AI and Services-Led Growth: Evidence from Indian Job Adverts*

Alexander Copestake[1], Max Marczinek[2], Ashley Pople[2], and Katherine Stapleton[3]

[1]International Monetary Fund

[2]University of Oxford

[3]World Bank

February 26, 2023

[Link to latest version]

Abstract

We document near-exponential growth in the demand for artificial intelligence (AI)-related skills in India's services sector since 2016, using a new dataset of online vacancies from its largest jobs website. We evaluate the impact of demand for AI skills on establishment-level non-AI postings in the short term using an event study, and medium term using a shift-share design that exploits variation in exposure to new AI inventions. We find negative effects on postings and wage offers for non-AI postings. The effects are most pronounced in high-skilled managerial and professional occupations and non-routine work, particularly complex analytical and communication tasks.

Keywords: Artificial Intelligence, Labour Markets, Wages, Services, Development
JEL Classification Codes: J23, O33

*Copestake email: acopestake@imf.org, Marczinek email: max.marczinek@economics.ox.ac.uk, Pople email: ashley.pople@economics.ox.ac.uk, Stapleton email: kstapleton@worldbank.org. For support and comments on this project, we are extremely grateful to Abi Adams-Prassl, Richard Baldwin, Stefan Dercon, Paolo Falco, Lukas Freund, Sanjay Jain, Lawrence Katz, Simon Quinn, Ferdinand Rauch and Chris Woodruff, as well as various seminar participants. We are indebted to Sanjay Jain, Sanjeev Bikhchandani, Pawan Goyal, Abhishek Shyngle and Shweta Bajad for help with the vacancy data, and to Rob Seamans for sharing the measure from Felten et al. (2019). This document is an output from the research initiative 'Structural Transformation and Economic Growth' (STEG), a programme funded by the Foreign, Commonwealth & Development Office (FCDO). We also gratefully acknowledge financial support by the Oxford Centre for the Study of African Economies (CSAE). The views expressed herein are those of the authors and should not be attributed to the FCDO or any of the institutions with which the authors are affiliated.

1 Introduction

Rapid advances in machine learning have spurred an intense debate about the labour market consequences of artificial intelligence (AI).[1] Online job adverts show that demand for AI-related skills has grown almost exponentially and concurrently in several countries around the world since 2015 (Figure 1.1). Yet detailed empirical evidence on the extent of AI deployment and its distributional impacts remains scarce, particularly beyond a handful of advanced economies. For low- and middle-income countries, the use cases and impacts of AI need not be the same as for advanced economies. AI could have important consequences for their development trajectory, especially for countries promoting services-led growth. In India, for example, many of the services industries that have driven structural change, productivity growth and job creation, such as Business Process Outsourcing (BPO), are highly exposed to machine learning-based automation, raising questions over the future viability of a services-led development model and its promise of promoting widespread prosperity.

In this paper, we aim to fill this knowledge gap by shedding light on the labour market impacts of AI in India – the archetypical pioneer of a services-led development model. We investigate these effects in the predominantly urban, white-collar services sector using a new dataset of online job adverts posted from 2010 to 2019 on the country's largest online jobs platform. The platform has an estimated market share of around 60 percent of online job postings in India. Following a growing literature including Rock (2019) and Acemoglu et al. (2022), we use the demand for AI-related skills, as observed in the text of posted job descriptions, as a proxy for AI deployment. The basic idea of this approach is that, in the absence of detailed administrative data on firm-level AI adoption, we can induce the demand for AI by studying which firms are hiring machine learning engineers, deep learning specialists and other related staff.

Using the job adverts data, we first document several patterns in AI-related hiring in India. We see a rapid take-off in the rate of 'AI demand' (shorthand for the demand for AI-related skills in job posts) after 2016, albeit growing from a low base. AI demand increased from 0.37% of all job vacancies in 2015 to 1.03% in 2019, coinciding with an increase in demand for specific 'deep learning' skills, along with 'natural language processing' to a lesser extent. Take-up

[1] To fix definitions, we consider artificial intelligence (AI) 'the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages' (Oxford English Dictionary 2020). Machine Learning, the sub-field responsible for many of the recent commercial applications of AI, comprises 'the statistical techniques that enable computers and algorithms to learn, predict and perform tasks from large amounts of data without being explicitly programmed' (Acemoglu & Restrepo 2019). We henceforth use 'AI' as an umbrella term encompassing machine learning.

Figure 1.1: AI share of online job postings, by country

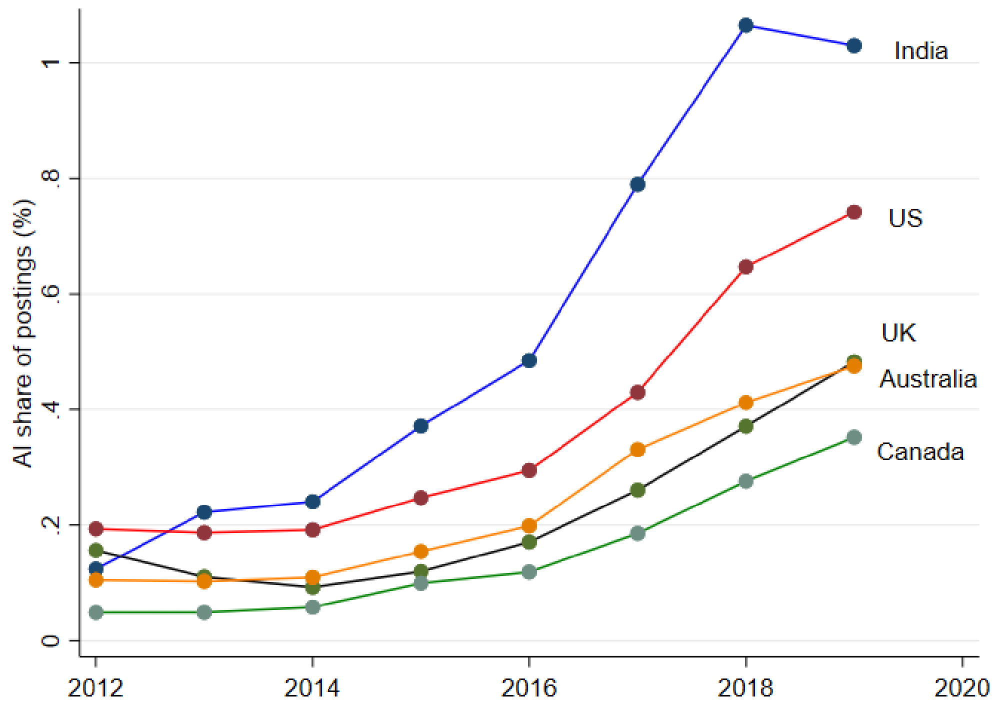

Notes: This graph shows the share of all online vacancies that specify particular AI skills, with these skills defined as described in the text. Data for India is that used in this paper; data for all other countries is from Lightcast, which does not cover India.

is particularly pronounced in the IT, finance and professional services industries. AI roles tend to require substantially more education, particularly graduate degrees, while also paying significantly higher wages. Even after controlling for a host of fixed effects, posts demanding AI skills still pay a 13-17% salary premium, which is similar to the 11% estimate found in the US (Alekseeva et al. 2020). Consistent with Babina et al. (2021) and Zolas et al. (2021), we show that AI roles are heavily concentrated in the largest firms and a few key technology clusters – particularly Bangalore, Mumbai, Hyderabad, Pune, Chennai and Delhi. In line with this spatial clustering, we find evidence of local agglomeration in diffusion: after the first firm posts an advert demanding AI skills in a given industry and region, other firms in the same industry and region are, on average, more likely to start demanding AI skills, even after taking into account industry and region trends.

We then turn to evaluating the impact of demand for AI skills on labour demand within 'establishments', defined as firm-city pairs. The theoretical impact of AI on labour demand is ambiguous. Advances in machine learning have been conceptualised in the literature as reducing the cost, or improving the quality, of the task of 'prediction', which is prevalent in many occupations (Agrawal et al. 2018).[2] While this implies displacement of labour, improvements

[2]For example, a back office employee of a multinational bank takes the input of scrawled handwriting on a

in the task of prediction could also expand labour demand by reducing costs or increasing quality of production, and hence raising productivity.[3] In addition, AI could complement human labour, create new tasks or incentivise changes in organisational structure; there is growing evidence that AI is a general-purpose technology (GPT), an 'invention of a method of invention' (Brynjolfsson et al. 2017, Cockburn et al. 2018, Klinger et al. 2018, Agrawal et al. 2021, Goldfarb et al. 2023).[4] Emerging economies like India could also stand to benefit from new AI value chains, capitalising on their technical engineering workforce and comparative advantage in IT outsourcing (Baldwin & Forslid 2020). Indeed, revenues in India's BPO sector nearly tripled over the past ten years (NASSCOM 2018).

We first evaluate the impacts of AI over the medium term. Using a long difference specification between 2010-12 and 2017-19, we investigate the effects of growth in the demand for AI skills on the growth of non-AI job postings and wage offers at the establishment level. To isolate the impact of AI *demand*, rather than AI *production*, we exclude AI 'producing' sectors from our analysis – specifically IT and education, which are responsible for the vast majority of AI patents (Klinger et al. 2020). We exploit establishment-level variation in their workforce's compatibility in 2010 with future capabilities of AI, as measured by the occupation-level AI exposure measure of Webb (2020). This measure captures the degree of overlap between occupations' tasks and the tasks that patented AI technologies are designed to perform. We combine these occupational AI 'shocks' with the establishment-level occupation vacancy shares at baseline and then use the combined shift-share measure as an instrument for the demand for AI skills. The key idea behind this instrument is that firms in 2010 with a high share of workers conducting tasks that later become feasible to automate with AI, such as civil engineers or actuaries, are more likely to start deploying AI and thus hire new staff with AI skills.

In the first stage, we find that establishments more exposed to AI *ex ante* do indeed see a relative increase in their demand for AI skills in online vacancy posts. Turning to the second stage, growth in AI demand has a significant negative effect on growth in non-AI and total postings at the establishment level. A 1% increase in the AI vacancy growth rate results in a 3.61 percentage point decrease in establishment non-AI vacancy growth between 2010-12 and

mortgage application form and then generates the correctly spelled name of the applicant as predicted output.

[3]Some early research has therefore modelled machine learning in a comparable way to other forms of automation, such as industrial robots (e.g. Webb (2020) and Acemoglu et al. (2022)). These papers build on the canonical framework of Acemoglu & Restrepo (2018) in which task structure determines adoption. Beyond the boundary of the adopting firm, AI could also have broader indirect effects, as workers reallocate across occupations – as explored in detail by Humlum (2019) in the case of robots. Here we focus on direct within-firm effects.

[4]Specifically, GPTs are widely used across sectors; have inherent potential for technical improvement; and spawn further innovation in application sectors (Bresnahan & Trajtenberg 1995).

2017-19, controlling for region, industry and firm size fixed effects. Growth in total establishment vacancies (AI plus non-AI job postings) falls by a similar 3.57 percentage points, reflecting that the increase within the small set of AI posts is far outweighed by the displacement effect in the larger set of non-AI vacancies.

These negative effects on vacancy growth are most pronounced for higher-skilled professional and managerial occupations, notably engineering professionals and general and corporate managers. Using the seminal classification of Autor et al. (2003), we find that AI lowers demand for occupations that are typically non-routine task intensive, in contrast to previous recent waves of technological change that lowered demand for routine tasks. The negative impact on non-routine task-intensive occupations holds both overall and within the affected managerial and professional occupation groupings. We find similar negative impacts within and between occupation groups for 'abstract' task intensity, as defined in Autor & Dorn (2013). We additionally adopt a more granular and flexible approach to measuring task content, following Michaels et al. (2018) by counting verbs in job descriptions and classifying them according to meaning using Roget's Thesaurus. We find that AI adoption reduces demand for verbs related to 'intellectual faculties'. In particular, there is a reduction in the frequency of verbs related to: 'precursory conditions', such as investigate, scrutinize, research, explore, examine; 'extension of thought', such as predict, forecast, anticipate, memorize, recall; and those related to 'means of communicating ideas', such as narrate or describe. The same also holds within the highest paying roles.

How does this displacement affect wage offers for new hires? We estimate that a 1% higher growth rate in AI vacancies reduces the growth rate of non-AI median wage offers by 2.6 percentage points between 2010-12 and 2017-19, instrumenting with AI exposure and controlling for region, industry and firm size. As with vacancy growth, the negative effects of AI demand on wage offer growth are also present when considering all posts, inclusive of AI vacancies. The decline in average wage growth primarily reflects the changing occupational composition as the relative frequency of skilled managerial jobs declines. When controlling for occupation shares, we find that only the top 1% of wages at baseline see a statistically significant decrease in wage offer growth.

We next further consider the impact of hiring AI-related skills in the years immediately after AI hiring starts, using an event-study design combined with propensity score matching. Following Koch et al. (2021), we match AI adopters to never adopters. Specifically, we run a probit regression and construct propensity scores, such that conditional on these, treatment is orthogonal to establishment characteristics. We then run an event study with a balanced panel and establishment and year fixed effects. We find a significant negative effect of AI adoption on the demand for non-AI workers in the first year after adoption, which further increases in magnitude over the following two years.

Finally, we explore wider effects beyond establishments by aggregating the online vacancies dataset to the district level and using data from administrative labour force survey datasets – the National Sample Survey and Periodic Labour Force Survey – on employment and wages. We run the same long-difference instrumental variables and propensity score matching event study specifications at the district level. When we do so, we find no statistically significant effects of AI demand on total hiring or employment at the district level in the short or medium term and only weak negative effects on wages in the medium term.

Taken together, these results suggest that the demand for AI-related skills has already had important effects on Indian service sector firms, altering the distribution of demand and wage offers across occupations and tasks. However, the limited district-level results could reflect that either negative within-establishment effects in AI 'using' industries are offset by hiring growth in new or AI 'producing' firms in the same district, or that growth in AI demand is not yet large enough to impact the wider economy.

This paper makes several contributions to the literature documenting the diffusion of AI and its impacts. First, we offer new insights on the distributional impacts of AI – as proxied through the demand for machine learning – on labour demand and wages, with a particularly granular perspective on the types of jobs affected. Although India has a different labour market structure to high-income countries, our findings of negative within-establishment effects of AI echo those of Grennan & Michaely (2020) and Acemoglu et al. (2022) for the US. Our findings that these effects are driven by higher-skilled occupations, including engineers and verbs relating to analytical tasks, are also consistent with US evidence from Grennan & Michaely (2020) and Webb (2020). The remarkably similar patterns may reflect our focus on a subset of the Indian labour force: the white collar and largely urban services sector where we would expect the highest demand for machine learning. However, our findings challenge the result by Acemoglu et al. (2022) that AI exposure predicts *increased* demand for skill families relating to engineering, analysis, marketing, finance and IT in the US, although they do not examine actual impacts by occupation or skill level. Our evidence that AI negatively affects non-routine task intensive occupations also stands in contrast to findings for previous waves of technological change that were shown to lower demand for routine task intensive occupations (for instance, Autor et al. 2003, Goos & Manning 2007, Goos et al. 2014).

Second, discussions about the impact of AI for development have been largely theoretical to date, so this paper starts to bridge the evidence gap on the consequences of AI adoption for low- and middle-income countries. For instance, Baldwin (2019) and Baldwin & Forslid (2020) have conjectured that machine learning, along with online platforms and software robots, could benefit developing countries by increasing offshoring of services. Korinek & Stiglitz (2021) take an alternative view that developing countries will be negatively affected, because AI devalues their comparative advantage in lower-cost labour and natural resources. Nevertheless, it remains to be seen whether AI has even been deployed within emerging economies. We focus on the domestic angle to these broader questions and show that the growth in the demand for AI skills within India has been remarkably similar, and even more rapid, to that documented in advanced economies. India experienced a similar take-off around 2016, highlighting the global diffusion in AI. Our finding that the demand for AI skills has rapidly grown in professional services could be consistent with a rise in outsourcing of AI work to India. However, while the negative impacts of AI demand on higher skilled jobs could be an equalising force on the labour market in the short term, it remains to be seen whether these effects reflect temporary transition costs associated with new technological adoption or whether they could have perverse longer-term implications. Our negative findings also only concern within-firm effects for incumbent firms and within 'AI-using' industries. Avenues for future research include assessing potentially offsetting effects, such as through firm creation or in 'AI-producing' industries, particularly IT.

Finally, our paper adds to a growing literature which uses online vacancy postings to investigate labour market effects more broadly (e.g. Deming & Kahn 2018, Adams et al. 2020, Javorcik et al. 2020). We contribute through a large new dataset of job postings in India, stretching back to 2010.

The rest of this paper proceeds as follows. Section 2 introduces the data, and Section 3 presents detailed descriptives on AI demand in the Indian white-collar services sector. The next two sections then present our main empirical findings on the impact of AI adoption in hiring on non-AI labour demand and wage offers. Specifically, Section 4 examines medium-term effects of AI using Bartik-type long-difference estimates, while Section 5 provides event-study results on the short-term impact of AI adoption. Section 6 provides robustness checks and extensions. Section 7 concludes. The online appendix provides further detail on the construction of our dataset, additional descriptives and results, and further robustness checks.

2 Data

2.1 Vacancy data

Our primary dataset is vacancy data posted on India's largest online job platform between 2010 and 2019.[5] The site serves primarily as an advertising platform for firms to post vacancies, with subsequent recruitment and hiring processes taking place directly with firms. The platform estimates that they had approximately 60 percent of the market share of Indian online job vacancies in 2020. They shared a randomly-selected 80 percent of all posts over the period 2010 to 2019. We focus on the services sector, for which the data is most representative of overall job vacancies, dropping posts from the manufacturing and agriculture sectors. Our primary dataset hence includes data from around 15.5 million service sector job postings, equating to roughly 1.5 million per year, but skewed towards later years. This compares to an estimated 19 million people formally employed in the service sector in India in 2021, as published by the Ministry of Labour and Employment.[6]

When submitting a vacancy on the platform, firms are required to upload information into a standardised template. Hence, all posts include information on the job title, industry, role category, location, skills required, salary and experience ranges and educational requirements. The job postings also include an open text section for the job description. We manually map industries and occupations into the National Industrial Classification (NIC) at the two-digit level and National Classification of Occupations (NCO) at the four-digit level, covering 99% of all vacancies. We also harmonise city names and add geolocations, separating out overseas job postings. Using the geolocations, we match cities to districts using the 2011 census. Firm names are removed for client anonymity but replaced with a consistent panel identifier.[7]

The two key advantages of the vacancy data relative to administrative datasets are its breadth and detail. In 2018, the vacancy data includes more than two million services vacancies from 40,000 unique firms, compared to 12,000 workers directly surveyed in the 2017-18 Periodic Labour Force Survey and 2,000 firms recorded in Prowess.[8] The representative sample surveys only took place in 2011-12 and 2017-18, so provide no information on short-term fluctuations or more recent developments in the Indian services sector. Prowess, while useful for studying the largest firms, only contains a limited selection. Our vacancy dataset has roughly 30 times as

[5] The company requested to remain undisclosed.

[6] https://static.pib.gov.in/WriteReadData/specificdocs/documents/2022/jan/doc20221104101.pdf

[7] We also focus on full-time jobs and so drop the small number of part-time and non-permanent positions from our sample.

[8] Full observation totals by sector and dataset are shown in Appendix Table A.1.

many firms as Prowess (see Appendix Figure A.2), and neither Prowess nor the labour surveys offer a way to measure AI adoption, whereas we can directly observe the demand for AI skills in the text of online job descriptions.

However, this richness of the vacancy data also comes with certain shortcomings, notably that online vacancies are not representative of all vacancies and online job postings only proxy for firm hiring behavior. Broadly speaking, our vacancy data best represents urban, white-collar service sector jobs. We provide a detailed overview of the coverage and representativeness of the data in Appendix A.2, where we benchmark the vacancy data relative to nationally-representative labour surveys and firm-level data from Prowess. Compared to the administrative data, the vacancy data has a higher share in the IT & BPO, Finance, Insurance and Real Estate, Professional and Business Services industries, and they are under-representative of all other industries. Vacancies posted on the site are disproportionately concentrated in urban centres, but at least one post appears for nearly every district in India.[9] Although the share of formal jobs is lower in India compared to advanced economies, we find a ratio of 0.08 of annual average job postings to total formal employment, which is similar to the ratio of approximately 0.09 for annual US job postings from Burning Glass Technologies relative to total US employment from 2010-2018 as reported in Acemoglu et al. (2022).

2.2 Measuring AI demand

Despite the prominence of the topic of AI in popular discussion, firm-level data on AI adoption remains scarce (for instance, see the discussion in Raj & Seamans 2018). In the absence of data on the adoption of specific technologies, a growing body of work has started using technology-related human capital to proxy for technology adoption. For example, Rock (2019) and Benzell et al. (2019) use LinkedIn profiles to construct firm-level measures of engineering and IT talent, while Harrigan et al. (2020) use the firm-level employment share of 'technology workers' in French matched worker-firm data as a measure of technology adoption.

Human capital is one of the key inputs for deploying an AI system. It is well recognised that one of the primary obstacles to widespread adoption of AI is the available labour supply, with top-tier scientists earning extremely high salaries and being bought out of academic positions. It would be expected that firms wishing to implement an AI-driven automation project would, at least to some degree, need to hire individuals with AI skills or experience. Alternative options would be to rely on external consultants, contract out the process to a third party software

[9]See Appendix Figure B.2 for the detailed geographical distribution of posts.

provider or retrain existing staff to develop AI skills.

There are a number of reasons to believe that the dominant channel is external hiring. Work by McKinsey Global Institute (2019) surveying around 2000 companies globally found that the primary method for sourcing AI talent and capabilities was to hire externally and that the majority of companies built their AI capabilities in house, as opposed to buying or licensing capabilities from large technology companies. Additionally, even if firms were to subcontract AI services, it would be expected that they would still require at least some related human capital in-house to oversee and manage the process. We therefore assume that AI skills demand and actual AI deployment within a firm will be positively correlated and follow this emerging literature in using the demand for AI skills as a proxy for the extent of AI adoption within firms.

Online job vacancy data are well suited to the measurement of demand for very specific technology-related human capital owing to the detailed text data on the skills demanded for specific roles. To measure firm demand for AI skills, we classify job postings based on the text in the job description or skills requirements. Our main classification is the 'narrow' measure employed by Acemoglu et al. (2022), which categorises a post as an AI vacancy if it includes any word from a list of specific AI terms.[10] By using this narrow measure of AI skills, we reduce measurement error, although our estimates of demand for AI skills are likely to be a lower bound of the true level of adoption.

2.3 Verb categorisation

In addition to measuring actual AI-related terms, we also follow Michaels et al. (2018) in using verbs mentioned in the text of job adverts as a proxy for task demand. We use the same list of 1,665 English verbs and the meaning of verbs from Roget's Thesaurus, which classifies words according to their underlying concepts and meanings. The Thesaurus is organized into 6 classes and 38 sections. The 6 classes are: Abstract Relations (ideas such as number, order and time); Space (movement, shapes and sizes); Matter (the physical world and humankind's perception of it by means of the five senses); Intellect (the human mind); Volition (the human will); and Emotion, Religion, and Morality (the human heart and soul). More details can be found in

[10]Specifically, a post is categorised as AI-related if any of the following terms appear in either the 'job description' or 'skills required' fields: Machine Learning, Computer Vision, Machine Vision, Deep Learning, Virtual Agents, Image Recognition, Natural Language Processing, Speech Recognition, Pattern Recognition, Object Recognition, Neural Networks, AI ChatBot, Supervised Learning, Text Mining, Support Vector Machines, Unsupervised Learning, Image Processing, Mahout, Recommender Systems, Support Vector Machines (SVM), Random Forests, Latent Semantic Analysis, Sentiment Analysis / Opinion Mining, Latent Dirichlet Allocation, Predictive Models, Kernel Methods, Keras, Gradient boosting, OpenCV, Xgboost, Libsvm, Word2Vec, Chatbot, Machine Translation and Sentiment Classification.

Michaels et al. (2018).

3 Demand for AI Skills

In this section, we present six descriptive findings about the demand for AI skills in the Indian white-collar services sector, based on our job postings data.

1. AI demand increased rapidly after 2016, particularly in the the IT, finance, education and professional services sectors.

The rate of demand for AI skills increased rapidly after 2016, rising from 0.37% of all job vacancies in 2015 to 1.03% in 2019. Figure 3.1 Panel (a) shows the share of posts that are tagged as AI posts and the top five AI-related phrases which appear. Panel (b) shows the share of AI posts over time and within the top five industries by AI share. Overall we see steady growth from our base year of 2010, which accelerates after 2016, especially in financial services and IT and software. This growth is driven largely by demand for general 'machine learning', with the sub-field 'deep learning' rising rapidly from relative obscurity to being the second most common AI term from 2017 onwards.[11]

Figure 3.1: Trends in AI demand

(a) Most demanded AI skills (b) AI share of posts, by industry

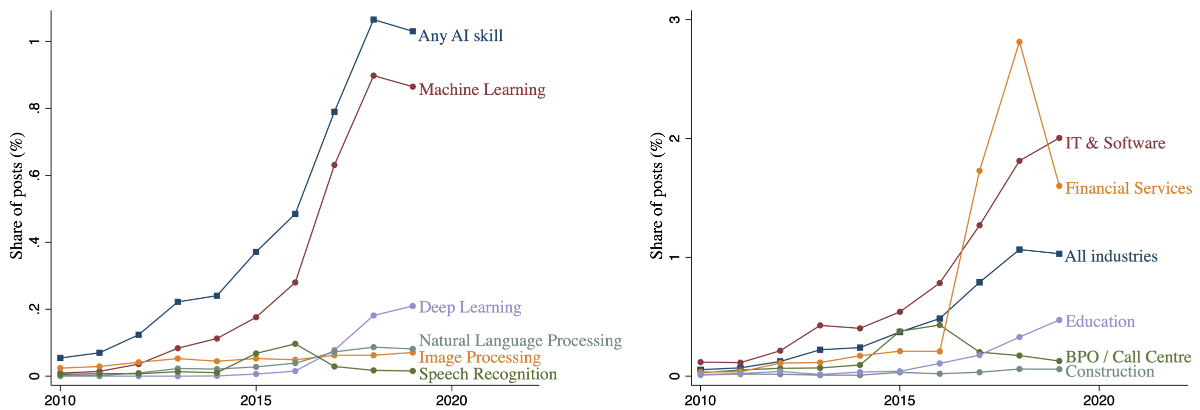

Notes: Panel (a) shows the share of all vacancies that specify particular AI skills, for the top five most demanded skills. Panel (b) shows the share of vacancies that are AI vacancies, both for all industries together and within each of the top five industries by AI share.

The breakdown of demand for AI skills across sectors is also striking, suggesting a wider

[11]For discussion of possible causes of the rapid acceleration, such as the open-source release of Google's TensorFlow software library in 2015, see Stapleton & O'Kane (2020).

diffusion of AI beyond AI-producing sectors. AI demand grew steadily in the IT sector since 2011, whereas AI demand in the financial sector started from a low base then grew ten-fold between 2016 and 2018. In contrast, the business process outsourcing and call centre sector saw a small boom in AI demand in earlier years, corresponding to a spike in demand for 'speech recognition', before returning to a lower share.

2. AI roles require more education and offer substantially higher wages than other white-collar services jobs.

What are these AI roles, and how do they compare to the rest of the vacancies advertised? By far the most common AI role title is 'Software Developer', followed by other technical roles such as 'Data Analyst', 'Technical Lead' and 'Technical Architect' (Appendix Figure B.6). AI skills are also required in technical management roles, with titles as 'Analytics Manager', 'VP - Analytics and BI', and 'Project Manager-IT/Software' also appearing in the top 20 AI-related roles. Yet there is also a long tail of more generalist roles, including 'Business Analyst', 'Trainee', 'Program Manager' and 'Product Manager'. The size of the 'Other' category grouping all other vacancy posts (25%) indicates how widespread the hiring of AI skills is across multiple job titles, albeit each with a small share of overall posts.

Moreover, we see that AI-hiring firms are seeking candidates who are slightly more experienced and substantially more educated than average – and for that they are willing to pay a substantial salary premium (Figure 3.2). AI vacancies are almost twice as likely as non-AI vacancies to require a master's degree, and more than seven times more likely to require a doctorate. They post a median salary of ₹250,000 (approximately US$3,333, without adjusting for PPP), twice the median non-AI salary of ₹125,000 (US$1,666). This 'AI wage premium' remains high (19%) even when controlling for experience, education and firm fixed effects.[12] Even when adding occupation or role fixed effects, to control for AI being used in different types of jobs, a premium of 13-17% remains.[13]

3. AI roles are highly concentrated in a few key technology clusters, particularly Bangalore.

AI demand is highly concentrated in large cities, particularly the major technology clusters

[12]See Models (1) and (2) of Appendix Table B.3 in Appendix B.
[13]See Models (3) and (4) of Appendix Table B.3 in Appendix B.

Figure 3.2: Hiring profile of AI vs. non-AI vacancies

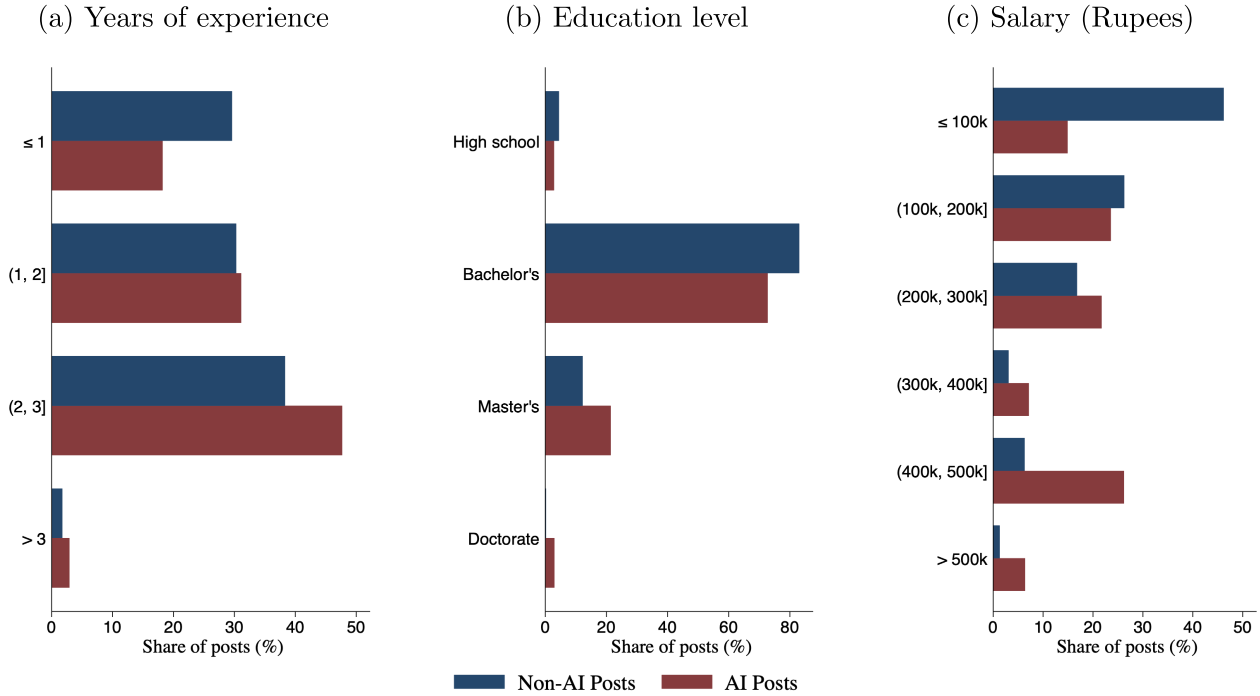

Notes: These graphs compare the distribution of posts, for AI and non-AI vacancies, across experience, education and salary. This information is reported directly in the online jobs platform. For experience and salary, the vacancy posts record a minimum and maximum value, so we take the midpoint of the specified range. AI posts are classified based on keywords, as described in Section 2.2.

around Bangalore, Mumbai, Hyderabad and Delhi. Bangalore alone has more than 30% of all AI vacancies across India. Figure 3.3 compares cities' shares of all posts with their shares of AI posts. AI shares at the top end are larger than general post shares, showing that AI vacancies are even more spatially concentrated than vacancies generally. Shares of AI demand in cities have been remarkably constant over the last decade, except for a prominent increase in AI activity in Mumbai as AI demand took off in the financial sector.[14]

4. AI roles are highly concentrated in the largest firms.

Which firms hire AI skills? We proxy for firm size by the number of vacancies posted on the platform. Figure 3.4 plots the cumulative share of AI posts against the corresponding cumulative share of all posts. This traces out a Lorenz-type curve, where the deviation from the 45° line shows the extent to which AI vacancies are disproportionately posted by the largest firms. Inspecting the top right corner reveals that the largest 14 firms are responsible for 10% of all

[14]See Appendix Figure B.4 in Appendix B for the change over time in the distribution of AI posts across cities.

Figure 3.3: Shares of posts across cities

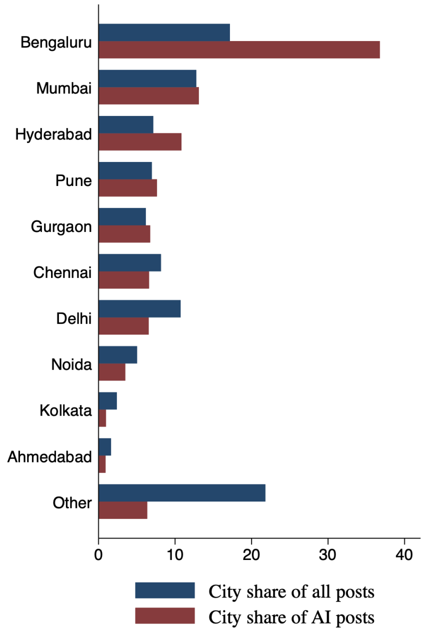

Notes: Bars show the shares of all posts and AI posts across cities, for the entire period 2010 to 2019.

Figure 3.4: AI posts by firm size

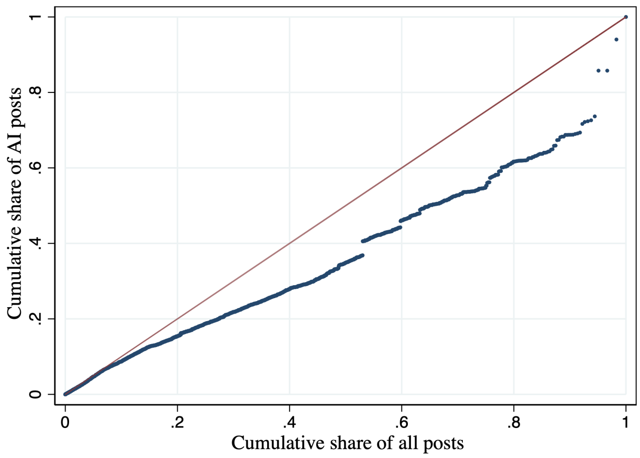

Notes: We plot the cumulative share of AI posts against the corresponding cumulative share of all posts, for the whole period 2010-2019. The red 45° line indicates a one-for-one increase in the share of AI posts relative to all posts. The deviation of our scatter plot from the 45° line shows the extent to which AI vacancies are disproportionately posted by the largest firms.

vacancies, with each posting at least 50,000 vacancies, and these account for 31% of all AI posts. While there are some smaller firms that post a disproportionate number of AI posts, the largest AI-hiring firms are also the largest hirers in general. AI posts are also more concentrated than other vacancies in each individual year, with this excess concentration increasing rapidly from 2015.[15] The take-off in AI demand in Figure 3.1 thus coincided with increased concentration in the hiring of AI skills.

5. Initial AI adoption is associated with agglomeration in AI demand, over and above industry and region trends, particularly in the IT sector.

How does AI diffuse across industries and cities? The vast majority of industry-city pairs had zero AI posts at the start of our period, so we can examine the correlation between the first AI post in a given city and industry and subsequent hiring by other local establishments. We first construct the difference between the current time and adoption of AI for each industry-city pair, K_{irt}. We then construct the outcome AI share$_{irt,-F}$ by pooling all posts in the industry-city *except* those from the first adopter F. We then run the event study specification at the level of

[15] See Appendix Figure B.7 in Appendix B, which plots the trends in the Herfindahl Index for AI and non-AI vacancies over time.

Figure 3.5: Local AI diffusion

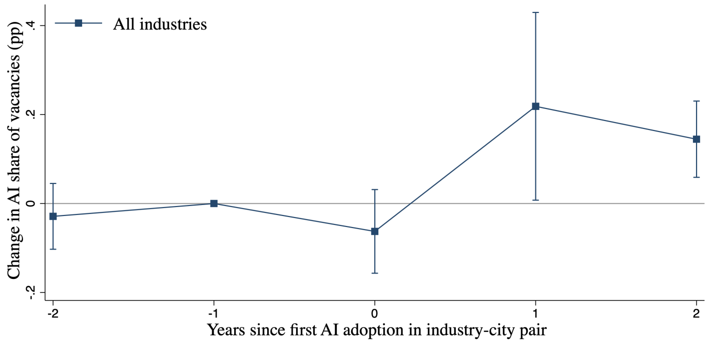

Notes: We plot the change in the industry-city-wise share of vacancies that are AI vacancies, for years before and after the first adoption of AI in an industry-city pair. AI posts by the initial AI adopter are excluded to focus on diffusion of AI posting to other firms in the industry-city pair.

Table 3.1: Verbs in AI posts

	Less common	**More common**
1	Call	Experience
2	Manage	Develop
3	Shift	Build
4	Plan	Program
5	Account	Design
6	Tar	Work
7	Look	Predict
8	Recruit	Deliver
9	Apply	Use
10	Report	Advance

Notes: We count verbs in job descriptions of AI and non-AI job posts and form verb shares. This table shows the verbs with the largest difference in shares between AI and non-AI job posts. Positive (negative) differences imply that the corresponding verbs are more (less) likely to be included in AI posts.

industry-city pairs ir with two leads and lags:

$$\text{AI share}_{irt,-F} = \alpha_{ir} + \alpha_{it} + \alpha_{rt} + \sum_{k=-2\backslash -1}^{2} \beta_k \cdot 1(K_{it} = k) + \epsilon_{irt}. \tag{3.1}$$

This gives the descriptive coefficients β_k, which reflect the average percentage point increase in the AI share of vacancies posted in each year k years from first adoption of AI in the city-industry pair. Crucially, this association is that which remains even after controlling for the broader industry- and city-level trends.

We find that there is a significant positive relationship between initial AI adoption and the share of AI postings by other local firms in subsequent years (Figure 3.5). In the first year after the first AI post within an industry and city, the AI share is more than 0.2 percentage points higher ($p = 0.042$) than that in the absence of an AI adopter. This is a substantial difference considering that the average AI share of posts across all industries was only 1% by 2019 (see Figure 3.1). We also investigate heterogeneity in the diffusion of AI adoption across industries, and find substantial dispersion in the magnitude of the local effect, with by far the strongest relationship in IT and Software.[16] We conclude that while local influence can be relevant for AI diffusion, its importance varies substantially across industries.

6. AI job postings include more complex, creative, and data-driven tasks.

[16]Specifically, we estimate separate βs across industries by interacting $1(K_{it} = k)$ in equation 3.1 with industry dummies. Results are shown in Appendix Figure B.8 in Appendix B.

How do AI jobs differ from non-AI jobs using verbs in the text of job descriptions? Extracting the verbs in AI and non-AI jobs ads using Roget's Thesaurus, we calculate the share of each verb relative to all verbs, and compare verb shares in AI job posts to non-AI job posts. Table 3.1 suggests that AI jobs contain a higher share of complex tasks (e.g. experience and advance), creative tasks (e.g develop, design and experience), and data-driven tasks (e.g program and predict).

4 Medium-Term Impacts of AI

4.1 Empirical approach

Our main specification assesses the impact of increased AI demand on the change in non-AI demand between our baseline period (2010 to 2012) and our endline period (2017 to 2019) – spanning the take-off in AI demand in 2016.[17] We take a long-difference approach common in the automation and labour markets literature to ensure effects do not merely reflect short-term fluctuations in labour demand. Our primary unit of analysis is 'establishments', defined as firm-city pairs, which we use because many firms report postings in several different cities. Our main estimation sample therefore contains almost 25,000 establishments together posting approximately two million vacancies on the platform within our baseline and endline periods.[18] We run:

$$\Delta y_{fr,t-t_0} = \beta \cdot \Delta Adoption_{fr,t-t_0} + \alpha_r + \alpha_i + \alpha_{f10} + \epsilon_{fr,t-t_0}, \tag{4.1}$$

where $\Delta y_{fr,t-t_0}$ is the change in the inverse hyperbolic sine of outcome Y_{fr} between 2010-2012 and 2017-19; $\Delta Adoption_{fr,t-t_0}$ is the change in the inverse hyperbolic sine of the number of AI posts by an establishment between 2010-12 and 2017-19; α_i and α_r are two-digit industry and city fixed effects; and α_{f10} is a firm decile fixed effect, where firm deciles are calculated over the baseline period 2010-2012. We cluster standard errors at the firm level to account for common shocks across establishments within the same parent firm. The key variables $\Delta y_{fr,t-t_0}$

[17] We pool within these periods in order to improve precision and maximise the probability that a firm advertises on the job postings platform during both time periods.

[18] By definition, these establishments are part of relatively large, incumbent firms – as they existed already in 2010-12, and were still operational in 2017-19. We focus on this sub-set of all establishments to allow us to assess the medium-run *establishment-level* impacts of AI, recognizing that other channels (such as new AI-focused startups beginning in the mid-2010s) will also impact *aggregate* labour market outcomes. We discuss these potential channels further in Section 7.

and $\Delta Adoption_{fr,t-t_0}$ hence approximate the growth in establishment outcomes and AI demand between 2010-12 and 2017-19.[19] The coefficient of interest β then approximates the elasticity of each outcome with respect to AI demand: increasing the growth rate of AI demand by 1% between 2010-12 and 2017-19 generates a β percentage point rise in the growth rate of the outcome variable across the same time period.

AI demand is likely to be endogenous with respect to establishment-level outcomes.[20] We therefore instrument AI demand with 'AI exposure' to isolate changes in AI adoption resulting from supply-side technical advances in the capabilities of AI. Specifically, we first take the occupation-level AI exposure measure developed by Webb (2020), which measures the extent to which workers' tasks can be performed by AI technologies using the degree of overlap between the text of AI patents and the text of O*NET job-task descriptions.[21] Occupations with a higher share of tasks that are capable of automation by AI are assigned a higher exposure value. We use publicly-available crosswalks to map the Webb (2020) exposure measure to the Indian National Classification of Occupations (NCO) 2004 at the four-digit level. We then aggregate this measure to the establishment level by weighting across baseline establishment occupation shares, to capture establishment-wise exposure to AI-based automation.[22] Specifically, we calculate:

$$Exposure_{fr,t_0} = \sum_{o} PostShare_{fro}^{t_0} \cdot WebbExposure_o \tag{4.2}$$

where o represents occupations. We then standardize $Exposure_{fr,t_0}$ to have a mean of zero and a standard-deviation of one, and estimate the first stage:

$$\Delta Adoption_{fr,t-t_0} = \gamma \cdot Exposure_{fr,t_0} + \alpha_r + \alpha_i + \alpha_{f10} + \epsilon_{fr,t-t_0}, \tag{4.3}$$

The first stage coefficient γ in Equation 4.3 therefore approximates the proportional change in AI posts between 2010-12 and 2017-19 that is associated with a one standard deviation rise in establishment-wise AI exposure.

[19] Mathematically, for growth rate g defined by $Y_t = (1+g)Y_{t_0}$, and using the approximations that $\ln(1+g) \approx g$ for small g and that $\ln Y_t - \ln Y_{t_0} \approx arcsinhY_t - arcsinhY_{t_0}$, we have $g = \ln Y_t - \ln Y_{t_0} = arcsinhY_t - arcsinhY_{t_0} = \Delta y_{t-t_0}$.

[20] For example, more innovative managers are more likely to hire more AI workers, but they are also more productive and grow the business more quickly, hence also increasing non-AI labour demand.

[21] These task descriptions are based on US occupations. While Indian occupations in general may have very different task compositions, the white-collar service sector is likely to be more similar. To the extent that this is not the case, it would also merely count against the strength of our first stage.

[22] Appendix Figure B.9 in Appendix B shows the distribution of exposure scores across occupation-wise wage offer percentiles. This reveals that AI exposure rises with wage offers up to a peak around the 80th percentile, before falling thereafter.

The establishment-level exposure variables are thus constructed in the shift-share or Bartik (1991) style. Here the 'shares' are establishment posting shares by occupation in 2010-12, and the 'shocks' are occupation varying measures of exposure to AI patents. Recent literature (e.g. Borusyak, Hull & Jaravel (2021), Goldsmith-Pinkham et al. (2020)) has pointed to the conditions under which shift-share instruments are valid. In our setup, we view the case for causal identification as stemming primarily from the plausible exogeneity of the baseline posting shares. Following Goldsmith-Pinkham et al. (2020) we provide several robustness checks to test the validity of this instrument discussed in Section 6 below and Appendix C.1.

We also note that India is not a significant producer of new AI research and lags far behind other major AI research hubs on per capita terms, predominantly the USA, China and Singapore.[23] We therefore do not expect advances in AI patenting globally to be affected by hiring patterns in Indian firms, meaning our shocks could be plausibly exogenous as well. Nevertheless, to ensure that our results are not affected by growth in AI demand in India affecting global patents, in our non-descriptive analysis we also follow Acemoglu et al. (2022) in excluding vacancy posts from AI-producing sectors in order to focus on AI-demanding sectors.[24]

Turning to the relevance of the instrument, we find that AI exposure does indeed predict AI demand. A one standard deviation higher establishment AI exposure score is associated with a significant 1.93% increase ($p<0.01$) in the number of AI vacancies between 2010-12 and 2017-19, after controlling for region, industry and firm decile fixed effects.[25] This relationship is illustrated in Figure 4.1 Panel (a), while Panel (b) confirms that this differential appeared at the same time that machine learning techniques became widely used. For instance, the AI share of vacancies posted by the most exposed quintile of establishments was relatively similar until 2016, before rapidly diverging to reach almost 8% in 2018.

[23]Despite strengths in applied computer science and engineering, India is not a significant producer of new AI patents (Perrault et al. 2019). Appendix Figure B.5 in Appendix B ranks each country on a wide range of AI progress metrics. In terms of the number of AI patents, the USA is dominant, while the USA, China and Singapore are all significant producers of AI conference papers and journal articles.

[24]Specifically, we drop education, IT, internet and e-commerce, telecom and internet service providers, which make up 34.8% of the sample.

[25]Full results are provided in Appendix Table B.4 in the Appendix.

Figure 4.1: Impact of AI exposure on the AI share of establishments' posts

(a) Long differenced AI share vs. exposure

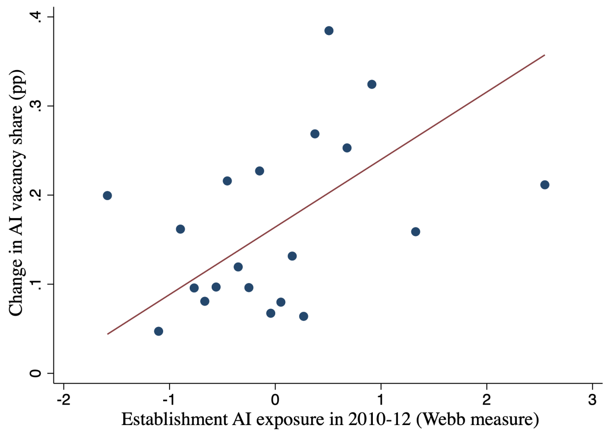

(b) Annual AI share by exposure quintile

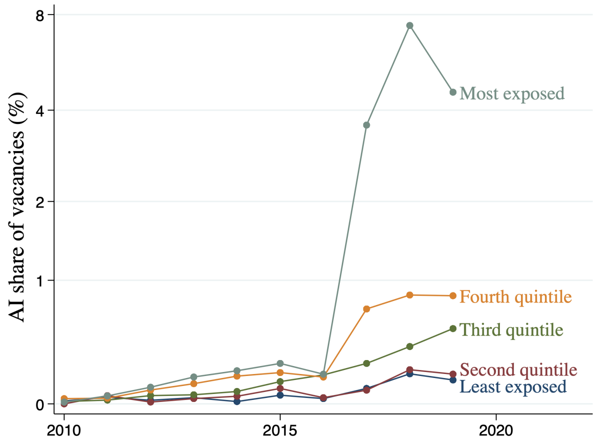

Notes: These graphs show the relationship between AI exposure and the AI share of establishments' posts. The binned scatter plot in (a) summarizes the relationship between baseline AI exposure and establishments' change in AI vacancy share between 2010-12 and 2017-19, after partialling out region, industry and firm-decile fixed effects. Panel (b) plots the time variation in this relationship, using an inverse hyperbolic sine scale for the y-axis.

4.2 Impacts of AI on labour demand

We now turn to the second stage to examine the effects of AI demand on non-AI vacancies. Table 4.1 shows the impact of faster growth in AI vacancies on the growth of non-AI vacancies and total vacancies, instrumenting with the Webb (2020) AI exposure measure.[26] The growth in AI demand reduces the growth in non-AI demand: in our main specification with region, industry and firm-decile fixed effects, a 1% increase in the growth rate of AI vacancies results in a 3.6 percentage point decrease ($p<0.01$) in non-AI vacancy growth at the establishment level between 2010-12 and 2017-19. There is a similarly sized decrease of 3.57 percentage points in the growth rate of total vacancies, highlighting that AI vacancies crowd out other white-collar services-sector vacancies. We find that the median growth rate in total and non-AI vacancies is 24.9%[27].

4.2.1 Mechanisms

To dig deeper into the heterogeneity of effects across the labour market, we explore the impacts on different occupations and tasks within occupations, with the following findings.

[26] To avoid spurious correlation, we use the number of AI posts, rather than the share. For instance, regressing total posts on the AI share would likely have a mechanical negative relationship, as demand shocks for non-AI workers would affect both the denominator of AI share and the outcome variable.

[27] Appendix Figure B.10 in the Appendix examines the dynamics of this effect by repeating Figure 4.1 Panel (b) using non-AI vacancy shares, and again finds that the divergence between more and less exposed establishments coincides with the take-off in machine learning demand from 2016.

Table 4.1: Second stage: Impact of AI adoption on establishment non-AI vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-3.574***	-5.942***	-3.605***	-3.534***	-5.909***	-3.566***
	(1.168)	(1.624)	(1.139)	(1.166)	(1.624)	(1.137)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	26.06	26.31	27.17	26.06	26.31	27.17
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

AI reduces growth in demand for higher-skilled occupations. We first study the effects of AI on postings growth at the occupation level, following India's NCO2004 classification of one-digit and two-digit occupations. Table 4.2 shows that the decline in demand hits higher-skilled occupations: the categories of 'Professionals' and 'Managers' suffer large reductions in their respective growth rates. A 1% increase in the establishment growth rate of AI vacancies results in a 6.2 percentage point decrease in the growth rate of non-AI vacancies for 'Professionals' and a 12.19 percentage point decrease in the growth rate of non-AI vacancies for 'Managers'. In contrast, we see a small increase in the demand for lower-skilled workers, such as 'Personal, Sales and Security', 'Clerks' and 'Associate Professionals'. In the baseline (2010-2012), 'Professionals' and 'Manager' occupations make up 18% and 48% of all postings, respectively. The reduced growth rate for vacancies in these occupations therefore plays a critical role in explaining our aggregate result of lower non-AI vacancies growth. Our finding that AI hiring hits growth in demand for high-skilled occupations aligns with the findings of Webb (2020), who notes that the occupations most exposed to AI are disproportionately high-skilled jobs, i.e. those involving pattern-detection, judgement and optimization, such as clinical laboratory technicians, chemical engineers, optometrists, and power plant operators.

Disaggregating this result further within these two groups using two-digit occupation codes, Appendix Table B.5 shows that increased AI hiring reduces the growth rate of non-AI vacancies for 'Engineering Professionals', 'General Managers', and particularly strongly for 'Corporate

Managers'. In the baseline, 24% of all postings are for 'Engineering Professionals', 6% for 'General Managers', and 12% for 'Corporate Managers'. Reduced vacancies growth for these three occupations consequently plays a large role in our aggregate results.

Table 4.2: Second stage: Impact of AI adoption on establishment non-AI vacancies, by occupation group

	Growth in Non-AI Vacancies				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Growth in AI Vacancies	2.094***	1.092***	5.121***	-6.222***	-12.19***
	(0.487)	(0.354)	(1.252)	(1.581)	(2.632)
Fixed Effects:					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	27.17	27.17	27.17	27.17	27.17
Observations	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Occupation groups are 1-digit occupation groups from the NCO04. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

AI reduces growth in demand for non-routine task intensive occupations. We next aim to understand more about the relationship between impacts on occupations and their task content, following the seminal literature (e.g. Autor et al. (2003) and Acemoglu & Autor (2011*a*)) in examining routine and non-routine task intensity of occupations. We map the measures of Acemoglu & Autor (2011*a*) onto India's NCO2004 classification of occupations and standardize them. We then analyse the impact of AI hiring on the change in IHS-transformed routine and non-routine scores at the establishment-level. As for the main results presented above, we are instrumenting establishment AI hiring by AI exposure in the baseline. Table 4.3 documents the significant negative effect of increased AI hiring on establishment-level growth in non-routine task intensity. A 1% higher growth rate in AI hiring at the establishment level leads to a reduction in the growth rate of non-routine task intensity of 6-7 standard deviations. For routine tasks, we do not find consistently significant effects of AI. Appendix Table B.9 finds similar results for abstract and routine tasks following Autor & Dorn (2013), with a negative impact on abstract tasks and no discernible impact on routine tasks.

Table 4.3: Second stage: Impact of AI adoption on establishment routine and non-routine tasks

	Growth in Non-Routine Tasks			Growth in Routine Tasks		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-5.871***	-7.200***	-5.701***	0.298	0.599**	0.349
	(1.179)	(1.432)	(1.126)	(0.216)	(0.283)	(0.219)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	26.06	26.31	27.17	26.06	26.31	27.17
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. We average standardised routine and non-routine O*NET task contents by occupation, and form establishments' routine and non-routine task demand by weighting occupations by their standardised routine and non-routine scores. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

AI lowers demand for non-routine and abstract tasks within occupation groups. The above result of declining demand for non-routine tasks could reflect changing task demand within occupations or a shift in demand between occupations. We hence next ask how increased AI hiring affects establishment-level non-routine task intensity within occupation groupings, as shown in Table 4.4. We find that increased AI demand also reduces demand growth for non-routine tasks within all broad occupation categories. We find the strongest reduction in non-routine task growth for the category of 'Managers', which makes up 48% of job posts in the baseline. Appendix Table B.10 repeats Table 4.4 for abstract tasks following Autor & Dorn (2013), finding similar results of a decline in growth in demand for abstract tasks within all occupational groupings except Clerks and a coefficient with the largest order of magnitude for Managers.

AI reduces demand for analytical and complex communication tasks. The above measures of task intensity of occupations relied on time-invariant measures of the task content of occupations from O*NET. AI could also impact the task content within occupations over time. We hence next aim to take a further, more granular, approach by studying the verbs mentioned within the text of job adverts. We use the verb classification of Roget's Thesaurus, as discussed above. We count verbs in all job postings, dividing by the total number of verbs counted in each job post. These are then aggregated at the establishment level for the baseline

Table 4.4: Second stage: Impact of AI adoption on establishment non-routine tasks, by occupation group

	Personal, Sales and Security	Clerks	Associate Professionals	Professionals	Managers
	Growth in Non-Routine Tasks				
Growth in AI Vacancies	-0.920***	-0.785***	-2.801***	-0.832***	-4.870***
	(0.263)	(0.241)	(0.487)	(0.277)	(1.067)
Fixed Effects:					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	27.17	27.17	27.17	27.17	27.17
Observations	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. We average standardised routine and non-routine O*NET task contents by occupation, and form establishments' routine and non-routine task demand by weighting occupations by their standardised routine and non-routine scores. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

and endline, such that we are able to run similar regressions to those in Table 4.1, with the dependent variable now changing verb usage, which we interpret as changing task demand.

Figure 4.2a shows that AI has a negative effect on the share of verbs in one category: those relating to the 'Intellectual Faculties'. A 1% higher AI hiring growth rate leads to a 13.8 percentage point lower growth in the share of verbs in the category of 'Intellectual Faculties' in job postings. Figure 4.2b shows that within this category AI hiring has statistically significantly negative effects on the share of verbs in three sections: 'Precursory Conditions', 'Means of Communicating Ideas' and to a lesser extent 'Extension of Thought'. The first involves analytical tasks like 'investigate', 'research' and 'explore'. The second involves complex communication tasks like 'narrate' and 'describe'. The final involves tasks like 'predict' and 'forecast'. These results are in line with the notion that machine learning reduces the cost or improves the quality of the task of 'prediction' (Agrawal et al. 2018). We also find that these results also hold when keeping only the top 1% highest paid jobs per establishment, suggesting that AI also reduces demand for these same tasks within the highest paid occupations. These results are found in Appendix Figures B.12 and B.13.[28]

[28]When keeping only the top 1%, results are only significant at the 10% level, perhaps due to the smaller sample size.

Figure 4.2: Impact of 1% higher establishment AI hiring growth on verb usage by class and section

(a) By verb class

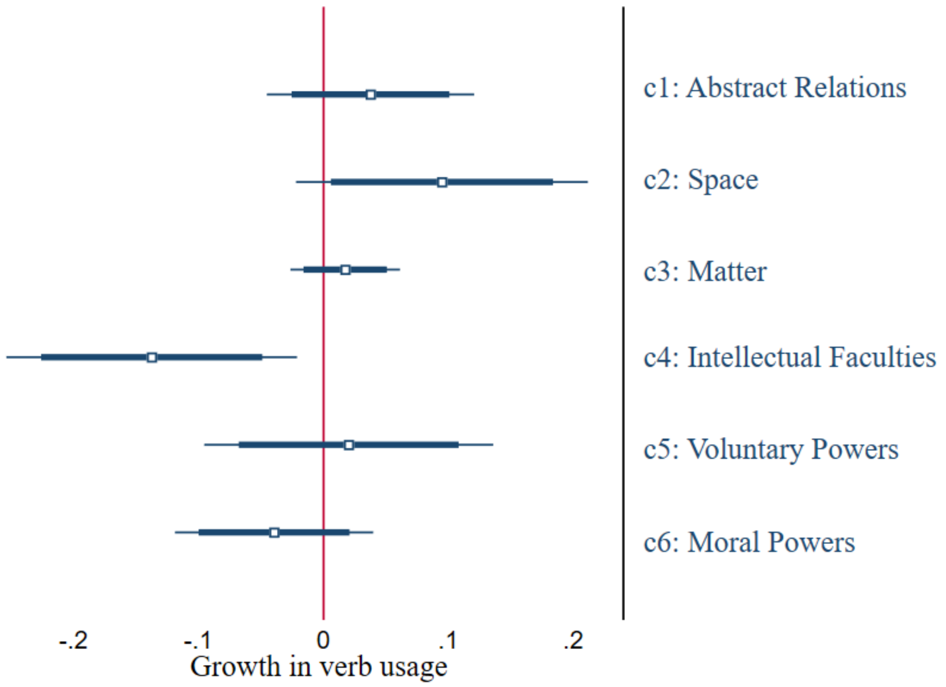

(b) By verb section within Intellectual Faculties, and example verbs

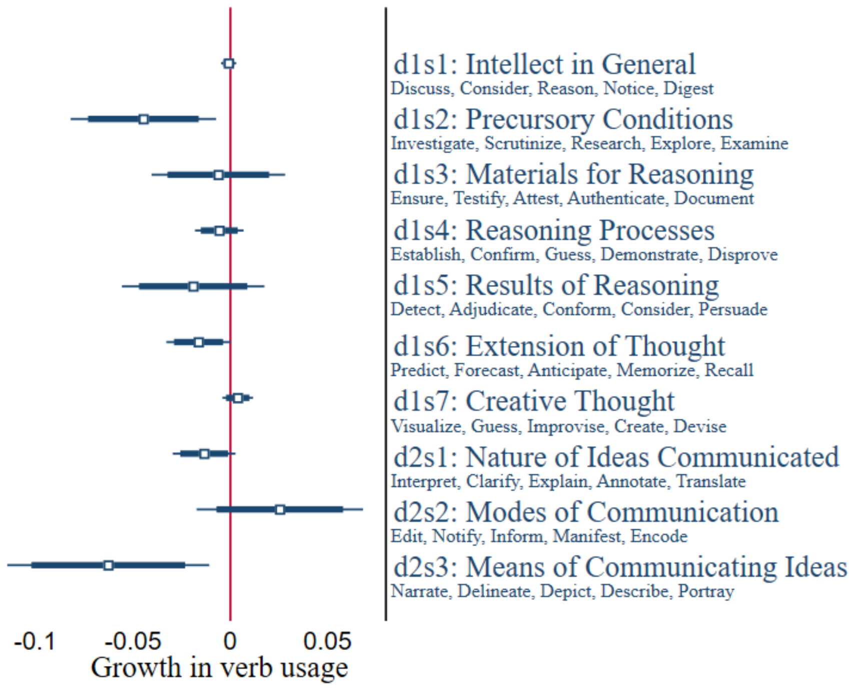

Notes: These coefficient plots show the impact of increased establishment AI demand on verb share growth between 2010-2012 and 2017-2019, where verb shares are formed from counting verbs in job descriptions of job ads. Point estimates accompanied by 95% and 90% confidence intervals. Each coefficient is from a regression of type (2) in Table 4.1. Here, the outcome variable is growth in the IHS-transformed share of verbs from the respective section or class. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on verb share growth. AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects.

4.3 Wage results

Turning to wage offers, we take advantage of the fact that firms complete a standard template to advertise on the platform, providing us with wage offer data for all vacancies. Table 4.5 documents the impact of higher growth in AI vacancies on the growth of median wages for non-AI postings and all job postings. A 1% higher growth rate in AI vacancies between 2010-12 and 2017-19 reduces the growth rate of non-AI wage offers by 2.6 percentage points ($p<0.01$) across the same time period, again instrumenting with AI exposure and controlling for region, industry and firm size. As with vacancy growth, the negative effects of AI demand are hardly changed when considering all posts, inclusive of AI postings.

Table 4.5: Second stage: Impact of AI adoption on establishment non-AI wages

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-2.703***	-3.101***	-2.599***	-2.632***	-3.017***	-2.527***
	(0.799)	(0.895)	(0.758)	(0.770)	(0.862)	(0.730)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	25.32	25.64	26.39	26.61	26.84	27.71
Observations	22,064	22,064	22,064	22,071	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

The reduction in the growth rate of non-AI wage offers in response to increased AI demand occurs across the entire wage offer distribution after the 20th percentile. Figure 4.3 illustrates the percentage point impact of a 1% higher growth rate in AI demand on the growth rate of a given percentile of the wage offer distribution at establishment level between 2010-12 and 2017-19, instrumented with AI exposure and controlling for region, industry and firm size fixed effects. Across the wage offer distribution from the 20th percentile onwards, we observe a statistically significant reduction (at the 5 % level) in establishment level wage offers for non-AI jobs over time, ranging from -1.3 to 2 percentage points ($p<0.05$). Growth in wage offers at the

10th percentile is weakly statistically significant, at the 10% level, with a t-statistic of -1.92.[29][30] Mid-wage offers are most affected by the changes in wage growth over time, although these results are not statistically significantly different from the changes to low and high-wage offers.

Figure 4.3: Impact of AI demand on the wage offer distribution in non-AI posts

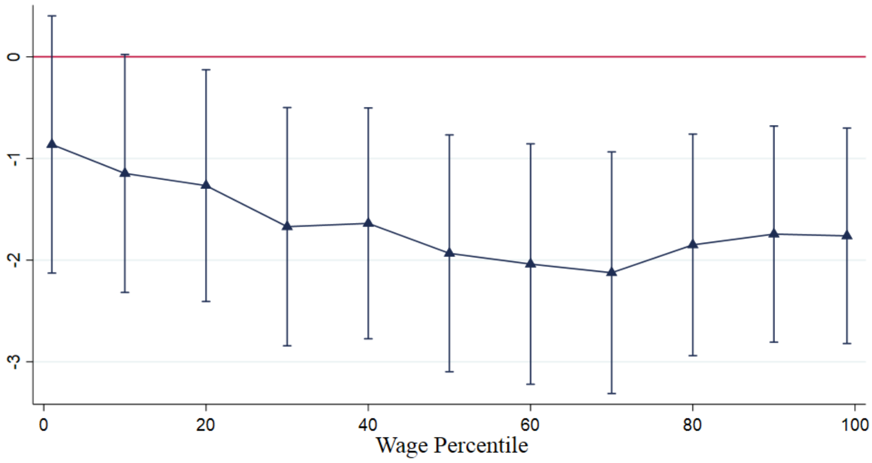

Notes: This coefficient plot shows the impact of increased establishment AI demand on wage growth over time across the distribution of establishment wage offers. Each coefficient is from a regression of type (2) in Appendix Table B.7. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on wage growth over time for a given percentile of the wage offer distribution. We report the 1st and 99th percentile of the wage offer distribution and deciles in between the two extremes, alongside 95% confidence intervals. As in Appendix Table B.7, AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects. Since AI posts make up only a small share of all roles in most establishments, the pattern is very similar across the distributions for all posts and for non-AI posts only.

These findings of lower wages could be driven by either between occupation effects, with AI changing the occupational composition and position of the median wage offer, or within occupation effects, with AI affecting wage offers for the same occupations. We showed in Table 4.2 that AI lowers growth in demand for the highest paid occupations, and raises demand for the lowest paid occupations. This likely explains part of the downward shifting wage distribution. Splitting the sample by occupation group results in small sample sizes and so it is challenging to evaluate wage impacts within occupation groups.[31] However, we additionally control for changing occupation group shares and find evidence for within occupation effects driving these

[29]The negative effect on the wage growth at the 1st percentile of the wage offer distribution is the only non-statistically significant result at conventional levels of significance.

[30]While we observe a fall in wage offers for each decile over time, we remain agnostic about the change in the composition of jobs at each decile.

[31]Although we do so in Appendix Figure B.6, finding weak negative effects

results in Figure 4.4. Accounting for the changing occupational composition, only the top 1% highest paid jobs see a decline in wage offer growth. This effect is statistically significant at the 10% level.

Figure 4.4: Impact of AI demand on the non-AI wage offer distribution, holding occupational composition fixed

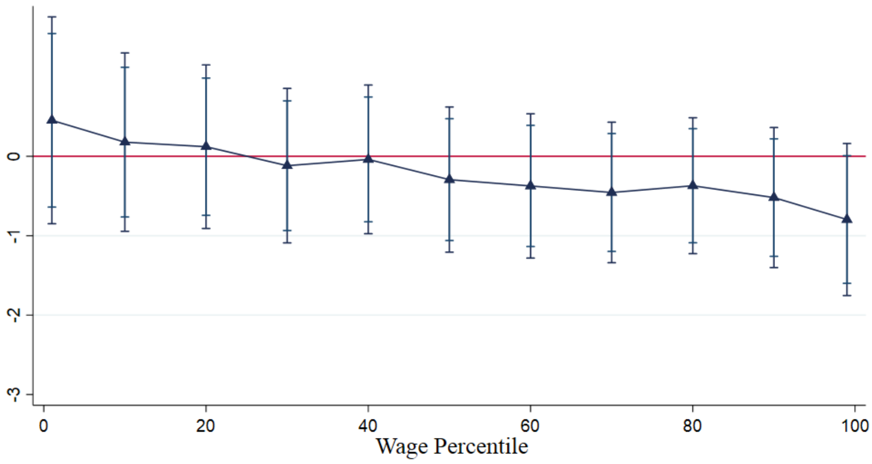

Notes: This coefficient plot shows the impact of increased establishment AI demand on wage growth over time across the distribution of establishment wage offers, controlling for the change in shares of 1-digit NCO04 occupations, leaving out occupation 2 ('Professionals'), as it is the largest occupation category in the baseline period. Each coefficient is from a regression of type (2) in Appendix Table B.7. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on wage growth over time for a given percentile of the wage offer distribution. We report the 1st and 99th percentile of the wage offer distribution and deciles in between the two extremes, alongside 95% confidence intervals. As in Appendix Table B.7, AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects.

We also explore whether these wage offer results hold when additionally controlling for changes in the job requirements specified in the vacancy posts. We find that the reduction in the growth of non-AI wage offers persists, even when controlling for changes in education and experience requirements over time. Appendix Table B.7 shows the effect of the growth of AI vacancies on the growth of the median wage offers of non-AI posts and all posts, after controlling for the growth in the average experience and education levels. Even when controlling for changes in job requirements over time, a 1% higher growth rate in AI vacancies reduces the growth rate in the median wage offers of non-AI posts by 1.93 percentage points (Column 2) and of all posts by 1.89 percentage points (Column 5) between 2010-12 and 2017-19, both precisely estimated at the 1% level of significance. These coefficients are, however, slightly lower than those in Table 4.5, suggesting that both potential explanations play a role. Greater AI demand both changes the type of workers hired, and reduces the wage offer conditional upon worker profiles.

4.4 Wider effects

How does AI demand affect broader outcomes beyond establishments? We attempt to answer this question in two ways. First, we aggregate the job postings to the district level and run the same regression specification. With the same cross-sectional long difference specification the sample size is now substantially smaller than that for establishments. We still find that AI exposure predicts AI demand (Appendix Table B.11 column (1)), but we do not find any statistically significant results for non-AI or total vacancies growth (columns (2) and (3)). We do find weakly significant (at the 10% level) negative results for the effect of AI demand on non-AI and overall wage offer growth.

Second, we also use alternative datasets for the outcome variables of employment and wages from the 2011-12 National Sample Survey (NSS) and 2017–18 Periodic Labour Force Survey (PLFS) administrative datasets, which are nationally representative labour force surveys. We combine our existing district-level first stage with a second stage using these labour force survey derived outcome variables. Similarly, we still find no statistically significant effects for employment or wages. Results are shown in Appendix Table B.12.

One interpretation of these results could be that while AI has had some negative within-establishment effects, these have been offset by positive effects in other establishments in the same district. However, given the small sample size for these results, more work will need to be done to understand in more detail these wider effects.

4.5 Taking stock

Our key finding is that rising demand for AI skills within establishments lowers growth in non-AI vacancies and total vacancies. We showed that this negative impact was driven by lower growth in demand for higher-skilled jobs (professional and managerial occupations), while AI has a positive impact on growth in demand for personal, sales and security occupations and clerks. In terms of the content of work being affected, the negative effects are also underpinned by falling demand for non-routine task intensive occupations, both across the board and within our five broad occupational groupings. Within the professional and managerial categories, AI concurrently raises demand for routine-task intensive occupations. In terms of the specific verbs being mentioned in the job postings, AI lowers mentions of verbs relating to analytical tasks and complex communication. This also holds both across the board and within the top 1 percent highest paying jobs in each establishment.

For wages, we showed that rising demand for AI skills within establishments also lowered wage offers in non-AI posts and wages at each percentile of the establishment wage distribution. Most of this effect appears to be driven by the impacts of AI on changing the occupational composition within establishments, but we also find some evidence of lower wage offers for the highest paying roles. The reduction in the growth of non-AI wage offers also persists, even when controlling for changes in education and experience requirements over time, suggesting that AI both changes the type of workers hired, and reduces wage offers conditional upon worker profiles.

5 Short-Term Impacts of AI

In addition to the medium-terms impacts of AI, we also explore the impacts of AI immediately after demanding AI skills, using an event study design. We showed in Section 3 that AI adopters are larger and pay higher wages than non-AI adopters and so AI adoption is unlikely to be random. Following Koch et al. (2021), we therefore match AI adopters to similar non-adopters using propensity scores. We construct these by running a probit regression of AI adoption on a range of firm characteristics, and hence deriving predicted adoption probabilities.[32] These characteristics control for the fact that AI adopters are generally larger and more mature companies with higher productivity.[33] Conditional on our propensity scores, AI adoption is orthogonal to observable establishment characteristics.[34]

We then run an annual propensity score weighted event study regression at the establishment level with three lags and leads. We estimate at the level of establishments i:

$$Y_{it} = \alpha_i + \alpha_t + \sum_{k=-3\backslash-1}^{2} \beta_k \cdot 1(K_{it} = k) + \beta_{3+} \cdot 1(K_{it} \geq 3) + \epsilon_{it}, \tag{5.1}$$

where Y_{it} is the number of job postings, using an inverse hyperbolic sine (IHS) transformation. α_i and α_t respectively are establishment and time fixed effects. K_{it} is the time difference between the current year and adoption of AI. ϵ_{it} is the error term. The parameters β_k are the outcomes of interest. We include three lags and leads, omitting the first lead to set the baseline level.

[32] Specifically, we include lags of firm size decile, hiring, salary, experience levels, firm age, the standard deviation of salaries and experience, and a number of interaction terms.

[33] Appendix Table B.13 affirms that more productive establishments with a greater share of high-skilled workers are more likely to adopt AI.

[34] Appendix B.3 provides further details on the propensity scores, including results from the probit regression.

Figure 5.1: Event study for non-AI postings following AI adoption

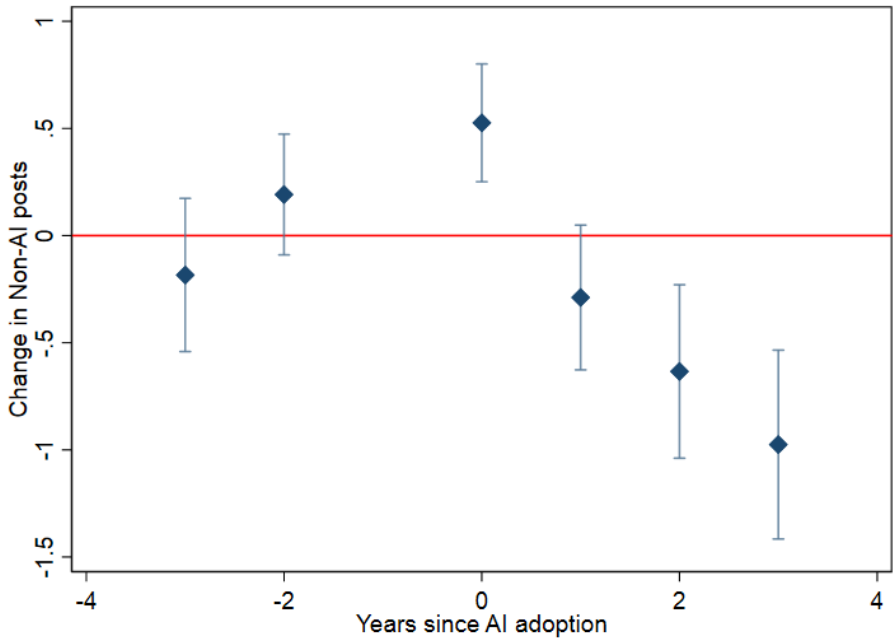

Notes: We use two way fixed effects on a balanced panel. The outcome variable is IHS-transformed establishment non-AI hiring. Adopters are matched to never adopters by propensity score weighting, with propensity scores from a probit regression of establishment characteristics on AI adoption, documented in Appendix Table B.13. We use three leads and lags, leaving out the first lead (t-1) as the base period, and cluster standard errors on establishments.

Further, we account for non-hiring following adoption by balancing the panel, i.e., by imputing zero hiring where no vacancy posts are observed.

We find that, following an initial positive impact of AI adoption on non-AI hiring, non-AI hiring is significantly lower (at the 5% level) in the second year after adoption. This negative impact increases in magnitude in the third year following AI adoption. Using the IHS transformation, coefficients can be interpreted as semi-elasticities: non-AI hiring is 0.7% lower for adopters in the second year after adoption, and 1% lower three years after adoption. These results are robust to the imputation estimator of Borusyak, Jaravel & Spiess (2021). Similarly to the medium-term results, however, when we aggregate to the district level we do not find significant results, as is shown in Appendix Figure B.14 in Appendix B.3.[35]

[35] For wage offers, we cannot balance the panel using the same imputation method, as vacancies that were never posted have no wage offer.

6 Robustness

6.1 Shift-share validity

The instrument used in our long-difference specification is constructed in a Bartik or shift-share format. The 'shares' are establishment posting shares by occupation in 2010-12, and the 'shocks' are occupation varying measures of exposure to future AI patents in the USA. Recent literature (e.g. Borusyak, Hull & Jaravel (2021), Goldsmith-Pinkham et al. (2020)) has pointed to the conditions under which shift-share instruments are valid. In our setup, we view the case for causal identification as stemming from the plausible exogeneity of the baseline posting shares. Following Goldsmith-Pinkham et al. (2020) we provide three robustness checks to test the validity of this instrument, which are discussed in more detail in Online Appendix C.1. First, we investigate the correlates of the shares, finding that the instrument does not appear to be correlated with baseline controls. Second, we test for pre-trends by investigating whether the baseline occupation shares predict year-on-year growth in employment or wages, finding low predictive power. Finally, we compare a range of estimators and run over-identification tests, finding similar results across estimators that further allay concerns.

6.2 Alternative exposure measures

For our main specifications, we use the AI exposure measure proposed by Webb (2020), as it measures which tasks overlap with capabilities outlined in AI patents. Webb (2020) also validates the measure against previous IT and robotic trends. However, alternative AI exposure measures have also been proposed in the literature to date. Therefore, we examine whether our results remain robust to alternative definitions of our instrument.

We first consider the AI exposure measure proposed by Felten et al. (2018). Their AI Occupational Impact measure draws on data from the AI Progress Measurement project from the Electronic Frontier Foundation. The data identify nine application areas in which AI has made progress since 2010. Felten et al. (2018) crowdsource assessments on the applicability of these applications to 52 O*NET ability scales using Amazon MTurk. The AI Occupational Impact assigns an AI exposure score to each O*NET occupation as the weighed sum of the 52 O*NET ability assessments, where the weights are equal to the O*NET-reported prevalence and importance of each ability within each occupation. We map the Felten et al. (2018) measure to Indian NCO using a publicly available crosswalk (see Appendix A).

The results on wages remain robust to the use of the Felten et al. (2018) AI exposure instrument. We first observe that the AI exposure predicts AI demand in the first stage (Appendix Table C.4). Turning to the second stage, we observe that the negative effects on the growth of wage offers in response to increased AI demand remains robust to the use of Felten et al. (2018) as an instrument. A 1% higher growth in AI demand results in a 1.51% decrease (p<0.05) in the growth rate in wage offers between 2010-12 and 2017-19 (Appendix Table C.8). Moreover, we similarly observe a negative effect on the growth rate across the entire wage offer distribution, except for the very lowest percentiles (Appendix Figure C.1). Although we do not observe any significant effects on the growth of non-AI vacancies for the full sample (Appendix Table C.5), we do find similar negative effects for professional and managerial occupations (although now positive effects for associate professionals) and for non-routine task intensity at the establishment level.

We also consider the Suitability for Machine Learning (SML) methodology from Brynjolfsson et al. (2018), which uses surveys to score O*NET direct work activities against a rubric of suitability for machine learning (e.g. inputs and outputs are machine-readable, feedback is immediate, task is principally concerned with matching or prediction, etc.). We use an India-specific version of the SML index created by Mani et al. (2020), who interviewed more than 3000 Indian employees using the SML rubric and mapped a SML score onto every occupation in the 2004 NCO at the four-digit level. However, the SML exposure measure fails to predict firm demand for AI skills using our job vacancy data. One explanation for these differences could be that the Webb (2020) measure is based on current patented technological capabilities, whereas the SML measure is more forward-looking in its predictions. This result that the SML index does not predict AI demand was also found in Acemoglu et al. (2022) suggesting that the limited predictive power is not limited to India only.

6.3 Alternative controls and specifications

We also explore several alternative specifications, with results provided in Appendix C. We first address concerns that establishments experiencing a rise in demand for machine learning skills are also software-engineering intensive firms or firms more affected by computerization by adding additional controls for the baseline establishment posting shares for software engineers and sales and administrative professionals. Results are shown in Figures C.10 and C.12, demonstrating that the key conclusions still hold when adding these controls.

One further concern could be that the number of AI postings by an establishment is too noisy a proxy for the intensive margin of AI demand. We hence also explore the impact of the extensive margin of AI demand where demand is a binary variable. Tables C.13 to C.16 use an AI adoption dummy as the instrument, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. Table C.13 shows that establishment-level AI exposure predicts this AI adoption dummy. The results in Table C.14 maintain a significant negative impact of AI adoption on growth in non-AI and overall vacancies while tables C.15 and C.16 maintain similar results for non-AI and overall median wage growth. The interpretation of the estimated coefficients is different for the adoption dummy: adopters of AI have (approximately) 7% lower growth in non-AI vacancies and total vacancies, as seen in columns (2) and (5) of Table C.14. Columns (2) and (5) of tables C.15 and C.16 indicate an (approximately) 5% and 4% lower growth in non-AI and overall median wages, respectively, for adopters of AI between the baseline and endline. Tables C.17 and C.18 maintain similar results for a specification with the dependent variable in logs.

Another concern could be that our long difference specification is too restrictive in focusing only on the subset of firms that posted on the platform in both the baseline and endline. Tables C.19 to C.21 show the results from using a shorter difference, between 2013-15 and 2017-19. These draw on a larger sample, as fewer firms drop out over the shorter period, at the cost of not spanning the entire takeoff in AI hiring (Figure 3.1). Nonetheless, the results are similar to those in the main specifications. Increased machine learning hiring leads to significant and substantial reductions in non-machine learning hiring (Table C.19) and wages (Table C.20), including when controlling for job profiles (Table C.21). This shorter long-difference also allows us to differentiate between newer firms ('start-ups') and older firms in the sample from the start ('incumbents').[36] We classify an establishment as an incumbent if it posts vacancies in the years 2010-2012, and as a start-up if it only starts posting vacancies after that period. In Tables C.22 and C.23, we show that the negative employment and wage results are driven by incumbent establishments.

The wage results are robust to using mean rather than median wage offers (Table C.24). Our results are also robust to weighting by baseline establishment size, proxied by number of job postings advertised between 2010 and 2012, with the top 5% winsorised (Tables C.25, C.26).

[36] In the long-difference before, every establishment present had to be posting in the baseline period in order to be included, such that all establishments are, in that sense, incumbents.

7 Conclusion

Rapid innovations in machine learning could reshape many jobs, raising questions about the distributional impacts of AI. While the extent of AI diffusion has been measured in high-income countries and particularly the US, there is little evidence on its deployment in low- and middle-income countries. AI could have important implications for a services-led development model, given the potential for recent advances in machine learning to automate many of the constituent tasks of many white-collar services occupations.

In this paper, we use a new dataset of online vacancy posts from India's largest jobs website to shed light on the demand for AI skills in the predominantly urban, white-collar services sector. There was a rapid take-off in demand for AI-related skills after 2016, particularly in the IT, finance and professional services industries, closely mirroring patterns found for advanced economies (Grennan & Michaely 2020, Acemoglu et al. 2022). We evaluate the labour market impacts of establishment-level demand for AI skills, as a proxy for AI deployment. We use a long-difference shift-share specification that exploits variation in establishments' exposure to patented advances in AI capabilities (Webb 2020).

We find that growth in AI demand has a significant negative impact on the growth in non-AI postings and average wage offers by establishments. These negative effects on vacancy growth are most pronounced for higher-skilled professional and managerial occupations, notably engineering professionals and general and corporate managers. Using the classification of Acemoglu & Autor (2011*b*), we find that AI lowers demand for occupations that are typically non-routine task intensive, both overall and within the affected managerial and professional occupation groupings. This stands in sharp contrast to findings for computerisation and robotics, which have been shown to lower demand for routine tasks (Autor et al. 2003, Goos & Manning 2007, Goos et al. 2014). Taking an even more granular approach and classifying verbs in the text of the job adverts using Roget's Thesaurus, we find that AI adoption reduces demand for verbs related to 'intellectual faculties', particularly those relating to investigation, prediction and narration. These results are in line with the notion that machine learning reduces the cost or improves the quality of the task of 'prediction' (Agrawal et al. 2018). When we unpack the negative effects on wage offers and control for the shifting occupational distribution, we find negative wage offer effects only for the top 1% highest paid jobs, suggesting that our wage results are primarily explained by a change in the composition of hiring.

We use an event study design combined with propensity score matching to evaluate the immediate effects on vacancy posting after firms advertise AI-related roles. We find that the demand for AI skills reduces demand for non-AI roles in the subsequent three years. We also explore wider effects at the district level using the long difference shift-share empirical strategy. In contrast to the findings at the establishment level, we find no evidence of negative impacts of AI hiring on total labour demand and only weak negative effects on wage offers at the district level.

Taken together, we find sizable impacts of AI on high-skilled, non-routine, analytical work within establishments in India's predominantly urban, white collar service sector. These effects of AI in its early years of adoption contrast with those for the computer and robot revolutions, where routine work lost out. AI jobs pay a substantial wage premium, but these opportunities are highly concentrated in certain industries, cities and large firms, providing benefits for a small group with AI skills at the expense of demand for other types of high-skilled worker. However, these displacement effects are driven by older 'incumbent' firms that are not AI 'producers'. When we look for wider effects at the district level, we find little evidence of displacement. This could suggest that negative within-establishment effects are offset by growth in new startups, including those focused on AI production, or that the aggregate impacts of AI are not yet economically meaningful. Tracing the aggregate impacts of future advances in, and deployment of, AI on innovation and employment will an important task for future research.

References

Acemoglu, D. & Autor, D. (2011*a*), Chapter 12 - skills, tasks and technologies: Implications for employment and earnings, Vol. 4 of *Handbook of Labor Economics*, Elsevier, pp. 1043–1171.

Acemoglu, D. & Autor, D. (2011*b*), Skills, Tasks and Technologies: Implications for Employment and Earnings, *in* 'Handbook of Labor Economics', Vol. 4, Elsevier, pp. 1043–1171.

Acemoglu, D., Autor, D., Hazell, J. & Restrepo, P. (2022), 'Artifial Intelligence and Jobs: Evidence from Online Vacancies', *Journal of Labor Economics* **40**.

Acemoglu, D. & Restrepo, P. (2018), 'The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment', *American Economic Review* **108**(6), 1488–1542.

Acemoglu, D. & Restrepo, P. (2019), 'The wrong kind of AI? Artificial intelligence and the future of labour demand', *Cambridge Journal of Regions, Economy and Society* **13**(1), 25–35. **URL:** *https://doi.org/10.1093/cjres/rsz022*

Adams, A., Balgova, M. & Qian, M. (2020), 'Flexible Work Arrangements in Low Wage Jobs: Evidence from Job Vacancy Data', p. 46.

Adão, R., Kolesár, M. & Morales, E. (2019), 'Shift-Share Designs: Theory and Inference*', *The Quarterly Journal of Economics* **134**(4), 1949–2010.
URL: *https://doi.org/10.1093/qje/qjz025*

Agrawal, A., Gans, J. & Goldfarb, A. (2018), *Prediction Machines: The Simple Economics of Artificial Intelligence*, illustrated edition edn, Harvard Business Review Press, Boston, Massachusetts.

Agrawal, A. K., Gans, J. S. & Goldfarb, A. (2021), AI Adoption and System-Wide Change, Technical Report w28811, National Bureau of Economic Research.

Alekseeva, L., Azar, J., Gine, M., Samila, S. & Taska, B. (2020), 'The demand for AI skills in the labour market'.

Autor, D. H. & Dorn, D. (2013), 'The growth of low-skill service jobs and the polarization of the us labor market', *American Economic Review* **103**(5), 1553–97.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.103.5.1553*

Autor, D. H., Levy, F. & Murnane, R. J. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *The Quarterly Journal of Economics* **118**(4), 1279–1333. Publisher: Oxford University Press.
URL: *http://www.jstor.org/stable/25053940*

Babina, T., Fedyk He, X. & Hodson, J. (2021), 'Artificial intelligence, firm growth, and product innovation', *Management Science* . Forthcoming.

Baldwin, R. (2019), *The Globotics Upheaval: Globalisation, Robotics and the Future of Work*, W&N.

Baldwin, R. & Forslid, R. (2020), Globotics and Development: When Manufacturing is Jobless and Services are Tradable, Working Paper 26731, National Bureau of Economic Research.

Bartik, T. (1991), 'Who Benefits from State and Local Economic Development Policies?', *Upjohn Press* .
URL: *https://research.upjohn.org/up$_p$ress/77*

Benzell, S., Lagarda, G. & Rock, D. (2019), Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine-Biased Technological Change, 2000-2016, SSRN Scholarly Paper ID 3427396, Social Science Research Network, Rochester, NY.

Borusyak, K., Hull, P. & Jaravel, X. (2021), 'Quasi-Experimental Shift-Share Research Designs', *The Review of Economic Studies* **89**(1), 181–213.
URL: *https://doi.org/10.1093/restud/rdab030*

Borusyak, K., Jaravel, X. & Spiess, J. (2021), 'Revisiting event study designs: Robust and efficient estimation'.
URL: *https://arxiv.org/abs/2108.12419*

Bresnahan, T. F. & Trajtenberg, M. (1995), 'General purpose technologies 'Engines of growth'?', *Journal of Econometrics* **65**(1), 83–108.

Brynjolfsson, E. & Mitchell, T. (2017), 'What can machine learning do? Workforce implications', *Science* **358**(6370), 1530–1534.

Brynjolfsson, E., Mitchell, T. & Rock, D. (2018), 'What Can Machines Learn, and What Does It Mean for Occupations and the Economy?', *AEA Papers and Proceedings* **108**, 43–47.

Brynjolfsson, E., Rock, D. & Syverson, C. (2017), Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics, Technical Report w24001, National Bureau of Economic Research.

Cockburn, I. M., Henderson, R. & Stern, S. (2018), The Impact of Artificial Intelligence on Innovation, Technical Report w24449, National Bureau of Economic Research.

Deming, D. & Kahn, L. B. (2018), 'Skill requirements across firms and labor markets: Evidence from job postings for professionals', *Journal of Labor Economics* **36**(S1), S337–S369.
URL: *https://doi.org/10.1086/694106*

Felten, E. W., Raj, M. & Seamans, R. (2018), 'A Method to Link Advances in Artificial Intelligence to Occupational Abilities', *AEA Papers and Proceedings* **108**, 54–57.

Felten, E. W., Raj, M. & Seamans, R. (2019), The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization, SSRN Scholarly Paper ID 3368605, Social Science Research Network, Rochester, NY.

Goldfarb, A., Taska, B. & Teodoridis, F. (2023), ‘Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings’, *Research Policy* **52**(1), 104653.
URL: *https://www.sciencedirect.com/science/article/pii/S0048733322001743*

Goldsmith-Pinkham, P., Sorkin, I. & Swift, H. (2020), ‘Bartik instruments: What, when, why, and how’, *American Economic Review* **110**(8), 2586–2624.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20181047*

Goos, M. & Manning, A. (2007), ‘Lousy and lovely jobs: The rising polarization of work in britain’, *Review of Economics and Statistics* **89**(1), 118–133.

Goos, M., Manning, A. & Salomons, A. (2014), ‘Explaining job polarization: Routine-biased technological change and offshoring’, *American Economic Review* **104**(8), 2509–2526.

Grennan, J. & Michaely, R. (2020), ‘Artificial Intelligence and High-Skilled Work: Evidence from Analysts’, *SSRN Electronic Journal* .
URL: *https://www.ssrn.com/abstract=3681574*

Harrigan, J., Reshef, A. & Toubal, F. (2020), ‘The March of the Techies: Job Polarization Within and Between Firms’, *Research Policy* p. 104008.
URL: *http://www.sciencedirect.com/science/article/pii/S0048733320300871*

Humlum, A. (2019), ‘Robot Adoption and Labor Market Dynamics’, p. 79.

Javorcik, B., Kett, B., Stapleton, K. & O’Kane, L. (2020), ‘Unravelling Trade Integration: Local Labour Market Effects of the Brexit Vote’, p. 40.

Klinger, J., Mateos-Garcia, J. & Stathoulopoulos, K. (2018), ‘Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology’, *arXiv:1808.06355 [cs, econ]* .

Klinger, J., Mateos-Garcia, J. & Stathoulopoulos, K. (2020), ‘A narrowing of AI research?’, *arXiv:2009.10385 [cs]* .

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

Korinek, A. & Stiglitz, J. E. (2021), ‘Artificial Intelligence, Globalization, and Strategies for Economic Development’, p. 41.

Mani, D., Tomar, S., Madan, N. & Bhatia, A. (2020), The Impact of AI on the Indian Labour Market, Technical report, Srini Raju Centre for IT & the Networked Economy, Indian School of Business.

McKinsey Global Institute (2019), AI adoption advances, but foundational barriers remain, Technical report, McKinsey.

Michaels, G., Rauch, F. & Redding, S. J. (2018), 'Task Specialization in U.S. Cities from 1880 to 2000', *Journal of the European Economic Association* **17**(3), 754–798.

NASSCOM (2018), 'Facts & Figures', https://nasscom.in/knowledge-centre/facts-figures.

Oxford English Dictionary (2020), *Oxford English Dictionary*, Oxford University Press.

Perrault, Shoham, Brynjolfsson, Clark, Etchemendy, Grosz, Lyons, Manyika, Mishra & Niebles (2019), The AI Index 2019 Annual Report, Technical report, AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA.

Raj, M. & Seamans, R. (2018), *Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data*, University of Chicago Press, pp. 553–565.
URL: *http://www.nber.org/chapters/c14037*

Rock, D. (2019), Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence, SSRN Scholarly Paper ID 3427412, Social Science Research Network, Rochester, NY.

Stapleton, K. & O'Kane, L. (2020), 'Artificial Intelligence and Service-sector Offshoring', p. 20.

Webb, M. (2020), 'The Impact of Artificial Intelligence on the Labor Market', *SSRN Electronic Journal* .

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L. & Dinlersoz, E. (2021), Advanced technologies adoption and use by us firms: Evidence from the annual business survey, Technical Report w25727, National Bureau of Economic Research.
URL: *https://www.nber.org/papers/w25727*

AI and Services-Led Growth: Evidence from Indian Job Adverts

Alexander Copestake[1], Max Marczinek[2], Ashley Pople[2], and Katherine Stapleton[3]

[1]International Monetary Fund

[2]University of Oxford

[3]World Bank

February 26, 2023

ONLINE APPENDIX

This supplementary online appendix contains three sections, in turn providing further details on the construction of the data, additional descriptives and results, and further robustness checks. Appendix A lays out how we construct our dataset, benchmarks it against administrative data and discusses representativeness. Appendix B provides additional descriptives on demand for AI skills in Indian job postings, and provides additional results for the short-term and medium-term impacts of AI, including at wider levels of aggregation. Finally, Appendix C provides detailed robustness checks, including tests of our shift-share instrument and results using alternative measures of exposure to AI, alternative specifications, and alternative data sources.

Appendix A Data Details

This paper uses four main datasets: vacancy data from India's largest jobs site; balance sheet data from Prowess, which contains longitudinal balance sheet data on all publicly-listed and many large private Indian firms; and nationally representative labour surveys conducted in 2011-2012 (the National Sample Survey) and in 2017-2018 (the Periodic Labour Force Survey). Table A.1 summarises the number of observations across these datasets.

In this Appendix, we provide more details on the composition of the data and the construction of the variables used. We first describe how we classify the occupations, industries and locations in the vacancy data. We then assess the representativeness of the vacancy data by benchmarking it against Prowess and the nationally-representative labour surveys.

Table A.1: Number of observations by data source

Online vacancy postings 2010-2019	#Firms	#Posts
Agriculture	13,811	463,675
Manufacturing	57,980	2,543,995
Services*	167,969	15,481,330
— *Financial*	*17,805*	*1,815,798*
— *Information*	*72,057*	*5,834,878*
— *Professional*	*38,533*	*834,932*
— *Other*	*106,798*	*6,995,722*
Prowess (balance sheets)	#Firms	#Observations
Agriculture	123	590
Manufacturing	2,276	11,257
Services	3,675	16,722
— *Financial*	*1,020*	*4,830*
— *Information*	*516*	*2,557*
— *Professional*	*199*	*811*
— *Other*	*1,940*	*8,524*
Surveys (demographics)	#Districts	#Households
NSS 2012	626	101,725
PLFS 2018	646	102,063

Notes: Some services firms post in multiple sub-sectors, hence the total number of services firms is less than the sum of all firms posting in the sub-sectors.

A.1 Construction of the vacancy dataset

The largest online job postings platform in India scraped and shared 80% of all job postings (randomly sampled) from 2010 to 2019. All posts include text data on the job title, industry, role category, location, skills required, salary and experience ranges and educational requirements. We manually map 99% of role titles to the 2004 Indian National Classification of Occupations (NCO) at the four-digit level. We also manually map all industries to the 2008 Indian National Industrial Classification (NIC) at the two-digit level. We clean 95% of city names and add geo-locations, separating out oversees job postings. Using the geolocations, we match cities to districts, using the 2011 census.

We also use publicly-available crosswalks to translate the AI exposure measures to the Indian context. We map the 2000 Standard Occupation Classification used by Webb (2020) to the 2004 Indian National Classification of Occupations (NCO) via the 1988 International Standard Classification of Occupations (ISCO), at the four-digit level. For the Felten et al. measure, we map the 2008 ISCO to the 1988 ISCO, before again mapping onto the 2004 NCO.

A.2 Representativeness of the vacancy data

In this section we evaluate the representativeness of our vacancy data in relation to the broader Indian labour market by benchmarking against widely-used administrative datasets and labour surveys. Using the January 2022 Quarterly Employment Survey by the Indian Labour Ministry, India's services sector is estimated to formally employ 18.9 million workers.[1] The QES gives formal employment figures by sector excluding agriculture. Our paper excludes manufacturing so we focus on services throughout. The breakdown by sub-industry, compared to our dataset of online services job adverts, is shown in Figure A.1. Compared to the QES our vacancies data is over-representative of IT BPOs and Financial Services. On the other hand, it is under-representative of Hospitality, Education, Health, Transportation, and Trade.

The industry distribution of services firms in the vacancy data and the Prowess firm dataset are shown in Figure A.2 Panel (a). The distribution of vacancies is shown in Panel (b), alongside the distribution of equivalent white-collar service sub-industries in the pooled National Sample Survey (NSS) and Periodic Labour Force Survey (PLFS).[2] The vacancy data has relatively fewer finance, insurance and real estate firms than Prowess, but a greater share in that sector relative to the representative labour surveys. The national surveys also report relatively more workers in education and transportation, likely because they include public sector workers, whereas the vacancies and Prowess balance sheet data include only private firms. Panel (c)

[1] https://static.pib.gov.in/WriteReadData/specificdocs/documents/2022/jan/doc20221104101.pdf

[2] We define white-collar services workers in the NSS context as salaried workers in divisions 1-5 of the 2004 Indian National Classification of Occupations, i.e. excluding agricultural, fishery, craft, manufacturing, elementary and unclassified workers.

Figure A.1: Distribution of formal workers in the QES compared with vacancy posts

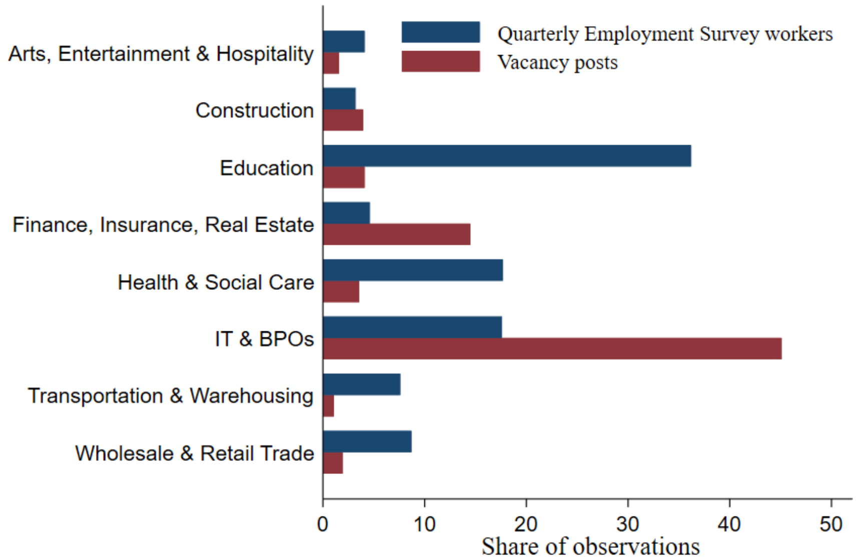

Notes: Excludes the sectors 'Other Services' and 'Administrative' which feature in the vacancies data but do not match to an equivalent sector in the Quarterly Employment Survey. Agriculture, fishery, and manufacturing are excluded in both.

shows the distribution of occupations in the vacancy data in contrast to the national surveys. As would be expected, the vacancy data is over-representative of high-skill white-collar jobs and under-representative of lower-skilled jobs, such as shop assistants or security guards, which are more typically filled through referrals and offline hiring.

Figure A.2: Comparison of vacancy data with Prowess firm-level data and labour force surveys

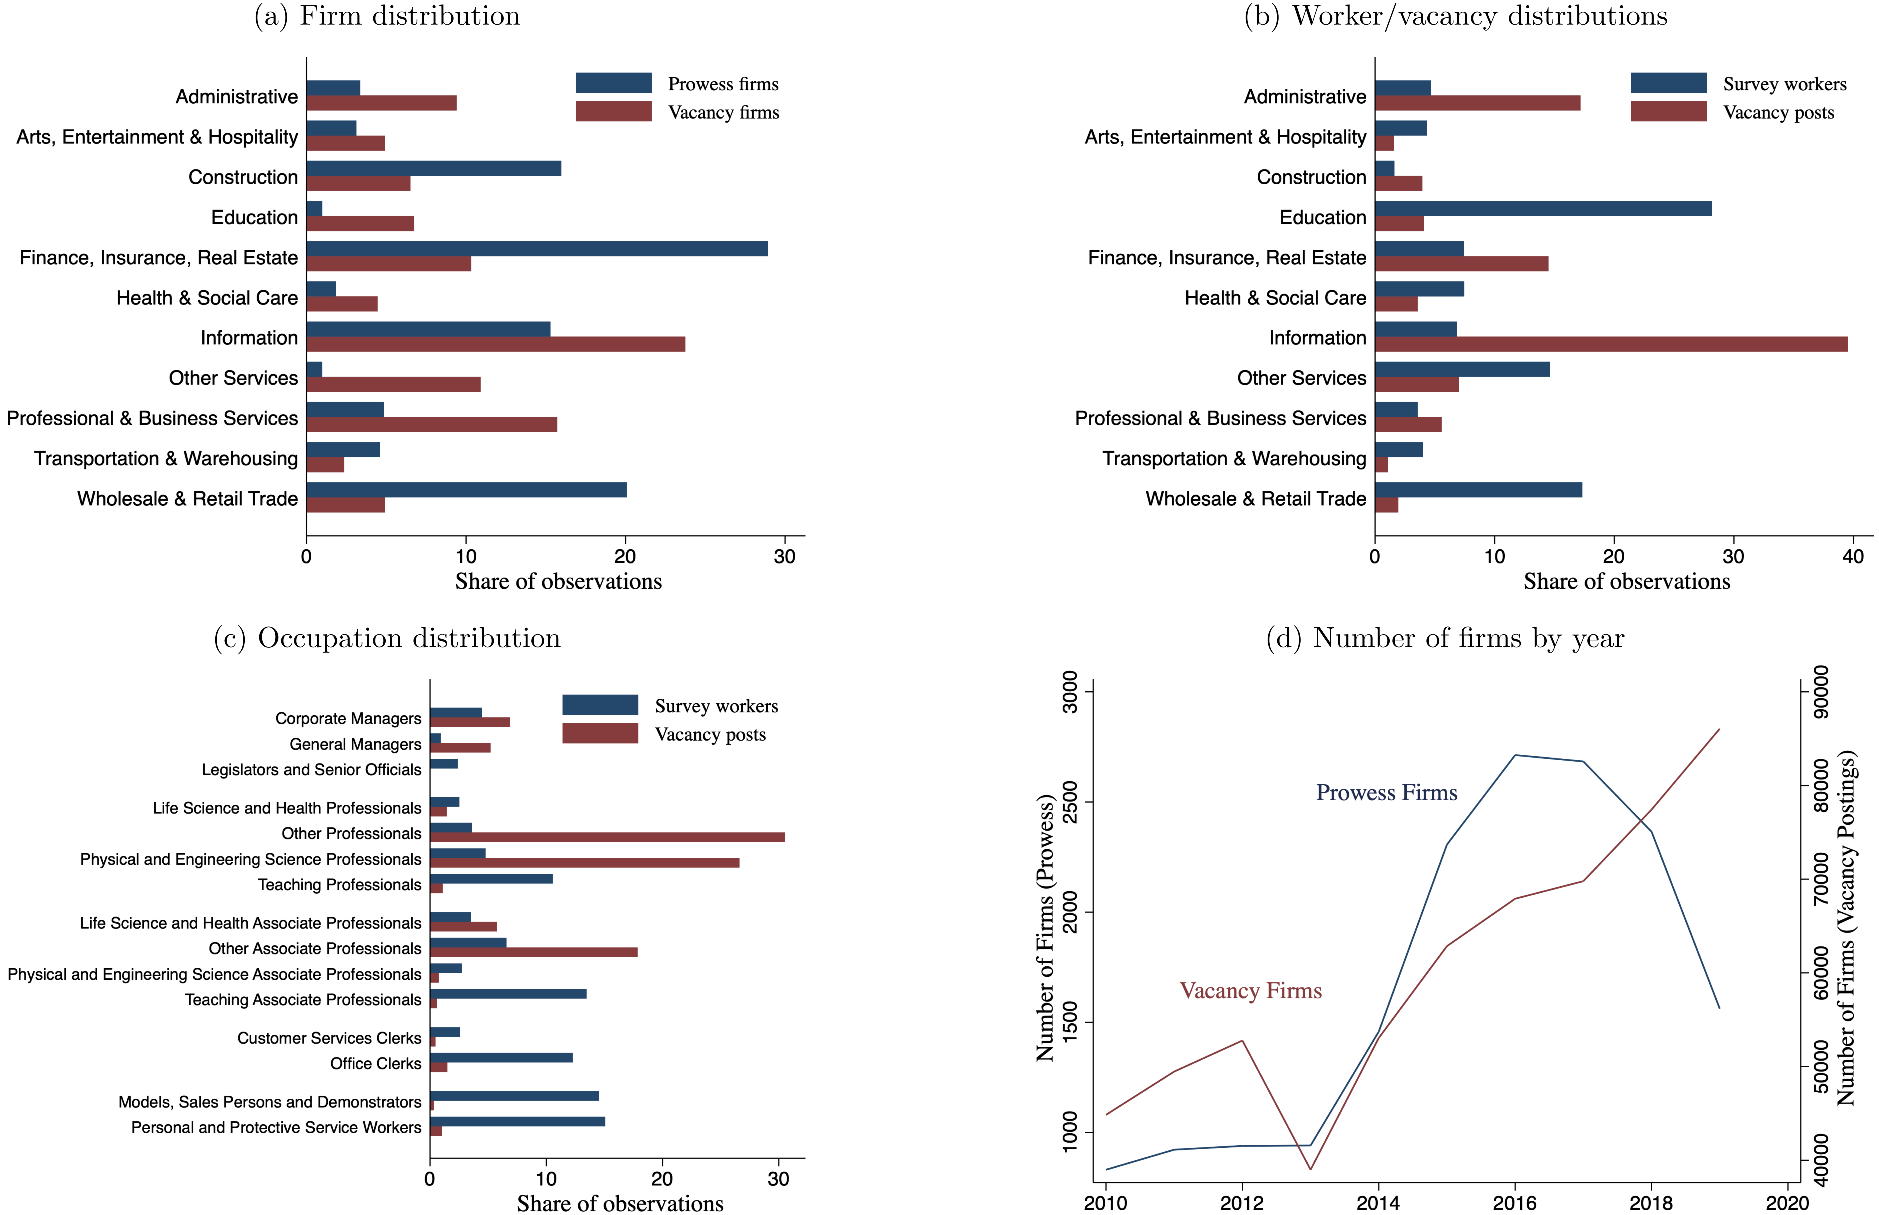

Notes: These figures compare the composition of our vacancy dataset (red) to that of available administrative datasets (blue). Panel (a) shows the distribution of firms across industries relative to Prowess. Panel (b) compares the distribution of vacancies to that of workers in the NSS and PLFS. Panel (c) shows the distribution of white-collar services occupations relative to NSS and PLFS. Panel (d) compares the number of firms in the vacancy data to that in Prowess.

Appendix B Additional Descriptives and Results

B.1 Descriptives

Details on the restricted sample: Table B.1 shows key descriptives across postings in the restricted sample. This sample includes only establishments which post vacancies in both the baseline (2010-2012) and endline (2017-2019), and only posts in service sector, AI-demanding industries. Figure B.1a shows how shares of all postings for broad occupation groups evolve over time, and Figure B.1b shows how shares of postings for AI in broad occupation groups evolve over time.

For the restricted sample, we pool 2010-2012 as the baseline and 2017-2019 as the endline. Accordingly, the growth rates in Table B.2 are defined for the endline over the baseline, and not year on year.

Calculating the AI wage premium: When including industry-region, industry-time and region-time fixed effects, we find that AI posts on average offer 30% higher wages than non-AI posts (see Model (1) of Figure B.3). However, this may be driven by the highest-paying firms also disproportionately hiring AI roles. Therefore, we add firm fixed effects to control for differences between firms in Model (2). Even in this case, AI posts pay 19% more relative to the average non-AI post. Finally, posts that require AI skills may simply be different types of jobs. Models (3) and (4) therefore include fixed effects for the occupation and role, using respectively, the NCO 2004 classification codes and the more granular role label built into the online jobs site. A substantial AI premium of 13-17% remains.[3] Table B.3 shows the AI wage premium controlling for job characteristics.

Geographic breakdown of postings: In Figures B.2 and B.3, we show the distribution of total posts and the share of AI jobs across districts. The distribution of cities' shares of AI posts over time is shown in Figure B.4.

[3]The interpretation of the control variables is as follows. An extra year of experience is associated with a more than 35% higher salary (at least within the predominantly early-career jobs posted on the site – see Figure 3.2), while having a Master's degree is associated with up to 10% higher salary. In this sample, having only a high school education is associated with wage offers 3-6% below the baseline of having an undergraduate degree, though this figure is likely a dramatic underestimate of the effect, given the major under-representation of lower-skilled professions on the platform. The relationship between wage offers and having a doctoral degree is expressed predominantly through the firm- and role-effects: conditional on firm and occupation/role, there is no significant relationship to salary, but without such conditioning salaries are 7-13% higher. This is consistent with the wage offer premium for workers with doctorates being driven by taking higher-skilled jobs at more advanced firms.

Table B.1: Key descriptives of vacancy data in restricted sample, all jobs and AI jobs only

Characteristic	All Jobs	AI Jobs
# Job Posts	4,101,323	38,230
Ratio of # Job Posts in 2019 to # Job Posts in 2010	2	61
% Posts in Baseline (2010-2012)	36	3
% Permanent Positions	99.6	99.75
# Establishments	104,578	2,140
% Posts by 5 Most Active Establishments	1	19
% Posts by 100 Most Active Establishments	12	60
Mean # Posts Per Establishment	167	18
Median # Posts Per Establishment	29	3
90th Percentile Posts Per Establishment	312	20
10th Percentile Posts Per Establishment	5	1
Mean # Posts per Establishment and Year (2010, 2019)	(29, 51)	(2, 11)
# Posting Establishment Per Year (2010, 2019)	(14,000, 17,000)	(107, 1,222)
Mean # Reposts	3.8	.8
Median # Reposts	4	0
90th Percentile Reposts	6	2
10th Percentile Reposts	1	0
# Firms	5,605	1,123
% Posts by 5 Most Active Firms	10	51
% Posts by 100 Most Active Firms	43	81
# Cities	442	52
% Posts in Bengaluru	21	43
% Posts in Top 5 Cities	64	84
% Posts in Top 10 Cities	85	98
# Industries	63	39
% Posts in IT Services	32	49
% Posts in BPO and Call Centres	11	2
% Posts in Banking and Financial Services	10	31
% Posts in Research and Analytics	1	5
# 4-digit Occupations (NCO 2004)	96	59
% Posts for Computer Programmers	17	36
% Posts for Business Professionals	15	13
% Posts for Technical and Sales Representatives	13	1
% Posts for Accountants	5	1
% Posts for Computer Professionals	5	22
% Posts for Computer Systems Designers and Analysts	2	6
% Posts for Research and Development Managers	1	3
% Requiring Undergraduate Education	81	12
% Requiring Postgraduate Education	15	85
Mean Required Work Experience	1.87 years	2.1 years
Median Required Work Experience	2 years	2.5 years
90th Percentile Work Experience	2.5 years	2.5
10th Percentile Work Experience	1 year	1
Mean Salary	195,000 rupees	351,000 rupees
Median Salary	137,500 rupees	300,000 rupees
90th Percentile Salary	425,000 rupees	500,000 rupees
10th Percentile Salary	37,500 rupees	105,000 rupees

Notes: Descriptive statistics overall and within AI jobs only.

Table B.2: Growth descriptives of vacancy data in restricted sample, all jobs and AI jobs only

Characteristic	All Jobs	AI Jobs
Mean Growth in # Posts	22	.65
90th Percentile Growth in # Posts	50	0
Mean Growth in AI Share	.17%	-
Growth in Mean Salary	17,000 rupees	23,000 rupees
Growth in Median Salary	16,000 rupees	18,000 rupees
Growth in 10th Percentile Salary	6,000 rupees	-16,000 rupees
Growth in 90th Percentile Salary	27,000 rupees	77,000 rupees
Growth in Mean Required Experience	.15 years	.15 years
Growth in Median Required Experience	.18 years	.18 years
Median Growth in Postgraduate Share	0	0
90th Percentile Growth in Postgraduate Share	22%	71%

Notes: Descriptive statistics overall and within AI jobs only.

Further descriptives on AI demand and exposure: Figure B.5 displays the global AI vibrancy index from Perrault et al. (2019), including India's performance relative to other countries. The top 20 roles demanding AI skills in our analysis are listed in Figure B.6. Figure B.7 shows the concentration of AI posts in firms. Figure B.8 shows differences by industry in the diffusion of AI. Figure B.9 displays exposure to AI by wage offer.

Table B.3: Wages in AI vs. non-AI roles

	log Annual Salary			
	(1)	(2)	(3)	(4)
AI post	0.318***	0.198***	0.128***	0.174***
	(0.0484)	(0.0358)	(0.0220)	(0.0422)
Experience Required (Years)	0.470***	0.411***	0.386***	0.351***
	(0.00693)	(0.00797)	(0.00787)	(0.00818)
High School	0.00481	-0.0644***	-0.0408**	-0.0395**
	(0.0788)	(0.0192)	(0.0185)	(0.0172)
Master's	0.104***	0.0774***	0.0448***	0.0198**
	(0.0144)	(0.00990)	(0.00781)	(0.00814)
Doctorate	0.131**	0.0741*	0.0132	0.00218
	(0.0588)	(0.0417)	(0.0325)	(0.0339)
Fixed Effects:				
– Industry-Region	✓	✓	✓	✓
– Industry-Year	✓	✓	✓	✓
– Region-Year	✓	✓	✓	✓
– Firm		✓	✓	✓
– Occupation Code			✓	
– Role Label				✓
R^2	0.343	0.535	0.556	0.577
Observations	14012499	13976759	13275348	13976757

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. All regressions include industry-region, industry-time and region-time fixed effects, and models (2)-(4) also include firm fixed effects. *AI post* is a dummy such that the coefficient is the percentage increase in annual salary associated with posts requiring AI skills, after accounting for the control variables and fixed effects. Similarly, *Experience* is measured in years, so the coefficient reflects the percentage salary increase associated with an additional year of experience. The education variables are dummies, with the baseline category being a Bachelor's degree; for instance, *High School* reflects the percentage salary decrease associated with posts that only require a high school education. The *Occupation Code* fixed effect also accounts for variation across India's 4-digit National Classification of Occupations codes, while the more granular *Role Label* fixed effect accounts for variation across the self-selected role classifications built into the jobs portal.

Figure B.1: Share of job posts by broad occupation group over time.

(a) All Job Posts

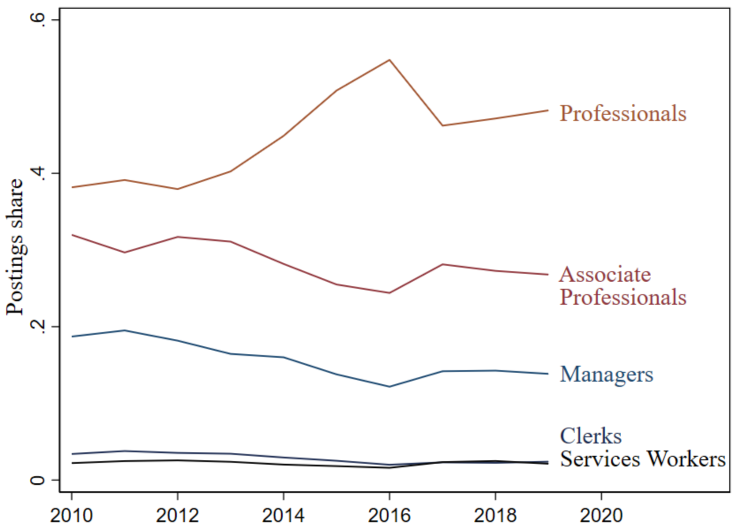

(b) Only AI Job Posts

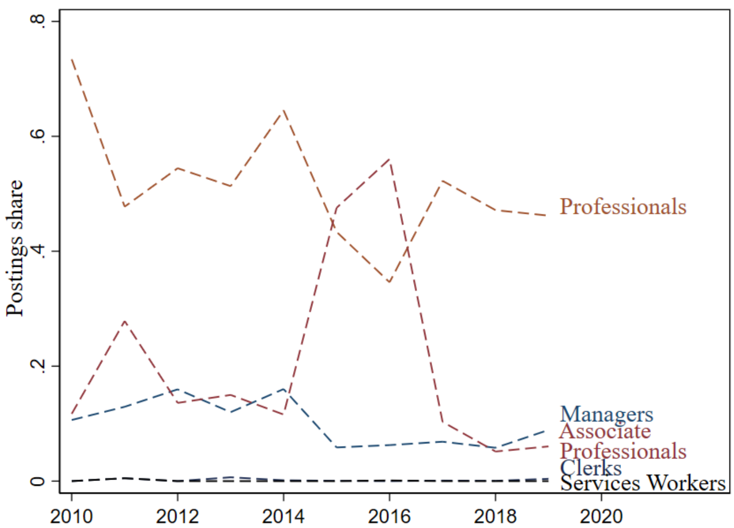

Notes: Share of 1 digit NCO04 occupations overall and within AI jobs only.

Figure B.2: Total posts by district, 2010-2019

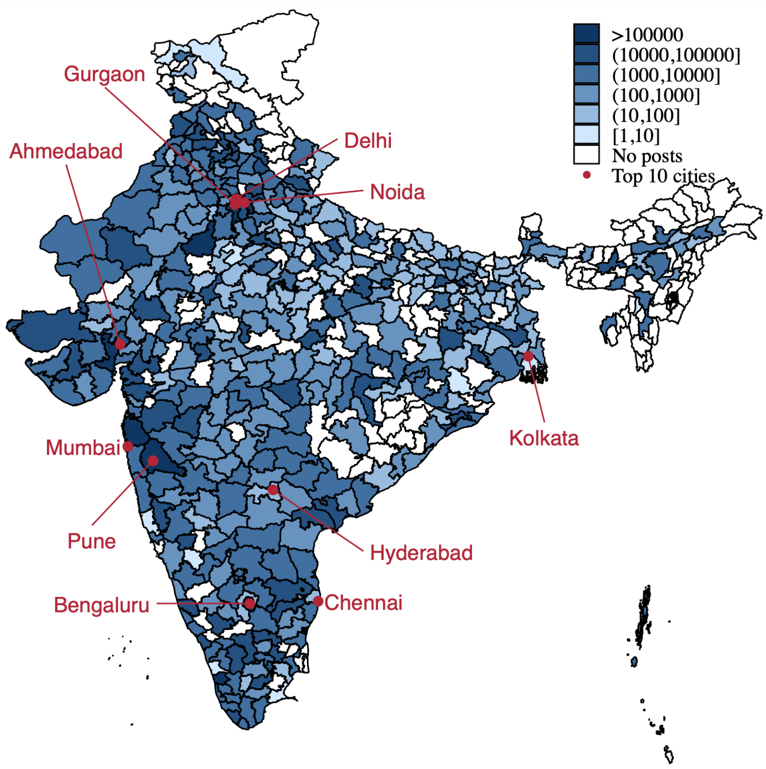

Notes: This map shows the distribution of our online vacancy posts across Indian districts for the entire period 2010-2019. Labels are shown for the ten cities with the largest numbers of posts.

Figure B.3: Share of all AI posts by district, 2010-2019

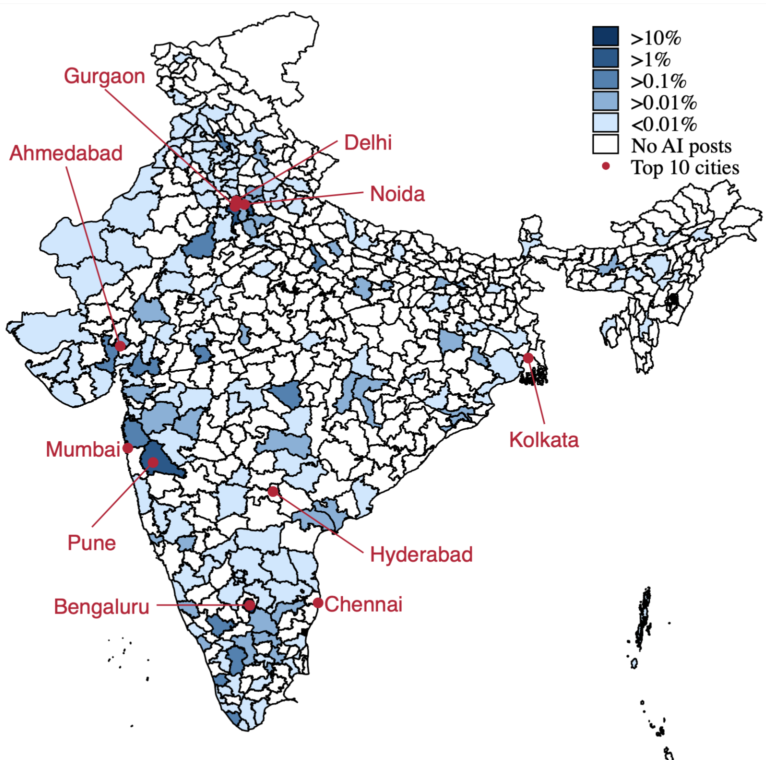

Notes: The map shows the distribution of the share of all AI posts by particular districts, for the entire period 2010 to 2019. Labels are shown for the top ten cities with the most AI posts. The majority of districts have few AI posts, since hiring is clustered in the largest cities.

Figure B.4: Cities' shares of AI posts over time

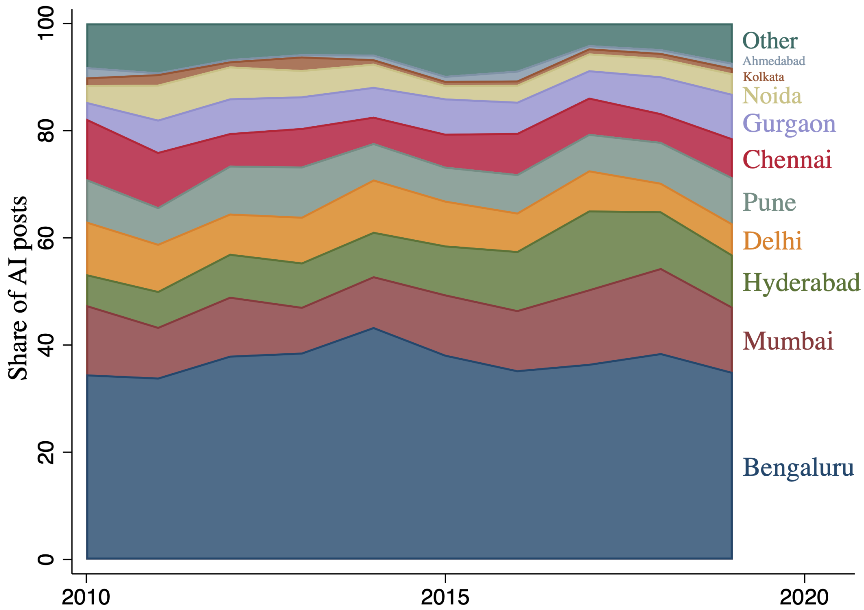

Notes: This graph shows the distribution of AI posts across cities over time. Each year reflects the share of all AI vacancies in that year which were in each city. Shares have been remarkably constant. Bangalore's share peaked at just over 40% in 2014, then Mumbai's share in particular has risen subsequently as AI demand increased in finance (see Figure 3.1).

Figure B.5: Global AI Vibrancy, from Perrault et al. (2019)

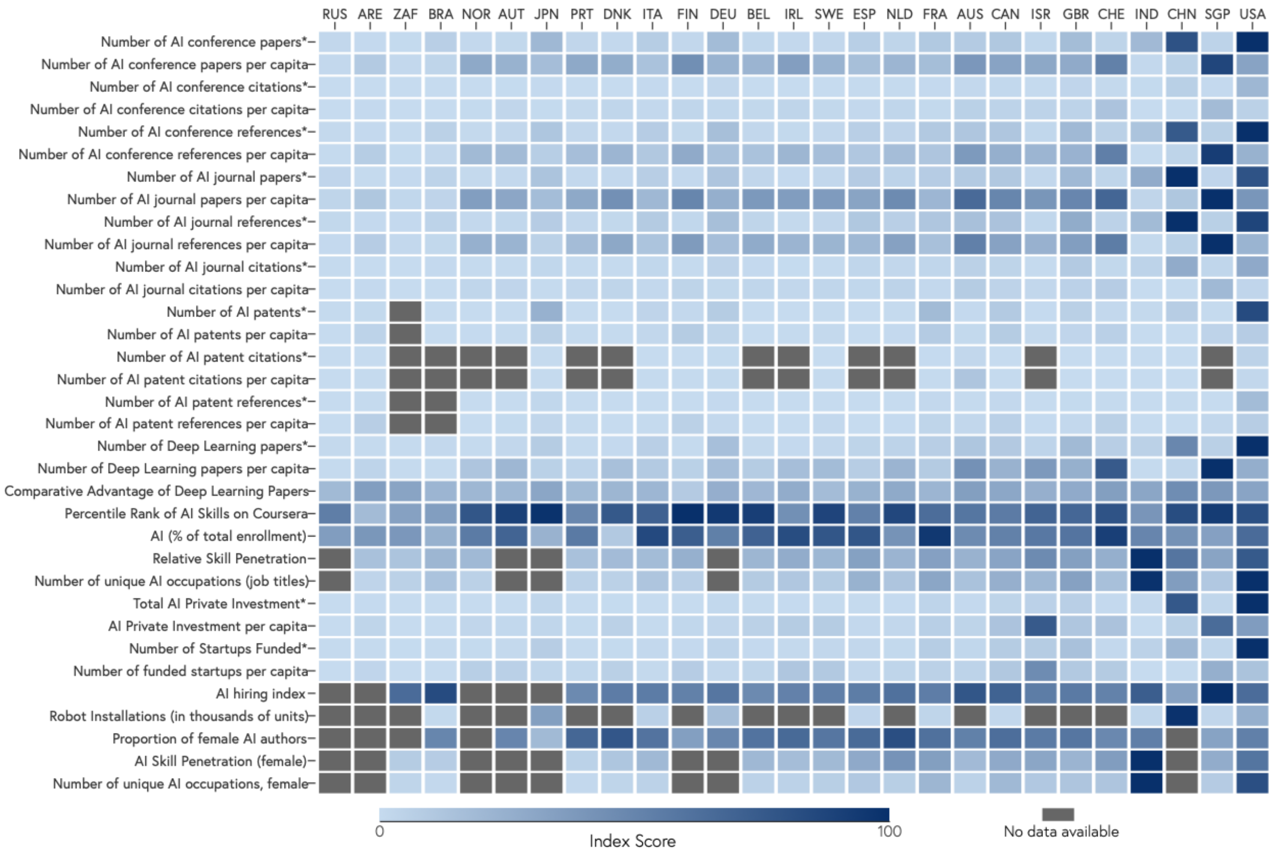

Notes: This chart shows relative country scores on a wide range of AI progress metrics. India (fourth from right) scores highly only on skill penetration (the average share of AI skills among all the top 50 skills in each occupation, across all occupations in the country) and number of unique AI occupations (those that have any AI skills in their top 50 skills). These are both calculated using LinkedIn data, which is far less representative in India than in developed countries due to low coverage. Skill penetration is thus likely an overestimate, while the number of AI occupations is largely driven by India's population size.

Figure B.6: Top 20 roles demanding AI skills, 2010-2019

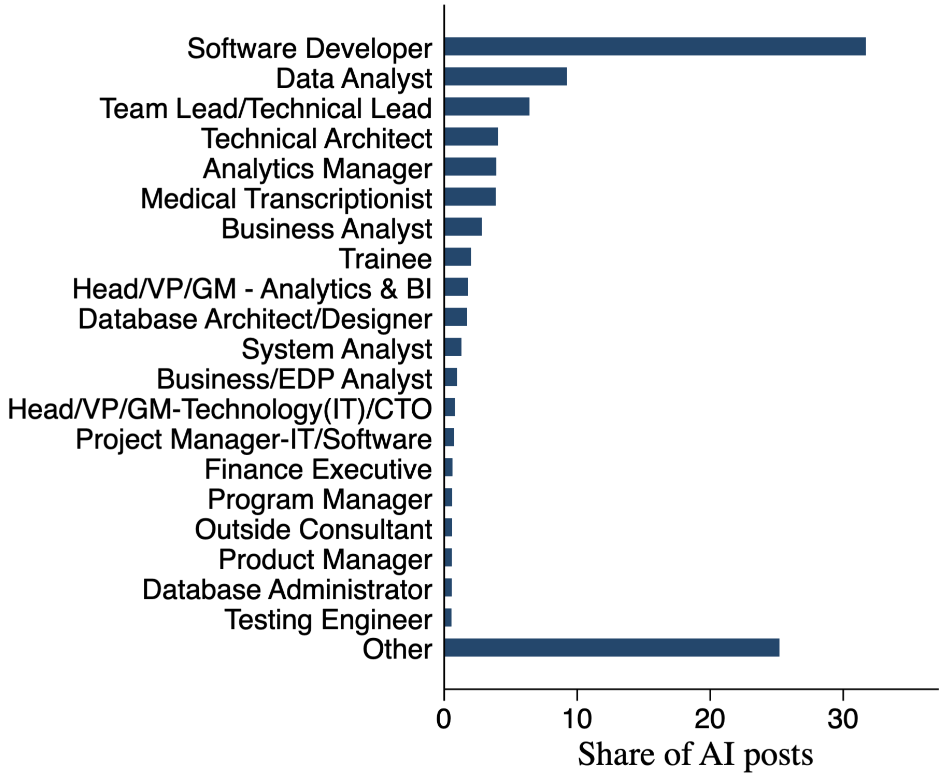

Notes: We rank the top roles demanding AI skills by their share of AI posts. All other roles hiring AI skills are grouped in the 'Other' category.

Figure B.7: Firm concentration of AI posts, 2010-2019

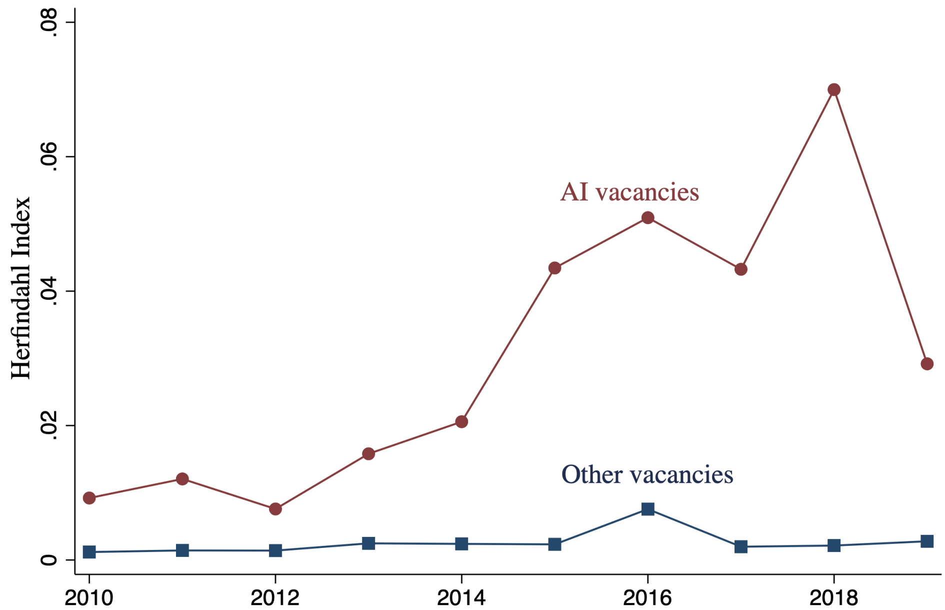

Notes: We plot the trend in the Herfindahl Index for AI and non-AI vacancies over time. These are calculated for each year as the sum of squared firm market shares of all AI or non-AI posts, respectively.

Figure B.8: Heterogeneity in AI diffusion across industries

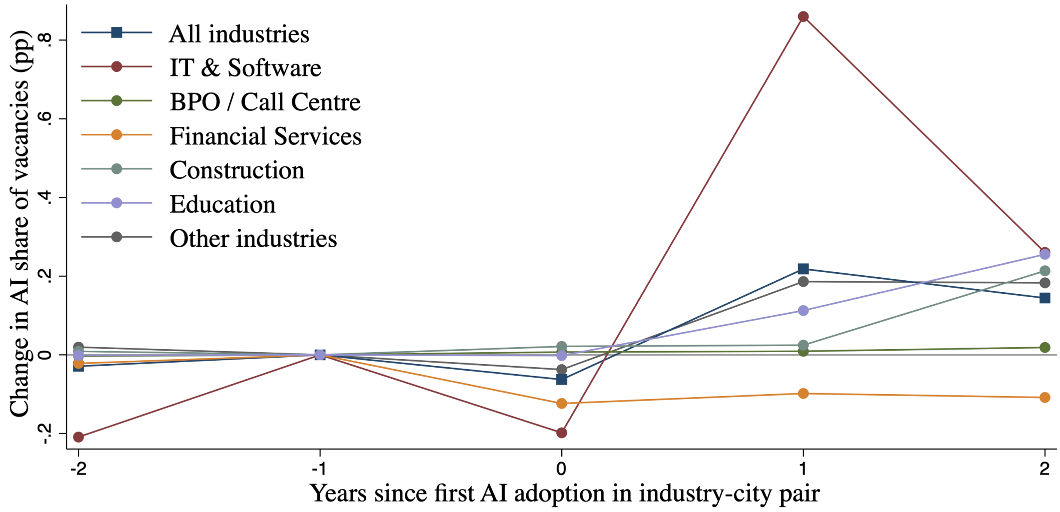

Notes: We plot the change in the share of vacancies that are AI vacancies, for years before and after the first adoption of AI in an industry-city pair, for each of the top five industries by AI adoption. AI posts by the initial AI adopter are excluded in order to focus on diffusion of AI posting to other firms in the industry-city pair.

Figure B.9: AI exposure by occupation wage offers

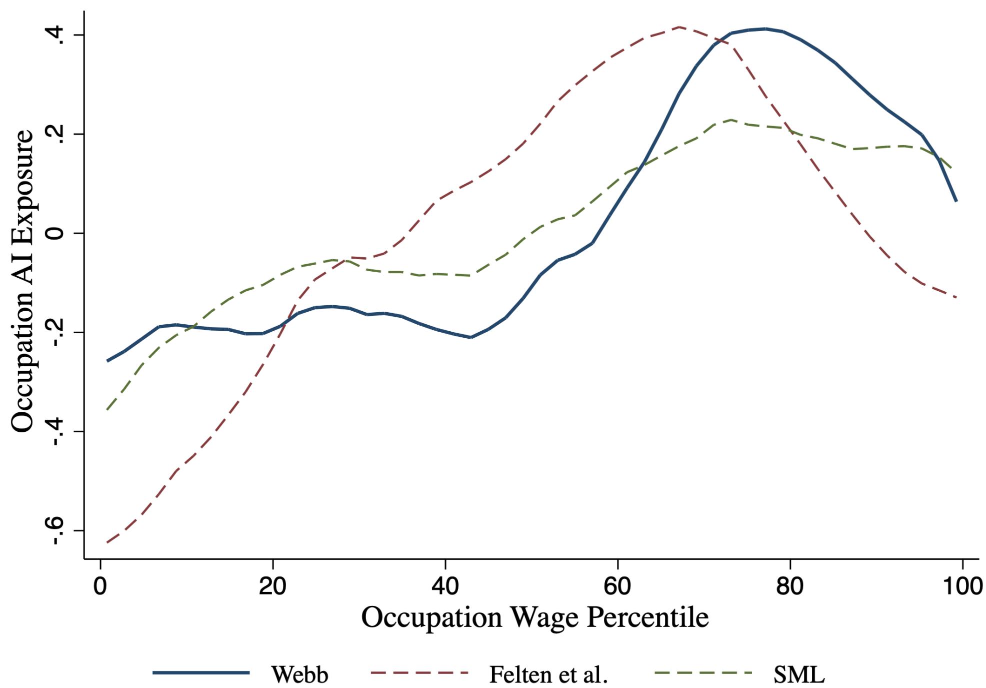

Notes: This graph shows a smoothed local polynomial regression of the Webb AI exposure measure on occupational wage offers. We first rank occupations by their average salary across all vacancy posts 2010-2019. We then plot the AI exposure associated with each, smoothing across a bandwidth 10 of percentage points. In addition to our main measure, from Webb (2020), we also show analogous results for the alternative measures (Felten et al. 2018, Mani et al. 2020) which we use in robustness checks in Appendix C.

B.2 Medium-term impacts

This appendix includes further results for the medium-term impacts section. In Table B.4, we show the first stage. Figure B.10 shows the relationship between AI exposure and the non-AI share of establishment's posts. Figure B.11 extends the wage distribution graph to all postings. Table B.5 studies the impact of AI on 2-digit occupations. Table B.6 shows wage growth results by 1-digit occupations. Table B.7 additionally controls for job profiles in the wage growth regressions and Table B.8 shows the impact of AI on non-AI education and experience. Table B.9 shows task growth results for abstract and routine tasks following Autor & Dorn (2013) and confirms the findings of Table 4.3. Similarly, Table B.10 repeats Table 4.4 for abstract and routine tasks following Autor & Dorn (2013). Figures B.12 and B.13 restrict the regressions of AI demand growth on verb demand growth to the 1% highest paid jobs in order to study the top end of the wage distribution, for which we find within occupation wage growth reductions when controlling for the shifting occupational distribution.

Table B.4: First stage: Impact of AI exposure on establishment AI adoption

	Growth in AI Vacancies	
	(1)	(2)
Establishment AI Exposure	0.0170***	0.0193***
	(0.00331)	(0.00370)
Fixed Effects:		
– Region	✓	✓
– Firm Decile	✓	✓
– Industry		✓
R^2	.0341	.049
Observations	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The dependent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. The independent variable is establishment AI exposure, calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Each coefficient therefore represents the proportional impact on AI hiring of a one-standard deviation rise in AI exposure.

Table B.5: Second stage: Impact of AI on establishment non-AI postings, by detailed occupation group

	Growth in Non-AI Vacancies					
	Professionals				Managers	
	Engineering Professionals	Health Professionals	Teaching Professionals	Other Professionals	Corporate Managers	General Managers
Growth in AI Vacancies	-4.951***	0.548*	0.284***	-2.687***	-12.18***	-2.403***
	(1.198)	(0.332)	(0.107)	(0.926)	(2.592)	(0.827)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓	✓
First Stage F-Stat	27.17	27.17	27.17	27.17	27.17	27.17
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes.
Occupation groups are 2-digit occupations within Professionals and Managers from the NCO04.

Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table B.6: Second stage: Impact of AI adoption on establishment non-AI wages, by occupation group

	Growth in Non-AI Median Wage				
	(1)	(2)	(3)	(4)	(5)
Growth in AI Vacancies	0.297	0.474*	-0.393	-0.460	-0.684*
	(0.858)	(0.277)	(0.348)	(0.413)	(0.415)
Fixed Effects:					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	5.22	12.34	30.04	14.21	16.7
Observations	981	2,059	13,128	9,296	8,003

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Occupation groups are 1-digit occupation groups from the NCO04. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table B.7: Second stage: Impact of AI adoption on establishment non-AI wages, controlling for job profiles

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-2.132***	-1.933***	-2.088***	-1.891***
	(0.674)	(0.594)	(0.652)	(0.575)
Growth in Experience	0.836***	0.824***	0.836***	0.823***
	(0.0299)	(0.0284)	(0.0297)	(0.0282)
Growth in High School share	-0.0662	-0.0830	-0.0692	-0.0860
	(0.0903)	(0.0848)	(0.0883)	(0.0830)
Growth in Master's share	0.254***	0.257***	0.252***	0.255***
	(0.0355)	(0.0356)	(0.0355)	(0.0356)
Growth in Doctorate share	2.669**	2.385**	2.624**	2.345**
	(1.282)	(1.116)	(1.253)	(1.090)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	26.1	26.84	27.31	28.16
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table B.8: Impact of AI adoption on establishment non-AI education and experience

	Growth in Non-AI Postgraduate Vacancy Share		Growth in Non-AI Years of Experience	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-0.225	-0.319	-1.065***	-0.691**
	(0.224)	(0.206)	(0.379)	(0.302)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat				
Observations	25.11994	25.87232	25.11994	25.87232
N	22,244	22,244	22,244	22,244

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). *Non-AI Postgraduate Vacancy Share* is defined as the establishment-level share of non-AI posts requiring either a Master's or a Doctorate.

Table B.9: Second stage: Impact of AI adoption on establishment abstract and routine tasks following Autor & Dorn (2013)

	Growth in Abstract Tasks			Growth in Routine Tasks		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-1.124***	-1.530***	-1.140***	-0.203	-0.0587	-0.138
	(0.329)	(0.409)	(0.326)	(0.189)	(0.234)	(0.187)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	26.06	26.31	27.17	26.06	26.31	27.17
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. We use the occupational task scores for abstract, routine, and manual tasks from Autor & Dorn (2013) (based on data from the Dictionary of Occupational Titles 1977) and map occ1990dd occupations to NCO04 occupations. Scores are standardised as in Acemoglu & Autor (2011*a*). Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table B.10: Second stage: Impact of AI adoption on establishment abstract and routine tasks following Autor & Dorn (2013), by occupation group

	Personal, Sales and Security	Clerks	Associate Professionals	Professionals	Managers
	Growth in Abstract Tasks				
Growth in AI Vacancies	-1.467***	-0.346	1.731***	-0.649***	-5.232***
	(0.327)	(0.213)	(0.335)	(0.218)	(1.133)
Fixed Effects:					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	27.17	27.17	27.17	27.17	27.17
Observations	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. We use the occupational task scores for abstract, routine, and manual tasks from Autor & Dorn (2013) (based on data from the Dictionary of Occupational Titles 1977) and map occ1990dd occupations to NCO04 occupations. Scores are standardised as in Acemoglu & Autor (2011*a*). Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table B.11: AI adoption at the district level

	AI Vacancies	Non-AI Vacancies	Vacancies	Non-AI Wages	Wages
	(1)	(2)	(3)	(4)	(5)
First stage:					
AI Exposure	0.147***				
	(0.0402)				
Second stage:					
Growth in AI Vacancies		0.143	0.144	-0.436*	-0.433*
		(0.337)	(0.337)	(0.230)	(0.230)
First Stage F-Stat		13.34	13.34	13.57	13.57
Observations	399	399	399	399	399

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the district level. The dependent variables are the growth between 2010-12 and 2017-19 in district AI vacancies, non-AI vacancies, total vacancies, non-AI wages, and total wages, each approximated by the change in the inverse hyperbolic sine. The independent variable for the first stage is district AI exposure, calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the district posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). The first two coefficients therefore represent the percentage point impact on AI hiring of a one-standard deviation rise in AI exposure. The independent variable for the second stage is the growth in district AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. The latter four coefficients therefore represent the percentage point impact upon the outcome variable of a one percent increase in district AI hiring, instrumented by district AI exposure.

Table B.12: AI adoption at the district level, administrative employment and wage data

	AI Vacancies	Vacancies		Wages	
	(1)	(2)	(3)	(4)	(5)
First stage:					
AI Exposure	0.147***				
	(0.0402)				
Second stage:					
Growth in AI Vacancies		-0.376	-0.944	-0.234	0.0337
		(0.470)	(0.590)	(0.261)	(0.296)
Fixed Effects:					
– State			✓		✓
First Stage F-Stat		13.34	13.33	13.34	13.33
Observations	399	399	398	399	398

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the district level. The dependent variables are the growth between 2010-12 and 2017-19 in district AI vacancies (from the vacancies data), vacancies (from administrative datasets), and total wages(likewise from administrative datasets), each approximated by the change in the inverse hyperbolic sine. The independent variable for the first stage is district AI exposure, calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the district posts vacancies in 2010-12 (from the vacancies data), weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). The first coefficient therefore represents the percentage point impact on AI hiring of a one-standard deviation rise in AI exposure. The independent variable for the second stage is the growth in district AI vacancies between 2010-12 and 2017-19 (from the vacancies data), approximated by the change in the inverse hyperbolic sine. The latter four coefficients therefore represent the percentage point impact upon the outcome variable of a one percent increase in district AI hiring, instrumented by district AI exposure.

Figure B.10: Impact of AI exposure on the non-AI share of establishments' posts

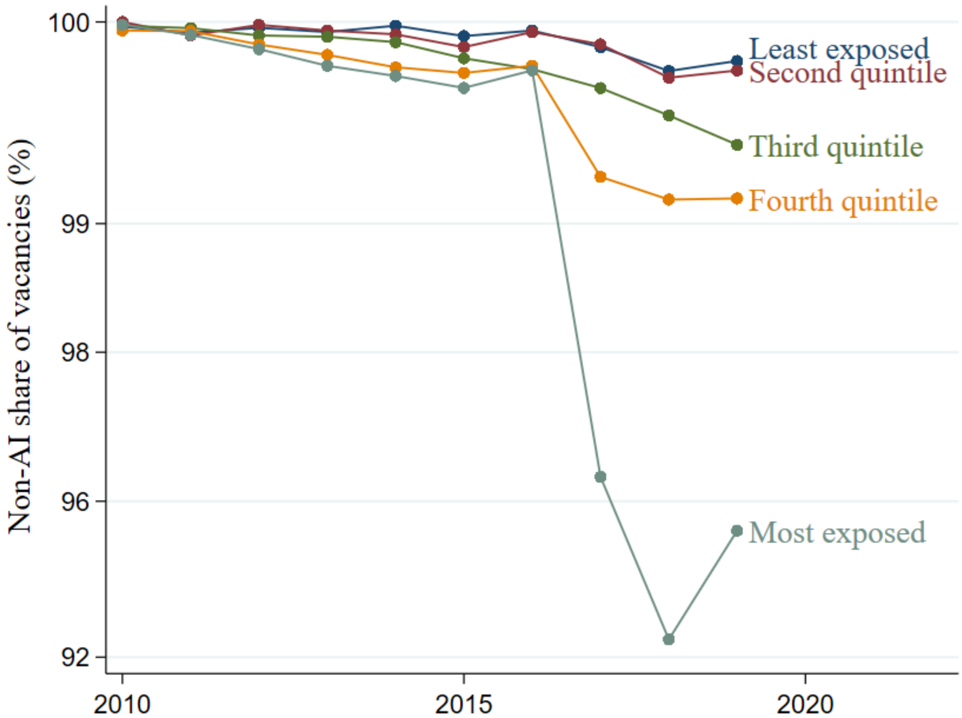

Notes: This graph shows the relationship between AI exposure and the non-AI share of establishments' posts, using an inverse hyperbolic sine scale for the y-axis.

Figure B.11: Impact of AI demand on the wage offer distribution in all posts

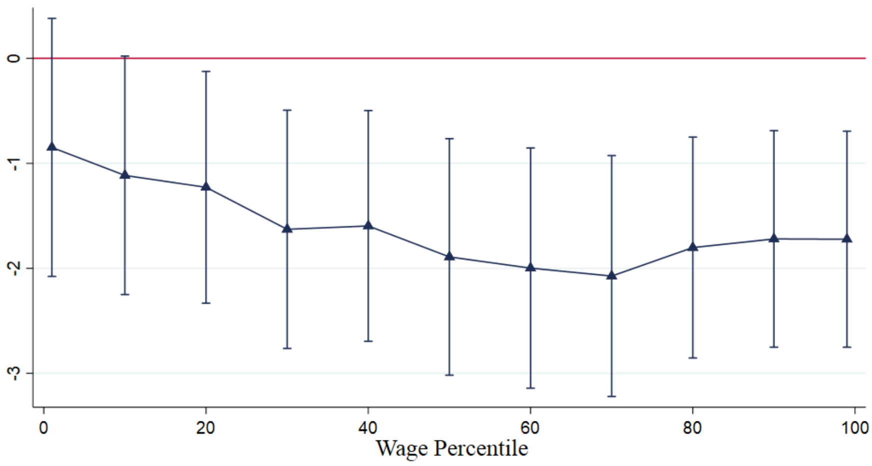

Notes: This coefficient plot shows the impact of increased establishment AI demand on wage growth over time across the distribution of establishment wage offers. Each coefficient is from a regression of type (5) in Appendix Table B.7. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on wage growth over time for a given percentile of the wage offer distribution. We report the 1st and 99th percentile of the wage offer distribution and deciles in between the two extremes, alongside 95% confidence intervals. As in Appendix Table B.7, AI demand is instrumented with AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects. Since AI posts make up only a small share of all roles in most establishments, the pattern is very similar across the distributions for all posts and for non-AI posts only.

Figure B.12: Impact of 1% higher establishment AI hiring growth on verb usage by class, keeping only top 1% highest paid jobs within establishments

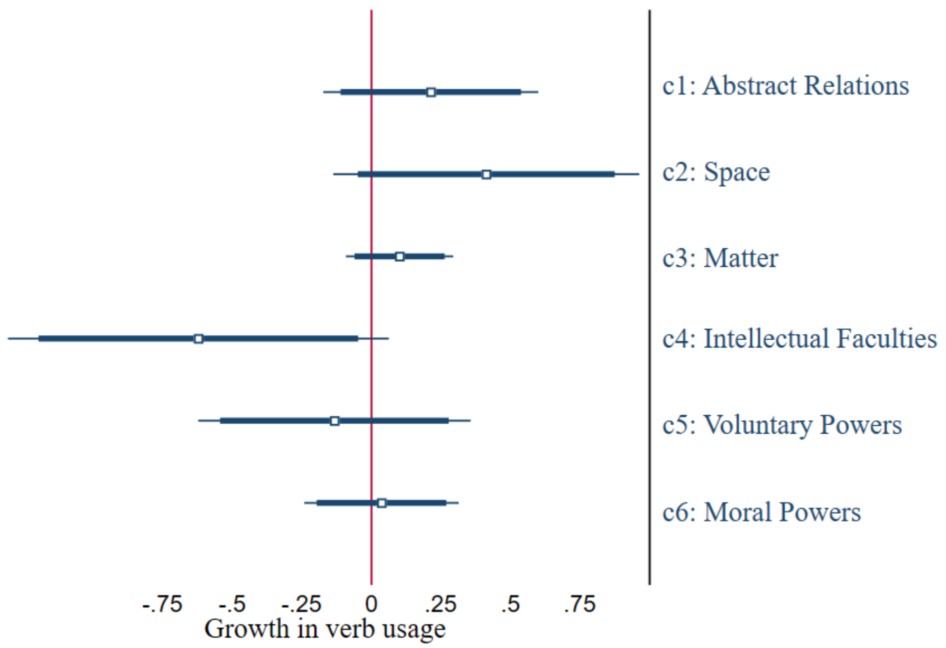

Notes: This coefficient plot shows the impact of increased establishment AI demand on verb share growth between 2010-2012 and 2017-2019, where verb shares are formed from counting verbs in job descriptions of job ads. Point estimates accompanied by 95% and 90% confidence intervals. Each coefficient is from a regression of type (2) in Table 4.1. Here, the outcome variable is growth in the IHS-transformed share of verbs from the respective section or class. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on verb share growth. AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects.

Figure B.13: Impact of 1% higher establishment AI hiring growth on verb usage by section within Intellectual Faculties, keeping only top 1% highest paid jobs within establishments

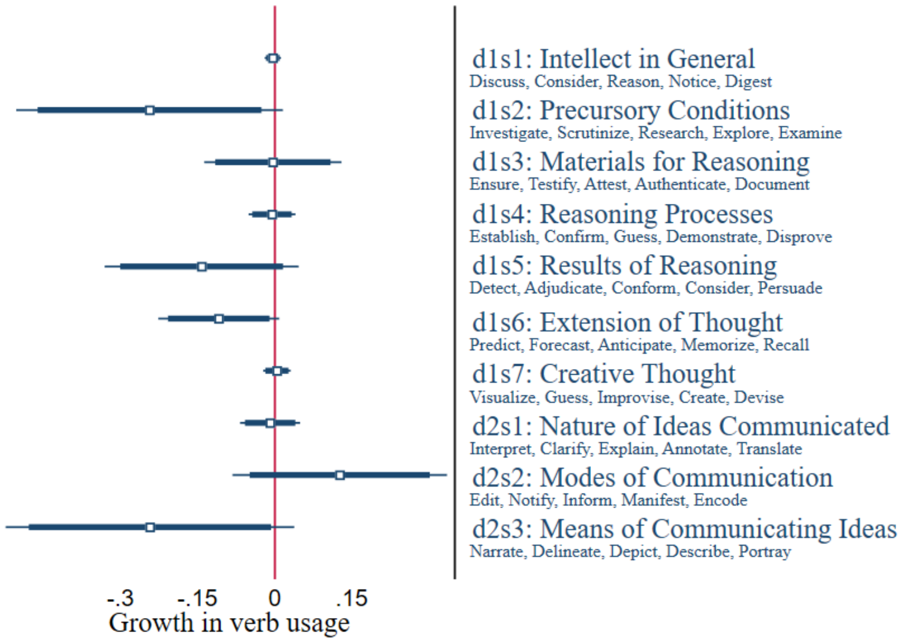

Notes: This coefficient plot shows the impact of increased establishment AI demand on verb share growth between 2010-2012 and 2017-2019, where verb shares are formed from counting verbs in job descriptions of job ads. Point estimates accompanied by 95% and 90% confidence intervals. Each coefficient is from a regression of type (2) in Table 4.1. Here, the outcome variable is growth in the IHS-transformed share of verbs from the respective section or class. In other words, each coefficient represents the percentage point impact of a 1% higher growth in establishment AI demand on verb share growth. AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects.

B.3 Short-term impacts

Table B.13 shows the probit results: we regress lagged establishment characteristics on AI adoption and form predicted values as propensity scores. Adopters are larger and more mature establishments with a higher share of highly-skilled workers hired in the past. Figure B.14 repeats the propensity score matching event study for hiring at the district level.

Table B.13: Probit regression on establishment AI adoption

	AI adoption
Firmsize Decile	-0.0675***
	(0.0121)
Firm age	0.0367***
	(0.00647)
Hiring	0.259***
	(0.0153)
Postgraduate Share	9.895***
	(3.067)
Median Salary	-10.75*
	(5.775)
Median Salary Growth	-0.100***
	(0.0285)
90th Percentile of Salary	0.118**
	(0.0589)
Growth of 90th Percentile of Salary	0.0591
	(0.0394)
99th Percentile of Salary	0.429***
	(0.0497)
Growth of 99th Percentile of Salary	-0.188***
	(0.0330)
Salary Dispersion	-0.00000106***
	(0.000000231)
Median Experience	-0.556***
	(0.0624)
Growth of Median Experience	0.457***
	(0.141)
90th Percentile of Experience	-0.308***
	(0.0999)
Growth of 90th Percentile of Experience	0.0819
	(0.0741)
99th Percentile of Experience	-2.024***
	(0.626)
Growth of 99th Percentile of Experience	-0.159**
	(0.0639)
Experience Dispersion	0.170***
	(0.0493)
N	111044

Notes: Results of a probit regression to compute propensity scores when matching AI adopters to never adopters as described in the short-term results section. All independent variables are lagged by one year. Included but not displayed is a set of year dummies for all years except 2019, and the following interactions: Square of Postgraduate Share, Square and Cube of Median Salary, Square of Median Salary Growth, Square of Growth of 90th Percentile of Salary, Square and Cube of Salary Dispersion, Growth of 99th Percentile of Salary x Salary Dispersion, Median Salary x 99th Percentile of Experience, Growth of Median Experience x 99th Percentile of Experience, Growth of Median Experience x 99th Percentile of Experience x Median Salary. Standard errors clustered at the establishment-level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure B.14: Non-AI hiring following AI adoption at the district level

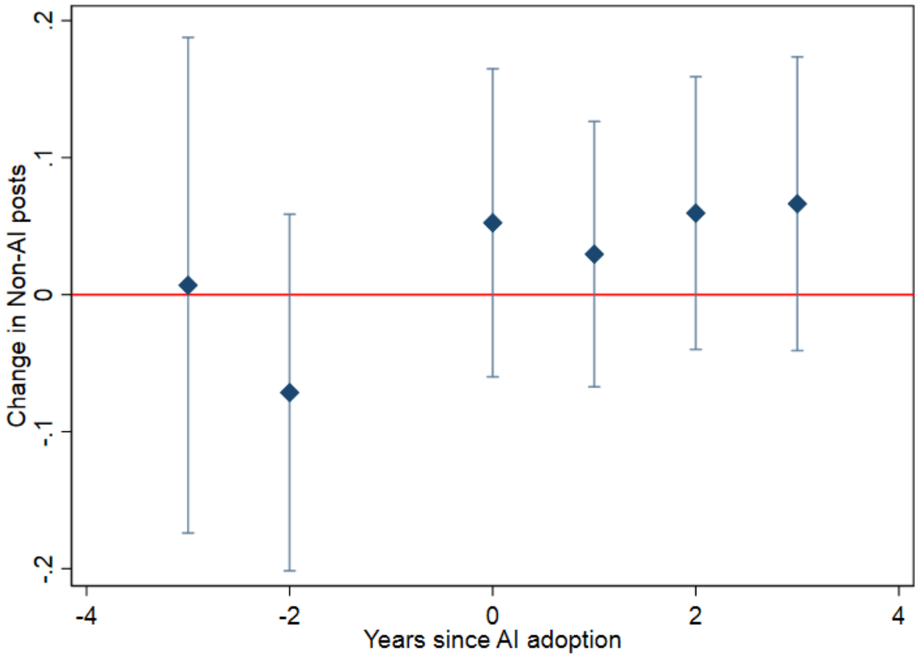

Notes: Two way fixed effects on a balanced panel. The outcome variable is IHS-transformed non-AI hiring by region and district, respectively. At the region level, propensity scores from a probit regression on lagged total hiring, lagged median salary growth, firm age, and year dummies for 2013, 2015, 2016, and 2018. At the district level, propensity scores from a probit regression on lagged median salary, firm age, median salary growth, and year dummies for 2023 and 2013. We use three leads and lags, leaving out the first lead (t-1) as the base period, and cluster standard errors on region and district, respectively. AI adoption leads to reduced non-AI hiring on the region level, but has no effects on the district level.

Appendix C Robustness

In this section, we demonstrate that our results are generally robust to alternative AI exposure measures and various controls and alternative specifications. We also find some similar findings in other administrative datasets, as discussed in the main text.

C.1 Shift-share validity and inference

We construct our instrument from baseline (2010-2012) occupation shares at the establishment level and their respective exposure to AI according to Webb (2020):

$$Exposure_{fr,t_0} = \sum_o PostShare^{t_0}_{fro} \cdot ExposureMeasure_o \tag{C.1}$$

This is a Bartik style instrument with occupation shares in the pre-AI baseline that capture an establishment's exposure to a common shock: occupation-level advances in AI. We can test for the plausible exogeneity of the baseline shares following Goldsmith-Pinkham et al. (2020), who propose several validity checks by analogy with GMM and DiD: investigating correlates of shares, examing pre-trends, and comparing different estimators and running over-identification tests. We find that all three provide support for the validity of our instrument.

Test #1: Investigating correlates of shares. We investigate the extent to which the baseline shares correlate with baseline controls, which could themselves affect hiring and wage offer trends. To this end, we regress the instrument on baseline controls (the structure of required education, experience, and wage offers in an establishment). Table C.1 shows the results, demonstrating that this does not appear to be an issue for the overall instrument. Some individual occupation shares warrant the inclusion of controls, in particular experience, and we thus show robustness to including these controls in our main specification.

Test #2: Examining pre-trends. Most of our results are derived from the long-difference specification discussed above. Therefore, we do not have a pre-period and cannot test for pre-trends. This corresponds to the first empirical example given in Goldsmith-Pinkham et al. (2020), where the shares are fixed in a time period from which we are forming the first difference, such that there is no pre-period. We can, however, ask whether our instrument, which is based on baseline occupation shares, predicts year-on-year employment or salary growth. We regress annual employment and wage growth from 2014 onwards (so that the first differences do not contain the baseline years, 2010-2012, from whose occupation shares the instrument

Table C.1: Investigating correlates of shares

VARIABLES	(1) Instrument	(2) Instrument
Share of Highschool Education	-0.166	-0.166
	(0.204)	(0.204)
Share of Undergraduate Education	-0.232	-0.232
	(0.204)	(0.204)
Share of Postgraduate Education	-0.221	-0.221
	(0.204)	(0.204)
Mean Salary	4.86e-09	4.86e-09
	(4.34e-09)	(4.34e-09)
Mean Experience	-0.00217	-0.00217
	(0.00355)	(0.00355)
Constant	0.635***	0.635***
	(0.204)	(0.204)
Observations	22,201	22,201

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Baseline controls (education, experience, salary) do not correlate significantly with the overall instrument.

is constructed) on the instrument. The results are shown in Table C.2: we do not find any indication of pre-trends: baseline exposure to AI does not predict differential growth rates. This remains the case when including the set of fixed effects included in our main regressions.

Table C.2: Examining pre-trends for the instrument

	(1) Growth in Non-AI Vacancies	(2) Growth in Total Vacancies	(3) Growth in Non-AI Median Wage	(4) Growth in Overall Median Wage
Instrument	-0.00885	-0.00833	0.0184	0.0185
	(0.0130)	(0.0130)	(0.0298)	(0.0298)
Constant	-0.124***	-0.123***	-0.411***	-0.410***
	(0.0164)	(0.0164)	(0.0344)	(0.0344)
Observations	296,730	296,730	296,730	296,730

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: No pre-trends: instrument constructed from baseline occupation shares does not predict employment or salary growth after the baseline. Baseline shares (2010-2012) pooled over three years, as in all regressions. Outcome variables in year-on-year growth for 2013-2019.

Test #3: Alternative estimators and over-identification tests. We next compare a range of estimators (OLS, a range of IV estimators, a machine learning estimator and a Fuller-like estimator) and run over-identification tests. Following Goldsmith-Pinkham et al. (2020), we compare Bartik to OLS, over-identified TSLS, using each share as a separate instrument, the Modified Bias-corrected TSLS (MBTSLS) estimator, the Limited Information Maximum Likelihood (LIML) estimator, and the HFUL estimator. Similarity in results between HFUL and LIML on the one hand, and MBTSLS and over-identified TSLS on the other hand supports the validity of our instrument. Bartik estimates are similar to LIML estimates when including establishment controls. Results from HFUL and MBTSLS are also similar, further supporting our instrument. The comparison of alternative estimators suggests validity of our instrument as we find estimates to be quite similar.

We then run over-identification tests for the HFUL, LIML, and over-identified TSLS estimators, where the null hypothesis is the validity of the over-identifying restrictions. These tests do not reject the null hypothesis when including controls. For misspecification tests, we test whether Bartik is sensitive to the inclusion of controls. Similarity in estimates would support our instrument, and indeed we find support for our instrument's validity.

Adjusted standard errors. In addition to validity, a further issue with shift-share instruments concerns standard errors that are correlated. Table C.3 addresses this concern by computing standard errors according to the correction developed by Adão et al. (2019) and finds that our results are robust.

Table C.3: Second stage: Impact of AI adoption on establishment non-AI vacancies. Adão, Kolesár, and Morales (2019) standard errors.

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-3.574**	-5.942*	-3.605**	-3.534**	-5.909*	-3.566**
	(1.666)	(3.436)	(1.479)	(1.663)	(3.437)	(1.475)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	26.06	26.31	27.17	26.06	26.31	27.17
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors calculated as in Adão, Kolesár, and Morales (2019) in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

C.2 Alternative exposure measures and alternative specifications

This section provides the key results repeated using the alternative exposure measures, as well as a series of alternative specifications, as discussed in Section 6.

Table C.4: First stage: Impact of AI exposure on establishment AI adoption – alternative exposure measures

	Growth in AI Vacancies			
	(1)	(2)	(3)	(4)
AI Exposure	0.0202***	0.0142***	-0.0151***	-0.0102***
	(0.00342)	(0.00308)	(0.00265)	(0.00276)
Exposure Measure	Felten et al.	Felten et al.	SML	SML
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
R^2	.0349	.0481	.0338	.0476
Observations	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The dependent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. The independent variable is establishment AI exposure, calculated as the standardized average of occupation AI exposure (from either Felten et al. 2018, or Mani et al. 2020 building on Brynjolfsson & Mitchell 2017), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Each coefficient therefore represents the proportional impact on AI hiring of a one-standard deviation rise in AI exposure.

Table C.5: Second stage: Impact of AI adoption on establishment non-AI vacancies – Felten et al. exposure measure

	Growth in Non-AI Vacancies		Growth in Total Vacancies	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	1.076	0.698	1.095	0.714
	(0.746)	(1.089)	(0.744)	(1.087)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
– Firm				
First Stage F-Stat	34.97	21.25	34.97	21.25
Observations	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Felten et al. 2018), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.6: Second stage: Impact of AI adoption on establishment non-AI vacancies, by occupation group – Felten et al. exposure measure

	Growth in Non-AI Vacancies				
	(1)	(2)	(3)	(4)	(5)
Growth in AI Vacancies	9.059***	-1.298*	7.243***	-8.499***	-3.633***
	(2.035)	(0.670)	(1.821)	(2.256)	(1.385)
Fixed Effects:					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	21.25	21.25	21.25	21.25	21.25
Observations	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Occupation groups are 1-digit occupation groups from the NCO04. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Felten et al. 2018), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.7: Second stage: Impact of AI adoption on establishment routine and non-routine tasks – Felten et al. exposure measure

	Growth in Non-Routine Tasks			Growth in Routine Tasks		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-1.526***	-1.338***	-1.529***	0.851**	0.289	0.813**
	(0.541)	(0.404)	(0.557)	(0.335)	(0.213)	(0.337)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	22.17	34.97	21.25	22.17	34.97	21.25
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. We average standardised routine and non-routine O*NET task contents by occupation, and form establishments' routine and non-routine task demand by weighting occupations by their standardised routine and non-routine scores. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Felten et al. 2018), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.8: Second stage: Impact of AI adoption on establishment non-AI wages – Felten et al. exposure measure

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-1.108**	-1.512**	-1.133**	-1.567**
	(0.441)	(0.675)	(0.452)	(0.698)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
– Firm				
First Stage F-Stat	36.02	22.15	35.05	21.22
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Felten et al. 2018), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.9: Second stage: Impact of AI adoption on establishment non-AI vacancies, controlling for baseline share of software engineers

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-4.245***	-7.235***	-4.196***	-4.205***	-7.202***	-4.157***
	(1.410)	(2.100)	(1.371)	(1.409)	(2.100)	(1.369)
Covariates:						
Share of Software Engineers	✓	✓	✓	✓	✓	✓
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	22.24	20.91	22.96	22.24	20.91	22.96
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. As an additional control, the baseline share of vacancies for software engineers in each establishment is included. Software engineers are captured by the NCO04 4-digit occupations 2131 (Computer Systems Designers and Analysts), 2132 (Computer Programmers), and 2139 (Computer Professionals, n.e.c.). Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.10: Second stage: Impact of AI adoption on establishment non-AI wages, controlling for baseline share of software engineers

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-3.080***	-3.645***	-2.994***	-2.999***	-3.544***	-2.912***
	(0.950)	(1.119)	(0.911)	(0.914)	(1.075)	(0.876)
Covariates:						
Share of Software Engineers	✓	✓	✓	✓	✓	✓
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	21.67	20.41	22.36	22.82	21.45	23.54
Observations	22,064	22,064	22,064	22,071	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. As an additional control, the baseline share of vacancies for software engineers in each establishment is included. Software engineers are captured by the NCO04 4-digit occupations 2131 (Computer Systems Designers and Analysts), 2132 (Computer Programmers), and 2139 (Computer Professionals, n.e.c.). Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.11: Second stage: Impact of AI adoption on establishment non-AI vacancies, controlling for baseline share of sales & admin vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-7.877**	-16.97*	-7.881**	-7.819**	-16.91*	-7.825**
	(3.273)	(8.718)	(3.219)	(3.268)	(8.710)	(3.214)
Covariates:						
Share of Sales & Admin	✓	✓	✓	✓	✓	✓
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	8.529	4.447	8.784	8.529	4.447	8.784
Observations	22,251	22,251	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. As an additional control, the baseline share of vacancies in each establishment belonging to the broad occupations of either sales or administration is included. We use the occ1990dd occupation classification (by Autor & Dorn 2013) in defining retail sales and clerical jobs as sales and administrative. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.12: Second stage: Impact of AI adoption on establishment non-AI wages, controlling for baseline share of sales & admin vacancies

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-2.991**	-4.922*	-2.932**	-2.862**	-4.597*	-2.796**
	(1.513)	(2.960)	(1.460)	(1.420)	(2.652)	(1.367)
Covariates:						
Share of Sales & Admin	✓	✓	✓	✓	✓	✓
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	7.875	4	8.153	8.556	4.527	8.872
Observations	22,064	22,064	22,064	22,071	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. As an additional control, the baseline share of vacancies in each establishment belonging to the broad occupations of either sales or administration is included. We use the occ1990dd occupation classification (by Autor & Dorn 2013) in defining retail sales and clerical jobs as sales and administrative. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.13: First stage: Impact of AI exposure on establishment AI adoption dummy

	Adoption of AI	
	(1)	(2)
Establishment AI Exposure	0.00965***	0.0106***
	(0.00149)	(0.00157)
Fixed Effects:		
– Region	✓	✓
– Firm Decile	✓	✓
– Industry		✓
R^2	.0434	.062
Observations	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The dependent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The independent variable is establishment AI exposure, calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Each coefficient therefore represents the impact on the propensity to adopt AI of a one-standard deviation rise in AI exposure.

Table C.14: Second stage: Impact of AI adoption dummy on establishment non-AI vacancies

	Growth in Non-AI Vacancies		Growth in Total Vacancies	
	(1)	(2)	(3)	(4)
Adoption of AI	-12.02***	-7.534***	-11.95***	-7.453***
	(2.845)	(2.095)	(2.851)	(2.097)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	41.83	45.62	41.83	45.62
Observations	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of the adoption of AI between baseline and endline. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.15: Second stage: Impact of AI adoption dummy on establishment non-AI wages

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Adoption of AI	-6.351***	-5.514***	-6.089***	-5.273***
	(1.630)	(1.423)	(1.566)	(1.366)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	42.91	46.86	43.12	47.02
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of the adoption of AI between baseline and endline. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.16: Second stage: Impact of AI adoption dummy on establishment non-AI wages, controlling for job profiles

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Adoption of AI	-4.398***	-4.124***	-4.244***	-3.968***
	(1.272)	(1.150)	(1.223)	(1.105)
Growth in Experience	0.829***	0.819***	0.829***	0.818***
	(0.0286)	(0.0274)	(0.0284)	(0.0272)
Growth in High School share	-0.0916	-0.106	-0.0953	-0.110
	(0.0804)	(0.0779)	(0.0782)	(0.0758)
Growth in Master's share	0.251***	0.254***	0.250***	0.252***
	(0.0357)	(0.0356)	(0.0356)	(0.0355)
Growth in Doctorate share	2.562**	2.305**	2.489**	2.235**
	(1.185)	(1.060)	(1.147)	(1.025)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	43.98	47.61	44.13	47.72
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of the adoption of AI between baseline and endline. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.17: Second stage: Impact of AI adoption dummy on establishment non-AI vacancies, logs

	Growth in Non-AI Vacancies		Growth in Total Vacancies	
	(1)	(2)	(3)	(4)
Adoption of AI	-12.90***	-8.064***	-12.47***	-7.840***
	(3.092)	(2.282)	(2.959)	(2.181)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	41.58	45.43	41.83	45.62
Observations	22,244	22,244	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The dependent variables are the change in the log of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of the adoption of AI between baseline and endline. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.18: Second stage: Impact of AI adoption dummy on establishment non-AI wages, logs

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Adoption of AI	-6.351***	-5.514***	-6.089***	-5.273***
	(1.630)	(1.423)	(1.566)	(1.366)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	42.91	46.86	43.12	47.02
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is a dummy, which equals 1 for establishments that did not post AI vacancies in the baseline period, but posted AI vacancies in the endline period. The dependent variables are the change in the log of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of the adoption of AI between baseline and endline. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.19: Second stage: Impact of AI adoption on establishment non-AI vacancies, 2013-15 to 2017-19

	Growth in Non-AI Vacancies		Growth in Total Vacancies	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-5.708***	-3.741**	-5.696***	-3.722**
	(2.065)	(1.627)	(2.072)	(1.632)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat				
Observations	24.882	23.11134	24.882	23.11134
N	38,490	38,490	38,490	38,490

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2013-15 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2013-15, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.20: Second stage: Impact of AI adoption on establishment non-AI wages, 2013-15 to 2017-19

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-3.717***	-3.590***	-3.921***	-3.791***
	(1.172)	(1.115)	(1.269)	(1.210)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat				
Observations	28.3624	26.17782	24.97242	23.13824
N	38,249	38,249	38,281	38,281

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2013-15 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2013-15, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.21: Second stage: Impact of AI adoption on establishment non-AI wages, controlling for job profiles, 2013-15 to 2017-19

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-3.548***	-3.448***	-3.738***	-3.638***
	(1.135)	(1.062)	(1.224)	(1.150)
Growth in Experience	0.753***	0.752***	0.757***	0.755***
	(0.0358)	(0.0337)	(0.0363)	(0.0342)
Growth in High School share	0.0171	-0.00221	0.0224	0.00248
	(0.105)	(0.0994)	(0.108)	(0.102)
Growth in Master's share	0.191***	0.192***	0.194***	0.196***
	(0.0373)	(0.0376)	(0.0384)	(0.0387)
Growth in Doctorate share	1.447***	1.435***	1.545***	1.521***
	(0.496)	(0.472)	(0.598)	(0.563)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat	29.25	26.83	25.8	23.73
Observations	38,249	38,249	38,281	38,281

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2013-15 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2013-15, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.22: Employment results for 'incumbents' and 'startups'

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-3.043***	-2.530**	-2.998*	-3.035***	-2.520**	-2.983*
Incumbents	(1.146)	(1.027)	(1.808)	(1.150)	(1.030)	(1.811)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	24.51	24.33	7.454	24.51	24.33	7.454
Observations	17,348	17,348	14,729	17,348	17,348	14,729
Growth in AI Vacancies	-8.088	-17.32	-8.887	-8.053	-17.32	-8.853
Start-ups	(7.710)	(13.90)	(7.827)	(7.741)	(13.96)	(7.858)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	2.637	2.469	2.801	2.637	2.469	2.801
Observations	21,085	21,085	21,085	21,085	21,085	21,085

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2013-15 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. In order to distinguish between start-ups and incumbents, we look at the shorter long difference between 2013-15 and 2017-19. A start-up is an establishment that did not post in the baseline, 2010-12, and only started posting in 2013-15. An incumbent posted vacancies already in the baseline, 2010-12. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.23: Wage results for 'incumbents' and 'startups'

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-1.781***	-1.813***	-4.630**	-1.824***	-1.858***	-4.645**
Incumbents	(0.622)	(0.619)	(1.926)	(0.640)	(0.638)	(1.931)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	25.64	25.58	7.519	24.48	24.35	7.529
Observations	17,259	17,259	14,648	17,266	17,266	14,652
Growth in AI Vacancies	-9.946*	-11.88*	-9.754*	-12.26	-14.77	-11.93
Start-ups	(5.697)	(6.913)	(5.478)	(8.323)	(10.31)	(7.880)
Fixed Effects:						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	4.131	4.12	4.326	2.668	2.558	2.837
Observations	20,934	20,934	20,934	20,959	20,959	20,959

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2013-15 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. In order to distinguish between start-ups and incumbents, we look at the shorter long difference between 2013-15 and 2017-19. A start-up is an establishment that did not post in the baseline, 2010-12, and only started posting in 2013-15. An incumbent posted vacancies already in the baseline, 2010-12. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.24: Impact of AI adoption on establishment non-AI mean wages

	Growth in Non-AI Mean Wage		Growth in Overall Mean Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-2.606***	-1.785***	-2.531***	-1.746***
	(0.726)	(0.544)	(0.698)	(0.526)
Controls for Experience & Education		✓		✓
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry	✓	✓	✓	✓
First Stage F-Stat	26.39	26.93	27.71	28.25
Observations	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Models (2) and (4) also control for changes in establishment job profiles over the period, specifically the mean number of years of experience required and the shares of posts requiring different levels of education.

Table C.25: Second stage: Impact of AI adoption on establishment non-AI vacancies, weighted (top 5% winsorized)

	Growth in Non-AI Vacancies		Growth in Total Vacancies	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-1.523**	-0.968*	-1.500**	-0.941*
	(0.628)	(0.498)	(0.627)	(0.495)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat				
Observations	18.79682	16.22949	18.79682	16.22949
N	22,251	22,251	22,251	22,251

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. Establishments are weighted by baseline number of posts, with the top 5% winsorized. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Table C.26: Second stage: Impact of AI adoption on establishment non-AI wages, weighted (top 5% winsorized)

	Growth in Non-AI Median Wage		Growth in Overall Median Wage	
	(1)	(2)	(3)	(4)
Growth in AI Vacancies	-0.512**	-0.362**	-0.507**	-0.357*
	(0.210)	(0.184)	(0.208)	(0.182)
Fixed Effects:				
– Region	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓
– Industry		✓		✓
First Stage F-Stat				
Observations	18.53311	15.95353	18.78698	16.24418
N	22,064	22,064	22,071	22,071

Notes: Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01. Standard errors clustered at the firm level. Establishments are weighted by baseline number of posts, with the top 5% winsorized. The independent variable is the growth in establishment AI vacancies between 2010-12 and 2017-19, approximated by the change in the inverse hyperbolic sine. Likewise the dependent variables are the change in the inverse hyperbolic sine of the respective establishment-level outcomes. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent increase in establishment AI hiring. The latter is instrumented by establishment AI exposure. This is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022).

Figure C.1: Impact of establishment AI adoption on the wage offer distribution – Felten et al. measure

(a) Non-AI posts only

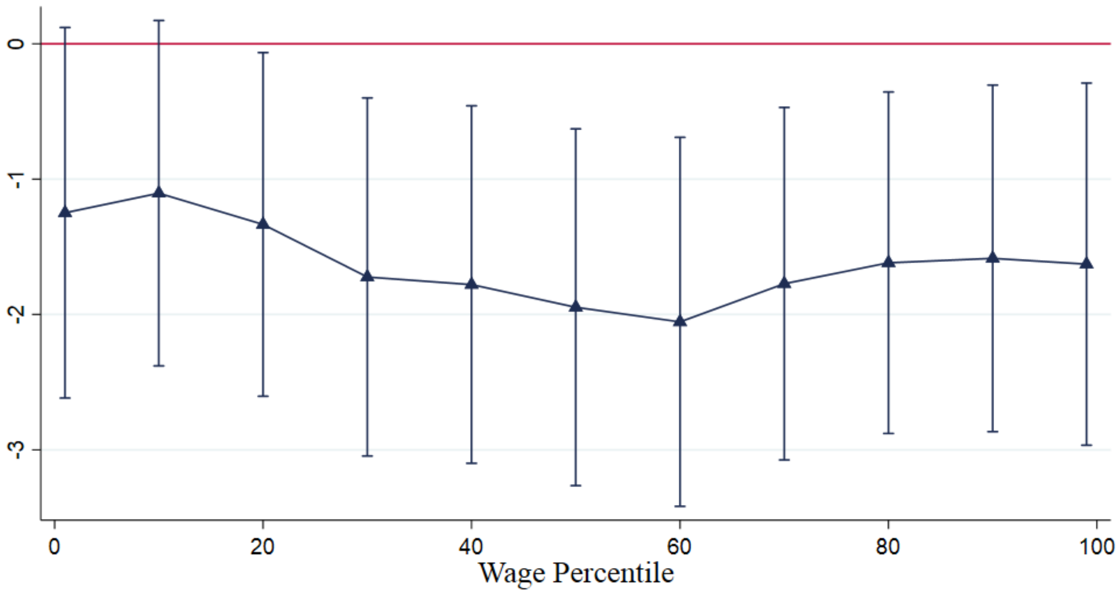

(b) All posts

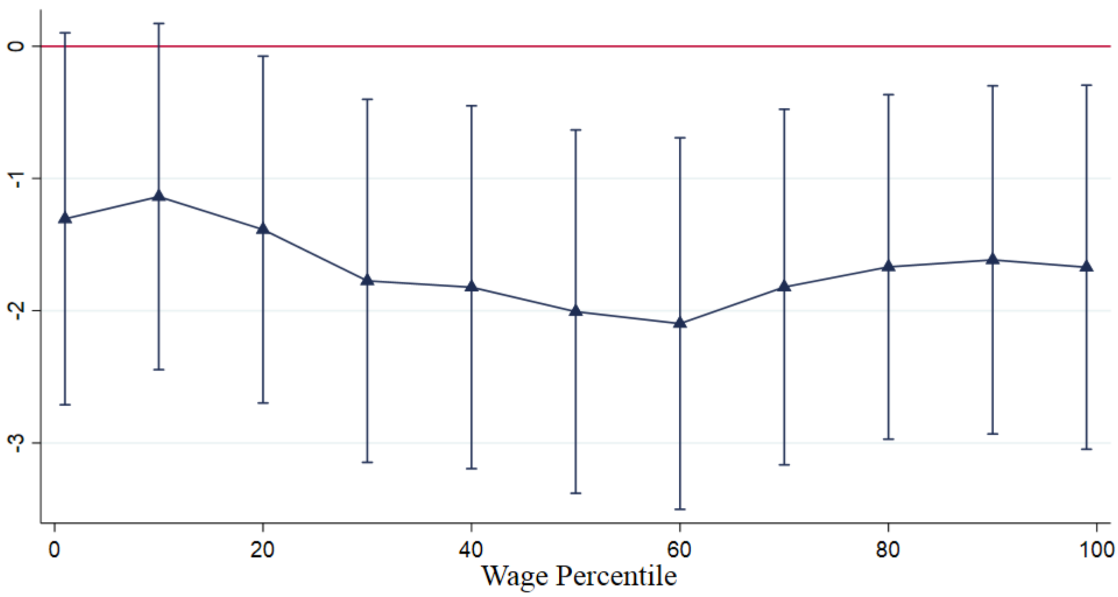

Notes: These coefficient plots show the impact of establishment AI adoption on the distribution of establishment wage offers. Each coefficient in Panel (a) is from a regression of type (2) in Table B.7, and likewise each coefficient in Panel (b) is from a regression of type (5). In other words, each coefficient represents the percentage point impact of a 1% increase in establishment AI demand upon a given percentile of the wage distribution. As in Table B.7, AI demand is instrumented by AI exposure. This is calculated as the standardized average of occupation AI exposure (from Felten et al. 2018), over the occupations for which the establishment posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). Standard errors are clustered at the firm level, and we include region, firm decile and industry fixed effects. Since AI posts make up only a small share of all roles in most establishments, the pattern is very similar across the two distributions.