

Bonfiglioli, Alessandra; Crinò, Rosario; Gancia, Gino; Papadakis, Ioannis

Working Paper

Artificial Intelligence and Jobs: Evidence from US Commuting Zones

CESifo Working Paper, No. 10685

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Bonfiglioli, Alessandra; Crinò, Rosario; Gancia, Gino; Papadakis, Ioannis (2023) : Artificial Intelligence and Jobs: Evidence from US Commuting Zones, CESifo Working Paper, No. 10685, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/282373>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Artificial Intelligence and Jobs: Evidence from US Commuting Zones

Alessandra Bonfiglioli, Rosario Crinò, Gino Gancia, Ioannis Papadakis

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Artificial Intelligence and Jobs: Evidence from US Commuting Zones

Abstract

We study the effect of Artificial Intelligence (AI) on employment across US commuting zones over the period 2000-2020. A simple model shows that AI can automate jobs or complement workers, and illustrates how to estimate its effect by exploiting variation in a novel measure of local exposure to AI: job growth in AI-related professions built from detailed occupational data. Using a shift-share instrument that combines industry-level AI adoption with local industry employment, we estimate robust negative effects of AI exposure on employment across commuting zones and time. We find that AI's impact is different from other capital and technologies, and that it works through services more than manufacturing. Moreover, the employment effect is especially negative for low-skill and production workers, while it turns positive for workers at the top of the wage distribution. These results are consistent with the view that AI has contributed to the automation of jobs and to widen inequality.

JEL-Codes: J230, J240, O330.

Keywords: artificial intelligence, automation, displacement, labor.

Alessandra Bonfiglioli
University of Bergamo / Italy
alessandra.bonfiglioli@unibg.it

Gino Gancia
University of Milan Bicocca / Italy
ginoalessandro.ganci@unimib.it

Rosario Crinò
University of Bergamo / Italy
rosario.cрино@unibg.it

Ioannis Papadakis
University of Sussex / United Kingdom
I.Papadakis@sussex.ac.uk

September 2023

We thank Daron Acemoglu for useful comments. The usual caveat applies.

1 INTRODUCTION

Artificial Intelligence (AI) is often viewed as one of the most transformative and disruptive technologies of recent times. Thanks to improvements in machine learning techniques and the growing availability of vast amounts of digital data, the last two decades have witnessed a tremendous increase in the use of AI applications, which include web search engines, targeted advertising, recommendation systems, self-driving cars, generative or creative tools and chat-bots. A pressing policy question is how these advances will affect labor markets and especially employment.

On the one hand, intelligent tools promise to enhance human capabilities and create new demand for certain skills. For instance, AI can substitute existing capital (Bresnahan, 2019) and boost labor productivity (McKinsey Global Institute, 2017). On the other hand, the fear is that AI will surpass workers in a growing set of tasks, making them redundant. For instance, AI-assisted machines can be used to automate jobs (Acemoglu, 2022, Acemoglu and Johnson, 2023). Whether AI will complement or substitute workers is therefore an empirical question, for which there is still very little systematic evidence.

This paper is one of the first attempts at filling this gap by studying the effect of AI on employment across US Commuting Zones (CZs) over the period 2000-2020. Our analysis covers the years of the rise of the digital economy. All the major companies involved in big data collection were founded by the early 2000s: Amazon (1994), Google (1998), LinkedIn (2003), Facebook (2004), Twitter (2006). There is also evidence that the diffusion of AI technologies speeded up after 2010 (e.g., Taddy, 2019, Alekseeva et al., 2021, and Acemoglu et al. 2022). Yet, precise metrics for quantifying AI adoption are at present lacking. To overcome this limitation, we recognize that using AI necessitates skills that only professionals in a narrow range of occupations possess. Taking advantage of a novel section of the O*NET database, we classify AI-related occupations as those whose job postings most frequently require specialized software used for machine learning and data analysis. Then, we detect AI adoption from the growth in the relative importance of these AI-related occupations.

A second challenge in identifying causal effects is that AI adoption might be correlated with other shocks hitting a CZ. To overcome this problem, we use a shift-share instrument that combines industry-level AI adoption for the US with CZ-level employment shares across industries. Then, guided by a simple model, we identify CZs more exposed to AI as those specialized in industries that experienced faster growth in AI-related occupations.

With this data at hand, we first document some patterns about the change in the employment share of AI-related occupations across industries, CZs and time. This preliminary

analysis confirms that AI-adoption varies significantly and that it accelerated after 2010. We then estimate the effect of AI-adoption on employment using 2SLS stacked first-differences models for the decades 2000-2010 and 2010-2020. To control for other characteristics of a CZ that may influence AI adoption and employment, we include a wealth of fixed effects and covariates. We also control for other technologies such as ICT and industrial robots, and follow various approaches to account for underlying trends and unobserved shocks. In all cases, we estimate robust negative effects of AI exposure on employment across CZs and time. Interestingly, we find that the 2SLS coefficient is more strongly negative than its OLS counterpart, which is consistent with the view that contemporaneous shocks may induce an upward bias.¹

Finally, we dig deeper into the effect of AI adoption. We start by comparing AI with other shocks studied in the literature. Then, we explore the mechanisms through which the effect of AI unfolds. Differently from industrial robots, we find that AI adoption is concentrated in the service sector, but it exerts a negative effect on employment both in services and in manufacturing. Finally, we study how the employment response to AI varies by gender, age, skill and occupation. We show that the negative employment effects are largest for low-skill and production workers, while they turn strongly positive for workers at the top of the wage distribution. This suggests that AI adoption may have contributed to increasing inequality.

This paper is related to the large and growing literature on the labor market effects of new technologies and especially automation. Particularly close is the seminal paper by Acemoglu and Restrepo (2020), which studies the effects of industrial robots, from the International Federation of Robotics (IFR), on US CZs. We differ in several important respects. First, we focus on an entirely different technology, the adoption of AI, using a novel measure. Second, our measure of AI adoption has several advantages compared to their proxy for automation. The IFR data, first used by Graetz and Michaels (2018), is available for 19 broad sectors only, while our variable has a much finer variation (188 industries). Moreover, the IFR data are not available at the CZ level, which makes it impossible to test the mechanisms through which US-level exposure operates across localities. In terms of results, we find evidence that AI adoption, similarly to robots, displaces workers. However, we also find that AI adoption, unlike robots, operates mostly through the service sector.

A recent strand of the literature studies the evolution of occupations that are at risk of being displaced by AI. Various measures have been proposed: the AI occupational impact measure by Felten et al. (2018, 2019, 2021); the Suitability for Machine Learning (SML) index

¹For instance, Bonfiglioli et al. (2022a) argue that positive demand shocks trigger investment in automation and also increase employment.

by Brynjolfsson et al. (2018, 2019); and Webb’s (2020) AI Exposure score. Combining these measures with information from job postings in the US, Acemoglu et al. (2022) find that, consistently with our approach to measuring AI adoption, establishments with AI-suitable tasks increase their demand for AI-related skills and reduce their overall hiring. Albanesi et al. (2023) study instead the evolution of occupations exposed to AI in a panel of European countries. None of these papers detect a negative relationship between displacement risk and employment at the occupation or industry level, suggesting that existing measures of displacement risk might be correlated with other shocks. Our paper is the first to identify negative aggregate effects of AI adoption on employment. Our results are consistent with Hui, Reshef and Zhou (2023), who document a negative impact of generative AI on the employment of free-lancers in an online labor market, and with Grennan and Michaely (2020), who find negative effects for financial analysts.

While most of the literature focuses on occupations that compete with AI, a few recent papers study instead the jobs that are involved in the creation of AI. Hanson (2021) selects AI-related occupations based on keywords such as "computer", "data" or "software" in job titles, while we use specific software requirements included in job postings. More importantly, he uses his classification to study a very different question, namely, the determinants of regional specialization in AI-related activities and not the labor market effects of AI. Alekseeva et al. (2021), instead, provide some descriptive evidence on the demand for AI skills in the US across occupations and industries. Finally, relative to using vacancies to measure the demand for AI-related skills (as in Acemoglu et al., 2022), occupational data have the advantage of being available for a broader sample.

The remainder of the paper is organized as follows. Section 2 presents a simple model of the effects of AI on employment to guide the empirical analysis. Section 3 describes the data and the main variables. Section 4 discusses the empirical specification and the identification strategy. Section 5 presents some stylized facts and preliminary evidence. Section 6 contains the main empirical results. Section 7 considers possible threats to identification. Section 8 compares the results to other technologies and explores the adjustment mechanisms. Section 9 concludes.

2 AI AND LOCAL LABOR DEMAND

This section presents a simple partial-equilibrium, task-based, model of the effects of AI on the demand for labor across CZs similar to Acemoglu and Restrepo (2020). The role of the model is to show that AI can automate jobs or complement workers, to illustrate how

to estimate its effect using variation in local exposure to AI and to clarify how to measure it. The analysis is kept deliberately simple to make the link with the empirical section as transparent as possible. Partial equilibrium is adopted both for simplicity and for consistency with the econometric specifications.

2.1 A SIMPLE MODEL OF AI AND EMPLOYMENT

The economy consists of C commuting zones. Each CZ $c \in C$ has identical preferences over I industries. For simplicity, trade is free across CZs and we denote with p_i the price of the output of industry $i \in I$. Each industry produces output by combining a specific capital with a continuum of tasks indexed by $z \in [0, 1]$, each of which can be performed by AI or labor. The production function for industry i in CZ c is:

$$y_{ci} = \varphi_{ci} \left[\exp \left(\int_0^1 \ln x_{ci}(z) dz \right) \right]^\alpha K_{ci}^{1-\alpha}, \quad (1)$$

where $x_{ci}(z)$ is the output of task z and K_{ci} is CZ c 's endowment of the specific capital used in industry i . Differences in the endowment of specific capital generate differences in the industrial composition of employment across CZs.

Denote with ω_c and w_c the cost of AI and the wage in CZ c . We assume that $\omega_c < w_c$. We identify the capabilities of AI with the set of tasks it can perform. Workers can perform all tasks. However, since AI is cheaper than labor, workers will not be employed in tasks that can be performed by AI. Assuming that AI can be used in the set of tasks $[0, \kappa_{ic}]$, we will have $x_{ci}(z) = A_{ci}/\kappa_{ci}$, where A_{ci} is the AI input used in CZ c in industry i . Labor performs the remaining tasks $[\kappa_{ic}, 1]$, with $x_{ci}(z) = L_{ci}/(1 - \kappa_{ci})$, where L_{ci} is employment in CZ c in industry i . Substituting these quantities into (1), industry output becomes:

$$y_{ci} = \varphi_{ci} \left(\frac{A_{ci}}{\kappa_{ci}} \right)^{\alpha \kappa_{ci}} \left(\frac{L_{ci}}{1 - \kappa_{ci}} \right)^{\alpha(1-\kappa_{ci})} K_{ci}^{1-\alpha}. \quad (2)$$

The demand for AI services and for labor from industry i in CZ c are:

$$\omega_c A_{ci} = \alpha \kappa_{ci} p_i y_{ci}, \quad (3)$$

and

$$w_c L_{ci} = \alpha (1 - \kappa_{ci}) p_i y_{ci}. \quad (4)$$

Substituting (3)-(4) into (2) yields:

$$y_{ci} = (\varphi_{ci}\alpha)^{\frac{1}{1-\alpha}} p_i^{\frac{\alpha}{1-\alpha}} \left(\frac{1}{\omega_c}\right)^{\frac{\alpha\kappa_{ci}}{1-\alpha}} \left(\frac{1}{w_c}\right)^{\frac{\alpha(1-\kappa_{ci})}{1-\alpha}} K_{ci}.$$

We define an improvement of the AI technology as an increase in the set of tasks it can perform: $d\kappa_{ci} > 0$. We also allow technological progress in AI to boost productivity of all factors: $d\varphi_{ci}/d\kappa_{ci} = \gamma \geq 0$. This assumption may capture phenomena such as the use of AI as a general purpose technology complementing all existing factors or the creation of new tasks. These are mechanisms that have been emphasized in part of the literature. Since AI services are cheaper than labor and may raise productivity, industry output necessarily expands with $d\kappa_{ci} > 0$:

$$d \ln y_i = \left(\frac{\alpha \ln \pi_c}{1 - \alpha} + \frac{\gamma}{1 - \alpha} \right) d\kappa_{ci}, \quad (5)$$

where $\pi_c \equiv w_c/\omega_c > 1$ is the cost saving of AI relative to labor.

The partial-equilibrium effects on labor demand follow from differentiating (4):

$$d \ln(w_c L_{ci}) = -\frac{d\kappa_{ci}}{1 - \kappa_{ci}} + d \ln(p_i y_{ci}). \quad (6)$$

This equation illustrates the two key possible effects of AI. The first term captures the negative effect on labor demand when AI displaces workers in some tasks previously performed by humans. The second term is the increase in labor demand when AI raises productivity and hence total revenue of the industry. Eq. (6) clarifies that the displacement effect of AI is the only one responsible for any negative impact on labor demand.

Aggregating the industry-level implications yields the effect on local labor demand:

$$d \ln(w_c L_c) = -\sum_{i \in I} \frac{L_{ci}}{L_c} \frac{d\kappa_{ci}}{1 - \kappa_{ci}} + \sum_{i \in I} \frac{L_{ci}}{L_c} d \ln(p_i y_{ci}). \quad (7)$$

2.2 AI AND LOCAL LABOR DEMAND: EMPIRICAL SPECIFICATION

Equations (5) and (7) summarize the effects of advances in AI on labor demand, through the displacement of workers and the increase in productivity. In order to derive an estimation equation for employment in the most transparent way, we now make some additional simplifying assumptions. First, we assume that the supply of labor is fully elastic and set the wage as the numeraire, $w_c = 1$. For instance, this would be the case if labor is freely mobile across CZs or if there is a binding wage floor. We discuss later what happens when labor

supply is not fully elastic. Second, we assume that AI services are also in fully elastic supply. In particular, the AI input is a bundle of intermediate inputs I_{ic} (such as computing power and data) and specialized workers H_{ci} (e.g., programmers and data scientists), which must be combined in fixed proportions:

$$A_{ci} = \min \{I_{ic}, \nu H_{ci}\}.$$

Third, as in Acemoglu and Restrepo (2020), we consider an initial equilibrium in which $\kappa_{ci} \approx 0$. This assumption is appropriate to study the adoption of an entirely new technology such as AI.

Under these assumptions, we can detect technological progress in AI from the increase in AI-specific workers. Differentiating (3) and rearranging yields:

$$d\kappa_{ci} = \frac{1}{\pi} \frac{dA_{ci}}{L_{ci}} = \frac{\nu}{\pi} \frac{dH_{ci}}{L_{ci}}. \quad (8)$$

This equation shows that technological progress in AI in industry i can be measured by the increase in AI employment over raw labor. Finally, using $\kappa_{ci} \approx 0$, $w_c = 1$, (5) and (8) into (7) yields:

$$\frac{dL_c}{L_c} = \frac{\nu}{\pi} \left(\frac{\alpha \ln \pi}{1 - \alpha} + \frac{\gamma}{1 - \alpha} - 1 \right) \sum_{i \in I} \frac{L_{ci}}{L_c} \frac{dH_{ci}}{L_{ci}}. \quad (9)$$

Equation (9) shows that the effect of AI can be estimated by regressing changes in employment on changes in AI-related jobs at the CZ level.

However, the main concern is that shocks to AI employment in a CZ might be correlated with other local shocks that have a direct effect on employment. For instance, an increase in local demand may trigger AI adoption and simultaneously raise labor demand. Ideally, we want to use changes in AI technology that are exogenous to other labor market shocks in CZ c . To do so, we adopt a classic Bartik design and instrument dH_{ci}/L_{ci} with its national counterpart, dH_i/L_i , where dH_i is the change in AI-related jobs in industry i in the US. We defer a detailed discussion of the identification strategy and the possible threats to identification to Section 4.

Before proceeding, we pause to briefly discuss what happens when labor supply is not fully elastic. In this case, employment in CZ c must satisfy the market-clearing condition $\sum_{i \in I} L_{ci} = L_c$. Adding a labor supply equation of the form $L_c = \eta w_c^\phi$, it is possible to solve simultaneously for L_c and w_c . An increase (decrease) in labor demand will then translate into an increase (decrease) in both employment and wages.

3 DATA AND VARIABLES

Our sample consists of 722 CZs covering the entire mainland of the US.² The time period of our analysis spans the last two decades, given that the surge of AI is a recent phenomenon. Specifically, we observe each CZ at the endpoints of each decade, i.e., in the years 2000, 2010 and 2020. We now present the data sources and explain the construction of the main variables. Descriptive statistics are provided in Appendix Table B1.

3.1 EMPLOYMENT, POPULATION AND OTHER CHARACTERISTICS OF CZS

For each CZ, we measure employment, population and other characteristics using micro-level data from two sources: the decennial Census for the year 2000 and the American Community Survey (ACS) for the years 2010 and 2020. Both data sources are extracted from IPUMS (Ruggles et al., 2013). To increase sample size, we follow Autor, Dorn and Hanson (2013) and Acemoglu and Restrepo (2020) in measuring 2010 variables using pooled five-year ACS data for 2011 (2007-2011). Similarly, we measure 2020 variables from pooled five-year ACS data for 2021 (2017-2021).³

We construct total CZ-level employment using sample weights, considering working-age individuals (age 16+) who are not unpaid family workers, do not reside in institutional group quarters, and have reported being employed over the previous year (Autor and Dorn, 2013). In later sections, we also consider disaggregations of employment along various dimensions. Specifically, we distinguish between employment in private and public sectors as well as in primary, secondary and tertiary industries. We also disaggregate employment between high- and low-skill workers, production and non-production workers, male and female workers, occupations at different points of the wage distribution, and by age.⁴ Using data from the

²CZs are clusters of counties with strong commuting ties within them and weak commuting ties among them (Tolbert and Sizer, 1996). As such, CZs may approximate local labor markets. The CZs in our sample are the same as in Autor and Dorn (2013), Autor, Dorn and Hanson (2013) and Acemoglu and Restrepo (2020).

³The Census and the ACS are 5% and 1% samples, respectively, of the US population, and are representative at the level of micro-regions known as Public Use Microdata Areas (PUMAs). We map PUMAs to CZs using a crosswalk developed by Autor and Dorn (2013).

⁴The public sector comprises transportation, communication, other public utilities, public administration, and armed forces. The private sector is made up of all non-public sector industries. The primary sector comprises agriculture, forestry, fisheries and mining. The secondary sector comprises construction and manufacturing. The tertiary sector consists of the public sector as well as wholesale trade, retail trade, finance, insurance, real estate, business and repair services, personal services, entertainment and recreation services, professional and related services. High-skill workers are those who have completed at least a bachelor's degree, low-skill workers those who have completed less than a bachelor's degree. Production occupations comprise construction, extraction, installation, maintenance, repair and other production occupations.

Census and the ACS, we also construct CZ-level population and numerous proxies for initial demographic and industrial composition (details in Section 4).

3.2 MEASURING AI ADOPTION WITHIN CZs

Precise metrics for quantifying AI adoption are presently lacking. To make progress, we build a novel proxy for AI adoption in each CZ by exploiting a specific feature of AI technologies. Specifically, adopting these technologies demands a wide spectrum of specialized software, which is essential for a firm to: execute existing machine learning algorithms, which extract patterns from large datasets by exploiting results from statistics and data science; combine and adapt these algorithms to solve complex problems specific to the distinct needs of the firm; generate, update and assemble the input datasets in real time; train the algorithms and govern their learning process. In turn, operating this specialized software necessitates a distinct set of skills, which are primarily possessed by professionals engaged in a narrow range of occupations pertaining to the domains of computer science, mathematics and statistics.

Our proxy for AI adoption leverages the specific requirement of specialized software and advanced technical skills that is typical of AI. In a nutshell, we first identify a set of AI-related occupations using data on the software knowledge required to workers in each job. Then, following the theoretical model, we proxy for AI adoption in a CZ using the increase in the relative importance of AI-related occupations in that locality.

To identify AI-related occupations, we take advantage of a novel section of the O*NET database called "Hot Technologies". The latter reports the software requirements that are most frequently included in all current employer job postings in the US. The list of software includes 157 titles, spanning from software with general applications like Microsoft Excel, to advanced programming languages like Python and C++. With the help of computer scientists, we narrow down the list to 54 software that are normally used for data collection and generation; for execution and adaptation of machine learning algorithms; and to feed these algorithms with large (structured and unstructured) datasets. The list of software is reported in Appendix Table B2.

Using the "Hot Technology" section of O*NET, we identify occupations for which each software is "in demand". These are occupations to which knowledge of a given software is typically required in job postings. This yields 82 occupations, defined according to the 2018 version of the Standard Occupational Classification (SOC). We refine the list by applying two sequential filters. First, we restrict to occupations for which at least two software are "in demand". This excludes 21 occupations that use a single software in their daily activities.

Examples are "Special Effects Artists and Animators", who use Python, or "Commercial and Industrial Designers", who use JavaScript. Second, we use occupational definitions from the SOC classification, together with information on each occupation's main tasks provided in O*NET, to further exclude occupations whose main activities fall outside of the domains of computer science, mathematics and statistics. This filter eliminates occupations like "Actuaries". While both Python and SQL are "in demand" for this job, the main role of an actuary is to "analyze statistical data, such as mortality, accident, sickness, disability, and retirement rates and construct probability tables to forecast risk and liability for payment of future benefits."

The final list of AI-related occupations comprises 19 titles, which are listed in Table 1. Using correspondence tables from the US Bureau of Labor Statistics, we track these occupations back in time across the revisions of the SOC classification occurred during the sample period. Then, we use information on each worker's SOC occupation (available in the Census and in the ACS) to match our consistent set of AI-related occupations with the micro-level data. With this information, we measure employment in AI-related occupations in each CZ and year. Finally, following (9), we construct our proxy for AI adoption as follows:

$$AIado_{ct} = \frac{L_{c\tau_1}^{AI} - L_{c\tau_0}^{AI}}{L_{c2000}}, \quad (10)$$

where $L_{c\tau_0}^{AI}$ and $L_{c\tau_1}^{AI}$ denote employment in AI-related occupations in CZ c in the first year (τ_0) and last year (τ_1) of each decade t , while L_{c2000} is total employment (across all occupations) in CZ c in 2000. Hence, for each CZ, $AIado_{ct}$ measures the decadal change in the relative importance of AI-related occupations, as proxied by the employment of these professions relative to initial employment in the CZ.

Before proceeding, we pause to discuss two potential limitations of this variable. First, not all employment in AI-related occupations needs to be devoted to AI adoption. Second, employment growth in AI-related occupations may reflect a wider usage of ICT independently of AI. We believe both issues to have a modest influence on $AIado_{ct}$. On the one hand, the list of AI-related occupations is obtained starting from a restricted list of specialized software, and is further narrowed down by our sequential filters. Second, because ICT have rapidly spread all over the US during previous decades, the scope for further diffusion of ICT over our sample period is likely to be limited.

Nevertheless, in the empirical analysis, we specifically tackle these two issues. In a robustness check, we construct $AIado_{ct}$ using a single AI-related occupation, "Data Scientists",

Table 1: AI-Related Occupations

Computer and Information Research Scientists	Mathematicians
Computer Network Architects	Network and Computer Systems Administrators
Computer Network Support Specialists	Operations Research Analysts
Computer Occupations, All Other	Software Developers
Computer Programmers	Software Quality Assurance Analysts and Testers
Computer Systems Analysts	Statistical Assistants
Computer User Support Specialists	Statisticians
Data Scientists	Web and Digital Interface Designers
Database Administrators	Web Developers
Database Architects	

Occupations are classified according to the 6-digit level of the 2018 Standard Occupational Classification.

whose tasks are limited to the typical domains of AI.⁵ The resulting indicator is certainly too narrow to capture the actual size of AI adoption in the US, but the patterns it delivers are identical to those of our broader proxy, both qualitatively and quantitatively. Second, we present robustness checks controlling for various proxies for ICT exposure. Controlling for these variables does not affect our main results. Interestingly, we also find ICT to have smaller, and sometimes opposite, effects on employment compared to AI.

4 EMPIRICAL SPECIFICATION AND IDENTIFICATION STRATEGY

In this section, we present the empirical specification and illustrate our identification strategy.

4.1 REGRESSION EQUATION

Guided by the theoretical model (eq. 9), our empirical analysis relates changes in employment to AI adoption across CZs. To this purpose, similarly to Autor, Dorn and Hanson (2013) and Acemoglu and Restrepo (2020), we estimate specifications of the following form:

$$\Delta(L/P)_{c dt} = \alpha_d + \alpha_t + \beta \cdot AIado_{c dt} + \mathbf{X}'_{c dt} \cdot \boldsymbol{\gamma} + \varepsilon_{c dt}, \quad (11)$$

⁵According to the definition provided in the SOC classification, Data Scientists "develop and implement a set of techniques or analytics applications to transform raw data into meaningful information using data-oriented programming languages and visualization software. Apply data mining, data modeling, natural language processing, and machine learning to extract and analyze information from large structured and unstructured datasets. Visualize, interpret, and report data findings. May create dynamic data reports".

where c indexes a CZ, d denotes the Census Division to which the CZ belongs, t stands for time, and $\varepsilon_{c dt}$ is an error term.⁶ We estimate (11) by stacking two first-differences corresponding to changes over 2000-2010 and 2010-2020, the two decades spanned by our sample. The main outcome variable, $\Delta(L/P)_{c dt}$, is the change in the employment-to-population ratio of CZ c over decade t .⁷ To shed light on the mechanisms, we also consider alternative outcomes, such as unemployment and non-participation rates, and disaggregate employment across different sectors. To study heterogeneity, we further split employment across occupations, skill levels, genders and age groups. The main explanatory variable is $AIado_{c dt}$, the proxy for the adoption of AI technologies in CZ c over decade t introduced in Section 3. Given the definition of $AIado_{c dt}$ in (10), eq. (11) is a changes-on-changes regression.

The specification in (11) includes a wealth of fixed effects and covariates, which control for other characteristics of the CZ that may influence AI adoption and employment. Specifically, we control for Census Division fixed effects, α_d , to absorb heterogeneous trends in labor market conditions across groups of contiguous states. We also include decade fixed effects, α_t , to soak up time-varying shocks hitting all CZs simultaneously. Moreover, we control for a large set of covariates, included in the vector $\mathbf{X}_{c dt}$. The latter contains two types of variables. First, it includes several proxies for initial characteristics of each CZ: (i) size, measured by log population; (ii) demographic composition, proxied by the population shares of female, college-educated, White and old individuals (age 65+); and (iii) composition of economic activities, captured by the share of manufacturing in total employment, the share of females in manufacturing employment, the share of light manufacturing in total manufacturing employment, and the employment share of workers in routine-intensive occupations.⁸ These variables account for heterogeneous trends across CZs characterized by different initial demographic and industrial structures. Second, the vector $\mathbf{X}_{c dt}$ includes proxies for two main labor market shocks considered in the recent literature: the change in import competition from China (Autor, Dorn and Hanson, 2013) and exposure to industrial robots (Acemoglu and Restrepo, 2020). Both variables are measured in each decade.⁹

⁶Census Divisions are defined by the US Census Bureau and subdivide the country into nine groups of states: New England, Middle Atlantic, East North Central, West North Central, South Atlantic, East South Central, West South Central, Mountain, and Pacific.

⁷While the model yields sharp predictions for non-AI employment, we prefer to focus on total employment to be more general and conservative, also given that AI employment is still small as shown in Section 5.

⁸Light manufacturing is made up of textile mill products; apparel and other finished textile products; paper and allied products; printing, publishing and allied industries. The employment share of routine-intensive occupations is defined as the share of hours worked in occupations with routine-task intensity at the top tercile of the distribution (Autor and Dorn, 2013) and is constructed using data from the 2000 Census.

⁹Import competition from China is a Bartik measure of changes in imports from China per worker across industries, as in Autor, Dorn and Hanson (2013). Exposure to industrial robots is a Bartik measure of

The coefficient of interest is β . Our empirical design implies that this coefficient measures the relative difference in the growth of employment relative to population across CZs that have similar initial conditions, face similar trade and automation shocks, but experience different rates of AI adoption over time. Although the rich set of controls and fixed effects absorb most observable confounders, the OLS estimate of β need not have a causal interpretation, due to unobservable factors that may simultaneously affect AI adoption and employment. We now turn to discussing the key identification concerns and illustrate our strategy to estimate causal effects.

4.2 IDENTIFICATION STRATEGY

Variation in $AIado_{cdt}$ across CZs could reflect CZ-specific unobservables that also influence employment. In particular, demand shocks in a CZ may lead firms to hire more workers and use more AI technologies, inducing a spurious positive correlation between $AIado_{cdt}$ and $\Delta(L/P)_{cdt}$. To account for this issue, we construct an instrument that is meant to isolate variation in $AIado_{cdt}$ not due to demand shocks within CZs. Following a long tradition initiated by Bartik (1991) and Blanchard and Katz (1992), and recently applied to the automation literature by Acemoglu and Restrepo (2020), we use a shift-share (Bartik) instrument, which combines nation-wide industry *shifts* with local industry *shares*. The instrument is constructed as follows:

$$AIexp_{cdt} = \sum_{i=1}^I \lambda_{cdi1980} \times \left(\frac{L_{i\tau_1}^{AI} - L_{i\tau_0}^{AI}}{L_{i2000}} \right), \quad (12)$$

where $L_{i\tau_0}^{AI}$ and $L_{i\tau_1}^{AI}$ denote employment of AI-related occupations in industry i , at the national level, in the first year (τ_0) and last year (τ_1) of decade t , respectively; L_{i2000} is total employment in industry i , at the US level, in 2000; and $\lambda_{cdi1980} \equiv \frac{L_{cdi1980}}{L_{cd1980}}$ is industry i 's share in total employment of CZ c in 1980.

The intuition behind this instrument is the following. As technological progress lowers the cost of AI and/or increases its productivity, AI adoption grows, especially in industries whose activities are more amenable to the use of these technologies. These industry-level AI adoption shifts have asymmetric effects across CZs, depending on historical differences in their industrial specialization, as captured by the employment shares $\lambda_{cdi1980}$. In the baseline version of the instrument, we measure these shares in 1980, i.e., 20 years before the

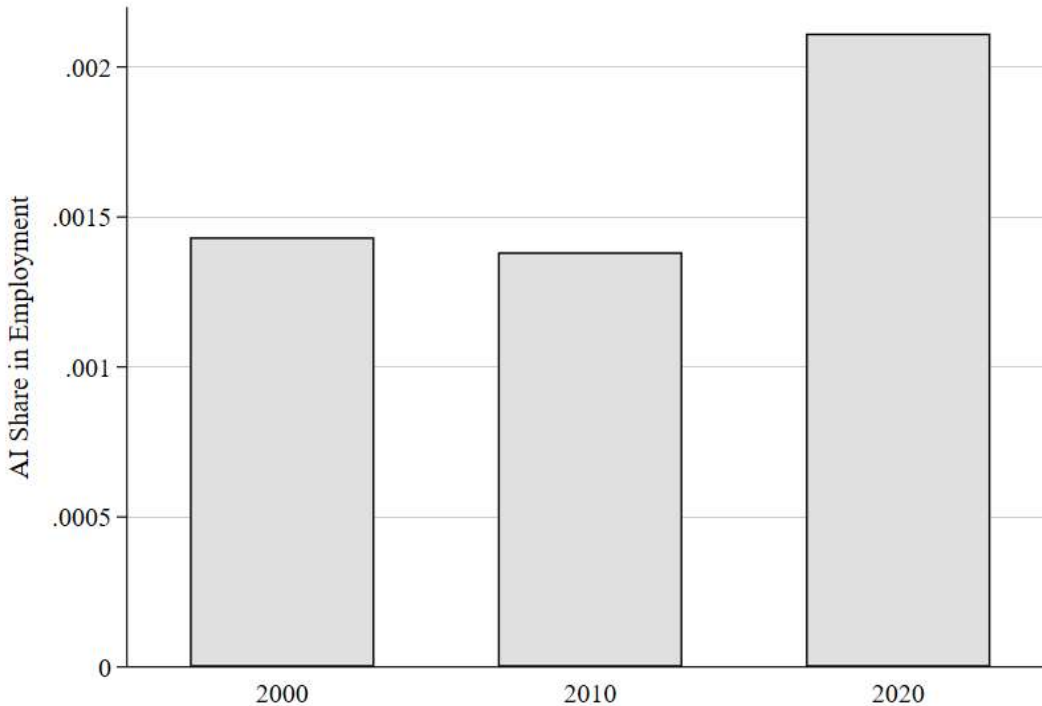
changes in robot density (the number of installed robots per worker) across 29 sectors, and is constructed as in Acemoglu and Restrepo (2020) using data on robot installments from the IFR.

beginning of the sample period. Because AI was largely inexistent at that time, this choice mitigates the concern that the industrial composition of a CZ might have been influenced by the anticipation of AI developments in future decades. We keep the employment shares constant to avoid endogenous and serially correlated changes in $AIexp_{czt}$ in the context of our stacked first-differences specification.

To compute the industry-level shifts, we use the micro-level data from the Census and the ACS, which contain information on each worker’s industry of employment. The employment shares are computed using data from the County Business Patterns, which we slightly aggregate to match the industry detail of the micro-level data. The final sample includes 188 NAICS industries, mostly at the 3- or 4-digit level, spanning all sectors of the economy. It is worth noting that our instrument exploits significantly larger cross-industry variation than the typical Bartik measures used in the recent automation literature. For instance, the proxies for exposure to industrial robots based on data from the IFR aggregate industry-level shifts for 19 broad sectors.

The identifying assumption behind our approach is that the industry-level shifts are exogenous to shocks occurring in individual CZs. We believe this to be a reasonable assumption, given that most CZs are tiny relative to the overall size of the US economy. Moreover, our specification controls for a large set of fixed effects and covariates, which absorb a wealth of potential observable confounders. Yet, our identification strategy would be endangered in two cases. First, if some contemporaneous shocks remained that correlate with the outcome $\Delta(L/P)_{czt}$ and the instrument $AIexp_{czt}$. Second, if some remaining underlying trends at the CZ-level influenced employment independently of AI.

In Section 7, we use various approaches to account for underlying trends and contemporaneous shocks more flexibly. We find that these potential confounders are unlikely to drive the results. Moreover, we implement a falsification test showing that current AI adoption does not explain past changes in employment. This further suggests that our results are driven by period-specific AI adoption rather than secular trends that predate the rise of AI. This extensive sensitivity analysis gives further credibility to our results but might not eliminate all concerns with a violation of the exclusion restriction. Hence, we also use an approach developed by Conley, Hansen and Rossi (2012) to study how strong a violation would have to be for inference on β to become uninformative about the causal effect of AI. We find that inference would remain informative even under implausibly large violations.



Source: US Census and American Community Survey. Each bar corresponds to the employment share of AI-related occupations in the US in a given year.

Figure 1: Employment Share of AI-Related Occupations in the US

5 STYLIZED FACTS AND PRELIMINARY EVIDENCE

We now illustrate the main patterns of AI adoption emerging from our data and provide preliminary evidence on the relationship between AI and employment in the US. Figure 1 plots the nation-wide employment share of AI-related occupations in 2000, 2010 and 2020. This share has remained fairly constant, at around 0.14%, over the 2000s but has rapidly increased thereafter, exceeding 0.2% in 2020. This pattern confirms existing evidence according to which the use of AI was quite limited in the early 2000s but has significantly accelerated after 2010 (e.g., Taddy, 2019, Alekseeva et al., 2021). Our stacked first-differences specification exploits the differential AI adoption between the two decades.

The process of AI adoption is not uniform across industries. Rather, our data reveal a substantial degree of heterogeneity across the 188 industries in our sample, a crucial aspect for the identification strategy. Specifically, the standard deviation of the industry-level AI adoption shifts is 0.3%, almost five times the mean, and the difference between the industries at the top and bottom percentile of the distribution is 1.2%. Another distinguishing feature of AI adoption is the type of industries that are most interested by the phenomenon. Table

Table 2: Top and Bottom Industries in Terms of AI Adoption in the US

Top 10 Industries		Bottom 10 Industries	
Computer Systems Design and Related Services	0.0241	Household Appliance Stores	-0.0048
Software Publishers	0.0156	Computer and Peripheral Equipment Manufacturing	-0.0037
Executive Offices and Legislative Bodies	0.0078	Other General Government and Support Activities	-0.0018
Management of Companies and Enterprises	0.0071	Printing and Related Support Activities	-0.0011
Electric Power Generation, Transmission and Distribution	0.0030	Utilities (excl. Electricity)	-0.0009
National Security and International Affairs	0.0030	Commercial and Service Industry Machinery Manufacturing	-0.0007
Scientific Research and Development Services	0.0028	Structural Clay Products	-0.0006
Management, Scientific and Technical Consulting Services	0.0023	Apparel Accessories and Other Apparel	-0.0006
Finance and Insurance	0.0019	Beverage Manufacturing	-0.0006
Other Professional, Scientific, and Technical Services	0.0017	Electronic and Precision Equipment Repair and Maintenance	-0.0005

Industries are classified according to the NAICS classification. The second and fourth columns report the AI adoption shift in each industry (the term in round brackets in eq. 12) averaged between the decades 2000-2010 and 2010-2020.

2 lists the top ten and bottom ten industries in terms of the average shifts over the sample period. The top industries consist of advanced services such as information services; professional, scientific and technical services; business services; and financial services. They also comprise some utilities (electricity) and some of the most advanced branches of the public sector (national security and international affairs; executive offices and legislative bodies). Conversely, the bottom industries include traditional manufacturing sectors like beverage, apparel and structural clay products; utilities other than electricity; retail services; and the public administration.

Interestingly, AI adoption involves very different industries compared to the adoption of industrial robots. Indeed, Acemoglu and Restrepo (2020) show that automation is mostly concentrated in manufacturing, especially in highly mechanized industries such as automotive; in the chemical and pharmaceuticals sectors; and in heavy activities such as production of metal products. This suggests that AI adoption and automation are distinct shocks, which are likely to have different implications for labor demand. Our empirical results will support this view.

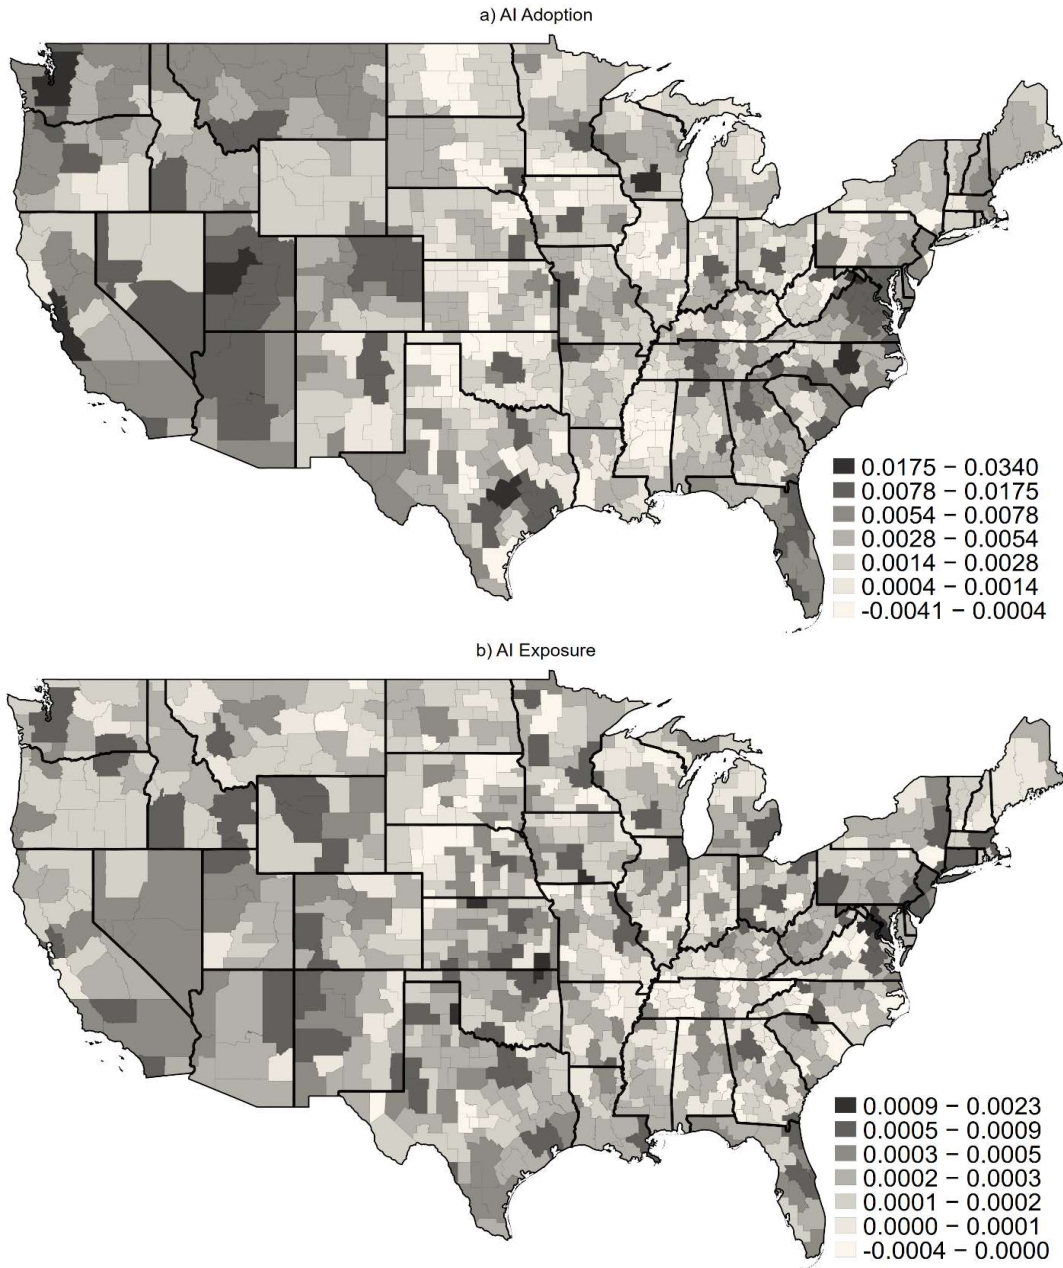
We now discuss regional variation. Figure 2a shows the average value of $AIado_{cdt}$, computed between the two decades, in each CZ; Figure 2b displays the corresponding values of $AIexp_{cdt}$. Four main patterns stand out. First, negative values of $AIado_{cdt}$ are very rare (only 6% of CZs), implying that AI adoption has been a widespread phenomenon in the US over the last two decades. Second, because CZs differ in their historical industrial specialization, the cross-industry variation in the shifts translates into significant differences in AI exposure across localities. $AIexp_{cdt}$ is particularly high in some CZs on the West Coast—especially in states like California and Washington—and in South-Central US—e.g., in states like Texas and Arizona. It is also high in CZs comprising large cities on the East Coast and the Great Lakes region, like Boston, New York, Miami, Philadelphia, Washington D.C. and Chicago. On the contrary, $AIexp_{cdt}$ is relatively low in Northern states and in most of the Midwest.

Third, the spatial variation in $AIado_{cdt}$ largely resembles that in $AIexp_{cdt}$. This suggests that the combination of nation-wide industry-level AI adoption shifts with spatial historical differences in industrial structure is a good predictor of actual AI adoption across CZs in the US. Fourth, there are notable exceptions to this pattern. For instance, some CZs in Northern states score higher in the ranking of $AIado_{cdt}$ than they do in terms of $AIexp_{cdt}$. This is consistent with the idea that AI adoption could be driven by CZ-specific factors, independently of actual exposure to these technologies. This is exactly the endogenous variation in AI adoption our identification strategy aims to get rid of.

Figure 3 illustrates our empirical strategy and provides a preliminary, visual, representation of the main results. The figure contains four scatterplots. In each of them, the hollow circles correspond to CZ \times decade pairs, for a total of 1,444 observations. Plot a) illustrates the relationship between $\Delta(L/P)_{cdt}$ and $AIado_{cdt}$. The relationship is negative (coeff. -0.345 , s.e. 0.062), implying that CZs with larger AI adoption experience relatively faster declines in employment as a share of population. Plot b) shows the first-stage relationship between $AIado_{cdt}$ and $AIexp_{cdt}$. The plot underscores the strong predictive power of our instrument in explaining variation in AI adoption across CZs (coeff. 7.782 , s.e. 0.616). Plot c) displays the reduced-form relationship between $\Delta(L/P)_{cdt}$ and $AIexp_{cdt}$. The strong negative association between the two variables (coeff. -12.889 , s.e. 1.015) suggests that CZs with higher AI exposure experience significantly larger reductions in the employment-to-population ratio compared to less exposed CZs. Plot d) finally shows the relationship between $\Delta(L/P)_{cdt}$ and \widehat{AIado}_{cdt} , the fitted value of AI adoption from the first-stage regression. The association between the two variables is strongly negative (coeff. -1.656 , s.e. 0.130), suggesting that differences in AI adoption across CZs, driven by variation in exposure to AI, have a strong negative effect on employment as a share of population. Interestingly, the relationship in Plot d) is stronger than its counterpart in Plot a), consistent with demand shocks inducing an upward bias in the OLS estimates.

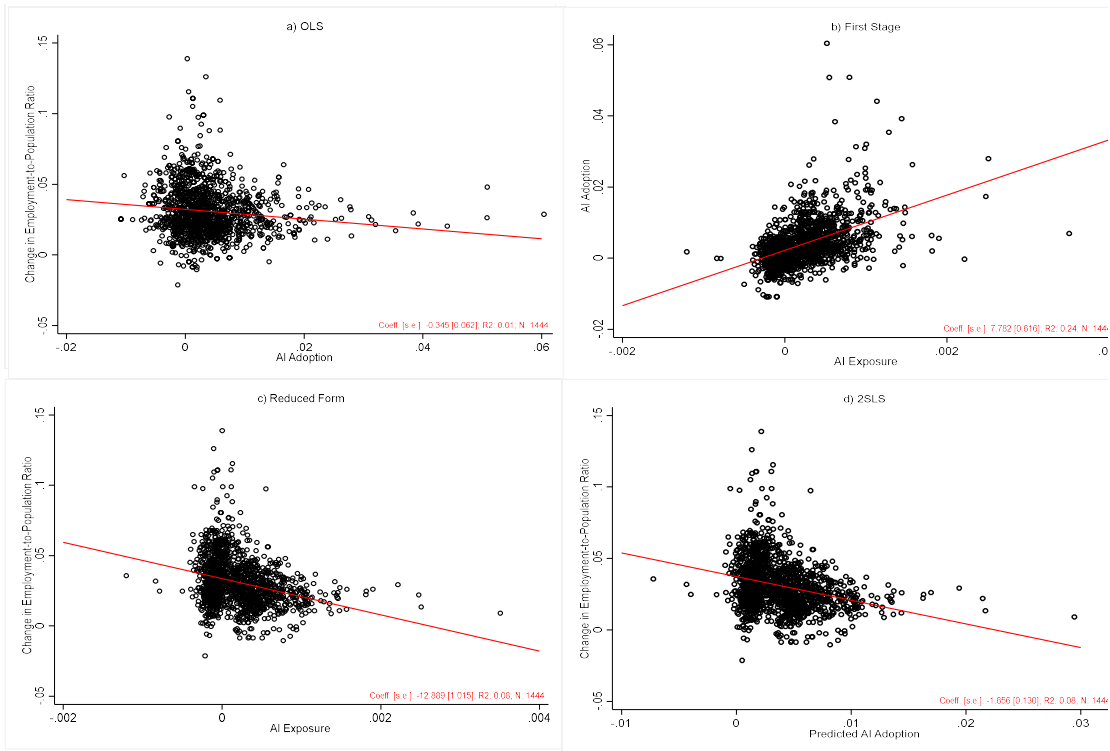
6 MAIN RESULTS AND ROBUSTNESS

In this section, we present the baseline results and perform an extensive sensitivity analysis to assess their robustness.



Source: US Census and American Community Survey. The top map plots the average value of $AIado_{czt}$ (see eq. 10) in each CZ between the decades 2000-2010 and 2010-2020. The bottom map plots the corresponding values of $AIexp_{czt}$ (see eq.12).

Figure 2: AI Adoption and AI Exposure in US Commuting Zones



The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. In each plot, an observation is a CZ x decade pair. *AI Adoption* and *AI Exposure* are defined in eq. (10) and (12), respectively. *Predicted AI Adoption* is the fitted value of *AI Adoption* from the first-stage regression in Plot b).

Figure 3: AI Adoption, AI Exposure and Employment in US Commuting Zones

6.1 BASELINE ESTIMATES

The baseline estimates of (11) are reported in Table 3. Standard errors are corrected for clustering at the state level to account for residual correlation across CZs within each state, and observations are weighted by the initial-period share of each CZ in total population. The first four columns correspond to a parsimonious specification including only Census Division and decade fixed effects. Column (1) reports the OLS estimate of β , which is negative and very precisely estimated (coeff. -0.253 , s.e. 0.080), confirming that AI adoption and changes in the employment-to-population ratio are negatively correlated across CZs. Columns (2) and (3) show, respectively, the first-stage and reduced-form coefficients on the instrument $AIexp_{cdt}$. The first-stage estimate is positive, large and highly statistically significant (coeff. 9.695 , s.e. 1.128), confirming the strong predictive power of the instrument.¹⁰ At the same time, the reduced-form estimate is negative and very precisely estimated (coeff. -6.507 , s.e. 1.447), confirming that CZs more exposed to AI experience a relatively larger reduction in employment as a share of population. The results in columns (2) and (3) imply a negative 2SLS estimate of β (coeff. -0.671 , s.e. 0.157), as shown in column (4). Also in this case, the 2SLS coefficient is more negative than its OLS counterpart, consistent with OLS estimates being upward biased due to unobserved CZ-specific shocks.

The last four columns of Table 3 report estimates from our preferred specification, which includes the full list of controls described in Section 4. The results are confirmed, suggesting that our evidence depends neither on initial differences in CZ characteristics nor on other important shocks occurred over the sample period, namely, the increase in Chinese import competition and the diffusion of industrial robots. The 2SLS coefficient reported in column (8) implies that a 1 standard deviation (s.d.) higher $AIado_{cdt}$ causes a reduction in $\Delta(L/P)_{cdt}$ by 1 percentage point (p.p.), roughly 0.56% of a s.d.. To have a sense of the magnitude of the effect, if the CZ with average AI adoption over the sample period (0.004) had hypothetically had no AI adoption at all, its employment-to-population ratio would have grown by 0.6 p.p. more, i.e., 20% faster than observed growth. While these numbers cannot be interpreted as a counterfactual exercise, they nevertheless suggest that AI adoption has contributed to slowing down job creation in the US over the last two decades.

¹⁰The Kleibergen-Paap F -statistic equals 73.8 and thus exceeds the value of 10 normally considered as a rule-of-thumb threshold for instrument relevance.

Table 3: Baseline Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	First Stage	Reduced Form	Second Stage	OLS	First Stage	Reduced Form	Second Stage
$AIado_{cdt}$	-0.253*** [0.080]			-0.671*** [0.157]	-0.238** [0.104]			-1.594*** [0.377]
$AIexp_{cdt}$		9.695*** [1.128]	-6.507*** [1.447]			9.599*** [1.813]	-15.298*** [3.279]	
Census Division FE	yes	yes	yes	yes	yes	yes	yes	yes
Decade FE	yes	yes	yes	yes	yes	yes	yes	yes
Control Variables	no	no	no	no	yes	yes	yes	yes
Kleibergen–Paap F -stat.	-	-	-	73.8	-	-	-	28.0
Obs.	1444	1444	1444	1444	1444	1444	1444	1444
R2	0.33	0.58	0.33	-	0.42	0.63	0.45	-

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variable is the change in the employment-to-population ratio in each CZ over each decade. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12) respectively. *Control variables* include the following initial characteristics of each CZ: log population; the population shares of female, college-educated, White and old individuals (age 65+); the share of manufacturing in total employment; the share of females in manufacturing employment; the share of light manufacturing in total manufacturing employment; and the employment share of workers in routine-intensive occupations. *Control variables* also include the change in import competition from China and exposure to industrial robots in each CZ over each decade. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

6.2 ROBUSTNESS CHECKS

We study the robustness of the baseline results along three dimensions: the presence of influential observations; the use of different corrections for the standard errors; and the adoption of alternative definitions for the main variables.

6.2.1 Outliers

In Table 4, we report 2SLS estimates of β obtained on various sub-samples, which exclude extreme observations in $\Delta(L/P)_{cdt}$ or $AIado_{cdt}$. We start by dropping observations for which $\Delta(L/P)_{cdt}$ is above or below the sample mean by more than two standard deviations (column 1), or for which $\Delta(L/P)_{cdt}$ falls in the top or bottom 1% (column 2) or 5% (column 3) of the distribution. In the last three columns, we use the same approaches to exclude extreme observations in $AIado_{cdt}$. In all cases, the coefficient β remains negative and very precisely estimated. In fact, excluding extreme observations either leaves the point estimate close to the baseline or makes it larger, suggesting that our results are not driven by outliers.

Table 4: Robustness Checks: Outliers

	(1)	(2)	(3)	(4)	(5)	(6)
	Excluding Outliers in $\Delta(L/P)$			Excluding Outliers in $AIado$		
	2std from Mean	1% Tails	5% Tails	2std from Mean	1% Tails	5% Tails
<u>2nd Stage Regression</u>						
$AIado_{cdt}$	-1.119*** [0.333]	-1.533*** [0.370]	-1.125*** [0.334]	-2.940*** [0.642]	-2.395*** [0.616]	-3.190*** [0.777]
<u>1st Stage Regression</u>						
$AIexp_{cdt}$	9.306*** [1.969]	9.556*** [1.852]	9.237*** [1.999]	7.169*** [1.150]	7.467*** [1.076]	5.752*** [1.096]
Kleibergen–Paap F -stat.	22.6	26.6	21.3	38.9	48.1	27.5
Obs.	1386	1414	1295	1387	1409	1298

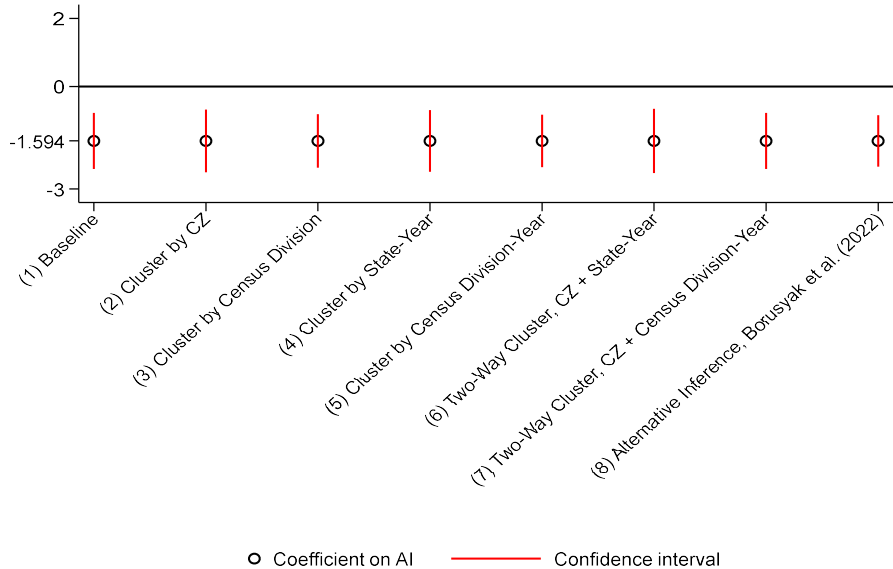
The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variable is the change in the employment-to-population ratio in each CZ over each decade. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. The specifications in columns (1)-(3) exclude observations for which the change in the employment-to-population ratio is above or below the sample mean by more than two standard deviations (column 1), or falls in the top or bottom 1% (column 2) or 5% (column 3) of the distribution. The specifications in columns (4)-(6) exclude the corresponding observations of $AIado_{cdt}$. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

6.2.2 Inference

In Figure 4, we plot the baseline 2SLS estimate of β along with confidence intervals corresponding to alternative corrections for the standard errors. Confidence interval (1), labeled "Baseline", is based on standard errors corrected for clustering at the state level. Confidence intervals (2)-(5) are based on standard errors corrected for clustering at: the CZ level; the Census Division level; the state-year level; and the Census Division-year level. These clustering structures allow for correlation in the residuals: within each CZ over time; across all CZs within the same Census Division and over time; and across all CZs within the same state, or within the same Census Division, in a given year. Confidence intervals (6) and (7) combine the latter two clustering structures with clustering at the CZ level (two-way clustering) to allow for residual correlation both across CZs and over time within a CZ. Finally, confidence interval (8) is based on Borusyak et al. (2022) inference approach. The latter takes account of the fact that, with a shift-share instrument, standard inference may be invalid because observations with similar industry shares may have correlated residuals. Reassuringly, all confidence intervals are very similar to the baseline one, suggesting that our results are not sensitive to the specific approach used for inference.

6.2.3 Variables Definitions

In Table 5, we consider alternative ways of defining the main variables. We start by tackling the concern that our definition of AI-related occupations might be too broad, including work unrelated to AI. To this purpose, we reconstruct both $AIado_{cdt}$ and $AIexp_{cdt}$ using a narrow



Hollow circles correspond to the baseline 2SLS coefficient on $AIado_{cdt}$ (Table 3, column 8). Confidence interval (1) refers to standard errors corrected for clustering at the state level. Confidence intervals (2)-(7) correspond to standard errors corrected using the alternative clustering schemes indicated on the horizontal axis. Confidence interval (8) is based on Borusyak et al. (2022) inference approach. All confidence intervals are at the 95% level.

Figure 4: Robustness Checks: Inference

set of AI-related occupations, which only includes "Data Scientists". Employment in this occupation has started to increase only recently and is still very low (0.01% at the national level in 2020). Nevertheless, the estimate of β reported in column (1) is negative and very precisely estimated also in this case. Quantitatively, the effect is essentially identical to the baseline: a 1 s.d. higher $AIado_{cdt}$ would cause a fall in $\Delta(L/P)_{cdt}$ of 1 p.p., or 0.56% of its s.d..

In column (2), we reconstruct $AIexp_{cdt}$ using industry shares for 1990 rather than 1980. It is reassuring that β hardly changes, suggesting that $AIexp_{cdt}$ exploits long-lasting differences in the industrial specialization of CZs, rather than period-specific developments in their industrial structure that could result from transitory shocks. In column (3), we adjust $AIado_{cdt}$ for cross-industry differences in the growth rate of employment within each CZ, and $AIexp_{cdt}$ for cross-industry differences in the growth rate of employment at the national level. While the first-stage relationship is slightly weaker, our main results are qualitatively unchanged and quantitatively stronger. This suggests that our evidence is not driven by changes in industries' relative size but reflects variation in AI intensity within industries.

In column (4), we estimate (11) for the log change in employment, adding the log change in population among the controls. The coefficient implies that a 1 s.d. higher $AIado_{cdt}$ would

Table 5: Robustness Checks: Variables Definitions

	(1)	(2)	(3)	(4)	(5)	(6)
	Only Data Scientists	Ind. Shares for 1990	Adj. for Ind. Growth	Log Emp.	Private Sector Emp.	Public Sector Emp.
<u>2nd Stage Regression</u>						
$AIado_{c,d,t}$	-17.340*** [4.395]	-1.138*** [0.289]	-3.443** [1.466]	-2.143*** [0.510]	-1.593*** [0.425]	-0.000 [0.194]
<u>1st Stage Regression</u>						
$AIexp_{c,d,t}$	12.286*** [1.835]	8.620*** [2.225]	1.288*** [0.458]	8.337*** [1.467]	9.599*** [1.813]	9.599*** [1.813]
Kleibergen–Paap F -stat.	44.8	15.0	7.9	32.3	28.0	28.0
Obs.	1444	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZ, Census Divisions and decades, respectively. The dependent variable is: the change in the employment-to-population ratio in each CZ over each decade (columns 1-3); the log change in employment in each CZ over each decade (column 4); and the change in private sector employment (column 5) or public sector employment (column 6) relative to population in each CZ over each decade. $AIado_{c,d,t}$ and $AIexp_{c,d,t}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. In column (1), $AIado_{c,d,t}$ and $AIexp_{c,d,t}$ are based on a narrow definition of AI-related occupations, which only includes "Data Scientists". In column (2), $AIexp_{c,d,t}$ is constructed using CZ-level industry shares for 1990 rather than 1980. In column (3), $AIado_{c,d,t}$ and $AIexp_{c,d,t}$ are adjusted for cross-industry differences in employment growth rates at the CZ and national level, respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. The specification in column (4) also includes the log change in population. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

lower employment by 1.3%. Finally, in columns (5) and (6), we revert to the employment-to-population ratio as the dependent variable but split the numerator into private and public sector employment. The results show that the effect of AI adoption is entirely concentrated in the private sector, where the bulk of the phenomenon is currently taking place.

7 THREATS TO IDENTIFICATION

Identification requires that, after controlling for Census Division fixed effects, decade fixed effects, initial characteristics of CZs and shocks to trade and automation, no unobservable remains that correlates with the instrument and influences employment across localities. As mentioned in Section 4, the exclusion restriction is threatened by two types of confounders, namely, underlying trends and contemporaneous shocks that might affect local labor markets independent of AI adoption. We now use alternative approaches for accommodating the possible influence of these confounders and study how the coefficient β is affected.

In Table 6, we deal with contemporaneous shocks. To start with, we allow for the possibility that CZs with similar changes in employment or similar AI adoption might be hit by similar unobservable shocks. To accommodate them, we divide CZs into ten bins corresponding to the deciles of $\Delta(L/P)_{cdt}$ or $AIado_{cdt}$ over 2000-2020. We then interact a dummy for each bin with decade dummies. We add these interactions either individually (columns 1 and 2) or jointly (column 3). These interactions absorb shocks hitting all CZs with comparable changes in our key variables. Accordingly, identification exploits the remaining variation across CZs falling in the same bin within each decade. In column (4), we take a complementary approach and exclude CZs with the highest values of $AIado_{cdt}$. These are CZs falling in the top decile of the distribution in a given decade. The coefficient β remains negative, precisely estimated and in the same ballpark as the baseline estimates across all specifications.

A related concern is that the industry shifts that make up the instrument might be driven by shocks occurring in specific CZs. If these shocks also affected the local labor market, the exclusion restriction would be violated. We believe this concern to be assuaged by the small size of most CZs relative to the US as a whole. Nevertheless, in column (5), we estimate (11) using a "leave-one-out" instrument, in which the industry-level shifts are computed after excluding the CZ to which the instrument refers. Our main results are preserved also in this case.

Unobserved shocks may also play out at the industry level. For example, industries with large AI adoption might experience other shocks that are relevant to the labor market of some CZs. Our specification already controls for trade and automation shocks, by means of Bartik

Table 6: Threats to Identification: Contemporaneous Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Bins of $\Delta(L/P)$	Bins of $AIado$	Bins of $\Delta(L/P)$ and $AIado$	No CZs with Top $AIado$	Leave-One-Out $AIexp$	No Ind. with Top Shifts	Bartik for Δ Ind. Emp.
<u>2nd Stage Regression</u>							
$AIado_{cdt}$	-1.131*** [0.229]	-2.965*** [0.940]	-2.421*** [0.810]	-2.536* [1.274]	-2.222* [1.132]	-1.945*** [0.656]	-1.345*** [0.320]
<u>1st Stage Regression</u>							
$AIexp_{cdt}$	9.276*** [1.345]	4.520*** [1.217]	4.462*** [1.330]	4.754*** [1.047]	0.315*** [0009]	8.402*** [2.235]	10.610*** [2.214]
Kleibergen–Paap F -stat.	47.5	13.8	11.3	20.6	12.3	14.1	22.9
Obs.	1444	1444	1444	1286	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variable is the change in the employment-to-population ratio in each CZ over each decade. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. In column (5), $AIexp_{cdt}$ is constructed excluding the CZ to which this variable refers. In column (6), $AIado_{cdt}$ and $AIexp_{cdt}$ are constructed excluding industries whose shifts (the round bracket terms of eq. 12) fall in the top decile of the distribution. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. The specification in column (1) controls for full sets of interactions between decade dummies and dummies for deciles of the change in the employment-to-population ratio over 2000-2020. The specification in column (2) controls for full sets of interactions between decade dummies and dummies for deciles of $AIado_{cdt}$ over 2000-2020. The specification in column (3) jointly includes the two sets of interactions used in columns (1) and (2). The specification in column (4) excludes CZs falling in the top decile of the distribution of $AIado_{cdt}$ in a given decade. The specification in column (7) includes a Bartik measure of the change in log industry employment. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

measures of changes in Chinese import penetration and in robot density. Moreover, in the next section, we will enrich the specification with Bartiks for different forms of technical change, and our results will turn out to be insensitive to these additional controls. Nevertheless, we now use two complementary approaches to further address concerns with industry-specific shocks.

In a first exercise, reported in column (6), we reconstruct both $AIado_{cdt}$ and $AIexp_{cdt}$ excluding industries whose AI adoption shifts fall in the top decile of the distribution. The coefficient of interest remains close to the baseline estimate. Interestingly, this exercise also serves two additional purposes. On the one hand, it further helps mitigating concerns with outliers. On the other hand, because some of the industries with the largest shifts could be AI-producing sectors, excluding them may better isolate the effect of AI adoption. In a second exercise, shown in column (7), we augment the specification with a Bartik measure of changes in log industry employment. This variable serves as a synthetic proxy for all shocks resulting in differential changes in employment across industries. Controlling for this measure is largely inconsequential for the results.

We now discuss underlying trends. Our specification controls for linear trends both across Census Divisions—through the Census Division fixed effects, α_d —and across CZs with similar initial characteristics within each Census Division—through the start-of-period controls included in the vector \mathbf{X}_{cdt} . In Table 7, we allow for richer sets of fixed effects to accommodate underlying trends more flexibly. In column (1), we replace the Census Division fixed effects

with state fixed effects, thereby allowing for heterogeneous linear trends across individual states rather than across groups of them. The coefficient of interest is essentially unchanged. In column (2), we go one step further by taking advantage of our stacked first-differences design, which allows us to include CZ fixed effects. In this case, identification relies on a different source of variation, namely, changes in $\Delta(L/P)_{cdt}$ and $AIado_{cdt}$ within each CZ between the two decades. This specification fully exploits the acceleration in AI adoption occurred after 2010, but is extremely demanding because the sample comprises only two observations for each CZ.¹¹ The estimate of β is largely unaffected. In a last exercise, we replace the Census Division and decade fixed effects with Census Division \times decade fixed effects (column 3) and state \times decade fixed effects (column 4), thereby allowing for non-linear trends. The main results are preserved.

In the remainder of this section, we use two additional approaches to further raise trust in our 2SLS estimates. In column (5) of Table 7, we implement a falsification test by regressing changes in the employment-to-population ratio prior to the beginning of the sample (over 1980-1990 and 1990-2000) on current AI adoption (in 2000-2010 and 2010-2020). We include the same controls as in the baseline specification, and instrument $AIado_{cdt}$ using $AIexp_{cdt}$. If current AI adoption explained past changes in employment, our evidence could reflect labor market trends that antedate the rise of AI, or long-lasting CZ characteristics that jointly affect innovation and employment. Reassuringly, however, the coefficient on $AIado_{cdt}$ is very small and statistically not significant. In column (6), we include the pre-sample change in the employment-to-population ratio among the controls. Consistently with column (5), this variable has no bearing on the coefficient β . Overall, these results help strengthening the view that our evidence captures the effects of current AI adoption rather than other time-varying confounders.

In a second approach, we change perspective and start from the premise that our instrument might be correlated with the error term due to some confounding factor. Then, following Conley, Hansen and Rossi (2012), we study how strong a violation of the exclusion restriction would have to be for inference on β to become uninformative about the causal effect of AI adoption. We illustrate the approach of Conley, Hansen and Rossi (2012) in Appendix A. Here, we present the main results, which are summarized in Figure 5. The latter plots the 90% confidence interval around β corresponding to different violations of the exclusion restriction, as captured by the parameter δ . When $\delta = 0$, we are in the benchmark case in which the exclusion restriction is satisfied. When $\delta > 0$, instead, the exclusion re-

¹¹Controlling for CZ fixed effects also accounts for possible mean reversion after the bust of the digital technology bubble.

Table 7: Threats to Identification: Underlying Trends

	(1)	(2)	(3)	(4)	(5)	(6)
	State Fixed Effects	CZ Fixed Effects	Cens. Div. x Decade FE	State x Decade FE	Past $\Delta(L/P)$ as Dep. Var.	Past $\Delta(L/P)$ as Control
<u>2nd Stage Regression</u>						
$AIado_{cdt}$	-1.529*** [0.389]	-1.299*** [0.438]	-1.400*** [0.317]	-1.202*** [0.330]	0.253 [0.171]	-1.494*** [0.363]
<u>1st Stage Regression</u>						
$AIexp_{cdt}$	10.398*** [1.924]	15.997*** [2.281]	9.591*** [1.836]	9.898*** [2.229]	9.599*** [1.813]	9.515*** [1.798]
Kleibergen–Paap F -stat.	29.2	48.9	27.3	19.7	28.0	28.0
Obs.	1444	1444	1444	1440	1444	1444

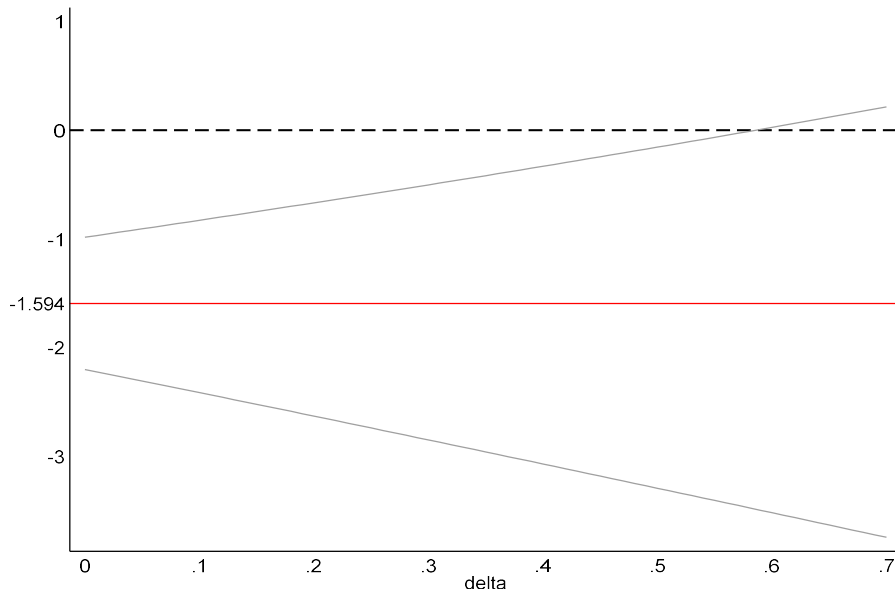
The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. Except for column (5), the dependent variable is the change in the employment-to-population ratio in each CZ over each decade; in column (5), the dependent variable is the change in the employment-to-population ratio in each CZ over the pre-sample decades 1980-1990 and 1990-2000. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. The specifications in columns (1)-(4) control for state fixed effects, CZ fixed effects, Census Division x decade fixed effects and state x decade fixed effects, respectively. The specification in column (6) controls for the change in the employment-to-population ratio in each CZ over the pre-sample decades 1980-1990 and 1990-2000. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

striction is violated. Specifically, $\delta = x > 0$ corresponds to a violation such that a change in $AIexp_{cdt}$ by one interquartile range has a direct effect on $\Delta(L/P)_{cdt}$ equal to a change in $AIado_{cdt}$ by x interquartile ranges.

When $\delta = 0$, the confidence interval around β is $[-2.202, -0.985]$. As δ departs from this benchmark, the confidence interval progressively widens. However, thanks also to the strong predictive power of the instrument, the reduction in precision is slow and the confidence interval starts including zero only when $\delta \approx 0.6$. Hence, for our parameter of interest to become uninformative, the direct effect of $AIexp_{cdt}$ on $\Delta(L/P)_{cdt}$ would have to be at least 60% as large as the effect of a commensurate exogenous change in $AIado_{cdt}$. Such a change is twice the median value of $AIado_{cdt}$ in our sample. More concretely, it roughly corresponds to the difference in average AI adoption between the CZ of Los Angeles and that of New Albany. Overall, these figures suggest that even substantial, and likely implausible, relaxations of the exclusion restriction would leave inference informative about the employment effect of AI adoption.

8 ADDITIONAL EVIDENCE

In this section, we dig deeper into the effect of AI adoption. We start by comparing AI with other shocks studied in the literature. Then, we explore the mechanisms through which the



The figure plots 90% confidence intervals around the baseline 2SLS coefficient on $AIado_{it}$ (Table 3, column 8) for different priors about a potential violation of the exclusion restriction. Priors are described by the parameter δ reported on the horizontal axis: $\delta=0$ implies that the exclusion restriction is satisfied; $\delta=x>0$ corresponds to a violation of the exclusion restriction such that a change in $AIexp_{it}$ by 1 interquartile range has a direct effect on the employment-to-population ratio equal to a change in $AIado_{it}$ by x interquartile ranges. The confidence intervals are based on standard errors corrected for clustering at the state level.

Figure 5: Threats to Identification: Sensitivity of Inference to Violations of the Exclusion Restriction

effect of AI unfolds. Finally, we study how the employment response to AI adoption varies by gender, age, skill and occupation.

8.1 AI ADOPTION AND OTHER SHOCKS

Our model shows that AI adoption affects labor demand in two ways: by replacing workers in some tasks (displacement effect) and by increasing efficiency (productivity effect). Accordingly, the effects of AI adoption should differ from those of other shocks that do not have a displacement effect. We now compare AI adoption with various shocks of this type. The results are reported in Table 8. In columns (1)-(6), we augment (11) with proxies for different shocks; in column (7), we include all these proxies together. The specification in column (1) includes a proxy for capital deepening, which is a Bartik measure of the change in capital intensity (capital-to-labor ratio) across industries.¹² The coefficient β is largely unchanged, suggesting that our results are not capturing the effect of capital deepening. At the same time, the coefficient on the new control is very small and positive, consistent with the effect

¹²This and the other Bartik measures used in Table 8 are constructed using data from the Production Accounts Tables of the US Bureau of Economic Analysis.

Table 8: Controls for Other Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<u>2nd Stage Regression</u>							
$AIado_{cdt}$	-1.597*** [0.377]	-1.590*** [0.376]	-1.575*** [0.377]	-1.744*** [0.428]	-1.675*** [0.379]	-1.833*** [0.503]	-2.185*** [0.660]
$CapInt_{cdt}$	0.000* [0.000]						0.000** [0.000]
$SoftInt_{cdt}$		-0.000 [0.000]					-0.001* [0.000]
$CompInt_{cdt}$			-0.001 [0.001]				0.003 [0.002]
$CommInt_{cdt}$				0.003** [0.001]			0.003 [0.002]
VA_{cdt}					0.088*** [0.021]		0.084*** [0.023]
$Offsh_{cdt}$						0.060* [0.035]	0.084* [0.045]
<u>1st Stage Regression</u>							
$AIexp_{cdt}$	9.603*** [1.818]	9.595*** [1.812]	9.468*** [1.778]	8.987*** [1.803]	9.636*** [1.821]	7.984*** [1.886]	7.011*** [1.879]
Kleibergen–Paap F -stat.	27.9	28.0	28.3	24.8	28.0	17.9	13.9
Obs.	1444	1444	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variable is the change in the employment-to-population ratio in each CZ over each decade. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. $CapInt_{cdt}$, $SoftInt_{cdt}$, $CompInt_{cdt}$ and $CommInt_{cdt}$ are Bartik measures of the change in, respectively, the capital-to-labor ratio, the software capital-to-labor ratio, the computer capital-to-labor ratio and the communication equipment-to-labor ratio, across industries. VA_{cdt} is a Bartik measure of the change in log industry value added. $Offsh_{cdt}$ is the initial employment share of offshorable occupations in each CZ. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

of capital deepening being different from that of AI.

Next, we study the implications of ICT. In columns (2)-(4), we add three Bartik measures of changes in (i) software, (ii) computer and (iii) communication equipment intensities, respectively. In all specifications, the effect of AI adoption remains very close to the baseline estimate. This reassures against the concern that $AIado_{cdt}$ might capture the effect of computers, software and other high-tech capital. Also in this case, the coefficients on the new controls are small and, when they are precisely estimated, their sign is positive. This pattern suggests that AI adoption has different labor demand consequences compared to ICT.¹³

¹³This finding is consistent with Gaggl and Wright (2017) and Blanas, Gancia and Lee (2019), who find positive associations between ICT and labor demand.

In column (5), we consider other productivity-enhancing shocks. To remain agnostic about the precise nature of the shock, we add a broad proxy, which is obtained as a Bartik measure of the change in log industry value added. This proxy enters with a positive and statistically significant coefficient, but its inclusion is inconsequential for our coefficient of interest. Hence, AI adoption differs from productivity-enhancing shocks, and has distinct effects on the labor market. Finally, in column (6) we compare AI adoption with offshoring. To this purpose, we add the initial employment share of offshorable occupations, constructed using data from Autor and Dorn (2013). The coefficient β is largely unchanged. Moreover, while offshoring could also displace workers, the results show that AI adoption has markedly different labor market effects.¹⁴

8.2 CHANNELS

So far, our evidence highlights a negative impact of AI adoption on employment across CZs. We now explore some of the mechanisms underlying this effect. The results are reported in Table 9. We start by analyzing the sectors that contribute the most to the overall effect. As shown in Section 5, the leading industries in terms of AI adoption belong to the service sector, while manufacturing is still lagging behind. It is conceivable, therefore, that AI adoption in services might currently have larger effects on labor demand than AI adoption in manufacturing. To study the role of the two sectors, we split both $AIado_{cdt}$ and $AIexp_{cdt}$ in two separate variables, constructed as in (10) and (12) on the subsets of manufacturing and non-manufacturing industries, respectively. We then estimate (11) using the sector-specific variables in place of the aggregate variables. The results are reported in column (1) for manufacturing and in column (2) for non-manufacturing. Consistently with the different speed of AI diffusion in the two sectors, the estimates show that the service sector makes up the lion's share of the overall effect.

A related issue has to do with the sectors that experience the largest changes in employment as a consequence of AI adoption. The previous results do not imply that the effect is entirely concentrated in the service sector, because industries are linked to each other by upstream or downstream relationships and because workers may move across sectors in response to shocks. To investigate the employment effects in different sectors, we divide the numerator of the dependent variable into employment in primary, secondary and tertiary industries. We also consider employment in manufacturing, which makes up the secondary sector together with the construction industry. We then estimate (11) using employment in each branch of

¹⁴See Bonfiglioli et al. (2022b) for more evidence on the relationship between automation and offshoring.

Table 9: Channels

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	AIado and AIexp for Mfg	AIado and AIexp for Nmfg	Primary Sector Emp.	Secondary Sector Emp.	Mfg Sector Emp.	Tertiary Sector Emp.	Unemployment	Not in Labor Force
<u>2nd Stage Regression</u>								
AIado _{edt}	1.127 [1.878]	-1.405*** [0.422]	0.252** [0.101]	-0.866*** [0.259]	-0.725*** [0.253]	-0.979*** [0.291]	0.390* [0.211]	1.203*** [0.241]
<u>1st Stage Regression</u>								
AIexp _{edt}	2.296*** [0.650]	5.089*** [1.269]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]
Kleibergen–Paap <i>F</i> -stat.	12.5	16.1	28.0	28.0	28.0	28.0	28.0	28.0
Obs.	1444	1444	1444	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. In columns (1)-(2), the dependent variable is the change in the employment-to-population ratio in each CZ over each decade. In columns (3)-(6), the dependent variables are the changes in primary sector employment, secondary sector employment, manufacturing sector employment and tertiary sector employment, respectively, as a share of population in each CZ over each decade. In columns (7)-(8), the dependent variables are the changes in the number of unemployed workers and in the number of individuals out of the labor force, respectively, relative to population in each CZ over each decade. $AIado_{c,d,t}$ and $AIexp_{c,d,t}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. In columns (1)-(2), these variables are constructed on the subsets of manufacturing and non-manufacturing industries, respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

the economy, relative to population, as the dependent variable. The results are reported in columns (3)-(6). The largest effect is found in the tertiary sector, which accounts for roughly 60% of the impact of AI adoption on overall employment. In addition, negative effects are evident also in the secondary sector, primarily in manufacturing, which accounts for 45% of the aggregate impact. Finally, the results show a small positive effect on employment in the primary sector, which probably absorbs a small fraction of workers laid off from the other branches.

The last results raise the question of what happens to displaced workers that are not re-employed. To answer this question, in the last two columns, we study the response of unemployment and non-participation rates. Specifically, we estimate (11) using two different dependent variables: unemployment as a share of population (column 7) and the share of population out of the labor force (column 8). The results indicate that AI adoption raises both unemployment and non-participation. However, non-participation absorbs a much larger share of the overall reduction in employment compared to unemployment (75% vs. 25%). This fraction includes both workers who temporarily leave the labor force to update their skills and discouraged workers who drop out of the labor force. In the next section, we study heterogeneity in the effect of AI adoption, with the aim of shedding light on the possible winners and losers from this new technology.

Table 10: Heterogeneity: Gender and Age

	(1)	(2)	(3)	(4)	(5)
	Male	Female	Emp.	Emp.	Emp.
	Emp.	Emp.	16-24 Yrs	25-44 Yrs	45+ Yrs
<u>2nd Stage Regression</u>					
$AIado_{cdt}$	-0.928*** [0.260]	-0.665*** [0.181]	-0.296* [0.166]	-1.401*** [0.366]	0.103 [0.207]
<u>1st Stage Regression</u>					
$AIexp_{cdt}$	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]
Kleibergen–Paap F -stat.	28.0	28.0	28.0	28.0	28.0
Obs.	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. In columns (1)-(2), the dependent variables are the changes in the employment of males and females, respectively, as a share of population in each CZ over each decade. In columns (3)-(5), the dependent variables are the changes in the employment of workers aged 16-24, 25-44 and 45+, respectively, as a share of population in each CZ over each decade. $AIado_{cdt}$ and $AIexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, * indicate significance at the 1, 5 and 10% level, respectively.

8.3 HETEROGENEITY

In columns (1) and (2) of Table 10, we investigate how the effect of AI adoption varies by gender. We find negative effects on both male employment and female employment as a share of population. While the coefficient is somewhat larger for men, the difference is not statistically significant. Hence, our results do not clearly point to a strong gender bias in the effect of AI adoption.

Different conclusions are reached about other dimensions of heterogeneity. A first aspect to play a role is age. Columns (3)-(5) report the effects of AI adoption on employment for three groups of workers: younger (age 16-24), middle-age (age 25-44) and older (age 45+) workers. The largest negative effect is found for middle-age workers, which account for almost 88% of the effect on overall employment. The impact of AI adoption is also negative, albeit much smaller, on younger workers, while it is small and imprecisely estimated on older workers. One possible interpretation of these results is that younger workers are better equipped to cope with the challenges posed by AI, as they are capable of swiftly updating their skills. Older workers are instead shielded from the AI shock possibly due to their tendency to hold more stable jobs. Notably, these advantageous traits are often less pronounced among middle-aged

Table 11: Heterogeneity: Skills and Occupations

	(1)	(2)	(3)	(4)	(5)	(6)
	HS Emp.	LS Emp.	HS-NP Emp.	HS-P Emp.	LS-NP Emp.	LS-P Emp.
<u>2nd Stage Regression</u>						
$AIado_{c dt}$	0.048 [0.326]	-1.641*** [0.450]	0.115 [0.326]	-0.067*** [0.023]	-1.021*** [0.340]	-0.621*** [0.206]
<u>1st Stage Regression</u>						
$AIexp_{c dt}$	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]	9.599*** [1.813]
Kleibergen–Paap F -stat.	28.0	28.0	28.0	28.0	28.0	28.0
Obs.	1444	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000-2010 and 2010-2020. Each observation is a CZ x decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variables are the changes in the employment of high-skill workers (column 1), low-skill workers (column 2), high-skill/non-production workers (column 3), high-skill/production workers (column 4), low-skill/non-production workers (column 5) and low-skill/production workers (column 6), respectively, as a share of population in each CZ over each decade. $AIado_{c dt}$ and $AIexp_{c dt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

workers.

A second important dimension of heterogeneity is education. As shown in columns (1) and (2) of Table 11, the negative effect of AI adoption is entirely concentrated on low-skill workers. For high-skill workers, the estimated coefficient is positive albeit imprecisely estimated. Occupations also play a role. As shown in columns (3)-(6), the effect of AI adoption is always negative for production workers. On the contrary, for non-production workers, the effect is negative only among low-skill ones. These results strengthen the view that the implications of AI adoption differ from those of industrial robots, whose effects are mostly concentrated on production workers and largely independent of education (Acemoglu and Restrepo, 2020). This difference may be due to two factors. On the one hand, the two technologies have spread in different industries: while AI is prevalent in advanced service industries, automation is more prominent in highly mechanized manufacturing sectors. On the other hand, the skill requirement of AI technologies is higher than that of industrial robots.

In conclusion, we consider another dimension of heterogeneity, which is the focus of an important stream of literature. It is well documented that, both in the US and in other industrialized countries, the labor market has undergone a process of job polarization: employment

Table 12: Heterogeneity: Quintiles of Initial Wage Distribution

	(1)	(2)	(3)	(4)	(5)
	Emp. 1st Quintile	Emp. 2nd Quintile	Emp. 3rd Quintile	Emp. 4th Quintile	Emp. 5th Quintile
<u>2nd Stage Regression</u>					
$AIado_{c dt}$	-0.750*	-0.916***	-0.262**	-0.283**	0.616***
	[0.415]	[0.213]	[0.122]	[0.137]	[0.199]
<u>1st Stage Regression</u>					
$AIexp_{c dt}$	9.599***	9.599***	9.599***	9.599***	9.599***
	[1.813]	[1.813]	[1.813]	[1.813]	[1.813]
Kleibergen–Paap F -stat.	28.0	28.0	28.0	28.0	28.0
Obs.	1444	1444	1444	1444	1444

The sample consists of 722 CZs observed over two decades, 2000–2010 and 2010–2020. Each observation is a CZ \times decade pair. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. The dependent variables are the changes in employment, as a share of population, for occupations in the first quintile (column 1), second quintile (column 2), third quintile (column 3), fourth quintile (column 4) and fifth quintile (column 5) of the initial wage distribution in each CZ over each decade. $AIado_{c dt}$ and $AIexp_{c dt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively. All specifications include Census Division fixed effects, decade fixed effects and the control variables used in Table 3. Observations are weighted by the initial-period share of each CZ in total population. Standard errors are corrected for clustering at the state level. ***, **, *: indicate significance at the 1, 5 and 10% level, respectively.

has shrunk among occupations in the middle of the wage distribution and expanded in those at the upper and lower tails. The hollowing-out of the wage distribution has proceeded over the period of our analysis. To see this, we follow Autor and Dorn (2013) and compute the change in employment across occupations belonging to each quintile of the initial wage distribution. We find a decline in the employment share of occupations in the third (middle) quintile (-1.45 p.p.) between 2000 and 2020, accompanied by an increase in the employment shares of occupations in the first and second (bottom) quintiles ($+0.51$ and $+0.49$ p.p., respectively) as well as of occupations in the fourth and fifth (top) quintiles ($+0.11$ and $+0.34$ p.p., respectively).

The most accredited explanation for job polarization is routine-biased technical change, i.e., the diffusion of technologies that substitute labor in standardized and codifiable activities, while complementing it in abstract and manual-intensive tasks. Our interest is to study whether AI adoption has contributed to job polarization over the last two decades. The previous results suggest that this might not be the case: as previously shown, the effects of AI adoption on overall employment are different from those of software and computers, which are the main culprits of job polarization according to the literature (e.g., Autor and Dorn, 2013).

To provide direct evidence, in Table 12, we estimate (11) using as the dependent variable the change in employment, as a share of population, for occupations in each quintile of the initial wage distribution. The results are inconsistent with a contribution of AI adoption to job polarization. Indeed, there is no evidence of a U-shaped pattern in the estimated coefficients. Rather, the coefficients follow a seemingly monotonic trend: they are strongly negative for the bottom two quintiles, mildly negative for the third and fourth quintiles, and strongly positive at the top of the distribution. This pattern is broadly in line with our previous evidence by skill group, as long as an occupation’s position in the wage distribution is correlated with its skill requirements. Overall, these results suggest that AI adoption poses threats to low- and medium-skill workers, while offering new opportunities to the most educated individuals. Hence, AI adoption may be a source of rising inequality in the US labor market.

9 CONCLUSIONS

Recent improvements in the field of AI have triggered much hype. The ongoing debate highlights the fact that AI is a flexible technology with the potential of turning dreamlike scenarios or nightmares into reality. Nobody can predict the direction that future innovations and applications will take. However, to inform policy decision, it is important to understand the consequences of these technologies so far. The goal of this paper has been to study the effect of AI adoption on labor demand as measured by changes in employment. Since the deployment of AI can potentially increase productivity but also automate work, its impact on employment is a still unanswered empirical question.¹⁵

Using data across US CZs over the period 2000-2020, a novel measure of AI adoption based on the growth of AI-related jobs and a shift-share empirical strategy to identify causal effects, we were able to estimate robust negative effects of AI exposure on employment. We also found that AI’s impact is different from other forms of capital and technologies, such as robots or ICT, and that it works through services more than manufacturing. Moreover, the employment effect is especially negative for low-skill and production workers, while it turns positive for workers at the top of the wage distribution. Overall, these results are consistent with the view that AI, so far, has contributed to the automation of jobs and to widen inequality. Finally, while the focus of this paper has been on employment so as to best capture displacement effects, AI adoption is likely to have affected wages as well. In the

¹⁵See Varian (2018) and Acemoglu (2022) for a discussion of policy questions pertaining to other aspects of AI.

interest of space, we left a detailed analysis of this possibility to future research.

We conclude by discussing some policy implications of our findings. The existing literature has already pointed out that automation, in the form of industrial robots, is partly to blame for falling manufacturing employment (e.g., Acemoglu and Restrepo 2020, Dauth et al. 2021, and Bonfiglioli et al. 2022a,b). Our alarming result is that service employment, which constitutes the lion's share in developed countries, may not be immune to automation through AI. It is therefore crucial that governments put in place measures aimed at alleviating adverse labor-market consequences. Since the negative effects are concentrated among low-skill workers, a first remedy could be to help employees acquire new skills. Facilitating job-to-job transitions and improving the flexibility of the labor market can also ease reallocation costs. Since top earners seem to actually benefit from these new technologies, appropriate transfer schemes could also be used to ensure that the gains are more broadly shared. Finally, incentive schemes could be designed to redirect innovation towards applications aimed at improving human capabilities rather than labor savings. To this end, collecting better data and developing methodologies to identify the complementarities between AI applications and jobs seem an important step for future research.

APPENDIX A THE CONLEY, HANSEN AND ROSSI (2012) APPROACH

In this section, we illustrate the main idea behind the approach of Conley, Hansen and Rossi (2012) using our set-up. Consider the following version of (11):

$$\Delta(L/P)_{c dt} = \alpha_d + \alpha_t + \beta \cdot AIado_{c dt} + \mathbf{X}'_{c dt} \cdot \boldsymbol{\gamma} + \lambda \cdot AIexp_{c dt} + \varepsilon_{c dt},$$

where λ is a parameter measuring the size of a violation of the exclusion restriction. The baseline results presented in the text are based on the standard IV assumption that $\lambda = 0$. However, if the exclusion restriction was not satisfied, i.e., if $\lambda \neq 0$, inference on β could still be performed, using alternative priors about λ and conditional on this parameter. This can be done by estimating the following specification

$$\Delta(L/P)_{c dt} - \lambda \cdot AIexp_{c dt} = \alpha_d + \alpha_t + \beta \cdot AIado_{c dt} + \mathbf{X}'_{c dt} \cdot \boldsymbol{\gamma} + \varepsilon_{c dt}$$

with 2SLS, instrumenting $AIado_{c dt}$ with $AIexp_{c dt}$. Varying the prior about λ allows assessing how inference on β would be influenced by different degrees of violation of the exclusion restriction. Because the sensitivity of the 2SLS estimator to violations of the exclusion restriction inversely depends on the strength of the instrument, the same value of λ induces a smaller decrease in precision the stronger is the first-stage relationship.

We set λ to be a function of a parameter δ , which we progressively raise (by intervals of 0.01 starting from 0) to generate increasingly larger violations of the exclusion restriction. Specifically, we set $\lambda \equiv -1.594 \times 11 \times \delta$, where -1.594 is the baseline 2SLS estimate of β and the interquartile range of $AIado_{c dt}$ is approximately 11 times that of $AIexp_{c dt}$. For each value of λ , we estimate the confidence interval of β for both the lower and the upper end of the support $[-\lambda, \lambda]$, and compute the final confidence interval as the union of the two confidence intervals.¹⁶

APPENDIX B DATA APPENDIX

¹⁶Besides this “union of confidence intervals” approach, Conley, Hansen and Rossi (2012) discuss other strategies that use more prior information about λ . By imposing additional parametric restrictions, these alternative approaches tend to yield narrower confidence intervals around the treatment parameter, and thus may be less conservative.

Table B1: Summary Statistics

	Obs.	Mean	Median	Std. Dev.
<u>Outcomes</u>				
$\Delta(L/P)_{cdt}$	1444	0.0309	0.0279	0.0178
$\Delta(L/P)_{cdt}$ (pre-sample: 1980-1990, 1990-2000)	1444	0.0055	0.0056	0.0136
$\Delta(L/P)_{cdt}$ (private sector)	1444	0.0269	0.0246	0.0196
$\Delta(L/P)_{cdt}$ (public sector)	1444	0.0040	0.0037	0.0113
$\Delta(L/P)_{cdt}$ (primary sector)	1444	0.0018	0.0022	0.0127
$\Delta(L/P)_{cdt}$ (secondary sector)	1444	-0.0101	-0.0069	0.0233
$\Delta(L/P)_{cdt}$ (manufacturing sector)	1444	-0.0123	-0.0086	0.0216
$\Delta(L/P)_{cdt}$ (terziary sector)	1444	0.0393	0.0368	0.0254
$\Delta(L/P)_{cdt}$ (high-skill)	1444	0.0349	0.0339	0.0177
$\Delta(L/P)_{cdt}$ (low-skill)	1444	-0.0041	-0.0070	0.0269
$\Delta(L/P)_{cdt}$ (high-skill, non-production)	1444	0.0338	0.0327	0.0174
$\Delta(L/P)_{cdt}$ (high-skill, production)	1444	0.0011	0.0010	0.0023
$\Delta(L/P)_{cdt}$ (low-skill, non-production)	1444	0.0088	0.0071	0.0283
$\Delta(L/P)_{cdt}$ (low-skill, production)	1444	-0.0128	-0.0111	0.0180
$\Delta(L/P)_{cdt}$ (male)	1444	0.0137	0.0137	0.0129
$\Delta(L/P)_{cdt}$ (female)	1444	0.0172	0.0159	0.0147
$\Delta(L/P)_{cdt}$ (age 16-24)	1444	0.0013	0.0016	0.0119
$\Delta(L/P)_{cdt}$ (age 25-44)	1444	-0.0134	-0.0115	0.0342
$\Delta(L/P)_{cdt}$ (age 45+)	1444	0.0430	0.0473	0.0422
$\Delta(L/P)_{cdt}$ (1st quintile of initial wage distribution)	1444	0.0098	0.0089	0.0300
$\Delta(L/P)_{cdt}$ (2nd quintile of initial wage distribution)	1444	0.0082	0.0084	0.0128
$\Delta(L/P)_{cdt}$ (3rd quintile of initial wage distribution)	1444	0.0000	-0.0002	0.0126
$\Delta(L/P)_{cdt}$ (4th quintile of initial wage distribution)	1444	0.0061	0.0051	0.0124
$\Delta(L/P)_{cdt}$ (5th quintile of initial wage distribution)	1444	0.0054	0.0046	0.0208
$\Delta \ln L_{cdt}$	1444	0.0606	0.0512	0.0858
$\Delta(U/P)_{cdt}$	1444	-0.0054	-0.0072	0.0153
$\Delta(NILF/P)_{cdt}$	1444	-0.0255	-0.0223	0.0226
<u>AI Adoption</u>				
$\Delta Iado_{cdt}$	1444	0.0038	0.0026	0.0063
$\Delta Iado_{cdt}$ (data scientists only)	1444	0.0003	0.0001	0.0006
$\Delta Iado_{cdt}$ (manufacturing industries)	1444	-0.0001	0.0000	0.0012
$\Delta Iado_{cdt}$ (non-manufacturing industries)	1444	0.0039	0.0029	0.0059
$\Delta Iado_{cdt}$ (adjusted for industry employment growth)	1444	0.0010	0.0012	0.0053
$\Delta Iado_{cdt}$ (excluding top decile industries)	1444	0.0033	0.0024	0.0061
<u>AI Exposure</u>				
$\Delta Iexp_{cdt}$	1444	0.0002	0.0002	0.0004
$\Delta Iexp_{cdt}$ (data scientists only)	1444	0.0000	0.0000	0.0000
$\Delta Iexp_{cdt}$ (manufacturing industries)	1444	0.0000	0.0000	0.0004
$\Delta Iexp_{cdt}$ (non-manufacturing industries)	1444	0.0004	0.0003	0.0006
$\Delta Iexp_{cdt}$ (adjusted for industry employment growth)	1444	0.0006	0.0002	0.0031
$\Delta Iexp_{cdt}$ (excluding top decile industries)	1444	0.0001	0.0001	0.0003
$\Delta Iexp_{cdt}$ (1990 industry shares)	1444	0.0003	0.0002	0.0005
$\Delta Iexp_{cdt}$ (leave-one-out)	1444	0.0017	-0.0026	0.0650

All statistics are computed on a sample of 722 CZs observed over two decades, 2000-2010 and 2010-2020. The subscripts c , d and t denote CZs, Census Divisions and decades, respectively. L , U , $NILF$ and P denote employment, unemployment, not in labor force and population, respectively. $\Delta Iado_{cdt}$ and $\Delta Iexp_{cdt}$ are the measures of AI adoption and AI exposure defined in eq. (10) and (12), respectively.

Table B2: Software used to Identify the AI-Related Occupations

Amazon Redshift	GitHub	Oracle PL/SQL
Amazon Simple Storage Service S3	Go	PHP
Amazon Simple Storage Service S4	JavaScript	Perl
Amazon Web Services AWS CloudFormation	JavaScript Object Notation JSON	PostgreSQL
Amazon Web Services AWS software	Jenkins CI	Python
Ansible Software	Kubernetes	Ruby
Apache Hadoop	Microsoft .NET Framework	Scala
Apache Hive	Microsoft Azure software	Selenium
Apache Kafka	Microsoft PowerShell	ServiceNow
Apache Spark	Microsoft SQL Server	Splunk Enterprise
Atlassian Confluence	Microsoft SQL Server Reporting Services SSRS	Spring Boot
Atlassian JIRA	MongoDB	Spring Framework
Bash	NoSQL	Structured query language SQL
C	Node.js	Transact-SQL
C#	Objective C	TypeScript
C++	Oracle Database	UNIX
Docker	Oracle Java	Vue.js
Git	Oracle Java 2 Platform Enterprise Edition J2EE	jQuery

REFERENCES

- [1] Acemoglu, Daron (2022) “Harms of AI,” in Justin B. Bullock and others (eds), *The Oxford Handbook of AI Governance* (online edn, Oxford Academic).
- [2] Acemoglu, Daron, Autor, David, Hazell, Joe, Restrepo, Pascual (2022) “AI and Jobs: Evidence from Online Vacancies,” *Journal of Labor Economics*, 40(S1), S293-S340.
- [3] Acemoglu, Daron and Simon Johnson (2023) *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, John Murray Press.
- [4] Acemoglu, Daron, and Pascual Restrepo (2020) "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128(6): 2188-2244.
- [5] Adao, Rodrigo, Kolesar, Michal, and Morales, Eduardo (2019) "Shift-share designs: Theory and inference." *The Quarterly Journal of Economics*, 134(4), 1949-2010.
- [6] Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel (2019) “What Is the Impact of Automation on Employment? New Evidence from France,” CEPR Discussion Paper 14443.
- [7] Albanesi, Stefania, Antonio Dias da Silva Juan Francisco Jimeno Ana Lamo Alena Wabitsch (2023) “New Technologies and Jobs in Europe”, CEPR Press Discussion Paper No. 18220.
- [8] Alekseeva, Liudmila, Jose Azar, Mireia Gine, Sampsa Samila, and Bledi Taska (2021) "The Demand for AI Skills in the Labor Market." *Labour Economics*, 71.
- [9] Autor, David and David Dorn (2013) “The Growth of Low Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103(5), 1553-1597.
- [10] Autor, David, David Dorn and Gordon Hanson (2013) “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103(6), 2121-2168.

- [11] Bartik, Timothy (1991) “Who Benefits from State and Local Economic Development Policies?” W.E. Upjohn Institute.
- [12] Blanas, Sotiris, Gino Gancia, and Sang Yoon (Tim) Lee (2019) “Who is Afraid of Machines?” *Economic Policy*, 34(100), 627–690.
- [13] Blanchard, Olivier and Lawrence Katz (1992) “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1:1992.
- [14] Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger and Gino Gancia (2022a) “Robot Imports and Firm-Level Outcomes,” CEPR Discussion Paper 14593.
- [15] Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia and Ioannis Papadakis (2022b) “Robots, Offshoring and Welfare,” in “Robots and AI: a New Economic Era” (eds. Lili Yan Ing and Gene M. Grossman) Oxon and New York: Routledge, pp. 40-81.
- [16] Borusyak, Kirill, Peter Hull and Xavier Jaravel (2022) "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies*, 89(1), 181-213.
- [17] Bresnahan, Timothy (2019) "Artificial Intelligence Technologies and Aggregate Growth Prospects." Unpublished manuscript, Stanford.
- [18] Brynjolfsson, Erik, Tom Mitchell and Daniel Rock (2018) “What can machines learn, and what does it mean for occupations and the economy?” *AEA Papers and Proceedings*, 108, 43–47.
- [19] Brynjolfsson, Erik, Mitchell, Tom, & Rock, Daniel. (2019) “Machine Learning and Occupational Change,” Unpublished manuscript, MIT.
- [20] Conley, Timothy G., Christian B. Hansen and Peter E. Rossi (2012) “Plausibly Exogenous,” *The Review of Economics and Statistics*, 94(1), 260-272.
- [21] Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum and Nicole Woessner (2021) “The Adjustment of Labor Markets to Robots,” *Journal of the European Economic Association*, 19: 3104-3153.
- [22] Felten, Edward, Manav Raj and Robert Seamans (2018) “A method to link advances in artificial intelligence to occupational abilities,” *AEA Papers and Proceedings*, 108, 54–57.
- [23] Felten, Edward, Manav Raj, and Robert Seamans (2019) "The Effect of Artificial Intelligence on Human Labor: An Ability-based Approach," *Academy of Management Proceedings* Vol. 1.
- [24] Felten, Edward, Manav Raj, and Robert Seamans (2021) “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses,” *Strategic Management Journal*, 42, 2195–2217.
- [25] Gaggl, Paul and Greg C. Wright (2017) “A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive,” *American Economic Journal: Applied Economics*, 9(3), 262-94.

- [26] Graetz, Georg and Guy Michaels (2018). "Robots at Work," *Review of Economics and Statistics* 100(5): 753–68.
- [27] Grennan, Jillian, and Roni Michaely (2020) "Artificial Intelligence and High-Skilled Work: Evidence from Analysts," Swiss Finance Institute Research Paper No. 20-84.
- [28] Hanson, Gordon (2021) "Immigration and regional specialization in AI," NBER WP 28671.
- [29] Hui, Xiang, Oren Reshef and Luofeng Zhou (2023) "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market," CESifo WP 10601.
- [30] McKinsey Global Institute (2017) "Artificial Intelligence: The Next Digital Frontier?" Discussion Paper.
- [31] Ruggles, Steven, Sarah Flood, Matthew Sobek, Danika Brockman, Grace Cooper, Stephanie Richards and Megan Schouweiler (2013) "IPUMS USA: Version 13.0 [US Census of Population, 1990]," Minneapolis, MN: IPUMS.
- [32] Taddy, Matt (2018) "The technological elements of artificial intelligence," NBER WP 24301.
- [33] Tolbert Charles M. and Molly Sizer (1996) "US Commuting Zones and Labor Market Areas: A 1990 Update," Staff Reports 278812, United States Department of Agriculture, Economic Research Service.
- [34] Varian, Hal (2018) "Artificial intelligence, economics, and industrial organization," NBER WP 24839.
- [35] Webb, Michael. (2020) "The Impact of Artificial Intelligence on the Labor Market," Unpublished manuscript, Stanford.