

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

Alexander Copestake[†] Max Marczinek[‡]
Ashley Pople[§] Katherine Stapleton[¶]

September 2024

[\[Link to latest version\]](#)

Abstract

We document near-exponential growth in the demand for artificial intelligence (AI)-related skills in India's services sector from 2015, 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, using a shift-share design that exploits variation in exposure to new AI inventions. We find that AI adoption significantly reduces growth in non-AI posting volumes and wage offers, particularly for highly skilled managerial and professional occupations, non-routine work, and analytical and communication tasks.

Keywords: Artificial Intelligence, Labor Markets, Wages, Services, Development

JEL Classification Codes: J23, O33

*First version: May 2021. 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. (2018). 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. Declarations of interest: none.

[†]International Monetary Fund, acopestake@imf.org.

[‡]Corresponding author. University of Oxford, max.marczinek@economics.ox.ac.uk. Nuffield College, New Road, OX1 1NF, Oxford, UK.

[§]University of Oxford, ashley.pople@economics.ox.ac.uk.

[¶]World Bank, kstapleton@worldbank.org.

1 Introduction

Rapid advances in artificial intelligence (AI) have spurred an intense debate about its labor market consequences.¹ Online job adverts show that demand for AI-related skills grew almost exponentially and concurrently in several countries around the world from 2015 to 2019 (Figure 1.1).² Yet detailed empirical evidence on the extent of AI deployment and its distributional impacts in developing countries remains scarce. The use cases and impacts of AI need not be the same in developing countries as in advanced economies, and AI could have important implications for development pathways. Consider two examples: machine translation (between languages) and machine transcription (from handwriting to digital text). Improvements in the former could reduce barriers to exporting services and increase employment in developing countries (Brynjolfsson et al. 2019, Baldwin & Forslid 2020), while improvements in the latter could automate an existing back-office process and decrease employment. Understanding the net impact of such technologies on employment opportunities is therefore important for developing countries considering whether to promote a services-led development model.


In this paper, we shed light on the labor market impacts of AI in India – the archetypical pioneer of a services-led development model. We investigate these effects in the urban, white-collar services sector using a new dataset of online job adverts posted from 2010 to 2019 on India’s largest jobs platform, which hosts an estimated 60 percent of the country’s online vacancies. Following a growing literature including Rock (2019) and Acemoglu et al. (2022), we use the demand for AI-related skills, observed in the text of posted job descriptions, as a proxy for AI deployment.

Given the scarcity of evidence on adoption of AI in developing countries, we first provide four stylized facts on the demand for AI-related skills (hereafter ‘AI demand’) in India using the job adverts data. First, we observe a rapid take-off in the rate of AI demand after 2015, rising from 0.37 percent of all job postings that year to 1.03 percent in 2019, with take-up particularly pronounced in the IT, finance and business process outsourcing (BPO) industries. This growth coincides with the take-off in developed countries and reflects an increase in demand for specific

¹To 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. However, this paper focuses on advances in AI before recent progress in large language models and generative AI.

²While online job postings in different countries and from different data providers may not be directly comparable, we see a sharp and simultaneous increase in AI-related postings in all cases.

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



Notes: This graph shows the share of all online vacancies – across all sectors – that specify particular AI skills, with these skills defined as described in Section 2.2. Data for India is that used in this paper; data for all other countries is from Lightcast, which does not cover India.

‘deep learning’ skills. Demand for AI skills is largely driven by technical roles such as software development, but there is also substantial demand for AI skills in generalist and managerial roles. Second, AI roles are heavily concentrated in the largest firms and a few key technology clusters – particularly Bengaluru, Mumbai, Hyderabad, Pune, Chennai and Delhi – similar to patterns observed in the USA by Babina et al. (2024) and McElheran et al. (2024). Third, adverts for AI roles tend to feature more verbs associated with complex, creative, and data-driven tasks. Fourth, 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 percent salary premium, which is similar to the 11 percent estimate found in the USA (Alekseeva et al. 2020).

We then investigate the impact of this AI adoption on demand for non-AI labor at the establishment level. In theory, the impact is ambiguous. Advances in machine learning have been conceptualized as reducing the cost, or improving the quality, of the task of ‘prediction’, which is prevalent in many occupations (Agrawal et al. 2018).³ While this suggests displacement of labor, such improvements could also expand labor demand by reducing costs or increasing

³For example, a back office employee of a multinational bank takes the scrawled handwriting on a mortgage application form as a visual input, then generates a typed name of the applicant as predicted output.

the quality of production and hence raising productivity (Webb 2020, Acemoglu et al. 2022, Acemoglu 2024). In addition, AI could create new tasks, spawn further innovation, or incentivize changes in organisational structure (Brynjolfsson et al. 2017, Cockburn et al. 2018, Klinger et al. 2018, Goldfarb et al. 2023, Agrawal et al. 2024). Finally, AI could reduce barriers to services trade, potentially expanding the set of tasks that are outsourced (Brynjolfsson et al. 2019, Baldwin & Forslid 2020).

To measure the net impact empirically, we use a long difference and shift-share instrumental variables strategy, examining how AI adoption relates to growth in AI and non-AI job postings and wage offers between 2010-12 and 2017-19. AI adoption is measured as the year in which an establishment (defined as a firm-city pair) first posts an AI vacancy. We focus on mature ‘incumbent’ establishments that posted on the platform at both baseline and endline, allowing us to observe their labor demand before and after the mid-decade surge in AI adoption. To focus on AI usage rather than AI production, we exclude AI-producing industries, as in Acemoglu et al. (2022).⁴ We exploit plausibly exogenous variation in firms’ *ex ante* exposure to advances in AI to account for unobservable differences between AI and non-AI adopters. Specifically, we first measure the extent to which the workforce in 2010-12 within establishments performed tasks later automatable by AI, using the occupation-level metric developed by Webb (2020) that assesses overlap between job tasks and patented AI capabilities. We then create a shift-share instrument for AI adoption by combining these occupational AI ‘shocks’ with establishment-level occupation vacancy shares at baseline. This instrument assumes that, within cells of similar establishments formed by our granular fixed effects, those more exposed to AI-automatable tasks are more likely to adopt AI, independent of other characteristics. Our first stage results confirm higher AI adoption in more exposed establishments. We present *ex ante* arguments (as described in Borusyak, Hull & Jaravel 2021) and *ex post* robustness tests (following Goldsmith-Pinkham et al. 2020) in support of our instrument’s validity.

Overall, we find a significant negative impact of AI adoption on establishment-level labor demand. A one percentage point increase in the probability of an establishment adopting AI between 2010-12 and 2017-19 – which corresponds to approximately one standard deviation higher *ex ante* exposure – results in an 8.1 percentage point reduction in the growth of non-AI vacancy postings. The effect on total vacancies (including AI vacancies) is similarly substantial, with a 7.8 percentage point decrease. The small difference between these two numbers indicates that the growth in vacancies constituted by new AI posts is far outweighed by the displacement

⁴In particular, we drop IT and education, which are responsible for the vast majority of AI patents (Klinger et al. 2020).

effect in the much larger set of non-AI vacancies. These effects are also substantial when compared to the median growth in postings of 24.9 percent over the same period.

To better understand the mechanisms driving these effects and identify which types of jobs and tasks are most impacted by AI adoption, we analyze the heterogeneous impacts across occupations, skill requirements, and tasks. These negative impacts primarily reflect a reduction in demand for skilled managerial and professional occupations, non-routine work, and analytical and communication tasks. First, examining the impact on establishment-level vacancy composition by occupation group, we find a substantial reduction in higher-skilled professional and managerial occupations, notably engineering professionals and general and corporate managers. Second, using the classification of Acemoglu & Autor (2011), we find that AI demand lowers demand for occupations that are typically non-routine task intensive. This contrasts with previous recent waves of technological change that primarily lowered demand for routine tasks. Importantly, this negative impact on non-routine task-intensive occupations persists within the affected managerial and professional occupation groupings. We find similar negative impacts on ‘abstract’ task intensity, as defined in Autor & Dorn (2013), both within and between occupation groups. Finally, we leverage the richness of our text data to generate a more granular and flexible measure of task content following Michaels et al. (2018). We count verbs in job descriptions and classify them according to meaning using Roget’s Thesaurus, and 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). Such verbs align closely with the capabilities of new machine learning algorithms, suggesting that AI technologies are reducing the need for human labor in tasks involving analysis, prediction, and complex communication.

How does this displacement affect wage offers for new hires? We estimate that a one percentage point increase in the probability of an establishment adopting AI between 2010-12 and 2017-19 reduces the growth rate of median non-AI wage offers by 5.5 percentage points, and that of all wage offers by 5.3 percentage points. This is driven almost entirely by a composition effect across occupations: when controlling for occupation shares, we do not see statistically significant reductions across the wage offer distribution. Our findings thus suggest that AI adoption reduced the availability of relatively high-wage managerial and professional occupations within incumbent Indian white-collar services firms.

Finally, we investigate wider effects of AI adoption by aggregating across establishments, first

within firms and second within districts.⁵ At the firm level, we again find significant negative impacts of AI adoption on labor demand and wage offers. However, when aggregating instead to the district level, we do not detect a significant overall impact of AI adoption. A relatively small share of establishments adopted AI within our sample period; thus, it is likely that the impacts were not yet large enough to be observable at this broader geographic scale.

Taken together, our results suggest that the demand for AI-related skills has already had important effects on Indian service-sector establishments and firms, altering the distribution of labor demand and wage offers across occupations and tasks. Despite limited district-level effects at present, the wider economic implications of AI adoption are likely to grow as usage of the technology diffuses further throughout the economy.

Related literature: Our primary contribution is to the nascent literature on the impacts of AI on developing countries. To date, this has been primarily theoretical due to the scarcity of data on AI adoption in developing countries. Baldwin (2019) and Baldwin & Forslid (2020) have conjectured that machine learning, along with online platforms and software robots, could benefit developing countries by facilitating the offshoring of white-collar services. In contrast, Korinek & Stiglitz (2021) take an alternative view that developing countries will be negatively affected, because AI devalues their comparative advantage in lower-cost labor and natural resources. To the best of our knowledge, we are the first to use data on AI deployment in a developing country to provide causal evidence on such hypotheses.⁶ Our evidence provides partial support for both perspectives. On one hand, we document rapid increases in labor demand for AI skills in India, particularly in the IT and financial services sectors, consistent with the upskilling of existing offshored white-collar services functions. On the other hand, we find a significant negative impact of AI deployment on non-AI labor demand and wages within establishments and firms, consistent with labor displacement. More broadly, our findings that these negative impacts are concentrated in relatively highly skilled managerial and professional occupations align with concerns about the potential for services-led growth to generate good jobs in a world where advances in AI enable the automation of an increasing share of service sector tasks.

We also contribute to the broader literatures on AI diffusion and its consequences. First, we

⁵Firms are national in our dataset, so this corresponds to collapsing respectively the cross-geography and cross-firm dimensions of our data.

⁶Notably, Brynjolfsson et al. (2019) provide evidence on the causal impact of AI on US-Latin America trade using evidence from AI deployment on eBay, a multinational e-commerce platform. Several other studies provide descriptive evidence on AI exposure, as opposed to observed AI adoption, using data on the cross-sectional distribution of occupations (e.g., Pizzinelli et al. 2023, Gmyrek et al. 2023). Alonso et al. (2022) consider the impact of automation technologies, broadly defined, on developing countries, using data on industrial robots.

offer new insights on the diffusion of AI skills in the important context of India. In contrast to evidence on slow diffusion within countries (e.g., Bloom et al. 2024), as well as theories of slow diffusion to low-income countries (e.g., Benhabib & Spiegel 2005), we find that growth in the demand for AI skills within India has been remarkably similar to, and even faster than, that documented in some advanced economies, highlighting the rapid global diffusion of AI capabilities. At the same time, we also find a high degree of clustering, with our finding that AI hiring is highly concentrated within a small number of firms and tech hubs aligning with patterns of diffusion in the US documented by Babina et al. (2024) and McElheran et al. (2024).

Turning to the impacts of AI, we find both similarities and differences to findings in advanced economies. Although India has a different labor market structure to high-income countries, our findings of negative within-establishment effects of AI on labor demand echo those of Grennan & Michaely (2020) and Acemoglu et al. (2022) in the US. Our findings that these negative effects are driven by higher-skilled occupations and analytical tasks are also consistent with Grennan & Michaely (2020) and Webb (2020), yet contrast with the results of Acemoglu et al. (2022) that AI exposure predicts *increased* demand for skill families relating to engineering, analysis, marketing, finance and IT in the US. Our findings thus highlight the potentially nuanced impacts of AI across contexts, with occupation-level impacts differing in India despite similar establishment-level displacement effects. Our evidence that AI negatively affects non-routine task-intensive occupations also stands in contrast to findings for previous waves of technological change, such as computerization, that lowered demand for routine task-intensive occupations (for instance, Autor et al. 2003, Goos & Manning 2007, Goos et al. 2014).⁷

Finally, we add granularity to the existing literature on the labor demand impacts of AI by leveraging the rich data available on our partner’s platform. Building on Michaels et al. (2018), who study the historic occupational distribution in the US, we use the specific verbs in job descriptions to show how AI disproportionately displaces particular tasks, such as forecasting and prediction. Our dataset also includes detailed information on wage offers – in contrast to commonly used platforms in advanced economies, where wage offer information is missing for the majority of posts (Batra et al. 2023) – allowing us to investigate the implications of AI adoption for wage offers. In doing so, we also contribute to the growing literature that uses online vacancy postings to explore labor market effects (e.g. Deming & Kahn 2018, Adams et al. 2020, Javorcik et al. 2020, Babina et al. 2023, 2024).

The rest of this paper proceeds as follows. Section 2 introduces the data, and Section 3

⁷This finding that the pattern of impact of machine learning is very different to previous waves is also consistent with emerging evidence on the impacts of generative AI (e.g., Brynjolfsson et al. (2023), Noy & Zhang (2023)).

presents detailed descriptives on AI demand in the Indian white-collar services sector. Section 4 lays out our empirical strategy, and Section 5 presents our main findings on the impact of AI adoption in hiring on non-AI labor demand and wage offers. Section 6 provides several robustness checks and extensions, including exploring effects beyond the establishment level. Section 7 concludes. The online appendix provides further detail on the construction of our dataset, as well as additional results and robustness checks.

2 Data

2.1 Vacancy data

Our primary dataset is a random sample of 80 percent of job vacancies posted on India’s largest online job platform between 2010 and 2019.⁸ Firms primarily advertise on the platform and conduct subsequent recruitment and hiring offline. The platform estimates a 60 percent market share of Indian online job vacancies in 2020. We focus on the white-collar services sector, for which the data is most representative of overall job vacancies, dropping posts from the manufacturing and agriculture sectors. Our dataset includes approximately 15.5 million service sector job postings, equating to an average of 1.5 million per year, skewed towards later years. This compares to an estimated 19 million people formally employed in the service sector in India in 2021, according to the Ministry of labor and Employment.⁹

Firms submitting job vacancies on the platform are required to use a standardized 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 feature an open text section for the job description, exclusively in English as they are within the formal services sector. 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 percent 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. We focus on full-time jobs, which make up 96 percent of the sample, and drop the small number of part-time and non-permanent positions.

Our dataset has several advantages relative to administrative datasets. First, and most importantly, we can directly observe demand for AI skills in the text of job descriptions. In contrast, national surveys during the period (specifically, the National Sample Survey in 2011-12

⁸The company requested to remain undisclosed.

⁹<https://static.pib.gov.in/WriteReadData/specificdocs/documents/2022/jan/doc20221104101.pdf>

and the Periodic Labour Force Survey in 2017-18) only record broad occupational categories, which are insufficient to identify the use of AI skills.¹⁰ Similarly, firm-level datasets such as Prowess only include variables commonly reported in periodic financial statements, which are insufficiently granular to identify specific AI-related investment. Second, our vacancy dataset has broader coverage, with more than two million services vacancies in 2018, from 40,000 unique firms, compared to 12,000 workers surveyed in the 2017-18 Periodic Labour Force Survey and 2,000 firms recorded in Prowess (see Online Appendix A.2 for details). Lastly, our data are at a higher frequency than the representative sample surveys, which only took place in 2011-12 and 2017-18 and so can provide only limited insight into employment dynamics in the Indian services sector around the rapid take-off in AI demand from 2015.¹¹

However, the richness of the vacancy data comes with certain shortcomings, notably that online vacancies are not representative of all vacancies and only proxy for firm hiring behaviour. Broadly speaking, our vacancy data best represents urban, white-collar service sector jobs, with a greater representation of the IT & BPO, Finance, Insurance and Real Estate, Professional and Business Services industries relative to administrative datasets. We provide a detailed overview of the coverage and representativeness of the data in Online Appendix A.2, where we benchmark the vacancy data relative to nationally-representative labor surveys and firm-level data from Prowess. Vacancies are disproportionately concentrated in urban centres, but at least one post appears for nearly every district in India.¹² Although the share of formal jobs is lower in India than in advanced economies, we find a ratio of 0.08 of annual average job postings to total formal employment. This is similar to the ratio of 0.09 for annual US job postings from Burning Glass Technologies relative to total US employment from 2010-2018 (Acemoglu et al. 2022).

2.2 Measuring AI demand

Despite the prominence of AI in popular discussion, firm-level data on AI adoption remains scarce (Raj & Seamans 2018). In the absence of more granular data, a growing body of work uses technology-related human capital to proxy for technology adoption.¹³ Human capital is a key input for deploying AI systems and the scarcity of skilled AI professionals is well recognized as

¹⁰For instance, the closest categories to a machine learning engineer in the National Occupational Classification are the broad codes ‘2132 – Computer Programmers’ and ‘3122 – Computer Assistants’.

¹¹Specifically, we receive daily data, which we aggregate to an annual frequency in our regressions.

¹²The full geographic distribution of posts is shown in Online Appendix Figure A.1.

¹³For 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. Trefer & Sun (2022) build a measure of AI deployment for mobile apps.

a primary obstacle to widespread adoption of AI, with top-tier scientists earning extremely high salaries and frequently being bought out of academic positions. A global survey of around 2,000 companies by McKinsey Global Institute (2019) found that the primary method for sourcing AI talent and capabilities was external hiring. Most companies build their AI capabilities in house rather than buying or licensing capabilities from large technology companies. Even when firms subcontract AI services, they typically require some in-house AI-related human capital to oversee and manage the process. We therefore follow this emerging literature in using the demand for AI skills as a proxy for the extent of AI adoption within firms, assuming the two are positively correlated.

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.¹⁴ 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

To shed light on the tasks that accompany AI-related skills, we 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 human perception of it); Intellect (the human mind); Volition (the human will); and Emotion, Religion, and Morality.

¹⁴Specifically, a post is categorized 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, Supervized 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.

3 Descriptives

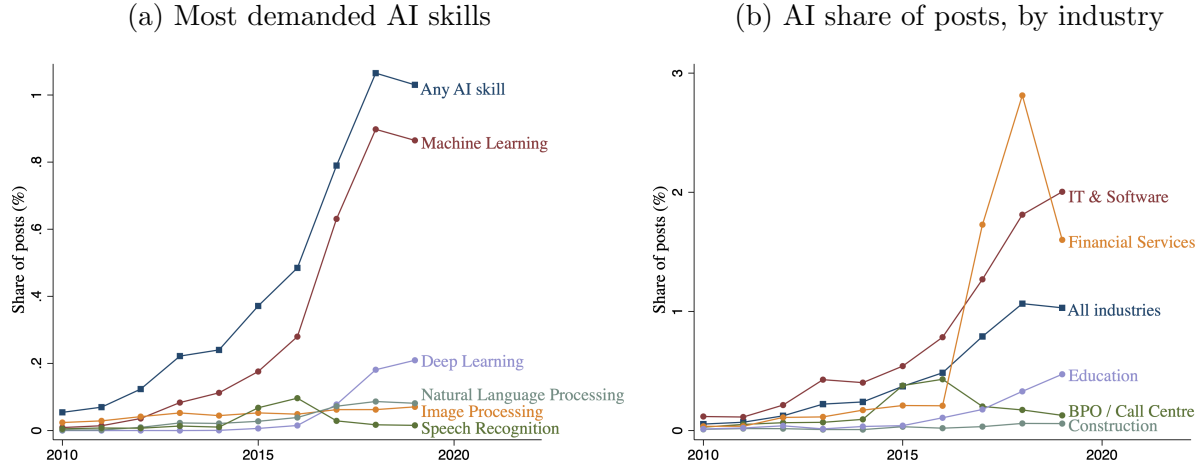
Given the dearth of evidence on the adoption of AI in developing countries, this section presents four descriptive findings from our job postings data, documenting the trends in and composition of posts listing AI skills in the white-collar services sector in India. Here, and throughout the paper, we focus on non-remote, full-time and permanent jobs, which make up 96 percent of the sample.

1. AI demand increased rapidly after 2015, with highest demand for technical data and management roles.

We first observe that the demand for AI skills increased rapidly after 2015, rising from 0.37 percent of all job postings that year to 1.03 percent in 2019 (Figure 3.1). This growth is largely driven 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. AI demand grew steadily in the IT sector since 2011. However, after 2015, demand for AI-related skills also grew rapidly in ‘AI-using’ sectors. For instance, the financial sector saw a ten-fold growth in AI demand between 2016 and 2018, starting from a low base. 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’. The growth in AI demand in India coincides with the take-off in high-income countries, as shown in Figure 1.1.

The most common AI role title by far is ‘Software Developer’, followed by other technical roles such as ‘Data Analyst’, ‘Technical Lead’ and ‘Technical Architect’ (Online Appendix Figure A.8). AI skills are also required in technical management roles, with titles such 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’. Indeed, more than 25 percent of all AI vacancies are fragmented across other roles that each account for less than 1 percent of the total, highlighting that demand for AI skills is spread widely across many occupations.

Figure 3.1: Trends in AI demand



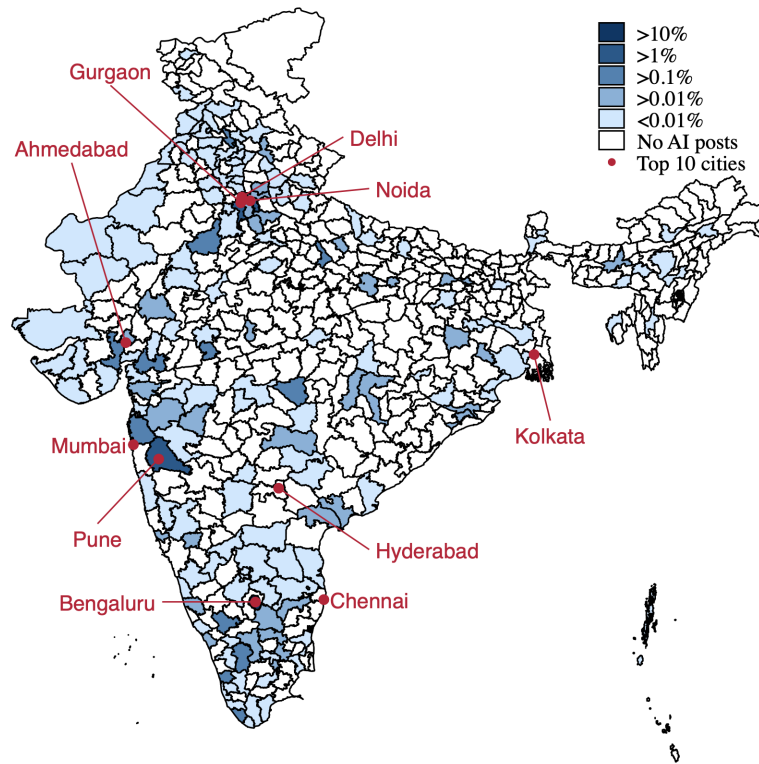
Notes: Panel (a) shows the share of all vacancies that specify particular AI skills, for the top five most demanded skills. Panel (b) shows the share of vacancies that are AI vacancies, both for all industries together and within each of the top five industries by AI share.

2. AI roles are highly concentrated in the largest firms and a few key technology clusters, particularly Bengaluru.

AI roles are highly concentrated in the largest firms. We proxy for firm size by the number of vacancies posted on the platform and compare the cumulative share of AI posts against the corresponding share of all posts. We find that the largest 14 firms are responsible for 10 percent of all vacancies, with each posting at least 50,000 vacancies. These account for 31 percent of all AI posts (Online Appendix Figure A.5). 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 demand is also highly concentrated in large cities, particularly the major technology clusters around Bengaluru, Mumbai, Hyderabad and Delhi. Bengaluru alone has more than 30 percent of all AI vacancies across India (Online Appendix Figure A.6). 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 (Online Appendix Figure A.7).

Figure 3.2: 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.

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

The text of AI job postings contains more verbs associated with complex, creative, and data-driven tasks. Extracting verbs from the job descriptions in AI and non-AI ads, we calculate the share of each verb relative to all verbs, and compare these shares in AI job posts to non-AI job posts. Table 3.1 shows the verbs with the most significant frequency difference between AI and non-AI posts. Compared to non-AI posts, AI posts contain a higher share of complex tasks (e.g. develop and advance), creative tasks (e.g design and build), and data-driven tasks (e.g compute and predict).

Table 3.1: Verbs in AI posts

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


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.

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

AI roles list more advanced educational requirements and offer substantially higher wages than other white-collar services jobs advertized on the platform (Figure 3.3). 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 CPI-deflated median salary of ₹162,000 (approximately US\$3,333, without adjusting for PPP), twice the median non-AI salary of ₹84,000 (US\$1,666). This ‘AI wage premium’ remains high even when controlling for experience, education and firm fixed effects (20 percent), and occupation or role fixed effects (13-17 percent).¹⁵

¹⁵See Online Appendix Table A.4.

Figure 3.3: 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. Salaries are CPI deflated to 2010 levels. 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.

In the rest of the paper, we build on these descriptive findings to assess the causal impact of AI adoption on labor demand. Our long-difference specification exploits the rapid increase in AI demand in the middle of our sample period (descriptive finding #1). Noting the concentration within particular cities and firms (descriptive finding #2), we do so using an establishment-level approach that controls for both city and firm-size fixed effects. Finally, given the specialization of AI roles in particular tasks (descriptive finding #3), and the ‘AI wage premium’ (descriptive finding #4), we examine the implications of AI adoption for the establishment-level composition of occupations and tasks and the wage offer distribution.

4 Empirical Strategy

4.1 Main specification

Our main specification assesses the impact of increased AI demand on changes in the demand for non-AI roles between our baseline period (2010 to 2012) and our endline period (2017 to

2019), spanning the take-off in AI demand from 2015.¹⁶ We take a long-difference approach common in the automation and labor markets literature to focus on structural changes rather than short-term fluctuations in labor demand. Our primary unit of analysis is ‘establishments’, defined as firm-city pairs, as many firms report postings in several different cities. To focus on the effects of AI usage, rather than AI production, we exclude AI-producing industries as in Acemoglu et al. (2022).¹⁷ Our main estimation sample contains almost 25,000 establishments posting approximately two million vacancies on the platform within our baseline and endline 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) could also impact aggregate labor market outcomes.

Our main specification is

$$\Delta y_{fr,t-t_0} = \beta \cdot \text{AdoptsAI}_{fr,t-t_0} + \alpha_r + \alpha_i + \alpha_{f10} + \epsilon_{fr,t-t_0}, \quad (4.1)$$

where: $\Delta y_{fr,t-t_0}$ is the change in the natural logarithm of outcome Y_{frt} between 2010-2012 and 2017-19; $\text{AdoptsAI}_{fr,t-t_0}$ is a dummy taking value one for establishments that do not post an AI role in the baseline but do post an AI role in the endline; α_i and α_r are two-digit industry and city fixed effects; and α_{f10} is a firm size decile fixed effect, where firm size deciles are calculated over firm posting volumes in the baseline period.¹⁹ We cluster standard errors at the firm level to account for common shocks across establishments within the same parent firm.

The primary identification concern is that AI adoption is systematically related to unobserved establishment-level factors that also increase non-AI labor demand, such as manager quality.²⁰ To address such potential endogeneity, we instrument AI adoption with ‘AI exposure’ to isolate changes resulting from supply-side technical advances in the capabilities of AI. We first take the occupation-level AI exposure measure developed by Webb (2020), which measures the extent to

¹⁶We 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.

¹⁷Specifically, we drop education, IT, internet and e-commerce, telecom and internet service providers, which make up 34.8 percent of the sample.

¹⁸Appendix A.3 provides summary statistics for this sample.

¹⁹The variables $\Delta y_{fr,t-t_0}$ approximate the growth in establishment outcomes between 2010-12 and 2017-19. Mathematically, for growth rate g defined by $Y_t = (1 + g)Y_{t_0}$, and using the approximation that $\ln(1 + g) \approx g$ for small g , we have $g = \ln Y_t - \ln Y_{t_0} = \Delta y_{t-t_0}$.

²⁰For example, more innovative managers may be more likely to hire AI workers, but are also more productive and grow the business more quickly, hence also increasing non-AI labor demand. Appendix Table B.35 shows OLS results that support this suspected upward bias.

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 map the Webb (2020) exposure measure to the Indian National Classification of Occupations (NCO) 2004 at the four-digit level using publicly-available crosswalks. To capture establishment-wise exposure to AI-based automation, we then aggregate this measure to the establishment level by weighting across baseline establishment occupation shares.²² Specifically, we calculate:

$$Exposure_{fr,t_0} = \sum_o PostingShare_{fro}^{t_0} \cdot WebbExposure_o \quad (4.2)$$

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:

$$AdoptsAI_{fr,t-t_0} = \gamma \cdot Exposure_{fr,t_0} + \alpha_r + \alpha_i + \alpha_{f10} + e_{fr,t-t_0}, \quad (4.3)$$

This is a linear probability model, such that the first stage coefficient γ estimates the increase in the probability of an establishment adopting AI between 2010-12 and 2017-19 that is associated with a one standard deviation rise in establishment exposure. When instrumenting equation 4.1 with equation 4.3, the coefficient β estimates the percentage point increase in the growth rate of outcome Y_{frt} that is associated with a 1 percent rise in the predicted probability of a firm adopting AI between 2010-12 and 2017-19.

We discuss several variations of this main specification in Section 6, including alternatives to a linear probability model in the first stage (Online Appendix Tables B.32 and B.33) and alternatives to using an AI adoption dummy (Online Appendix Tables B.26 to B.29).

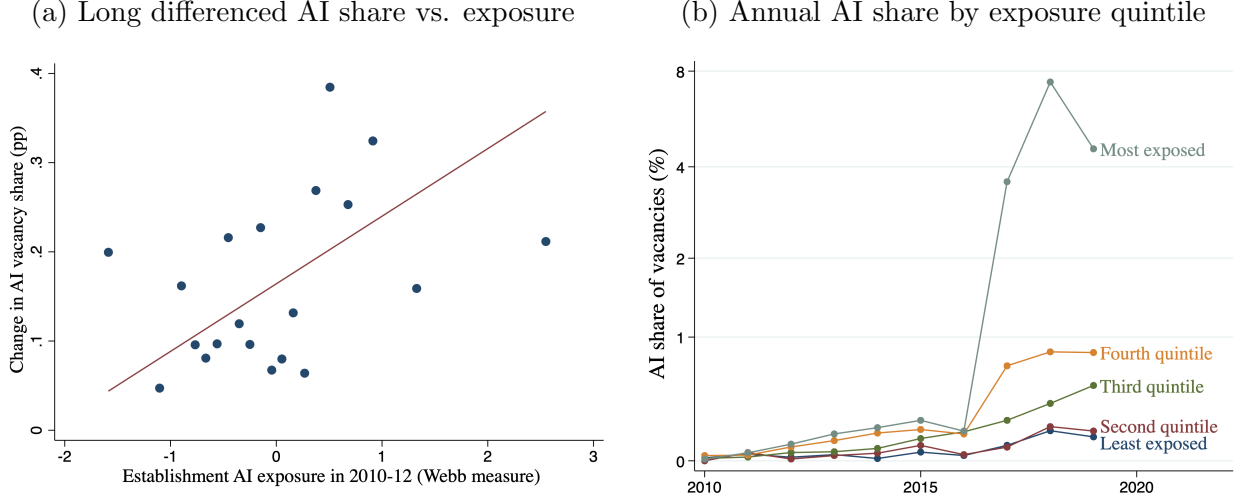
4.2 Identification assumptions

Our case for a causal interpretation stems from the baseline posting 'shares' by occupation in 2010-12 ($PostingShare_{fro}^{t_0}$) being exogenous with respect to the outcome variables (primarily changes in establishment-level postings between 2010-12 and 2017-19) after conditioning on

²¹These task descriptions are based on US occupations. While Indian occupations in general may have different task compositions, the white-collar service sector – which contains many multinational firms – 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.

²²Online Appendix Figure A.9 shows the distribution of exposure scores across occupation-wise wage offer percentiles. AI exposure rises with wage offers up to a peak around the 80th percentile, before falling thereafter.

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



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 following Cattaneo et al. (2023). Panel (b) plots the time variation in this relationship, using an inverse hyperbolic sine scale for the y-axis.

our fixed effects. Our Bartik instrument combines these shares with the 'shocks' of occupation-varying measures of exposure to AI patents, $WebbExposure_o$.

Instrument relevance requires that the instrument must correlate with the endogenous variable, $\Delta Adoption_{fr,t-t_0}$, conditional on region, industry, and firm size decile fixed effects, $\alpha_r, \alpha_i, \alpha_{f10}$. We find that AI exposure does indeed predict AI demand (Figure 4.1 Panel (a)). A one standard deviation higher establishment AI exposure score is associated with a significant 1.06 percent increase ($p < 0.01$) in the probability of an establishment adopting AI between 2010-12 and 2017-19, after controlling for region, industry and firm size decile fixed effects.²³ Panel (b) of Figure 4.1 shows the relationship between AI exposure and AI adoption over time and confirms that the 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 to other exposure quintiles until 2016, before rapidly diverging to reach almost 8 percent in 2018.

Instrument validity requires that, conditional on region, industry and firm size decile fixed effects, our AI exposure instrument is exogenous with respect to the error term. We require:

$$\mathbb{E} \left[Exposure_{fr,t_0} \cdot \epsilon_{fr,t-t_0} | \alpha_r, \alpha_i, \alpha_{f10} \right] = 0. \quad (4.4)$$

This assumption has a more intuitive interpretation. Goldsmith-Pinkham et al. (2020) highlight

²³Full results are provided in Online Appendix Table B.4.

that this exogenous shares approach is equivalent to combining many difference-in-differences designs. Conditional on fixed effects, we need to argue that each of these is valid. For each occupation o , we require:

$$Cov\left[\epsilon_{fr,t-t_0}, PostingShare_{fro}^{t_0} | \alpha_r, \alpha_i, \alpha_{f10}\right] = 0. \quad (4.5)$$

As our outcome is measured in first differences, this is a parallel trends assumption. Thus, we claim that establishments with different baseline posting shares would have had parallel trends in labor demand and wage offers in the absence of the occupation-level AI shock. While this exogeneity assumption cannot be directly tested, we suggest both *ex-ante argument* and *ex post* robustness tests.

First, we argue that, conditional on fixed effects, baseline occupation shares are exogenous to future trends in labor demand and wage offers. $\epsilon_{fr,t-t_0}$ is a residual trend in labor demand and wage offers between the baseline (2010-12) and endline (2017-19) periods. Baseline occupation shares, $PostingShare_{fro}^{t_0}$, are predetermined by definition. Thus, for these shares to be influenced by $\epsilon_{fr,t-t_0}$, establishments would have had to predict the AI boom and adjust their baseline occupation shares preemptively. We contend that this is highly unlikely, since the AI boom was sudden and unexpected. The boom was driven by significant advances in computer science, such as the AlexNet breakthrough in late 2012, which were rapidly commercialized and disseminated, for instance through the open-source release of TensorFlow in 2015. We argue that these breakthroughs were not predictable during the baseline period and thus should not have influenced the occupation structure of establishments at that time.

Could $\epsilon_{fr,t-t_0}$ be correlated with baseline posting shares for other reasons? While we cannot rule out all potential unobserved shocks affecting both baseline posting shares and subsequent labor demand trends, our fixed effects already account for varying trends by region, industry and firm size. Moreover, the rapid adoption of machine learning techniques, as documented in Section 3, represent a large, generalized shock to the services sector, with no contemporaneous parallels of which we are aware. This limits the potential for our estimates to be biased by other unobserved technology shocks.

Following Goldsmith-Pinkham et al. (2020), we provide three robustness checks to affirm that our instrument is valid, which are discussed in more detail in Online Appendix B.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

no predictive power. Finally, we compare a range of estimators and run over-identification tests, finding similar results across estimators. All three tests thus support the validity of our instrument.

4.3 Other potential identification concerns

Our approach also addresses two other potential identification concerns. First, it mitigates the risk of mean reversion bias when relating total labor demand to AI exposure. In a reduced-form setting, an idiosyncratic rise in hiring of AI-exposed workers during the baseline period could artificially inflate an establishment’s measured exposure and subsequently reduce the total number of posts as the effect dissipates. This could lead to a spurious negative relationship between AI exposure and vacancy growth. However, our specification uses AI adoption between 2010-12 and 2017-19 as the main regressor, with AI exposure an instrument. For the same bias to occur, the idiosyncratic increase in 2010-12 posting of AI-exposed occupations would need to be systematically related to a higher probability of adopting AI between 2010-12 and 2017-19 – which is not the case for a truly idiosyncratic increase. Our approach in contrast captures the true effect: high baseline susceptibility to AI capabilities increases the probability of adopting AI in subsequent years, which in turn lowers total labor demand.

A second potential concern is that the hiring patterns of Indian firms may have influenced AI innovation, biasing our results. Building on our previous discussion of identification assumptions, we offer two additional observations. First, India was not a significant producer of new AI research during our study period, lagging behind major hubs like the USA and China, despite strengths in applied computer science (Perrault et al. 2019). Thus, global advances in AI patenting are unlikely to be affected by hiring patterns in Indian firms.²⁴ Second, by excluding establishments in AI-producing sectors and focusing on AI usage, we drop the establishments where such concerns would be most relevant.

5 Main Results

5.1 Impacts of AI on labor demand

We first examine the effects of AI adoption on non-AI vacancies. Table 5.1 shows the impact of AI adoption on the growth of non-AI vacancies and total vacancies, instrumenting with the Webb (2020) AI exposure measure. AI adoption reduces the growth in non-AI demand: a one

²⁴Indeed, this means our shocks could plausibly be exogenous as well.

percent increase in the predicted probability of adopting AI (which corresponds to approximately a one-standard deviation increase in AI exposure in the first stage) results in a 8.1 percentage point decrease ($p < 0.01$) in non-AI vacancy growth at the establishment level between 2010-12 and 2017-19 in our main specification with region, industry and firm-decile fixed effects. There is a similarly sized decrease of 7.8 percentage points in the growth of total vacancies, highlighting that AI vacancies crowd out other white-collar services-sector vacancies.²⁵ Considering that the median growth rate in total and non-AI vacancies is 24.9 percent, these are substantial effects.

Table 5.1: 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)	(5)	(6)
Adoption of AI	-7.975*** (2.350)	-12.90*** (3.092)	-8.064*** (2.282)	-7.737*** (2.245)	-12.47*** (2.959)	-7.840*** (2.181)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	43.7	41.58	45.43	44.06	41.83	45.62
Observations	22,244	22,244	22,244	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 log of the respective establishment-level outcomes. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 reported F statistic is the cluster-robust Kleibergen-Paap Wald F -statistic.

5.2 Mechanisms and distributional impacts

To understand the drivers of these effects and identify which types of jobs and tasks are most impacted by AI adoption, we explore the heterogeneous impacts of AI adoption across occupations, skill requirements and tasks. We find the following results:

AI adoption reduces growth in demand for higher-skilled occupations. We first study the effects of AI on postings at the occupation level, using India’s NCO2004 classification of one-digit and two-digit occupations. We regress AI adoption on the change in the share of non-AI vacancies within each occupation group. Table 5.2 shows that the decline in overall demand is

²⁵In both these results, the standard errors increase to 1.38 when implementing the adjustment by Lee et al. (2022), leaving the finding significant ($p < 0.05$).

accompanied by a shift away from higher-skilled occupations: the categories of ‘Professionals’ and ‘Managers’ suffer large reductions in their respective shares. For every one percent increase in the predicted probability of adopting AI, we observe a 10.6 percentage point decrease in the growth of non-AI vacancies for ‘Managers’ and a 3.6 percentage point decrease for ‘Professionals’. These impacts are significant, given that these occupations comprised 18 percent and 48 percent of all postings in the 2010-12 baseline period respectively. Conversely, we note increases in shares for lower-skilled workers, such as ‘Personal, Sales and Security’, ‘Clerks’ and ‘Associate Professionals’. These findings align with Webb (2020) who found that AI disproportionately affects high-skilled jobs involving pattern detection, judgement and optimization, such as clinical laboratory technicians, chemical engineers and optometrists.

Disaggregating this result further within these two groups using two-digit occupation codes, Appendix Table B.13 shows that increased AI demand reduces the share of non-AI vacancies for ‘Engineering Professionals’, ‘General Managers’, and particularly strongly for ‘Corporate Managers’. In the baseline period, 24 percent of all postings are for ‘Engineering Professionals’, 6 percent for ‘General Managers’, and 12 percent for ‘Corporate Managers’. Reduced vacancies growth for these three occupations consequently plays a large role in our aggregate results.

Table 5.2: Second stage: Impact of AI adoption on establishment non-AI vacancy shares, by occupation group

	Change in Non-AI Vacancy Shares				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Adoption of AI	2.074*** (0.385)	1.324*** (0.272)	10.46*** (1.718)	-3.637*** (0.717)	-10.59*** (1.709)
<i>Fixed Effects:</i>					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	45.43	45.43	45.43	45.43	45.43
Observations	22,244	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 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 occupation shares. Occupation groups are 1-digit occupations from the NCO04. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 reported F statistic is the cluster-robust Kleibergen-Paap Wald F-statistic.

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 literature examining routine and non-routine task intensity of occupations (e.g. Autor et al. (2003) and Acemoglu & Autor (2011)). We map the measures of Acemoglu & Autor (2011) onto India’s NCO2004 classification of occupations and standardize them.²⁶ We then analyze the impact of instrumented AI demand on the change in logged routine and non-routine scores at the establishment-level. Table 5.3 documents the significant negative effect of increased AI demand on establishment-level growth in non-routine task intensity. A one percent higher predicted probability of AI adoption leads to 3.3 percentage point lower growth in non-routine task intensity. In contrast, the intensity of routine tasks is not significantly affected.²⁷

This result of declining demand for non-routine tasks could reflect changing task demand within occupations or a shift in demand between occupations. We hence also ask how AI adoption affects establishment-level non-routine task intensity within occupation groups, as shown in Table 5.4. We find that AI adoption also reduces growth in demand for non-routine tasks within higher-skilled occupations. We find the strongest reduction in non-routine task growth for the category of ‘Managers’.²⁸

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. But AI could also impact the task content within occupations over time. We therefore adopt a further, more granular, approach by studying the verbs mentioned within the text of job adverts. For each establishment, we count both the number of appearances of each verb and the total number of appearances of all verbs at baseline and endline. We thus construct a ‘verb share’ for each verb, as well as its change over time. We also repeat this procedure for groupings of verbs with similar meanings, using the verb classification of Roget’s Thesaurus as discussed above.

To analyze changing demand in tasks, we run regressions of a similar form to those in Table 5.1, using changes in verb shares as the dependent variable. Figure 5.1a shows that

²⁶To allow us to consider percentage impacts using logs, we transform the standardized values x to $x + x_2$, where x_2 is the second lowest observed value, such that all transformed values are positive.

²⁷Online Appendix Table B.14 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.

²⁸The estimated effect differs significantly across occupation groups, with the hypothesis of equal coefficients rejected at the 1 percent level (p value=0.00). Differences are also generally statistically significant for pairwise comparisons at the 1 percent level, but insignificant for comparisons of ‘Professionals’ and ‘Personal’, ‘Professionals’ and ‘Clerks’, and ‘Personal’ and ‘Clerks’. Online Appendix Table B.15 repeats Table 5.4 for abstract tasks following Autor & Dorn (2013), finding similar results of a decline in growth in demand for abstract tasks within professional and managerial occupation groups, with the latter impact especially large.

Table 5.3: 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)
Adoption of AI	-3.395*** (0.889)	-3.281*** (0.894)	-3.258*** (0.849)	0.287 (0.841)	0.243 (0.846)	0.273 (0.804)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	52.87	59.01	56.31	52.87	59.01	56.31
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 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. We also control for the change in the log of total posts. The dependent variables are the change in the log of the respective establishment-level outcomes, where we have shifted the mean of the distributions of routine and non-routine scores to be able to take logarithms. 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 standardized routine and non-routine scores. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 reported F statistic is the cluster-robust Kleibergen-Paap Wald F -statistic.


AI demand negatively impacts the share of verbs related to ‘Intellectual Faculties’ in Roget’s Thesaurus. Specifically, a one percent increase in predicted AI adoption probability leads to a 0.29 percentage point decline in the share of those verbs. Figure 5.1b further breaks down this effect, showing statistically significantly negative effects on three sections within the ‘Intellectual Faculties’ category: i) ‘Precursory Conditions’, involving analytical tasks (e.g. ‘investigate’, ‘research’ and ‘explore’), ii) ‘Means of Communicating Ideas’ involving complex communication tasks (e.g. ‘narrate’ and ‘describe’) and to a lesser extent iii) ‘Extension of Thought’, involving prediction (e.g. ‘predict’ and ‘forecast’). These verbs align closely with known AI capabilities and support the idea that AI reduces the cost or improves the quality of the task of ‘prediction’ prevalent across many occupations (Agrawal et al. 2018).

Table 5.4: Second stage: Impact of AI adoption on establishment non-routine tasks, by occupation group

	Growth in Non-Routine Tasks				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Adoption of AI	0.113 (0.462)	-1.363** (0.552)	-6.378*** (1.646)	-5.564*** (1.612)	-10.76*** (2.052)
<i>Fixed Effects:</i>					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	47.03	46.82	49.29	48.36	47.19
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 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, where we have shifted the mean of the distributions of routine and non-routine scores to be able to take logarithms. We average standardized routine and non-routine O*NET task contents by occupation, and form establishments' routine and non-routine task demand by weighting occupations by their standardized routine and non-routine scores. We control for the change in the respective occupation's vacancy share. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 reported F statistic is the cluster-robust Kleibergen-Paap Wald F -statistic.

Figure 5.1: Impact of 1 unit higher establishment AI demand growth on verb usage by class and section



Notes: These coefficient plots show the impact of AI adoption 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 percent and 90 percent confidence intervals. Each coefficient is from a regression of type (3) in Table 5.1. Here, the outcome variable is growth in the share of verbs from the respective section or class. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 size decile and industry fixed effects.

5.3 Impacts of AI on wage offers

We analyze wage offers by taking advantage of the platform’s standardized job postings template, requiring firms to list wage offers for all vacancies.²⁹ Table 5.5 documents the impact of AI adoption on the growth of median wages for both non-AI postings and all job postings. We find that a one percent higher predicted probability of AI adoption reduces the growth of non-AI wage offers by 5.5 percentage points between 2010-12 and 2017-19, instrumenting with AI exposure and controlling for region, industry and firm size fixed effects. This negative effect remains largely unchanged when considering all posts, including AI postings, despite the higher wage offers associated with AI posts.

Table 5.5: 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)
Adoption of AI	-5.696*** (1.485)	-6.351*** (1.630)	-5.514*** (1.423)	-5.452*** (1.424)	-6.089*** (1.566)	-5.273*** (1.366)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	45.14	42.91	46.86	45.46	43.12	47.02
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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 reported F statistic is the cluster-robust Kleibergen-Paap Wald F-statistic.

The negative impact of AI demand on non-AI wage offers is evident across the wage offer distribution, with effects statistically significantly different from zero from the 20th percentile onwards. Figure 5.2 illustrates the impact of a one percent higher predicted probability of AI adoption on wage offer growth at various percentiles of the establishment-level distribution between 2010-12 and 2017-19, using the same AI exposure instrument and fixed effects as above. We observe a statistically significant reduction (at the 5 percent level) in establishment-level wage offers for non-AI jobs over time, ranging from -2.5 to -4.4 percent across the distribution.

²⁹This contrasts with commonly used platforms in advanced economies, where wage offer information is missing for the majority of posts (Batra et al. 2023).

Figure 5.2: Impact of AI adoption on the wage offer distribution in non-AI posts




Notes: This coefficient plot shows the impact of a one percent rise in the predicted probability of adopting AI on wage growth over time across the distribution of establishment wage offers. As in Online Appendix Table B.17, AI adoption is instrumented by AI exposure. Each coefficient is from a regression of type (3) in Appendix Table B.17 and represents the impact of a one percent rise in the predicted probability of adopting AI on wage growth over time for a given percentile of the wage offer distribution. Standard errors are clustered at the firm level, and we include region, firm size decile and industry fixed effects. The outer (inner) error bars reflect 95 percent (90 percent) confidence intervals. 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.

The observed lower wage growth associated with AI demand could stem from two potential sources: between-occupation or within-occupation effects. In other words, AI could change the occupational composition, thereby shifting the median wage offer, or it could reduce wage offers for individual occupations. Our findings in Table 5.2 showed that AI demand lowers hiring growth for the highest paid occupations, which likely contributes to the downward shift in overall wage offers. However, evaluating wage offer effects within occupation groups is challenging due to small sample sizes when splitting the data by occupation group.³⁰ When we add controls for changing occupation group shares, we find no significant declines in wage offer growth at the 5 percent level of significance, suggesting that composition effects are the primary driver of the results.

³⁰We nevertheless report heterogeneous results on wage offers by occupation group in Online Appendix Table B.16, finding weak negative effects for higher-skilled occupation groups.

Figure 5.3: Impact of AI adoption on the non-AI wage offer distribution, holding occupational composition fixed



Notes: This coefficient plot shows the impact of a one percent rise in the predicted probability of adopting AI 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 ‘Professionals’. As in Online Appendix Table B.17, AI adoption is instrumented by AI exposure. Each coefficient is from a regression of type (3) in Online Appendix Table B.17 and represents the impact of a one percent rise in the predicted probability of adopting AI on wage growth over time for a given percentile of the wage offer distribution. Standard errors are clustered at the firm level, and we include region, firm size decile and industry fixed effects. The outer (inner) error bars reflect 95 percent (90 percent) confidence intervals.

6 Robustness and Extensions

In this section, we first show that our findings are robust to several variations of our main specification and an event-study approach. We then investigate wider effects of AI, beyond the establishment level, and find significant impacts on (national) firms that adopt AI, but no robust evidence yet for aggregate district-level effects.

6.1 Alternative specifications

We show that our results are robust to several alternative specifications, with results provided in Online Appendix B. First, we show that our results are robust to adding additional controls for the baseline establishment posting shares of software engineers and sales and administrative professionals, following Acemoglu et al. (2022). By controlling for two broad occupations that see a decline due to computerization (Autor & Dorn 2013), we account for the possibility that firms experiencing increased demand for machine learning skills might also be software-engineering intensive or more affected by computerization. Our results remain robust with these additional controls (Tables B.18 to B.21).

Second, we show that our results are robust to using a shorter time period between baseline (2013-15) and endline (2017-19) in Tables B.22 to B.23. This approach draws on a larger sample

of firms, as firms are more likely to post in this shorter timeframe, but at the cost of not spanning the entire takeoff in AI demand (Figure 3.1). Nonetheless, the results are similar to those in the main specification. This shorter long-difference also allows us to differentiate between firms newly appearing in the data (termed ‘startups’ for convenience) and older firms (‘incumbents’) appearing in the data from the start.³¹ In Tables B.24 and B.25, we show evidence of negative impacts on labor demand and wage offers in both groups of firms in the shorter long-difference specification, albeit less significant for newer firms where there is a weaker first stage.

Third, we demonstrate the robustness of our results using an alternative dependent variable that goes beyond our main ‘AI demand’ measure by taking into account the full quantity of AI vacancies posted by an establishment, rather than simply whether it is positive. In Tables B.26 to B.29, we replace $AdoptsAI_{fr,t-t_0}$ in equation 4.1 with either the log of one plus the number of AI posts by an establishment or the inverse hyperbolic sine transformation of the same. These measures introduce an intensive margin (albeit at a cost to interpretability, as discussed in Chen & Roth (2023)). Nevertheless, the consistent findings across specifications provide reassurance that our results are not driven by our choice of independent variable.

Lastly, we confirm the robustness of our results to several other specification decisions. Our findings remain unchanged when weighting by baseline establishment size (Tables B.30 and B.31), where this is proxied by the total number of vacancies posted in the baseline, with the upper 5 percent winsorized to reduce the dominance of the top firms as highlighted in Section 3. Likewise, our results are robust to using a two-step consistent estimator accounting for endogenous treatment instead of a one-step linear probability model (Tables B.32 and B.33) following Wooldridge (2010). Our wage results are robust to using mean rather than median wage offers (Table B.34). Finally, when using the alternative AI exposure measure of Felten et al. (2018), we find similar negative impacts on wage offers and relative demand for professionals and managers (Appendix B.4).

6.2 Event-study approach

Our long-difference approach focuses on a fixed window (from 2010-12 to 2017-19) centered on the national (and global) take-off in AI demand in the middle of the 2010s. To provide an alternative perspective, we also take an event study approach, centering our analysis on the specific moment each establishment first posts an AI vacancy. This in turn allows the observation window to vary across establishments. In this case, the ideal experiment would

³¹Specifically, we classify an establishment as an incumbent if it posts vacancies in the years 2010-2012 and 2017-19, and as a startup if it posts vacancies in 2013-15 and 2017-19 but not 2010-12.

compare the post-adoption outcomes of each AI-adopting establishment to those of an otherwise identical establishment that instead did not adopt AI. However, AI adopters are unlikely to resemble the average firm – for instance, they tend to be larger, as shown in Section 3.

To approach the ideal ‘otherwise identical’ counterfactual, we therefore follow Koch et al. (2021) in matching AI adopters to similar non-adopters using propensity scores. We construct these by running a Probit regression of AI adoption on a range of lagged establishment characteristics and hence deriving predicted adoption probabilities.³² By accounting flexibly for a wide range of observables, this approach aims to control for unobservable differences between establishments. Rosenbaum & Rubin (1983) show that the propensity score is a balancing score, such that conditional on the propensity score, AI adoption is indeed orthogonal to observable establishment characteristics.

We then run an annual propensity score-weighted event-study regression at the level of establishments i :

$$Y_{it} = \alpha_i + \alpha_t + \sum_{k=-3 \setminus -1}^2 \beta_k \cdot 1(K_{it} = k) + \beta_{3+} \cdot 1(K_{it} \geq 3) + \epsilon_{it} \quad (6.1)$$

where Y_{it} is the inverse hyperbolic sine (IHS)-transformed number of job postings, α_i and α_t respectively are establishment and time fixed effects, and K_{it} is the time difference between the current year and adoption of AI.³³ We include three lags and leads, omitting the first lead to normalize the difference between adopters and similar non-adopters to zero in the year before adoption. The estimates β_k then measure the impact of AI adoption on the IHS of Y_{it} , relative to the level in similar non-adopters.


Figure 6.1 plots the resulting estimates. We find that, following an initial positive impact of AI adoption on non-AI vacancies, non-AI vacancies are significantly lower (at the 5 percent level) in the second year after adoption. This negative impact increases in magnitude in the third year following AI adoption. Non-AI vacancies are 0.7 IHS points lower for adopters in the second year after adoption, and one IHS point lower three years after adoption.³⁴

³²Specifically, we include lags of firm size decile, establishment size, salary, experience levels, firm age, the standard deviation of salaries and experience, and several interaction terms, as shown in Online Appendix Table B.36.

³³We take the inverse hyperbolic sine transformation to balance the panel such that firms that do not post within a year are not dropped from the sample.

³⁴These findings are also robust to using the imputation estimator of Borusyak, Jaravel & Spiess (2021).

Figure 6.1: Event study for non-AI postings following AI adoption



Notes: Two way fixed effects on a balanced panel. The outcome variable is the IHS-transformed number of establishment non-AI vacancies. Adopters are matched to never adopters by propensity score weighting, with propensity scores from a Probit regression of establishment characteristics on AI adoption (see Online Appendix Table B.36 for details). We leave out the first lead as the base period and cluster standard errors on establishments.

6.3 Wider effects

Does AI demand affect broader labor market outcomes, beyond the establishment level? First, we consider impacts at the firm level by aggregating nationally across a firm's constituent establishments. Tables 6.1 and 6.2 respectively show the results for the number of vacancies and the median wage offer. While the sample size drops substantially due to the aggregation and the F -statistic is correspondingly much weaker, we nonetheless find significant negative impacts of AI adoption on the growth in firms' overall posting of vacancies and their average wage offers.

Second, we consider impacts at the district level by aggregating across establishments in a given region (Table 6.3), with a significant reduction in sample size. We nonetheless find that AI exposure predicts AI adoption (column 1), but we do not find a significant impact on growth in non-AI or total vacancies (columns 2 and 3), and at most weak evidence ($p < 0.11$) for a negative effect on non-AI and overall wage offer growth (columns 4 and 5). Given the weak first stage in these long-difference regressions, we also explore results from a district-level event-study approach (Appendix Figure B.2), but again do not detect a significant impact of AI adoption on labor demand. Given that only 3.9 percent of establishments in our study used AI between 2010 and 2019, it is unsurprising that the impact of AI adoption on labor demand in

AI-using industries is not detectable in district-level aggregates. This suggests that the effect, while present at the establishment level, was too small to influence broader labor market trends during our study period.

Table 6.1: Second stage: Impact of AI adoption on firm non-AI vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-4.603*	-12.05**	-3.140	-4.377**	-10.44**	-3.111*
	(2.445)	(5.934)	(2.079)	(2.186)	(4.722)	(1.860)
<i>Fixed Effects:</i>						
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	18.38	7.268	18	19.7	8.475	19.16
Observations	5,701	5,702	5,701	5,703	5,704	5,703

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the firm posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). The reported F statistic is the cluster-robust Kleibergen-Paap Wald F -statistic.

Table 6.2: Second stage: Impact of AI adoption on firm non-AI wages

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-2.983** (1.204)	-5.211** (2.507)	-3.034** (1.228)	-2.690** (1.089)	-4.365** (2.044)	-2.703** (1.105)
<i>Fixed Effects:</i>						
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	18.65	7.356	18.23	19.97	8.565	19.38
Observations	5,671	5,672	5,671	5,673	5,674	5,673

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure is calculated as the standardized average of occupation AI exposure (from Webb 2020), over the occupations for which the firm posts vacancies in 2010-12, weighted by the number of vacancies posted per occupation, as in Acemoglu et al. (2022). The reported F statistic is the cluster-robust Kleibergen-Paap Wald *F*-statistic.

Table 6.3: AI adoption at the district level

	AI Adoption	Non-AI Vacancies	Vacancies	Non-AI Wages	Wages
	(1)	(2)	(3)	(4)	(5)
<i>First stage:</i>					
AI Exposure	0.0253** (0.0107)				
<i>Second stage:</i>					
Adoption of AI		0.601 (1.995)	0.610 (1.996)	-2.430 (1.519)	-2.411 (1.514)
First Stage F-Stat		5.569	5.569	6.300	6.300
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 in log differences. 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 coefficient therefore represents the impact on the probability to adopt AI of a one-standard deviation rise in AI exposure. The independent variable for the second stage is a dummy for district AI adoption between 2010-12 and 2017-19. AI adoption is instrumented by district AI exposure. The latter four coefficients therefore represent the percentage impact on the outcome variable of a one percent rise in the predicted probability of adopting AI. The reported F statistic is the cluster-robust Kleibergen-Paap Wald *F*-statistic.

7 Conclusion

AI could have important implications for a services-led development model, given the potential for recent advances in machine learning to automate many service sector 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, and implications of, AI capabilities in the white-collar services sector. We document a rapid take-off in demand for AI-related skills from 2015, particularly in the IT, finance and professional services industries, closely mirroring patterns found for advanced economies (Grennan & Michaely 2020, Acemoglu et al. 2022). AI jobs pay a substantial wage premium and are highly concentrated in certain industries, cities and large firms.

We assess the labor market effects of establishment-level demand for AI skills, as a proxy for AI deployment. Employing a long-difference shift-share specification based on exposure to patented advances in AI capabilities (Webb 2020), we find that AI adoption significantly reduces growth in non-AI postings and average wage offers. This decline primarily reflects a reduction in demand for skilled managerial and professional occupations, non-routine work, and analytical and communication tasks. This contrasts with previous findings on computerisation and robotics, which primarily affected 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 description. 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).

Our findings underscore AI’s potential to reduce the availability of high-skilled, high-wage white-collar jobs in developing countries — traditionally the target of services-led development strategies. However, while we observe significant negative impacts of AI adoption on labor demand and wage offers within firms, these effects do not translate to significant changes in district-level aggregates. This discrepancy likely stems from the relatively small proportion of establishments adopting AI during our sample period, such that their adoption decisions did not have a detectable impact on aggregate regional outcomes. As AI usage diffuses more widely, understanding the impacts on local labor markets in developing countries will be an important task for future research.

References

- Acemoglu, D. (2024), ‘The Simple Macroeconomics of AI’, *NBER Working Paper Series* .
- Acemoglu, D. & Autor, D. (2011), 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), ‘Artificial Intelligence and Jobs: Evidence from Online Vacancies’, *Journal of Labor Economics* **40**.
- 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.
- 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.
- 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., Gans, J. S. & Goldfarb, A. (2024), ‘Artificial intelligence adoption and system-wide change’, *Journal of Economics & Management Strategy* **33**(2), 327–337.
- Alekseeva, L., Azar, J., Gine, M., Samila, S. & Taska, B. (2020), ‘The demand for AI skills in the labour market’.
- Alfaro, L. & Chari, A. (2009), ‘India transformed? insights from the firm level 1988-2005’, *Harvard Business School Working Paper Series* .
- Alonso, C., Berg, A., Kothari, S., Papageorgiou, C. & Rehman, S. (2022), ‘Will the AI revolution cause a great divergence?’, *Journal of Monetary Economics* **127**(C), 18–37.
- 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.
- 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.
- Babina, T., Fedyk, A., He, A. & Hodson, J. (2024), ‘Artificial intelligence, firm growth, and product innovation’, *Journal of Financial Economics* **151**, 103745.
- Babina, T., Fedyk, A., He, A. X. & Hodson, J. (2023), Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition, in ‘Technology, Productivity, and Economic Growth’, University of Chicago Press.
- 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.
- Batra, H., Michaud, A. M. & Mongey, S. (2023), ‘Online Job Posts Contain Very Little Wage Information’, *Opportunity and Inclusive Growth Institute Working Papers* . Number: 083 Publisher: Federal Reserve Bank of Minneapolis.

- Benhabib, J. & Spiegel, M. M. (2005), Human Capital and Technology Diffusion, Vol. 1 of *Handbook of Economic Growth*, Elsevier, pp. 935–966. ISSN: 1574-0684.
- 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.
- Bloom, N., Hassan, T. A., Kalyani, A., Carvalho, M., Lerner, J. & Tahoun, A. (2024), The Diffusion of New Technologies, SSRN Scholarly Paper 3875827, 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.
- Borusyak, K., Jaravel, X. & Spiess, J. (2021), ‘Revisiting event study designs: Robust and efficient estimation’.
- Brynjolfsson, E., Hui, X. & Liu, M. (2019), ‘Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform’, *Management Science* **65**(12), 5449–5460.
- Brynjolfsson, E., Li, D. & Raymond, L. R. (2023), Generative ai at work, Technical report, National Bureau of Economic Research.
- 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.
- Cattaneo, M. D., Crump, R. K., Farrell, M. H. & Feng, Y. (2023), ‘On binscatter’.
- Chen, J. & Roth, J. (2023), ‘Logs with Zeros? Some Problems and Solutions*’.
- 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.
- 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.
- Gmyrek, P., Berg, J. & Bescond, D. (2023), Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality, SSRN Scholarly Paper 4584219, Social Science Research Network, Rochester, NY.
- Goldberg, P., Khandelwal, A., Pavcnik, N. & Topalova, P. (2009), ‘Trade liberalization and new imported inputs’, *American Economic Review* **99**(2), 494–500.
- 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.
- Goldsmith-Pinkham, P., Sorkin, I. & Swift, H. (2020), ‘Bartik instruments: What, when, why, and how’, *American Economic Review* **110**(8), 2586–2624.

- 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*.
- Harrigan, J., Reshef, A. & Toubal, F. (2020), ‘The March of the Techies: Job Polarization Within and Between Firms’, *Research Policy* p. 104008.
- 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.
- Lee, D. S., McCrary, J., Moreira, M. J. & Porter, J. (2022), ‘Valid t-ratio inference for iv’, *American Economic Review* **112**(10), 3260–90.
- 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.
- McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. & Zolas, N. (2024), ‘AI adoption in America: Who, what, and where’, *Journal of Economics & Management Strategy* **33**(2), 375–415.
- 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.
- Noy, S. & Zhang, W. (2023), ‘Experimental evidence on the productivity effects of generative artificial intelligence’, *Science* **381**(6654), 187–192.
- 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.
- Pizzinelli, C., Panton, A., Mendes Tavares, M., Cazzaniga, M. & Li, L. (2023), Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications, SSRN Scholarly Paper 4612697, Social Science Research Network, Rochester, NY.

- Raj, M. & Seamans, R. (2018), *Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data*, University of Chicago Press, pp. 553–565.
- 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.
- Rosenbaum, P. R. & Rubin, D. B. (1983), ‘The central role of the propensity score in observational studies for causal effects’, *Biometrika* **70**(1), 41–55.
- Shimer, R. (2012), ‘Reassessing the Ins and Outs of Unemployment’, *Review of Economic Dynamics* **15**(2), 127–148.
- Trefler, D. & Sun, R. (2022), AI, Trade and Creative Destruction: A First Look, NBER Working Papers 29980, National Bureau of Economic Research, Inc.
- Webb, M. (2020), ‘The Impact of Artificial Intelligence on the Labor Market’, *SSRN Electronic Journal* .
- Wooldridge, J. M. (2010), *Econometric Analysis of Cross Section and Panel Data*, The MIT Press.

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

Alexander Copestake Ashley Pople
Max Marczinek Katherine Stapleton

September 2024

FOR ONLINE PUBLICATION

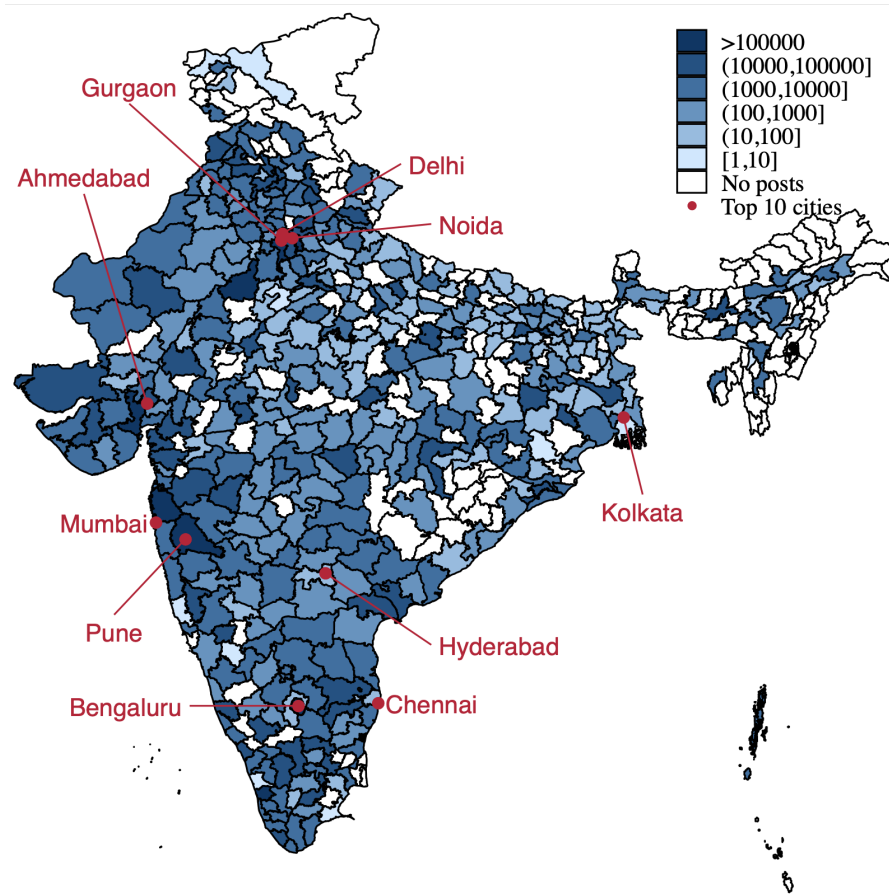
This supplementary online appendix contains two sections. Appendix A lays out how we construct our dataset, benchmarks it against administrative data, discusses representativeness, and provides additional descriptives. Appendix B provides an array of additional results related to our shift-share instrument, our main specification, and alternative specifications.

Appendix A Data and Descriptives

A.1 Construction of the vacancy dataset

The largest online job postings platform in India shared 80 percent of all job postings on its site (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 percent 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 percent of city names and add geo-locations, separating out overseas job postings. Using the geolocations, we match cities to districts, using the 2011 census. Figure A.1 shows the distribution of posts across districts.

Figure A.1: 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.

Table A.1 summarises the number of observations in our main dataset, as well as in three other datasets that we use to gauge its representativeness (see the next section). 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. (2018) measure, we map the 2008 ISCO to the 1988 ISCO, before again mapping onto the 2004 NCO.

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
— <i>Financial</i>	17,805	1,815,798
— <i>Information</i>	72,057	5,834,878
— <i>Professional</i>	38,533	834,932
— <i>Other</i>	106,798	6,995,722
Prowess (balance sheets)	#Firms	#Observations
Agriculture	123	590
Manufacturing	2,276	11,257
Services	3,675	16,722
— <i>Financial</i>	1,020	4,830
— <i>Information</i>	516	2,557
— <i>Professional</i>	199	811
— <i>Other</i>	1,940	8,524
Surveys (demographics)	#Districts	#Households
National Sample Survey 2012	626	101,725
Periodic Labour Force Survey 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.2 Representativeness of the vacancy data

In this section we evaluate the representativeness of our vacancy data in relation to the broader Indian labor market by benchmarking against widely-used administrative datasets and labor surveys. First, we consider the number of vacancies in the dataset. The Indian Labour Ministry estimates that India’s services sector formally employs approximately 18.9 million workers.¹ Our unrestricted sample for 2010-19 includes approximately 15.5 million vacancies in the services sector. To compare these numbers, we need to turn the flow variable of hiring into the stock of employment. Shimer (2012) documents a job separation rate of 3.4 percent for the US. If we also assume this number for India, this implies 6.4 million job ads over 10 years. This disregards

¹See the January 2022 Quarterly Employment Survey, available at <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2022/jan/doc20221104101.pdf>

the growth of India’s service sector, however, and if we assume that one quarter of job posts are replacement and three quarters are employment growth, this implies about 25.7 million job ads should be expected in the formal services sector over the course of 10 years. For this rough estimate, our coverage would be approximately 60 percent. While this back-of-the-envelope estimate is necessarily imprecise due to the lack of comprehensive worker censuses over our full sample period, we consider it nonetheless reassuring that our data are likely to reflect a substantial share of the market that we aim to understand.


We next consider the composition of the data relative to three widely used administrative datasets: firm balance sheet data from Prowess, and nationally representative labor surveys conducted in 2011-2012 (the National Sample Survey) and in 2017-2018 (the Periodic Labour Force Survey). 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 services sub-industries in the pooled National Sample Survey (NSS) and Periodic Labour Force Survey (PLFS).² The vacancy dataset has relatively fewer finance, insurance and real estate firms than Prowess, but a greater share of vacancies in that sector relative to the representative labor 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) 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. Panel (d) compares the number of firms over time in Prowess and the vacancy data. Prowess is not a firm census but is instead based on firms’ published financial reports; it significantly over-represents India’s largest firms (Goldberg et al. 2009).³ In contrast, the barrier to appearing in our vacancy dataset is low, since the process for setting up an account and posting an online job advert is simple and does not require accounting expertise or other specialized training. We therefore see substantially more firms in the vacancy data than in Prowess, albeit with the total number of firms following a similar trajectory over time.

²We 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.


³Specifically, Prowess contains longitudinal balance sheet data on all publicly listed and many large private Indian firms, as long as they submit financial statements to the government. About a third of Prowess firms are publicly listed, and even among those, some do not submit financial statements to the government (Alfaro & Chari 2009). There are many firms that both are not listed and do not submit financial statements, so do not appear in the dataset.

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


(a) Firm distribution




(b) Worker/vacancy distributions



(c) Occupation distribution



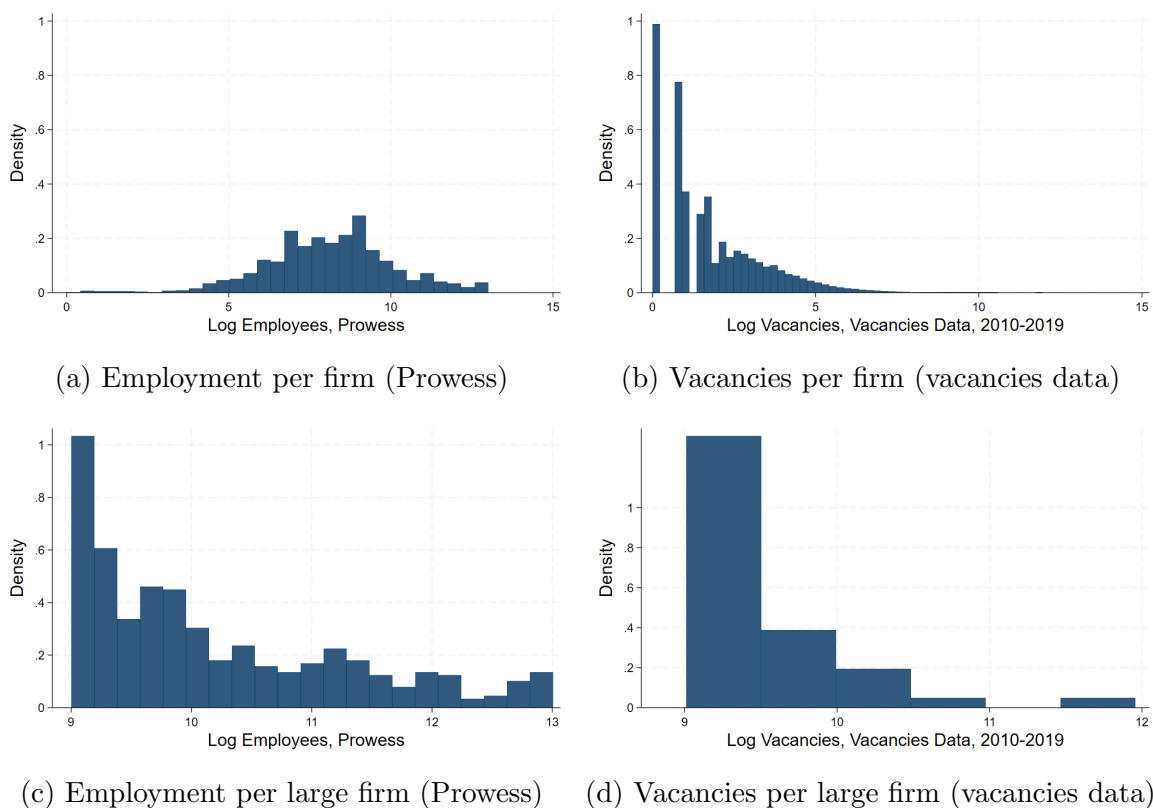
(d) Number of firms by year



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.

Finally, we examine the firm size distribution in our data by comparing the distribution of vacancies per firm to the distribution of employees per firm in Prowess. Figure A.3 shows the results; in both cases, we focus on firms active between 2010 and 2019 and in the services sector. Comparing Panels (a) and (b) illustrates that Prowess is missing the left-tail of the distribution, as expected by its focus on large firms that report financial statements. When analysing employment per firm and vacancies per firm for the subset of firms larger than the median in Prowess (about 3,500 employees), we show in Panels (c) and (d) that our vacancies dataset has a firm size distribution very similar to Prowess.

Figure A.3: Histogram of employment per firm (Prowess) and vacancies per firm (vacancies data)

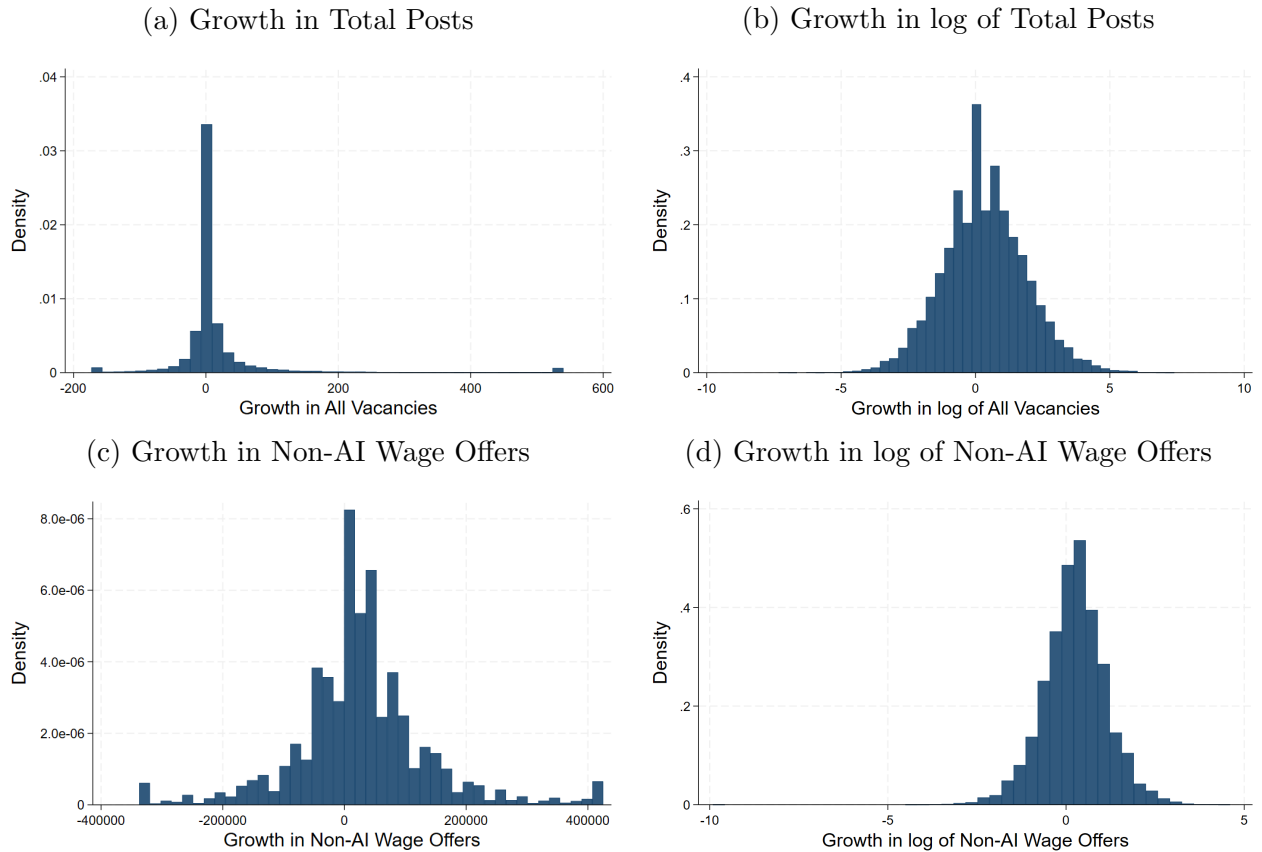


Notes: Panel (a) shows a histogram for employment per firm in Prowess, while Panel (b) shows a histogram for posted vacancies data per firm in our vacancies data. Panel (c) restricts the employment per firm in Prowess to values above the median and Panel (d) restricts vacancies per firm in our vacancies data by the same cut-off, i.e. the firm size median in Prowess. For both histograms, we restrict the sample to the services sector and to firms active in our sample period, 2010-2019.

A.3 Details on the restricted sample

This section provides further details on the ‘restricted sample’, i.e., the sample on which we run our main long-difference regressions. This sample reflects three sets of restrictions. First, as with all analysis in the paper, we focus on non-remote, full-time, permanent jobs located in India in the white-collar services sector. Second, and as with all causal analysis in the paper (i.e., excluding the descriptives in Section 3), we focus on ‘AI-using’ industries only (as opposed to also including ‘AI-producing’ industries). Finally, and in contrast to the event study in Section 6.2, we focus only on establishments that post vacancies in both the baseline (2010-2012) and endline (2017-2019). Table A.2 shows summary statistics for this sample (pooling 2010-12 and 2017-19, unless otherwise stated), while Table A.3 provides further summary statistics for differences between the baseline and endline. Lastly, Figure A.4 shows histograms of key variables and of the log of these variables.

Figure A.4: Histograms of key variables



Notes: This figure shows histograms for growth in vacancies and growth in non-AI wage offers. The left hand side shows growth in levels and the right hand side shows growth in the log of the respective variable. For growth in levels, we display the variables winsorized at the 1 percent level to exclude outliers and improve the legibility of these graphs.

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

	All Jobs	AI Jobs
# Job Posts	1,902,969	16,827
% Posts in Baseline (2010-2012)	36	2.6
# Establishments	24,485	944
% Posts by 5 Most Active Establishments	4.4	63.8
% Posts by 100 Most Active Establishments	11.2	84.9
Mean # Posts Per Establishment	78	18
Median # Posts Per Establishment	16	2
90th Percentile Posts Per Establishment	140	11
10th Percentile Posts Per Establishment	4	1
Mean # Posts per Establishment and Year (2010, 2019)	(16, 25)	(2, 10)
# Posting Establishment Per Year (2010, 2019)	(11,898, 14,998)	(37, 457)
Mean # Reposts	3.3	.6
Median # Reposts	3	0
90th Percentile Reposts	6	2
10th Percentile Reposts	1	0
# Firms	5,605	525
% Posts by 5 Most Active Firms	10	73
% Posts by 100 Most Active Firms	29	81
# Cities	442	32
% Posts in Bengaluru	16	43
% Posts in Top 5 Cities	64	92
% Posts in Top 10 Cities	85	99
# Industries (NIC 2008 features 86)	29	24
% Posts in BPO and Call Centres	25	5
% Posts in Banking and Financial Services	21	71
% Posts in Research and Analytics	2	11
# 4-digit Occupations (NCO 2004 features 96)	94	49
% Posts for Computer Programmers	2	16
% Posts for Business Professionals	22	12
% Posts for Technical and Sales Representatives	18	0.4
% Posts for Accountants	7	2
% Posts for Computer Professionals	1	11
% Posts for Computer Systems Designers and Analysts	0.4	1.3
% Posts for Research and Development Managers	0.2	3
% Requiring Undergraduate Education	11	8
% Requiring Postgraduate Education	86	88
Mean Required Work Experience, Years	1.86	2.2
Median Required Work Experience, Years	2	2.5
90th Percentile Work Experience, Years	2.5	2.5
10th Percentile Work Experience, Years	1	1
Mean Annual Salary, Rupees	167,000	389,000
Median Annual Salary, Rupees	100,000	500,000
90th Percentile Annual Salary, Rupees	325,000	500,000
10th Percentile Annual Salary, Rupees	37,500	125,000

Notes: This table shows descriptives for the restricted sample, which includes only establishments that post vacancies in both the baseline (2010-2012) and endline (2017-2019), and only posts from these years, and only posts in service sector, AI-using (not producing) industries. These are thus the vacancy posts underlying our data set for the main regressions. The number of NIC 2008 industries in our data is smaller than the total number of these industries. These are 2-digit industries, called divisions in the NIC 2008. The 86 divisions include manufacturing, agriculture, and mining, whereas we only include the services sector.

Table A.3: Difference (2017-19 over 2010-12) descriptives of vacancy data in restricted sample, all jobs and AI jobs only

Difference in	All Jobs	AI Jobs
Mean # Posts	22	.65
90th Percentile # Posts	50	0
Mean AI Share	.17%	-
Mean Annual Salary, Rupees	17,000	23,000
Median Annual Salary, Rupees	16,000	18,000
90th Percentile Annual Salary, Rupees	27,000	77,000
10th Percentile Annual Salary, Rupees	6,000	16,000
Mean Required Work Experience, Years	.15	.15
Median Required Work Experience, Years	.18	.18
Median Postgraduate Share	0	0
90th Percentile Postgraduate Share	22%	71%

Notes: Descriptive statistics overall and within AI jobs only.

A.4 Calculating the AI wage premium

This section provides details on our calculation of the AI wage premium. When including industry-region, industry-time and region-time fixed effects, we find that AI posts on average offer 30 percent higher wages than non-AI posts (see model (1) of Table A.4). 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 percent 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 percent remains.⁴

⁴The interpretation of the control variables is as follows. An extra year of experience is associated with a more than 35 percent higher salary (at least within the predominantly early-career jobs posted on the site – see Figure 3.3), while having a Master’s degree is associated with up to 10 percent higher salary. In this sample, having only a high school education is associated with wage offers 3-6 percent 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 percent 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 A.4: Wages in AI vs. non-AI roles


	log Annual Salary			
	(1)	(2)	(3)	(4)
AI post	0.318*** (0.0484)	0.198*** (0.0358)	0.128*** (0.0220)	0.174*** (0.0422)
Experience Required (Years)	0.470*** (0.00693)	0.411*** (0.00797)	0.386*** (0.00787)	0.351*** (0.00818)
High School	0.00481 (0.0788)	-0.0644*** (0.0192)	-0.0408** (0.0185)	-0.0395** (0.0172)
Master's	0.104*** (0.0144)	0.0774*** (0.00990)	0.0448*** (0.00781)	0.0198** (0.00814)
Doctorate	0.131** (0.0588)	0.0741* (0.0417)	0.0132 (0.0325)	0.00218 (0.0339)
<i>Fixed Effects:</i>				
– 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. 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 firm-specified role categories built into the jobs portal. Standard errors clustered at the firm level.

A.5 Further descriptives on AI demand and exposure


This section collects additional descriptives on AI demand and exposure. As with Section 3, these figures all cover both the AI-producing and AI-using sectors. Figure A.5 plots the cumulative share of AI posts against the corresponding share of all posts. The distribution of cities' shares of AI posts and all posts is shown in Figure A.6, while the city-wise distribution of AI posts over time is shown in Figure A.7. The top 20 roles demanding AI skills in our analysis are listed in Figure A.8. Finally, Figure A.9 displays exposure to AI by wage offer.

Figure A.5: 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.

Figure A.6: Shares of posts by city




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

Figure A.7: 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. Bengaluru's share peaked at just over 40 percent in 2014, then Mumbai's share in particular has risen subsequently as AI demand increased in finance (see Figure 3.1).

Figure A.8: 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 A.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 of 10 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 B.

Appendix B Additional Results and Robustness

This section collects additional results related to our shift-share instrument, our main specification, and various robustness checks, as referenced in the main text.

B.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):

$$Exposure_{fr,t_0} = \sum_o PostingShare_{fro}^{t_0} \cdot ExposureMeasure_o \quad (B.1)$$

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 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, examining 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 B.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 confirm 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 is constructed) on the instrument. The results are shown in Table B.2: we do not find any

Table B.1: Investigating correlates of shares

	Overall Instrument			
	(1)	(2)	(3)	(4)
Share of Highschool Education	-0.166 (0.192)	0.0203 (0.0929)	-0.0615 (0.146)	0.0253 (0.0926)
Share of Undergraduate Education	-0.232 (0.194)	0.00812 (0.0915)	-0.122 (0.146)	0.0131 (0.0912)
Share of Postgraduate Education	-0.221 (0.195)	0.0403 (0.0933)	-0.0999 (0.147)	0.0454 (0.0928)
Mean Salary	4.86e-09 (4.59e-09)	4.96e-09 (4.85e-09)	4.19e-09 (4.37e-09)	4.97e-09 (4.87e-09)
Mean Experience	-0.00217 (0.00823)	0.00524 (0.00442)	0.00334 (0.00590)	0.00512 (0.00445)
<i>Fixed Effects:</i>				
– Region		✓	✓	✓
– Industry		✓		✓
– Firm Decile			✓	✓
Observations	22,201	22,052	22,052	22,052

Notes: Standard errors in parentheses. Standard errors clustered at the firm level. The dependent variable is 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 independent variables are baseline controls.

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 B.2: Examining pre-trends for the instrument

	Growth in Non-AI Vacancies				Growth in Non-AI Median Wage			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Instrument	0.749 (2.273)	-0.492 (0.703)	-0.101 (1.885)	-0.495 (0.703)	-116.1 (611.1)	-439.5 (464.0)	-88.87 (616.1)	-438.3 (463.1)
<i>Fixed Effects:</i>								
– Region		✓	✓	✓		✓	✓	✓
– Industry		✓		✓		✓		✓
– Firm Decile			✓	✓			✓	✓
Observations	296,730	296,730	296,730	296,730	296,730	296,730	296,730	296,730

Notes: Standard errors in parentheses. Standard errors clustered at the firm level. The independent variable is 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 dependent variables are 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 B.3 presents results when instead computing standard errors according to the correction developed by Adão et al. (2019), and finds that our results are robust.

Table B.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)
Adoption of AI	-7.975** (3.654)	-12.90* (7.549)	-8.064*** (3.265)	-7.737** (3.659)	-12.47* (7.564)	-7.840*** (3.269)
<i>Fixed Effects:</i>						
– 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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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).

B.2 Additional results from main specification

This section presents additional results from our main specification. As this concerns the causal identification of the effects of AI demand on labor demand and wage offers, we are excluding AI-producing sectors from this analysis, and keep only white-collar services jobs in AI-using industries. In Table B.4, we show the first stage. Tables B.5 and B.6 repeat the main regressions on hiring and wage offers when dropping the ten largest firms. Tables B.7 and B.8 repeat the main regressions when dropping the three establishments accounting for most of the 2016-17 growth in AI vacancies. Similarly, Tables B.9 and B.10 drop the 5 establishments accounting for almost two thirds of AI job posts as per Table A.2. Tables B.11 and B.12 repeat the main regressions when excluding establishments in the finance sector. We show that this sector, which is a major user of AI in India, is not solely driving our findings.

Figure B.1 extends the wage distribution graph to all postings. Table B.13 studies the impact of AI on 2-digit occupations.⁵ Table B.14 shows task growth results for abstract and routine tasks following Autor & Dorn (2013) and confirms the findings of Table 5.3. Similarly, Table B.15 repeats Table 5.4 for abstract tasks following Autor & Dorn (2013). Table B.16

⁵The estimated effect differs significantly across occupation groups, with the hypothesis of equal coefficients rejected at the 1 percent level (p value=0.00). Differences are also generally statistically significant for pairwise comparisons at the 1 percent level, but insignificant for the comparisons of ‘General Managers’ and ‘Other Professionals’, and of ‘Health Professionals’ and ‘Teaching Professionals’.

shows wage growth results by 1-digit occupations.⁶ Finally, Table B.17 shows wage growth results when controlling for job profiles.

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

	Adoption of AI		
	(1)	(2)	(3)
Establishment AI Exposure	0.0103*** (0.00156)	0.00965*** (0.00149)	0.0106*** (0.00157)
<i>Fixed Effects:</i>			
– Region	✓	✓	✓
– Firm Decile		✓	✓
– Industry	✓		✓
R ²	.0547	.0434	.062
Observations	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 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 probability of adopting AI of a one-standard deviation rise in AI exposure according to a linear probability model.

⁶The estimated effect does not differ significantly across occupation groups (p value=0.0949).

Table B.5: Second stage: Impact of AI adoption on establishment non-AI vacancies, excluding ten largest firms

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-8.297*** (2.366)	-13.38*** (3.187)	-8.339*** (2.304)	-8.041*** (2.259)	-12.91*** (3.044)	-8.098*** (2.199)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	43.01	39.73	44.35	43.38	40.02	44.57
Observations	21,595	21,595	21,595	21,602	21,602	21,602

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 ten largest firms in terms of baseline (2010-2012) hiring are excluded.

Table B.6: Second stage: Impact of AI adoption on establishment non-AI wages, excluding ten largest firms

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-5.549*** (1.483)	-6.287*** (1.661)	-5.399*** (1.431)	-5.307*** (1.422)	-6.019*** (1.593)	-5.157*** (1.373)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	44.44	41.01	45.76	44.78	41.27	45.94
Observations	21,417	21,417	21,417	21,424	21,424	21,424

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 ten largest firms in terms of baseline (2010-2012) hiring are excluded.

Table B.7: Second stage: Impact of AI adoption on establishment non-AI vacancies, excluding the three establishments accounting for most of the 2016-17 growth in AI vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-8.205*** (2.411)	-13.23*** (3.176)	-8.295*** (2.340)	-7.958*** (2.301)	-12.78*** (3.036)	-8.062*** (2.233)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	42.7	40.61	44.43	43.05	40.87	44.62
Observations	22,241	22,241	22,241	22,248	22,248	22,248

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 three establishments accounting for most of the 2016-17 growth in AI vacancies are excluded.

Table B.8: Second stage: Impact of AI adoption on establishment non-AI wages, excluding the three establishments accounting for most of the 2016-17 growth in AI vacancies

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-5.854*** (1.522)	-6.513*** (1.671)	-5.664*** (1.458)	-5.599*** (1.459)	-6.239*** (1.604)	-5.411*** (1.398)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	44.13	41.93	45.85	44.45	42.15	46
Observations	22,061	22,061	22,061	22,068	22,068	22,068

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 three establishments accounting for most of the 2016-17 growth in AI vacancies are excluded.

Table B.9: Second stage: Impact of AI adoption on establishment non-AI vacancies, excluding the 5 establishments contributing two thirds of AI job posts

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-8.233*** (2.415)	-13.28*** (3.187)	-8.327*** (2.344)	-7.985*** (2.305)	-12.83*** (3.046)	-8.093*** (2.237)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	42.6	40.5	44.32	42.95	40.75	44.51
Observations	22,239	22,239	22,239	22,246	22,246	22,246

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 5 establishments contributing two thirds of AI job posts as per Appendix Table A.2 have been excluded.

Table B.10: Second stage: Impact of AI adoption on establishment non-AI wages, excluding the 5 establishments contributing two thirds of AI job posts

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-5.864*** (1.525)	-6.531*** (1.676)	-5.674*** (1.461)	-5.608*** (1.462)	-6.256*** (1.609)	-5.420*** (1.401)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	44.02	41.81	45.74	44.35	42.03	45.9
Observations	22,059	22,059	22,059	22,066	22,066	22,066

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 5 establishments contributing two thirds of AI job posts as per Appendix Table A.2 have been excluded.

Table B.11: Second stage: Impact of AI adoption on establishment non-AI vacancies, excluding finance sector

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-14.42*** (4.008)	-17.58*** (4.233)	-14.47*** (3.950)	-14.65*** (4.112)	-17.85*** (4.349)	-14.70*** (4.056)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	28.19	32.57	28.82	27.27	31.49	27.81
Observations	16,558	16,558	16,558	16,561	16,561	16,561

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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). Establishments in the finance sector have been excluded.

Table B.12: Second stage: Impact of AI adoption on establishment non-AI wages, excluding finance sector

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-8.346*** (2.436)	-8.257*** (2.285)	-8.217*** (2.377)	-8.477*** (2.496)	-8.378*** (2.339)	-8.344*** (2.437)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	28.32	32.47	28.9	27.39	31.39	27.9
Observations	16,421	16,421	16,421	16,424	16,424	16,424

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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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). Establishments in the finance sector have been excluded.

Table B.13: Second stage: Impact of AI adoption on establishment non-AI vacancy shares, by granular occupation group

	Change in Non-AI Vacancy Shares					
	Engineering Professionals	Health Professionals	Teaching Professionals	Other Professionals	Corporate Managers	General Managers
Adoption of AI	-2.689*** (0.494)	0.130 (0.120)	0.212*** (0.0748)	-1.290*** (0.409)	-9.964*** (1.589)	-0.626** (0.299)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓	✓
First Stage F-Stat	45.43	45.43	45.43	45.43	45.43	45.43
Observations	22,244	22,244	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 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 occupation shares. Occupation groups are 2-digit occupations within Professionals and Managers from the NCO04. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.14: Second stage: Impact of AI adoption on establishment abstract and routine tasks

	Growth in Abstract Tasks			Growth in Routine Tasks		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-5.753*** (1.094)	-7.208*** (1.348)	-5.772*** (1.073)	-0.219 (0.214)	-0.193 (0.193)	-0.217 (0.213)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	44.06	41.83	45.62	44.06	41.83	45.62
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 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, where we have shifted the mean of the distributions of scores to be able to take logarithms. 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 standardized as in Acemoglu & Autor (2011). AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.15: Second stage: Impact of AI adoption on establishment abstract tasks, by occupation group

	Growth in Abstract Tasks				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Adoption of AI	-0.0477** (0.0225)	-0.00733 (0.0121)	0.296*** (0.114)	-1.651** (0.780)	-29.15*** (4.863)
<i>Fixed Effects:</i>					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	45.62	45.62	45.62	45.62	45.62
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 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, where we have shifted the mean of the distributions of scores to be able to take logarithms. 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 standardized as in Acemoglu & Autor (2011). AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.16: Second stage: Impact of AI adoption on establishment non-AI wages, estimated separately by occupation group

	Growth in Non-AI Median Wage				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Adoption of AI	0.716 (2.060)	1.816* (1.067)	-0.840 (0.729)	-1.073 (0.942)	-2.037* (1.194)
<i>Fixed Effects:</i>					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	6.443	11.67	61.28	23.57	25.89
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 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. Occupation groups are 1-digit occupation groups from the NCO04. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.17: 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)	(5)	(6)
Adoption of AI	-4.147*** (1.191)	-4.347*** (1.278)	-4.049*** (1.152)	-3.997*** (1.144)	-4.196*** (1.229)	-3.898*** (1.107)
Growth in Experience	0.409*** (0.0132)	0.409*** (0.0138)	0.405*** (0.0133)	0.409*** (0.0132)	0.409*** (0.0136)	0.405*** (0.0132)
Growth in High School share	-0.0883 (0.0724)	-0.0896 (0.0740)	-0.102 (0.0716)	-0.0919 (0.0705)	-0.0928 (0.0720)	-0.105 (0.0697)
Growth in Master's share	0.222*** (0.0320)	0.220*** (0.0318)	0.222*** (0.0318)	0.221*** (0.0320)	0.219*** (0.0318)	0.222*** (0.0317)
Growth in Doctorate share	2.090** (0.988)	2.283** (1.085)	2.045** (0.968)	2.029** (0.957)	2.219** (1.052)	1.985** (0.937)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	45.86	43.58	47.26	46.13	43.73	47.38
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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.1: Impact of AI adoption on the wage offer distribution in all posts



Notes: This coefficient plot shows the impact of a one percent rise in the predicted probability of adopting AI on wage growth over time across the distribution of establishment wage offers. Each coefficient is from a regression of type (6) in Appendix Table B.17. As in Appendix Table B.17, AI adoption is instrumented with AI exposure. In other words, each coefficient represents the impact of a one percent rise in the predicted probability of adopting AI 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 percent confidence intervals. Standard errors are clustered at the firm level, and we include region, firm size 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.

B.3 Alternative specifications

This section provides the key results repeated for a series of alternative specifications, as discussed in Section 6, along with additional results for the event-study approach.

Table B.18: 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)
Adoption of AI	-9.656*** (2.944)	-16.10*** (4.206)	-9.582*** (2.865)	-9.353*** (2.801)	-15.51*** (3.997)	-9.293*** (2.725)
<i>Covariates:</i>						
Share of Software Engineers	✓	✓	✓	✓	✓	✓
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	32.3	28.77	33.05	33.01	29.4	33.69
Observations	22,244	22,244	22,244	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 log 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.). AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.19: 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)
Adoption of AI	-6.620*** (1.833)	-7.642*** (2.121)	-6.482*** (1.776)	-6.326*** (1.749)	-7.307*** (2.024)	-6.186*** (1.695)
<i>Covariates:</i>						
Share of Software Engineers	✓	✓	✓	✓	✓	✓
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	33.58	29.9	34.32	34.27	30.52	34.94
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 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. 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.). AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.20: 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)
Adoption of AI	-17.96** (6.976)	-40.93** (20.67)	-18.08*** (6.916)	-16.82*** (6.335)	-36.83** (17.49)	-16.96*** (6.270)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	10.49	4.644	10.69	11.17	5.282	11.35
Observations	22,244	22,244	22,244	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 log 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.21: 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)
Adoption of AI	-6.393** (3.005)	-10.85* (6.168)	-6.318** (2.933)	-6.393** (3.005)	-10.85* (6.168)	-6.318** (2.933)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	10.48	4.705	10.76	10.48	4.705	10.76
Observations	22,064	22,064	22,064	22,064	22,064	22,064

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. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.22: 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)	(5)	(6)
Adoption of AI	-6.735** (2.878)	-10.40*** (3.404)	-7.029** (2.782)	-6.682** (2.916)	-10.38*** (3.445)	-6.989** (2.816)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	41.13	45.05	44.32	39.45	43.08	42.62
Observations	38,458	38,458	38,458	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 a dummy, which equals 1 for establishments that did not post AI vacancies in the 2013-15 period, but posted AI vacancies in the endline period. The dependent variables are the change in the log of the respective establishment-level outcomes. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.23: 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)	(5)	(6)
Adoption of AI	-7.114*** (2.002)	-6.920*** (2.015)	-6.829*** (1.911)	-7.185*** (2.038)	-6.966*** (2.045)	-6.883*** (1.940)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	41.17	45.17	44.28	39.49	43.19	42.57
Observations	38,249	38,249	38,249	38,281	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 a dummy, which equals 1 for establishments that did not post AI vacancies in the 2013-15 period, but posted AI vacancies in the endline period. The dependent variables are the change in the log of the respective establishment-level outcomes. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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).

Figure B.2: Non-AI vacancies following AI adoption at the district level



Notes: Two way fixed effects on a balanced panel. The outcome variable is the IHS-transformed number of non-AI vacancies 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 vacancies on the region level, but has no effects on the district level.

Table B.24: Vacancies results for ‘incumbents’ and ‘startups’

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI Incumbents	-5.420*** (2.065)	-6.204*** (2.171)	-5.162*** (1.938)	-5.392*** (2.087)	-6.202*** (2.200)	-5.147*** (1.960)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	34.59	35.31	38.14	33.72	34.22	37.11
Observations	17,341	17,341	17,341	17,348	17,348	17,348
Adoption of AI Start-ups	-11.20 (8.499)	-21.30* (11.96)	-12.30 (8.430)	-11.17 (8.689)	-21.26* (12.15)	-12.29 (8.605)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	9.133	9.889	9.725	8.54	9.245	9.138
Observations	21,060	21,060	21,060	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 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. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.25: Wage results for ‘incumbents’ and ‘startups’

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI Incumbent	-3.853*** (1.250)	-3.596*** (1.189)	-3.650*** (1.163)	-3.913*** (1.267)	-3.655*** (1.207)	-3.703*** (1.178)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	34.71	35.28	38.23	33.84	34.19	37.19
Observations	17,259	17,259	17,259	17,266	17,266	17,266
Adoption of AI Start-ups	-16.27** (7.103)	-17.74** (7.778)	-15.81** (6.760)	-16.48** (7.347)	-17.74** (7.949)	-15.97** (6.966)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	9.135	10.08	9.757	8.538	9.432	9.161
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 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. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.26: Second stage: Impact of AI adoption on establishment non-AI vacancies, intensive margin of AI demand

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth of AI Vacancies	-4.611*** (1.557)	-7.704*** (2.187)	-4.632*** (1.517)	-4.555*** (1.513)	-7.564*** (2.115)	-4.584*** (1.475)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	22.76	23.26	23.74	23.9	24.34	24.9
Observations	22,244	22,244	22,244	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 logarithm plus 1. The dependent variables are the change in the log of the respective establishment-level outcomes. AI vacancies growth is instrumented by establishment AI exposure. AI exposure 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.27: Second stage: Impact of AI dummy on establishment non-AI wages, intensive margin of AI demand

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth of AI Vacancies	-3.302*** (0.997)	-3.803*** (1.119)	-3.174*** (0.946)	-3.218*** (0.962)	-3.703*** (1.079)	-3.090*** (0.912)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	23.22	23.71	24.18	24.37	24.8	25.36
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 logarithm plus 1. The dependent variables are the change in the log of the respective establishment-level outcomes. AI vacancies growth is instrumented by establishment AI exposure. AI exposure 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.28: Second stage: Impact of AI adoption on establishment non-AI vacancies, IHS of AI vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-3.777*** (1.252)	-6.285*** (1.745)	-3.794*** (1.219)	-3.727*** (1.215)	-6.165*** (1.686)	-3.751*** (1.184)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	24.79	25.12	25.87	26.06	26.31	27.17
Observations	22,244	22,244	22,244	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 log. 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 on the outcome variable of a one IHS unit increase in establishment AI demand. 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.29: Second stage: Impact of AI adoption on establishment non-AI vacancies, IHS of AI vacancies

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in AI Vacancies	-2.703*** (0.799)	-3.101*** (0.895)	-2.599*** (0.758)	-2.632*** (0.770)	-3.017*** (0.862)	-2.527*** (0.730)
<i>Fixed Effects:</i>						
– 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 log. 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 on the outcome variable of a one IHS unit increase in establishment AI demand. 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.30: Second stage: Impact of AI adoption on establishment non-AI vacancies, weighted

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-2.907** (1.298)	-4.142*** (1.437)	-2.894** (1.264)	-2.846** (1.279)	-4.077*** (1.419)	-2.839** (1.246)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	49.54	48.42	51.54	50.7	49.31	52.73
Observations	22,244	22,244	22,244	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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.31: Second stage: Impact of AI adoption on establishment non-AI wages, weighted

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-1.135** (0.491)	-1.382*** (0.481)	-1.088** (0.474)	-1.111** (0.484)	-1.364*** (0.475)	-1.064** (0.466)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	49.58	48.36	51.54	50.74	49.26	52.74
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. Establishments are weighted by baseline number of posts, with the top 5% winsorized. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.32: Second stage: Impact of AI adoption on establishment non-AI vacancies, two-step estimation method

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-2.567*** (0.831)	-5.112*** (0.929)	-3.072*** (0.830)	-2.728*** (0.790)	-5.108*** (0.890)	-3.258*** (0.801)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat						
Observations	22,392	22,392	22,392	22,399	22,399	22,399

Notes: Standard errors in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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). We use the two-step estimation method implemented using etregress in Stata.

Table B.33: Second stage: Impact of AI adoption on establishment non-AI wages, two-step estimation method

	Growth in Non-AI Median Wage			Growth in Overall Median Wage		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	-3.381*** (0.520)	-3.540*** (0.523)	-3.363*** (0.518)	-3.211*** (0.489)	-3.353*** (0.492)	-3.183*** (0.487)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat						
Observations	22,214	22,214	22,214	22,221	22,221	22,221

Notes: Standard errors in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01. 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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). We use the two-step estimation method implemented using etregress in Stata.

Table B.34: Second stage: 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)	(5)	(6)
Adoption of AI	-5.703*** (1.403)	-6.626*** (1.580)	-5.530*** (1.345)	-5.450*** (1.343)	-6.344*** (1.516)	-5.281*** (1.289)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	45.14	42.91	46.86	45.46	43.12	47.02
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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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.35: OLS: Regression of AI adoption on establishment non-AI vacancies

	Growth in Non-AI Vacancies			Growth in Total Vacancies		
	(1)	(2)	(3)	(4)	(5)	(6)
Adoption of AI	1.021*** (0.0882)	1.253*** (0.0852)	1.160*** (0.0856)	1.053*** (0.0874)	1.283*** (0.0842)	1.189*** (0.0848)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat						
Observations	24,312	24,312	24,312	24,322	24,322	24,322

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.

Table B.36: Probit regression on establishment AI adoption

	AI adoption
Firm Size Decile	-0.0675*** (0.0121)
Firm Age	0.0367*** (0.00647)
Total Posts	0.259*** (0.0153)
Postgraduate Share	9.895*** (3.067)
Median Salary	-10.75* (5.775)
Median Experience	-0.556*** (0.0624)
Growth in Median Salary	-0.100*** (0.0285)
Growth in Median Experience	0.457*** (0.141)
90th Percentile of Salary	0.118** (0.0589)
90th Percentile of Experience	-0.308*** (0.0999)
Growth of 90th Percentile of Salary	0.0591 (0.0394)
Growth of 90th Percentile of Experience	0.0819 (0.0741)
99th Percentile of Salary	0.429*** (0.0497)
99th Percentile of Experience	-2.024*** (0.626)
Growth of 99th Percentile of Salary	-0.188*** (0.0330)
Growth of 99th Percentile of Experience	-0.159** (0.0639)
Salary Dispersion	-0.00000106*** (0.000000231)
Experience Dispersion	0.170*** (0.0493)
<i>N</i>	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.4 Alternative exposure measures

For our main specifications, we use the AI exposure measure proposed by Webb (2020), which estimates the overlap between occupations’ constituent tasks and the capabilities described in AI patents. The advantage of this approach is that it is relatively objective – relying on no ‘ad hoc’ expert judgements of the capabilities of particular technologies – and flexible, allowing Webb to validate it against previous developments in software and robots. Nonetheless, we also explore alternative AI exposure measures that have been proposed in the literature.

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) then crowdsource assessments on the applicability of these applications to 52 O*NET ability scales using Amazon MTurk. Their measure 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) and repeat our main specifications. The results are shown in Tables B.37 to B.39. While we do not observe any significant impacts on the growth of non-AI vacancies for the full sample, we do find similar negative effects on the share of professional and managerial vacancies and on median wage offers.

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). We use an India-specific version of the SML index created by Mani et al. (2020), who interviewed a sample of Indian employees using the SML rubric and mapped a SML score onto every occupation in the 2004 NCO at the four-digit level. However, unlike the Webb or Felten et al. measures, we find that the SML exposure measure fails to positively predict establishment AI adoption (Table B.40). This may reflect that the SML measure is more forward-looking in its predictions. The 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.

Table B.37: 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)	(5)	(6)
Adoption of AI	0.784 (2.541)	2.315 (1.682)	1.508 (2.569)	0.966 (2.669)	2.496 (1.739)	1.750 (2.709)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	23.91	42.65	22.6	19.53	36.89	18.34
Observations	22,244	22,244	22,244	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 log of the respective establishment-level outcomes. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.38: Second stage: Impact of AI adoption on establishment non-AI vacancy shares, by occupation group – Felten et al. exposure measure

	Growth in Non-AI Vacancies				
	Personal, Sales & Security	Clerks	Associate Professionals	Professionals	Managers
Adoption of AI	7.981*** (1.809)	-0.919*** (0.289)	7.608*** (1.780)	-10.31*** (2.261)	-4.039*** (1.002)
<i>Fixed Effects:</i>					
– Region	✓	✓	✓	✓	✓
– Industry	✓	✓	✓	✓	✓
– Firm Decile	✓	✓	✓	✓	✓
First Stage F-Stat	22.6	22.6	22.6	22.6	22.6
Observations	22,244	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 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 occupation shares. Occupation groups are 1-digit occupations from the NCO04. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.39: 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)	(5)	(6)
Adoption of AI	-3.300** (1.507)	-2.421** (0.958)	-3.480** (1.560)	-3.551** (1.662)	-2.534** (1.020)	-3.753** (1.729)
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Industry	✓		✓	✓		✓
– Firm Decile		✓	✓		✓	✓
First Stage F-Stat	23.71	42.77	22.52	19.35	36.99	18.27
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 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. AI adoption is instrumented by establishment AI exposure. Each coefficient therefore represents the percentage point impact upon the outcome variable of a one percent rise in the predicted probability of adopting AI. AI exposure 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 B.40: First stage: Impact of AI exposure on establishment AI adoption – alternative exposure measures

	Adoption of AI					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Exposure	0.00704*** (0.00159)	0.0103*** (0.00170)	0.00679*** (0.00159)	-0.00509*** (0.00144)	-0.00922*** (0.00147)	-0.00581*** (0.00147)
Exposure Measure	Felten et al.	Felten et al.	Felten et al.	SML	SML	SML
<i>Fixed Effects:</i>						
– Region	✓	✓	✓	✓	✓	✓
– Firm Decile		✓	✓		✓	✓
– Industry	✓		✓	✓		✓
R ²	.0536	.0439	.0607	.0531	.0433	.0604
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 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 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 impact on the probability of adopting AI of a one-standard deviation rise in AI exposure according to a linear probability model.