

NBER

NATIONAL BUREAU of
ECONOMIC RESEARCH

National
Bureau of
Economic
Research

Studies in
Income
and
Wealth
Volume 83

Technology, Productivity, and Economic Growth

Edited by
Susanto Basu,
Lucy Eldridge,
John Haltiwanger,
and Erich Strassner

Technology, Productivity, and Economic Growth

NBER

NATIONAL BUREAU *of*
ECONOMIC RESEARCH

Studies in Income and Wealth

Volume 83

NBER

NATIONAL BUREAU of
ECONOMIC RESEARCH

Technology, Productivity, and Economic Growth

Edited by

**Susanto Basu, Lucy Eldridge,
John Haltiwanger, and Erich Strassner**

The University of Chicago Press

Chicago and London

The University of Chicago Press, Chicago 60637
The University of Chicago Press, Ltd., London
© 2025 by National Bureau of Economic Research
All rights reserved. No part of this book may be used or reproduced
in any manner whatsoever without written permission, except in the
case of brief quotations in critical articles and reviews. For more
information, contact the University of Chicago Press, 1427 E. 60th St.,
Chicago, IL 60637.
Published 2025
Printed in the United States of America

34 33 32 31 30 29 28 27 26 25 1 2 3 4 5

ISBN-13: 978-0-226-83908-0 (cloth)
ISBN-13: 978-0-226-83909-7 (e-book)
DOI: <https://doi.org/10.7208/chicago/9780226839097.001.0001>

Library of Congress Cataloging-in-Publication Data

Names: Basu, Susanto editor | Eldridge, Lucy editor | Haltiwanger,
John editor | Strassner, Erich editor

Title: Technology, productivity, and economic growth / edited
by Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich
Strassner.

Other titles: Studies in income and wealth v. 83.

Description: Chicago : The University of Chicago Press, 2025. | Series:
Studies in income and wealth ; volume 83 | Includes bibliographical
references and index.

Identifiers: LCCN 2024056936 | ISBN 9780226839080 cloth |
ISBN 9780226839097 ebook

Subjects: LCSH: Technological innovations | Research, Industrial |
Industrial productivity

Classification: LCC HD45 .T39683 2025 | DDC 658.4/063—dc23/
eng/20250207

LC record available at <https://lcn.loc.gov/2024056936>

© This paper meets the requirements of ANSI/NISO Z39.48-1992
(Permanence of Paper).

National Bureau of Economic Research

Officers

Peter Blair Henry, <i>Chair</i>	Kelly Horak, <i>Controller and Assistant Corporate Secretary</i>
Karen G. Mills, <i>Vice Chair</i>	Alterra Milone, <i>Corporate Secretary</i>
James M. Poterba, <i>President and Chief Executive Officer</i>	Denis Healy, <i>Assistant Corporate Secretary</i>
Robert Mednick, <i>Treasurer and Representative of the American Institute of Certified Public Accountants</i>	

Directors at Large

Kathleen B. Cooper	Karen N. Horn	Douglas Peterson
Charles H. Dallara	Lisa Jordan	Andrew Racine
Jessica P. Einhorn	Karin Kimbrough	John S. Reed
Mohamed El-Erian	William M. Lewis, Jr.	Hal Varian
Diana Farrell	John Lipsky	Mark Weinberger
Helena Foulkes	Laurence C. Morse	Martin B. Zimmerman
Esther George	Michael H. Moskow	
Peter Hancock	Alicia H. Munnell	

Directors by University Appointment

Timothy Bresnahan, <i>Stanford</i>	Richard Schmalensee, <i>Massachusetts Institute of Technology</i>
Alan V. Deardorff, <i>Michigan</i>	Christopher Sims, <i>Princeton</i>
Benjamin Hermalin, <i>California, Berkeley</i>	Richard Steckel, <i>Ohio State</i>
R. Glenn Hubbard, <i>Columbia</i>	Ann Huff Stevens, <i>Texas, Austin</i>
Samuel Kortum, <i>Yale</i>	Lars Stole, <i>Chicago</i>
George J. Mailath, <i>Pennsylvania</i>	Ingo Walter, <i>New York</i>
Angelo Melino, <i>Toronto</i>	David B. Yoffie, <i>Harvard</i>
Joel Mokyr, <i>Northwestern</i>	
John Pepper, <i>Virginia</i>	

Directors by Appointment of Other Organizations

Timothy Beatty, <i>Agricultural and Applied Economics Association</i>	Maureen O'Hara, <i>American Finance Association</i>
Constance Hunter, <i>National Association for Business Economics</i>	Dana M. Peterson, <i>The Conference Board</i>
Arthur Kennickell, <i>American Statistical Association</i>	Peter L. Rousseau, <i>American Economic Association</i>
Anne McCants, <i>Economic History Association</i>	Gregor W. Smith, <i>Canadian Economics Association</i>

Directors Emeriti

George A. Akerlof	Jacob Frenkel	Robert T. Parry
Peter C. Aldrich	Martin Gruber	Andrew Postlewaite
Jagdish Bhagwati	Robert S. Hamada	John J. Siegfried
Don R. Conlan	Saul H. Hymans	Craig Swan
George C. Eads	Marjorie B. McElroy	Marina v. N. Whitman
Ray C. Fair	Laurence H. Meyer	

Relation of the Directors to the Work and Publications of the NBER

1. The object of the NBER is to ascertain and present to the economics profession, and to the public more generally, important economic facts and their interpretation in a scientific manner without policy recommendations. The Board of Directors is charged with the responsibility of ensuring that the work of the NBER is carried on in strict conformity with this object.

2. The President shall establish an internal review process to ensure that book manuscripts proposed for publication DO NOT contain policy recommendations. This shall apply both to the proceedings of conferences and to manuscripts by a single author or by one or more co-authors but shall not apply to authors of comments at NBER conferences who are not NBER affiliates.

3. No book manuscript reporting research shall be published by the NBER until the President has sent to each member of the Board a notice that a manuscript is recommended for publication and that in the President's opinion it is suitable for publication in accordance with the above principles of the NBER. Such notification will include a table of contents and an abstract or summary of the manuscript's content, a list of contributors if applicable, and a response form for use by Directors who desire a copy of the manuscript for review. Each manuscript shall contain a summary drawing attention to the nature and treatment of the problem studied and the main conclusions reached.

4. No volume shall be published until forty-five days have elapsed from the above notification of intention to publish it. During this period a copy shall be sent to any Director requesting it, and if any Director objects to publication on the grounds that the manuscript contains policy recommendations, the objection will be presented to the author(s) or editor(s). In case of dispute, all members of the Board shall be notified, and the President shall appoint an ad hoc committee of the Board to decide the matter; thirty days additional shall be granted for this purpose.

5. The President shall present annually to the Board a report describing the internal manuscript review process, any objections made by Directors before publication or by anyone after publication, any disputes about such matters, and how they were handled.

6. Publications of the NBER issued for informational purposes concerning the work of the Bureau, or issued to inform the public of the activities at the Bureau, including but not limited to the NBER Digest and Reporter, shall be consistent with the object stated in paragraph 1. They shall contain a specific disclaimer noting that they have not passed through the review procedures required in this resolution. The Executive Committee of the Board is charged with the review of all such publications from time to time.

7. NBER working papers and manuscripts distributed on the Bureau's web site are not deemed to be publications for the purpose of this resolution, but they shall be consistent with the object stated in paragraph 1. Working papers shall contain a specific disclaimer noting that they have not passed through the review procedures required in this resolution. The NBER's web site shall contain a similar disclaimer. The President shall establish an internal review process to ensure that the working papers and the web site do not contain policy recommendations, and shall report annually to the Board on this process and any concerns raised in connection with it.

8. Unless otherwise determined by the Board or exempted by the terms of paragraphs 6 and 7, a copy of this resolution shall be printed in each NBER publication as described in paragraph 2 above.

Contents

	Prefatory Note	ix
	Introduction	1
	Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner	
1.	Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey	13
	Daron Acemoglu, Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas	
2.	Similarities and Differences in the Adoption of General Purpose Technologies	57
	Ajay Agrawal, Joshua S. Gans, and Avi Goldfarb	
3.	Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition	75
	Tania Babina, Anastassia Fedyk, Alex He, and James Hodson	
4.	The Strange and Awful Path of Productivity in the US Construction Sector	119
	Austan Goolsbee and Chad Syverson	

5.	Digital Concrete: Productivity in Infrastructure Construction	139
	Diane Coyle and Rehema Msulwa	
6.	After Redefinition TFP Accounting	169
	Jon D. Samuels	
7.	Data, Intangible Capital, and Productivity	193
	Carol Corrado, Jonathan Haskel, Massimiliano Iommi, Cecilia Jona-Lasinio, and Filippo Bontadini	
8.	Valuing the US Data Economy Using Machine Learning and Online Job Postings	245
	José Bayoán Santiago Calderón and Dylan G. Rassier	
9.	An Occupation and Asset-Driven Approach to Capital Utilization Adjustment in Productivity Statistics	285
	Josh Martin and Kyle Jones	
10.	Opening the Black Box: Task and Skill Mix and Productivity Dispersion	323
	G. Jacob Blackwood, Cindy Cunningham, Matthew Dey, Lucia Foster, Cheryl Grim, John Haltiwanger, Rachel Nesbit, Sabrina Wulff Pabilonia, Jay Stewart, Cody Tuttle, and Zoltan Wolf	
	Author Index	361
	Subject Index	367

Prefatory Note

This volume contains revised versions of the papers presented at the Conference on Research in Income and Wealth titled “Technology, Productivity, and Economic Growth,” held in Washington, DC, on March 17–18, 2022.

We are grateful for support for this project from the Bureau of Economic Analysis. Support for the general activities of the Conference on Research in Income and Wealth is provided by the following agencies: Bureau of Economic Analysis, Bureau of Labor Statistics, the Census Bureau, the Board of Governors of the Federal Reserve System, the Statistics of Income/Internal Revenue Service, and Statistics Canada.

We thank Susanto Basu, Lucy P. Eldridge, John C. Haltiwanger, and Erich H. Strassner, who served as conference organizers and as editors of the volume.

Executive Committee, February 2023

Katharine G. Abraham (chair)	Ron S. Jarmin
John M. Abowd	J. Bradford Jensen
Vipin Arora	Barry Johnson
Susanto Basu	Andre Loranger
Ernst R. Berndt	Aysegül Şahin
Alberto Cavallo	Peter K. Schott
Carol A. Corrado	Daniel E. Sichel
Lucy P. Eldridge	William Wascher
John C. Haltiwanger	

Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition

Tania Babina, Anastassia Fedyk, Alex He,
and James Hodson

The arrival of new general purpose technologies (GPT) is a key driver of economic growth (Romer 1990; Aghion and Howitt 1992; Kogan et al. 2019, 2017). Yet as firms adapt their production processes and organization in response to technological changes, this shift raises major concerns about the impact on workers. For example, computer software and robots have displaced low- and medium-skilled workers (Autor et al. 2003; Acemoglu and Restrepo 2020), while the arrival of cotton-spinning machinery and electricity led to a complete re-organization of production processes within firms (Fizsbein et al. 2020; Juhász et al. 2024). In recent years, the focus has shifted to a new technological wave: artificial intelligence and related “big data” technologies. AI is a prediction technology, and predictions are at the heart of decision-making under uncertainty (Agrawal et al. 2019), making AI applicable to solve a variety of business problems with different potential effects on labor. On the one hand, firms might use AI to automate high-skilled tasks (Webb 2020).¹ On the other hand, investments in AI so far have been primarily associated with product innovation (Babina et al.

Tania Babina is an associate professor at the University of Maryland, an affiliate of the Centre for Economic Policy Research, and a faculty research fellow of the National Bureau of Economic Research.

Anastassia Fedyk is an assistant professor of finance at the Haas School of Business at the University of California, Berkeley.

Alex He is an assistant professor of finance at the University of Maryland.

James Hodson is the chief executive officer of the AI for Good Foundation and a scientific advisor at Cognism.

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

2024), which can require complementary investments and more educated workforces. Indeed, Babina et al. (2024) document that AI-investing firms experience increases in their overall employment—and it is an open question whether (and how) this employment growth is associated with changes in labor composition.

In this paper, we examine whether firms that invest more heavily in AI technologies experience changes in labor composition and workforce organization. To date, in-depth empirical understanding of the relationship between firms' investments in new AI technologies and firms' labor composition has remained elusive due to two key challenges: the difficulty of measuring firm-level AI investments and the lack of granular data on firms' labor composition and labor organization (Seamans and Raj 2018; Frank et al. 2019). Several recent papers make progress in overcoming the first challenge by using firms' job postings and worker resumes to identify the hiring and stock of AI-skilled labor (Acemoglu et al. 2022a; Babina et al. 2024; Alekseeva et al. 2020; Abis and Veldkamp 2024).² The main contribution of this paper lies in overcoming the second challenge. We use matched employer-employee US data based on worker resumes, including detailed information on both individual jobs and employees' educational backgrounds. Using these data, we construct firm-level measures of labor composition and workforce organization, including employees' educational backgrounds and hierarchical positions, and link them to firm-level AI investments.

We find that firms that invest more in AI significantly shift toward more educated workforces, with greater emphasis on STEM degrees. At the same time, AI-investing firms become less top-heavy in terms of their organizational structure, with increasing shares of junior employees and less emphasis on middle-management and senior roles. Overall, our findings suggest that investments in AI are associated with major changes in firms' labor composition and organization, translating into a broader shift toward more junior employees with high educational attainment and technical expertise.

To construct workforce composition measures and assess their relationships to firms' AI investments, we leverage a unique combination of data sets that capture both the *stock* of current employees and the *demand* for new employees among US firms. The stock of current employees at each point in time comes from a resume data set provided by Cognism Inc., which offers job histories for 535 million individuals globally. Cognism resume data offer a complementary perspective to granular administrative firm-worker matched US data, which contain individual workers' wages but do not feature comprehensive information on individual workers' educational backgrounds or job characteristics. Cognism resume data, while representing more than 63 percent of full-time US employment as of 2018, offer detailed

2. These measures are complementary to surveys of AI use by firms (Zolas et al. 2020; Acemoglu et al. 2022b) and data on AI-related patents (Alderucci et al. 2020).

job titles and descriptions (from which Cognism infers hierarchical positions) and educational backgrounds including degree-granting institutions and majors. We complement the resume data with information on firms' demand for new workers from the job postings data provided by Burning Glass, which capture 180 million online job vacancies. While job postings data have been instrumental in understanding how firms target their new hiring, the resume data provide a full picture of what happens to the overall workforce within firms—including new hires and potential displacement. Since our goal is to study the impact of AI on AI-using firms rather than AI-producing firms, we exclude firms in the tech sectors, which are likely to be producers of new AI tools.

We begin our analysis by describing how firm ex-ante labor composition predicts future growth in firms' AI investments. We adopt the novel measure of AI investments proposed by Babina et al. (2024), based on firms' AI-skilled human capital. The human-capital-based approach is motivated by the heavy reliance of AI implementation on human expertise. The method first identifies skills that are empirically related to principal AI technologies (machine learning, computer vision, and natural language processing) from the Burning Glass job postings data and then uses the identified highly AI-related skills to classify AI-related workers in the Cognism resume data. At the firm level, growth in AI investments is more pronounced among firms that initially have more workers with doctoral degrees and STEM majors. This is in line with the evidence in Babina et al. (2024), who find that firms with more technical workers and more educated workers are able to attract AI talent more easily. The hierarchical structure of firms' labor organization—as measured by the shares of employees in junior, middle-management, and senior roles—does not significantly predict growth in AI investments.

We next address our main question: whether AI investments are associated with changes in labor composition and workforce organization. We consider three sets of outcomes related to firms' workforce composition and organization. First, motivated by the literature on technologies and firm organization (Acemoglu et al. 2007), we examine changes in firms' organizational structure using measures of organizational structure from the resume data. The relationship between technological investments and the relative weights of different hierarchical levels is ex ante ambiguous. As highlighted by theoretical work (Garicano and Rossi-Hansberg 2006; Bloom et al. 2014), different types of technologies can have opposing effects on the need for managerial layers. Second, we look at both workers' education levels in the resume data and educational requirements in the job postings data to test whether AI facilitates skill-biased technological change (Autor et al. 1998; Machin and Van Reenen 1998) or replaces high-skilled labor as predicted by Webb (2020). Here, too, the predicted effect of AI is ex ante ambiguous, and we provide the first systematic evidence of its direction. The shifts in

labor composition are likely to go hand-in-hand with changes to organizational structure, as Caroli and Van Reenen (2001) point out that flattening hierarchical structures require higher human capital from each individual employee. Finally, we use detailed information on workers' majors and required skills to study how AI changes firms' demand for different types of labor.

Our main empirical specification is a long-differences regression of changes in labor outcomes from 2010 to 2018 on changes in the firm-level share of AI workers during the same period, following the standard approach in settings with slow-moving processes like technological change (e.g., Acemoglu and Restrepo 2020). As shown in Babina et al. (2024), AI investments accumulate gradually over time and generate effects that are not immediate, making the long-differences strategy well suited for our setting. Furthermore, by taking first differences in independent and dependent variables, the long-differences specification ensures that time-invariant firm characteristics do not drive the results. To bolster the causal interpretations of the results, we include a rich set of controls featuring industry fixed effects and firm-, industry-, and commuting-zone-level characteristics in 2010. All of our coefficient estimates are remarkably consistent across specifications with and without these detailed controls. Moreover, none of our results are driven mechanically by the hiring of AI workers, and excluding those workers from the calculation of dependent variables produces similar results. Finally, we show that our results are robust to using an instrumental variable strategy based on an instrument that isolates the variation in firms' AI investments that is driven by the supply of AI-skilled labor.

In terms of hierarchical structure, we provide evidence that AI investments are associated with firms becoming flatter, with higher shares of employees in entry-level or single-contributor roles and fewer employees in either middle-management or senior roles. Specifically, a one-standard-deviation change in the share of AI workers at a firm is associated with a 1.6 percent increase in the share of junior employees from 2010 to 2018, while middle management declines by 0.8 percent and senior management by 0.7 percent. This result is consistent with the channel suggested by Garicano and Rossi-Hansberg (2006) and explored by Bloom et al. (2014), where reductions in costs of accessing knowledge through improved data processing, such as AI technology, result in increased problem-solving ability of employees at all levels, leading to increased span of control and less reliance on top-heavy hierarchical structures.

In terms of labor composition, we observe a general upskilling trend associated with larger AI investments. Firms that invest more in AI tend to increase their shares of workers with bachelor's, master's, and doctoral degrees (correspondingly decreasing the share of workers without college education). A one-standard-deviation increase in the firm's share of AI workers translates into a 3.7 percent increase in the share of workers whose

maximal educational attainment is an associate or bachelor's degree, a 2.9 percent increase in the share of workers whose maximal educational attainment is a master's degree, and a 0.6 percent increase in doctoral degrees. These increases in educated workers correspond to a 7.2 percent decline in the share of workers without college education. The upskilling shifts in education are also observed in the firms' explicit labor demand in the Burning Glass job postings, which feature both required education and required number of years of prior experience for prospective job applicants. For example, a one-standard-deviation increase in the share of AI workers is associated with a 0.5 additional year of required education in the firm's new job openings.

The additional demand for educated workers in AI-investing firms tends to concentrate in technical fields. Leveraging the information on majors of the most recent degree for each individual employee in the resume data, we observe that AI investments are associated with a significant increase in the share of employees with majors in STEM degrees and a corresponding decline in the share of employees with degrees in social science fields. Similarly, the skill requirements in Burning Glass job postings reveal that AI-investing firms experience a significant increase in demand for employees with skills in data analysis and IT, while decreasing their search for employees with skills in traditional operational fields such as finance and maintenance.

Our work contributes to the recent literature on the impact of AI technologies on the labor market. Previous literature has conjectured that AI has the potential to displace some human tasks, including high-skilled tasks (Acemoglu and Restrepo 2019a; Webb 2020; Frank et al. 2019; Mihet and Philippon 2019). Empirically, prior work made progress in measuring exposure to AI at the occupation level (Felten et al. 2018; Brynjolfsson et al. 2018; Webb 2020) and the impact of AI on overall labor demand and employment at the firm level (Alderucci et al. 2020; Rock 2019; Babina et al. 2024). Several papers look at the impact of AI on labor in specific settings, such as financial analysts and startups (Abis and Veldkamp 2024; Cao et al. 2021; Gofman and Jin 2022; Grennan and Michaely 2022).³ Our work is closest to Acemoglu et al. (2022a), who use firms' occupational structure to proxy for *exposure* to potential displacement by AI and explore how this exposure relates to firms' labor demand. By contrast, we study the effect of firms' overall AI *investments*—including applications of AI that aim to displace workers and those that do not (e.g., AI-fueled product innovation in Babina et al. 2024). Furthermore, we complement prior work that looks at labor demand by examining changes in actual worker composition, including new hires and departures.

3. Relatedly, other papers compare decision-making of humans and AI algorithms in various settings (e.g., D'Acunzio et al. 2019; Fuster et al. 2020; Jansen et al. 2023; Erel et al. 2021; Lyonnet and Stern 2022).

Our paper is the first to document the relationship between the use of AI technologies and workforce composition at the firm level. While Babina et al. (2024) show that AI investments increase total firm employment, our evidence further shows that this increase is concentrated in highly educated workers and high-skill workers with STEM backgrounds and IT skills. A potential explanation for our findings is that AI-fueled product innovation—the main channel through which AI investments power firm growth (Babina et al. 2024)—increases firms' demand for complementary skilled labor. These results contribute to the literature on the labor market effects of general purpose technologies, which shows that technologies like IT and electricity favor high-skilled labor but displace medium-skilled workers (Autor et al. 1998, 2003; Fizsbein et al. 2020; Zator 2019; Acemoglu and Restrepo 2022). Bessen et al. (2022) find that IT investments are associated with an increase in the returns to skill at the firm level. We show that AI investments are associated with an overall increase in firms' hiring of skilled labor, but these effects are heterogeneous. Demand for some high-skilled labor (e.g., STEM majors, IT skills) rises, while demand for other medium-skilled or high-skilled labor (e.g., finance, maintenance) declines.

Our evidence on firms' hierarchical structures also contributes to the literature on technology adoption and firm organization (Hitt 1999; Acemoglu et al. 2007; Bloom et al. 2014; McElheran and Forman 2019). We find that firms investing in AI technologies become less top-heavy, which is similar to the previously documented effect of IT but opposite to the effect of communication technologies. The combination of our results linking AI to flatter organizational structure and increased demand for skilled labor echoes the finding in Caroli and Van Reenen (2001) that there are complementarities between organizational change and skilled employees. Our results therefore support the notion that new technologies such as AI can be an important driver of skill-biased organizational change.

Methodologically, we provide a new measure of firms' organizational structure based on text descriptions of jobs in worker resumes. This measure can be applied to all firms across all industries and complements previous survey-based measures of firm hierarchy (e.g., Rajan and Wulf 2006; Bloom et al. 2012). More broadly, our method contributes to the growing literature that uses textual analysis to measure job tasks and skills (Kogan et al. 2019; Fedyk and Hodson 2023; Jiang et al. 2024; Hansen et al. 2021).⁴

The remainder of the paper is organized as follows. We introduce our data in section 3.1 and detail our methodology for measuring AI investments and

4. Our methodology also contributes to the literature that uses information about firms' employees to proxy for intangible capital. See Eisfeldt and Papanikolaou (2013); Crouzet and Eberly (2018); Peters and Taylor (2017) for recent examples of papers on intangible capital.

workforce composition in section 3.2. Section 3.3 explores how firms' initial workforce composition predicts AI investments, while section 3.4 presents our main results on the relationship between AI investments and changes in workforce composition and organization. Section 3.5 concludes.

3.1 Data

To investigate how the composition and structure of firms' workforces changes in firms that invest more heavily in AI, we bring together two data sets. First, we take advantage of a unique matched employer-employee data set built from resumes and featuring individual employees' detailed job descriptions and educational backgrounds. Second, we supplement the resume data with a comprehensive data set of job postings revealing firms' demand for education and skills.

3.1.1 Employment profiles from Cognism

We leverage the employment profile (resume) data set from Cognism, which offers matched employer-employee data covering approximately 535 million individuals globally. These data are introduced in detail in Fedyk and Hodson (2023) and Babina et al. (2024) and bring several key advantages that complement existing administrative data. First, Cognism offers broad coverage in the United States. Appendix figure A.1 compares the coverage of US full-time employment in the Cognism data against official numbers from the Bureau of Labor Statistics.⁵ Cognism captures 42 percent of all US employment in 2010 (the beginning of our sample), and the coverage steadily increases to 63 percent in 2018 (the end of our sample), with the average coverage being 53 percent across these years.⁶

Second, while Cognism does not have information on wages (as would be included, for example, in the US Census Bureau's Longitudinal Employer-Household Dynamics program), the Cognism data provide detailed information on individual workers' occupations, job tasks, and educational backgrounds—the kind of information that is not available in administrative data. Specifically, for each individual, we observe the start and end dates of each job, the job title (often along with a detailed job description), each job's company name and location, the individual's educational record (with university names, degrees, and majors), as well as any patents, awards, or publications that the individual chooses to include on the resume. This allows us to examine how firms' investments in new technologies, such as AI, interplay

5. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

6. Although our Cognism data snapshot is from July 2021, we follow Tambe et al. (2020) and Babina et al. (2024) and only use the years through 2018 to avoid potential noise from workers updating their resumes with a delay.

with granular changes in their workforce composition, including employee educational attainment, specialization, and seniority.

Finally, the Cognism data also bring advantages relative to the job postings data that have been previously used to understand the impact of AI on the labor markets (e.g., Alekseeva et al. 2020; Acemoglu et al. 2022a). Working directly with employee resumes enables us to see who is actually working at each firm, rather than only firms' demand for employees. As a result, we are able to capture changes in workforce composition that occur outside of new hiring (e.g., promotions, onboarding of new employees through acquisitions, or layoffs of existing employees).

Cognism's AI Research department leverages techniques from machine learning and natural language processing, including named entity disambiguation and graph-based modeling methods, to further enrich the resume data by normalizing job titles and occupations, associating employees with functional divisions and teams within each firm, and identifying institutions, degrees, and majors from education records. We match employer names in the Cognism data to the names of publicly traded firms in the Compustat data set using the approach developed in Fedyk and Hodson (2023). The matching of individual resumes to firm entities is performed dynamically to account for acquisitions and divestitures. We limit our attention to public firms with data in Compustat in order to link individual employees to firms and include detailed controls for other firm characteristics (e.g., sales, cash reserves, R&D expenditures, and markups). The data cover 657 million US-based person-firm-year observations between 2007 and 2018, of which 120 million (18 percent) are matched to US public firms. This is consistent with prior statistics showing that publicly listed firms account for approximately 26 percent of overall US employment (Davis et al. 2006). The sample of 120 million person-firm-years matched to US public firms corresponds to 19 million distinct individual employees.

We benchmark the resulting sample of Cognism employees at public firms against these firms' employment in the Compustat database. Appendix figure A.2 presents the median and interquartile range of firm-level coverage rates for each year from 2010 to 2018.⁷ It is important to note that our analyses are performed on US-based employees in the Cognism data; however, US public firms do not report the number of US-based employees and instead only report their global employment numbers. Therefore, the coverage rate in appendix figure A.2 is lower than that in appendix figure A.1, where the denominator consists of US-based employment numbers. The coverage rate in appendix figure A.2 is stable throughout our sample period, with the median firm having around 30 percent of its global employment captured in the Cognism data set, and the interquartile range being 15–70 percent.

7. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

Furthermore, appendix table A.1 reports the median coverage rate in each industry sector for our Compustat-Cognism matched data. The industry median coverage rates range from 18 percent (Health Care) to 52 percent (Finance and Insurance).

3.1.2 Job Postings from Burning Glass

The second data set we use covers over 180 million job postings in the United States in 2007 and 2010–2018. The data set is provided by Burning Glass Technologies (BG), which examines over 40,000 online job boards and company websites, collects the job postings data, parses them into a machine-readable form, and uses the data to construct labor market analytics products. BG employs a sophisticated deduplication algorithm to avoid double-counting vacancies that post on multiple job boards. BG data are quite comprehensive, covering approximately 60–70 percent of all vacancies posted in the United States, either online or offline.⁸ The data contain detailed information for each job posting, including the job title, location, occupation, and employer name. Most importantly for our paper, the job postings are tagged with (i) thousands of specific skills standardized from the open text in each job opening, and (ii) specific requirements such as years of education and experience.

We focus on jobs with non-missing employer names (approximately 65 percent of all job postings) and at least one required skill (which corresponds to 93 percent of all job postings). Since we are interested in the composition of a firm's core workforce, we drop job postings that are internships. We match the employer firms in the remaining job postings to Compustat firms using fuzzy matching after stripping out common endings such as "Inc." and "L.P." For observations that do not match exactly on firm name, we manually assess the top ten potential fuzzy matches based on the firm name, industry, and location. Out of 112 million job postings with non-missing employer names and skills, 42 million (38 percent) are matched to Compustat firms. This slightly over-represents employees of publicly listed firms, which constitute just over one-fourth of US employment in the non-farm business sector (Davis et al. 2006).

3.1.3 Additional Data Sources

We merge the Cognism resume data and the Burning Glass job postings data to several additional data sources. We collect commuting-zone-level wage and education data from the Census American Community Surveys (ACS) and industry-level wages and employment data from the Census Quarterly Workforce Indicators (QWI). Firm-level operational variables (e.g., sales, cash, assets) come from Compustat.

8. See Hershbein and Kahn (2018) for a detailed description of the BG data, including their representativeness, which is stable over time at the occupation level.

3.2 Methodology and Descriptive Statistics

3.2.1 AI Investments

We leverage the methodology proposed by Babina et al. (2024) to measure firms' investments in AI based on their intensity of AI-skilled hiring. The intuition is that successful implementation and use of AI technologies by firms requires employees with expertise in AI methods. Since other inputs to AI, such as data and computing infrastructure, are complementary to AI-skilled labor, our human-capital based measure allows us to capture the relative intensity of AI investments across firms.⁹

In order to identify AI expertise, we take advantage of (i) the detailed information on required skills in the job postings data and (ii) new, data-driven methodology for identifying AI-related jobs. Previous methods for classifying job postings based on the presence of key terms from a pre-specified list (e.g., Hershbein and Kahn 2018; Alekseeva et al. 2020) are likely to suffer from both Type I (incorrectly labeling tangentially related employees as AI-related) and Type II (missing real AI skills that did not make the initial dictionary) errors due to the arbitrariness of the list of keywords. This is especially relevant in a quickly evolving domain such as AI, where new emerging skills can be easily missed. The methodology from Babina et al. (2024) circumvents these challenges by learning the AI-relatedness of each of approximately 15,000 unique skills directly from the job postings data, based on their empirical co-occurrence (within required lists of skills across job postings) with unambiguous core AI skills. We then take the skills that are empirically most related to the core AI skills and search for those in our resume data. Finally, we aggregate the worker-level data to the firm-year level by calculating the share of the firm's employees who are AI-skilled.

More precisely, we start by measuring the AI-relatedness of each skill in the job postings data by calculating that skill's co-occurrence with Artificial Intelligence (AI) and its three main sub-fields: machine learning (ML), natural language processing (NLP), and computer vision (CV):

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

Intuitively, this measure captures how correlated each skill s is with the core AI skills. For example, the skill "Recurrent Neural Network" has a value of 0.965, which means that 96.5 percent of job postings that list "Recurrent Neural Network" as a required skill also require one of the core AI skills or contain one of the core AI skills in the job title. Thus, a "Recurrent Neu-

9. It is possible that external AI software and solutions (e.g., IPSoft Amelia) may substitute for the hiring of AI-skilled labor. However, Babina et al. (2024) show that the use of external AI software solutions tends to be complementary to internal AI hiring and incorporating those in the measure of AI investments yields similar results.

ral Network” requirement in a job posting is highly indicative of that job being AI-related. On the other hand, the AI-relatedness measure of the skill “Microsoft Office” is only 0.003. In appendix table A.2, we list the skills with the highest AI-relatedness measures—namely, the skills that co-occur with the core AI skills in at least 70 percent of all job posting.¹⁰

In the Cognism resume data, we identify AI-skilled employees as those whose job positions directly involve AI. We begin with the set of 67 keywords in table A.2, which have the highest skill-level AI-relatedness measures in the job postings data. We then consider every employment record of each individual in the resume data and identify whether any of these AI-related terms appear in: (i) the job title or description; (ii) any patents obtained during the year of interest or the two following years (to account for the time lag between the work and the patent grant); (iii) any publications during the year of interest or the following year; or (iv) any awards received during the year of interest or the following year. If any of these conditions are met, then we classify that employee at that firm in that year as AI-skilled. For example, jobs with titles such as “senior **machine learning** developer” or publications such as “A new cluster-aware regularization of **neural networks**” are identified as AI jobs.

To aggregate to the firm level, we use the number of AI-related employees and the number of total employees at each firm in each year and compute the fraction of employees of that firm in that year who are classified as AI-skilled. Given that our empirical analyses focus on US-listed firms, our firm-level measure focuses on the employees who are based in the United States. Babina et al. (2024) provide a detailed discussion of this measure, perform multiple validation exercises, and offer detailed case studies of AI investments by individual firms in our sample. For the sake of brevity, we do not reproduce that analysis in this paper.

3.2.2 Labor Composition

We use the resume data to examine three aspects of firms’ workforces: (i) educational attainment in terms of college and post-graduate degrees; (ii) specialization in terms of college majors (e.g., STEM vs. humanities vs. social science); and (iii) hierarchical structure in terms of the composition of employees across different levels of seniority. We describe the construction of each of these variables in turn below.

3.2.2.1 Educational Attainment

Cognism uses the educational information from the resumes to classify each individual at each point in time based on that individual’s highest educational attainment to date. The categories are: (i) no secondary education; (ii) associate’s degree; (iii) bachelor’s degree; (iv) master’s degree other

10. Reproduced from Babina et al. 2024, <http://www.nber.org/data-appendix/c14753/appendix.pdf>.




Fig. 3.1 Distribution of education levels in the Cognism resume data

This figure shows the fraction of workers in each education level (no college, college, master’s, and doctoral) in the Cognism data. The four education levels are mutually exclusive and each worker is counted once for their highest level of education. The sample includes all workers in Compustat firms between 2010 and 2018.

than an MBA; (v) MBA; and (vi) doctoral degree (including Ph.D. and J.D. degrees). For each firm in our sample, we compute four educational attainment variables: (i) the share of employees in each year who have a college degree (either a bachelor’s or an associate’s); (ii) the share of employees who have at least a master’s degree; (iii) the share of employees who have a doctoral degree; and (iv) the share of employees who do not have a college degree. Figure 3.1 plots the mean of these four shares in the resume data for the sample of Compustat firms.

3.2.2.2 Educational Specialization

Cognism extracts major information from individuals’ education records and groups majors into broad categories of (i) Humanities, (ii) Social Sciences, (iii) Science, Technology, Engineering, and Mathematics (STEM), (iv) Fine Arts, and (v) Medicine. We take these broad categories and compute, for each firm in each year, the share of current employees whose most recent degrees fall in each category. Figure 3.2 plots the distribution of majors based on the most recent degree in the resume data for the sample of Compustat firms.

3.2.2.3 Seniority

The Cognism data are enriched with state-of-the-art machine learning techniques to identify employees’ departments and seniority. First, over 20,000 individual job titles are classified manually based on markers of seniority and department. The remaining job titles are then classified into

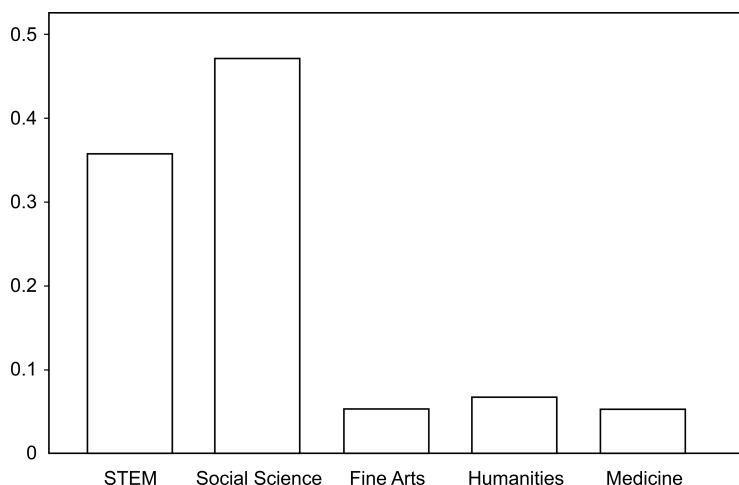


Fig. 3.2 Distribution of majors in the Cognism resume data

This figure shows the fraction of workers in each major (STEM, social science, fine arts, humanities, and medicine) in the Cognism data. STEM includes engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., math, physics, chemistry, computer science, statistics), and biological sciences (e.g., biology, pharmacology). The majors are mutually exclusive; for each worker, we record the major of the most recent degree earned. The sample includes all workers in Compustat firms between 2010 and 2018.

departments using a probabilistic language model and into seniority levels using an artificial neural network. There are six levels of seniority in total: (1) entry-level positions where individuals start straight out of undergraduate or high school education; (2) experienced staff in roles such as individual senior contributor but not managing others; (3) team leads who manage others but have little to no company-level decision-making responsibility; (4) middle management roles that oversee several smaller teams; (5) leadership positions that head larger departments or business segments; and (6) executive-level leadership such as the Chief Executive Officer and Chief Operating Officer. Fedyk et al. (2022) perform an evaluation of Cognism's seniority classification on the sample of accounting firms by assessing the model's output against a manually reviewed sample of over 10,000 positions. They find that Cognism's seniority classification has an accuracy rate of over 95 percent. In this paper, we group the seniority levels into three broader bands: low (consisting of entry-level positions and experienced individual contributors), medium (team leads and middle management), and high (leadership and the executive level). Figure 3.3 plots the shares of workers in each of these three seniority levels based on the resume data for the sample of Compustat firms. Overall, over 60 percent of employees are in junior-level or non-supervisory roles, with approximately 25 percent in mid-tier and 10–15 percent in senior management.

Other measures of firm organization structure, such as Rajan and Wulf

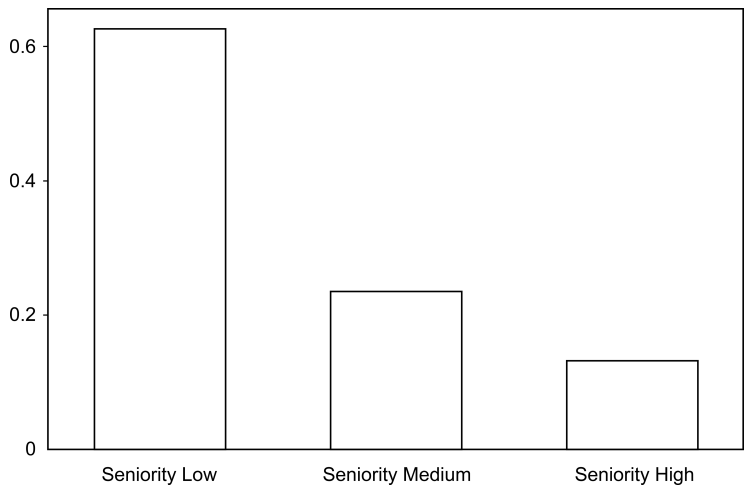


Fig. 3.3 Distribution of seniority levels in the Cognism resume data

This figure shows the fraction of workers in each seniority level (low, medium, and high) in the Cognism data. The seniority levels are described in text. The sample includes all workers in Compustat firms between 2010 and 2018.

(2006) and Guadalupe and Wulf (2010), are typically based on surveys of firm managers and focus on the breadth and depth of the hierarchy among managers. In particular, breadth is represented as the CEO’s span of control, and depth is represented by the number of levels between the CEO and divisional managers.¹¹ Our measure is related to these measures: for example, flatter organizational hierarchy (i.e., more entry-level workers and fewer senior managers) is associated with more breadth and less depth. In addition, our measure offers several advantages. First, it is based on the actual organizational structure of the firm, complementing survey-based measures. Second, our measure captures the composition of workers across all hierarchical levels, which provides a more holistic picture of the firm’s organizational structure than the number of levels between the CEO and lower-level managers or workers.

3.2.3 Labor Demand

We use the job postings data from Burning Glass to measure two aspects of firms’ labor demand: (1) required education and experience and (2) required skills. Since these measures are calculated from firms’ job postings, they only measure firms’ labor demand—the types of workers firms wish to hire—instead of the types of workers working at each firm.

11. Other surveys, such as the World Management Survey (<https://worldmanagementsurvey.org/>), also measure the number of levels between the CEO and entry-level workers.

3.2.3.1 *Required Education and Experience*

For each job posting, Burning Glass codes the minimum years of required education and the minimum years of required experience. 59 percent of job postings specify an educational requirement, which averages 14.5 years of school. 52 percent of job postings specify a requirement for prior work experience in related fields, which averages 4 years. Figure 3.4 plots the distribution of the number of years of minimum education required and the number of years of minimum experience required (using job postings that specify a given requirement). Hershbein and Kahn (2018) show that average educational requirements in Burning Glass align well with the education levels of employed workers at the occupation and MSA levels.

3.2.3.2 *Skill Clusters*

Burning Glass groups all skills into one of 28 skill clusters. Skill clusters are groupings of skills that have similar functionality, can be trained together, and/or frequently appear together in job postings. For example, the skill “Python” belongs to the “Information Technology” skill cluster, and the skill “Machine Learning” belongs to the “Analysis” skill cluster. Table 3.1 presents the top five skills (i.e., skills appearing in the largest number of job postings) for each skill cluster.

For each job posting, we calculate the share of required skills that fall within each skill cluster. For example, if a job posting requires “Python” and “Machine Learning,” then the share of the “Information Technology” skill cluster and the share of the “Analysis” skill cluster are both 50 percent. We then average these shares across all job postings of a given firm in a given year. This results in a weighted share of job postings that require skills in each skill cluster.¹² The shares of all 28 skill clusters add up to one.

3.2.4 *Descriptive Statistics*

We present summary statistics for each of our measures of worker composition. We start by showing the evolution of education levels and specialization of workers over time. Figure 3.5 plots the share of workers in four education levels (undergraduate, master’s, doctoral, and less than college) based on the Cognism resume data over time. Figure 3.6 plots the share of workers in five major fields (STEM, social science, fine arts, humanities, and medicine) based on the Cognism resume data over time. For comparison, we also plot the distribution of education levels and majors of US workers in the Census American Community Survey (ACS), which is a 1 percent random sample of the US population, in each year between 2010

12. This is equivalent to weighting each skill required by a job posting by the inverse of the total number of skills required by the job posting. We do not directly compute the share of job postings requiring skills in a skill cluster, because generic skills like “communication” are required by most job postings, although they constitute a small part of the job requirements for each job posting.

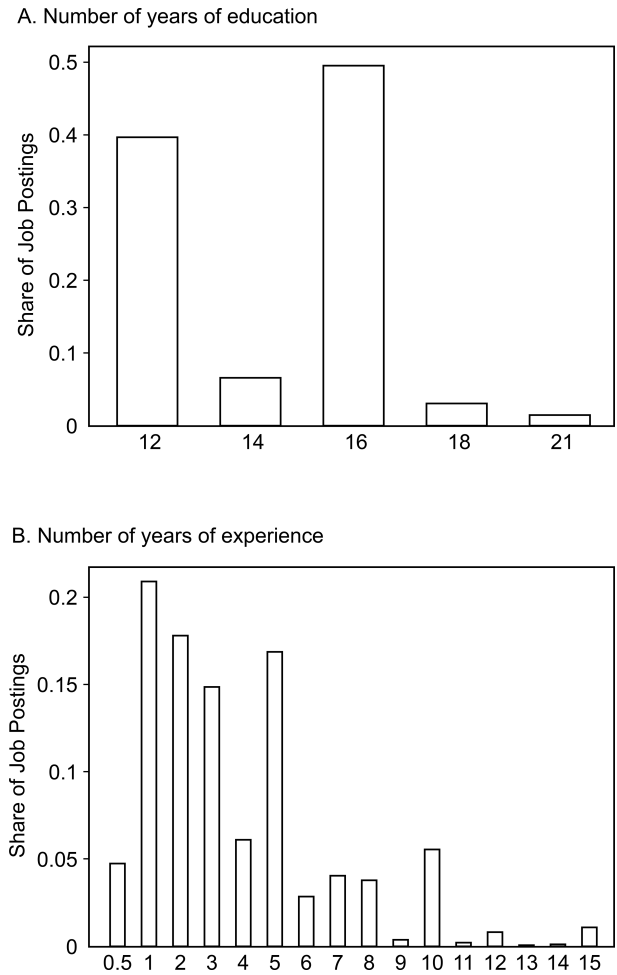


Fig. 3.4 Distribution of required years of education and experience in the Burning Glass job postings data

This figure shows the fraction of job postings with the number of years of required education (Panel A) or the number of years of required work experience (Panel B) in the Burning Glass job postings data. The sample includes all job postings of Compustat firms between 2010 and 2018.

and 2018 in appendix figures A.3 and A.4.¹³ In both Cognism and ACS, the share of workers in each education level is mostly flat over time, with an intuitive slight upward trend in the share of workers with bachelor’s or master’s degrees and slight downward trend in the share of workers with no post-secondary education. In both data sets, there is a small increase in

13. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

Table 3.1 Top skills in each skill cluster in the Burning Glass job posting data This table reports the top five skills required by the largest number of job postings in each skill cluster in the Burning Glass job posting data.						
Administration	Analysis	Business	Customer Service	Engineering	Finance	Health Care
Scheduling	Data Analysis	Project Management	Customer Service	Mechanical Engineering	Budgeting	Patient Care
Administrative Support	Data Collection	Staff Management	Customer Contact	AutoCAD	Accounting	Cardiopulmonary Resuscitation (CPR)
Data Entry	Business Intelligence	Quality Assurance and Control	Basic Mathematics	Computer Engineering	Customer Billing	Lifting Ability
Appointment Setting	SAS	Supervisory Skills	Cash Handling	Simulation	Financial Analysis	Treatment Planning
Record Keeping	Statistics	Business Process	Customer Checkout	Civil Engineering	Financial Reporting	Advanced Cardiac Life Support
Human Resources	Information Technology	Legal	Marketing	Sales	Science	Supply Chain
Occupational Health and Safety	Microsoft Office	Legal Compliance	Marketing	Sales	Chemistry	Store Management
Onboarding	Microsoft PowerPoint	Litigation	Social Media	Product Sales	Biology	Purchasing
Recruiting	SQL	Government Regulations	Packaging	Merchandising	Physics	Forklift Operation
Employee Training	Java	Legal Documentation	Client Base Retention	Business Development	Experiments	Procurement
Personal Protective Equipment	Software Development	Criminal Justice	Facebook	Sales Goals	Laboratory Testing	Inventory Management
(continued)						

Table 3.1

(cont.)

Administration	Analysis	Business	Customer Service	Engineering	Finance	Health Care
Agriculture	Construction	Design	Economics	Education	Utilities	Environment
Snow Removal	Estimating	Adobe Photoshop	Economics	Teaching	Natural Gas	HAZMAT
Lawn Care	Carpentry	Microsoft Visio	Public administration	Training Programs	Energy Management	Hazardous Waste
Fertilizers	Cost Estimation	Graphic Design	Economic Development	Training Materials	Power Distribution	Environmental Science
Agronomy	Construction Management	Adobe InDesign	Social Studies	Technical Training	Power Generation	Water Treatment
Agribusiness	Interior Design	Adobe Acrobat	Policy Analysis	Special Education	Energy Efficiency	Natural Resources
Industry Knowledge	Maintenance	Manufacturing	Media	Personal Care	Public Safety	Religion
Retail Industry Knowledge	Hand Tools	Product Development	Journalism	Cooking	Asset Protection	Youth Ministry
Information Technology Industry Knowledge	Plumbing	Machinery	Music	Food Safety	Surveillance	Student Ministry
Biotechnology	Predictive/ Preventative Maintenance	Welding	Preparing Proposals	Child Care	Loss Control/ Prevention	Children's Ministry
Industrial Engineering Industry Expertise	HVAC	Six Sigma	Proposal Writing	Food Preparation	Handling of Crisis	Family Ministry
Asset Management Industry Knowledge	Schematic Diagrams	Manufacturing Processes	Content Management	Food Service Experience	Emergency Services	Religious Education

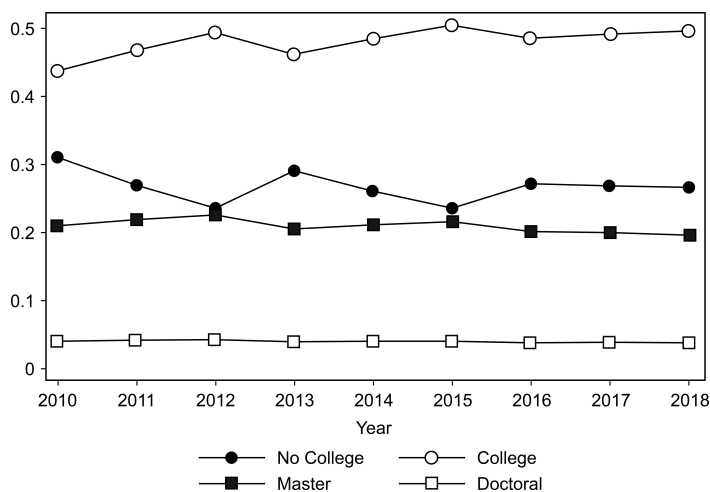


Fig. 3.5 Time series of workers' education levels in the Cognism resume data

This figure shows the time series of workers' education levels. Each line is the fraction of all employees (across all public firms) with each highest education level (no college, college, master's, or doctoral) in the Cognism resume data in a given year from 2010 to 2018.

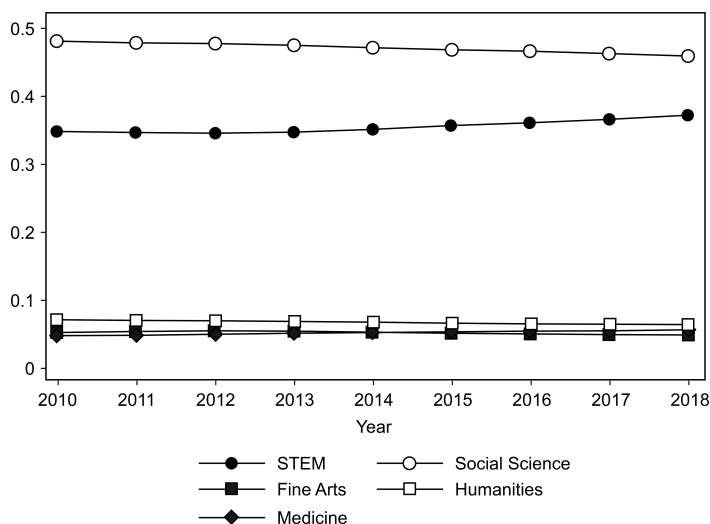


Fig. 3.6 Time series of workers' majors in the Cognism resume data

This figure shows the distribution of workers' majors over time. Each line is the fraction of all employees (across all public firms) in each major (STEM, social science, fine arts, humanities, and medicine) in the Cognism resume data in a given year from 2010 to 2018. STEM includes engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., math, physics, chemistry, computer science, statistics), and biological sciences (e.g., biology, pharmacology). The majors are mutually exclusive; for each worker, we record the major of the most recent degree earned.

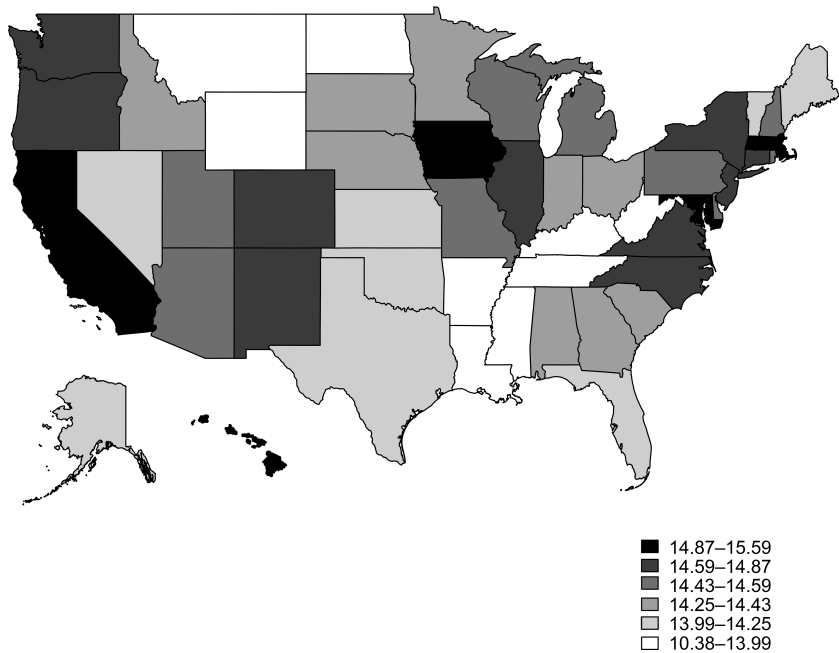


Fig. 3.7 Map of average required years of education in the Burning Glass job postings data

This figure shows a heat map of the average required years of education across U.S. states. It plots the average required years of education of job postings of public firms in each state from 2010 to 2018.

the share of workers with STEM majors and a small decrease in the share of workers with social science majors. While Cognism offers slightly more comprehensive coverage of more educated workers and workers in STEM fields, our benchmarking exercises suggest stable representativeness of the Cognism resume data across education categories and majors. Importantly, there is no differential over-representation of highly educated workers in some periods versus others.

Next, we document the variation of workers' education levels across geographic areas. Figure 3.7 considers the job postings data and shows the average required number of years of education in each state. Intuitively, states such as Massachusetts and California have the highest demand for educated workers.

We then look at the distribution of workers' education levels and specialization across industries. Figure 3.8 plots the average share of workers with undergraduate, master's, and doctoral degrees in the Cognism resume data for public firms in each of the 2-digit NAICS sectors. Firms in the "Education Services" sector and the "Professional and Business Services"

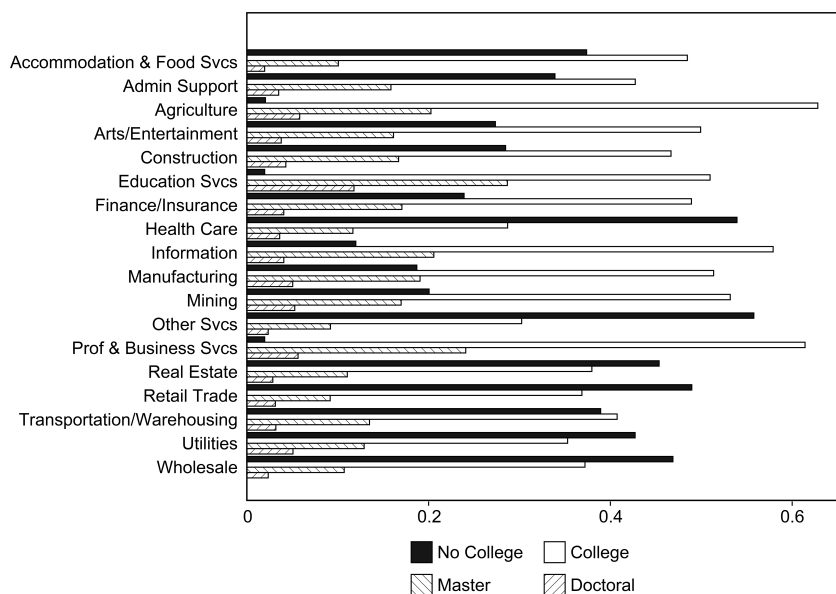


Fig. 3.8 Education level by industry sector in the Cognism resume data

This figure presents the share of workers in each highest level of education at the industry level, based on the sample of public firms. For each sector (based on NAICS 2-digit industry codes), we compute the share of workers with the highest level of education being less than college, college, master's, or doctoral in the Cognism resume data between 2010 and 2018.

sector have the highest shares of workers with undergraduate degrees, master's degrees, and doctoral degrees. Figure 3.9 considers the distribution of workers' educational majors across industries. We see intuitive trends that help validate Cognism's classification of educational majors: the tech sectors ("Professional and Business Services" and "Information") and the "Manufacturing" sector have high shares of workers with STEM majors, "Finance/Insurance" and "Real Estate" have the highest shares of workers with social science majors (which include all business school degrees such as MBAs), and the "Health Care" sector has the highest share of workers with degrees in medicine fields.

Finally, figure 3.10 considers the distribution of workers' seniority levels across industries. We observe that all industries have a pyramid structure, with the majority of workers in low-seniority levels and a small percentage of workers in high-seniority levels. The only exceptions are "Arts/Entertainment" and "Health Care," which are top-heavy with more workers in senior positions than in mid-level positions. Given the relative homogeneity of hierarchical structures across diverse industry sectors, even small changes in the proportion of employees in different levels is a meaningful indicator of shifts in a firm's organizational structure.




Fig. 3.9 Distribution of majors by industry sector in the Cognism resume data

This figure presents the share of workers in each major, based on the sample of public firms. For each sector (based on NAICS 2-digit industry codes), we compute the share of workers in each major (STEM, social science, fine arts, humanities, and medicine) in the Cognism resume data between 2010 and 2018. STEM includes engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., math, physics, chemistry, computer science, statistics), and biological sciences (e.g., biology, pharmacology). The majors are mutually exclusive; for each worker, we record the major of the most recent degree earned.

3.3 Does Ex Ante Labor Composition Predict Growth in AI Investments?

We consider the determinants of firms' investments in AI technologies and whether firms' initial labor composition can predict future AI investments. Theoretically, firms' initial labor composition could affect both their demand for AI investments and their ability to invest in AI by attracting AI talent. For example, Bresnahan (2019) and Agrawal et al. (2024) argue that the degree of modularity in the organizational structure of a firm could impact the firm's ease of AI adoption. When modularity is high, tasks are more independent, and there is less need for coordination; as a result, it is easier to implement AI and change decision-making in one part of the organization, as it does not require changes elsewhere. In terms of employee specialization, Acemoglu et al. (2022a) show that establishments with occu-

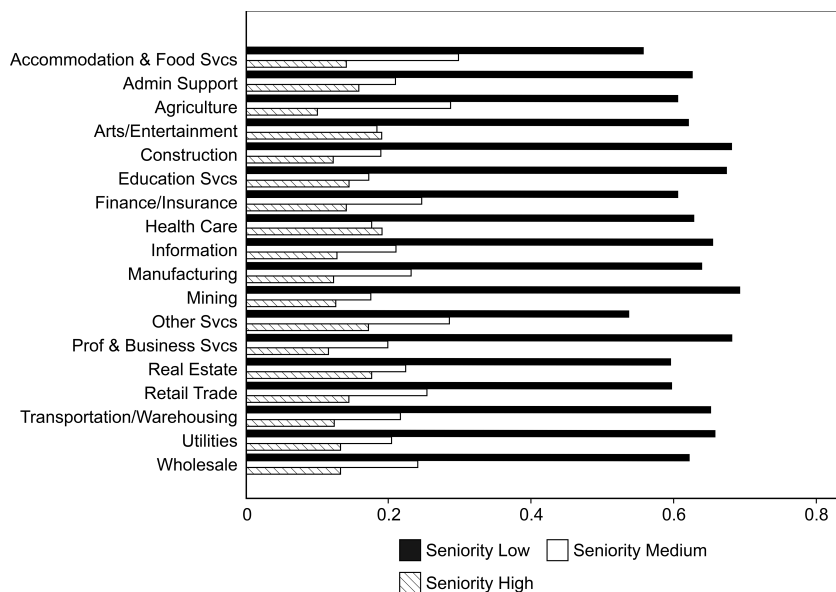


Fig. 3.10 Seniority level by industry sector in the Cognism resume data

This figure presents the share of workers in each seniority level at the industry level, based on the sample of public firms. For each sector (based on NAICS 2-digit industry codes), we compute the share of workers in each seniority level (low, medium, or high) in the Cognism resume data between 2010 and 2018.

pations that are more exposed to AI technologies have a higher demand for AI workers. And in terms of firms' workforce education, Babina et al. (2024) document that firms with alumni connections to universities that are historically strong in AI research invest more in AI by being able to attract AI-trained students from those universities.

We are interested in understanding the *use* of AI technologies by a wide range of firms. In order to not conflate this with the invention of new AI tools, we exclude firms in the tech sector (2-digit NAICS 51 or 54) from our empirical analyses. In table 3.2, we examine how ex ante worker composition predicts future growth in firm-level AI investments by estimating the following cross-sectional specification:

$$(1) \quad \Delta \text{ShareAIWorkers}_{i,[2010,2018]} = \beta \text{WorkerComposition}_{i,2010} + \text{IndustryFE} + \epsilon_i,$$

where $\Delta \text{ShareAIWorkers}_{i,[2010,2018]}$ denotes the change in the share of firm i 's AI-skilled employees from 2010 to 2018, standardized to mean zero and standard deviation one to streamline the economic interpretation. All regressions include 2-digit NAICS industry fixed effects. The explanatory variables

include the following measures of ex ante worker composition measured as of 2010: the shares of workers in each seniority level in column 1, the shares of workers in each education level in column 2, and the shares of workers in each major in column 3. To avoid multi-collinearity, we omit the share of workers in the high seniority level in column 1, the share of workers with no college degree in column 2, and the share of workers with medicine majors in column 3. Column 4 considers seniority, educational attainment, and college majors simultaneously. All continuous variables are winsorized at 1 percent and 99 percent to limit the influence of outliers. We weight the estimating equation by each firm's total number of employees in the Cognism resume data in 2010 to account for potential differences in precision in the measurement of AI investments across firms with different coverage.¹⁴

The results in table 3.2 highlight that firms with more workers with doctoral degrees and more workers with STEM majors invest more in AI going forward. This is broadly consistent with the evidence in Babina et al. (2024) that firms with more educated workforces and alumni connections to AI-strong universities are able to attract AI talent more easily. The hierarchical structure of the firm does not significantly predict AI investments.

3.4 Firms' AI Investments and Changes in Labor Composition

We explore how the key aspects of firms' labor composition change with firms' investments in AI. Firms that invest more in AI shift toward more educated workforces, with more emphasis on STEM degrees and skills in analysis and IT. At the same time, AI-investing firms become less top-heavy in terms of their hierarchies, with increasing shares of junior employees and less emphasis on middle-management and senior roles.

3.4.1 AI Investments and Employee Seniority

We begin the analysis by examining whether firms that invest in AI become more top-heavy or bottom-heavy in terms of their hierarchical structure. The direction of this shift is ex ante ambiguous and an open empirical question. On the one hand, AI contributes to firm growth (Babina et al. 2024), and as enterprises have grown in size over the 20th century, the share of employees in managerial positions has risen dramatically (Radner 1992). Thus, AI-fueled growth may result in the continuation of this trend toward increased organizational complexity and increased need for middle- and top-level managerial positions. For example, Caliendo et al. (2015) show that many firms expand by adding layers of management. On the other hand, Garicano and Rossi-Hansberg (2006) present a theoretical model where reductions in costs of accessing knowledge through improved data

14. Since the numbers of worker resumes are correlated with firm size, this weighting scheme also roughly weights firms in accordance with their contribution to the economy. Our results are also robust to weighting by Cognism's coverage rate of Compustat employment in 2010.

Table 3.2 Initial worker composition and AI investments

	Δ Share of AI Workers 2010–2018			
	(1)	(2)	(3)	(4)
Share of Workers with Low Seniority, 2010	–0.211 (0.393)			–0.394 (0.414)
Share of Workers with Medium Seniority, 2010	–0.443 (0.407)			–0.152 (0.425)
Share of Workers with College Degree, 2010		–0.489** (0.232)		–0.457* (0.245)
Share of Workers with Master’s Degree, 2010		0.891** (0.402)		0.590* (0.310)
Share of Workers with Doctoral Degree, 2010		1.766** (0.819)		2.933*** (0.906)
Share of Workers with STEM Major, 2010			0.741** (0.326)	1.238*** (0.447)
Share of Workers with Social Science Major, 2010			–0.219 (0.286)	0.563** (0.232)
Share of Workers with Fine Arts Major, 2010			0.185 (0.661)	1.021 (0.734)
Share of Workers with Humanities Major, 2010			1.321* (0.694)	0.979 (0.675)
Industry Sector FE	Y	Y	Y	Y
Adj R-Squared	0.096	0.134	0.125	0.158
Observations	1,218	1,218	1,216	1,216

Note: This table reports the coefficients from regressions of cross-sectional changes in AI investments by US public firms (in non-tech sectors) from 2010 to 2018 on the following ex ante firm characteristics measured in 2010: share of workers in each seniority level in column 1, share of workers in each education level in column 2, and share of workers in each major (based on the highest degree earned) in column 3. Column 4 includes all firm characteristics from columns 1 to 3. The dependent variable is the growth in the share of AI workers from 2010 to 2018 using the resume data from Cognism. Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. The dependent variable is normalized to have a mean of zero and a standard deviation of one. Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

processing (which is arguably the main effect of AI technology) result in increases in the problem-solving ability of employees at all levels, leading to increased span of control and less reliance on top-heavy hierarchical levels.

Previous technologies also had differing impacts on firm organization. Acemoglu et al. (2007) show that firms investing in new information technologies are more likely to favor decentralization. Bloom et al. (2014) find that information technology and communication technology have opposing effects on firm organization: information technology is a decentralizing force allowing workers and lower-level managers to handle more problems, while communication technology decreases autonomy and is associated with more centralization.

Empirically, we measure hierarchical flatness as the share of a firm’s over-

all employees who are in more junior versus more senior positions: if firms become more top-heavy, increasing their middle-management and senior roles, then the share of employees in senior positions will rise, and vice versa. We link the changes in the shares of employees across levels to firms' AI investments using long-differences regressions, which are standard in settings analyzing slow-moving processes like technological progress (Acemoglu and Restrepo 2020) and especially well suited to study AI investments, which are gradual over time and have non-immediate effects (Babina et al. 2024). Specifically, we regress firm-level changes in the share of junior, middle-ranked, and senior employees from 2010 to 2018 on changes in AI investments proxied by the growth in the share of AI workers. By taking first differences in independent and dependent variables, the long-differences specification ensures that time-invariant firm characteristics do not drive the results. In table 3.3, we report the estimates from the following regression:

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

where the main independent variable, $\Delta \text{ShareAIWorkers}_{i,[2010,2018]}$, captures the change in the share of AI workers in firm i from 2010 to 2018, standardized to mean zero and standard deviation of one as in table 3.2. *IndustryFE* are 2-digit NAICS fixed effects. As in section 3.2.4, we focus on firms in non-tech sectors, and weigh the estimating equation by each firm's total number of employees in the Cognism resume data in 2010. In columns 1, 3, and 5 we include only industry fixed effects to examine the unconditional relationship between changes in AI investments and employee seniority. In columns 2, 4, and 6, we add a rich set of controls proposed by Babina et al. (2024) and measured at the start of the sample period in 2010: (i) firm-level characteristics (log sales, cash/assets, R&D/Sales, log markup, and the log of the firm's total number of jobs); (ii) characteristics of the commuting zones (CZ) where each firm is located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers); and (iii) the log industry-average wage.¹⁵ Summary statistics on key variables for the regression sample are provided in appendix table A.3.¹⁶

In columns 1 and 2 of table 3.3, the dependent variable is the firm-level change in the share of junior employees (i.e., employees in entry-level and single-contributor positions) from 2010 to 2018. In columns 3 and 4, the

15. When firms span multiple commuting zones, we calculate commuting-zone-level variables as the weighted average, using numbers of Burning Glass job postings in each commuting zone as weights.

16. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

Table 3.3 AI investments and workers' seniority levels

	Δ Share Seniority Low		Δ Share Seniority Middle		Δ Share Seniority High	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Share AI Workers	0.015*** (0.004)	0.016*** (0.004)	-0.007*** (0.002)	-0.008*** (0.002)	-0.007** (0.003)	-0.008** (0.003)
Industry Sector FE	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y
Adj R-Squared	0.175	0.334	0.170	0.232	0.170	0.314
Observations	1,218	1,218	1,218	1,218	1,218	1,218

Note: This table reports the coefficients from long-differences regressions of the change in the share of workers in each seniority level from 2010 to 2018 on the contemporaneous firm-level changes in AI investments among US public firms (in non-tech sectors). The independent variable is the growth in the share of AI workers from 2010 to 2018 based on the Cognism resume data, standardized to mean zero and standard deviation of one. The dependent variables are the changes in the share of workers in each seniority level (low in columns 1 and 2; medium in columns 3 and 4; and high in columns 5 and 6) in the Cognism resume data. Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. Columns 2, 4, and 6 also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, log markup, and log number of jobs), log industry wage, and characteristics of the commuting zones where the firms are located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

dependent variable is the firm-level change in mid-level employees (i.e., team leads and middle managers), and columns 5 and 6 consider the firm-level change in senior employees (department heads and top-level leadership). The results reveal that AI-investing firms become flatter (or, more precisely, more bottom-heavy and less top-heavy). A one-standard-deviation increase in AI investments is associated with a 1.6 percent increase in the share of junior employees, accompanied by a 0.8 percent decline in both mid-level employees and senior management. Importantly, the results are almost identical with and without the inclusion of detailed ex ante firm-level, location-level, and industry-level controls in even columns, despite the adjusted *R-Squared* rising significantly (nearly doubling for both the share of junior employees and the share of senior employees). This makes it unlikely that the results are driven by ex ante omitted firm characteristics (Altonji et al. 2005; Oster 2019). In appendix table A.4, we alternatively weight the regressions by the coverage rate in 2010 and find similar results.¹⁷

One concern is that firms' hierarchical structure may change as a result of firm growth and not AI investments per se. For example, if the adjustment costs of hiring junior workers are lower, fast-growing firms will likely expand

17. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

by hiring junior employees. In appendix table A.5, we directly control for firms' sales growth from 2010 to 2018 in the regressions.¹⁸ While faster firm growth is indeed correlated with an increase in the share of junior workers and a decrease in the share of senior workers, the effects of AI investments remain unchanged conditional on firm growth. This suggests that the effects are not mechanically driven by firm growth.

The magnitude of the results in table 3.3 is economically meaningful, given that the cross-sectional average change in the share of junior employees over the sample period is only 0.18 percent (see table A.3), and there has been no overall trend toward more junior employees across the cross-section of US public firms. AI-investing firms experience fast shifts toward more junior employees and reductions in the share of senior-level employees. This is consistent with the theoretical prediction of Garicano and Rossi-Hansberg (2006) and empirical evidence in Bloom et al. (2014) that technologies that improve prediction and decision-making, such as AI, will give lower-level workers more autonomy and require fewer managerial layers in firms.

3.4.2 AI Investments and Employee Educational Attainment

We leverage the detailed individual-level information in the resumes to study the association between changes in firm-level AI investments and the upskilling of the firms' workforces in terms of the employees' educational attainment. Educational attainment is a particularly relevant trend to investigate in the context of firms' AI investments, given the extensive labor economics literature on skill-biased technological change (Autor et al. 1998; Acemoglu and Autor 2011; Autor et al. 2003; Katz and Murphy 1992). On the one hand, previous technologies such as IT have increased the relative demand for college graduates. In the case of AI, Babina et al. (2024) show that AI-investing firms engage in more product innovation, which may further increase firms' demand for skilled labor (Bresnahan et al. 2002). On the other hand, Webb (2020) predicts that AI is more likely to replace high-skilled tasks performed by highly educated workers than previous technologies such as software and robots. Grennan and Michaely (2022) study the impact of AI on a particular group of high-skilled workers—financial analysts—and find that AI replaces some technical tasks but increases the importance of soft skills. Furthermore, changes to organizational hierarchies may induce shifts in labor composition, as flatter hierarchical structures could require higher human capital from each individual employee (Caroli and Van Reenen 2001).

In table 3.4, we investigate the extent to which AI, as a technology, is associated with labor shifts toward more educated workers. To do this, we estimate the regression in equation 2 using the same independent variable and controls as in table 3.3, but looking at the changes in the share of

18. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

Table 3.4 AI investments and workers' education levels

	Δ Share College		Δ Share Master's		Δ Share Doctoral		Δ Share No College	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Share AI Workers	0.033*** (0.012)	0.037*** (0.010)	0.027*** (0.004)	0.029*** (0.004)	0.007*** (0.001)	0.006*** (0.001)	-0.068*** (0.014)	-0.073*** (0.014)
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.099	0.199	0.107	0.216	0.155	0.211	0.115	0.215
Observations	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218

Note: This table reports the coefficients from long-differences regressions of the change from 2010 to 2018 in the share of workers in each highest education level on the contemporaneous firm-level changes in AI investments among US public firms (in non-tech sectors). The independent variable is the growth in the share of AI workers from 2010 to 2018 based on the Cognism resume data, standardized to mean zero and standard deviation of one. The dependent variables are measured using the Cognism resume data and represent the changes in the share of workers whose maximal attainment is a college degree in columns 1 and 2, the share of employees whose maximal attainment is a master's degree in columns 3 and 4, the share of employees with doctoral degrees in columns 5 and 6, and the share of employees with no college degree in columns 7 and 8. Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. Columns 2, 4, 6, and 8 also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, log markup, and log number of jobs), log industry wage, and characteristics of the commuting zones where the firms are located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

workers in each education level as the outcome variable. To motivate our regression specification for examining changes in firms' low versus high-skilled workers, suppose that firms' production function takes the following form: $Q = AF(L^H, L^L, AI, X)$, where Q is output, A is productivity, L^H and L^L are high-skilled and low-skilled labor, respectively, AI is AI capital, and X is other inputs, including physical capital. Assume that labor input L^j ($j = H, L$) is supplied flexibly at wage W^j , and other inputs are quasi-fixed. Following Caroli and Van Reenen (2001), we consider a translog short-run cost function, and Shepard's Lemma implies that the share of high-skilled labor in total labor costs is: $SHARE^H = \alpha \ln(W^H/W^L) + \beta AI + \gamma \ln Q$. Assuming that the production function is homothetic (which can be relaxed), after taking the differences, we get: $\Delta SHARE_{it}^H = \beta \Delta AI_{it} + X_i + \Delta u_{it}$.¹⁹ Therefore, our long-differences specification approximately corresponds to the factor demand function, with a positive (or negative) coefficient β in a regression of high-skilled workers on AI investment indicating that AI and high-skilled labor are complements (or substitutes).

The results in table 3.4 show that larger AI investments are associated with educational upskilling of the workforce. Based on estimates from even columns when all controls are included, a one-standard-deviation increase in AI investments is associated with a 3.7 percent increase in the share of college-educated workers, a 2.9 percent increase in employees with master's degrees, and a 0.6 percent increase in employees with doctoral degrees. Correspondingly, the share of employees with no college education declines by a substantial 7.3 percent. Therefore, our results suggest that AI investments and high-skilled labor are complements. As with the results on seniority, the results on educational attainment are nearly identical with (even columns) and without (odd columns) the inclusion of detailed firm-level, location-level, and industry-level controls, indicating that these findings are likely not driven by omitted firm, industry, or geographic characteristics. The results are also interesting in light of relatively slow shifts in the educational makeup of the workforce in general, as shown earlier in figure 3.5. The share of workers with advanced degrees (master's and doctoral) has remained practically flat from 2010 to 2018 in the overall workforce. By contrast, the share of employees with advanced degrees (both master's and doctoral degrees) has risen significantly in firms that have been investing in AI, suggesting that there is a reallocation of highly educated workers away from non-AI investing firms and toward firms that invest more heavily in AI.

The reason for the shift toward more educated workforces in AI-investing firms appears to be, at least in part, increasing *demand* for educated and experienced employees on the firms' side. In table 3.5, we complement the results from the Cognism resume data with an analysis of labor-related outcomes measured using job postings data from Burning Glass. We estimate

19. We can write the relative wages $\ln W^H/W^L$ as composed of industry-year dummies, firm fixed effects, and idiosyncratic shocks.

Table 3.5 AI investments and required education and experience in the job postings data

	Δ Years of Education		Δ Years of Experience	
	(1)	(2)	(3)	(4)
Δ Share AI Workers	0.476** (0.204)	0.516** (0.217)	0.137* (0.072)	0.061 (0.078)
Industry Sector FE	Y	Y	Y	Y
Controls	N	Y	N	Y
Adj R-Squared	0.397	0.458	0.161	0.235
Observations	1,060	1,060	1,059	1,059

Note: This table reports the coefficients from long-differences regressions of the change from 2010 to 2018 in the average required education and experience on the contemporaneous firm-level changes in AI investments among US public firms (in non-tech sectors). The independent variable is the growth in the share of AI workers from 2010 to 2018 based on the Cognism resume data, standardized to mean zero and standard deviation of one. The dependent variables are the average required years of education in the Burning Glass job postings data in columns 1 and 2, and average required years of experience in the Burning Glass job postings data in columns 3 and 4. Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. Columns 2 and 4 also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, log markup, and log number of jobs), log industry wage, and characteristics of the commuting zones where the firms are located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

equation 2 using the same independent variable and controls as in tables 3.3 and 3.4, but with the dependent variables being: (i) the change in the average number of years of education required in the firm's job postings from 2010 to 2018 (columns 1 and 2), and (ii) the change in the average number of years of experience required in the firm's job postings from 2010 to 2018 (columns 3 and 4).²⁰ We observe that firms that invest more in AI look for more educated and more experienced workforces. For example, a one-standard-deviation increase in the share of AI workers from 2010 to 2018 is associated with additional 0.52 years of educational experience (column 2), reinforcing the increased educational attainment of actual workers that we observe in table 3.4.

We perform a series of robustness tests in the appendix. First, in appendix table A.6, we exclude AI workers when calculating the share of workers in each education level (columns 1–6).²¹ This has little impact on the results, suggesting that the educational upskilling is not mechanically due to the

20. The sample size is smaller than in table 3.4, because not all firms in the Cognism resume data are matched to job postings in Burning Glass data.

21. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

hiring of AI workers. Second, appendix table A.7 shows that the results are robust when we alternatively weight the regressions by our Cognism data coverage rate in 2010. Finally, to account for potential non-homotheticity of the production function (i.e., the ratio of high-skilled to low-skilled labor changes as firm expands), we control for firm sales growth from 2010 to 2018 in appendix table A.8 and find similar results.

3.4.3 AI Investments and Employee Specialization and Skills

We consider one additional aspect of the changing workforce of firms and its relationship to AI investments: the importance of technical and non-technical skills. A number of recent studies point out that technical employee skills are uniquely important to modern firms. For example, Agrawal et al. (2021b) document that engineers and scientists are among the employees whose net flows (arrivals and departures) are most predictive of the firm's stock returns. Fedyk and Hodson (2023) show that technologies such as IT in the early 2000s and data analysis in 2010s can even be overvalued by corporate investors.

We use the resume data to observe whether AI investments are associated with broader changes in the technical specialization of AI-investing firms. Specifically, in table 3.6, we re-estimate equation 2 using the same independent variable and controls as in tables 3.3 and 3.4, but with dependent variables being: (i) the change in the share of employees whose most recent degree was in a STEM field in columns 1 and 2; (ii) the change in the share of employees whose most recent degree was in social science in columns 3 and 4; (iii) the change in the share of employees whose most recent degree was in fine arts in columns 5 and 6; (iv) the change in the share of employees whose most recent degree was in humanities in columns 7 and 8; and (v) the change in the share of employees whose most recent degree was in medicine in columns 9 and 10.

The results reveal that increased AI investments are associated with a general trend toward more technically skilled employees at the firm level. When all controls are included, a one-standard-deviation increase in the share of AI workers at the firm is associated with a 1.9 percent increase in the share of employees whose most recent degree is in STEM.²² This increase is offset by declines in the shares of employees with backgrounds in social science (a decline of 1.1 percent), and medicine (a smaller decline of 0.4 percent).

Once again, we supplement our resume-based results with firms' demand from Burning Glass to see whether firms that invest more in AI start increasing their demand for technical skills more generally. AI has been highlighted as a technology that can shift the skill requirements of the workforce by Acemoglu et al. (2022a), who also consider job postings and find that estab-

22. This increase is not driven by the hiring of AI workers, who account for a small fraction of firms' employees. In appendix table A.6 columns 7 and 8 (<http://www.nber.org/data-appendix/c14753/appendix.pdf>), we exclude AI workers and find a similar increase in the share of workers in STEM fields.

Table 3.6 AI investments and workers' majors

	Δ Share STEM		Δ Share Social Science		Δ Share Fine Arts		Δ Share Humanities		Δ Share Medicine	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Share AI Workers	0.019*** (0.007)	0.019*** (0.006)	-0.011*** (0.003)	-0.011*** (0.003)	-0.004 (0.005)	-0.003 (0.003)	-0.000 (0.002)	-0.001 (0.001)	-0.003*** (0.001)	-0.004*** (0.001)
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.122	0.216	0.092	0.208	0.127	0.206	0.218	0.258	0.110	0.148
Observations	1,216	1,216	1,216	1,216	1,216	1,216	1,216	1,216	1,216	1,216

Note: This table reports the coefficients from long-differences regressions of the change in the share of workers in each major from 2010 to 2018 on the contemporaneous firm-level changes in AI investments among US public firms (in non-tech sectors). The dependent variables are the share of employees whose most recent degree was in a STEM field in columns 1 and 2, the share of employees whose most recent degree was in social science in columns 3 and 4, the share of employees whose most recent degree was in fine arts in columns 5 and 6, the share of employees whose most recent degree was in humanities in columns 7 and 8, and the share of employees whose most recent degree was in medicine in columns 9 and 10. The majors are mutually exclusive; for each worker, we record the major of the most recent degree earned. STEM includes engineering (e.g., electrical, chemical, mechanical), physical sciences (e.g., math, physics, chemistry, computer science, statistics), and biological sciences (e.g., biology, pharmacology). The independent variable is the growth in the share of AI workers from 2010 to 2018 based on the Cognition resume data, standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Cognition resumes in 2010. All specifications control for industry sector fixed effects. Columns 2, 4, 6, 8, and 10 also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, log markup, and log number of jobs), log industry wage, and characteristics of the commuting zones where the firms are located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

lishments with more occupations that are highly exposed to AI are associated with both increased redundancies in existing skills and more requirements of new skills. We focus on skill clusters, and for each skill cluster in Burning Glass, we estimate equation 2 with the same independent variable and controls as in tables 3.3–3.6, but with the dependent variable being the share of job postings within each specific skill cluster. The results, reported in table 3.7, show that the main skill shifts associated with firms' AI investments are (i) increased demand for data analysis skills, (ii) increased demand for IT skills, and (iii) lower demand for maintenance skills. For example, a one-standard-deviation increase in the share of AI workers at a given firm from 2010 to 2018 corresponds to that firm increasing the share of its job postings requiring IT skills by 1.2 percent. These results suggest that AI skills are complementary to IT skills and can substitute for maintenance skills. Interestingly, firms that invest more heavily in AI do not reduce their demand for some of the skill groups that are most often predicted to be replaced by AI, such as customer service, HR, and legal skills.²³

Overall, we find that firms' AI investments are associated with changing workforces and increased importance of technical skills. Babina et al. (2024) find that AI-fueled innovation is the main channel through which AI investments seem to power firm growth to date. Our findings are consistent with AI-fueled product innovation increasing firms' demand for complementary skilled labor in STEM and IT jobs, which are necessary to structure, store, and process data—a crucial input for AI applications.

3.4.4 Instrumental Variable Estimates

In this section, we use an instrumental variable strategy based on firms' ex ante exposure to universities' supply of AI graduates. The instrument was first used in Babina et al. (2024). The instrument isolates variation in firms' AI investments driven by the supply of AI-skilled labor, which is a key input to AI. This mitigates concerns regarding reverse causality and potential bias from unobserved demand shocks driving both firms' AI investments and changes in firm organization or workforce composition. The estimates are informative about the effects on firm organization or workforce composition if firms invest more in AI due to increased access to AI-skilled labor. In particular, we instrument firm AI investments using the variation in firms' ex ante exposure to the supply of AI talent from universities that are historically strong in AI research. As argued in Babina et al. (2024), academic research in AI has been ongoing for much longer than the commercial interest in AI, and universities that are historically strong in AI research have been able to train more AI-skilled graduates in recent years. As a result, firms' preexisting

23. This may be due to challenges in adopting AI for these tasks. For example, Tambe et al. (2019) identify complexity of HR phenomena, small data sets, and fairness and legal constraints as the main challenges in using AI techniques for HR tasks.

Table 3.7 AI investments and required skills in the job postings data

	Δ Share of Jobs w/ Administration Skill		Δ Share of Jobs w/ Analysis Skill		Δ Share of Jobs w/ Business Skill		Δ Share of Jobs w/ Customer Service Skill		Δ Share of Jobs w/ Engineering Skill		Δ Share of Jobs w/ Finance Skill		Δ Share of Jobs w/ Healthcare Skill	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Δ Share AI Workers	-0.001 (0.001)	-0.000 (0.002)	0.005*** (0.001)	0.004*** (0.001)	0.001 (0.003)	-0.003 (0.004)	-0.006 (0.005)	0.004 (0.006)	-0.000 (0.002)	-0.002 (0.002)	-0.007 (0.005)	-0.011** (0.005)	-0.003 (0.003)	-0.002 (0.004)
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.062	0.094	0.264	0.306	0.094	0.146	0.222	0.349	0.038	0.123	0.059	0.178	0.046	0.100
Observations	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099
	Δ Share of Jobs w/ HR Skill		Δ Share of Jobs w/ IT Skill		Δ Share of Jobs w/ Legal Skill		Δ Share of Jobs w/ Marketing Skill		Δ Share of Jobs w/ Sales Skill		Δ Share of Jobs w/ Science Skill		Δ Share of Jobs w/ Supply Chain Skill	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Δ Share AI Workers	0.001 (0.001)	-0.000 (0.002)	0.016** (0.007)	0.012* (0.007)	-0.001 (0.001)	-0.000 (0.001)	0.003* (0.002)	0.001 (0.003)	0.004 (0.005)	0.008 (0.006)	-0.000 (0.001)	-0.001 (0.001)	-0.006*** (0.002)	-0.002 (0.003)
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.057	0.105	0.227	0.271	0.034	0.119	0.016	0.066	0.332	0.386	0.107	0.141	0.069	0.145
Observations	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099

(continued)

Table 3.7 (cont.)

	Δ Share of Jobs w/ Agriculture Skill	Δ Share of Jobs w/ Construction Skill	Δ Share of Jobs w/ Design Skill	Δ Share of Jobs w/ Economics Skill	Δ Share of Jobs w/ Education Skill	Δ Share of Jobs w/ Utilities Skill	Δ Share of Jobs w/ Environment Skill							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Δ Share AI Workers	-0.000 (0.000)	-0.000* (0.000)	-0.001 (0.000)	-0.000 (0.001)	-0.001 (0.001)	-0.002 (0.001)	0.000 (0.001)	0.000 (0.000)	-0.001 (0.001)	-0.002* (0.001)	-0.001 (0.001)	-0.000 (0.001)	0.001 (0.001)	0.001 (0.001)
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Adj R-Squared	0.074	0.107	0.055	0.083	0.169	0.352	0.081	0.176	0.068	0.103	0.023	0.081	0.181	0.192
Observations	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099
Δ Share of Jobs w/ Industry Knowledge Skill														
	Δ Share of Jobs w/ Maintenance Skill			Δ Share of Jobs w/ Manufacturing Skill			Δ Share of Jobs w/ Media Skill		Δ Share of Jobs w/ Personal Care Skill			Δ Share of Jobs w/ Public Safety Skill		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)		
Δ Share AI Workers	-0.001 (0.003)	0.002 (0.003)	-0.007*** (0.002)	-0.006*** (0.002)	-0.002* (0.001)	-0.002 (0.001)	0.001 (0.001)	0.001* (0.001)	0.001 (0.001)	-0.000 (0.002)	-0.000 (0.000)	0.000 (0.001)		
Industry Sector FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y		
Controls	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y		
Adj R-Squared	0.117	0.166	0.121	0.168	-0.002	0.028	0.099	0.108	0.161	0.239	0.154	0.248		
Observations	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099	1,099		

Note: This table reports the coefficients from long-differences regressions of the change in the share of job postings requiring each skill cluster from 2010 to 2018 on the contemporaneous firm-level changes in AI investments among US public firms (in non-tech sectors). The dependent variables are the change in the average share of required skills in each skill cluster across all the job postings of the firm. The independent variable is the growth in the share of AI workers from 2010 to 2018 based on the Cognism resume data, standardized to mean zero and standard deviation of one. Regressions are weighted by the number of Cognism resumes in 2010. All specifications control for industry sector fixed effects. All even-numbered columns also include the baseline controls all measured as of 2010: firm-level characteristics (log sales, cash/assets, R&D/sales, log markup, and log number of jobs), log industry wage, and characteristics of the commuting zones where the firms are located (the share of workers in IT-related occupations, the share of college-educated workers, log average wage, the share of foreign-born workers, the share of routine workers, the share of workers in finance and manufacturing industries, and the share of female workers). Standard errors are clustered at the 5-digit NAICS industry level and reported in parentheses. *, **, and *** denote statistical significance at the 10 percent, 5 percent, and 1 percent levels, respectively.

connections to AI-strong universities offer a plausibly exogenous source of variation in firms' access to the supply of AI talent during the 2010s boom in commercial interest in AI.

The instrument for firm i is: $IV_i = \sum_u s_u^{2010} AIstrong_u$, where s_u^{2010} is the share of STEM workers in firm i in 2010 who graduated from university u , and $AIstrong_u$ equals one if university u is identified as an AI-strong university based on pre-2010 publications.²⁴ Two key concerns with our instrument are that AI-strong universities may also be strong in computer science (CS) outside of AI or generally be strong (highly-ranked) universities, which might affect firm outcomes through channels other than AI investments. To address these concerns, we control for firms' ex ante exposure to CS-strong universities $\sum_u s_u^{2010} CSstrong_u$ and top-ranked universities $\sum_u s_u^{2010} Top10_u$, where $CSstrong_u$ is the average pre-2010 share of (non-AI) CS researchers at university u , and $Top10_u$ equals one if a university is among the top 10 universities ranked by the US News & World Report.

Appendix tables A.9, A.10, and A.11 present the IV estimates.²⁵ In each table, the odd columns control for industry fixed effects and exposure to CS-strong and top-10 universities, and the even columns additionally control for (i) baseline controls (firm-, industry-, and commuting-zone-level controls), (ii) pre-period firm sales and employment growth between 2000 and 2008 to address unobservable firm characteristics that might simultaneously drive firms' growth trajectories and their hiring of AI workers, and (iii) state fixed effects to control for local labor market characteristics that might drive both firms' AI hiring and their growth. In all specifications, the first-stage F-statistics are above 10. Consistent with our main results using the long-differences specification, we find that instrumented firm AI investments are associated with the flattening of organizational hierarchy and upskilling of firms' workforce.²⁶ We also find that AI-investing firms experience an increase in the share of STEM workers and a decrease in the share of workers in social sciences, although these IV estimates are not statistically significant (appendix table A.11). These results suggest that the patterns we document are not driven by unobserved demand shocks simultaneously affecting AI investments and firm organization or workforce, but rather reflect the effects of AI investments due to supply shocks.

3.5 Conclusion

In this paper, we study the relationship between the use of AI technologies and workforce composition and organization at the firm level. We find that firms that initially have a more educated workforce and higher emphasis on

24. Babina et al. (2024) describe the data used to construct the IV and perform validation exercises in more detail.

25. See <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

26. Appendix tables A.9 and A.10, <http://www.nber.org/data-appendix/c14753/appendix.pdf>.

STEM workers are more likely to invest in AI. At the same time, firm-level growth in AI investments is associated with an increasingly flatter hierarchical structure, an increase in the share of workers with college degrees and advanced degrees, and a further increase in the share of workers with STEM majors. Data on job postings reveal that firms investing in AI technologies increase their demand for workers with more years of education and workers with data analysis and IT skills.

Our evidence of major changes in firms' workforce composition and organization accompanying AI investments contributes to our understanding of how AI can transform firms' organization and production processes. As a predictive technology, AI improves individual employees' ability to make predictions and decisions, which increases the autonomy of workers and reduces the demand for managerial positions. However, unlike previous automation technologies that displaced routine tasks, AI investments are not associated with the reduction in demand for high-skilled workers performing prediction tasks, instead increasing the share of high-skilled labor at the firm level. Further understanding the interactions between AI technology, production processes, and firm organization would be a fruitful area for future work.

Our evidence also helps to shed light on the impact of AI on the labor market. At the firm level, AI is associated with increased demand for skilled labor and does not seem to displace tasks that are commonly predicted to be replaced by AI, such as customer service, human resources, and legal jobs. However, it remains an open question how these effects aggregate to the labor market level, and it is possible that AI displaces jobs in non-AI-investing firms. Furthermore, our results imply that there is a reallocation of skilled labor from non-AI-investing firms to AI-investing firms, which might have important implications for sorting and between-firm inequality in the labor market.

Finally, while the research in this paper examines the impact of firms' AI use before the popularization of generative AI tools like DALL·E and ChatGPT, the upskilling of employees in AI-investing firms is likely to accelerate with the adoption of generative AI tools. This upskilling effect is associated with AI-powered product innovation (Babina et al. 2024), and many firms that have already invested heavily in AI technologies are starting to use tools like ChatGPT to further improve their products and services.²⁷ Alternatively, the spread of AI as a software via tools like ChatGPT could

27. For example, (<https://www.theverge.com/2023/3/9/23632312/microsoft-azure-openai-chatgpt-feature-available>) Microsoft makes ChatGPT available in its Azure OpenAI service. This allows developers and businesses to integrate OpenAI's ChatGPT model into their own services. A broad range of companies like Morningstar (<https://newsroom.morningstar.com/newsroom/news-archive/press-release-details/2023/Mo-an-AI-Chatbot-Powered-by-Morningstar-Intelligence-Engine-Debuts-in-Morningstar-Platforms/default.aspx>) and Mercedes-Benz (<https://www.cnbc.com/2023/06/15/mercedes-benz-microsoft-to-test-chat-gpt-in-vehicles.html>) have already partnered with Microsoft to incorporate ChatGPT into their products and services.

expose more workers to automation from AI, especially those who are highly skilled (Eloundou et al. 2023, Felten et al. 2023, Eisfeldt et al. 2023). Finally, early evidence suggests that generative AI can improve worker productivity across the skill distribution (Brynjolfsson et al. 2023; Casal and Kessler 2023; Dell’Acqua et al. 2023; Agarwal et al. 2023; Noy and Zhang 2023), and Agrawal et al. (2023) suggest that AI advances can potentially open up employment opportunities for workers with generic skills. More work is needed to understand the evolving impacts of firms’ AI use on workers and labor composition.

Acknowledgments

The authors are grateful to Lucy Eldridge, John Van Reenen (discussant), Janis Skrastins (discussant), Robert Seamans (discussant), Bryan Seegmiller (discussant), and participants at the AEA, Federal Reserve Digitization Collab Week, the Jozef Stefan Institute Artificial Intelligence Lab, the Labor and Finance Conference, Microsoft Research, the NBER CRIW Conference on Productivity, Technology and Economic Growth, the Organisation for Economic Co-operation and Development (OECD), Stanford University Digital Innovation Lab, the US Government Accountability Office (GAO), the WFA, and Yeshiva University for helpful suggestions. The authors thank Cognism Ltd. for providing the employment data and Burning Glass Technologies for providing the job postings data. Peter Tong and Derek Luan provided excellent research assistance. For acknowledgments, sources of research support, and disclosure of the authors’ material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/technology-productivity-and-economic-growth/firm-investments-artificial-intelligence-technologies-and-changes-workforce-composition>.

References

- Abis, Simona, and Laura Veldkamp. 2024. “The Changing Economics of Knowledge Production.” *Review of Financial Studies* 37 (1): 89–118. See <https://doi.org/10.1093/rfs/hhad059>.
- Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labour Economics* 4: 1043–1171.
- , and Pascual Restrepo. 2019a. “Artificial Intelligence, Automation, and Work.” University of Chicago Press.
- , and ———. 2019b. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.
- , and ———. 2020. “Robots and Jobs: Evidence from U.S. Labor Markets.” *Journal of Political Economy* 128 (6): 2188–2244.

- , and ———. 2022. “Tasks, Automation, and the Rise in US Wage Inequality.” *Econometrica* 90 (5). See <https://doi.org/10.3982/ECTA1981>.
- , David Autor, Jonathon Hazell, and Pascual Restrepo. 2022a. “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” *Journal of Labor Economics* 40 (S1): S293–S340.
- , Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff et al. 2022b. “Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey.” Technical Report 2022.
- , Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti. 2007. “Technology, Information, and the Decentralization of the Firm.” *Quarterly Journal of Economics* 122 (4): 1759–1799.
- Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. 2023. “Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology.” NBER Working Paper 31422. Cambridge, MA: National Bureau of Economic Research.
- Aghion, Philippe and Peter Howitt. 1992. “A Model of Growth through Creative Destruction.” *Econometrica* 60: 323–51.
- Agrawal, Ajay K., Joshua S. Gans, and Avi Goldfarb. 2019. “Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction.” *Journal of Economic Perspectives* 33 (2): 31–50.
- , Joshua S. Gans, and Avi Goldfarb. 2024. “AI Adoption and System-Wide Change.” *Journal of Economics and Management Strategy* 33(2): 327–337.
- , Joshua S. Gans, and Avi Goldfarb. 2023. “The Turing Transformation: Artificial Intelligence, Intelligence Augmentation, and Skill Premiums.” NBER Working Paper 31767. Cambridge, MA: National Bureau of Economic Research.
- Agrawal, Ashwini, Isaac Hacamo, and Zhongchen Hu. 2021. “Information Dispersion across Employees and Stock Returns.” *Review of Financial Studies* 34 (10): 4785–4831.
- Alderucci, Dean, Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas. 2020. “Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata.” *Allied Social Science Associations*.
- Alekseeva, Liudmila, José Azar, Mireia Gine, Sampsa Samila, and Bledi Taska. 2020. “The Demand for AI skills in the Labor Market” CEPR Discussion Paper No. DP14320.
- Altonji, Joseph, Todd Elder, and Christopher Taber. 2005. “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools.” *Journal of Political Economy* 113 (1), 151–84.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *Quarterly Journal of Economics* 118 (4): 1279–1333.
- , Lawrence Katz, and Alan Krueger. 1998. “Computing Inequality: Have Computers Changed the Labor Market?” *Quarterly Journal of Economics* 113 (4): 1169–213.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics* 151: 103745.
- Bessen, James E., Erich Denk, and Chen Meng. 2022. “The Remainder Effect: How Automation Complements Labor Quality.” SSRN Scholarly Paper ID 4042317, Social Science Research Network, Rochester, NY, February.
- Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen. 2014. “The Distinct Effects of Information Technology and Communication Technol-

- ogy on Firm Organization.” *Management Science* 60 (12): 2859–2885. Publisher: INFORMS.
- , Raffaella Sadun, and John Van Reenen. 2012. “The Organization of Firms Across Countries.” *Quarterly Journal of Economics* 127 (4): 1663–1705.
- Bresnahan, Timothy F. 2019. “Artificial Intelligence Technologies and Aggregate Growth Prospects.” *Prospects for Economic Growth in the United States*, edited by John W. Diamond and George R. Zodrow. Cambridge University Press.
- , Erik Brynjolfsson, and Lorin M. Hitt. 2002. “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence.” *Quarterly Journal of Economics* 117 (1): 339–76.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond. 2023. “Generative AI at Work.” NBER Working Paper 31161. Cambridge, MA: National Bureau of Economic Research.
- , 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.
- Caliendo, Lorenzo, Ferdinando Monte, and Esteban Rossi-Hansberg. 2015. “The Anatomy of French Production Hierarchies.” *Journal of Political Economy* 123 (4): 809–52.
- Cao, Sean, Wei Jiang, Junbo L Wang, and Baozhong Yang. 2021. “From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses.” *Journal of Financial Economics* 160: 103910.
- Caroli, Eve, and John Van Reenen. 2001. “Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments.” *Quarterly Journal of Economics* 116 (4): 1449–1492.
- Casal, J. Elliott, and Matthew Kessler. 2023. “Can Linguists Distinguish between ChatGPT/AI and Human Writing? A Study of Research Ethics and Academic Publishing.” *Research Methods in Applied Linguistics* 2 (3): 100068.
- Crouzet, Nicolas, and Janice C Eberly. 2018. “Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles.” *AEA Papers and Proceedings* 2018: 426–31.
- D’Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G Rossi. 2019. “The Promises and Pitfalls of Robo-advising.” *Review of Financial Studies* 32 (5): 1983–2020.
- Davis, Steven J., John Haltiwanger, Ron Jarmin, Javier Miranda, Christopher Foote, and Eva Nagypal. 2006. “Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms.” *NBER Macroeconomics Annual* 2006 21: 107–79.
- Dell’Acqua, Fabrizio, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim R Lakhani. 2023. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.” Harvard Business School Technology & Operations Mgt. Unit Working Paper 24–013.
- Eisfeldt, Andrea L., and Dimitris Papanikolaou. 2013. “Organization Capital and the Cross-section of Expected Returns.” *Journal of Finance* 68: 1365–1406.
- , Gregor Schubert, and Miao Ben Zhang. 2023. “Generative AI and Firm Values.” NBER Working Paper 31222. Cambridge, MA: National Bureau of Economic Research.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.” *arXiv preprint arXiv:2303.10130*.
- Erel, Isil, Léa H. Stern, Chenhao Tan, and Michael S Weisbach. 2021. “Selecting

- Directors Using Machine Learning.” *Review of Financial Studies* 34 (7): 3226–3264.
- Fedyk, Anastassia, and James Hodson. 2023. “Trading on Talent: Human Capital and Firm Performance.” *Review of Finance* 27: 1659–1698.
- , James Hodson, Natalya Khimich, and Tatiana Fedyk. 2022. “Is Artificial Intelligence Improving the Audit Process?” *Review of Accounting Studies* 27: 938–85.
- Felten, Edward W., Manav Raj, and Robert Seamans. 2018. “A Method to Link Advances in Artificial Intelligence to Occupational Abilities.” *AEA Papers and Proceedings* 108: 1–4.
- , Manav Raj, and Robert Seamans. 2023. “Occupational Heterogeneity in Exposure to Generative AI.” Available at SSRN 4414065.
- Fizsbein, Martin, Jeanne Lafortune, Ethan Lewis, and Jose Tessada. 2020. “Electrifying? How New Technologies Impact Productivity and Jobs.” Working Paper 28076.
- Frank, Morgan R., David Autor, James E. Bessen, Erik Brynjolfsson, Manuel Cebrian, David J. Deming, Maryann Feldman, Matthew Groh, João Lobo, Esteban Moro, Dashun Wang, Hyejin Youn, and Iyad Rahwan. 2019. “Toward Understanding the Impact of Artificial Intelligence on Labor.” *Proceedings of the National Academy of Sciences* 116 (14): 6531–6539.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. “Predictably Unequal? The Effects of Machine Learning on Credit Markets.” *Journal of Finance* 77 (1): 5–47.
- Garicano, Luis, and Esteban Rossi-Hansberg. 2006. “Organization and Inequality in a Knowledge Economy.” *Quarterly Journal of Economics* 121 (4): 1383–1435.
- Gofman, Michael, and Zhao Jin. 2022. “Artificial Intelligence, Education, and Entrepreneurship.” *Journal of Finance*, Forthcoming.
- Grennan, Jillian, and Roni Michaely. 2022. “Artificial Intelligence and High-Skilled Work: Evidence from Analysts.” Available at SSRN.
- Guadalupe, Maria, and Julie Wulf. 2010. “The Flattening Firm and Product Market Competition: The Effect of Trade Liberalization on Corporate Hierarchies.” *American Economic Journal: Applied Economics* 2 (4): 105–27.
- Hansen, Stephen, Tejas Ramdas, Raffaella Sadun, and Joe Fuller. 2021. “The Demand for Executive Skills.” NBER Working Paper 28959. Cambridge, MA: National Bureau of Economic Research.
- Hershbein, Brad, and Lisa B. Kahn. 2018. “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings.” *American Economic Review* 108 (7): 1737–72.
- Hitt, Lorin M. 1999. “Information Technology and Firm Boundaries: Evidence from Panel Data.” *Information Systems Research* 10 (2): 134–49. Publisher: INFORMS.
- Jansen, Mark, Hieu Nguyen, and Amin Shams. 2023. “Rise of the Machines: The Impact of Automated Underwriting.” *Management Science*, Forthcoming.
- Jiang, Wei, Yuehua Tang, Rachel (Jiqiu) Xiao, and Vincent Yao. 2024. “Surviving the Fintech Disruption.” *Journal of Financial Economics* (forthcoming).
- Juhász, Réka, Mara Squicciarini, and Nico Voigtländer. 2024. “Technology Adoption and Productivity Growth: Evidence from Industrialization in France.” *Journal of Political Economy* 132(10).
- Katz, Lawrence F., and Kevin M. Murphy. 1992. “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *Quarterly Journal of Economics* 107 (1): 35–78.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017.

- “Technological Innovation, Resource Allocation and Growth.” *Quarterly Journal of Economics* 132: 665–712.
- , Lawrence Schmidt, and Bryan Seegmiller. 2019. “Technological Change and Occupations over the Long Run,” Available at SSRN 3585676.
- Korinek, Anton, and Joseph E. Stiglitz. 2019. “Artificial Intelligence and Its Implications for Income Distribution and Unemployment.” *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Lyonnet, Victor, and Léa H. Stern. 2022. “Venture Capital (Mis)Allocation in the Age of AI.” SSRN Working Paper #4260882.
- Machin, Stephen, and John Van Reenen. 1998. “Technology and Changes in Skill Structure: Evidence from Seven OECD Countries.” *Quarterly Journal of Economics* 113 (4): 1215–1244.
- McElheran, Kristina, and Chris Forman. 2019. “Firm Organization in the Digital Age: IT Use and Vertical Transactions in U.S. Manufacturing.” SSRN Working Paper 3396116.
- Mihet, Roxana, and Thomas Philippon. 2019. “The Economics of Big Data and Artificial Intelligence.” *Disruptive Innovation in Business and Finance in the Digital World (International Finance Review)* 20: 29–43.
- Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381 (6654): 187–92.
- Oster, Emily. 2019. “Unobservable Selection and Coefficient Stability: Theory and Evidence.” *Journal of Business & Economic Statistics* 37 (2): 187–204.
- Peters, Ryan, and Lucien A. Taylor. 2017. “Intangible Capital and the Investment-q Relation.” *Journal of Financial Economics* 123: 251–72.
- Radner, Roy. 1992. “Hierarchy: The Economics of Managing.” *Journal of Economic Perspectives* 30 (5): 1382–1415.
- Rajan, Raghuram G., and Julie Wulf. 2006. “The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies.” *Review of Economics and Statistics* 88 (4): 759–73.
- Rock, Daniel. 2019. “Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence.” Available at SSRN 3427412.
- Romer, Paul M. 1990. “Endogenous Technological Change.” *Journal of Political Economy* 98 (5, Part 2): S71–S102.
- Seamans, Robert, and Manav Raj. 2018. “AI, Labor, Productivity and the Need for Firm-Level Data.” NBER Working Paper No. 24239. Cambridge, MA: National Bureau of Economic Research.
- Tambe, Prassana, Lorin Hitt, Daniel Rock, and Erik Brynjolfsson. 2020. “Digital Capital and Superstar Firms.” NBER Working Paper No. 28285. Cambridge, MA: National Bureau of Economic Research.
- , Peter Cappelli, and Valery Yakubovich. 2019. “Artificial Intelligence in Human Re- sources Management: Challenges and a Path Forward.” *California Management Review* 61 (4): 15–42. SAGE Publications Inc.
- Webb, Michael. 2020. “The Impact of Artificial Intelligence on the Labor Market.” SSRN Working Paper 3482150.
- Zator, Michal. 2019. “Digitization and Automation: Firm Investment and Labor Outcomes.” SSRN Working Paper 3444966.
- Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2020. “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey.” NBER Working Paper 28290. Cambridge, MA: National Bureau of Economic Research.