firm level for each year when job postings data are available, {2007,2010–2018}. 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 jobs in the Cognism resume dataset against the average continuous measure of AI-relatedness from the Burning Glass job postings data, $\omega_j^{NarrowAI}$. Panel 1 presents Pearson correlations, while Panel 2 tabulates Spearman (rank) correlations.

The correlations are quite high 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. The relationship between the AI investment measures computed from the two datasets is weaker in the earlier part of the sample: while the absolute numbers of AI jobs still display a Pearson correlation of 0.65 in 2007, the correlation for fractions is only 0.13. Given the low correlation in the measures in 2007, we limit all of our analyses to the BG data available from 2010 onward.

firms with headquarters outside of the U.S. from our baseline analysis sample.

4 Which Firms Invest in AI?

We document differential patterns of investments in AI technologies across firm characteristics. First, larger firms—with more employees, higher sales, and larger market share—see greater levels of AI investments between 2010 and 2018. Second, firms with higher markups tend to invest in AI more aggressively. Third, firms with higher cash reserves are also more likely to invest in AI technologies. Geographically, AI investments tend to concentrate in locations with higher wages and more educated workers.

4.1 AI Investments and Firm Characteristics

The firm-level AI investment patterns are presented in Table 2. Panel 1 displays the results for the changes in the number of *actual* AI employees using the Cognism resume data. Panel 2 presents the results for changes in the firms' *demand* for AI talent using the job postings data from Burning Glass. Since our focus is on understanding the impact of the general use of AI technologies on firms, we exclude firms in the information technology (IT) sector from our main empirical analyses in this and the following sections. The firms in IT sectors are likely to be inventors of new AI technologies or suppliers of AI solutions, in which case their hiring of AI employees would serve different purposes and lead to different dynamics than in other sectors.¹⁵ For each measure of AI investment, we estimate the following specification:

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

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 firm variables of interest: 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

¹⁵IT sectors include NAICS 2-digit industry codes 51 (Information) and 54 (Professional, Scientific, and Technical Services). In untabulated 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. \(2020\)](#).

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 8.¹⁶ Column 9 includes all variables in a multivariate specification, except for employment and market share, which are highly correlated with sales. All independent variables and controls are measured as of 2010. 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 resumes (job postings) in 2010.¹⁷

The results reported in Table 2 highlight that larger firms experience higher levels of AI investment. 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, which is consistent with contemporaneous work of [Alekseeva et al. \(2020\)](#), who use Burning Glass data to measure firms’ AI demand. While overall return on sales is not predictive of future AI investments, the COGS-based markups positively predict future AI investments in the Cognism data.

Importantly, the results for firm-level demand for AI talent measured with Burning Glass data are very consistent with the patterns using actual firm-level hiring of AI talent from 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.

4.2 AI Investments and Local Conditions

We now turn to examining how AI investment patterns relate to conditions at the local level, which further helps validate our measure of firm-level AI investments based on AI-skilled human

¹⁶We 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 $\mu_{it} = \theta_{st} \frac{Revenue_{it}}{VariableCost_{it}}$, 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 $\log(Revenue_{it}/VariableCost_{it})$.

¹⁷Since 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 or not using any weights.

capital. AI investment rates are higher in locations with a highly-educated and high-wage workforce, consistent with AI-skilled labor being the most critical input to successful deployment of AI programs (Bughin et al., 2018). This contrasts with investments in robotics, which concentrate in areas with larger shares of manufacturing employment (Acemoglu et al., 2020).

We document this empirically by estimating the following specification, which aggregates firm-level AI investments for all Compustat firms to a commuting-zone-level measure and links it to 2010 commuting zone characteristics:

$$\Delta ShareAIWorkers_{i,[2010,2018]} = \alpha + \beta CommutingZoneVariable_{i,2010} + \epsilon_i, \quad (3)$$

where $\Delta ShareAIWorkers_{i,[2010,2018]}$ measures the change in the share of AI workers in commuting zone i from 2010 to 2018.¹⁸ The independent variable, $CommutingZoneVariable_{i,2010}$ is either the log average wage or the share of college-educated workers for commuting zone i in 2010, calculated from the Census American Community Survey.

Figure 5 (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. The observations are weighted by the population as of 2010. We find a strong positive relationship between the average wage of a region and the local growth in AI jobs—the R-squared of the regression is 0.43. Figure 5 (b) plots the relationship between the change in the share of AI workers from 2010 to 2018 and the share of college-educated workers in 2010. We see a very similar pattern: growth in AI workers is concentrated in commuting zones with a large fraction of college-educated workers. These results highlight that regional variation in the available skilled labor is an important determinant of local AI investments, suggesting that it is important to control for the workforce composition of the local labor markets when examining the impact of AI investments on economic outcomes. Finally, Figure 6 displays a heat map of the growth in the average job-level AI measure from 2010 to 2018 and shows that there is significant variation in AI investments across commuting zones.

¹⁸The share of AI workers is based on job postings data from Burning Glass, because the coverage of job-level location information in job postings is better than in resumes, which tend to report an employee’s current location rather than the location of all prior job records.

5 AI Investment, Growth, and Industry Concentration

We present our main results on the effects of AI investments on firm- and industry-level growth. We document that firms and industries investing in AI technologies grow faster, and that this result is robust to using two distinct instruments for AI investments in an instrumental variables (IV) strategy. The positive effect of AI investments on growth is concentrated among the largest firms, which leads to an increase in industry concentration.

5.1 Firm Growth

We begin the analysis by exploring the impact of AI investments on the growth of individual firms, first in OLS regressions and then in the IV setting.

5.1.1 Ordinary Least Squares (OLS) Results

We examine the relationship between firms' investments in artificial intelligence from 2010 to 2018 and a number of measures of firm-level growth over the same time period. Given that we do not expect to observe the effects of AI investments on firms immediately, we use a long-differences specification. In Table 3, we report the estimates from the following regression:¹⁹

$$\Delta FirmVariable_{i,[2010,2018]} = \beta \Delta ShareAIWorkers_{i,[2010,2018]} + Controls'_{i,2010} \gamma + IndustryFE + \epsilon_i, \quad (4)$$

where the main independent variable, $\Delta ShareAIWorkers_{i,[2010,2018]}$, captures the change in the share of AI workers in firm i from 2010 to 2018, standardized to mean 0 and standard deviation 1. $IndustryFE$ 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 controls that are all measured at the start of the period in 2010: (a) 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) a rich set of 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

¹⁹In the baseline regressions, we include only firms that are observed in the Compustat sample both in 2010 and 2018. We consider entry and exit in industry-level results in Section 5.2.

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, which might independently affect the AI employment share or be correlated with investments in other technologies.²⁰

In columns 1 and 2 of Table 3, the dependent variable is the firm-level change in log sales from 2010 to 2018.²¹ Both measures of AI investment are associated with a significant and economically large increase in sales growth: a one-standard-deviation increase in the share of AI workers predicts an additional 12% to 16% growth in sales, depending on the specification. In columns 3 and 4, we examine how AI investments are associated with changes in employment. The effect of AI investments on employment is of particular interest, since its sign is *ex ante* ambiguous, with AI having the potential to displace a large fraction of jobs, as discussed in Section 1. 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 predicts an 11% to 16% increase in firm-level employment.²² 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.

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.4 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-AI-producing sectors. Over-

²⁰For example, the share of routine workers is correlated with adoption of automation technologies (Autor and Dorn, 2013), and the share of workers in IT-related occupations is correlated with general information technologies. When firms span multiple commuting zones, we calculate these variables as the weighted average, using the number of BG job postings in each commuting zone as weights. The results are similar in magnitude and economic significance if we only include firm-level controls enumerated in list (a).

²¹The sample size is smaller than the regressions in the previous section because: first, we estimate all columns on the sample with all non-missing control variables to keep the sample composition stable. Since this requires that firms have jobs in the BG data (to calculate the geographical controls), it leaves us with a sample of around 1100 firms; second, we include only firms that exist in Compustat in both 2010 and 2018, which reduces the number of firms further to around 800 for both measures. The results of the regressions without controls are similar if we estimate based on the entire available sample.

²²In unreported results, we use the number of workers in each firm in the Cognism resume data as an alternative measure of employment, and find a similar positive effect of 10%-15%.

all, we observe that investments in AI are associated with economically significant increases in firms' operations, and these effects are meaningful across all economic sectors.

5.1.2 Event Study Analysis

In this section, we conduct an event study analysis, which allows us to investigate the timing of the positive effects on sales and employment relative to the AI investments. In addition, this design helps to rule out reverse causality explanations and address concerns about AI-investing firms being on differential growth trajectories.

Firms tend to invest in AI on a continuous basis from year to year, rather than make a lumpy investment in a single year.²³ Therefore, we estimate the dynamic effects of AI investments using a distributed lead-lag model similar to [Aghion et al. \(2020\)](#) and discussed in [Stock and Watson \(2015\)](#). In particular, the model is specified as:

$$Y_{it} = \sum_{k=-2}^5 \delta_k \Delta \text{ShareAIWorkers}_{i,t-k} + \mu_i + \lambda_{st} + \epsilon_{it} \quad (5)$$

where $\Delta \text{ShareAIWorkers}_{i,t-k}$ 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 1, 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 firm size 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 (5) 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.²⁴ 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.

²³The percentage of AI-investing firms that only invest in a single year is 29.5%, compared to 70.6% for robots ([Humlum, 2019](#)).

²⁴Since the data on AI investments end in 2018, we include only two leads so that we can keep all firm-year observations up to 2016. We obtain similar results when including only one lead or no leads at all. Furthermore, we focus the event study analysis 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.

Figure 7 reports the coefficients from the event study regression. The top panel shows that sales increase following AI investments. The effect is not immediate, and it takes two 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 even five years after the investments. This is consistent with the long-differences estimates in Section 5.1.1.²⁵ 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 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 the potential set of confounders to the set of shocks that occur simultaneously with AI investments. We further validate the causal interpretation of the estimates using two IV strategies in the next section.

5.1.3 Instrumental Variables (IV) Results

In this section, we further explore the causal effect of AI on firm growth using two instrumental variables (IV) strategies. Although we control for factors that predict firms' investments in AI, the OLS estimates may still be biased if there are other omitted variables that correlate with both AI investments and firm growth. The direction of the bias is *ex ante* ambiguous. On the one hand, the OLS estimates may be biased upwards if the firms investing in AI are already on a faster growth trajectory or have more efficient managers who concurrently improve other aspects of the firm. On the other hand, firms that have fewer growth opportunities or anticipate negative shocks may have lower opportunity costs to engage in large-scale innovations and adopt new technologies (Bloom et al., 2013b), leading to a negative correlation between the endogenous choice to adopt AI and future growth. Furthermore, measurement error can attenuate the estimated coefficients in OLS regressions, which creates further downward bias. Leveraging two distinct instruments, our IV results are consistent with those from the OLS regressions, lending further credence to the estimated link between AI investments and firm growth.

²⁵In long-differences estimates, we find that a one-standard-deviation increase in AI investments over 8 years is associated with a 12-16% increase in sales, which is in line with the magnitude of the event study effects scaled up to 8 years.

Foreign Industry Instrument. We instrument for a firm’s change in the share of AI workers using the change from 2010 to 2018 in the share of AI workers of foreign firms in the firm’s 5-digit NAICS industry, which is the finest level at which our data provide good coverage of foreign firms. This instrument is similar to those used by [Autor et al. \(2013\)](#) and [Acemoglu and Restrepo \(2020\)](#), where a weighted average of US industry-level changes is instrumented by a weighted average of changes in corresponding foreign industries. The identifying assumption underlying this strategy is that the common within-industry component of the rising share of AI workers is due to differences in data availability and technological feasibility across industries. 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.²⁶ We focus on European firms for three reasons. First, similar to the U.S., Europe has experienced a surge in AI investment in recent years. Figure [A.6](#) shows that Europe has the largest increase in the share of AI workers between 2010 and 2018 outside of the U.S.²⁷ Second, Cognism has good coverage of resumes in Europe, while the coverage is sparser for other countries (e.g. China). Third, focusing on Europe helps to capture inherent cross-industry heterogeneity that is likely similar to that in the U.S. because, as a developed economy, Europe is likely to have similar technological development and data availability to the U.S. The main threat to the validity of our instrument is that industries with higher rates of AI investment (in the United States and Europe) could be on an upward trend because of, for example, increasing demand or other positive shocks. In Figure [A.7](#), we show that changes in U.S. industry-level prices from 2010 to 2018 are weakly negatively correlated with industry-level growth in the share of AI employees in Europe, which suggests that our measure is not correlated with positive demand shocks to industries. We also discuss pre-trend tests to further address this concern at the end of this subsection.

We begin by reporting the first stage of the IV regressions: the relationship between the instrument and our two measures of AI investments. The results are reported in columns 1 and 2 of Table [4](#). Panel 1 displays the results using the resume-based measure of AI investments, and

²⁶The European exchanges include 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. Similar to our US Cognism analysis, we include firms with at least 50 European employees in the Cognism data in at least one year. We exclude European workers of U.S. multinational firms that are listed in the U.S. and the non-European workers of European firms.

²⁷Europe also has the third largest number of AI researchers right behind the U.S. and China. See [here](#).

Panel 2 uses the job-postings-based measure. As in OLS regressions, we control for 2-digit NAICS industry fixed effects. The specification in column 1 includes only industry fixed effects, while column 2 also includes the same set of controls as in the OLS regressions. The instrument has a strong first stage for both AI measures, with F-statistics ranging from 18 to 39, depending on the specification.

Next, we report the instrumented effect of AI investments on firm growth in columns 3 through 6 of Table 4. We consider the same set of outcomes and controls as in the OLS regressions, with the outcome being the change in log sales in columns 3 and 4 and 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. A one-standard-deviation increase in the resume-based measure of AI investments corresponds to a 42% increase in sales and a 36% increase in employment. Similarly, a one-standard-deviation increase in the job-postings-based measure translates into a 19% increase in sales and a 23% increase in employment. Overall, the findings are consistent across the OLS and IV specifications with higher point estimates in the IV regressions potentially due to lower opportunity costs of innovation and lower growth prospects of firms investing in AI. It's important to note, however, that the OLS and IV coefficients are not statistically different. This suggests that the difference between the point estimates could also be driven by estimation error.

To address the concern that industries (in the U.S. and Europe) with larger investments in AI are on an upward trend prior to these investments, we evaluate the pre-trends in Table A.5. Specifically, we regress the changes in the outcome variables of interest (log sales and log employment) from 1999 to 2007 on subsequent changes in the two measures of AI investments from 2010 to 2018.²⁸ The coefficients are all statistically insignificant and flip signs across different AI measures, indicating that industries that invest more in AI are not on different pre-trends. Table A.6 goes one step further and directly controls for the 1999-2007 changes in sales and employment at the industry level and firm level on the right-hand side of our baseline IV specifications. Including these controls has little impact on the estimates of the effect of AI investments.

Geographic Shift-Share Instrument. In an alternative instrumental variables strategy, we employ a shift-share instrument for a firm's change in the share of AI workers using a weighted average of U.S. national industry-level changes in the share of AI workers, where the weights are

²⁸We 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.

given by the industry employment shares at the locations of the firm’s operations. Let s_{ic} denote the share of U.S. employment of firm i that falls within commuting zone c , approximated by the share of job postings of firm i in 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 similar to a Bartik-style shift-share instrument (Bartik, 1991), which would correspond to instrumenting firm-level AI investments by a weighted average of the national industries’ average AI investments. The difference from a canonical Bartik instrument is that our instrument not only assigns positive weights to the industries a firm operates in, but also assigns positive weights to the industries geographically close to the firm. For example, a technological shock that enables firms in the finance industry to adopt AI would not only affect financial firms in New York City, but also positively affect non-financial firms in New York City because of the increase in AI-skilled labor supply in the city.

The identifying assumption is that firms’ industry shares ($\sum_c s_{ic} \sum_j \theta_{cj}$) are uncorrelated with errors in changes in firm sales, employment, and market shares. In other words, whether a firm is geographically close to industries that subsequently invest in AI should be pre-determined and uncorrelated with changes in firm outcomes. Goldsmith-Pinkham et al. (2019) 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.7 and find no relationship between future shocks to AI investments and past changes in firm outcomes. Specifically, we regress the changes in the outcome variables of interest (log sales, log employment, and market share) from 1999 to 2007 on the subsequent weighted industry-level changes in the two AI investment measures from 2010 to 2018. Panel 1 of Table A.7 considers the resume-based measure of AI investments, while Panel 2 uses the job-postings-based measure. All pre-trends are statistically insignificant.

In Table 5, we report the results from the shift-share IV for both measures of AI investments (the resume-based measure in Panel 1 and the job-postings-based measure in Panel 2). 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 log sales in columns 3–4, changes in log employment in columns 5–6, and changes in market share within the 5-digit NAICS industry in columns 7–8. The results of the first stage regressions show that the instrument’s F-statistic ranges from 15 to 35 for our resume-based measure. The second stage regressions 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. The instrumented AI investments also have a positive effect on the firms’ market share. The magnitudes of the effects are similar to those from the foreign industry IV; for example, a one-standard-deviation change in the share of AI workers measured in the resume data translates into a 32% increase in sales and a 26% increase in employment. Table A.8 controls for the 1999–2007 changes in sales and employment at the industry level and firm level, yielding similar estimates of the effects of AI investments on firm growth.

Overall, the results from our second IV are consistent with the OLS results and the first IV specification: all three analyses show that AI investments predict robust and economically meaningful growth at the firm-level. Moreover, we document strikingly similar magnitudes of the effects across the two IV strategies, which are based on different identifying assumptions. In Table A.9, we use both IVs to instrument for AI investments simultaneously, and the over-identification test cannot reject that both IVs are valid in every single specification. This lends further credence to the IV estimation capturing the causal impact of AI investments.

5.1.4 Robustness

We perform several robustness tests around our main findings of increased firm growth stimulated by AI investments.²⁹ We first confirm that the results are not sensitive to different variable construction approaches. In Table A.10 and Table A.11, we show that the results are robust to using firm-level average continuous AI-relatedness measures of job postings, which are defined at the end of Section 3.1. In addition, since the mean of the share of AI workers is relatively low, we confirm that our results are not driven by firms going from zero AI employees in 2010 to a single AI employee in 2018 by excluding such firms in Table A.12. This test addresses a concern that, mechanically, faster expanding firms would employ an AI employee as part of increased employment across all occupations.

Importantly, our results reflect specifically investments in artificial intelligence, rather than

²⁹All robustness tests in this section report the results using the long-differences and the foreign industry IV specifications. In untabulated analyses, we also confirm that none of these robustness tests alter the results of the geographic shift-share IV specification.

other related technologies. In Table A.13, we estimate the relationship between AI investments and firm-level growth (changes in log sales, log employment, and market share), controlling for (i) investments in robots and (ii) investments in other information technologies (IT). Our measures of robot and IT investments are constructed from the job postings data and parallel the measure of AI investments: for each firm, we measure the percentage of job postings in each year requiring robotics- or IT-related skills. The estimated effects of AI investments remain very similar with the addition of these controls.

5.2 Industry Growth

To shed light on the aggregate effects of AI investments, we examine the relationship between industry-level variation in AI investments and industry growth. While AI-investing firms grow faster, the gain in sales and employment may be zero-sum if the rivalrous nature of data and technologies creates a business-stealing effect on the competitors, and signing the effect is an empirical question (Bloom et al., 2013a). 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).

To examine whether AI-fueled growth at the firm level translates into aggregate growth at the industry level, we estimate the following industry-level variant of our firm-level long-differences regression:

$$\Delta \ln y_{j,[2010,2018]} = \gamma \Delta \text{ShareAIWorkers}_{j,[2010,2018]} + \text{IndustryFE} + \epsilon_j \quad (6)$$

where $\Delta \ln y_{j,[2010,2018]}$ is the change in total sales or employment for all firms in our firm-level regression sample in industry j , and $\Delta \text{ShareAIWorkers}_{j,[2010,2018]}$ 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, regressions are weighted by the total number of resumes (or job postings) in each industry in 2010. We also use the first IV strategy described in Section 5.1.3. Specifically, we instrument US industry-level AI investments with European industry-level AI investments calculated from the resume data. All industry-level regressions are at the 5-digit NAICS level, which is the same level as the instrument.³⁰

Table 6 shows that AI investments are associated with a robust increase in employment and

³⁰We only consider the first IV strategy (foreign industry IV) for industry-level regressions. The second instrument (geographic shift-share IV) relies on firm-level variation and does not directly map to an industry-level instrumental variable.

sales at the industry level. We report the coefficients for the resume-based measure of AI investments in Panel 1 and the job-postings-based measure in Panel 2. In both panels, columns 1–3 present the OLS estimates, and columns 4–6 show the second stage IV results. 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. At the industry level, the instrument is generally strong with first-stage F-statistics above 10 for the resume-based measure. For the resume-based measure, a one-standard-deviation increase in the share of AI workers in an industry is associated with a 15% increase in sales and a 17% increase in employment (columns 2 and 4). Similarly to the firm-level IV results, the IV estimates are positive, statistically significant, and larger than the OLS estimates.

To take into account entry and exit of firms, in Table [A.14](#) we include all public firms when calculating total industry sales and employment in 2010 and 2018, including those that entered the sample after 2010 or exited before 2018. This yields similar positive and significant effects of AI investments on total industry sales and employment, suggesting that the negative spillovers on competitors, if any, are small in magnitude, at least within the sample of public firms. To evaluate the potential spillovers on firms outside of the Compustat sample, in Table [A.15](#) 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. The net effect is approximately zero, indicating that although there may be some negative spillovers from public to private firms, these spillovers do not dominate and, contrary to robots, AI does not reduce aggregate employment.

5.3 AI Investments and Industry Concentration

Our results so far show that larger firms invest more in AI, and that AI fuels firm and industry growth. We next examine whether AI investments are related to changes in industry concentration. As we discuss in Section 1, AI can reduce the barriers to growth among small firms or contribute to the growth of larger firms and higher industry concentration documented in [Autor et al. \(2020\)](#).

To explore these hypotheses, we begin by looking at how the effect of AI investments on firm growth varies along the initial firm size distribution. If AI leads to the highest growth among the ex ante smallest firms, then it will reduce overall concentration; on the other hand, if AI leads to greater growth among the ex ante largest firms, then it will increase overall concentration. Table

7 shows the relationship between AI investments and firm growth by firms' initial size, measured as firm employment in 2010. The independent variables are changes in the firm-level share of AI workers (resume-based in Panel 1 and job-postings-based in Panel 2) from 2010 to 2018 interacted with dummy variables indicating which size tercile (within the firm's 2-digit NAICS sector) the firm falls into in 2010. 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.

To examine whether the effect of AI on the growth of the largest firms is substantial enough to translate into 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-all dynamics, we also consider the fraction of sales accruing to the largest firm in each 5-digit NAICS industry among the Compustat firms. Table 8 estimates the relationship between changes in industry concentration and industry-level AI investments in both OLS and IV regressions, using the resume-based measure of AI investments in Panel 1 and the job-postings-based measure in Panel 2. Industry-level growth in AI investments leads to growth in both measures of industry concentration. For example, in OLS regressions with controls, a one-standard-deviation change in the industry share of AI workers measured using resume data corresponds to a 2.2% increase in the HHI and a 1.7% increase in the market share accruing to the top firm. It is worth noting that our concentration results are based on the sample of Compustat firms; to the extent that AI investments stimulate greater growth in this sample of firms than among non-Compustat firms (Table A.15), the overall effect on industry concentration is likely to be even greater. This result is consistent with [Bessen \(2017\)](#), who argues that investments in proprietary technology systems are likely 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.

6 Mechanisms

In this section, we examine the three non-mutually-exclusive mechanisms detailed in Section 1 that can explain both the AI-fueled growth and the increase in industry concentration. We find no evidence of higher productivity or market power. Instead, our results point to more nuanced effects of AI investments, with AI increasing the scale of the most productive firms by facilitating their expansion into new geographic and product markets.

6.1 Productivity

First, we explore whether the increase in firm growth from AI investments is driven by AI technologies making firms more productive (i.e. the productivity channel in Section 1.1). To test Prediction 1.1, we consider two measures of productivity: sales per worker (or labor productivity) and revenue Total Factor Productivity (TFP).³¹ Table 9 shows that investments in AI are associated with slightly lower sales per worker and revenue TFP, although the effects are not statistically significant. The lack of growth in labor productivity is consistent with the results in Section 5 that AI investments lead to similar increases in sales and employment and challenges the view that the primary use of AI is to replace human tasks and cut down labor costs. We also do not observe any evidence of a decrease in costs (Prediction 1.2): the growth in costs of goods sold (COGS) and operating expenses associated with AI investments is similar in magnitude to the growth in firm size. 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 in column 8 of Panel 1, comparable to the effect of AI on sales growth in Section 5.

As a result, we do not find evidence that investments in AI make firms more productive, at least in the short term. Brynjolfsson et al. (n.d.) point out that because complementary investments are necessary to obtain the full benefit of a new technology like AI, productivity growth will follow a J-curve, where firms investing in AI can have low measured productivity growth in the short run and high productivity growth in the future. In Table A.16, we look at the effect of AI investments during the first half of the period (2010–2014) on productivity growth through 2018 and do not find

³¹Revenue TFP is 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. For example, to the extent that AI changes the production function through a shift from variable to fixed costs, the empirical estimates would be biased in favor of finding large increases in sales per worker in response to AI investments, which is contrary to the null results we document.

any significant positive effect. This underscores the extent of the puzzle presented by [Brynjolfsson et al. \(2019\)](#): even with a lag of a few years, AI is not yet associated with productivity gains of AI-investing firms. If AI does bring productivity gains in the future, our results indicate that the time lag between adoption and productivity growth must be longer than a few years.

6.2 Market Power

As discussed in Section 1.2, an alternative explanation for AI-fueled growth is that artificial intelligence grants firms more market power in the product market, enabling firms to expand and increase markups. To test this mechanism, we look at the effect of AI investments on firms' markups (Prediction 2.1). We consider three measures of markups: the first two measures are based on [De Loecker et al. \(2020\)](#) and [Traina \(2018\)](#) and use Cost of Goods Sold (COGS) and Operating Expenses, respectively, as the measure of variable costs. For both of these measures, the markup is the log of revenues divided by the corresponding variable cost measure. For robustness, we also look at the operating profit rate, termed the "Lerner Index" ([Gutiérrez and Philippon, 2017](#)). The Lerner Index is defined as operating income before depreciation and amortization (OIBDA) minus depreciation, scaled by sales. As in previous analyses, we include industry fixed effects in all regressions to absorb sector-specific variation.

Table 10 shows that for all three measures, the effect of AI investments on market power is statistically insignificantly different from zero (except for a small negative effect on the Lerner Index in one of the specifications).³² This is inconsistent with the explanation that AI increases the market power of firms. It is, however, in line with the significant increase in variable costs (COGS or Operating Expenses) reported in Table 9. The increase in costs associated with AI investments has a similar magnitude to the increase in sales (both on the order of 15% per one-standard-deviation change in the share of AI workers in OLS regressions), suggesting that firms grow sales and variable costs at approximately the same rate.

³²Our measures of markups are robust to a wide range of production functions with constant returns to scale. However, the 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.

6.3 Scale Advantages

Lastly, we explore the possibility that investments in artificial intelligence allow already large firms to further scale up their operations, as outlined in Section 1.3. The results presented so far are consistent with this mechanism. First, Table 2 shows that large firms invest more in AI, consistent with Prediction 3.1. Second, the positive effects of AI investments are greatest among the ex ante largest firms (Table 7), consistent with Prediction 3.2.

We perform two additional tests to further investigate the scale advantages channel. First, in addition to the slicing by initial firm size performed in Table 7, we slice firms based on ex ante productivity, as also suggested by Prediction 3.2. Specifically, in Table 11, we group firms in each 2-digit NAICS industry into terciles based on revenue TFP measured as of 2010 and, within each group, examine the relationship between the changes in the share of AI workers and the growth in sales (columns 1 and 2), employment (columns 3 and 4), and market share (columns 5 and 6).³³ The results indicate that even though AI does not appear to improve firms' productivity, the growth fueled by AI is mostly concentrated among the most ex ante productive firms. For example, for the firms in the top tercile of ex ante productivity, a one-standard-deviation increase in the share of AI workers measured using the resume data predicts a 3.2 percentage point increase in market share from 2010 to 2018, compared to a much smaller (0.5 percentage point) and insignificant effect for firms in the bottom tercile of ex ante productivity. This result highlights that the increased industry concentration documented in Table 8 reflects the growth of the most ex ante productive firms.

Second, we look into the drivers of scale (Prediction 3.3) directly by exploring whether AI allows firms to expand their product offerings across markets in our Burning Glass job postings data. Specifically, we estimate the relationship between a firm's investment in AI and changes in: (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; (ii) the number of industries at the most granular (6-digit NAICS) level with at least 1% of the firm's job postings; and (iii) the number of product manager jobs posted by the firm. Since product managers align the firm's product capabilities to specific market segments and are responsible for product innovations, the number of product managers hired by a given firm is a direct measure of the firm's expansion of product offerings.

The results, reported in Table 12, are consistent with AI-investing firms expanding into new geographical and product markets. 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 (column

³³For robustness, Table A.17 shows similar results when using sales per worker to measure initial productivity.

2) and a 14-15% increase in job openings for product managers (column 6), both statistically significant.³⁴ The effect on the number of different industries spanned is milder at only 2% and not statistically significant, suggesting that AI investments are less strongly associated with expansion into new industries. Consistent with AI enabling more product innovations, investments in AI are also associated with increased R&D investments, both in absolute value and as a fraction of sales, as can be seen in columns 7–10 of Table 12. The juxtaposition of positive results on firms’ expansion and null results on productivity and markups points towards firms utilizing AI to enter new markets, without necessarily increasing efficiency in their existing markets or charging higher markups. This is consistent with evidence from prior technological innovations during industrialization, where new technologies have been shown to help firms expand both vertically and horizontally (Braguinsky et al., 2020).

7 Conclusion

We introduce a novel measure of AI investments at the firm level using two detailed datasets of 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 measure of AI investments takes advantage of the co-occurrence of different skills with fundamental concepts such as “machine learning” to empirically determine each skill’s AI-relatedness, which avoids relying on ex ante specified AI-related keywords. Our measure demonstrates the steep growth in AI over the last decade across the full landscape of industries.

We examine both the determinants and consequences of AI investments by firms. We find a positive feedback loop between AI investments and firm size: AI investments concentrate among the largest firms, and as a firm invests in AI, it grows larger, gaining sales, employment, and market share. This firm-level AI-fueled growth is concentrated among larger firms and is associated with increased concentration at the industry level. AI facilitates scalability of the ex ante largest and most productive firms, allowing them to expand product offerings and geographic reach.

Our results highlight that new technologies, such as AI, are an important factor contributing to the increase in industry concentration and the rise of “superstar” firms documented in recent

³⁴In unreported results, we also scale the number of product manager job postings by the number of all job postings at the same firm and find that a one-standard-deviation increase in the share of AI workers increases the share of product manager job postings by about 15% of the mean, although only statistically significant in some specifications.

papers (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 to new markets. Although we do not observe AI-investing firms obtaining measurable productivity gains in the short run, this can be reflective of the productivity J-curve proposed by Brynjolfsson et al. (2019). The positive effect on firm expansion and the null result on firm productivity are consistent with recent evidence in Juhász et al. (2020) and Braguinsky et al. (2020), who show—in the context of mechanized cotton spinning during industrialization—that technology adoption by firms is associated with a high degree of uncertainty in how to apply the new technology effectively. This high uncertainty necessitates experimentation and results in high dispersion with expansion of some firms and potentially slow accrual of productivity benefits for the average firm. Further understanding how AI affects production processes, corporate strategies, and product innovations of firms, and the distribution of gains from investing in AI technologies across firms and workers, are fruitful avenues for future research.

References

- Abis, Simona and Laura Veldkamp**, “The changing economics of knowledge production,” *Working Paper*, 2020.
- Acemoglu, Daron and Pascual Restrepo**, “The race between man and machine: Implications of technology for growth, factor shares and employment,” *American Economic Review*, 2018, 108 (6), 1488–1542.
- **and** – , “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.
- 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.
- Ark, Bart Van**, “The productivity paradox of the new digital economy,” *International Productivity Monitor*, 2016, (31), 3.
- Ater, Itai and Oren Rigby**, “The effects of mandatory disclosure of supermarket prices,” *Working Paper*, 2018.
- Athey, Susan, Kevin Bryan, and Joshua S. Gans**, “The allocation of decision authority to human and artificial intelligence,” *Working paper*, 2020.
- Autor, David and Anna Salomons**, “Is Automation labor-displacing? Productivity growth, employment, and the labor share,” *NBER Working paper w24871*, 2018.
- **and David Dorn**, “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 2013, 103 (5), 1553–97.

- , – , 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.
- Barkai**, **Simcha**, “Declining labor and capital shares,” *Journal of Finance*, 2020, *forthcoming*.
- 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. Upjohn Institute For Employment Research*, 1991.
- Bessen**, **James**, “Information technology and industry concentration,” *Boston University School of Law Law & Economics Paper Series*, 2017, (17-41).
- Bloom**, **Nicholas**, **Mark Schankerman**, and **John Van Reenen**, “Identifying technology spillovers and product market rivalry,” *Econometrica*, 2013, 81 (4), 1347–1393.
- , **Paul Romer**, **Stephen Terry**, and **John Van Reenen**, “A trapped-factors model of innovation,” *AEA Papers and Proceedings*, 2013, 103.
- Braguinsky**, **Serguey**, **Atsushi Ohyama**, **Tetsuji Okazaki**, and **Chad Syverson**, “Product innovation, product diversification, and firm growth: Evidence from Japan’s early industrialization,” *NBER Working Paper w26665*, 2020.
- Brunnermeier**, **Markus**, **Rohit Lamba**, and **Carlos Segura-Rodriguez**, “Inverse selection,” *Working Paper*, 2020.
- Brynjolfsson**, **Erik** and **Tom Mitchell**, “Track how technology is transforming work,” *Nature*, 2017, 544 (7650).
- , **Daniel Rock**, and **Chad Syverson**, “Artificial intelligence and the modern productivity paradox,” *The Economics of Artificial Intelligence: An Agenda*, 2019, p. 23.
- , – , and – , “The productivity J-curve: How intangibles complement general purpose technologies,” *NBER Working Paper w25148*.
- , **Lorin M Hitt**, and **Heekyung Hellen Kim**, “Strength in numbers: How does data-driven decision making affect firm performance?,” *Working Paper*, 2011.
- , **Tom Mitchell**, and **Daniel Rock**, “What can machines learn, and what does it mean for occupations and the economy?,” *AEA Papers and Proceedings*, 2018, 108.
- 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.
- 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.
- Cavallo, Alberto**, “Are online and oine prices similar? Evidence from large multi-channel retailers,” *American Economic Review*, 2017, 107, 283–303.
- 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.
- 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.
- 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.
- Cowgill, Bo**, “Bias and productivity in humans and machines,” *Upjohn Institute Working Paper 19-309*, 2019.
- Crouzet, Nicolas and Janice C Eberly**, “Understanding weak capital investment: The role of market concentration and intangibles,” *NBER Working Paper w25869*, 2019.
- 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.
- 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 H Stern, Chenhao Tan, and Michael S Weisbach**, “Selecting directors using machine learning,” *NBER Working Paper w24435*, 2019.
- 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.
- 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.
- Fujiwara, Ippei and Feng Zhu**, “Robots and labour: Implications for inflation dynamics,” *Working Paper*, 2020.
- Furman, Jason and Peter Orszag**, “A firm-level perspective on the role of rents in the rise in inequality,” *Presentation at “A Just Society” Centennial Event in Honor of Joseph Stiglitz, Columbia University*, 2015, 16.
- **and Robert Seamans**, “AI and the economy,” *Innovation Policy and the Economy*, 2019, 19, 161–191.
- 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,” *NBER Working Paper w24408*, 2019.
- 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.
- Hsieh, Chang-Tai and Esteban Rossi-Hansberg**, “The industrial revolution in services.,” *NBER Working Paper w25968*, 2019.
- Humlum, Anders**, “Robot adoption and labor market dynamics,” *Working Paper*, 2019.
- Juhász, Réka, Mara Squicciarini, and Nico Voigtländer**, “Technology adoption and productivity growth: Evidence from industrialization in France,” *Working Paper*, 2020.
- Karabarbounis, Loukas and Brent Neiman**, “The global decline of the labor share,” *Quarterly journal of economics*, 2014, 129 (1), 61–103.

- Korinek, Anton and Joseph E. Stiglitz**, “Artificial intelligence and its implications for income distribution and unemployment,” *NBER Working paper w24174*, 2017.
- Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard**, “Information technology and returns to scale,” *Working paper*, 2018.
- 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*.
- 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.
- Rock, Daniel**, “Engineering value: The returns to technological talent and investments in artificial intelligence,” *Working paper*, 2019.
- Seamans, Robert and Manav Raj**, “AI, labor, productivity and the need for firm-level data,” *NBER Working paper w24239*, 2018.
- Shiller, Benjamin Reed**, “Personalized price discrimination using big data,” *Working Paper #108, Brandeis University*, 2016.
- Stock, James H and Mark W Watson**, “Introduction to econometrics (3rd updated edition),” *Age (X3)*, 2015, 3 (0.22).
- Syverson, Chad**, “Macroeconomics and market power: Facts, potential explanations and open questions,” *Brookings Institution, Washington DC*, 2019.
- 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.

Varian, Hal, "Artificial intelligence, economics, and industrial organization," *NBER Working Paper w24839*, 2018.

W., Raj Manav Felten Edward and Robert Seamans, "A method to link advances in artificial intelligence to occupational abilities," *AEA Papers and Proceedings*, 2018, 108, 1–4.

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 from Job Posting Data

This figure shows the time series of the job-posting-based measure of AI investments. The figure reports averages for 2007 and 2010-2018, based on the sample of 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 ($w_j^{NarrowAI}$) (left y-axis), and the dashed line tracks the fraction of jobs with (narrow AI) continuous measure above 0.1 (right y-axis).

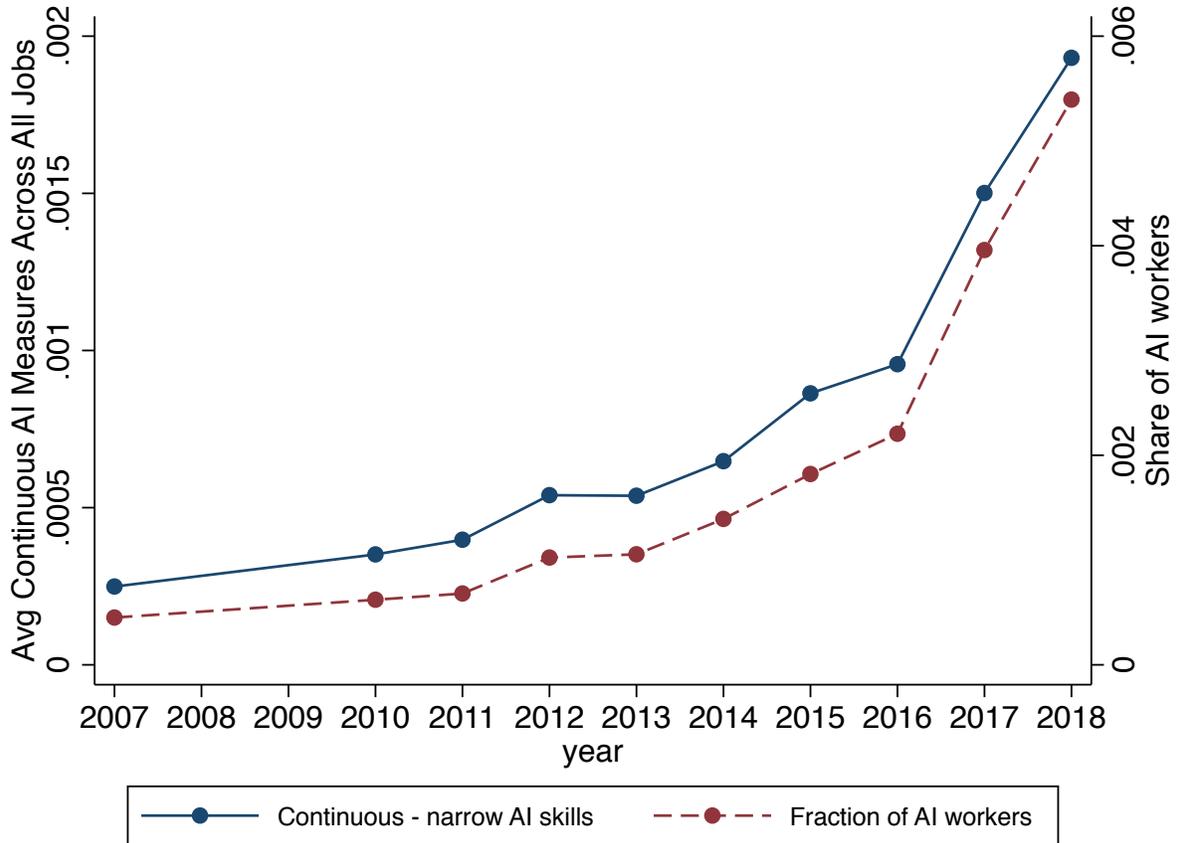


Figure 2. AI Investments by Industry Sector from Job Postings Data

This figure presents the job-posting-based measure of AI investments using the Burning Glass data (based on the sample of public firms) at the industry level. For each sector (based on NAICS-2 digit industry codes), we compute the average job-level continuous measure across all jobs posted by firms in that sector across two sub-periods: 2007–2014 and 2015–2018.

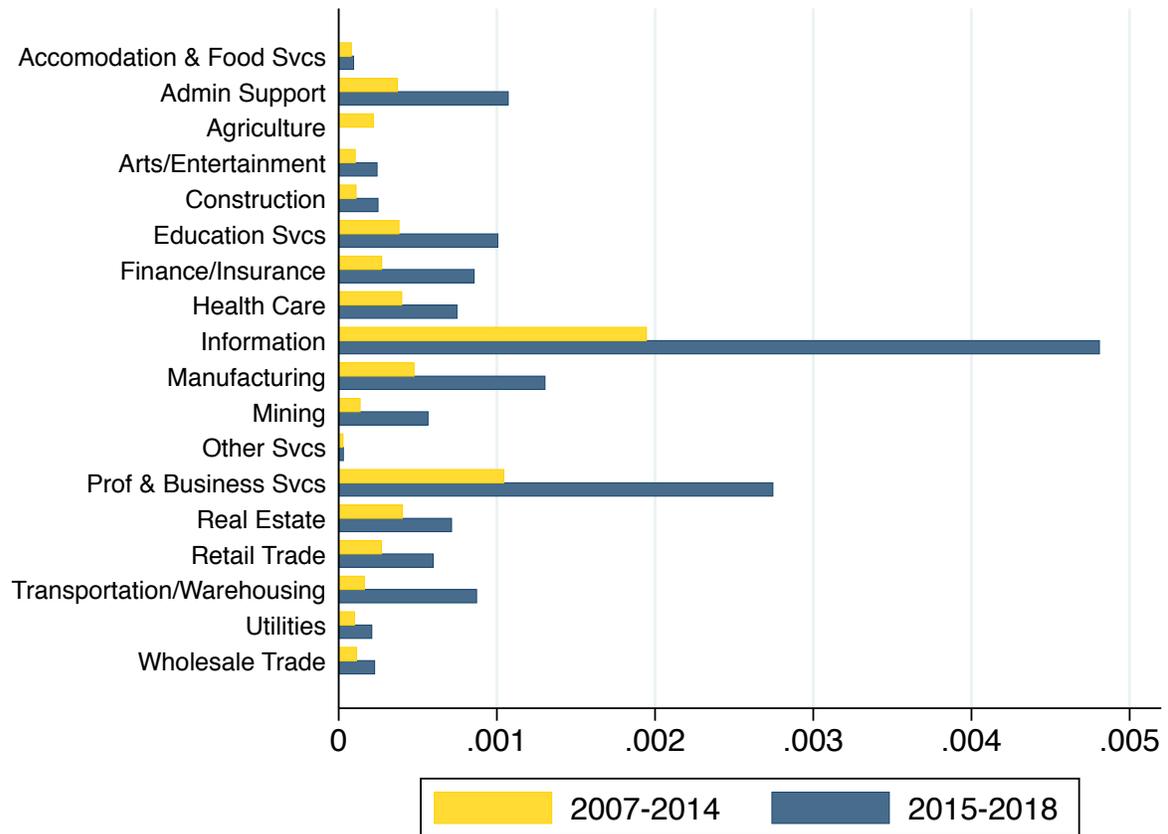


Figure 3. Time Series of the Share of AI Workers from Resume Data

This figure shows the time series of the resume-based measure of AI investments for the sample of public firms in Cognism data, computed as the fraction of all employees (across all firms) in a given year who are classified as holding directly AI-related positions. The figure reports the fraction for each year from 2007 to 2018.

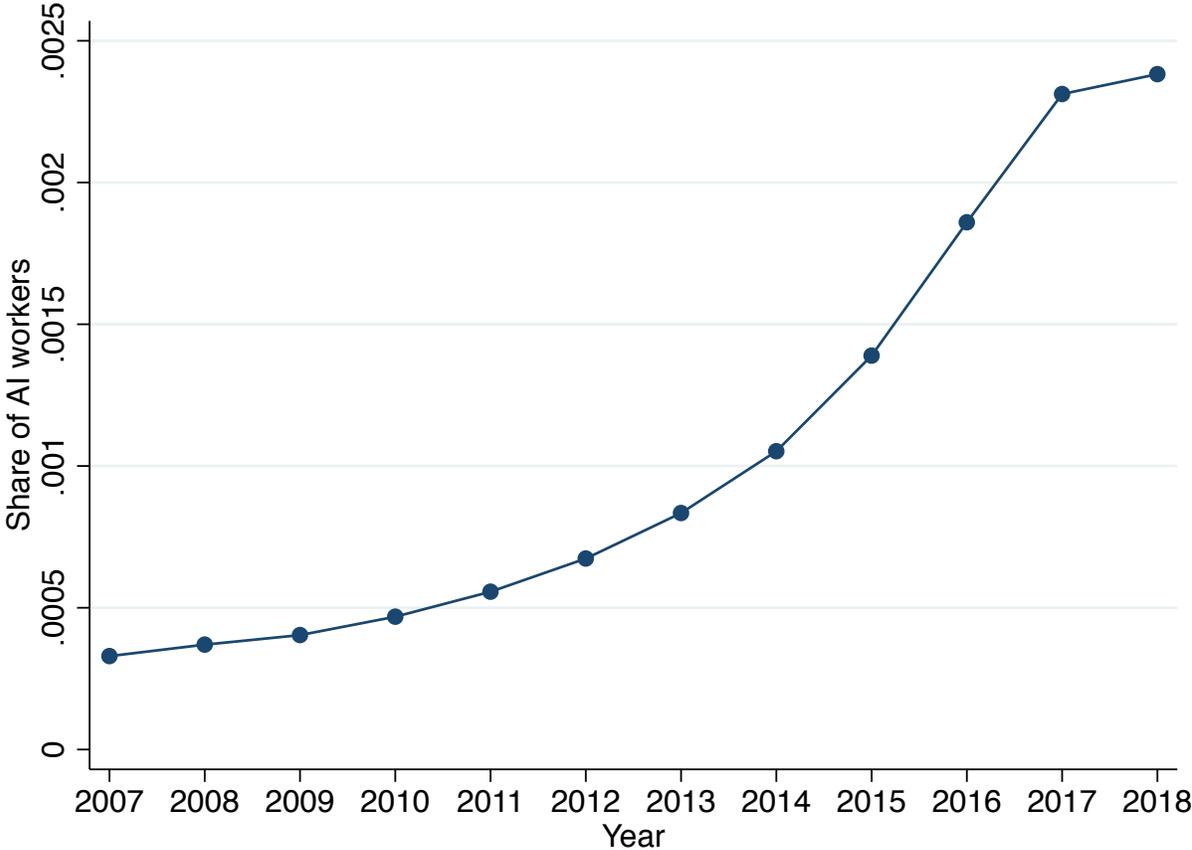


Figure 4. Share of AI Workers by Industry Sector from Resume Data

This figure presents the resume-based measure of AI investments using Cognism data (based on the sample of public firms) at the industry level. For each of the sectors (based on NAICS-2 digit industry codes), we compute the fraction of all individuals employed at the firms within that sector who are classified as AI-related employees. This is done separately for two sub-periods: 2007–2014 and 2015–2018.

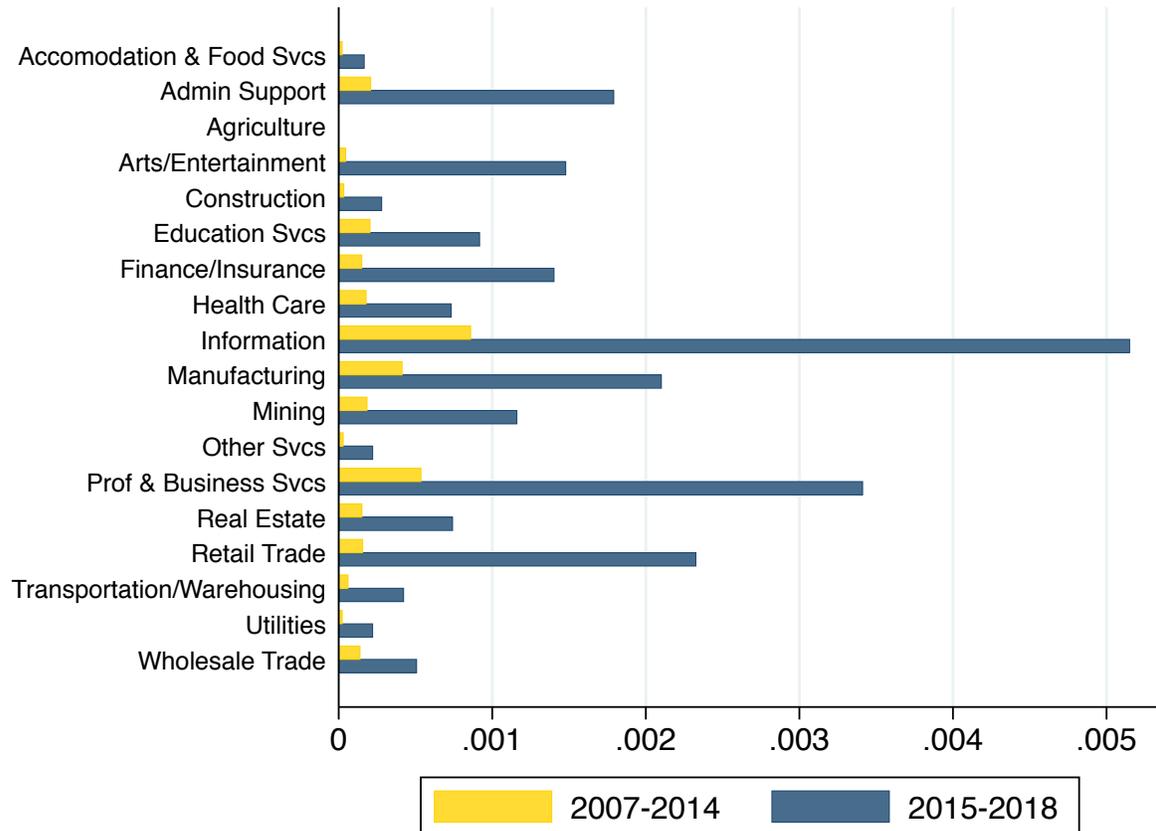
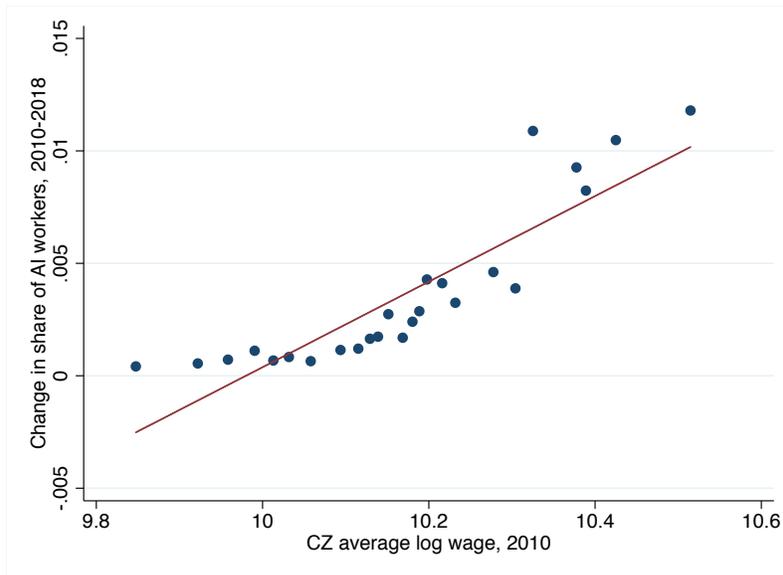
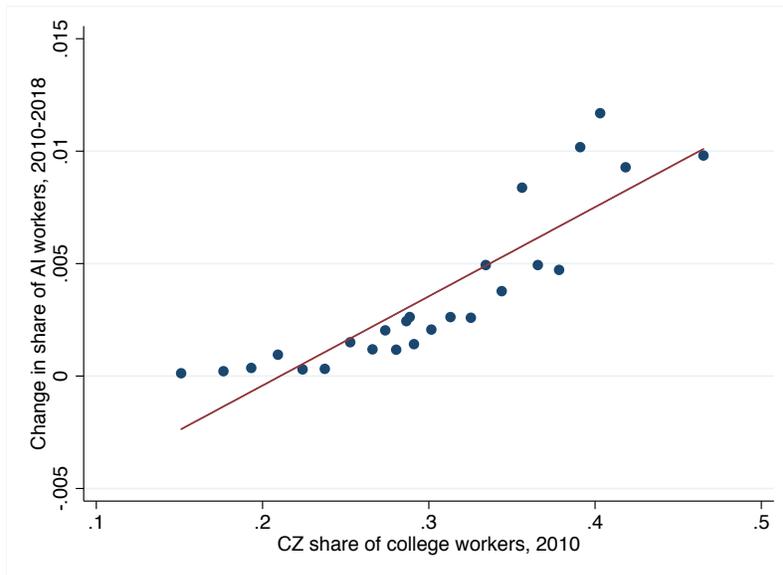


Figure 5. AI Investments and Local Conditions

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 the top figure is the average log wage of a commuting zone in 2010. The x-axis in the bottom figure 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 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.

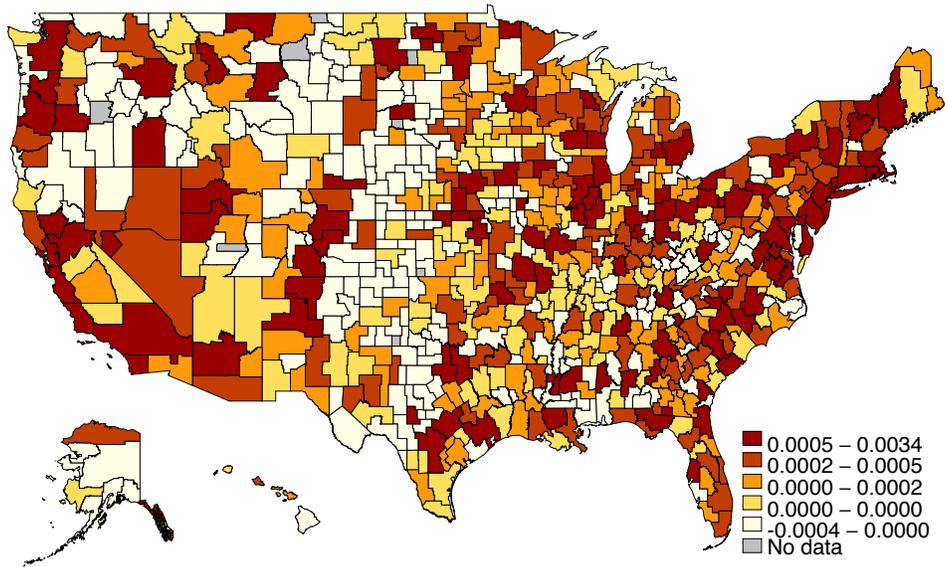
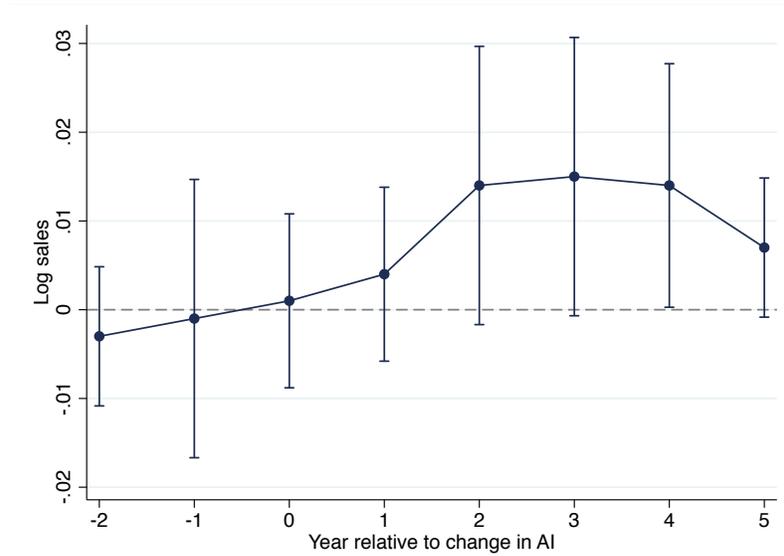
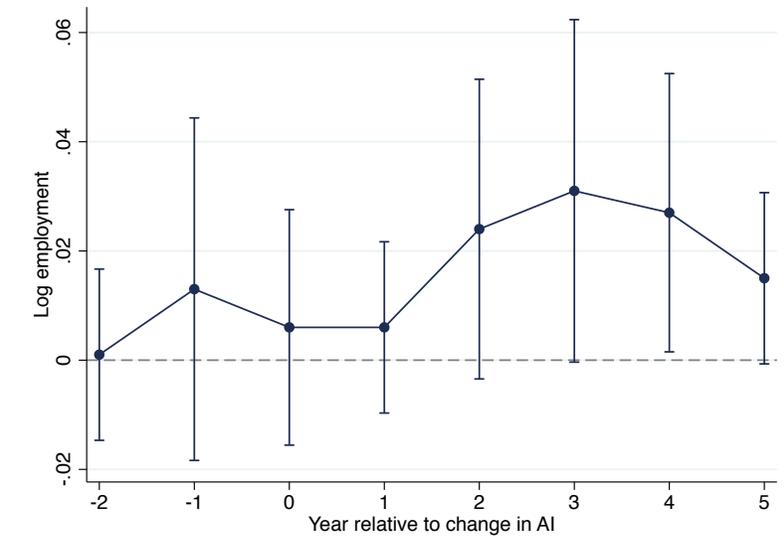


Figure 7. Effect of AI Investments on Firm Growth Over Time

This figure plots the coefficients of the distributed lead-lag model described in Section 5.1.2. The dependent variable is annual log sales in the top figure and log employment in the bottom figure. The independent variable is the annual change in the share of AI workers in Cognism resume data. The independent variables are standardized to have mean zero and standard deviation of one. All regressions include firm observations beginning in 2010 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 6-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 resume-based data 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 continuous $share_{f,t}^{NarrowAI}$ 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

Year	Correlations between:		
	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

Year	Correlations between:		
	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-IT 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.

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				2.933***					3.017***
				(0.729)					(0.675)
R&D/Sales 2010					2.846**				2.061*
					(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								1.389**	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***						0.169***
			(0.042)						(0.041)
Cash/Assets 2010				3.228***					2.975**
				(1.235)					(1.293)
R&D/Sales 2010					0.782*				3.130***
					(0.422)				(1.125)
ROS 2010						0.749*			1.164*
						(0.433)			(0.650)
Log Markup (COGS) 2010							0.341		-0.237
							(0.220)		(0.315)
Log Markup (Total Exp) 2010								1.229***	1.441**
								(0.468)	(0.683)
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	0.319
Observations	1,171	1,202	1,200	1,202	1,200	1,181	1,200	1,200	1,181

Table 3. Effect of AI Investments on Firm Growth: OLS results

This table reports the coefficients from long-differences regressions of changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous firm-level changes in AI investments. We consider three measures of firm size: log sales (columns 1 and 2), log employment (columns 3 and 4), and 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 6-digit NAICS industry level.

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*
	(0.071)	(0.057)	(0.079)	(0.062)	(0.012)	(0.008)
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
	(0.051)	(0.042)	(0.069)	(0.055)	(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. Effect of AI Investments on Firm Growth: Foreign Industry IV

This table estimates the relationship between AI investments and changes in firm size from 2010 to 2018 for U.S. public firms (in non-IT 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 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 6-digit NAICS industry level.

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*** (0.091)	0.281*** (0.050)				
Δ Share AI Workers			0.359*** (0.110)	0.418*** (0.106)	0.313* (0.168)	0.355* (0.190)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	19.1	32.1	19.1	32.1	19.1	32.1
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*** (0.131)	0.495*** (0.079)				
Δ Share AI Workers			0.224*** (0.066)	0.190*** (0.060)	0.253** (0.100)	0.232** (0.096)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	17.6	39.1	17.6	39.1	17.6	39.1
Observations	714	714	714	714	714	714

Table 5. Effect of AI Investments on Firm Growth: Bartik IV

This table estimates the relationship between AI investments and changes in firm size from 2010 to 2018 for U.S. public firms (in non-IT sectors), using a weighted average of national industry-level changes in the share of AI workers 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 headquartered (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 6-digit NAICS industry level.

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*** (0.109)	0.755*** (0.195)						
Δ Share AI Workers			0.343*** (0.087)	0.317*** (0.085)	0.297*** (0.088)	0.260*** (0.100)	0.036 (0.022)	0.024* (0.013)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	34.6	14.9	34.6	14.9	34.6	14.9	34.6	14.9
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*** (0.209)	0.489** (0.195)						
Δ Share AI Workers			0.230*** (0.053)	0.360** (0.154)	0.281** (0.125)	0.237 (0.252)	0.018 (0.017)	0.053 (0.040)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	13.7	6.3	13.7	6.3	13.7	6.3	13.7	6.3
Observations	849	849	849	849	849	849	849	849

Table 6. Effect of AI Investments on Industry-level Employment and Sales

This table reports the coefficients from industry-level long-differences regressions of the changes in sales and employment on contemporaneous changes in AI investments. Each observation is a 5-digit NAICS industry, and we exclude IT 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. Columns 1 to 4 are estimated by OLS, and in columns 5 to 8 the independent variable is instrumented by the contemporaneous growth in the share of AI workers in European firms in each industry. The dependent variables are changes in log total sales (columns 1, 2, 5, and 6) and log total employment (columns 3, 4, 7, and 8) at the industry level from 2010 to 2018. 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. Reported standard errors are robust against heteroskedasticity.

Panel 1: AI measure from resume data

	OLS				IV			
	Δ Log Sales		Δ Log Employment		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.134*** (0.046)	0.150*** (0.035)	0.143** (0.056)	0.167*** (0.047)	0.241*** (0.053)	0.289*** (0.053)	0.206* (0.107)	0.246** (0.095)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					13.6	20.5	13.6	20.5
Observations	233	233	233	233	152	152	152	152

Panel 2: AI measure from job postings data

	OLS				IV			
	Δ Log Sales		Δ Log Employment		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.150*** (0.055)	0.142*** (0.055)	0.154* (0.086)	0.138* (0.082)	0.235*** (0.065)	0.312*** (0.070)	0.252* (0.131)	0.343*** (0.112)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					7.9	12.6	7.9	12.6
Observations	243	243	243	243	158	158	158	158

Table 7. Heterogeneous Effects on Firm Growth by Initial Firm Size

This table reports the coefficients from long-differences regressions of changes in firm size from 2010 to 2018 on contemporaneous changes in AI investments among US public firms (in non-IT 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 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 6-digit NAICS industry level.

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 (0.036)	-0.000 (0.039)	0.037 (0.036)	0.008 (0.046)	-0.007 (0.005)	-0.007 (0.005)
Δ Share AI Workers*Size Tercile 2	0.090 (0.060)	0.043 (0.062)	0.139** (0.057)	0.116* (0.061)	0.009 (0.008)	0.008 (0.009)
Δ Share AI Workers*Size Tercile 3	0.149* (0.080)	0.173*** (0.067)	0.152* (0.088)	0.164** (0.070)	0.017 (0.014)	0.016* (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	2.7	6.5	1.8	4.3	2.8	6.6
T-test p value	0.103	0.011	0.185	0.038	0.093	0.011

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 (0.078)	-0.149*** (0.055)	-0.118 (0.093)	-0.207*** (0.079)	-0.008 (0.007)	-0.012 (0.009)
Δ Share AI Workers*Size Tercile 2	0.092* (0.049)	-0.000 (0.049)	0.064 (0.063)	0.018 (0.054)	0.004 (0.011)	-0.001 (0.012)
Δ Share AI Workers*Size Tercile 3	0.141*** (0.053)	0.126*** (0.043)	0.162** (0.070)	0.114** (0.057)	0.013* (0.007)	0.009 (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	4.4	16.6	5.6	11.3	4.7	6.1
T-test p value	0.037	0.000	0.019	0.001	0.031	0.014

Table 8. AI Investments and Changes in Industry 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 IT 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. Columns 1 to 4 are estimated by OLS, and in columns 5 to 8 the independent variable is instrumented by the contemporaneous growth in the share of AI workers in European firms in each industry. The dependent variables are the changes, from 2010 to 2018, in the Herfindahl-Hirschman Index (HHI) in columns 1, 2, 5, and 6 and in the market share of the top firm in an industry in columns 3, 4, 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.

Panel 1: AI measure from resume data

	OLS				IV			
	HHI		Top Firm Market Share		HHI		Top Firm Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.023*** (0.005)	0.022*** (0.006)	0.016*** (0.005)	0.017*** (0.006)	0.048*** (0.018)	0.053*** (0.018)	0.034** (0.014)	0.039*** (0.014)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					13.6	20.5	13.6	20.5
Observations	233	233	233	233	152	152	152	152

Panel 2: AI measure from job postings data

	OLS				IV			
	HHI		Top Firm Market Share		HHI		Top Firm Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.023*** (0.006)	0.026*** (0.007)	0.013** (0.006)	0.013* (0.008)	0.047*** (0.011)	0.054*** (0.011)	0.027* (0.016)	0.034*** (0.012)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					7.9	12.6	7.9	12.6
Observations	243	243	243	243	158	158	158	158

Table 9. Effect of AI Investments on 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-IT 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 standardized 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 6–digit NAICS industry level.

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 (0.030)	-0.009 (0.027)	-0.006 (0.032)	-0.003 (0.031)	0.142** (0.063)	0.142*** (0.051)	0.133** (0.067)	0.154*** (0.051)
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 (0.043)	-0.020 (0.039)	-0.017 (0.037)	-0.019 (0.035)	0.159*** (0.038)	0.124*** (0.028)	0.145*** (0.045)	0.121*** (0.034)
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 10. Effect of AI Investments on 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-IT 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 6-digit NAICS industry level.

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 (0.016)	0.013 (0.022)	-0.004 (0.006)	0.003 (0.009)	-0.007* (0.004)	-0.002 (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 (0.029)	-0.004 (0.031)	-0.008 (0.009)	0.001 (0.012)	-0.005 (0.006)	-0.001 (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 11. Heterogeneous Effects by Initial Firm Productivity (TFP)

This table reports the coefficients from long-differences regressions of changes in firm size and productivity from 2010 to 2018 on the contemporaneous changes in AI investments among US public firms (in non-IT 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 size: log sales (columns 1 and 2), log employment (columns 3 and 4), and 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 6-digit NAICS industry level.

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 (0.068)	0.069 (0.072)	0.149** (0.068)	0.152** (0.069)	0.010 (0.012)	0.005 (0.010)
Δ Share AI Workers*TFP Tercile 2	0.038 (0.030)	0.095*** (0.031)	0.021 (0.038)	0.076** (0.036)	0.002 (0.004)	0.010** (0.005)
Δ Share AI Workers*TFP Tercile 3	0.350*** (0.070)	0.346*** (0.075)	0.314*** (0.081)	0.294*** (0.095)	0.043** (0.020)	0.032** (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	8.7	8.0	2.6	1.8	1.8	2.0
T-test p value	0.003	0.005	0.106	0.175	0.179	0.160

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 (0.047)	0.023 (0.052)	0.169** (0.073)	0.140*** (0.051)	-0.004 (0.011)	-0.007 (0.011)
Δ Share AI Workers*TFP Tercile 2	0.051** (0.022)	0.045** (0.023)	0.045 (0.028)	0.024 (0.029)	-0.002 (0.005)	0.001 (0.005)
Δ Share AI Workers*TFP Tercile 3	0.252*** (0.045)	0.236*** (0.034)	0.234*** (0.080)	0.181*** (0.064)	0.032*** (0.010)	0.025** (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	12.7	14.9	0.3	0.2	4.2	4.3
T-test p value	0.000	0.000	0.565	0.642	0.042	0.040

Table 12. 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 and innovation efforts by U.S. public firms (in non-IT 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; in columns 5 and 6, the dependent variable is the change in the log number of job postings for product managers; in columns 7 and 8, the dependent variable is the change in log R&D investment; in columns 9 and 10, the dependent variable is the change in R&D expenditure as a fraction of sales. 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, 8, and 10 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 headquartered (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 6–digit NAICS industry level.

Panel 1: AI measure from resume data

	Δ Log Number of Counties		Δ Log Number of Industries		Δ Log Number of Product Managers		Δ Log R&D		Δ R&D/Sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Share AI Workers	0.081*** (0.019)	0.085*** (0.027)	-0.001 (0.016)	0.023 (0.022)	0.209*** (0.060)	0.137* (0.083)	0.227** (0.109)	0.227** (0.093)	0.008*** (0.002)	0.008*** (0.002)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.254	0.328	0.130	0.167	0.063	0.108	0.207	0.308	0.048	0.444
Observations	757	757	759	759	756	756	766	766	766	766

Panel 2: AI measure from job postings data

	Δ Log Number of Counties		Δ Log Number of Industries		Δ Log Number of Product Managers		Δ Log R&D		Δ R&D/Sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Share AI Workers	0.038* (0.023)	0.071*** (0.023)	0.003 (0.009)	0.017 (0.016)	0.193*** (0.041)	0.149** (0.074)	0.193** (0.079)	0.180** (0.084)	0.005*** (0.002)	0.006*** (0.002)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.441	0.610	0.076	0.106	0.119	0.190	0.258	0.360	0.009	0.738
Observations	847	847	849	849	849	849	849	849	849	849

A Online 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 forecast-

ing (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 Oberhelmer 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, Illinois, 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

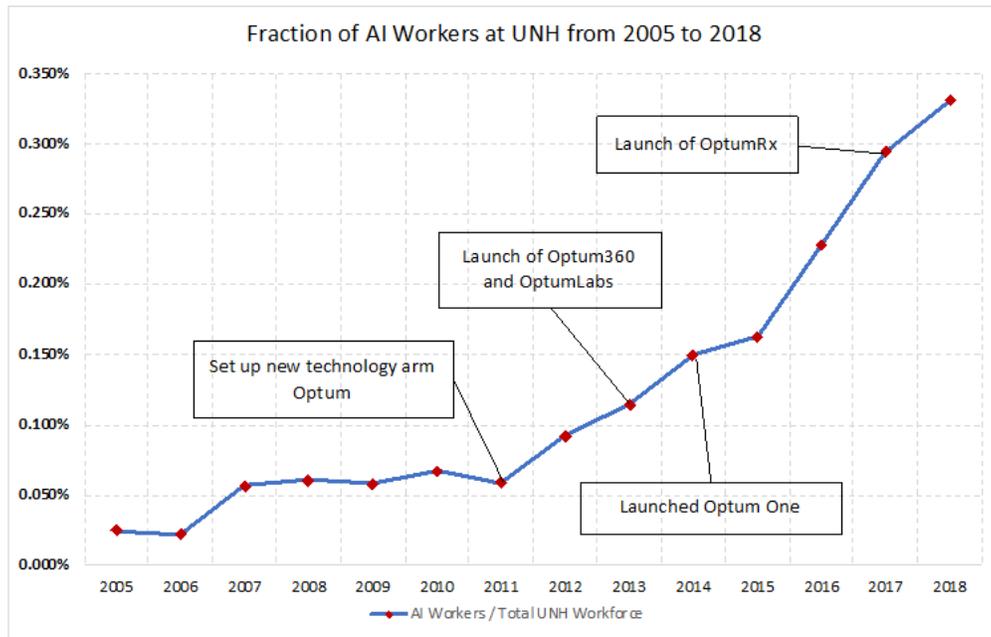


Figure A.2. Timeline of AI investments by JPMorgan Chase & Co

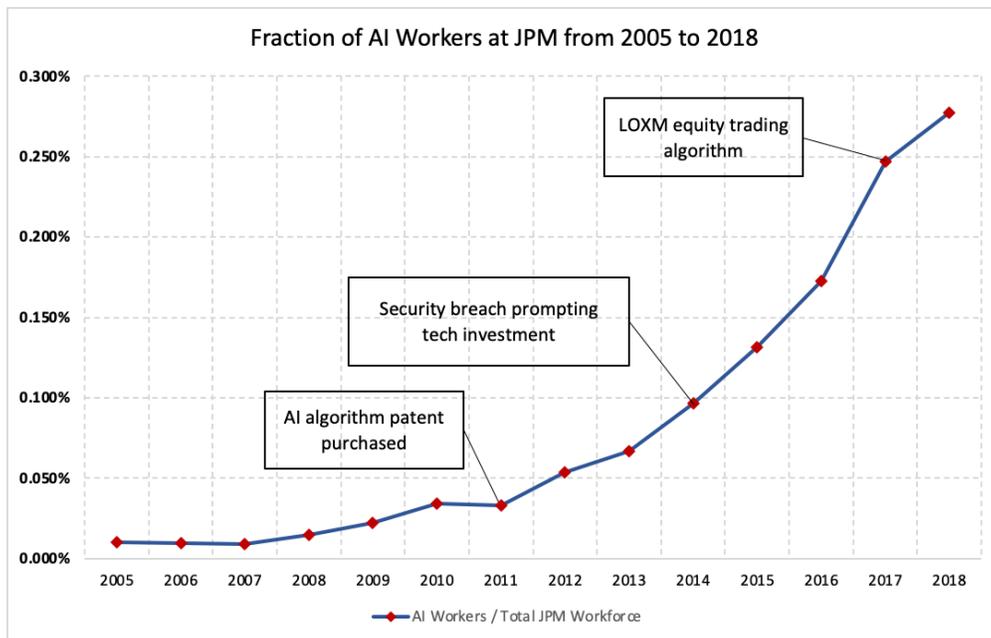


Figure A.3. Timeline of AI investments by Caterpillar Inc

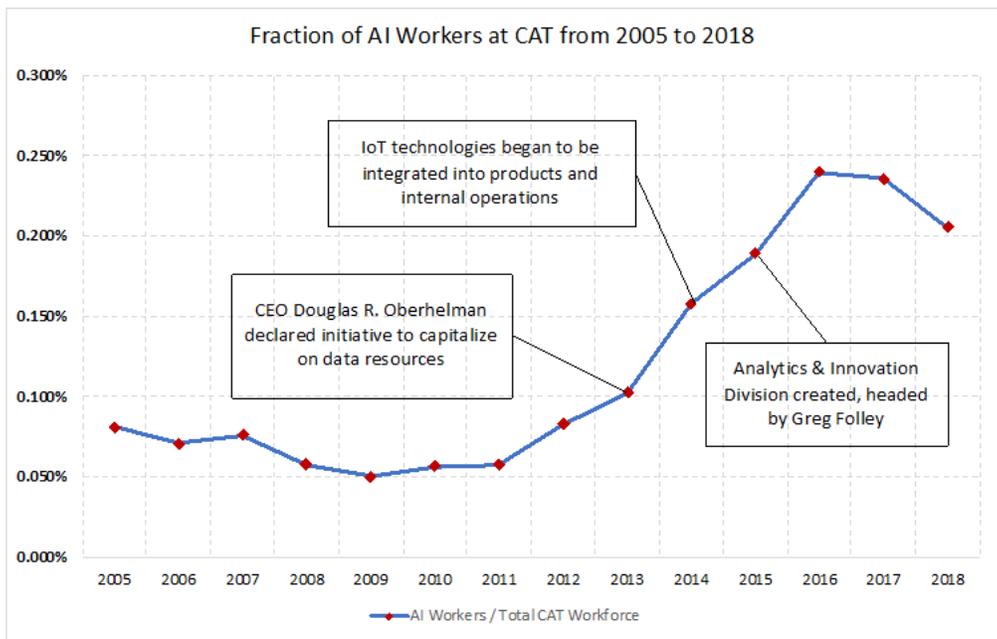
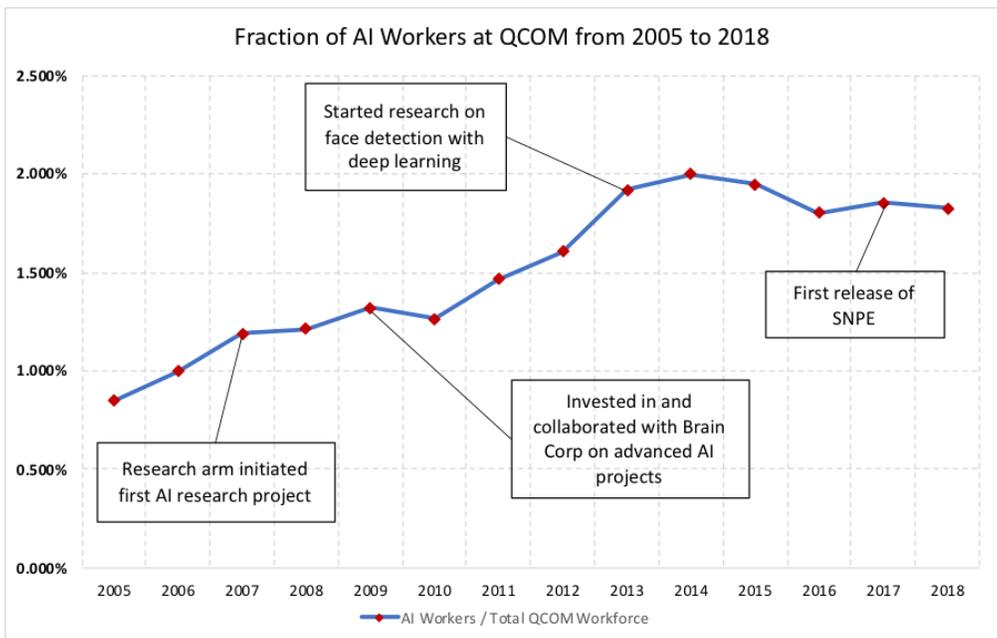


Figure A.4. Timeline of AI investments by Qualcomm Inc



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.

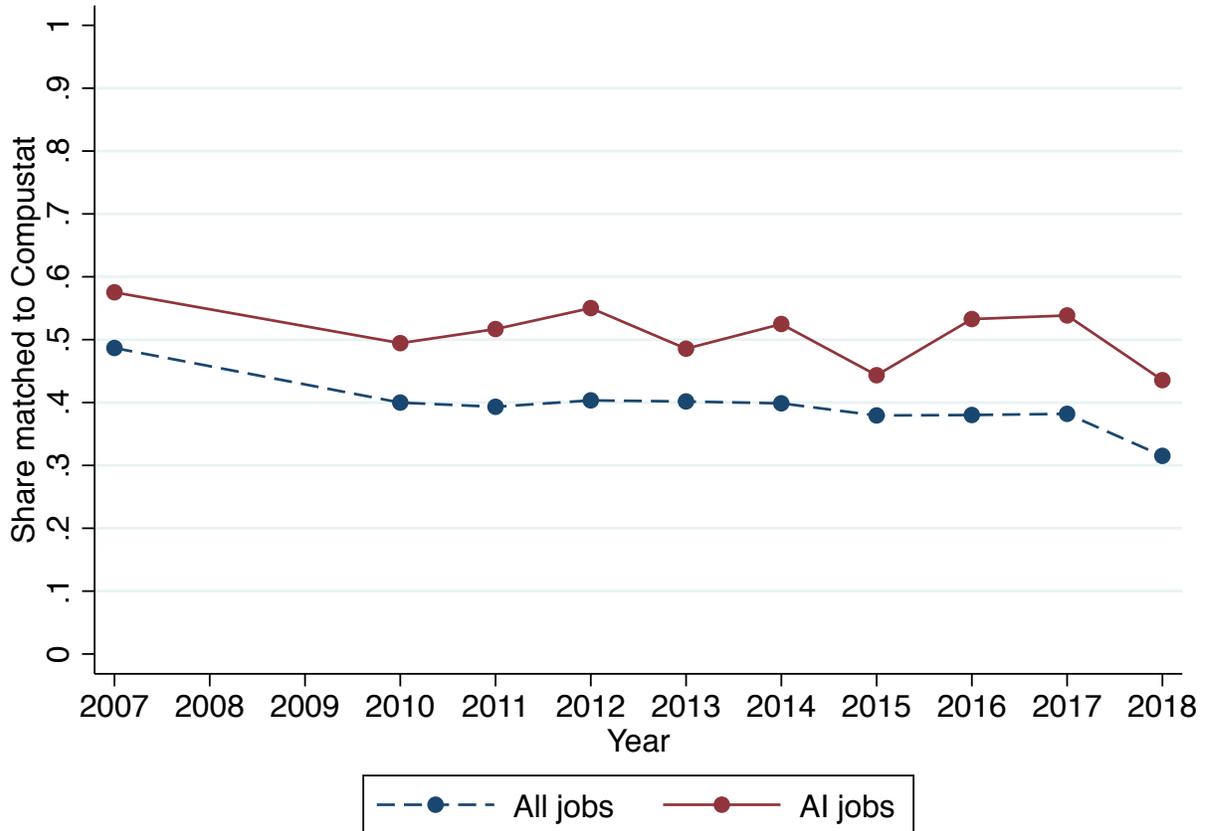


Figure A.6. Time Series of the Share of AI Workers in Foreign Countries from Resume Data

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 London Stock Exchange. The group “India” includes all firms listed on 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”.

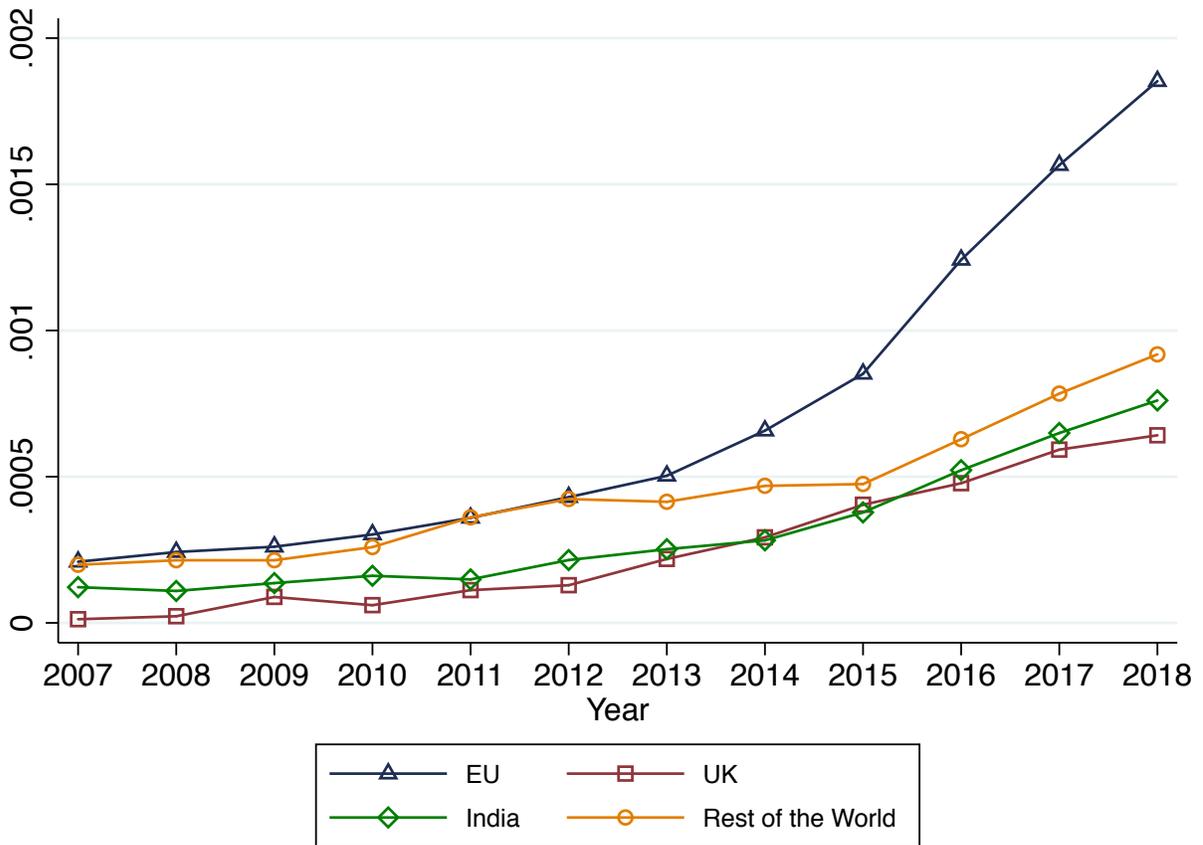


Figure A.7. Correlation Between Changes in the Producer Price Index (PPI) and Changes in the Share of AI Workers in Europe at the Industry Level

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 2018s, 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.

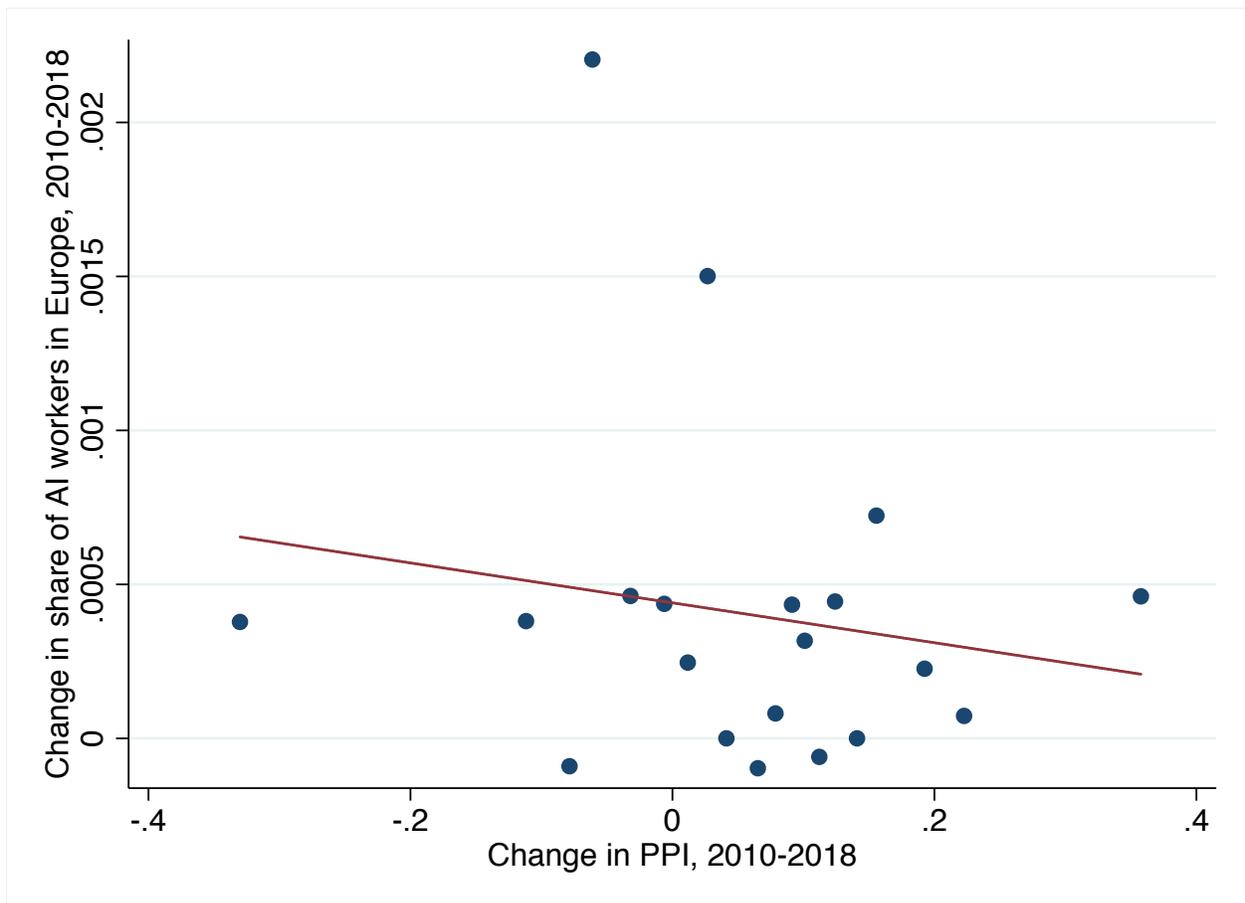


Table A.1. Job Titles with the Highest Average AI-relatedness Measure

This table reports the job titles in Burning Glass with the highest average AI measure $\omega^{NarrowAI}$. 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.2. Job Titles with the Highest Number of AI Jobs

This table reports the job titles in Burning Glass with the highest number of AI jobs. AI jobs are defined as jobs with continuous measure $\omega^{NarrowAI}$ above 0.1. We only include job postings that are matched to Compustat firms.

	cleantitle	number of AI jobs
1	Data Scientist	3,529
2	Senior Data Scientist	1,547
3	Software Engineer	665
4	Principal Data Scientist	434
5	Data Engineer	409
6	Senior Software Engineer	399
7	Research Scientist	398
8	Lead Data Scientist	358
9	Machine Learning Engineer	305
10	Big Data Engineer	239
11	Senior Data Engineer	230
12	Big Data Architect	197
13	Big Data Consultant	191
14	Data Analyst	176
15	Data Scientist, Senior	168
16	Data Scientist II	153
17	Hadoop Developer	153
18	Software Development Engineer	151
19	Data Science Engineer	147
20	Machine Learning Scientist	144
21	Big Data Developer	144
22	Software Engineer - Data Mining/Data Analysis/Machine Learning	140
23	Data Scientist, Mid	132
24	Senior Research Scientist	128
25	Research Engineer	128
26	Artificial Intelligence Consultant	126
27	Machine Learning Researcher	125
28	Research And Development Engineer - Data Mining/Data Analysis/Machine Learning	116
29	Applied Scientist	113
30	Lead Machine Learning Scientist - Enterprise Products	113
31	Software Engineer Ads & Data Mining	112
32	Software Engineer - Entry Level	111
33	Business Process Analyst	110
34	Artificial Intelligence Manager	109
35	Big Data Scientist	109
36	Big Data Engineer Consultant	108
37	Principal Software Engineer	106
38	Big Data Manager	105
39	Senior Applied Scientist	103
40	Software Developer	103
41	Principal Digital Product Manager	103
42	Senior Engineer II - Data Scientist	102
43	Artificial Intelligence Analyst	102
44	Senior Engineer I	102
45	F - Program Intel Threat Analyst	101
46	Staff Data Scientist	97
47	Engineering Manager - Feed Personalization Platform	96
48	Architect - Relevance Infrastructure	93
49	Software Engineer, Machine Learning	93
50	Computer Vision Engineer	91

Table A.3. Occupations with the Highest Number of AI Jobs

This table reports the names of the BLS occupations with the highest number of AI jobs in Burning Glass. AI jobs are defined as jobs with continuous measure $\omega^{NarrowAI}$ above 0.1. We only include job postings that are matched to Compustat firms.

	BLS Occupation Name	number of AI jobs
1	Computer and Information Research Scientists	21,273
2	Software Developers, Applications	20,977
3	Computer Occupations, All Other	12,692
4	Operations Research Analysts	6,980
5	Database Administrators	3,451
6	Marketing Managers	2,760
7	Managers, All Other	2,592
8	Architectural and Engineering Managers	1,559
9	Engineers, All Other	1,391
10	Computer Systems Analysts	1,212
11	General and Operations Managers	1,070
12	Management Analysts	1,047
13	Information Security Analysts	811
14	Mechanical Engineers	761
15	Detectives and Criminal Investigators	741
16	Statisticians	736
17	Web Developers	658
18	Computer Hardware Engineers	648
19	Electrical Engineers	628
20	Computer Network Architects	625
21	Financial Specialists, All Other	599
22	Sales Managers	594
23	Medical and Health Services Managers	539
24	Engineering Technicians, Except Drafters, All Other	524
25	Natural Sciences Managers	460
26	Market Research Analysts and Marketing Specialists	454
27	Computer Programmers	425
28	Medical Scientists, Except Epidemiologists	408
29	Sales Representatives, Wholesale and Manufacturing, Except Technical and Scientific Products	398
30	Computer and Information Systems Managers	357

Table A.4. Effect of AI on Firm Growth in Non-IT Sectors

This table reports the coefficients from long-differences regressions of the changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous changes in AI investments, separately by broad industry sectors. 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-IT 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 size: log sales in odd columns and log employment in even columns. The main dependent variables are measured as growth in the size measures from 2010 to 2018. 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 headquartered (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 6-digit NAICS industry level.

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** (0.062)	0.137** (0.063)	0.222** (0.094)	0.253* (0.128)	0.212* (0.103)	0.307*** (0.106)	0.125 (0.088)	0.096 (0.103)
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*** (0.049)	0.086** (0.038)	0.276*** (0.032)	0.325*** (0.032)	0.055** (0.020)	0.069** (0.031)	0.163* (0.092)	0.018 (0.114)
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.5. Robustness: Pre-trend of Foreign Industry IV

This table investigates pre-trends in the firm size variables of U.S. public firms (in non-IT sectors), measured over 1999-2007, prior to the 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 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 headquartered (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 6-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 (0.163)	-0.047 (0.174)
NAICS2 FE	Y	Y
Controls	Y	Y
F Statistic	28.4	26.7
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 (0.066)	0.040 (0.078)
NAICS2 FE	Y	Y
Controls	Y	Y
F Statistic	36.7	32.3
Observations	577	541

Table A.6. Robustness of Foreign Industry IV: Industry- and Firm-level Pre-trend

This table estimates the relationship between AI investments and changes in firm size from 2010 to 2018 for U.S. public firms (in non-IT 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 AI measures. In this table, we control for firm- and industry-level growth in the years before we measure changes in AI investments. The independent variable is 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 independent variable and the IV are standardized to mean zero and standard deviation of one. We consider changes in log sales in columns 1 and 2, and log employment in columns 3 and 4. All specifications 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 headquartered (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. Additionally, columns 1 and 3 control for the growth in industry-level total sales and employment of Compustat firms (at 5-digit NAICS level) from 1999 to 2007, and columns 2 and 4 also control for firm-level growth in sales and employment from 1999 to 2007. Standard errors are clustered at the 6-digit NAICS industry level.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.418*** (0.098)	0.472*** (0.109)	0.355** (0.177)	0.398** (0.182)
NAICS2 FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Industry pre-trend	Y	N	Y	N
Firm pre-trend	N	Y	N	Y
F Statistic	29.3	25.5	29.3	25.5
Observations	637	498	637	498

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.182*** (0.055)	0.197*** (0.053)	0.234** (0.098)	0.227*** (0.086)
NAICS2 FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Industry pre-trend	Y	N	Y	N
Firm pre-trend	N	Y	N	Y
F Statistic	41.1	31.9	41.1	31.9
Observations	709	541	709	541

Table A.7. Robustness: Pre-trend of Bartik IV

This table investigates pre-trends in the firm size variables of U.S. public firms (in non-IT sectors), measured over 1999-2007, prior to the 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 national industry-level changes in the share of AI workers. 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 headquartered (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 6-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 (0.184)	0.204 (0.165)	0.008 (0.023)
NAICS2 FE	Y	Y	Y
Controls	Y	Y	Y
F Statistic	13.1	12.4	13.1
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 (0.273)	-0.131 (0.301)	-0.019 (0.043)
NAICS2 FE	Y	Y	Y
Controls	Y	Y	Y
F Statistic	5.7	5.0	5.6
Observations	687	647	690

Table A.8. Robustness of Bartik IV: Industry- and Firm-level Pre-trend

This table estimates the relationship between AI investments and changes in firm size from 2010 to 2018 for U.S. public firms (in non-IT sectors), using a weighted average of national industry-level changes in the share of AI workers as an instrument for the two firm-level AI measures. In this table, we control for firm- and industry-level growth in the years before we measure changes in AI investments. The independent variable is 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 independent variable and the IV are standardized to mean zero and standard deviation of one. We consider changes in log sales in columns 1 and 2, log employment in columns 3 and 4, and market share within the 5-digit NAICS industry in columns 5 and 6. All specifications 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 headquartered (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. Additionally, columns 1, 3, and 5 control for the growth in industry-level total sales and employment of Compustat firms (at 5-digit NAICS level) from 1999 to 2007, and columns 2, 4, and 6 also control for firm-level growth in sales and employment from 1999 to 2007. Standard errors are clustered at the 6-digit NAICS industry level.

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	0.281*** (0.075)	0.252*** (0.075)	0.223** (0.094)	0.150 (0.097)	0.021 (0.013)	0.016 (0.013)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Industry pre-trend	Y	N	Y	N	Y	N
Firm pre-trend	N	Y	N	Y	N	Y
F Statistic	14.6	10.8	14.6	10.8	14.6	10.8
Observations	755	591	755	591	755	591

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	0.369** (0.156)	0.372** (0.185)	0.231 (0.250)	0.125 (0.141)	0.053 (0.041)	0.026 (0.038)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Industry pre-trend	Y	N	Y	N	Y	N
Firm pre-trend	N	Y	N	Y	N	Y
F Statistic	6.5	5.4	6.5	5.4	6.5	5.4
Observations	839	647	839	647	839	647

Table A.9. Robustness: Over-identification Test of the Two IVs

This table estimates the relationship between AI investments and changes in firm size from 2010 to 2018 for U.S. public firms (in non-IT sectors), using both instruments for firm-level growth in the share of AI workers: 1) the change in the share of AI workers in European public firms in the same NAICS 5-digit industry, and 2) a weighted average of national industry-level changes in the share of AI workers. 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 independent variable and the IV are standardized to mean zero and standard deviation of one. We consider changes in log sales in columns 1 and 2, and log employment in columns 3 and 4. 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 the firms are headquartered (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. We report the Hansen J Statistics for the over-identification test of the two IVs and its corresponding p values at the bottom of each panel, which all show that the null hypothesis that the overidentifying restrictions are valid cannot be rejected. Standard errors are clustered at the 6-digit NAICS industry level.

Panel 1: AI measure from resume data

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.338*** (0.079)	0.297*** (0.073)	0.304*** (0.091)	0.294*** (0.098)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
F Statistic	26.9	25.4	26.9	25.4
Hansen J Statistic	0.0	1.9	0.0	0.2
Over-id p value	0.82	0.17	0.94	0.67
Observations	643	643	643	643

Panel 2: AI measure from job postings data

	Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.192*** (0.051)	0.188*** (0.049)	0.271*** (0.098)	0.239** (0.101)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
F Statistic	14.5	27.1	14.5	27.1
Hansen J Statistic	0.5	0.0	0.0	0.0
Over-id p value	0.47	0.94	0.84	0.92
Observations	714	714	714	714

Table A.10. Robustness: Continuous Narrow AI Measure

This table reports the coefficients from long-differences regressions of the changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous changes in the average job-level continuous AI measure based on narrow AI skills ($\omega_j^{NarrowAI}$) across all job postings in Burning Glass . Panel 1 presents OLS results, Panel 2 shows IV results using the change in the share of AI workers in European public firms in the same NAICS 5-digit industry as instrument, and Panel 3 shows IV results using a weighted average of national industry-level changes in the share of AI workers as instrument. We consider changes in three measures of firm size: log sales (Panel 1: columns 1 and 2; Panel 2 and 3: columns 3 and 4), log employment (Panel 1: columns 3 and 4; Panel 2 and 3: columns 5 and 6), and market share within the 5-digit NAICS industry (Panel 1: columns 5 and 6; Panel 3: columns 7 and 8). The independent variable and the IV are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Burning Glass job postings in 2010. 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 headquartered (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 6-digit NAICS industry level.

Panel 1: OLS						
	Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.140*** (0.053)	0.128*** (0.043)	0.157** (0.069)	0.112* (0.058)	0.012* (0.007)	0.010 (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: Foreign IV						
	First Stage		Second Stage			
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Instrument	0.511*** (0.129)	0.443*** (0.068)				
Δ Share AI Workers			0.241*** (0.068)	0.208*** (0.064)	0.271*** (0.104)	0.257** (0.105)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	15.7	43.1	15.7	43.1	15.7	43.1
Observations	714	714	714	714	714	714

Panel 3: Bartik IV								
	First Stage		Second Stage					
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Instrument	0.793*** (0.190)	0.524*** (0.197)						
Δ Share AI Workers			0.232*** (0.055)	0.385*** (0.143)	0.282** (0.121)	0.256 (0.229)	0.019 (0.016)	0.048 (0.036)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	17.4	7.1	17.4	7.1	17.4	7.1	17.4	7.1
Observations	849	849	849	849	849	849	849	849

Table A.11. Robustness: Continuous All-skill AI measure

This table reports the coefficients from long-differences regressions of the changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous changes in the average job-level continuous AI measure based on all skills (ω_j^{AllAI}) across all job postings in Burning Glass. Panel 1 presents OLS results, Panel 2 shows IV results using the change in the share of AI workers in European public firms in the same NAICS 5-digit industry as instrument, and Panel 3 shows IV results using a weighted average of national industry-level changes in the share of AI workers as instrument. We consider changes in three measures of firm size: log sales (Panel 1: columns 1 and 2; Panel 2 and 3: columns 3 and 4), log employment (Panel 1: columns 3 and 4; Panel 2 and 3: columns 5 and 6), and market share within the 5-digit NAICS industry (Panel 1: columns 5 and 6; Panel 3: columns 7 and 8). The independent variable and the IV are standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Burning Glass job postings in 2010. 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 headquartered (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 6-digit NAICS industry level.

Panel 1: OLS						
	Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.136** (0.062)	0.123** (0.049)	0.147* (0.081)	0.092 (0.066)	0.014* (0.008)	0.014* (0.007)
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

Panel 2: Foreign IV						
	First Stage		Second Stage			
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Instrument	0.453*** (0.093)	0.373*** (0.060)				
Δ Share AI Workers			0.272*** (0.093)	0.247*** (0.084)	0.306** (0.134)	0.305** (0.136)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
F Statistic	23.6	38.3	23.6	38.3	23.6	38.3
Observations	714	714	714	714	714	714

Panel 3: Bartik IV								
	First Stage		Second Stage					
	Δ Share of AI Workers		Δ Log Sales		Δ Log Employment		Δ Log Market Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Instrument	0.672*** (0.150)	0.453*** (0.157)						
Δ Share AI Workers			0.274*** (0.073)	0.446*** (0.162)	0.333** (0.148)	0.296 (0.266)	0.023 (0.019)	0.055 (0.041)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic	20.1	8.3	20.1	8.3	20.1	8.3	20.1	8.3
Observations	849	849	849	849	849	849	849	849

Table A.12. Robustness: Excluding Firms With Only One AI Job

This table reports the coefficients from long-differences regressions of the changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous changes in the share of AI workers, excluding firms that had zero AI workers in 2010 and one AI worker in 2018s. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. The independent variable and the IV are standardized to mean zero and standard deviation of one. In each panel, columns 1-3 present OLS results, columns 4-5 show IV results using the change in the share of AI workers in European public firms in the same NAICS 5-digit industry as instrument, and columns 6-8 show IV results using a weighted average of national industry-level changes in the share of AI workers as instrument. We consider changes in three measures of firm size: log sales (columns 1, 4 and 6), log employment (columns 2, 5 and 7), and market share within the 5-digit NAICS industry (columns 3 and 8). 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 headquartered (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 6-digit NAICS industry level.

Panel 1: AI measure from resume data

	OLS			Foreign IV		Bartik IV		
	Δ Log Sales	Δ Log Employment	Δ Market Share	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Market Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.170*** (0.062)	0.158** (0.067)	0.015* (0.008)	0.470*** (0.119)	0.429** (0.210)	0.339*** (0.093)	0.272** (0.111)	0.028* (0.014)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
F Statistic				34.9	34.9	13.8	13.8	13.8
Observations	684	684	684	570	570	684	684	684

Panel 2: AI measure from job postings data

	OLS			Foreign IV		Bartik IV		
	Δ Log Sales	Δ Log Employment	Δ Market Share	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Market Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.122*** (0.040)	0.090* (0.052)	0.008 (0.006)	0.197*** (0.057)	0.256*** (0.090)	0.405*** (0.146)	0.419* (0.233)	0.056 (0.034)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
F Statistic				42.2	42.2	9.1	9.1	9.1
Observations	802	802	802	674	674	802	802	802

Table A.13. Robustness: Controlling for General IT and Robot Technologies

This table reports the coefficients from the long-differences regressions of the changes in firm size of U.S. public firms (in non-IT sectors) from 2010 to 2018 on the contemporaneous changes in each firm’s share of AI workers. Panel 1 considers the resume-based measure of the share of AI workers, while Panel 2 looks at the job-posting-based measure. All regressions also control for the 2010-2018 changes in the firm’s share of other IT jobs and the share of robot-related jobs measured using the Burning Glass data. An IT job is defined as a job for which at least 10% of the required skills are in the “Information Technology” skill cluster, and 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. The growth in the AI, IT, and robot measures and the IV are standardized to mean zero and standard deviation of one. In each panel, columns 1-3 present OLS results, columns 4-5 show IV results using the change in the share of AI workers in European public firms in the same NAICS 5-digit industry as instrument, and columns 6-8 show IV results using a weighted average of national industry-level changes in the share of AI workers as instrument. We consider changes in three measures of firm size: log sales (columns 1, 4 and 7), log employment (columns 2, 5 and 7), and market share within the 5-digit NAICS industry (columns 3 and 8). 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 6–digit NAICS industry level.

Panel 1: AI measure from resume data

	OLS			Foreign IV		Bartik IV		
	Δ Log Sales	Δ Log Employment	Δ Market Share	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Market Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.159*** (0.056)	0.152** (0.064)	0.015* (0.008)	0.419*** (0.117)	0.372* (0.202)	0.325*** (0.089)	0.271** (0.105)	0.027* (0.014)
Δ Share Other IT Workers	0.019 (0.039)	-0.001 (0.044)	0.005 (0.010)	0.080 (0.065)	0.053 (0.064)	0.033 (0.049)	0.008 (0.051)	0.005 (0.011)
Δ Share Robot Workers	-0.009 (0.050)	-0.023 (0.052)	-0.010 (0.010)	-0.060 (0.070)	-0.056 (0.067)	-0.035 (0.061)	-0.041 (0.060)	-0.012 (0.011)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
F Statistic				23.9	23.9	13.5	13.5	13.5
Observations	756	756	756	635	635	756	756	756

Panel 2: AI measure from job postings data

	OLS			Foreign IV		Bartik IV		
	Δ Log Sales	Δ Log Employment	Δ Market Share	Δ Log Sales	Δ Log Employment	Δ Log Sales	Δ Log Employment	Δ Market Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.108*** (0.035)	0.095* (0.049)	0.006 (0.006)	0.152** (0.067)	0.211* (0.111)	0.375** (0.170)	0.226 (0.266)	0.053 (0.041)
Δ Share Other IT Workers	0.046 (0.061)	0.056 (0.065)	0.022 (0.015)	0.079 (0.066)	0.057 (0.086)	-0.018 (0.088)	0.024 (0.101)	0.011 (0.019)
Δ Share Robot Workers	0.385** (0.171)	0.274 (0.188)	0.044 (0.038)	0.186 (0.134)	0.085 (0.184)	0.058 (0.270)	0.114 (0.388)	-0.014 (0.058)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y
F Statistic				31.4	31.4	6.1	6.1	6.1
Observations	849	849	849	714	714	849	849	849

Table A.14. Effect of AI Investments on Industry-level Employment and Sales Including Entry and Exit

This table reports the coefficients from industry-level regressions of the changes in total sales and employment for all firms in Compustat (including entrants and exits between 2010 and 2018) on contemporaneous changes in AI investments. Each observation is a 5-digit NAICS industry. 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 AI investments, 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. Columns 1 to 4 are estimated by OLS, and in columns 5 to 8 the independent variable is instrumented by the contemporaneous growth in the share of AI workers in European public firms in the same industry. The dependent variables are changes in log total sales (columns 1, 2, 5, and 6) and log total employment (columns 3, 4, 7, and 8) at the industry level from 2010 to 2018. 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.

Panel 1: AI measure from resume data

	OLS				IV			
	Δ Log Sales		Δ Log Employment		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.129** (0.056)	0.126*** (0.043)	0.124* (0.068)	0.136** (0.059)	0.184*** (0.056)	0.210*** (0.056)	0.210** (0.095)	0.256*** (0.086)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					13.6	20.5	13.6	20.5
Observations	233	233	233	233	152	152	152	152

Panel 2: AI measure from job postings data

	OLS				IV			
	Δ Log Sales		Δ Log Employment		Δ Log Sales		Δ Log Employment	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.163*** (0.062)	0.150*** (0.054)	0.155 (0.103)	0.136 (0.096)	0.201*** (0.057)	0.247*** (0.062)	0.307*** (0.106)	0.386*** (0.117)
NAICS2 FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
F Statistic					7.9	12.6	7.9	12.6
Observations	243	243	243	243	158	158	158	158

Table A.15. Effect of AI Investments on Aggregate 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 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. Panel 1 considers the resume-based measure of AI investments, 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. Columns 1 to 2 are estimated by OLS, and in columns 3 to 4 the independent variable is instrumented by the contemporaneous growth in the share of AI workers in European public firms in the same industry. 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.

Panel 1: AI measure from resume data				
	Dependent variable: Δ Log Employment			
	OLS		IV	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	-0.038 (0.038)	-0.031 (0.039)	-0.013 (0.031)	0.007 (0.022)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
F Statistic			34.1	37.9
Observations	149	149	130	130
Panel 2: AI measure from job postings data				
	Dependent variable: Δ Log Employment			
	OLS		IV	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	-0.003 (0.036)	-0.020 (0.039)	0.042 (0.039)	0.060 (0.042)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
F Statistic			8.6	16.7
Observations	155	155	132	132

Table A.16. Effect of AI Investments on 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-IT sectors) from 2010 to 2014. We consider two measures of productivity: log sales per worker and revenue TFP. 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 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 headquartered (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 6-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 (0.046)	-0.030 (0.040)	-0.021 (0.050)	-0.016 (0.046)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.212	0.265	0.159	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 (0.046)	-0.030 (0.040)	-0.021 (0.050)	-0.016 (0.046)
NAICS2 FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.212	0.265	0.159	0.223
Observations	766	766	720	720

Table A.17. Heterogeneous Effects by Initial Labor Productivity

This table reports the coefficients from long-differences regressions of changes in firm size from 2010 to 2018 on the contemporaneous changes in AI investments among U.S. public firms (in non-IT sectors), separately for each tercile of initial labor productivity. Firms in each 2-digit NAICS sector are divided into terciles based on sales per worker in 2010. We consider three measures of firm size: log sales (columns 1 and 2), log employment (columns 3 and 4), and market share within the 5-digit NAICS industry (columns 5 and 6). The independent variables are changes in the share of AI workers interacted with indicator variables for productivity terciles 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 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 and productivity tercile fixed effects. Regressions in 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 headquartered (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 6-digit NAICS industry level.

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*Productivity Tercile 1	-0.087 (0.109)	-0.094 (0.091)	0.023 (0.107)	0.048 (0.099)	0.011 (0.019)	0.009 (0.017)
Δ Share AI Workers*Productivity Tercile 2	-0.119 (0.079)	0.140 (0.107)	-0.139* (0.077)	0.118 (0.120)	-0.013 (0.013)	0.006 (0.019)
Δ Share AI Workers*Productivity Tercile 3	0.132 (0.081)	0.154** (0.065)	0.115 (0.085)	0.136** (0.069)	0.014 (0.014)	0.015* (0.009)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Productivity tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.195	0.349	0.158	0.287	0.236	0.288
Observations	766	766	766	766	766	766
T-test statistic	2.6	5.3	0.5	0.6	0.0	0.1
T-test p value	0.108	0.022	0.482	0.441	0.900	0.743

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*Productivity Tercile 1	0.086 (0.368)	-0.188 (0.240)	-0.434 (0.358)	-0.618 (0.386)	-0.039 (0.051)	-0.076 (0.061)
Δ Share AI Workers*Productivity Tercile 2	-0.029 (0.111)	0.189*** (0.070)	0.118 (0.128)	0.309** (0.135)	0.030* (0.017)	0.043*** (0.014)
Δ Share AI Workers*Productivity Tercile 3	0.137** (0.054)	0.116*** (0.044)	0.104 (0.070)	0.077 (0.060)	0.010 (0.008)	0.005 (0.006)
NAICS2 FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Productivity tercile FE	Y	Y	Y	Y	Y	Y
Adj R-Squared	0.229	0.407	0.432	0.487	0.170	0.299
Observations	849	849	849	849	849	849
T-test statistic	0.0	1.6	2.2	3.2	0.9	1.8
T-test p value	0.891	0.208	0.140	0.075	0.353	0.186