Maastricht University

AI at work

Citation for published version (APA):

Persyn, P. O. (2026). AI at work: how AI reshapes work and workers' labour market outcomes. [Doctoral Thesis, Maastricht University]. Maastricht University. https://doi.org/10.26481/dis.20261002pp

Document status and date:
Published: 02/10/2026

DOI:
10.26481/dis.20261002pp

Document Version:
Publisher's PDF, also known as Version of record

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy
If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Download date: 27 Nov. 2026

Doctoral thesis

AI AT WORK: HOW AI RESHAPES WORK AND WORKERS' LABOUR MARKET OUTCOMES

Pelin Özgül

2026

AI AT WORK: HOW AI RESHAPES WORK AND WORKERS' LABOUR MARKET OUTCOMES

Dissertation

To obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. Dr. Jan Smits,
in accordance with the decision of the Board of Deans,
to be defended in public
on Friday, October 2nd 2026, at 10.00 hours

by

Pelin Özgül

Supervisors:
Prof. Dr. Didier Fouarge, Maastricht University
Prof. Dr. Mark Levels, Maastricht University

Co-supervisor:
Dr. Marie Christine Fregin, Maastricht University

Assessment Committee:
Chair: Prof. Dr. Gerard Pfann, Maastricht University
Members:
Dr. Femke Cnossen, University of Groningen
Dr. Annemarie Künn-Nelen, Maastricht University
Prof. Dr. Anna Salomons, Utrecht University
Dr. Tania Treibich, Maastricht University

© Pelin Özgül, Maastricht 2026.
All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the author.

Published as part of the ROA dissertation series

Printed by Ipskamp Printing

To my dad

Acknowledgments

I always looked forward to writing this section. Reaching it meant the hard work was (almost) over and that I could finally look back and think of everyone, and all the memories, that brought me to this moment. Now that I am actually writing it, I realize it is one of the most difficult parts of the thesis to write down. Not because I can't think of what to say, but because the weight of the support I have received over these past four years turns out to be much harder to fit into words than I expected.

Someone once told me that a PhD can be a very personal and lonely journey, and advised me to embrace it fully. I am happy to say I have only experienced half of that. It was indeed deeply personal in ways I never anticipated, but not once was it lonely. This section is dedicated to the people who made sure it never was.

First, I would like to thank my promotors, Didier, Mark and Marie-Christine. Three supervisors might sound a lot on paper, but to me it meant something simple: constant learning from three people I look up to. Thank you for believing in me, for guiding the way and for giving me the opportunity to work with you. I am grateful.

Didier, the way you treat your work, your life and the people around you is admirable. I have learned as much from watching how you approach everything you do as from any method. Thank you for all the support you have given me over the years. You teach, you guide, you challenge and you lead, always exactly when each is needed. However busy you were, you always made time. Thank you, for reminding me what real expertise looks like, and why it is exactly the kind of thing no technology can ever replace.

Mark, thank you for always bringing out the best in me. Whenever we talked, whether it was a formal meeting or just a quick phone call, I always left feeling calmer, more confident, and more ready to take on whatever came next. You have a rare ability to make the overwhelming feel manageable, and you do it with a humility and positivity that I admire deeply. You always made sure

your students were the ones in the spotlight, giving us the credit, visibility, and the confidence that came with it. Thank you for showing me that the best research is collaborative, for pushing me when I needed it, and for never wavering in your support and your belief that I could do this.

Marie, you have been a true mentor to me. You helped me grow, not just professionally, but also personally. I remember the first call we had, when you offered me the opportunity to join the project, and how that conversation changed the course of my career and my life. Thank you for believing in me and for giving me the space to learn, to build, to contribute and to figure out the kind of researcher I want to be. I still don't know how, but you put together a team that is so complementary; one I felt incredibly lucky to be part of. What sets you apart is that the project was never just about the work; it was about the way we did it, together, through the failures and the successes. That is something I will take with me, always. Thank you for inspiring me to dream big, not to be afraid to be vulnerable, and for always reminding me of the strength that comes with it.

I also feel grateful to have had a defence committee that I have long looked up to. I am thankful to Gerard Pfann, Anna Salomons, Annemarie Künn-Nelen, Tania Treibich, Femke Cnossen and Melline Somers, whose research, and in many cases whose way of thinking and being, has shaped this thesis, and me, in more ways than they likely realize. Thank you for the time and expertise you gave to my work, and for the discussion I so look forward to.

Then there are my colleagues, and in particular those in the ai:conomics project, with whom I got to enjoy all the milestones (and the setbacks) along the way. Sanne, it feels natural to start with you because you were one of the first people I met, giving me an early glimpse of what the next 4 years would look like. I could go on endlessly about your data skills, your way with people or how you somehow managed every legal and administrative headache that would have buried the rest of us (well, at least me and Nick, for sure). Half the time we didn't even know there was a problem, because you had already solved it. But what I am most grateful for is everything you brought outside of office hours (and the sweet treats during office hours which successfully kept my blood sugar above my stress levels). You were the one who organized

every social gathering, every birthday celebration, every reason for us to come together and connect beyond work. Thank you, for all the drinks, the concerts, the good memories in Maastricht, Berlin and Mechelen, for always finding a way to get the funniest pictures of Nick, and for making sure I never missed a moment, even when I wasn't there.

Sander, I also met you on my very first week, when your compliment about my code left me prouder than I'd like to admit. But it wasn't until you joined the project that I realized the true joy of working with you. You changed the dynamic in a way that is hard to put into words, in your own quiet, steady (and let's admit, sometimes stubborn) way. Once you joined, you brought all the knowledge, skills, and the right questions we needed, but what I enjoyed most was watching you let us challenge you too. Working with you was like working with the best AI agent: I never got the answer to my question, left with more questions than I started with, and yet somehow that was exactly the point. You see what the rest of us miss, you think in ways we don't, and I am a better researcher for having worked with you. Thank you for letting me also push you into unfamiliar territory, for complaining about it, and for always coming back grateful.

Danique, joining a project that's already in motion is never easy. But you found your place so quickly and so naturally that it's hard to remember a time when you weren't part of the team. You cared about the project deeply, took on tasks that might have been outside your comfort zone, and blended into this team as if you had been there from the start. Thank you for bringing the energy that uplifts the room, for being positive, and more committed to the work than most people realize.

Lara, *canım*, what started as sharing an office and a project turned into a friendship in this whole journey. We shared so much more than work: our struggles, successes, coffee, and venting sessions, social events, weekend getaways, and the kind of conversations that gave me strength when I needed it most. You have a creativity that goes beyond your research: the way you write, formulate, and create while being completely unaware of how good you are, is remarkable. Thank you for being there through all of it, the good and the bad, for inspiring me in ways you probably don't realize, and for giving

me the rare comfort of speaking in my own native language for things that would otherwise get lost in translation, and also for the harmless gossip that made everything a little more fun.

Beyond the "engine room," I have had the opportunity to work alongside the senior minds in the project, whose guidance I am thankful for. On the IAB side: Simon, I have learned so much from you. You showed me the way at the IAB, literally and figuratively. You patiently answered my endless questions, helped me with my research, kept me (mostly) sane through the heaviest stretches of data work, and pushed me, always, to get better. If I can carry even half of your devotion into my own work, I will be very proud of the researcher I become. Thank you, for being a true academic, and someone I was lucky to learn from. Michael, you were a constant support through much of the same. You helped me find my way through the data and the research, kept me grounded when it all became too much, and were endlessly generous with your time and knowledge. Presenting at conferences was far less daunting knowing you were in the audience, and a lot more fun afterwards over drinks. Thank you for being a colleague first and a friend soon after, and for showing me that rigour and kindness are never a trade-off. On the ROA side: Andries, Rolf, Tim, Raymond, Katarina; your guidance, perspective and expertise brought the direction and clarity when it was most needed. Thank you for helping me see the bigger picture.

While a PhD can challenge you in every aspect, my two paranymphs got me through it, the two I could not have done it without. Nick, let me first thank you, in your own way: Thanks mate, it's been fun. Anna, I feel so lucky to have met you. From the moment we met, you became so much more than a colleague. We worked together, (part-time) lived together, shared a huge chunk of our lives and somehow never ran out of things to talk about. The conversations we had brought me so much joy, and kept me grounded more than you know. I admire the way you move through life, with kindness, intention, and a genuine desire to lift the people around you. Thank you for always lifting me up, for being the kind of support I can't imagine this journey without, and for making the last five years lighter simply by being my friend. Nick, I thanked you in your way, but you know me better than that. So here it is, now in my way. I don't think I could have asked for a better person to

do this PhD alongside. We complemented each other in ways that shouldn't have worked but always did. The IAB trips that tested our patience, the data that tested our sanity, and all the rest that tested both: I shared every bit of it with you. For the tallest person I know, you have been remarkably easy to lean on. This journey was much better, and much more fun, thanks to you. You have the kindest heart, a talent I'm not sure you fully realize, a way of seeing things that always put mine into perspective, and you are (most of the time) a joy to work with. Thank you for always listening, for your critical eye that always made my work better, for making me laugh without trying, and for being unapologetically, and chaotically, yourself.

It has been a privilege to be part of ROA's vibrant community. The support, the collaboration, and the warmth of the people made it more than just a place to work. I want to thank all my colleagues at ROA for making it such a pleasant and inspiring place. Thank you, not just for the research that taught me so much, but also for everything around it: the lunch breaks, the drinks, the parties, the costume competitions, the predictions, and all the other moments. A special thank you to the secretariat and the research data management team, for always being ready to help with every possible question and providing all the support. Thank you also to my colleagues at the IAB; Theresa, Verena and Milan, and colleagues at zukunft zwei for their support throughout.

Mantej, my friend, sharing an office with you and Nick was one of the best parts of this PhD. I am grateful for everything we shared; the discussions at the whiteboard, the practice runs for our presentations and all our insightful conversations that started about research and ended up being about life. Your enthusiasm and dedication for your research was contagious and often reminded me to enjoy the process. You are an incredible researcher, and even a better friend. Thank you for the laughs, the sharp feedback on my research, and for making it impossible to have a bad day in that office, except the times when I had to stare at that hideous painting you and Nick found on the street for free and proudly hung right in front of my desk. Evie, I feel so lucky to have gained not just a colleague, initially, but a friend who stayed long after. I admire your confidence, calmness, and the ambition you carry. From our summer holidays to every catch-up in between, you made this journey much more enjoyable. I cannot wait to see everything that is ahead of you and to

be there to celebrate with you. Thank you for being a great friend, for always picking up where we left off, and for making distance feel like nothing. Sandra, I am grateful for all our coffee breaks, catch-ups, and deep discussions. Thank you for your friendship, for always making time, and for sharing the special moments with me. Mariët, thank you for the amazing time in Prague, and for every conversation, coffee, and catch-up since.

I also want to thank my colleagues at FASOS; Toon, Michele, Luana and Darian, for welcoming me into their team, for giving me the space to contribute, and for building our project together. Working with you was a chance to be part of something new, and gave me the perspective and skills that strengthened me as a researcher. Thank you, for all the work we did together, for the stimulating discussions and for your continuous support throughout.

I have been incredibly fortunate to build friendships and connections that extended far beyond ROA. From Ankara to Groningen, Leuven, Nuremberg and Brussels, I have met people who challenged me, supported me, and shared wonderful memories with me. To all my former and current colleagues, and to all the friends who have been part of this journey, thank you for being in my life. I would also like to acknowledge the many women in economics and research whose example has inspired me over the years, often more than they may ever know. Some have been mentors, some collaborators, some colleagues and friends, and some role models from afar. Seeing brilliant women lead, question, create, and support those around them has shaped not only the kind of researcher I hope to become, but also how I hope to contribute to this field. I hope your influence continues to inspire every young researcher who follows.

No one succeeds alone, and I certainly did not. Every opportunity I have had, every challenge I have overcome, and every milestone I have reached has been built on the love, encouragement, and support of my family. To the families Goris and Persyn, thank you for welcoming me with open arms and for making Belgium my second home. The family weekends, Christmas celebrations, holidays, and all the gatherings brought me so much joy, and gave me the breaks I needed. Thank you for reminding me of life beyond work, and for filling mine with happiness, warmth, and a sense of home. To my *gezin*, Ingrid, Johan, Thomas and Eline, thank you for making me feel like part of

the family from the very beginning. You welcomed not only me, but also my culture, with immense curiosity and an open heart, always willing to try the occasional Turkish meal that I (and Bruno) put in front of you, embracing and making space for the little things that make me who I am. Thank you for your endless love, kindness, patience, and unconditional support. To my brother, *Murat*, and my mother, *Inci*, you two are my safest place, no matter where in the world I am. Thank you for being my foundation through everything, for every visit you made to be with me and for reminding me that I never have to face anything alone. I am incredibly grateful to have you by my side. My final thanks are reserved for the two most special people in my life: Bruno, I work in a field where we strive for randomness and variation. But in my life, you have been the one constant. Thank you for your incredible patience (you asked me to write that...), your endless support, and for carrying me in times when I found it difficult to carry myself. Thank you for embracing my every failure and achievement as if it were your own, for building this life with me and for being exactly the partner I always hoped for. Even though I know things will never be easy, I hope it remains simple. *Babacim*, thank you for believing in me long before I learned to believe in myself. I am where I am thanks to you. So I hope, wherever you are, you are watching, and feeling proud. This dissertation is as much yours as it is mine.

Pelin Özgül
Brussels
June 2026

Contents

Acknowledgments vii

1 Introduction 1
1.1 Motivation . . . 2
1.2 This Thesis . . . 6
1.3 Outline and Main Results . . . 7

2 Rethinking Skill and Routine Bias 13
2.1 Introduction . . . 15
2.2 Related Literature . . . 18
2.3 Data and Descriptive Statistics . . . 21
2.4 AI Exposure and Labour Market Outcomes . . . 32
2.5 Results . . . 37
2.6 Conclusion . . . 52
Appendix 2A: Additional Figures & Tables . . . 55
Appendix 2B: Linking administrative data with exposure scores . . 68

3 AI in High-Tech Manufacturing: Firm-Level Evidence from a Multi-Plant Rollout 69
3.1 Introduction . . . 71
3.2 Empirical Context . . . 75
3.3 Data . . . 81
3.4 Results . . . 87
3.5 Robustness Checks . . . 97
3.6 Conclusion . . . 101
Appendix 3A: Additional Figures & Tables . . . 104
Appendix 3B: Identification Challenges . . . 109

4 How AI-Augmented Training Improves Worker Productivity 111
4.1 Introduction . . . 113
4.2 The Customer Service Department . . . 117

4.3 The Training Program . . . 119
4.4 Identification and Estimation Strategy . . . 121
4.5 Data and Summary Statistics . . . 124
4.6 Results . . . 128
4.7 Conclusion . . . 147
Appendix 4A: Potential violation of SUTVA . . . 149
Appendix 4B: Alternative approaches to overcome the bias in classical TWFE estimation . . . 151
Appendix 4C: Treatment effects on time spent for different activities during the call . . . 152
Appendix 4D: Use of communication styles and call length . . . 154

5 Generating Anxiety, but not Policy Demands? AI Competition and the Policy Preferences of Tomorrow's Workforce 155
5.1 Introduction . . . 157
5.2 Young Citizens, Their Preferences, and a Changing Labour Market . . . 159
5.3 Experimental Design, Data and Methodology . . . 164
5.4 Results . . . 172
5.5 Conclusion . . . 182
Appendix 5A: Additional Figures & Tables . . . 184
Appendix 5B: AI-related Articles . . . 193
Appendix 5C: Survey Questionnaire . . . 203

6 Conclusion 207

Bibliography 215

Impact 231

Summary 237

Samenvatting 241

Published work 245

About the author 247

1

Introduction

1.1 Motivation

Advances in Artificial Intelligence (AI) have emerged as one of the most fundamental forces driving the reorganisation of labour markets over the past decade. Developments in machine learning (ML) technologies and AI models are reshaping occupations, task content and work processes across all sectors of the economy. From manufacturing to professional services, AI systems are increasingly deployed to perform cognitive non-routine tasks, ranging from quality inspection and data processing to text or image generation and supporting decision making (OECD, 2023). The economic and social implications of these developments extend beyond individual workers, affecting firm strategies, labour market institutions and public policy whilst raising fundamental questions about how the future of work will be organized and how labour markets will adjust.

The rapid diffusion and wide range of applications of AI technologies have also intensified public discourse on the opportunities and challenges brought about by AI, with concerns ranging from large-scale job losses, skill obsolescence and rising wage inequality to expectations of unprecedented productivity gains and economic growth. Public opinion surveys highlight this divide: a recent Eurobarometer study demonstrates that, across EU member states, majority of the respondents expect AI to improve productivity and economic performance, while simultaneously expressing concerns about job destruction, negative effects on workplace interactions, such as communication between colleagues and the need for regulation (European Commission, 2024). In the United States (US), Gallup data reveals that about 75% of US adults worry about the overall disruptive impact of AI on the job market and believe that AI will reduce the number of available jobs over the next 10 years (Marken and Nicola, 2023; Autor, 2024). Taken together, these perceptions underscore the economic and social relevance of AI and call for rigorous empirical investigation on how AI affects work, workers and labour markets.

Central to this debate is the dual prospect of AI as a technology platform, with the potential both to substitute for human labour and to complement worker capabilities. As explained in the task-based framework developed by Acemoglu and Autor (2011) and Acemoglu and Restrepo

(2019), AI can function as a labour-displacing, automating technology that substitutes for certain tasks that were previously performed by labour, potentially reducing labour demand and putting downward pressure on wages and employment for certain groups of workers. At the same time, AI can act as a labour-enhancing technology that augments workers' skills and raises productivity. Beyond its effects on labour demand and productivity, AI can also reshape the way work is performed by altering the nature and composition of tasks within jobs. It has the potential to free up time, enabling workers to focus more on creative and higher-level tasks. These complementarities can generate new employment opportunities, extend expertise to a broader set of workers and lead to higher wages for workers whose skills are well suited to work with the new technologies (Acemoglu and Restrepo, 2019; Autor, 2024; Autor et al., 2024). Yet despite heightened public anxiety and academic debate, determining which of these forces will prevail, or whether both dynamics will unfold simultaneously across different contexts, remains an open empirical question.

Understanding the Impact of AI on the Labour-Market

Against this background, the labour market implications of AI have become a salient topic in economic research. Economic theories of technological change have traditionally emphasized how automation, such as the adoption of information and communication technologies (ICT) and robotics, affect workers through skill-biased or routine-biased channels. These neoclassical frameworks have shown that technological progress disproportionately benefits high-skilled, high-educated workers while displacing workers performing routine tasks, thereby shaping employment and wage dynamics across occupations, leading to job polarization (Goldin and Katz, 2008; Acemoglu and Autor, 2011; Goos et al., 2014). While highly influential in explaining the impact of earlier waves of automation, these theories were developed in contexts where technologies primarily substituted routine and manual activities.[1] Recent advances in AI, however, are increasingly related to systems that perform cognitive tasks, such as absorbing tacit knowledge, recognizing patterns and making inferences from unstructured data (Autor, 2024).

[1] For a detailed discussion of these theoretical frameworks and their applicability in the context of AI, see Chapter 1.

These advances therefore raise questions about whether established theoretical frameworks adequately capture AI's potential labour-market impact and the range of jobs that could be affected.

Empirical research on AI and labour markets has evolved along several dimensions. To date, one of the most common approaches involves constructing exposure measures to identify which occupations are most susceptible to AI automation. These measures assess whether specific tasks can be performed by AI technologies and provide exposure metrics for various occupational groups (del Rio-Chanona et al., 2025). Existing studies rely on various data sources and methodologies, including expert assessments (Frey and Osborne, 2017b; Felten et al., 2018), patent-based measures of technological capabilities (Webb, 2020; Prytkova et al., 2024; Mann and Püttmann, 2023; Grimm and Gathmann, 2022) and more recently, evaluations based on large language models (LLMs) such as ChatGPT (Eloundou et al., 2024; Gmyrek et al., 2025). Exposure indices serve as important indicators for mapping the potential scope of AI's labour-market impact. By systematically assessing which occupations and tasks could theoretically be automated, these measures offer a prospective view of where AI technologies might be deployed, particularly in contexts where data on actual adoption are limited. They provide a useful starting point for understanding economy-wide patterns of susceptibility across occupational groups. At the same time, while exposure measures remain informative in assessing the feasibility of automating tasks, most indices do not indicate the likelihood that AI technologies are actually adopted, nor do they distinguish between substitution and complementarity in the use of AI at the task level. As a result, they offer limited insight into the pace, intensity and organizational dynamics of actual AI deployment as well as how adoption translates into changes in worker outcomes. [2]

More recently, researchers have started conducting randomized controlled trials (RCTs) to measure productivity effects of AI in narrow

[2] Another strand of the literature relies on online job vacancy data, such as postings from Lightcast and related providers to study how AI reshape firms' labour demand. By analyzing changes in the skill and task demand in online job postings, this line of research provides evidence on how AI adoption is reflected in job requirements across occupations and industries. For more detail, see Alekseeva et al. (2021); Peede and Stops (2024); Storm et al. (2025).

work domains. Overall, experimental evidence points out that AI and particularly generative AI applications (GenAI) can improve worker performance in tasks such as writing, coding, legal reasoning and medical diagnostics and facilitate human-machine collaboration (Noy and Zhang, 2023; Agarwal et al., 2023; Brynjolfsson et al., 2025b; Choi and Schwarcz, 2023; Rounding et al., 2025; Dell'Acqua et al., 2025).[3] Although they offer valuable worker-level insights into AI's potential as a labour-augmenting technology, these studies primarily examine direct usage of AI at the point of executing the task. By contrast, much less is known about the implications of indirect uses of AI solutions at workplaces, such as AI-enabled on-the-job training, feedback, or performance evaluation systems that can also shape worker skills, performance and learning over time. Understanding how these applications affect productivity, skill development and heterogeneity across workers remains an important gap in the literature.

Lastly, despite growing public concern about labour-market implications of AI, empirical evidence on how expectations about AI shape workers' behaviour and their policy preferences remains limited. Emerging early evidence from the US indicates that AI reduces employment opportunities for early-career workers (ages 22–25), highlighting that adverse effects of AI may materialize even before workers accumulate substantial labour-market experience (Brynjolfsson et al., 2025a). Yet little is known about how individuals form beliefs about future competition with AI and whether these beliefs influence attitudes toward regulation, labour-market protection, or upskilling strategies that are needed to develop skills for the future. This gap is particularly pronounced for prospective labour-market entrants, whose expectations about AI will shape educational choices, occupational sorting and policy preferences prior to any direct workplace exposure. Therefore, understanding how perceived AI risks and opportunities affect policy demands among future workers represents an important yet underexplored dimension of AI's labour-market consequences.

[3] Importantly, these studies demonstrate substantial heterogeneity in productivity gains from the use of AI tools.

1.2 This Thesis

This thesis addresses these gaps through four complementary empirical studies that examine the labour-market effects of AI at different levels of analyses. Written in collaboration with several co-authors, the four chapters provide new empirical evidence on how AI reshapes work and workers' labour market outcomes. Drawing on occupation-level exposure measures, worker and firm-level data and evidence on beliefs and perceptions on AI, this thesis aims to contribute to the empirical literature on technological change by studying the mechanisms through which AI affects labour markets.

The four chapters employ diverse empirical strategies, ranging from observational analyses of occupational exposure and wage dynamics to firm-level case studies of AI deployment, quasi-experimental evaluations of worker-level productivity effects and a survey experiment on policy preferences. Figure 1.1 illustrates the key elements of this thesis and demonstrates how these four chapters relate to the overarching theme of this thesis. Each chapter addresses a distinct but related dimension of AI's impact on the labour markets: Chapter 2 examines occupational exposure to AI and its implications for workers' wages and employment; Chapter 3 studies AI-driven automation and task substitution within a multinational firm; Chapter 4 investigates the introduction of an AI-augmented technology on worker productivity at a European call center; and Chapter 5 analyzes how perceived competition with AI shapes prospective workers' policy preferences.

The multi-level perspective employed throughout this thesis aims to highlight both the multifaceted impact of AI and the methodological diversity that is necessary to fully capture the breadth and depth of AI-driven transformation that shapes today's workforce. Rather than offering a single methodological lens, this thesis combines correlational and experimental evidence —from using early exposure metrics to studying firm-level AI adoption dynamics, worker productivity effects and more recently, investigating perceptions in the generative AI era —thereby mirroring the evolution of empirical AI research itself. Although each approach offers distinct advantages while facing limitations, employing a broad perspective is necessary to understand a phenomenon as com-

plex and rapidly evolving as AI. By doing so, this thesis aims to provide a comprehensive understanding of the potential implications of AI on labour markets and suggest that these implications can vary systematically across contexts, organizational settings as well as the design of AI systems.

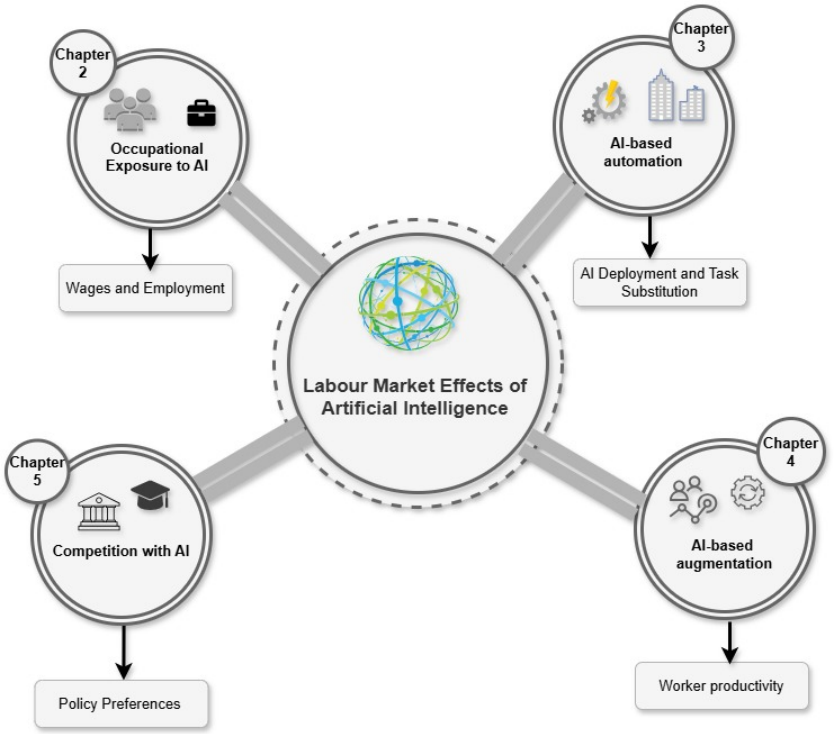


Figure 1.1: Outline of this thesis

1.3 Outline and Main Results

This thesis is structured around four empirical studies that study the labour-market effects of AI. The following paragraphs introduce the main research questions, empirical approaches and central findings of each chapter.

To what extent are workers exposed to AI and to what extent can AI

exposure explain workers' labour market outcomes?

Chapter 2 presents an observational study to examine which workers are most susceptible to AI and whether AI exposure is associated with observed labour-market outcomes. The chapter begins with a discussion of the existing literature on the labour-market effects of technological change. Building on canonical task-based models, it then provides descriptive evidence on how exposure to AI varies across occupations with different skill levels and task compositions and examines how AI exposure relates to patterns of wage growth.

Our analysis combines administrative employment biographies of workers in Germany with task-based measures of occupational exposure to AI, developed by Webb (2020), which enables us to capture heterogeneity in AI susceptibility within occupations between 2012-2019. We observe that AI susceptibility is most pronounced across workers in highly skilled, analytical and non-routine occupations. Moreover, we demonstrate that while workers in cognitive-routine occupations with higher AI susceptibility experience faster wage growth, wage premia at the top and bottom of the skill distribution appear largely unrelated to AI exposure. These findings suggest that AI affects labour markets in ways that might differ from earlier waves of automation and challenge established theoretical frameworks.

How do firms adopt AI technologies and what is the relation between AI deployment and labour demand?

Chapter 3 shifts the focus from potential exposure to actual adoption and studies how an AI-driven automation technology unfolds within a firm and how this relates to labour demand at the task level. Using detailed firm-level data from a multinational high-tech manufacturing company that introduces an AI-based visual inspection system, we document the dynamics of AI deployment across plants over time. Rather than observing discrete automation shocks we find a gradual expansion of AI automation, characterized by substantial heterogeneity across plants in both the level and pace of deployment. We demonstrate that AI automation expands primarily along the extensive margin, with plants deploying the technology to additional products. These adoption patterns reflect organizational constraints, data availability, customer

requirements and the need to build trust in AI systems in a high-stakes production environment.

We find a robust negative relation between the degree of AI automation and operator decision volumes, consistent with task-level substitution effects of AI within the production processes. Taken together, our findings highlight that while AI can displace specific tasks, its labour-demand effects can occur incrementally and suggest that its implications are shaped by organizational and technological feasibility rather than by immediate, large-scale workforce reductions.

How does introduction of a labour-augmenting AI tool into the training system impact worker productivity?

Chapter 4 moves beyond AI-driven automation and investigates whether AI can act as a labour-enhancing technology by studying how the introduction of an AI-enhanced training affects worker productivity. Using a staggered rollout of an AI-based feedback tool in a call-center setting, we examine the causal impact of AI-augmented training on performance outcomes, exploiting variation in the timing of adoption across teams. Importantly, the AI tool does not replace or directly assist workers; instead, it supports trainers by generating detailed, data-driven feedback that enables more targeted and individualized training.

We find that AI-augmented training significantly reduces call handling times and improves worker productivity, with particularly strong effects for short-tenured agents. At the same time, the analysis reveals productivity gains for more experienced workers, operating through different margins: while less-experienced agents benefit primarily from reduced hold times and improved handling of complex calls, longer-tenured agents improve their communication style. These productivity gains are accompanied by modest improvements in customer satisfaction. Overall, our findings highlight the role of AI as a complement to human labour through organizational channels such as training and learning and demonstrate that the productivity effects of AI are conditional on worker characteristics and the institutional context in which AI is deployed.

How do workers' perceptions of AI competition shape their preferences for labour market regulation?

Chapter 5 takes the findings of the previous chapters and asks, amid this balance between automation and augmentation, is there room for public policy? Focusing on prospective labour-market entrants, this chapter examines how perceived competition with AI shapes policy preferences before individuals acquire direct workplace exposure. Using a pre-registered survey experiment among university students in the UK and the US, we study how information about generative AI affects anxiety about future job prospects and demand for regulation.

The results show that exposure to information about generative AI increases concerns about labour-market risks. However, this heightened anxiety does not translate into stronger support for traditional welfare-based labour protection policies. Instead, respondents express a preference for regulating AI technologies directly and policy attitudes are strongly mediated by pre-existing ideological positions. These findings point to a disconnect between rising anxiety about AI-driven labour-market change and demand for conventional redistributive policies, highlighting the role of beliefs, expectations and political identity in shaping future policy debates around the governance of AI and the future of work.

Taken together, **Chapters 2-5** reveal that AI-driven labour-market change operates through multiple, interconnected channels: technological feasibility of task automation, organizational decisions about adoption and implementation, productivity effects through AI-augmented training and prospective workers' anticipatory policy responses. Across these dimensions, AI emerges neither as a purely labour-displacing nor a purely labour-augmenting technology, but as one whose heterogeneous effects will depend on task characteristics, organizational context and design choices. While Chapter 3 documents task substitution and Chapter 4 demonstrates complementarity between workers and AI, these seemingly contradictory findings indicate that both dynamics can unfold simultaneously. AI can replace specific tasks while enhancing worker productivity in others, with the balance shaped by how organizations design, deploy and integrate AI systems into work processes. Rather than operating through a single dominant mechanism, AI's labour-market

effects reflect the interaction of competing forces whose relative importance varies across tasks, organizational and institutional settings.

Lastly, **Chapter 6** provides an overall conclusion by reflecting on the empirical findings of the thesis, discussing their broader implications and outlining avenues for future research.

2

Rethinking Skill and Routine Bias: AI Susceptibility in Non-Routine, High-Skilled Jobs

Abstract

Artificial Intelligence (AI) will transform work by taking over specific job tasks, yet debate remains about which tasks are susceptible to automation, and how AI affects workers' wages. Using matched employer–employee data from Germany and combining information on job skills and task content with a measure of AI susceptibility, we document that high-skilled workers and workers in analytical non-routine jobs are among the most exposed to AI. We show that wage growth premia for low- and high-skilled workers are largely unrelated to AI exposure. We demonstrate that AI exposure is positively associated with wage growth for workers in occupations with predominantly routine tasks. Our findings suggest that AI may affect a broader set of workers and in ways that differ from the predictions of canonical economic theories of technological change.

This chapter is joint work with Marie-Christine Fregin, Michael Stops, Simon Janssen and Mark Levels, adapted from Ozgul et al., 2024. My individual contribution to this paper involved contributing to the development of the research question and methodology, accessing and preparing the register data (both on site and remote), performing the data analysis, contributing to the interpretation of results and writing the manuscript. This study is part of the research project ai:conomics funded by the German Federal Ministry of Labour and Social Affairs (Bundesministerium für Arbeit und Soziales (BMAS)/Denkfabrik Digitale Arbeitsgesellschaft) by resolution of the German Bundestag.

2.1 Introduction

The rapid acceleration of artificial intelligence (AI) has expanded the range of tasks that can be automated or augmented through machine-learning systems. However, while the exponential diffusion of AI technologies is transforming many sectors of the economy, the impact these developments will have on individual workers is unclear (Noy and Zhang, 2023). Although earlier automation technologies produced well-documented patterns of occupational restructuring and shifts in labour demand, the way AI's labour market implications will unfold remains unknown, raising a broader question that motivates this paper: *is this time different?*

Two complementary hypotheses have guided the literature on how technological change affects workers and employment. The skill-biased technological change (SBTC) hypothesis posits that technology substitutes for low-skilled workers while complementing high-skilled workers, increasing the relative demand for skills (Acemoglu and Autor, 2011; Goldin and Katz, 2008). The routine-biased technological change (RBTC) hypothesis emphasizes the displacement of routine-intensive tasks in middle-wage occupations during the computerization era (Autor et al., 2003; Autor and Dorn, 2013; Goos and Manning, 2007; Goos et al., 2009). Both hypotheses suggest that automation primarily threatens workers performing routine, codifiable tasks while protecting those in high-skill, non-routine occupations. While both frameworks have been central to explaining how previous waves of technological shifts have shaped labour market outcomes (Autor, 2015; Acemoglu and Restrepo, 2020b), whether they remain useful for understanding the impact of AI is not yet clear. As Goos et al. (2022) argue, observed labour-market outcomes are not technologically predetermined but depend critically on the interplay between the nature of technological progress itself and how new technologies interact with workers with different skill levels, making the applicability of established frameworks to AI an open empirical question.

This is because AI represents a different technological frontier that may depart from earlier automation patterns. Unlike robotics and traditional software which were largely confined to manufacturing and routine office tasks, AI systems increasingly perform cognitive tasks involving perception, prediction,

and contextual judgment (Felten et al., 2021; Webb, 2020; Autor, 2024). This means, if AI substitutes for tasks in occupations which were historically insulated from automation, its effects may differ substantially from what SBTC and RBTC predict.

An emerging body of literature has examined the labour-market effects of AI. Various findings suggest that high-income occupations, which are often assumed to be more automation-resilient, are among the most exposed to AI (Acemoglu et al., 2022; Webb, 2020; Felten et al., 2019; Grimm and Gathmann, 2022). At the aggregate level, research finds modest wage gains and no evidence of economy-wide displacement (Alekseeva et al., 2021; Georgieff and Hyee, 2022; Acemoglu et al., 2022; Peede and Stops, 2024; Kinder et al., 2025), though some recent evidence suggests early-career workers may be disproportionately affected (Brynjolfsson et al., 2025a). However, most of existing research relies predominantly on firm, occupation, industry or local labour data, obscuring how AI affects individual workers and the specific tasks they perform. Relatively little is known about the heterogeneity in AI exposure across workers within narrowly defined occupations.[1] Capturing the heterogeneity of AI exposure at the task level is critical to evaluate whether—and, if so, the extent to which— its labour-market implications diverge from those of earlier technological change. This distinction is also important because it shapes how we assess the impact of AI, whether entire high-wage occupations are at risk, or only specific task bundles within them, thereby shaping both policy responses and strategies for workforce adaptation.

This paper addresses these questions by providing empirical evidence on the worker-level impact of AI, leveraging within-occupation variation in AI exposure. We analyze individual-level employment and wage outcomes in Germany between 2012 and 2019, focusing on the early stages of AI diffusion and on machine-learning technologies prior to the advent of large language models (LLMs). We use the SBTC and RBTC hypotheses as interpretive frameworks to assess whether AI follows familiar patterns of previous technological

[1] A small but growing literature examines the impact of generative AI in experimental settings, typically finding productivity gains for very specific tasks or within narrow work contexts; see Noy and Zhang, 2023; Brynjolfsson et al., 2025b; Peng et al., 2023; Otis et al., 2024 for more detail.

advancements. Our empirical strategy exploits within-occupation variation using exposure scores for AI, software, and robotics developed by Webb (2020). These exposure scores are based on textual similarity between patent descriptions and occupational task content and enable direct comparison across distinct technological episodes.[2] Combining these technology exposure measures with individual labour-market biographies from Germany allows us to construct a unique dataset that captures differences in worker composition within occupations and to investigate how routine and skill content of occupational tasks shape the relationship between AI exposure and worker outcomes.

Our analysis proceeds in four steps. First, we descriptively examine the heterogeneity in occupations' susceptibility to AI automation.[3] Second, we analyze how susceptible occupational tasks are to AI automation and determine the extent to which job tasks' susceptibility to AI automation is related to their skill complexity and task routineness level. We compare these patterns with occupations' susceptibility to robotization and software exposure which reflect the effects of previous technological episodes. Third, we assess potential displacement effects by studying whether workers are more likely to leave occupations that are more strongly susceptible to AI. Fourth and last, conditional on observed employment outcomes, we investigate the association between AI exposure and wage trajectories among workers in different occupation groups.

Our results suggest that AI may represent a departure from earlier automation patterns. We find that workers in high-skilled, non-routine occupations are disproportionately exposed to AI, contrasting sharply with earlier automation waves that primarily targeted routine-intensive middle-wage jobs. Despite this high exposure, we observe no evidence of increased job separation risk for

[2]The exposure measures of Webb (2020) are well suited for our analysis for two reasons. First, Acemoglu et al. (2022) show a strong association between Webb's AI exposure scores and subsequent AI hiring, validating it as a credible proxy for adoption (Acemoglu et al., 2022). Second, by construction, these exposure scores enable direct comparison across different technological episodes, which is crucial for testing whether the predictions of SBTC and RBTC can be extended in the context of AI.

[3]Throughout this chapter, we use the terms AI susceptibility and AI exposure interchangeably. While conceptually distinct, both terms refer here to the same task-based measure capturing occupations' potential exposure to AI.

workers in AI-exposed occupations. We also find that AI exposure is positively associated with wage growth between 2012 and 2019, with heterogeneous effects across occupational groups. These findings indicate that, at least in the early stages of AI diffusion, AI may affect labour markets differently than what the SBTC and RBTC hypotheses predict and suggest that AI may reshape both the content and wages of high-skilled work through task restructuring rather than large scale job loss.

Our findings contribute to several strands of the literature. First, we provide new empirical evidence on how AI exposure relates to individual workers' employment and wage outcomes by exploiting within-occupation variation in task composition. Through a systematic task-level comparison of AI, robotics, and software exposure within occupations, we assess whether AI targets fundamentally different capabilities than earlier technologies. Second, by linking occupation-based exposure measures to individual administrative data, we shed light on granular worker-level responses to AI exposure—such as job mobility and wage dynamics—that are difficult to capture in studies relying on aggregate adoption measures. Third, building on the view that automation both substitutes and augments human tasks, we examine how skill complexity and routineness shape the relationship between AI exposure and wages, and evaluate whether canonical frameworks of technological change remain useful for understanding AI's labour-market implications.

The remainder of the paper is structured as follows: Section 2.2 provides related literature on technological change. Section 2.3 describes the data and presents descriptive statistics. Section 2.4 outlines the empirical strategy and Section 2.5 presents the findings. Section 2.6 concludes.

2.2 Related Literature

This section reviews the main models examining how technological change reshapes labour markets through changes in task content, skill demand, and occupational structure. Two canonical frameworks, the skill-biased technological change (SBTC) and routine-biased technological change (RBTC) have played a central role in organizing empirical and theoretical work in this area.

2.2.1 Skill-Biased Technological Change

The SBTC hypothesis argues that technological progress increases the relative demand for skilled workers by complementing tasks requiring higher levels of education, problem-solving ability, and abstract reasoning, while substituting for workers performing simpler tasks (Goldin and Katz, 2008; Acemoglu and Autor, 2011). As such, it predicts that the more skilled workers are, whether defined by educational attainment or by occupational status, the less likely it is that they will lose their jobs to machines, and the more likely it is that technology complements their skills. Under this framework, technology adoption raises wage inequality by disproportionately benefiting high-skilled workers and putting downward pressure on the wages and employment prospects of low-skilled workers (Autor et al., 1998). Although not without debate (Card and DiNardo, 2002), ample empirical scrutiny demonstrates that SBTC is largely consistent with a wide variety of empirical observations, at least to explain past technological developments. For example, the computerization wave of the 1990s is often cited as a period during which digital technologies increased the labour share and relative earnings of highly educated workers while reducing demand for less-skilled labour, leading to a rise in skill premium (Autor, 2014; Acemoglu, 2021). More broadly, SBTC has been used to explain why earlier industrial revolutions favoured workers with higher levels of formal education, as rising educational attainment allowed labour to complement new technologies (Goldin and Katz, 2008). Due to this impressive empirical validation, SBTC has long been viewed as the "canonical model" for explaining and predicting the impact of technology on labour.

2.2.2 Routine-Biased Technological Change and Task-Based Models

Starting in the early 2000s, labour markets in many Western economies experienced pronounced job polarization, characterized by employment growth at the top and bottom of the wage distribution alongside a decline in middle-income occupations — a pattern that could not be fully accounted for by SBTC alone (Autor and Dorn, 2013; Goos and Manning, 2007; Goos et al., 2009).

To explain this polarization on the labour market, labour economists have introduced a complementary hypothesis: routine-biased technological change (RBTC). RBTC shifts the focus from workers' skill levels in jobs to the task content of jobs and rests on two central assumptions: first, that technologies automate specific tasks rather than entire occupations; and second, digital technologies are well-suited to replace repetitive, routine tasks, but not non-routine tasks (Acemoglu and Restrepo, 2018; Autor and Dorn, 2013). Under this framework, computerization reduces the employment shares and demand for routine-intensive middle-wage occupations while increasing demand for non-routine tasks, offering a compelling explanation for observed patterns of job polarization in Western Europe between 1993 and 2010 (Autor et al., 2003; Goos et al., 2014).

However, historical precedent suggests caution in assuming that existing frameworks will adequately capture the implications of technological change. During the First Industrial Revolution, skilled artisans were displaced by machines operated by lower-skilled workers. As technology replaced the tasks of the skilled workers, it expanded the scope of work performed by the unskilled workers, leading to wage declines and social unrest.[4] This episode illustrates that the endogenous response of technological advancements to labour market conditions can produce patterns that canonical models struggle to explain. If AI affects job tasks differently than the technologies that shaped earlier industrial revolutions, existing models may prove inadequate for understanding the labour market consequences for affected workers.

Building on Acemoglu and Restrepo (2018) as well as on Acemoglu and Autor (2011) and Autor et al. (2003), Acemoglu and Restrepo (2019) provide a task-based framework for analyzing these patterns through three mechanisms by which technologies affect labour demand: the *displacement effect*, in which capital substitutes for labour in automated tasks and reduces the labour share; the *productivity effect*, in which automation increases demand for labour in non-automated tasks; and the *reinstatement effect*, in which the creation of new

[4]This episode is often associated with the Luddite movement of the early 19th century, in which skilled textile workers destroyed machinery, fearing that automation would replace their expert positions with low-wage labourers. See Acemoglu (2002) for a more elaborate discussion.

tasks and occupations increases the labour share. Net employment and wage outcomes therefore depend on how these countervailing mechanisms interact and whether they generate opportunities for further productivity gains.

The rapid and broad evolution of AI capabilities complicates predictions about how these mechanisms will operate in practice (Dell'Acqua et al., 2023). Estimates of automation risk vary widely across studies[5], reflecting uncertainty about the extent to which AI can perform non-routine, knowledge-intensive tasks. If such capabilities continue to advance, AI may affect a broader and more diverse set of workers than earlier technologies (Felten et al., 2021; Albanesi et al., 2025). These dynamics may further amplify wage differentials across occupations with different task requirements and increase labour-market returns to interpersonal, social, and communication skills (Deming, 2017; Brynjolfsson and Mitchell, 2017). There is, however, still limited understanding of which workers will be the winners and losers from the introduction of AI. Taken together, these considerations highlight the need for empirical evidence that captures how AI exposure varies across occupations and workers, and how this translates into employment and wage outcomes. The next section describes the data and estimation strategy used to uncover these relations at the individual worker level.

2.3 Data and Descriptive Statistics

To perform our analyses we combine various sources of data on workers' labour market information, job tasks and automation susceptibility. This section describes our data and main variables of interest. The first subsection describes the administrative worker data that contains detailed records on individual employment biographies. The second subsection explains the

[5]Various models of technological change and occupational structures have produced different predictions about the number and type of jobs at risk. Frey and Osborne (2013, 2017b) catalyzed the automation debate with their prediction that 47 percent of U.S. jobs would be susceptible to automation. Others reach to more modest conclusions: Arntz et al. (2017) and Borland and Coelli (2017) estimate that only 9 percent of jobs in the United States and Australia, respectively, face automation risk. In a comparative study of 32 OECD countries, Nedelkoska and Quintini (2018) find that while almost half of jobs are susceptible to automation, only 14 percent would face a 'high risk' (i.e., $\sim$ 70% probability).

technology exposure data by Webb (2020) that we rely on to measure workers' exposure to various automation technologies. The third subsection provides the first descriptive evidence.

2.3.1 Administrative Worker-Level Data

To get a comprehensive understanding on worker level outcomes, we analyze work biographies of workers in Germany between 2012-2019.[6] These register data come from a representative 25 percent Sample of Integrated Employment Biographies (IEB), provided by the German Institute for Employment Research (IAB). The IEB covers the near-universe labour market history of individuals subject to social security contributions in Germany in all firms, occupations and economic sectors. The data includes all employees (both regular and irregular [7]), benefit recipients, unemployed job seekers, and participants in active labour market policy programs with the exception of civil servants and self-employed workers (Genz et al., 2021; Bachmann et al., 2019). This high-quality register data contains comprehensive information on workers' socio-demographic characteristics such as their wage, employment and working time status, education levels, age, gender, residency and workplace. The data also contains detailed industry and occupational classifications, including the level of skill requirement to perform workers' respective occupation that is coded at the 5th digit of the German Classification of Occupations 2010 (Kldb 2010). Figure A2.1 provides a graphical overview of the data structure.

Sample construction

The register data offers a representative sample on individual labour market

[6] We limit our analysis to the period 2012-2019 to avoid any misallocation due to occupational code switch in 2010 in Germany and more importantly, to capture the post-2012 upsurge of AI patenting activities (see Acemoglu et al., 2022). This also ensures that the pandemic-related changes in the labour market from 2020 onwards are excluded from the analysis. As market diffusion of AI in Germany started around this period, any effects we report should be interpreted as lower bound effects.

[7] Irregular employment includes marginal employment, apprentices, and partial-retirement contract workers (Genz et al., 2021).

biographies and provides detailed information on workers' employment and unemployment episodes. In what follows, we adopt the procedures outlined in Dauth and Eppelsheimer (2020) and Antoni et al. (2016) to construct the panel dataset for our worker-level analysis.

Due to the nonlinearity in employment tracks of individuals (i.e changing jobs, working in multiple jobs or becoming unemployed) overlapping spells are splitted into episodes such that the resulting data contains parallel spells with same start and end dates. In order to remove all overlapping spells, we select the main episode for workers which is defined by the parallel spell with the longest tenure with a predefined cutoff date (i.e., June 30). This allows us to observe each individual only at each point in time with a unique worker id and date.[8]

Our main outcome variable is the change in real log daily wages. In the IEB, the workers' daily wages are imputed for those wages that exceed the upper earnings limit for statutory pension insurance and are top-coded at the social security contribution limit. This means that all wage values above the respective assessment ceiling— which varies by year and region—are replaced with the threshold value. To recover the censored wage information, we follow the imputation procedure proposed by Card et al. (2013), which is based on a series of Tobit regressions controlling for worker and regional characteristics. To ensure comparability of wages over time, we convert gross daily wages into real terms using the 2022 consumer price index (CPI) published by the German Federal Statistical Office, indexed to CPI = 100 for the year 2015. Lastly, to avoid outliers in the wage distribution, we exclude the top and bottom 1% from our analysis.

We complement our data with establishment-specific wage premia, commonly referred to as Abowd–Kramarz–Margolis (AKM) fixed effects, available for the period 2010–2017 at the establishment level.[9] The AKM effect represents establishment-level wage premia that reflect systematic differences in pay across firms after adjusting for observable and unobservable worker characteristics (Abowd et al., 1999; Card et al., 2013). A higher AKM value indicates

[8] See Dauth and Eppelsheimer (2020) and Antoni et al. (2016) for more detailed explanation of the spell splitting procedure and treatment of overlapping spells.

[9] See Bellmann et al. (2020) for more detail.

that a firm tends to pay wages above what would be expected based on its workforce composition alone. Previous research suggests that such firms are generally more productive (Abowd et al., 1999) and that workers with higher skill levels are disproportionately concentrated in these establishments (Card et al., 2013). In our analysis, we include the AKM effects to a control for non-random sorting of workers into firms.

In a final step, we extend our dataset by linking the explicit measure on the occupational task composition developed by Dengler et al., 2014 (henceforth DMP).[10] The DMP measure formalizes German occupations into five routine-intensive task categories that distinguish routine from non-routine content and separate non-routine cognitive tasks into analytical and interactive components. This design is analogous to the task-based approach of Autor et al. (2003) but is tailored to German classifications. We rely on this measure as it allows us to quantify both the skill and routine content of occupations within the same analytical framework.

2.3.2 Occupational exposure to robotics, software, and AI

One of the most common practices to capture the diffusion of technologies in the labour markets is to link information on skills, tasks, and job descriptions with patent data (Lane and Saint-Martin, 2021). Various methods have been proposed and empirically validated as a suitable way to capture the diffusion and potential impact of AI on work and workers.[11] We draw on a measure developed by Webb (2020) as it allows us to separately analyze different episodes of technological change.

Webb (2020) uses the text of job-task descriptions and the text of patents in a given technological field, to construct a measure of the exposure of tasks to

[10] DMP measure is constructed using the German expert database *BERUFENET*, a free online portal that documents for each occupation the tasks performed, tools and equipment used, working conditions, required qualifications, and relevant legal regulations (Dengler et al., 2014).

[11] See, for example, Georgieff and Hyee (2022); Grimm and Gathmann (2022); Mann and Püttmann (2023); Prytkova et al. (2024) for an overview.

automation.[12] Since patent data reveals which tasks a patented innovation enables machines to perform, Webb (2020) extracts verb–noun pairs from both text sources and quantifies the overlap between them. Depending on the frequency of overlaps - or in other words, prevalence of such pairs- he assigns a score to the task, and aggregates these task-level scores to the occupation level. With patents serving as indicators of technological progress and job descriptions as indicators of workers' tasks, the degree of overlap indicates how strongly patenting in a particular technology is directed at the tasks of a given occupation (see Figure A2.2 for an illustration of the construction of exposure scores). Occupations with a larger share of overlapping tasks are therefore classified as more exposed (Acemoglu et al., 2022).

After providing a numerical (raw) score, Webb (2020) establishes exposure percentiles to indicate each job's relative exposure compared to the average exposure of all occupations. As a result, the higher the exposure percentile in an occupation, the more likely it is that the occupation consists of capabilities described in the technology patents, and therefore the more likely to have automation potential.[13] Overall, this method links specific technologies to specific occupations and identifies which occupations can be automated by a particular technology. Webb (2020) applies this methodology to three different types of technologies: robots, software without embedded AI systems, and AI. A key advantage of this approach is its ability of direct matching between technological capabilities and occupational tasks, offering granular details of how AI might affect specific job components (del Rio-Chanona et al., 2025).

In the next step, we match these exposure scores with administrative worker-level data. Linking Webb's (2020) occupation–patent-based information to register data allows us to move beyond occupational averages and examine how AI susceptibility translates into employment and wage outcomes

[12] The data on job-tasks descriptions is based on the O*NET database of occupations and task. The patent data is based on Google Patents Public Data, provided by IFI CLAIMS Patent Services (Webb, 2020).

[13] If an occupation is assigned to a higher score, it means workers in these occupations perform tasks that AI (or robotics / software) could perform, thus they are more exposed to that certain technology. However, a higher exposure score does not necessarily mean that workers in these occupations are in danger of losing work or getting displaced. It is therefore important to note that these exposure scores do not describe the extent to which occupations have already been automated, but how these tasks have the potential to be automated (Webb, 2020).

across different worker groups. Appendix 2B provides further details on the crosswalk and merging procedure.

2.3.3 Descriptive Evidence

Table 2.1 provides summary statistics for the estimation sample, comparing the bottom quartile (<p25) and the top quartile (≥p75) of Webb's (2020) AI exposure measure. Workers in the top quartile of exposure have, on average, higher daily wages and a higher share of university degrees compared to those in the bottom quartile. The composition of gender also differs markedly: highly exposed occupations are more male dominated, while lower-exposed occupations have a greater share of female workers. While age differences are modest across quartiles, workers in more exposed occupations are slightly more concentrated in western Germany. In terms of occupational skill requirements, highly exposed occupations are disproportionately concentrated in complex and highly complex occupations, while unskilled and skilled occupations are more prevalent in the bottom quartile. In terms of task content, highly exposed occupations are characterized by a greater share of analytical non-routine tasks, but lower shares of interactive non-routine and manual non-routine tasks. These descriptive patterns suggest that exposure to AI is not evenly distributed across worker groups and occupational tasks, but is systematically related to workers' education levels, gender, skill requirements, and the type of work performed.

Occupations in all sectors of the labour market are susceptible to AI automation

In Figure 2.1, we start with showing the distribution of the AI exposure across broad industry categories in the German labour market. We observe that while AI exposure is highest for occupations in IT, agriculture and manufacturing industries, the exposure percentile is low for occupations in sales, tourism, and professional services (i.e., business organization, law, and administration). This heterogeneity likely reflects differences in the task composition of occupations across industries.

More importantly, the figure demonstrates that AI's reach extends well beyond

Table 2.1: Summary statistics along the AI exposure distribution

Variable	$<$ p25		$\geq$ p75	
	Mean	Std. Dev.	Mean	Std. Dev.
Log daily wage	4.169	0.961	4.637	0.779
Education				
No Vocational Degree	0.177	0.382	0.105	0.307
Apprenticeship Degree	0.680	0.466	0.626	0.484
University Degree	0.143	0.350	0.268	0.443
Gender				
Male	0.439	0.496	0.771	0.420
Female	0.561	0.496	0.229	0.420
Age	40.79	11.88	41.50	11.24
Workplace				
West	0.815	0.388	0.833	0.373
East	0.185	0.388	0.167	0.373
Occupational Skill Requirement				
Unskilled	0.220	0.414	0.055	0.228
Skilled	0.624	0.484	0.444	0.497
Complex	0.086	0.281	0.227	0.419
Highly Complex	0.070	0.254	0.274	0.446
Task Content				
Analytical non-routine	0.159	0.366	0.342	0.474
Interactive non-routine	0.167	0.373	0.019	0.137
Cognitive routine	0.296	0.457	0.267	0.442
Manual routine	0.106	0.308	0.241	0.428
Manual non-routine	0.271	0.444	0.131	0.338
Observations	45,164,291		15,444,722	

Notes: This table reports the summary statistics on the sample composition for the bottom and top quartiles of the Webb (2020) AI exposure measure.
Source: Webb (2020), IEB (2012–2019), Dengler et al. (2014), and author's own calculations.

traditional manufacturing sectors. Unlike earlier automation technologies, such as industrial robotics, advanced automation, and digitally controlled production systems which were primarily concentrated in manufacturing industry, AI demonstrates potential for widespread application across service industries. AI's capabilities to perform a wide-array of tasks enables deployment across diverse sectors, implying a broader reach across the economy.

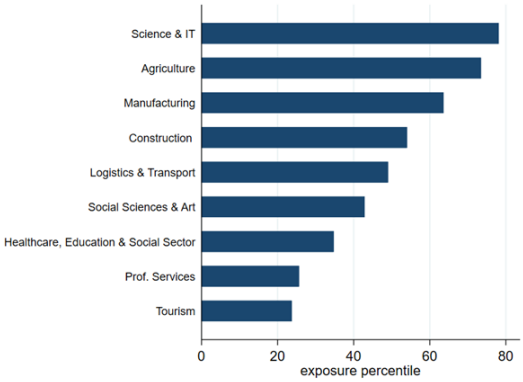


Figure 2.1: Distribution of AI exposure scores across broad industries

Notes: The exposure score percentile for robots, software and AI, averaged across all occupations and weighted by number of employees within each occupation.
Source: Webb (2020), IEB (2012–2019), author's own calculation.

Higher income occupations are more susceptible to AI automation

Figure 2.2 shows the distribution of AI, robotics, and software exposure scores across workers' wage deciles. The y-axis represents the susceptibility to exposure, and the x-axis shows the workers' wage decile distribution (with decile 1 being the lowest-paid workers and decile 10 being the highest-paid). We observe three regularities: Firs, consistent with SBTC hypothesis, susceptibility to robotization sharply decreases with workers' wages. Robot exposure is concentrated among low- and middle-income workers, reflecting the prevalence of manual and routine physical tasks in lower-wage occupations. Second, consistent with RBTC hypothesis, middle income jobs are strongly susceptible to robotization and software, reflecting the concentration of codifiable routine

cognitive tasks in clerical and administrative occupations. Third, and most notably, we observe that AI susceptibility increases with wages, indicating that workers in the top income distribution face the highest AI exposure. This pattern suggests that AI, unlike earlier automation technologies, targets tasks traditionally associated with high-skill, high-wage work and confirms recent empirical observations on AI (Acemoglu et al., 2022; Felten et al., 2019).

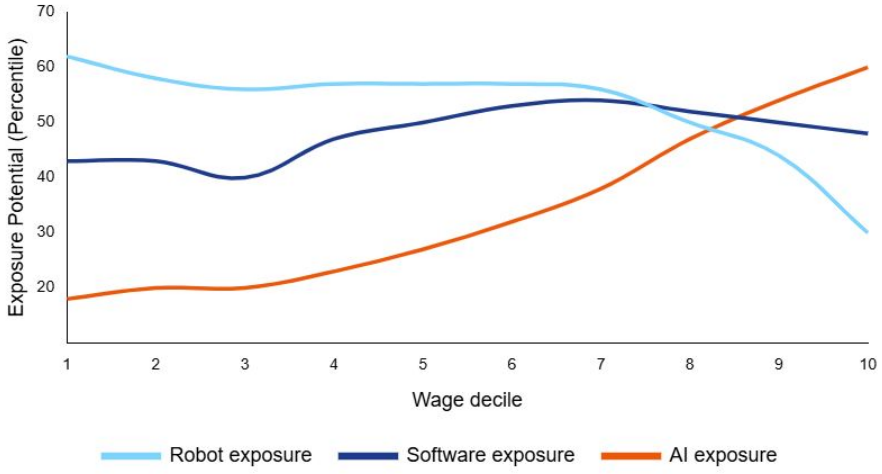


Figure 2.2: Distribution of exposure scores by wage deciles

Notes: The exposure score percentile for robots, software and AI, averaged across all occupations and weighted by number of employees within each occupation.
Source: Webb (2020), IEB (2012–2019), author's own calculation.

Complex and non-routine work is susceptible to AI automation

We now explore the extent to which jobs with different skill levels are susceptible to various forms of technology adoption. In Figure 2.3, we show how occupations' susceptibility to robotization, software, and AI relate to the task complexity of occupations. Task complexity in each occupation is defined by the skill level required to perform the occupation. Based on this, we distinguish workers' occupations into four categories which are defined as follows: (i) jobs with unskilled or semi-skilled tasks which require no formal qualification or only short-term training, such as e.g., helper/assistants, (ii) jobs with

skilled tasks which require a formal vocational education training of at least 2 years such as e.g., professionals, (iii) jobs with complex tasks which require a university degree or master craftsman's certificate such as e.g., specialists, and (iv) jobs with highly complex tasks which require a university degree or similar and, beyond that, profound professional experience or further formal highly specialized qualification certificates like a doctorate (Peede and Stops, 2024).

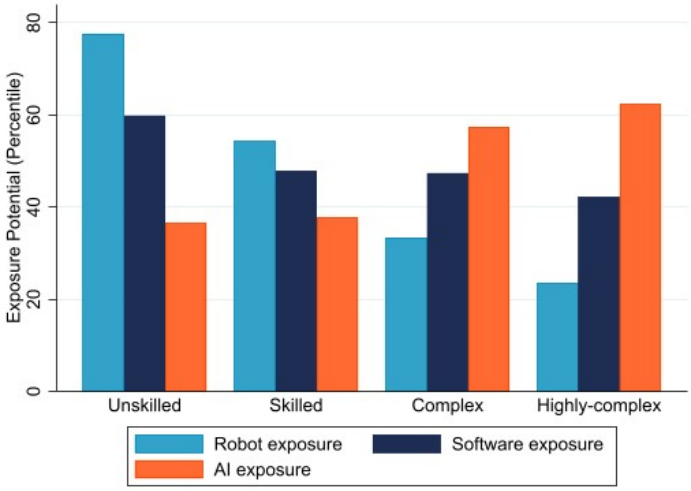


Figure 2.3: Distribution of exposure scores by required skill level of occupations

Notes: The exposure score percentile for robots, software and AI, averaged across all occupations and weighted by number of employees within each occupation.
Source: Webb (2020), IEB (2012–2019), author's own calculation.

From Figure 2.3 we observe that the more complex job tasks are, and thus the more highly skilled workers must be to perform these tasks, the less susceptible they are to robotization and software. However the more complex job tasks are, the more susceptible they are to AI automation. More highly complex job tasks require higher skill levels of workers. This is in line with findings from comparable analyses on US data and may explain why AI susceptibility appears larger in jobs with higher incomes (Acemoglu et al., 2022).

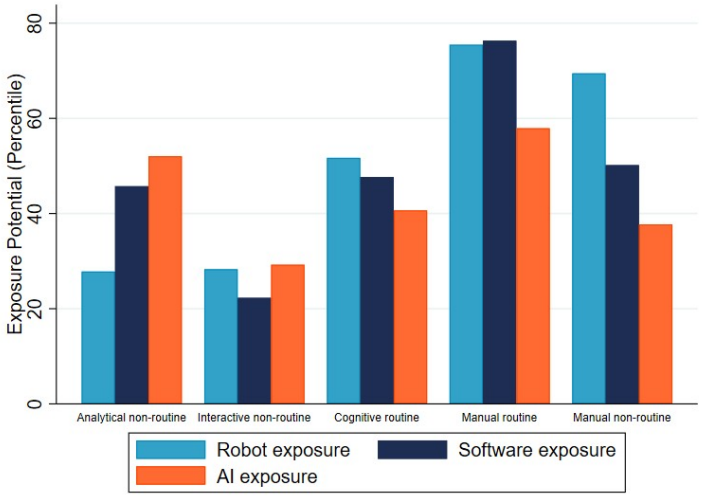


Figure 2.4: Distribution of exposure scores by task content of occupations

Notes: The exposure score percentile for robots, software and AI, averaged across all occupations and weighted by number of employees within each occupation.
Source: Webb (2020), IEB (2012–2019), Dengler et al. (2014) and author's own calculation.

Next, we examine how the level of routineness of job tasks is related to their susceptibility to technology.[14] When plotted against occupational routine intensity, Figure 2.4 reveals distinct exposure patterns. Manual-intensive jobs show the highest robot exposure, with manual routine and non-routine occupations demonstrating similar susceptibility. Software exposure presents a more nuanced picture: while interactive non-routine jobs show relatively low exposure, analytical non-routine tasks exhibit the highest software susceptibility, exceeding even cognitive routine work. This suggests software automation extends beyond routine tasks into complex analytical domains. AI exposure diverges markedly from these patterns, with analytical non-routine tasks having relatively high exposure while manual occupations show substantially lower susceptibility. Notably, interactive non-routine jobs demonstrate the lowest AI exposure across all task categories, suggesting that social skills and interpersonal capabilities may serve as critical bottlenecks to AI automation. This indicates AI's comparative advantage lies in cognitive rather than physical work, potentially threatening high-skill analytical occupations more than traditional manual labour.

2.4 AI Exposure and Labour Market Outcomes

Above evidence demonstrates that different technologies target different task types, resulting in heterogeneous exposure patterns across the occupational structure. More specifically, unlike earlier automation waves, high-skilled, high-educated workers and workers in occupations that require complex and highly complex, non-routine tasks are most exposed to AI, while their exposure to robotics and software is substantially lower.

[14] Based on Dengler et al. (2014), we distinguish between manual and non-manual jobs, and five levels of routineness: (i) analytical non-routine tasks, (ii) interactive non-routine tasks, (iii) cognitive routine tasks, (iv) manual routine tasks, and (v) manual non-routine tasks. Analytical non-routine tasks include for example management, planning, supervision, leadership, direction, controlling, monitoring, and analyzing. Interactive non-routine tasks include commercial and counseling activities such as marketing and advertising. Cognitive routine tasks include administrating, testing, inspecting, measuring, monitoring, and running diagnostics. Manual non-routine tasks include refurbishing, providing services, manual therapy, and bespoke crafts. Manual routine tasks include farming, construction tasks, manufacturing, production, harvesting, operating machines, setting up machines, and typesetting (Dengler et al., 2014).

This pattern raises a critical question: do these differences translate into changes in labour market outcomes? A central question in understanding AI's labour market impact is whether it operates as a complement or substitute to human labour—does AI augment workers' skills and enhance productivity, or does it substitute for workers and displace them from their occupations? In this section, we investigate how AI, robotics, and software exposure relate to workers' employment displacement and wage growth, with particular attention to whether AI's distinct exposure pattern produces different labour market implications than earlier automation technologies.

2.4.1 Empirical Strategy

To study the relation between technology exposure and labour market outcomes, we employ a series of ordinary least squares (OLS) regression specifications. Our empirical approach exploits variation in AI, robotics, and software exposure across occupations to identify their associations with workers' wage growth and employment transitions between 2012 and 2019. Our long-difference specification captures the cumulative effects of technology exposure rather than year-on-year change. This approach is particularly suitable for studying automation effects as firms gradually adopt new technologies and workers adjust to changing labour market conditions.[15]

Employment Transitions

We begin by examining whether technology exposure is related to risk of displacement at the worker level. Our baseline specification estimates the following equation:

$$Pr(Displacement)_i = \alpha + \beta Exposure_o + \mathbf{X}_i'\boldsymbol{\gamma} + \mathbf{Z}_f'\boldsymbol{\delta} + \theta_r + \mu_s + \lambda_g + \psi_{o2} + \varepsilon_i \tag{2.1}$$

[15] As a robustness check, we conduct our estimations for shorter time windows ranging from 2012-2013 *to* 2012-2019. Our main results remain unchanged across these alternative specifications (see Appendix Table A2.1)

where our main outcome of interest is a binary indicator equal to one if individual i experienced employment displacement, defined as transitioning from regular employment in 2012 to either marginal employment or unemployment in 2019, and zero otherwise. The key explanatory variable, $Exposure_o$, is the standardized technology exposure measure (AI, robotics, or software) in 2012 for the 5-digit occupation o of individual worker i. The parameter of interest β, measures the change in the probability of displacement associated with a one-standard-deviation increase in the occupational exposure measure.

The interpretation of our parameter of interest is threatened by a number of potentially confounding factors. First, occupational exposure may be correlated with individual characteristics that independently affect displacement risk. Second, establishment-level factors such as firm quality or size may drive both technology adoption and employment outcomes. Third, unobserved regional, industry, or occupational characteristics may confound the relation between exposure and displacement. To address these concerns, we include a comprehensive set of controls: individual-level controls $\mathbf{X}_i$ comprise baseline wages in 2012, education (no vocational degree, apprenticeship degree, or university degree), gender, age and age squared, part-time employment status, and exposure to the other two technologies (e.g., when estimating AI exposure, we control for robotics and software exposure scores); establishment-level controls, $\mathbf{Z}_f$, include firm size categories and AKM fixed effects to account for time-invariant establishment quality. We further control for unobserved heterogeneity through region fixed effects θ_r; 3-digit industry fixed effects μ_s; 2-digit occupation fixed effects ψ_{o2} and a dummy λ_g indicating whether the worker's occupation belongs to sales or administration, following Acemoglu et al. (2022). Lastly, standard errors are clustered at the 5-digit occupation level to account for within-occupation correlation.

Wage Growth

We next turn to wage dynamics to further distinguish between augmentation and substitution effects. Complementing existing research on job displacement, analyzing wage growth in response to exposure scores provides complementary evidence along a different adjustment margin: while employment displacement captures the extensive margin—whether workers exit

their occupations—wage changes capture the intensive margin, revealing how technology affects worker productivity among those who remain employed. Although the relation between individual productivity and wages is not uncontroversial, the positive association between wages and productivity is strong and well-established. Accordingly, we assume that wages are (at least partly) an expression of workers' productivity.[16] We further assume a positive relation between AI susceptibility and actual technology adoption in practice. Empirical evidence supports this assumption: Acemoglu et al. (2022) document a positive correlation between AI skills demand in job advertisements and AI exposure indices in the US, and Peede and Stops (2024) corroborate this pattern for Germany. Under this framework, if AI augments workers and enhances their productivity, we should observe faster wage growth in occupations with higher AI exposure. Conversely, if substitution effects dominate, we would expect downward pressure on wage growth in highly exposed occupations.

We estimate the following specification:

$$\Delta \ln(wage)_i = \alpha + \beta \cdot Exposure_o + \mathbf{X}_i'\boldsymbol{\gamma} + \mathbf{Z}_f'\boldsymbol{\delta} + \theta_r + \mu_s + \lambda_g + \psi_{o2} + \varepsilon_i \quad (2.2)$$

where $\Delta \ln(wage)_i$ denotes the change in log daily wages for individual i between 2012 and 2019. The coefficient β captures the change in wage growth associated with a one-standard-deviation increase in technology exposure. The control variables and fixed effects remain identical to those in Equation 2.1.

Heterogeneity by Required Skill Level

We extend Equation 2.2 to account for potential heterogeneity in exposure scores across different skill levels. To investigate whether wage outcomes vary by occupational skill requirements, we include skill-level indicators and their interaction effects:

[16] There is debate about the extent to which wages are a valid and reliable proxy of worker productivity, particularly in countries with strong institutional wage setting. Our assumption here is not that wages equal productivity, but that positive shocks to productivity are reflected, at least partially, in wage growth over a seven-year horizon.

$$\Delta \ln(wage)_i = \alpha + \sum_l \beta_l \cdot (Exposure_o \times Skill_l) + \sum_l \phi_l \cdot Skill_l + \mathbf{X}_i'\boldsymbol{\gamma} + \mathbf{Z}_f'\boldsymbol{\delta} + \theta_r + \mu_s + \lambda_g + \psi_{o2} + \varepsilon_i \tag{2.3}$$

where $Skill_l$ denotes the required skill level of individual i's occupation, classified into four categories: unskilled, skilled, complex and highly complex jobs. The interaction term $Exposure_o \times Skill_l$ allows the effect of technology exposure to vary across skill levels, with coefficients β_l capturing skill-specific technology effects. We use unskilled occupations as the reference category, such that β_l represents the differential effect of technology exposure for skill level l relative to unskilled workers. This specification allows us to test whether AI disproportionately affects high-skilled workers, as suggested by our descriptive evidence.

Heterogeneity by Task Content

We further examine heterogeneity across task types by estimating a specification analogous to Equation 2.3, replacing skill levels with task content categories:

$$\Delta \ln(wage)_i = \alpha + \sum_t \beta_t \cdot (Exposure_o \times Task_t) + \sum_t \phi_t \cdot Task_t + \mathbf{X}_i'\boldsymbol{\gamma} + \mathbf{Z}_f'\boldsymbol{\delta} + \theta_r + \mu_s + \lambda_g + \psi_{o2} + \varepsilon_i \tag{2.4}$$

where $Task_t$ denotes the primary task type of individual i's occupation: manual routine, cognitive routine, manual non-routine, analytical non-routine, or interactive non-routine. The coefficients β_t identify whether technology effects differ across occupations with distinct task profiles. This specification allows us to test, for example, whether AI disproportionately affects analytically intensive work relative to interactive or manual tasks, and whether social skills provide protection against automation as our descriptive analysis suggests. We use cognitive routine tasks as the reference category, such that β_t captures

the differential effect of exposure for task type t relative to cognitive routine occupations.

2.5 Results

Based on the empirical strategy outlined above, we begin by reporting the main regression results on the relation between technology exposure and employment displacement, examining whether workers in AI exposed occupations face heightened displacement risk. We then analyze how AI exposure relates to wage growth and whether these relations differ depending on occupational skill complexity and task content. Finally, we present a series of robustness checks to assess the validity of our results.

2.5.1 Employment in AI-Exposed Occupations

One potential mechanism through which AI may affect labour markets is by displacing workers whose tasks become automated. If AI serves primarily as a substitute for human labour, specific job tasks may be fully automated, and workers unable to adapt may exit their occupations, either by transitioning to different roles or by becoming unemployed.

Table 2.2 presents estimates from linear probability models examining the association between AI exposure and employment transitions. Our dependent variable is a binary indicator equal to one if a worker transitioned from regular full-time employment in 2012 to either marginal employment (i.e, low-paying jobs), unemployment, or becoming inactive in 2019, and zero otherwise. Column 1 shows the most parsimonious specification controlling only for workers' baseline wages in 2012. We then sequentially add control variables as described in section 2.4.1. Column 2 includes individual worker characteristics (education, gender, age, part-time status) and firm-level controls (firm size). Column 3 additionally controls for robotics and software exposure and AKM effects. Column 4 adds region and industry fixed effects as well as a dummy variable indicating whether the worker's occupation belongs

to sales or administration. Column 5 presents our full specification, including 2-digit occupation fixed effects.

The results indicate no systematic relation between AI exposure and displacement. In the most parsimonious specification (Column 1), a one-standard-deviation increase in occupational AI exposure is associated with a 0.8 percentage points increase in the probability of transitioning out of employment. However, this association disappears once we add worker and firm controls: the coefficient drops to 0.5 percentage points in Column 2 and becomes statistically insignificant. From Column 3 onwards, the coefficient changes sign, ranging from -0.6 percentage points to -0.3 and remains statistically insignificant across all specifications.

These findings are consistent with a growing body of evidence documenting null employment effects of AI exposure (Acemoglu et al., 2022; Georgieff and Hyee, 2022; Peede and Stops, 2024; Storm et al., 2025). While the initial positive association in Column 1 suggests that workers in AI-exposed occupations face slightly higher displacement risk, this relation is substantially attenuated once we account for worker characteristics, firm quality, and occupational structure. Taken together, the estimates suggest that, during the early stages of AI diffusion studied here, AI exposure does not systematically increase the probability of adverse employment transitions at the worker level.

Table 2.2: OLS regression of AI susceptibility on the transition out of employment, 2012–2019

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.008***	0.005	-0.006	-0.001	-0.003
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Software Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	6,065,738	6,065,640	5,798,898	5,798,898	5,798,898
Adjusted R^2	0.007	0.20	0.20	0.21	0.21

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and employment transitions between 2012–2019. The dependent variable is a binary indicator equal to one if a worker transitioned from regular full-time employment in 2012 into marginal employment, unemployment, or inactivity by 2019, and zero otherwise. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Column 1 contains only occupational AI exposure controlling for baseline wages in 2012. Column 2 includes worker and firm controls (education, gender, age, age squared, part-time work, and firm size). Column 3 additionally includes AKM effects and robot and software exposure measures. Column 4 includes region and industry (3-digit) fixed effects and a dummy variable indicating whether the worker's occupation belongs to sales or administration. Column 5 includes 2-digit occupation fixed effects. Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security, & Webb (2020).

2.5.2 Wages in AI-Exposed Occupations

Having established that AI exposure is associated with null employment effects, we now turn to examining wage dynamics among workers who remained in full-time employment throughout the observation period. Analyzing wage growth provides complementary evidence on whether AI primarily substitutes for or augments human labour. If AI substitutes for workers and reduces labour demand, we would expect downward pressure on wages in highly exposed occupations. Conversely, if AI augments workers by enhancing their skills and productivity, we would expect faster wage growth in occupations with higher AI exposure.

Table 2.3 presents OLS estimates of the relationship between AI exposure in 2012 and wage growth of German workers between 2012 and 2019. Column 1 shows the most parsimonious specification that controls only for workers' baseline wages in 2012. We include baseline wages to account for the mechanical relation whereby workers with higher initial wages tend to experience lower wage growth. Moreover, the baseline wage absorbs potential unobserved and time-constant heterogeneity that might correlate with workers' wage growth. The main coefficient of interest shows that a one-standard-deviation increase in AI exposure is associated with a 2.3 percentage point increase in workers' wage growth between 2012 and 2019. This finding is very similar to previous results from the United States.

Column 2 adds observable worker and firm characteristics and Column 3 includes AKM effects of the workers' baseline firm as a proxy for employer quality, as well as controls for robotics and software exposure. Column 4 adds region and industry (3-digit) fixed effects and a dummy variable indicating whether the worker's occupation is a sales or administrative occupation, to account for secular trends in the decline of employment shares in these occupations (Goldin and Katz, 2008; Acemoglu et al., 2022). Column 5 presents the full specification, adding a full set of baseline occupation dummies at the two-digit level.

The coefficient estimates remain robust across all specifications. After including worker and firm controls in Column 2, the coefficient declines to 1.1

percentage points but remains statistically significant. However, once we control for robotics and software exposure in Column 3, the coefficient increases to 1.8 percentage points, and in the most saturated specifications (Columns 4 and 5), which include region, industry, and occupation fixed effects, it reaches to approximately 2.2-2.3 percentage points.

These findings suggest that workers in AI-exposed occupations experienced faster wage growth and that this relationship remains robust even after accounting for observable characteristics that typically have substantial impacts on wage growth. Moreover, the fact that the coefficient remains positive and significant even when controlling for robotics and software exposure indicates that AI's relationship with wage growth is distinct from that of earlier automation technologies. This pattern makes alternative interpretations unlikely to hold. For example, the hypothesis that AI simply affects wage growth more in higher-paying occupations is not supported by our results, as the coefficient remains stable even after controlling for baseline wages, firm quality, and detailed occupation fixed effects.

These findings are consistent with the hypothesis that AI may serve as a productivity-enhancing technology that complements human labour. However, this pattern may conceal important heterogeneity across worker groups. To investigate this possibility, we next examine how the relationship between AI exposure and wage growth varies across occupational skill levels and task content.

Table 2.3: OLS regression of AI susceptibility on wage growth, 2012–2019

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.023***	0.011**	0.018***	0.022***	0.023***
	(0.005)	(0.005)	(0.007)	(0.005)	(0.003)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Software Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	3,358,470	3,358,428	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.070	0.106	0.111	0.126	0.129

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Column 1 contains only occupational AI exposure controlling for baseline wages in 2012. Column 2 includes worker and firm controls (education, gender, age, age squared, part-time work, and firm size). Column 3 additionally includes AKM effects and robot and software exposure measures. Column 4 includes region and industry (3-digit) fixed effects and a dummy variable indicating whether the worker's occupation belongs to sales or administration. Column 5 includes 2-digit occupation fixed effects. Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security, & Webb (2020).

The wage premium of high skills is unrelated to AI

Our findings thus far suggest that AI exposure is positively associated with wage growth. If this relation reflects productivity-enhancing augmentation effects, we might expect the SBTC hypothesis to hold: the more highly skilled workers are, the more likely it is that technologies will complement their skills and increase their productivity, which should be associated with higher wage growth. Under this framework, we would expect the positive relationship between AI exposure and wage growth to be strongest in occupations requiring high skill levels.

To explore whether such mechanisms operate, Table 2.4 presents results from models that extend our baseline specification by including interaction terms between AI exposure and four skill categories: unskilled (reference category), skilled, complex, and highly complex occupations. The table follows the same structure as Table 2.3, with Column 1 presenting the most parsimonious model and Column 5 the most saturated specification.

Rows two through five show the main effects of the skill categories. Since we standardize our AI exposure measure, these coefficients represent the wage growth premium for workers in occupations with different skill requirements at the average level of AI susceptibility. The estimates reveal substantial wage growth premiums associated with higher skill requirements: relative to unskilled occupations, workers in skilled occupations experience approximately 3 percentage points higher wage growth, workers in complex occupations experience approximately 7 percentage points higher wage growth, and workers in highly complex occupations experience approximately 10 percentage points higher wage growth.

Rows six through eight present the interaction terms between AI exposure and skill categories. These estimates indicate how the relationship between AI exposure and wage growth differs across skill levels relative to unskilled occupations. Across all specifications, the interaction terms are negative, indicating that the positive association between AI exposure and wage growth is actually weaker—not stronger—in higher-skilled occupations. The interaction for skilled occupations is consistently statistically significant and ranges from -3.4 to -1.4 percentage points across specifications. The interactions for

complex and highly complex occupations are smaller in magnitude and less consistently significant, though they follow the same pattern. The coefficient estimates remain robust across specifications that control for robotics and software exposure, worker and firm characteristics, and various fixed effects.

The findings indicate that AI susceptibility is positively related to wage growth, regardless of the required skill levels of jobs; it does not follow the skill-biased pattern that characterized earlier technological changes.

Table 2.4: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations requiring different skill levels

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.047**	0.030*	0.048***	0.034***	0.016*
	(0.016)	(0.017)	(0.017)	(0.011)	(0.009)
Unskilled Occupation	ref.	ref.	ref.	ref.	ref.
Skilled Occupation	0.044***	0.053***	0.032***	0.037***	0.030***
	(0.016)	(0.013)	(0.011)	(0.006)	(0.005)
Complex Occupation	0.107***	0.117***	0.068***	0.079***	0.068***
	(0.018)	(0.014)	(0.013)	(0.008)	(0.007)
Highly Complex Occupation	0.165***	0.146***	0.085***	0.103***	0.100***
	(0.035)	(0.021)	(0.024)	(0.017)	(0.014)
AI × Skilled	-0.034**	-0.036**	-0.041***	-0.024***	-0.014**
	(0.017)	(0.015)	(0.014)	(0.009)	(0.006)
AI × Complex	-0.034*	-0.029*	-0.036**	-0.018*	-0.006
	(0.018)	(0.016)	(0.015)	(0.009)	(0.008)
AI × Highly Complex	-0.049	-0.032	-0.037*	-0.024*	-0.010
	(0.033)	(0.021)	(0.020)	(0.014)	(0.011)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Task Routineness	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Software Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	3,358,470	3,358,428	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.078	0.113	0.114	0.129	0.131

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Occupational skill complexity is classified into four categories (unskilled, skilled, complex, highly complex) for the 5-digit KldB (2010) occupations. Column 1 contains only occupational AI exposure controlling for baseline wages in 2012. Column 2 includes worker and firm controls (education, gender, age, age squared, part-time work, and firm size) and task routineness. Column 3 additionally includes AKM effects and robot and software exposure measures. Column 4 includes region and industry (3-digit) fixed effects and a dummy variable indicating whether the worker's occupation belongs to sales or administration. Column 5 includes 2-digit occupation fixed effects. Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security, & Webb (2020).

AI-related wage growth is strongest in jobs with many routine tasks

Turning to task structure of occupations, RBTC hypothesis assumes that workers in jobs with more abstract tasks (i.e., higher share of non-routine tasks) should experience a notable increase in productivity from technology adoption, because information technology complements them, allowing them to specialize in non-routine tasks (such as flexibility and situational judgment). Following this logic, workers in non-routine task-intensive occupations should benefit from stronger complementarities between routine and non-routine tasks, elastic demand for their services, and inelastic labour supply over the short and medium term (Gathmann et al., 2024). Therefore, under the RBTC hypothesis, the more non-routine and non-manual tasks an occupation involves, the more likely it is that higher AI susceptibility would be associated with stronger wage growth.

Table 2.5 presents results from specifications that interact AI exposure with variables that measure routineness of job tasks. The table follows the same structure as our previous tables, with Column 1 presenting the most parsimonious specification and Column 5 the most saturated model.

The main effects of task categories reveal substantial wage differentials across task types at average levels of AI exposure. We observe that relative to cognitive routine occupations, workers that predominantly perform analytical non-routine and interactive non-routine tasks experience higher wage growth (approximately 1.1 percentage points), while workers in manual routine and manual non-routine occupations experience lower wage growth (approximately 0.7 and 1.1 percentage points, respectively).

The interaction terms demonstrate how the relationship between AI exposure and wage growth varies across task types. Contrary to RBTC predictions, the interaction for analytical non-routine tasks is negative and statistically significant across most specifications (ranging from -0.028 to -0.012). This indicates that workers in analytical non-routine occupations experience a weaker positive association between AI exposure and wage growth compared to those in cognitive routine occupations. Although the net effect for analytical non-routine workers in AI-exposed occupations remains positive (approx 0.027 - 0.012 = 0.015), this is smaller compared to workers in cognitive routine

tasks.

Similarly, the interaction for interactive non-routine tasks is negative in most specifications, though less consistently significant. The interaction for manual routine tasks is also negative, while manual non-routine shows a similar pattern. These findings suggest that AI exposure is positively associated with wage growth across all task types, but this association is strongest for workers in cognitive routine occupations, precisely the group that RBTC hypothesis predicts should be most vulnerable to automation.

This observed pattern challenges RBTC's predictions about task-based technological change. Rather than complementing non-routine tasks most strongly, AI appears to have its largest positive wage associations in routine-intensive occupations. One interpretation is that AI may increase productivity in routine cognitive tasks by augmenting rather than simply replacing them, or that workers in these occupations successfully adapt by shifting toward tasks that complement AI. Alternatively, selection effects may be at play: if AI displaces the lowest-performing workers in routine occupations (as suggested by our modest employment displacement findings), the remaining workers may experience wage gains as they take on expanded responsibilities.

Table 2.5: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations with different levels of task routineness

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.031***	0.035***	0.043***	0.030***	0.027**
	(0.006)	(0.013)	(0.013)	(0.011)	(0.011)
Cognitive Routine	ref.	ref.	ref.	ref.	ref.
Analytical non-routine	0.031***	0.022***	0.016**	0.009**	0.011**
	(0.009)	(0.007)	(0.007)	(0.004)	(0.005)
AI × Analytical non-routine	-0.028***	-0.022***	-0.019***	-0.010**	-0.012***
	(0.008)	(0.006)	(0.007)	(0.004)	(0.004)
Interactive non-routine	-0.001	0.001	-0.006	-0.000	0.007
	(0.005)	(0.006)	(0.007)	(0.003)	(0.005)
AI × Interactive non-routine	-0.013*	-0.002	-0.001	-0.003	-0.013***
	(0.007)	(0.006)	(0.006)	(0.003)	(0.004)
Manual routine	-0.013**	-0.026***	-0.024***	-0.021***	-0.007*
	(0.005)	(0.005)	(0.005)	(0.003)	(0.004)
AI × Manual routine	-0.014**	-0.004	-0.003	-0.001	-0.006**
	(0.006)	(0.005)	(0.005)	(0.002)	(0.002)
Manual non-routine	-0.021***	-0.023***	-0.023***	-0.024***	-0.011
	(0.007)	(0.006)	(0.007)	(0.005)	(0.007)
AI × Manual non-routine	-0.018**	-0.009	-0.007	-0.003	-0.007*
	(0.008)	(0.006)	(0.007)	(0.004)	(0.004)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Skill Complexity	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Software Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	3,358,470	3,358,428	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.082	0.116	0.118	0.132	0.130

Notes: This table presents OLS estimates of the effects of AI exposure on the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Occupational task routine measures are based on Dengler et al. (2014). Column 1 contains only occupational AI exposure controlling for baseline wages in 2012. Column 2 includes worker and firm controls (education, gender, age, age squared, part-time work, and firm size) and occupational skill complexity. Column 3 additionally includes AKM effects and robot and software exposure measures. Column 4 includes region and industry (3-digit) fixed effects and a dummy variable indicating whether the worker's occupation belongs to sales or administration. Column 5 includes 2-digit occupation fixed effects. Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security, Webb (2020) & Dengler et al. (2014).

2.5.3 Robustness Checks

To assess the robustness of our findings, we conduct a series of sensitivity analyses, including analyses exploring parameter sensitivity to model specification, estimation methods, and sampling.

Choice of time horizon Our main specifications employ long differences approach, analyzing wage growth and employment transitions over the 2012-2019 period. This approach has the advantage of capturing the cumulative effects of technological change, accounting for gradual technology diffusion across firms and worker adjustment to changing labour market conditions. However, to verify that our findings are driven by this methodological choice, we examine whether the patterns we observe hold across shorter time windows. Appendix tables A2.1 through A2.4 present results from specifications that vary the observation window from one year (2012-2013) to seven years (2012-2019), while maintaining our most saturated specification with the full set of controls and fixed effects.

For employment results (Table A2.2), the estimates remain small, negative, and insignificant across specifications, suggesting that displacement effects do not intensify over time. In contrast, for wage growth (Table A2.1), the positive association between AI exposure and wages strengthens from 1.0 percentage points in the one-year specification to 2.3 percentage points over the full seven-year period, indicating that effects accumulate over time.

The heterogeneity analyses by skill level (Table A2.3) and task content (Table A2.4) also show patterns consistent with our main results across all specifications. The key findings—negative interactions between AI exposure and skill levels, and stronger positive wage associations in cognitive routine occupations—persist across time windows, with effect magnitudes generally increasing over longer periods.

Robotics & Software exposure To verify that our findings are specific to AI rather than reflecting general patterns of automation, we replicate our main wage growth analysis using robotics and software exposure measures. Tables A2.5 and A2.6 in the Appendix present results using the same specification as our main models but replacing AI exposure with robotics and software

exposure, respectively.

The results reveal distinct patterns across technologies. For robotics exposure (Table A2.5), we observe a negative association with wage growth that strengthens over time, reaching -4.3 percentage points in the 2012-2019 specification. For software exposure (Table A2.6), the association is also negative but smaller in magnitude, reaching -1.9 percentage points over the full period. These findings support our interpretation that AI's labour market effects differ from those of earlier automation technologies. While robotics and software exposure are associated with wage penalties, AI exposure is associated with wage premiums, suggesting augmentation effects might occur for workers.

Education versus occupational skill requirements Our main analysis examines heterogeneity by occupational skill requirements, which capture the complexity of jobs. As a robustness check, we investigate whether results differ when using individual educational attainment instead. Specifically, we distinguish between low education (no formal vocational qualification), medium education (completed formal vocational education), and high education (tertiary or university degree). Table A2.7 replicates our skill heterogeneity analysis using education levels.

The results show similar patterns to our main skill-based analysis. Workers with higher education levels experience substantial wage growth premiums, but the interaction terms between AI exposure and education levels are negative or insignificant. This indicates that, as with occupational skill requirements, AI exposure does not generate stronger wage growth for more educated workers. The consistency between education-based and skill-based analyses strengthens our conclusions. Whether we measure worker capability through educational credentials or occupational skill requirements, we find that AI exposure is positively associated with wage growth across all groups, but does not disproportionately benefit higher-skilled or more educated workers.

Heterogeneity by gender To examine whether AI's effects differ by gender, we estimate our main specifications separately for male and female workers. Tables A2.8 through A2.13 present results for male workers (Tables A2.8 - A2.10) and female workers (Tables A2.11 - A2.13), covering baseline wage growth, skill and task heterogeneity analyses.

For both males and females, we observe that AI exposure is positively associated with wage growth (approximately 2.0–2.3 percentage points for males, 1.2–2.0 percentage points for females). The skill and task heterogeneity patterns observed in our main analysis hold for both workers separately: substantial skill wage premiums but negative or insignificant interactions with AI exposure, and no evidence that AI disproportionately benefits workers in non-routine occupations. While coefficients differ slightly in magnitudes between genders, the patterns remain consistent. Both groups show positive wage growth associations, strengthening our conclusion that AI's potential labour market effects are not driven by gender-specific occupational sorting.

2.6 Conclusion

In this paper we explored how AI may plausibly impact work and workers. Models based on SBTC and RBTC hypotheses predict that AI will affect a segmented part of the labour market and will do so in a more or less predictable way. Our analyses put the expectations derived from SBTC and RBTC to the test of the data.

While SBTC posits that high-skilled workers are less likely to be substituted by technological advancements and tend to benefit most from them, our findings reveal a more complex pattern. Workers in AI-exposed occupations are more susceptible to AI automation and experience positive wage growth, but this association does not strengthen with skill level; contrary to what SBTC would predict. RBTC hypothesis argues that routine tasks are most easily automated, and that workers performing non-routine tasks—such as abstract tasks found in professional and managerial occupations or tasks requiring situational adaptability and interpersonal interaction—are more likely to experience productivity gains, as these tasks are considerably more difficult to computerize (Autor et al., 2003).

Our analyses yield patterns that diverge from these predictions. While manual tasks are relatively resistant to AI automation, non-routine tasks related to management, planning, supervision, leadership, and analysis are among the most susceptible to AI exposure. Moreover, we find that workers in cognitive routine occupations experience the strongest positive wage growth associations with AI exposure, while interactive non-routine tasks show the weakest associations. This suggests that tasks requiring higher levels of interpersonal and social skills, in which labour has a comparative advantage, may indeed be more resistant to AI (Deming, 2017).

Our findings are consistent with conjectures that AI may serve as a productivity-enhancing complement to human labour by supporting expert judgment and decision-making (Autor, 2024). In this context, distinguishing between expert and non-expert tasks can provide a useful lens for capturing task complexity and understanding how AI may differentially affect job opportunities across occupations (Autor and Thompson, 2025). Our observa-

tions are also consistent with alternative frameworks that emphasize technical and organizational feasibility rather than task complexity alone as the key determinants of AI adoption (Brynjolfsson et al., 2018). Whether AI can be deployed depends not only on whether tasks are cognitive or routine, but also on factors such as availability of training data, codifiability of performance metrics, tolerance for errors, regulatory constraints, and organizational readiness. From this perspective, AI exposure reflects technical potential, not inevitability: most occupations contain some automatable tasks, but very few are fully automatable (Brynjolfsson et al., 2018).

An important limitation of our analysis is that we measure occupational exposure to AI technologies, not actual adoption or usage. Exposure captures technical feasibility, identifying which tasks could be, in principle, automated given current AI capabilities. However this does not indicate whether, how, or to what extent these technologies are implemented in actual workplaces. Furthermore, there is a substantial gap between exposure and real - world adoption and diffusion of AI: although we document widespread AI susceptibility across the occupational distribution, actual AI use in German workplaces during the period studied remained limited and highly uneven, being concentrated in a small subset of large, digitally advanced firms (Rammer et al., 2022; Storm et al., 2025).

This gap also reflects the fact that technology diffusion is not determined by technical feasibility alone. As emphasized in Autor (2015), assessments of automation that focus primarily on technological capability risk overlooking the economic, ethical, and legal boundaries that shape whether new technologies are actually deployed. Economic considerations (such as innovation-related expenses and expected returns), organizational factors (such as complementary investments, managerial capabilities, and workflow redesign), legal constraints (such as data protection) and workplace dynamics (such as worker acceptance, trust, and ethical considerations) shape whether and how firms adopt AI technologies. Understanding these adoption processes and in particular, the conditions under which AI technologies can be integrated into existing work practices remain essential for assessing the labour-market implications of AI.

The following chapter discusses these issues by examining AI adoption at the firm level, tracing how deployment unfolds gradually from initial deployment to broader implementation. In doing so, it also highlights the mechanisms through which actual AI deployment translates into changes in task allocation and labour demand.

Appendix 2A: Additional Figures & Tables

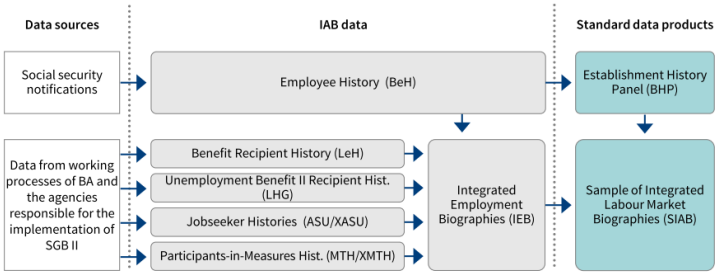


Figure A2.1: Structure of the Sample of Integrated Employment Biographies

Source: Frodermann et al., 2021

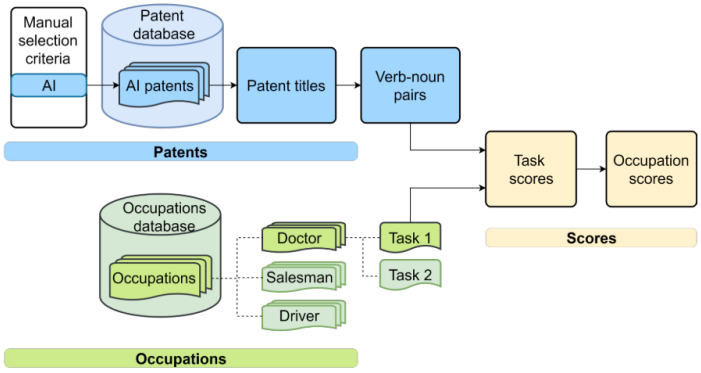


Figure A2.2: Illustration of the process for constructing technology exposure measures

Source: Webb, 2020

Table A2.1: OLS regression of AI susceptibility on wage growth, 2012 to 2019

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
AI Exposure$_{2012}$	0.010***	0.013***	0.017***	0.019***	0.021***	0.022***	0.023***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robot & Soft. Exp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.029	0.045	0.069	0.090	0.106	0.119	0.129

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.2: OLS regression of AI susceptibility on employment transitions, 2012 to 2019

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
AI Exposure$_{2012}$	-0.004*	-0.004*	-0.003	-0.003	-0.003	-0.003	-0.003
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robot & Soft. Exp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,798,898	5,798,898	5,798,898	5,798,898	5,798,898	5,798,898	5,798,898
Adjusted R^2	0.042	0.052	0.073	0.108	0.153	0.209	0.213

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and employment transitions between 2012–2019. The dependent variable is a binary indicator equal to one if a worker transitioned from regular full-time employment in 2012 into marginal employment, unemployment, or inactivity by the end of each observation window, and zero otherwise. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.3: OLS regression of AI susceptibility on wage growth, 2012 to 2019, for occupations requiring different skill levels

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
AI Exposure$_{2012}$	0.009***	0.014***	0.015***	0.016**	0.017**	0.017**	0.016*
	(0.003)	(0.005)	(0.005)	(0.007)	(0.008)	(0.008)	(0.009)
Unskilled Occupation	ref.	ref.	ref.	ref.	ref.	ref.	ref.
Skilled Occupation	0.018***	0.021***	0.022***	0.024***	0.025***	0.027***	0.030***
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.005)
Complex Occupation	0.033***	0.041***	0.047***	0.053***	0.059***	0.062***	0.068***
	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.007)	(0.007)
Highly Complex Occupation	0.048***	0.062***	0.073***	0.081***	0.089***	0.094***	0.100***
	(0.006)	(0.008)	(0.010)	(0.012)	(0.013)	(0.014)	(0.014)
AI × Skilled	-0.008***	-0.012***	-0.013***	-0.014***	-0.015***	-0.014**	-0.014**
	(0.002)	(0.003)	(0.004)	(0.005)	(0.005)	(0.006)	(0.006)
AI × Complex	-0.006**	-0.009**	-0.008	-0.008	-0.008	-0.007	-0.006
	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.007)	(0.008)
AI × Highly Complex	-0.007**	-0.011*	-0.010	-0.011	-0.011	-0.012	-0.010
	(0.004)	(0.006)	(0.007)	(0.008)	(0.009)	(0.010)	(0.011)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Task Routineness	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robot & Soft. Exp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.030	0.046	0.071	0.091	0.107	0.120	0.131

Notes: This table presents OLS estimates of the relationship between AI exposure and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Occupational skill complexity is classified into four categories (unskilled, skilled, complex, highly complex). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.4: OLS regression of AI susceptibility on wage growth, 2012 to 2019, for occupations with different levels of task routineness

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
AI Exposure$_{2012}$	0.015***	0.020***	0.022***	0.025***	0.026***	0.025***	0.027**
	(0.005)	(0.007)	(0.007)	(0.008)	(0.009)	(0.009)	(0.011)
Cognitive Routine	ref.	ref.	ref.	ref.	ref.	ref.	ref.
Analytical non-routine	0.006***	0.007**	0.009***	0.010***	0.010**	0.011**	0.011**
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.005)
AI × Analytical non-routine	-0.004***	-0.006***	-0.009***	-0.010***	-0.011***	-0.012***	-0.012***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
Interactive non-routine	0.003	0.003	0.004	0.006	0.006	0.005	0.007
	(0.002)	(0.003)	(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
AI × Interactive non-routine	-0.003**	-0.005**	-0.007***	-0.008***	-0.010**	-0.011**	-0.013***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
Manual routine	-0.002	-0.003	-0.006**	-0.007**	-0.007*	-0.007*	-0.007*
	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
AI × Manual routine	-0.002**	-0.002*	-0.003**	-0.004**	-0.005***	-0.006***	-0.006***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Manual non-routine	-0.006**	-0.010***	-0.011**	-0.013**	-0.014**	-0.014**	-0.011
	(0.003)	(0.004)	(0.005)	(0.005)	(0.006)	(0.006)	(0.007)
AI × Manual non-routine	-0.001	-0.003	-0.005**	-0.006*	-0.007**	-0.007**	-0.007*
	(0.001)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Skill Complexity	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Robot & Soft. Exp.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.029	0.046	0.070	0.091	0.107	0.120	0.130

Notes: This table presents OLS estimates of the relationship between AI exposure and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Occupational task routine measures are based on Dengler et al. (2014). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security, Webb (2020) & Dengler et al. (2014).

Table A2.5: OLS regression of robot exposure on wage growth, 2012 to 2019

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
Robot Exposure$_{2012}$	-0.014***	-0.022**	-0.028**	-0.033*	-0.038*	-0.040*	-0.043*
	(0.001)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.027	0.042	0.066	0.086	0.102	0.115	0.125

Notes: This table presents OLS estimates of the relationship between robot exposure in 2012 and the difference in log daily wage between 2012–2019. Occupational robot exposure is a standardized exposure measure over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.6: OLS regression of software exposure on wage growth, 2012 to 2019

	(1) 2012–13	(2) 2012–14	(3) 2012–15	(4) 2012–16	(5) 2012–17	(6) 2012–18	(7) 2012–19
Software Exposure$_{2012}$	-0.005***	-0.009**	-0.012**	-0.015**	-0.017**	-0.018**	-0.019*
	(0.001)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)
Baseline Wage	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AKM Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.025	0.040	0.063	0.082	0.098	0.110	0.120

Notes: This table presents OLS estimates of the relationship between software exposure in 2012 and the difference in log daily wage between 2012–2019. Occupational software exposure is a standardized exposure measure over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.7: OLS regression of AI susceptibility on wage growth, 2012–2019, for workers with different education levels

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.022	0.008	0.011	0.019	0.002
	(0.018)	(0.010)	(0.012)	(0.012)	(0.008)
Low Education	ref.	ref.	ref.	ref.	ref.
Medium Education	0.002	0.020**	0.019*	0.013	0.001
	(0.015)	(0.010)	(0.010)	(0.009)	(0.006)
High Education	0.107***	0.084***	0.083***	0.067***	0.054***
	(0.020)	(0.009)	(0.009)	(0.008)	(0.007)
AI × Medium Educ	0.002	-0.003	-0.003	-0.001	0.007
	(0.013)	(0.009)	(0.009)	(0.009)	(0.006)
AI × High Educ	-0.009	-0.015	-0.014	-0.012	-0.003
	(0.017)	(0.010)	(0.010)	(0.010)	(0.006)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Skill & Task Controls	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	3,358,470	3,358,428	3,221,085	3,221,085	3,221,085
Adjusted R^2	0.076	0.115	0.117	0.130	0.128

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Education is classified as low (no formal vocational qualification), medium (completed vocational education), and high (university degree). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time workers liable to social security & Webb (2020).

Table A2.8: OLS regression of AI susceptibility on wage growth, 2012–2019, for males

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.011	0.010*	0.022**	0.026***	0.023***
	(0.007)	(0.005)	(0.007)	(0.006)	(0.004)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,914,312	1,914,294	1,875,732	1,875,732	1,875,732
Adjusted R^2	0.075	0.110	0.116	0.133	0.136

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to male workers. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Source: IEB (2012–2019), restricted to full-time male workers liable to social security & Webb (2020).

Table A2.9: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations requiring different skill levels, males

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.020	0.035**	0.043**	0.029**	0.004
	(0.014)	(0.015)	(0.015)	(0.010)	(0.010)
Unskilled Occupation	ref.	ref.	ref.	ref.	ref.
Skilled Occupation	0.039***	0.046***	0.026**	0.033***	0.029***
	(0.014)	(0.012)	(0.011)	(0.006)	(0.004)
Complex Occupation	0.133***	0.131***	0.078***	0.086***	0.077***
	(0.015)	(0.013)	(0.016)	(0.010)	(0.008)
Highly Complex Occupation	0.173***	0.147***	0.089***	0.111***	0.112***
	(0.043)	(0.026)	(0.032)	(0.019)	(0.013)
AI × Skilled	-0.025	-0.034**	-0.034**	-0.014**	-0.006
	(0.017)	(0.013)	(0.015)	(0.007)	(0.005)
AI × Complex	-0.024	-0.034**	-0.030*	-0.008	0.004
	(0.015)	(0.014)	(0.013)	(0.007)	(0.006)
AI × Highly Complex	-0.019	-0.027	-0.025	-0.014	-0.001
	(0.026)	(0.025)	(0.023)	(0.014)	(0.009)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Task Routineness	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,914,312	1,914,294	1,875,732	1,875,732	1,875,732
Adjusted R^2	0.087	0.118	0.119	0.135	0.138

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to male workers. Occupational skill complexity is classified into four categories (unskilled, skilled, complex, highly complex). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time male workers liable to social security & Webb (2020).

Table A2.10: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations with different levels of task routineness, males

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.022	0.009	0.024*	0.008	0.005
	(0.005)	(0.011)	(0.010)	(0.010)	(0.011)
Cognitive Routine	ref.	ref.	ref.	ref.	ref.
Analytical non-routine	0.018*	0.014*	0.010	0.010*	0.011**
	(0.010)	(0.008)	(0.008)	(0.006)	(0.005)
AI × Analytical non-routine	-0.023***	-0.021***	-0.017**	-0.012**	-0.013**
	(0.009)	(0.006)	(0.007)	(0.004)	(0.004)
Interactive non-routine	-0.006	-0.009**	-0.016***	-0.002	-0.002
	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)
AI × Interactive non-routine	-0.009	-0.004	-0.002	-0.002	-0.010**
	(0.007)	(0.005)	(0.006)	(0.004)	(0.004)
Manual routine	-0.029***	-0.027***	-0.027***	-0.021***	-0.011**
	(0.006)	(0.005)	(0.005)	(0.003)	(0.005)
AI × Manual routine	-0.011*	-0.006	-0.004	-0.001	-0.005*
	(0.006)	(0.004)	(0.004)	(0.003)	(0.002)
Manual non-routine	-0.030***	-0.038***	-0.038***	-0.029***	-0.013
	(0.008)	(0.006)	(0.006)	(0.005)	(0.008)
AI × Manual non-routine	-0.008	0.002	0.005	0.001	-0.001
	(0.007)	(0.006)	(0.006)	(0.005)	(0.005)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Skill Complexity	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,914,312	1,914,294	1,875,732	1,875,732	1,875,732
Adjusted R^2	0.096	0.124	0.123	0.136	0.139

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to male workers. Occupational task routine measures are based on Dengler et al. (2014). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time male workers liable to social security, Webb (2020) & Dengler et al. (2014).

Table A2.11: OLS regression of AI susceptibility on wage growth, 2012–2019, for females

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.020*	0.012*	0.014	0.015**	0.020***
	(0.006)	(0.006)	(0.009)	(0.005)	(0.005)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,444,158	1,444,134	1,345,353	1,345,353	1,345,353
Adjusted R^2	0.069	0.101	0.106	0.121	0.124

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to female workers. Occupational AI, robot and software exposures are standardized exposure measures over the 5-digit KldB (2010) occupations, based on Webb (2020). Standard errors in parentheses are clustered at the 5-digit occupation level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Source: IEB (2012–2019), restricted to full-time female workers liable to social security & Webb (2020).

Table A2.12: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations requiring different skill levels, females

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.032**	0.042**	0.043*	0.028**	0.023*
	(0.011)	(0.017)	(0.023)	(0.012)	(0.012)
Unskilled Occupation	ref.	ref.	ref.	ref.	ref.
Skilled Occupation	0.074***	0.068***	0.036**	0.045***	0.035***
	(0.024)	(0.020)	(0.015)	(0.009)	(0.006)
Complex Occupation	0.121***	0.106***	0.056***	0.075***	0.064***
	(0.023)	(0.019)	(0.019)	(0.010)	(0.008)
Highly Complex Occupation	0.199***	0.143***	0.076***	0.101***	0.095***
	(0.029)	(0.023)	(0.029)	(0.019)	(0.017)
AI × Skilled	-0.013	-0.014	-0.034*	-0.025***	-0.018*
	(0.013)	(0.015)	(0.016)	(0.010)	(0.010)
AI × Complex	-0.015	-0.015	-0.038*	-0.022*	-0.017
	(0.015)	(0.018)	(0.019)	(0.012)	(0.012)
AI × Highly Complex	-0.041	-0.016	-0.024	-0.021	-0.014
	(0.025)	(0.020)	(0.022)	(0.016)	(0.014)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Task Routineness	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,444,158	1,444,134	1,345,353	1,345,353	1,345,353
Adjusted R^2	0.077	0.107	0.109	0.123	0.126

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to female workers. Occupational skill complexity is classified into four categories (unskilled, skilled, complex, highly complex). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time female workers liable to social security & Webb (2020).

Table A2.13: OLS regression of AI susceptibility on wage growth, 2012–2019, for occupations with different levels of task routineness, females

	(1)	(2)	(3)	(4)	(5)
AI Exposure$_{2012}$	0.029**	0.058***	0.059***	0.048***	0.047***
	(0.007)	(0.014)	(0.016)	(0.010)	(0.009)
Cognitive Routine	ref.	ref.	ref.	ref.	ref.
Analytical non-routine	0.046***	0.050***	0.025***	0.028***	0.011**
	(0.008)	(0.008)	(0.008)	(0.004)	(0.005)
AI × Analytical non-routine	-0.014	-0.012	-0.009	-0.003	-0.010**
	(0.010)	(0.008)	(0.009)	(0.004)	(0.004)
Interactive non-routine	0.014**	0.007	0.003	-0.001	0.013**
	(0.007)	(0.008)	(0.009)	(0.004)	(0.006)
AI × Interactive non-routine	0.001	0.006	0.006	0.000	-0.011***
	(0.009)	(0.008)	(0.008)	(0.003)	(0.003)
Manual routine	-0.031***	-0.034***	-0.031***	-0.027***	-0.012**
	(0.008)	(0.007)	(0.008)	(0.005)	(0.006)
AI × Manual routine	0.017**	0.015**	0.015*	0.006	-0.002
	(0.009)	(0.008)	(0.008)	(0.004)	(0.003)
Manual non-routine	-0.010	-0.012**	-0.013	-0.026***	-0.016*
	(0.010)	(0.008)	(0.008)	(0.007)	(0.009)
AI × Manual non-routine	-0.017**	-0.012	-0.012	-0.004	-0.011**
	(0.009)	(0.008)	(0.009)	(0.005)	(0.005)
Baseline Wage	Yes	Yes	Yes	Yes	Yes
Worker Controls	No	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Skill Complexity	No	Yes	Yes	Yes	Yes
AKM Effects	No	No	Yes	Yes	Yes
Robot & Soft. Exp.	No	No	Yes	Yes	Yes
Region FEs	No	No	No	Yes	Yes
Industry FEs	No	No	No	Yes	Yes
Sales & Admin	No	No	No	Yes	No
2-digit Occupation FEs	No	No	No	No	Yes
Observations	1,444,158	1,444,134	1,345,353	1,345,353	1,345,353
Adjusted R^2	0.083	0.111	0.111	0.123	0.125

Notes: This table presents OLS estimates of the relationship between AI susceptibility in 2012 and the difference in log daily wage between 2012–2019, restricted to female workers. Occupational task routine measures are based on Dengler et al. (2014). Standard errors in parentheses are clustered at the 5-digit occupation level.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Source: IEB (2012–2019), restricted to full-time female workers liable to social security, Webb (2020) & Dengler et al. (2014).

Appendix 2B: Linking administrative data with exposure scores

As the Webb (2020) exposure measures are based on US occupations (Standard Occupational Classification, SOC 2010) we rely on the crosswalk provided by Heß et al. (2023) to map Webb (2020) exposure measures to the German classification of occupations from 2010 at the 5-digit level (Kldb 2010) which provides an up-to-date classification of occupations for Germany (Paulus et al., 2013). The conversion table for the allocation of occupations in KldB 2010 to ISCO-08 is provided by the Federal Employment Agency. The KldB2010 is one of the baseline classifications for the fulfillment of the German employers' obligation to enroll employees with employment information in the social security register. Once we restrict our sample to all selected employment spells that include June 30thas the cutoff date this leaves us with a representative panel sample to capture the work biographies of workers over the period 2012-2019.

As some occupational classifiers and task definitions within occupations may differ between the US and the German labour markets, inaccuracies might occur during the crosswalk. These inaccuracies can be particularly significant when observing individual occupations. However, since our study refers to occupational groups with the aim of deriving generalized statements, these inaccuracies are likely to be less significant. We cannot rule out the possibility that business models or the standard of services have changed since the indicators were calculated and that occupational task structures might therefore have to be weighted differently in the meantime.

3

AI in High-Tech Manufacturing: Firm-Level Evidence from a Multi-Plant Rollout

Abstract

This paper examines AI adoption and its implications for labour demand using detailed firm-level data from a multinational electronics manufacturer. We study the deployment of an AI-based visual inspection system across 14 plants over 15 months, tracking decision-level data from a quality assurance (QA) task. We document two primary findings. First, AI adoption unfolds gradually rather than as discrete shocks, with plants expanding automation primarily along the extensive margin (deployment on new products) rather than the intensive margin (changes in the share of decisions automated once AI is already deployed). We observe substantial heterogeneity across plants in both the level and pace of adoption, reflecting how differences in organizational practices, data availability, and customer requirements shape the intensity of deployment. Second, using two-way fixed effects estimation, we find a one percentage point increase in AI automation is associated with approximately a 1.2% decline in operator decision growth, suggesting labour displacement. We also observe modest improvements in inspection efficiency, with the overall pass rate increasing without a deterioration in quality. While we cannot make causal claims by design, our findings highlight that even when AI systems are technically capable and rigorously validated, deployment proceeds cautiously due to trust, safety, and reliability considerations.

This chapter is joint work with Didier Fouarge, Marie-Christine Fregin and Nicholas Rounding. My individual contribution to this paper involved contributing to the development of the research question and methodology, performing the data analysis and interpretation of results, conducting qualitative interviews with firm representatives, and contributing to the writing of the manuscript. This study is part of the research project ai:conomics funded by the German Federal Ministry of Labour and Social Affairs (Bundesministerium für Arbeit und Soziales (BMAS)/Denkfabrik Digitale Arbeitsgesellschaft) by resolution of the German Bundestag. We thank Sanne Steens, Sander Dijksman, Danique Eijkenboom, Lara Fleck, and Evie Graus for their research support. We thank the anonymous company for giving us access to the data and allowing to interview employees.

3.1 Introduction

Recent advances in Artificial Intelligence (AI) have sparked numerous debates on the impact of automation on labour demand and allocation of worker tasks, igniting perennial anxieties about technological unemployment. While some AI technologies are predicted to be a General Purpose Technology with profound effects on productivity in labour markets, its employment consequences remain contested (Brynjolfsson et al., 2025a; Hampole et al., 2025; Humlum and Vestergaard, 2025). Despite a growing body of direct firm-level evidence suggesting that AI tools can improve labour productivity, including the evidence presented in Chapter 4 of this thesis (Brynjolfsson et al., 2025b; Dell'Acqua et al., 2025; Dillon et al., 2025), evidence on the direct labour displacing effects of AI adoption remains limited, leaving critical questions regarding its extent and scope unanswered.

Due to a lack of reliable adoption data, scholars have relied on indirect estimates of the impact of AI through occupational or firm-level exposures. An extensive literature has developed and utilised occupational exposure scores, mapping AI capabilities to occupational task descriptions [1] (Brynjolfsson and Mitchell, 2017; Gmyrek et al., 2025; Webb, 2020; Felten et al., 2021; Frey and Osborne, 2017b; Prytkova et al., 2024). While these exposure measures have proven valuable for assessing which occupations and tasks are most susceptible to AI-related changes, they are designed to capture potential exposure rather than actual adoption. By construction, they often assume that technological capability equates to substitution potential, and do not take into account the technological nor economic feasibility of task replacement. As a result, the exposure scores risk overestimating the potential impact of AI adoption on labour displacement (Svanberg et al., 2024). Furthermore other barriers to task displacement, such as trust in AI systems or regulatory constraints, could impede the progress of task displacement. Therefore, while exposure scores provide important insights into which parts of the labour market may be affected by AI, they cannot reveal when, how quickly, or to what extent firms actually adopt these technologies. To understand the true labour displac-

[1] Chapter 1 provides a detailed discussion of occupational AI exposure measures and presents empirical evidence to study worker-level outcomes.

ing effects of AI, we need direct evidence on how firms actually adopt and integrate these technologies into production.

This paper addresses this gap by providing detailed firm-level evidence on AI adoption and its implications for labour displacement in a high-tech manufacturing company. Specifically, we study how an AI-based visual inspection system is adopted and integrated into the quality-assurance (QA) process within a large multinational electronics manufacturer supplying advanced components to the automotive industry. The QA stage is critical for product reliability and safety and requires production workers to inspect printed circuit boards (PCBs) under strict accuracy, consistency, and throughput requirements. These features—high precision requirements, low adjustment costs, and strong organisational incentives for reliability—make this setting particularly well suited to studying the sequential nature of AI adoption and its consequences for task reallocation and labour demand in an industrial environment. Our analysis relies on firm-level data containing information on QA decisions made by either the AI tool or a human operator over a 15-month period from September 2022 to November 2023. This unique dataset allows us to investigate (i) how AI adoption unfolds over time and across plants, (ii) how firms expand AI automation along the extensive and intensive margins, and (iii) how increasing automation relates to changes in labour demand.

We document two findings. First, we find that AI adoption unfolds gradually and continues to increase long after initial deployment. Rather than discrete technological shocks, we observe plants progressively expanding automation along both the extensive margin (deploying AI on additional products and component libraries) and the intensive margin (adjusting decision thresholds and retraining algorithms to increase AI's share of decisions within existing deployments). To capture heterogeneity in deployment timing, we classify plants as early adopters if they began deployment more than 8 months before our observation period, and late adopters otherwise. We demonstrate that early adopters reach higher automation levels more quickly, while late adopters exhibit faster relative growth once the technology is introduced. Interestingly, expansion occurs primarily along the extensive margin, with plants deploying AI on new products rather than systematically increasing its intensity on already-enabled components.

Second, increasing automation is associated with changes in labour demand within the QA task. As the AI takes on a larger share of the inspection workload, we find that the volume of operator decisions declines. We find that a one percentage point increase in AI automation is associated with approximately a 1.2% decline in operator decision growth. This shift suggests both task reallocation and a reduction in labour input as automation intensifies. Taken together, these findings demonstrate that despite substantial global investment in AI capabilities, deployment within firms can occur incrementally, shaped by trust, safety, and organizational considerations rather than by discrete technological shocks.

Our study contributes to the growing literature on the labour market effects of automation and AI in several ways. First, while the task-based approach of Acemoglu and Autor (2011) and Acemoglu and Restrepo (2019) has provided the foundation for a rapidly expanding literature on the displacement effects of automation across industries and occupations, there remains limited empirical evidence about how such technologies are actually adopted and integrated into production settings. We build on this empirical task literature by studying the pace and extent of automation under which the AI deployment occurs. In doing so, we highlight how the deployment of AI unfolds over time within a high-stakes industrial setting, reshaping task allocation and share of labour input at the firm level.

Second, our study relates to the literature on the diffusion of new technologies and advances the empirical measurement of AI adoption at the firm level. Recent work has focused on measuring firm-level AI exposure by identifying AI-related skills in job advertisements (Acemoglu et al., 2022; Peede and Stops, 2024; Storm et al., 2025; Babina et al., 2024), while other studies have examined early-stage adoption patterns in knowledge work settings (Brynjolfsson et al., 2025a). However, vacancy postings represent hiring intentions rather than actual technology implementation and cannot capture the timing, intensity, or evolution of deployment once adoption occurs. Our study overcomes these limitations by directly observing AI deployment decisions as they unfold, allowing us to document not only whether a technology is adopted, but how extensively it is used and how its dynamics evolves over time. Moreover, a key insight from the technology diffusion literature is that different innovations

follow distinct adoption patterns. Prior automation technologies, such as robotics, typically involved lumpy capital-intensive investment spikes (Bessen et al., 2023; Domini et al., 2021, 2022). These discrete investment episodes often coincided with sharp shifts in labour demand, as firms reorganized production and substituted capital for labour in routine tasks (Acemoglu and Restrepo, 2020a; Bessen et al., 2023). Such episodes typically entail substantial adjustment costs—in terms of capital and labour—arising from new equipment installation, retraining, and task reallocation. In contrast, we analyze whether the diffusion of AI-based technologies follows a different pattern. Because AI tools are frequently integrated into existing production processes without large fixed investments or major reconfiguration of production lines, adoption can proceed sequentially and at relatively low adjustment cost (Asphjell et al., 2014). In our setting, we observe this gradual expansion of automation along both the extensive margin and the intensive margin. This sequential nature of adoption generates considerable heterogeneity across plants, reflecting differences in technological readiness, data availability, and organizational practices.

Finally, we highlight the central role of trust, safety, and reliability in the adoption of AI-based systems (Daly et al., 2025). Even when AI tools are technically capable of taking over tasks, firms are often reluctant to delegate decisions without sustained performance validation. In our setting—where quality assurance is safety-critical—we show how adoption proceeds cautiously, with AI initially used in decision support roles before gradually taking on autonomous tasks. This underscores the importance of organisational dynamics and confidence in shaping the trajectory of AI deployment, and suggests that even disruptive technologies may diffuse in an incremental and human-centric manner when deployed in high-stakes production environments.

The rest of the paper is structured as follows. In Section 3.2 we provide the company background and describe the AI adoption process within the firm. Section 3.3 outlines the data overview and presents descriptive statistics. Section 3.4 presents our main findings. We first analyze how AI deployment evolves across plants, capturing both the extensive and intensive margins. We then describe our empirical strategy and present results on the relation between AI adoption and changes in labour demand. Section 3.5 provides

robustness checks and sensitivity analyses. Section 3.6 concludes.

3.2 Empirical Context

3.2.1 Company Setting

We evaluate the effects of AI adoption in a large global international electronics manufacturer that develops, supplies and produces technologically advanced products and services for original equipment manufacturers in the automobile industry. The company is one of the largest in its market, with a global reach, around 100,000 employees, and sales figures in the tens of billions of euros. Despite this, competitive pressures are high, both for traditional competitors and new comers into the market. This pressure demands the company invest in continuous innovation and process improvement to maintain their market position. The successful adoption of AI is viewed by the company as crucial step in forward development.

Specifically, we evaluate the adoption of AI in a particular task within this company. We focus on a QA task that involves inspecting soldering points on Printed Circuit Boards (PCBs) to detect anomalies. Performed by quality control workers, this task is particularly demanding, requiring sustained visual attention and precision to identify minor defects at high production volumes. In what follows, we provide details on the company's production process, the nature of the QA task before AI, and the subsequent introduction and deployment of AI technologies that automate parts of this process.

3.2.2 Printed Circuit Board Assembly

The company assembles PCBs using surface-mount technology (SMT), a highly automated and technologically advanced production method with widespread use in PCB assembly. An exemplary representation of the SMT production environment and a circuit board are shown in Figure A3.1. In SMT production lines automated pick-and-place machines put electronic components—such as capacitors, diodes, resistors and transistors—onto circuit boards that have

had solder paste painted on at specific solder points. These boards are then passed through a reflow oven, melting the solder and permanently connecting the components to the board. For a more detailed overview of the process, see Prasad (2012).

Following the reflow soldering process, the assembled PCBs undergo visual inspection to detect manufacturing defects, such as faulty soldering points or misaligned components. Because PCBs serve as the central nervous system of all electronic devices, their precise and flawless performance is essential. Even minor defects, such as micro-cracks in solder joints or insufficient soldering, can lead to functional failures. In many of the company's end products, which include safety- and mission-critical applications (e.g., automotive safety and control systems), such failures can have severe consequences. As the SMT production process allows for smaller component sizes and higher circuit densities than alternative assembly processes (such as through-hole assembly), there is an increased risk of defective units. As a result, the QA process is a critical stage in the production process, crucial for identifying even minor deviations from manufacturing standards (Acciani et al., 2006). This ensures product reliability, operational safety, and compliance with stringent industry regulations and customer demands. In our case, customers impose particularly high accuracy requirements and detailed inspection protocols, ultimately shaping how extensively the AI system can be utilised in the QA task.

3.2.3 QA before AI

Prior to the introduction of AI, the visual inspection task of PCBs was carried out by human quality control operators and a rule-based Automated Optical Inspection (AOI) system. AOI systems capture high-resolution, multi-angle images of PCBs and compare them to a predefined "golden" template (Abd Al Rahman and Mousavi, 2020). If the captured images match exactly, then the component is marked as non-defective (or OK) and continues in production. If there are any detected deviations, components are flagged as defective. Because the rule-based algorithms operate with limited flexibility, they tend to generate a substantial share of false negatives, flagging components as

potentially defective even when they conform to production standards.[2] Since defective components must be either scrapped or manually repaired, this high false negative rate would introduce significant waste without further inspection. Human operators therefore play a vital role in improving the efficiency of the process: they review all components flagged by the AOI as defective, reclassifying OK components while catching any truly defective ones that must be removed from production.

Operators are selected from unskilled production line workers called assemblers, whose primary task is to load materials into the SMT production line. Assemblers can apply to become operators, though the position offers no pay premium and provides only the benefit of being seated during the task. The visual inspection task itself is cognitively demanding, requiring sustained visual attention and precision to identify the few defects within a high volume of non-defective images. Of the images seen by operators, 90-95% are eventually classified as OK. Due to these demand, workers must pass a simulation-based assessment at the beginning of each shift before being allowed to do the inspection task. Furthermore, operators are only assigned to this task for a portion of their shift, typically around 2 hours, after which they return to their regular assembly duties.

3.2.4 AI-Automated QA Decisions

The traditional QA workflow, while effective at catching defects, remains inefficient: due to stringent quality requirements, operators are trained to err on the side of caution. This results in a higher scrap rates than necessary. Developments in machine learning technologies created an opportunity for the company to tackle this inefficiency, while simultaneously reducing labour costs without compromising on quality. Based on an algorithm design taken from a Google white paper, the company's AI and data science division developed a

[2]We define false negatives as components flagged as defective that are actually OK. True negatives are components flagged as defective that are actually defective. True positives are components flagged as OK that are OK, false positives are components flagged as OK that are actually defective. Note that in quality assurance literature, this terminology is often inverted, with 'positive' referring to defect detection rather than non-defective classification, see: Abd Al Rahman and Mousavi (2020).

supervised learning classification algorithm for the QA task.

The AI system is designed to reduce the volume of images requiring human review by acting as an intermediate classifier between the AOI and human operators. Positioned after the rule-based AOI, the AI reviews components flagged as defective and autonomously reclassifies those it can confidently predict to be OK, allowing them to continue in production without operator review. These are actual share of AI-automated decisions as previously they would have required an operator to classify them as OK or defective. To ensure ongoing quality control, a sample of AI-classified OK components are also sent to operators for verification. Components that the AI flags as defective, or cannot classify with sufficient confidence, are forwarded to operators for further inspection. Figure 3.1 provides an overview of this process. The implementation of AI into this QA workflow changes the nature of the operator's task: rather than reviewing the full set of AOI-flagged components, operators now handle a filtered subset consisting primarily of genuinely ambiguous or defective cases, alongside a small subset of OK decisions for monitoring.

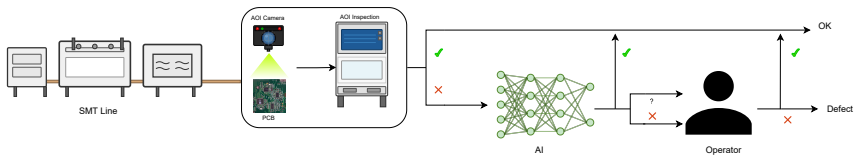


Figure 3.1: Quality Assurance Process Flow: AOI, AI, and Operator Decision Logic

Notes: This diagram demonstrates the Quality Assurance process flow once the AI has been enabled. First, the Automated Optical Inspection (AOI) system classifies components as OK or defective. Components classified as OK are forwarded directly to the back-end. Components classified as defective are then forwarded to the AI system. The AI autonomously decides on high-confidence cases, passing OK components to back-end production. If the AI classifies an image as defective, or the AI is unable to make a decision, it is forwarded to human operators for final verification. Operator decisions result in either acceptance to back-end production or component rejection.

Source: Author's own depiction.

3.2.5 Training and Deployment of the AI

The company's AI and data science division developed the algorithm, training procedures, and deployment pathways at a global level. Actual training and deployment, however, occur at the plant level due to differing production technologies across plants. Therefore, an AI model is trained for each plant using a locally curated test dataset. Each dataset contains samples of regular production data, as well as sufficient numbers of images of each known error class. These images are then labeled by certified PCB experts before being used to train the AI. While operator-generated labels can be incorporated into the training process, they are typically included only after being verified by experts, and the extent to which they are used depends on the maturity of the model at the respective plant.

The images used for AI training are grouped by component type. These groupings are referred to as component libraries, where a component library contains physically similar components sharing common visual features (e.g., all resistors of a specific size). Training is entirely supervised and component-specific, with no ongoing learning after deployment. Before deployment, the AI must achieve a false positive rate under 2%, and a false negative rate under 10%. As all components flagged as defective are eventually reviewed by operators, there is higher tolerance for false negatives. The false positive rate offers a substantial improvement over human operators, who are estimated by the company to have rates as high as 10 or 15%. Once the AI has exceeded these thresholds, plants can begin the deployment process.

Deployment occurs through a multi-step, gradual process contingent on the technical readiness of the AI model, customer approval, and plants' confidence in the technology. As a result, even highly capable AI models are not deployed uniformly or at full scale. This process proceeds along two margins: extensive and intensive. At the extensive margin, the AI is deployed on new products and component libraries. Once a model is trained and approved for release by the plant, it cannot be used on any product without customer approval. The company must therefore convince the customer that the AI provides a substantial quality improvement over the traditional workflow.

Once the AI has been deployed on a component, plants can adjust the intensive margin through several mechanisms. They may lower certainty thresholds over time, allowing the AI to handle a larger share of decisions. They may retrain models with updated data, improving the AI's confidence on previously marginal cases. Finally, they may reduce the sampling rate for operator verification of AI-approved decisions. These adjustments vary considerably across plants, resulting in substantial heterogeneity in how intensively the AI is utilized. We explore adjustments both at the extensive and intensive margin further below, in Section 3.4.1.

3.2.6 Role of human operators after AI deployment

Following AI deployment, human operators (i.e., assemblers assigned to the QA task) continue to play a crucial role in the QA process. They review all components for which no AI model has been trained, handle cases where the AI cannot decide with sufficient confidence, and inspect all components the AI flags as defective. Additionally, operators verify a sample of OK decisions made by the AI, allowing plants to monitor the AI's accuracy over time.

This changing role has significant potential implications for the work of operators. With a fully deployed AI, their task composition changes as they see a higher proportion of defective decisions, and potentially more challenging components to classify. They also become an important mechanism for ensuring the AI continues to operate reliably. Further, they are exposed to decreasing work demand in the verification task due to task replacement. This could either free up time for them to spend on the production line as assemblers, or reduce the overall demand for assemblers within the plant, potentially leading to job losses.

In the next section we investigate the relation between AI deployment on the QA process through the analysis of decision data provided to us by the company. Specifically, we address three research questions to elicit insights into the changing nature of the QA task. First, how does AI deployment unfold over time across plants, and along which margins—extensive or intensive—does automation expand? Second, how do changes in the intensity

of AI automation relate to labour demand within the QA task, as proxied by operator decision activity? Third, does increasing AI automation affect total decision throughput or the outcomes of the QA process, as measured by pass rates?

3.3 Data

3.3.1 Data Overview

The company provided us with access to customized QA decision data compiled specifically for this research project. The data cover 16 plants operating in 8 countries across Asia, the Americas, and Europe and include the date, time, and result of all QA decisions made by human operators and the AI. The data also include anonymized identifiers for the components inspected, the products on which they appear, and their associated component libraries.

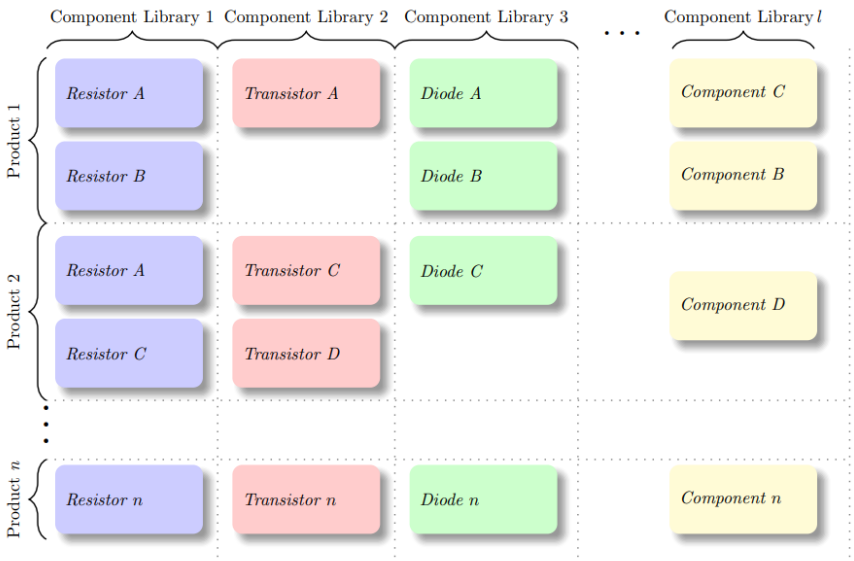
QA decisions are made on images of individual components. Each component belongs to both a specific product (the PCB being inspected) and a component library (a grouping of similar components, as described in Section 3.2.5). The relation between products, component libraries, and components is illustrated in Figure 3.2. Our primary analyses use a week-plant panel, though we also construct a week-component library panel for robustness checks.

Our main variable of interest is the **AI Automation Rate**. This is defined as the share of QA decisions actually automated by the AI relative to the total number of QA decisions. Automated decisions are those that are independently classified as "OK" by the AI without any operator review.[3] This measure captures the intensity of AI-driven automation in the QA process across plants and over time, reflecting adjustments at both the extensive and intensive margins, and forms the basis for the analyses presented in the subsequent sections.[4] Our

[3] Section 3.2.4 provides a detailed description of the AI system and the automated decision making process.

[4] It is important to note that this measure differs fundamentally from the occupation-level AI exposure measure used in Chapter 1, which captures potential susceptibility based on task content. Here, we observe actual AI adoption and usage within the firm, enabling us to study realized automation rather than potential exposure.

Figure 3.2: Relationship between Products, Component Libraries, and Components



Notes: This figure displays our data structure: Components exist at the intersection of Products and Libraries. Each product uses multiple components from multiple libraries, and each library contains components used across multiple products. Component libraries are groupings of similar types of components (here represented as Resistors, Transistors and Diodes). In practice, component libraries are more granular, e.g., grouping resistors by specific size and shape.

Source: Author's own depiction.

focus on "OK" decisions follows directly from how the AI system operates within the QA workflow. As described in Section 3.2.4, the AI is only allowed to make an autonomous decision when it classifies an image as non-defective. These cases constitute the part of the inspection task that the system can execute independently, without human verification. By contrast, any image the AI flags as defective, or is uncertain about, is routed to the human operator for review. We do not focus on the AI defect rate as the thresholds for the AI to make a decision are set significantly lower, and thus operators overturn up to 75% of AI defect decisions. Measuring the share of AI-approved OK decisions relative to all QA decisions therefore provides a direct indicator of the portion of the QA task that the AI can autonomously execute without human review.

Using this automation variable as a starting point, we introduce several measures in the following sections that allow us to (i) characterise how AI deployment unfolds over time and across plants, and (ii) examine how changes in automation intensity relate to labour demand within the QA task.

3.3.2 Sample Selection and Summary Statistics

We implement several sample restrictions to ensure data quality and consistency. First, we exclude two plants that contained no information on operator decisions, leaving us with 14 plants for analysis. Second, we address a data reporting issue that emerged toward the end of our observation period: across multiple plants, operator decision data were not recorded despite ongoing operator activity. To mitigate potential measurement error from this reporting gap, we truncate the sample, removing all observations from the affected period. This restriction eliminates 27.7% of plant-week observations, corresponding to 22.4% of total decisions. Our final estimation sample comprises 14 plants observed over 60 weeks, from September 2022 through November 2023. Although this cutoff helps us to account for a systematic data collection problem across plants, there remain persistent issues with the data throughout our observation period relating to missing product and decision information.

Within our final sample, an average of 9% of decisions per week are missing product and component identifiers. Thus we cannot assign them into compo-

nent libraries or identify the products they were made on. Missing product information occurs across all plants, throughout the observation period. Decisions with products missing have a 21% lower AI Automation Rate than non-missing, and are 5% less likely to be passing decisions. For our primary analysis, we remove those decisions with missing product information. This is likely to overestimate the AI Automation Rate, therefore we address this through a series of robustness checks in Section 3.5. We also run sensitivity analyses using our variable construction, but excluding weeks with high levels of missing product information. A further data issue arises from lines in the data missing decisions information, but containing product identifiers. On average, around 6% of data rows per week contain no decision information but have an assigned product identifier. To account for this, we identify weeks with a high level of missing product information and exclude them from our sample in a sensitivity analysis.

Table 3.1: Summary Statistics

	Mean	SD	Min	Max
Panel A: Plant Characteristics				
Weeks Observed	59	2	53	60
Unique Products	902	811	130	3303
Unique Component Libraries	1118	671	102	2404
Share of total QA Decisions (%)	7.1	10.6	0.2	42.0
Panel B: AI Automation & QA Outcomes				
AI Automation Rate (%)	23.1	20.6	1.8	67.9
Overall OK Rate (%)	95.3	4.2	86.7	99.8
Operator OK Rate (%)	93.8	5.2	83.4	99.7
Panel C: Week-to-Week Changes				
Δ AI Automation Rate (pp)	0.164	0.109	-0.038	0.374
Δ Total QA Throughput (Log)	0.008	0.018	-0.021	0.046
Δ Total Operator Decisions (Log)	0.006	0.018	-0.021	0.044
Δ Total AI Decisions (Log)	0.030	0.048	-0.024	0.167

Notes: This table reports plant-level sample statistics across the observation period. Panel A: observation weeks, unique products/libraries per plant, and each plant's share of total QA decisions. Panel B: mean AI Automation Rate, Overall OK Rate, and Operator OK Rate (plant-week averages). Panel C: mean week-to-week changes in AI Automation Rate (percentage points) and log changes in QA throughput, operator decisions, and AI decisions (plant-week averages).

Table 3.1 summarizes key characteristics of these variables across 14 plants and 60 weeks during the observation period in the plant-week panel. In Panel A, we observe that most plants have the full set of 60 week observations. However two plants have missing data resulting in only 53 and 57 week observations, thus averaging 59 week observations for plants. Plants make decisions across an average portfolio of 902 different products and 1,118 component libraries during the observation period. Plant size varies substantially: the smallest plant accounts for only under 0.2% of total QA decisions whereas the largest contributes to 42%. This heterogeneity in plant sizes presents challenges for our empirical strategy. On the one hand, we want to avoid results being driven disproportionately by the largest plants; on the other hand, we want to prevent very small plants having an outsized influence. To mitigate these concerns, we express outcomes in logs where appropriate and report both unweighted regressions and regressions weighted by the total number of QA decisions recorded in the plant at the plant–week level.

Panel B in Table 3.1 reports on the average AI Automation Rates and QA Outcomes. We observe substantial heterogeneity across plants, with average weekly AI Automation Rate ranging from 1.8% to 67.9%. The average AI Automation Rate across plants is 23.1%, thus on across plants roughly 75% of decisions are taken by human operators. For all QA decisions, regardless of whether they have been taken by the AI or a human operator, we find a high average OK rate of 95.3%, with some variation between plants ranging from 86.7% to 99.8%. The Operator OK Rate is slightly lower, at 93.8%. This discrepancy suggests that the AI influences the OK rate at the plant level. Operators may be more likely than the AI to flag components as defective. Alternatively, this pattern may reflect changes in task allocation, with operators handling a higher share of defective images as the AI automates primarily non-defective cases. We further explore these mechanisms in Section 3.4.4.

To assess the impact of AI Automation, we rely on weekly first differences, therefore in Panel C in Table 3.1, we report on week-to-week changes in QA decisions. We first report on our explanatory variable, AI Automation Rate. We see an average weekly growth in the AI Automation Rate of 0.164 percentage points across plants, across our observation window. Turning to Total Throughput, defined as the total number of decisions taken by operators and

AI, in general we observe small, positive average weekly changes of about 0.8%, suggesting that the plants increase the throughput during our observation window. Turning to Total Operator Decisions, we similarly demonstrate a slight increase of 0.6%. Whilst similar in magnitude, albeit slightly slower, this provides an early indication of substitution between labour and AI within this task. Measuring the increase in AI decisions, we see that the Total AI Decisions increases by about 3% per week, a higher growth rate than either Total QA Throughput or Total Operator Decisions. This pattern suggests that growth in AI decisions outpaces growth in overall inspection activity, implying a gradual shift in task allocation away from human operators.

3.4 Results

In this section, we first document how AI deployment unfolds across plants during our observation period, highlighting differences between early and late adopters and the relative contributions of the extensive and intensive margins. We then turn to labour demand, where we assess how week-to-week changes in the intensity of AI automation relate to changes in operator activity at the plant level and the empirical strategy guiding our analysis. Finally, we examine two additional margins of adjustment, overall QA volume and QA accuracy, to understand whether AI adoption affects productivity or the quality of the inspection process.

3.4.1 Mapping AI Deployment

We start by exploring the progression of AI deployment across plants during our observation period using the AI Automation Rate. We have precise start dates for all plants that began deployment within 15 months before the start of our observation period, with the remainder deploying earlier. Of the fourteen plants in our sample, nine had already begun AI deployment prior to the start of the observation period, while five deployed during our observation window. Among the nine that deployed before our sample, seven began more than 8 months prior, while two began 1-4 months before. The five plants

that deployed during the observation period began at various points in time: three within the first 3 months and two between 8-11 months after the sample start. Table 3.2 provides the full distribution of deployment timing. To capture heterogeneity in deployment phases, we classify plants as "early adopters" if they began deployment more than 8 months before our sample period (N=7), and as "late adopters" if they began deployment 8 months or less before the sample start or during the observation window (N=7).

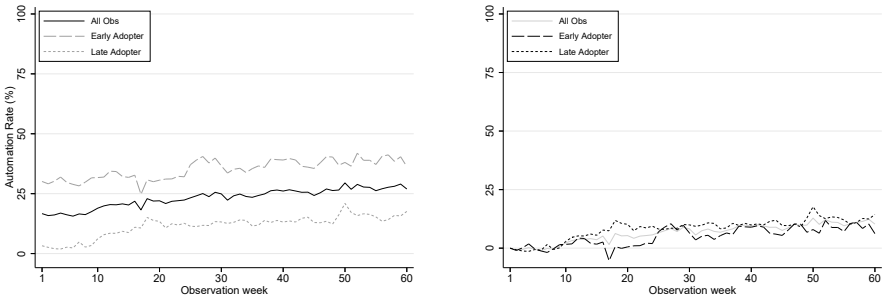
Table 3.2: Plant distribution by AI deployment timing

Deployment timing	Number of Plants
Early Adopters (>8 months before)	
>15 months before	4
8-15 months before	3
Late Adopters (<=8 months before/during)	
1-4 months before	2
0-3 months after start	3
8-11 months after start	2
Total number of plants	**14**

Notes: This table reports the number of plants by AI deployment timing relative to the observation window start. Early Adopters deployed AI more than 8 months before the observation period; Late Adopters deployed within 8 months before or during the observation period .

Figure 3.3 plots the AI Automation Rate averaged for all plants and split between early and late adopters. Figure 3.3a shows that the overall AI Automation Rate across plants remains relatively low, starting at around 16% at the beginning of the observation period and reaching 27% by the end. Between early and late adopters, we observe a persistent difference in levels, with early adopters maintaining higher automation rates throughout our observation period. Figure 3.3b compares the trajectories of early and late adopters by indexing to Week 1, revealing that both groups exhibit similar patterns of increase over time. This suggests that long after initial deployment, plants continue to gradually increase automation rather than reaching a stable steady

state. This indicates that AI implementation is an ongoing process rather than a discrete event. Figure A3.3a in Appendix shows individual plant trajectories, confirming that the group means are representative despite variation across plants within each adopter group.



(a) Automation Rate - Levels (b) Automation Rate -Indexed Growth

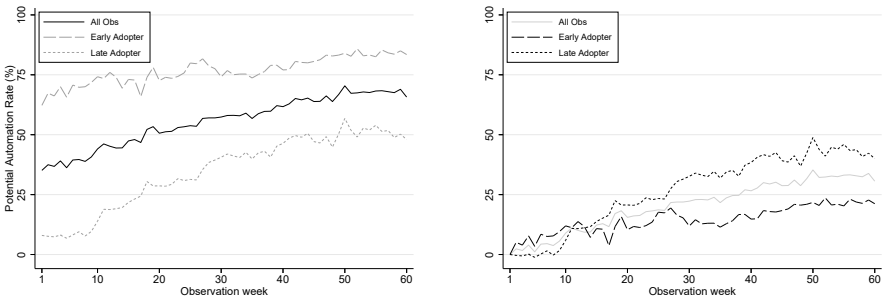
Figure 3.3: AI Automation Rate

Notes: Panel (a) shows the share of QA decisions automated by AI over the 60-week observation period for all plants, early adopters (deployed >8 months before sample start), and late adopters (deployed ≤8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

What drives this sustained automation growth? In Section 3.2.5 we distinguished two mechanisms: at the extensive margin, automation can increase as AI is progressively enabled on additional products or a newly trained algorithm is released on an additional component library. At the intensive margin, automation can rise as the thresholds at which the AI can make an autonomous decision are adjusted, or a retrained algorithm with additional images is deployed.

To examine the extensive margin of automation, Figure 3.4 plots the share of QA decisions made on components where the AI has been deployed, which we refer to as Potential AI Automation Rate. Unlike the AI Automation Rate, this measure contains both automated decisions and those taken by operators. Thus, it captures the expansion of AI deployment onto new products and new components but not necessarily the expansion of AI Automation. Figure 3.4a shows that the share rises from 30% to 70% across plants, through our ob-

servation period. Again, we observe a levels difference between early and late adopters, with a higher extensive margin for early adopters throughout. However, moving to Figure 3.4b, we begin to see differing patterns across the two groups. When indexed to the first week's value, we see that late adopters increased the extensive margin by 40 percentage points compared to just 20 percentage points for early adopters. Comparing the automation rate to expansions at the extensive margin, we see that the increase at the extensive margin is higher than the overall increase in the automation rate. This suggests that plants deploy the AI at a relatively low intensity.



(a) Potential AI Automation Rate - Levels (b) Potential AI Automation Rate Indexed

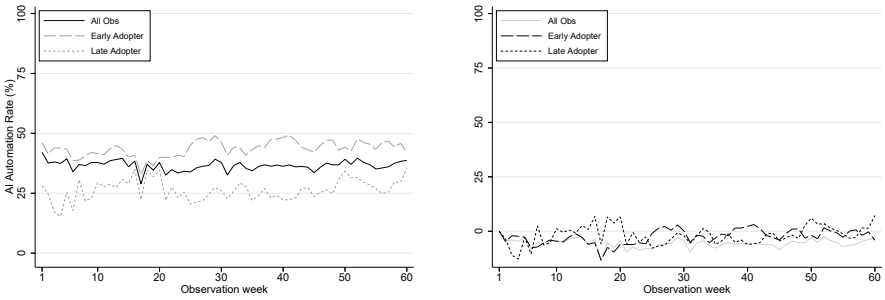
Figure 3.4: Potential AI Automation Rate

Notes: Panel (a) shows the share of QA decisions made that could be automated by AI over the 60-week observation period for all plants, early adopters (deployed >8 months before sample start), and late adopters (deployed ≤8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

To understand the intensity of AI use, we investigate the AI Automation Rate on decisions where the AI is deployed (i.e., the intensive margin). The average AI Automation Rate once the AI has been deployed on a component is around 36% for the entire period, showing that the AI only automates around one third of decisions leaving a substantial share for operators to review even after deployment. Comparing early and late adopters, we find a rate of 44% for early adopters and 28% for late adopters suggesting that plants who began using the AI earlier use it more intensely at both the extensive and intensive margins. However, even after reaching maturity, plants are still cautious about

how much to automate and require human oversight on over half of the AI's decisions.

This pattern is illustrated in Figure 3.5, which plots the AI Automation Rate, removing all decisions made on components where the AI is not deployed. Figure 3.5a reflects the averages described above; we find that on average there is higher *intensive* usage of the AI by early adopters compared to the late adopters. Looking into the indexed change, in Figure 3.5b we see that for both early and late adopters the adjustments at the intensive margin appear to be more random and show no systematic growth or decline over our observation period.



(a) AI Automation Rate - Levels (b) AI Automation Rate Indexed

Figure 3.5: Progression of AI Automation Rate on components where the AI has been deployed

Notes: Panel (a) shows the share of QA decisions automated by AI on components where the AI has been deployed, early adopters (deployed $>$8 months before sample start), and late adopters (deployed $\leq$8 months before sample start). Panel (b) shows the same series indexed to Observation Week 1 to highlight relative growth trajectories.

Taken together, these figures help us to understand the average progression of AI deployment. In plants, the AI Automation Rate increases, both across early and late adopters. These changes appear to be largely driven by the extensive margin, that is the deployment of AI on new products or component libraries as we are unable to see a systematic rise in the intensive margin over our observation period.

3.4.2 Labour Demand

3.4.2.1 Empirical Approach

We now turn to examining the relation between AI adoption and labour demand within the QA task. Two factors impede causal analysis. First deployment timing of the AI at product or component level is endogenous: plants are able to anticipate deployment and adjust their labour demand in advance. Second, products contain both treated and untreated components and operators work across lines with both treated and untreated components, creating spillovers that violate the Stable Unit Treatment Value Assumption (SUTVA) (see Appendix 3B for a detailed discussion on the challenges to causal identification).

We therefore aggregate to the plant-level and rely on a Two-Ways Fixed Effect (TWFE) estimation in a first difference framework to control for unobserved, time-invariant plant characteristics and estimate how week-to-week changes in the intensity of automation relate to changes in the demand for human labour. This allows us to capture the relation between AI and labour demand throughout the plant, including any potential spillover effects on products or components where the AI had not previously been enabled.

Figure 3.6 plots the weekly change in AI Automation Rate against the weekly change in the total operator decisions at the plant level, indicating a negative relation. To quantify this, we estimate the following model:

$$Y_{i,t} = \alpha + \beta \cdot \Delta\text{AI Automation Rate}_{i,t} + \gamma_t + \theta_i + \varepsilon_{i,t} \tag{3.1}$$

Here, $Y_{i,t}$ denotes the weekly change in the outcome of interest in plant i and week t e.g Δ Total Operator Decisions (Log). Our main explanatory variable, $\Delta\text{AI Automation}_{i,t}$, reflects the weekly change in the AI Automation Rate at the plant level. We include week fixed effects γ_t to control for common time shocks affecting all plants, such as the output growth. We also include plant fixed effects, θ_i, to account for plant specific growth trends. We cluster standard errors at the plant level. To account for the heterogeneity in plant

size, we run both unweighted and weighted regressions, weighting for the total number of decisions recorded in the plant in that week.

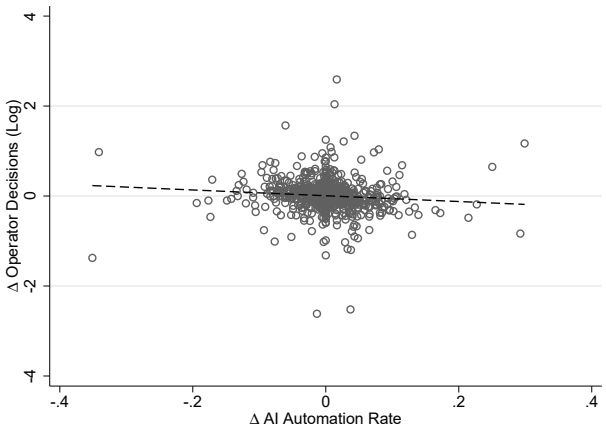


Figure 3.6: Change in operator decision and AI Automation Rate

Notes: This figure presents the relation between log changes in operator decisions and changes in the AI Automation Rate. Each dot represents a plant-week observation. The dashed line shows the line of best fit, indicating a negative relation between AI automation and operator decision growth.

3.4.2.2 Labour Demand Outcomes

We first explore the relation between an increase in the AI Automation rate and labour demand, as proxied by the growth rate of total operator decisions (i.e., Δ Total Operator Decisions (Log)). A limitation with this approach is that we rely on the number of decisions taken by operators as a proxy for total labour demand due to a lack of data on specific operators, including total employment numbers, working hours, and time spent per decision. It could be that the total demand for operators stays the same as they are able to spend more time on the decisions that remain after automation. However, discussions with the company suggest that such adjustments are unlikely to occur at scale and that the technology is labour-displacing. We therefore

interpret operator decisions as a conservative proxy for labour demand within the QA task — one that likely captures a lower bound of the true displacement effect, since we do not observe unmet labour demand or situations in which operator availability constrains the number of decisions processed.

Moving to the results, as reported in Table 3.3 we find a statistically significant and economically meaningful negative relation between weekly changes in the AI Automation Rate and the growth rate of operator decisions. The estimated coefficients, ranging from -1.20 in the unweighted specification to -1.16 in the weighted specification, imply that a 1 percentage point increase in the Automation Rate is associated with a 1.16%–1.2% lower growth rate in operator decisions. These results are consistent with the hypothesis that, as the AI is granted increasing autonomy in making QA decisions, the role of human operators diminishes accordingly.

A potential concern is that plants may anticipate the introduction of AI systems and adjust their production numbers in the weeks prior to a deployment. As demonstrated in Appendix 3B, there is endogeneity in deployment timing at the component library level. This could be replicated at the plant level, with reduced productivity in the weeks leading up to a large AI deployment. Thus, we may not be capturing the entirety of task reallocation. A related concern is that plants could use medium- and long-term forecasts to determine AI deployment patterns, thus increases in the Automation Rate may be correlated with anticipated shifts in customer demand.

We are unable to test for medium and long term endogeneity, however we can account for the possibility of short-run endogeneity. To achieve this, we re-estimate Equation 3.1 by adding one- and two- week leads of Δ AI Automation Rate. In other words, we test whether the growth in operator decisions in the current week is correlated with future growth in AI Automation Rate. The results in Table 3.3 show no evidence of such anticipation effects: future AI deployment is not systematically related to current operator decisions. Instead, the observed decline in operator decisions coincides with contemporaneous increases in AI decision shares, suggesting that the detected relation maps to increases in AI deployment.

Table 3.3: Relation between changes in AI Automation Rate and Main Outcomes

	Unweighted	Weighted
Panel A – Δ Total Operator Decisions (Log)		
Δ AI Automation Rate	-1.201**	-1.167**
	(0.413)	(0.396)
Observations	810	810
Anticipation Effect		
Δ AI Automation Rate (Current)	-1.066**	-1.053**
	(0.410)	(0.392)
Δ AI Automation Rate (lead 1 week)	0.420	0.398
	(0.261)	(0.246)
Δ AI Automation Rate (lead 2 weeks)	0.442	0.371
	(0.250)	(0.216)
Observations	800	800
Panel B – Δ Total QA Decisions (Log)		
Δ AI Automation Rate	0.374	0.425
	(0.397)	(0.382)
Observations	810	810
Anticipation Effect		
Δ AI Automation Rate (Current)	0.507	-0.356*
	(0.397)	(0.174)
Δ AI Automation Rate (lead 1 week)	0.384	0.304
	(0.265)	(0.385)
Δ AI Automation Rate (lead 2 weeks)	0.486*	0.124
	(0.256)	(0.226)
Observations	800	800
Panel C – Δ OK Rate		
Total OK Rate		
Δ AI Automation Rate	0.080**	0.076**
	(0.028)	(0.028)
Operator OK Rate		
Δ AI Automation Rate	0.013	0.009
	(0.020)	(0.019)
Observations	810	810

Notes: This table presents two-way fixed effects regressions of weekly changes in labour demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the week-plant level. Column (1) reports unweighted estimates; Column (2) presents results weighted by total weekly decisions. All regressions include plant and week fixed effects. Standard errors are clustered at the plant level. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

3.4.3 Total QA Volume (i.e., Productivity)

We next explore whether changes in AI adoption affect overall plant productivity. We proxy productivity with the growth rate of total QA decisions, estimating Equation (3.1) using Δ Total QA Decisions (Log) as the outcome variable. A limitation with this approach is that total QA decisions may not directly reflect productivity as the volume of total decisions can reflect unobserved characteristics of the AOI System, limiting the number of images seen by the AI or the human operator.

Panel B of Table 3.3 reports the estimation results. The baseline specification shows no statistically significant relationship between changes in AI Automation Rate and total throughput in either the unweighted or weighted specification. This suggests that increasing AI automation neither expands nor contracts the overall volume of QA activity at the week- plant level.

We again test for anticipation of AI deployment by including one- and two-week leads of Δ AI Automation Rate. The results provide no consistent evidence of anticipation effects. While the two-week lead shows marginal significance in the unweighted specification, this relationship disappears in the weighted specification and the one-week lead is insignificant in both. The weighted specification also yields a negative coefficient on the current Δ AI Automation Rate, though this pattern is not robust to the unweighted regression. Overall, we find no systematic relationship between AI deployment and total QA volume, suggesting that automation is associated with reallocation from operators to AI, rather than leading to increase in the number of QA decisions taken.

3.4.4 Total and Operator OK Rate

Next, we turn to the relation between AI Automation and the outcomes of the QA process. While we are unable to detect a relation between AI automation and total QA volume, productivity could be affected by making the QA process more efficient. To test this, we estimate Equation (3.1) using Δ OK Rate as the outcome variable. Panel C of Table 3.3 reports on the estimation results

for both the Overall OK Rate and Operator OK Rate. We detect a small but significant increase in the Total OK Rate of 0.08 percentage points in both the weighted and unweighted specification whilst we are unable to detect any change in the Operator OK Rate.

Interpreting these results, the effect of the increase in the Total OK Rate is somewhat ambiguous. Our data does not allow us to identify a ground truth for the components, therefore an increasing OK Rate could either reflect an increase in true positives or false positives. However, several pieces of evidence help us infer the most likely interpretation. First, operators very rarely overturn OK decisions made by the AI. Using operators as the ground truth, we detect a false positive rate of under 1%, indicating that the AI consistently and accurately detects OK components. Second, interviews with the company revealed that in their own internal assessments AI has a better false positive rate than operators. This leads us to conclude that the detected increase in the Total OK Rate does not represent an increase in defective components being flagged OK (i.e., false positives), but rather a decrease in OK components flagged as defective (i.e., false negatives). A decrease in false negatives increases the efficiency of the QA process, as fewer products are scrapped due to incorrect flagging.

3.5 Robustness Checks

Although our dataset is rich, allowing us to assess the relation between AI Automation and our outcomes, it also contains features that motivate a number of robustness and sensitivity analyses (as described in Section 3.3.1). First, certain plants contribute disproportionately to the sample. In the weighted regressions, the effects may be driven by the plants with many decisions, whilst the unweighted regressions may be affected by plants with low numbers of decisions. Compounding this, some plant-week observations record extreme swings in the AI Automation Rate, as can be seen in Figure 3.6, highlighting both swings of $\pm$ 30 percentage points. Second, as discussed in Section 3.3.2, persistent issues with missing product and decision information may affect our estimates. Third, our main estimation approach is conducted at the plant level.

This allows us to capture any potential spillover effects, however, detected relationships may be driven by compositional effects rather than changes in AI Automation. To address these concerns, we conduct a series of robustness checks to assess the validity of the results presented in the previous subsection.

3.5.1 Leaving-One-Plant-Out

First, to ensure that our findings are not driven by any particular plants, we re-estimate Equation (3.1) leaving out one plant at a time. This allows us to assess whether any single plant drives the estimated relationship between AI adoption and outcomes. Figure A3.2 displays the results from this leave-one-plant-out exercise. The red marker corresponds to the baseline estimate using the full sample, while the grey markers show the coefficients obtained when excluding each plant in turn. We find limited variation in the estimated coefficients across plants. Importantly, no single plant seems to drive the overall relationship between changes in AI Automation and the growth rates of our outcomes. This reinforces the interpretation that the estimated effect reflects a general pattern across plants rather than being specific to any individual production unit.

3.5.2 Sample Restriction

Second, we turn to the effects of our sample selection and issues relating to missing data. Table A3.1 reports on the results from various sample selections, using the weighted specification from Equation (3.1). First, in Column 1, we present the results from our preferred specification as shown in Table 3.3. In Column 2, we remove the largest positive and negative weekly swings of AI Automation. While the relationship with labour demand remains consistent, we observe that the relationship with the overall pass rate is no longer statistically significant and the coefficient of the Operator OK Rate becomes negative, which is consistent with findings at the Component Library level (see Section 3.5.3). In Column 3, we remove weeks where the number of lines missing decisions is at least 25% as large as the total number of decisions. In Column 4, we remove weeks where the number of lines missing product

information is at least 25% as large as the total number of decisions. Column 5 excludes weeks meeting both criteria. The last three columns demonstrate that the identified relationships between our outcomes and AI Automation are robust to weeks with missing information.

3.5.3 Component Library Analysis

To address the concerns relating to compositional effects driving our detected relationships we re-estimate Equation (3.1) at the component library level. Just under 80% of component libraries record under 200 decisions per week, yet this only corresponds to 8% of decisions with component library information. To reduce the effect of small component libraries at the week level, we construct a monthly panel at the component library level. To reduce the effect of shifts in our AI Automation Rate being driven by small library sizes, we remove all component library-month observations where we detect less than 200 decisions corresponding to 2.5% of total decisions.

Table A3.2 reports the results using our final estimation sample, applying the library size restrictions listed above. The results at the component library level are broadly consistent with our plant-level findings. In Panel A, we observe a statistically significant negative relationship between changes in AI Automation Rate and operator decisions, with coefficients of approximately -2.07 in both weighted and unweighted specifications. This is somewhat larger in magnitude than the plant-level estimates, suggesting that the negative effects of automation at the component level are offset by libraries with lower levels of AI Automation. When we include anticipation leads, the primary estimation remains negative and significant. However, we detect a positive relation between future AI deployment predicting current labour demand further representing the endogeneity of AI deployment timing at the component library level (see Appendix 3B for more detail).

In Panel B, examining total QA decisions, we find a small positive relationship at the component library level, contrasting with the null result at the plant level. This suggests that at the component level, AI deployment may be associated with modest throughput increases, though these effects do not

aggregate to significant plant-level productivity gains. Again, this may reflect the fact that plants anticipate AI deployments and increase production of particular products around deployment events. Turning to the anticipation effect in the unweighted specification, we see positive correlations between current throughput and concurrent changes in the AI Automation Rate, alongside the 1 and 2 month leads. In the weighted specification, this finding disappears. As with Table 3.3, the current change in AI Automation Rate is negatively correlated with the change in Total QA Decisions, while there is no detectable relationship with the 1 and 2 month leads. However, due to endogeneity in AI deployment, and a highly unbalanced panel we are unable to detect a robust relationship between QA throughput and changes in the AI Automation Rate.

Lastly, in Panel C, we observe that the Total OK Rate increases by approximately 0.05 percentage points per 1 percentage point increase in AI Automation Rate, consistent with our plant-level findings. We also observe a small negative relationship with the Operator OK Rate, in contrast to our results at the plant level. This is potentially an intuitive result, overall we are unable to detect an effect on the composition at the plant level, however, the AI takes a share of the OK decisions leaving operators with a pool of decisions with a higher defect rate. This indicates that as AI Automation progresses, the operator task may change. However, this association also must be viewed with caution as we are unable to test whether operators change their behavior following an AI deployment, nor whether this is affected by spillovers across products. Taken together, while the component library-level results are broadly consistent with the plant-level findings, some estimates vary across specifications, highlighting the sensitivity of results to compositional effects, weighting choices, and missing data. These patterns suggest that plant-level estimates capture both direct task substitution and compositional shifts across products and components.

3.6 Conclusion

This paper investigates the adoption of AI and its implications for labour demand within a high-tech manufacturing setting. Using detailed decision-level data from a multinational electronics manufacturer, we investigate the deployment of an AI system in a visual inspection task across 14 plants over 15 months. Our analysis reveals some key findings that provide insight into how firms adopt and integrate AI, raising questions about conventional narratives around rapid AI-driven displacement.

First, we document that AI adoption unfolds gradually and incrementally, even long after initial deployment. Rather than one-time technological shocks, plants expand automation progressively along both the extensive margin, by deploying AI on new products and component libraries, and the intensive margin, by adjusting decision thresholds and retraining the algorithm. Notably, growth occurs primarily through the extensive margin, with plants adding AI to new components rather than systematically intensifying its use on already-enabled products. This pattern holds across both early and late adopters, though we observe substantial heterogeneity in adoption levels.

Second, we document a significant negative relation between increases in AI automation and labour demand within the quality assurance task. A one percentage point increase in the AI Automation Rate is associated with approximately a 1.2% decline in the growth rate of operator decisions. Importantly, we find no evidence that plants anticipate AI deployment and adjust labour demand accordingly; rather, substitution occurs once the AI is deployed into the decision-making process. We also observe modest improvements in inspection efficiency, with the overall pass rate increasing slightly as AI automation expands, suggesting fewer false negatives without compromising quality.

What do these findings reveal beyond this specific context? Even in a setting where the technology is rigorously tested, meets strict performance thresholds and is proven capable of autonomous decision-making, adoption does not immediately scale to full potential. Despite the AI system's technical readiness and demonstrated reliability, plants deploy AI cautiously. While our quantitative data cannot fully capture all underlying mechanisms that drive adoption

decisions, qualitative insights from the firm highlight several key constraints. Responsibility for product quality is decentralized at the plant level, and plant managers are accountable for quality outcomes. As a result, adopting or expanding AI is not perceived as a default technological upgrade, but as part of a managerial decision that must be justified. Plants differ in how they evaluate whether AI deployment improves efficiency, reduces waste, or enhances performance relative to existing processes. In settings where managers perceive current quality outcomes as satisfactory, the incentive to expand AI adoption may be weaker, even when the technology is readily available. These governance structures, together with trust in algorithmic decision-making, regulatory and liability concerns, and customer-specific quality requirements act as institutional guardrails. These organizational constraints shape not only whether AI is adopted, but also the pace, scope, and extent of its deployment within production units.

These dynamics have broader implications for understanding AI diffusion in manufacturing environments where reliability, safety, and quality are paramount. Our findings suggest that employers do not simply deploy automating AI systems and displace labour. Instead, such systems hinge on ongoing human oversight, trust building, and sustained organizational commitment. In high-stakes production settings the path from initial deployment to full-scale automation is potentially neither automatic nor immediate. Rather, it is mediated by managerial judgment, institutional safeguards, and the gradual building of confidence in AI's performance.

More broadly, this study underscores that the labour-market effects of AI are shaped not only by technological capability, but by the interaction of technical performance with organizational governance and workplace culture. While AI can reduce labour demand within specific tasks, the timeline and extent of these changes depend both on the technical capability of technologies and how firms manage the transition to algorithmic decision-making. As AI continues to diffuse across industries, understanding these mechanisms will be essential for anticipating its broader economic and social consequences.

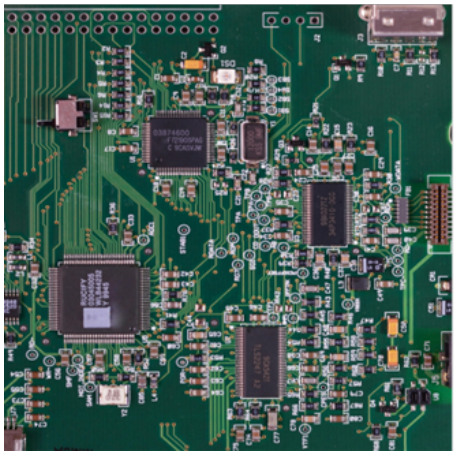
At the same time, AI's labour-market effects extend beyond displacement and substitution. When AI technologies are deployed in ways that complement

workers, such as by enhancing feedback, learning, and training processes, they can generate substantial productivity gains. The next chapter examines these productivity effects by studying how AI-augmented training influences worker performance in a customer contact center.

Appendix 3A: Additional Figures & Tables

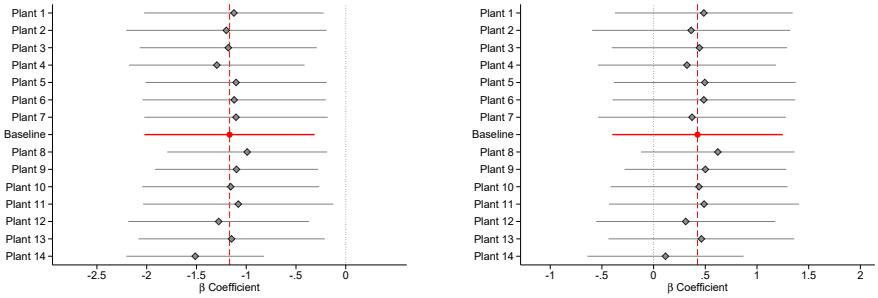


(a) Surface Mount Technology (SMT) Line



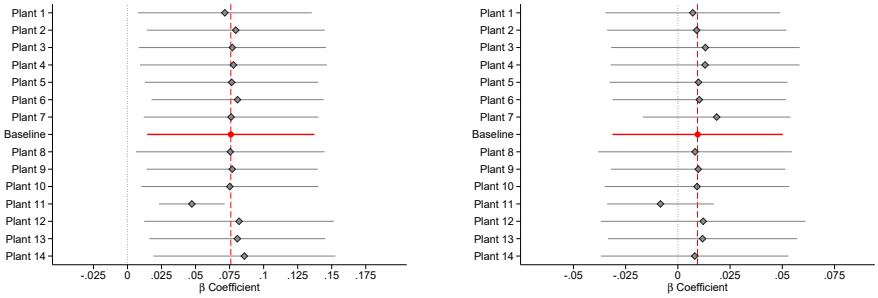
(b) Printed circuit boards (PCBs)

Figure A3.1: Exemplary representations of the production environment.
Source: Author's own depiction.



(a) Δ Total Operator Decisions (Log)

(b) Δ Total QA Decisions (Log)

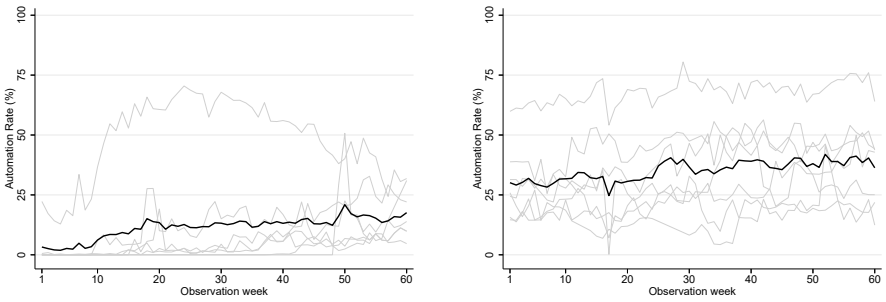


(c) Δ Total OK Rate

(d) Δ Operator OK Rate

Figure A3.2: Leave-One-Plant-Out Robustness Checks

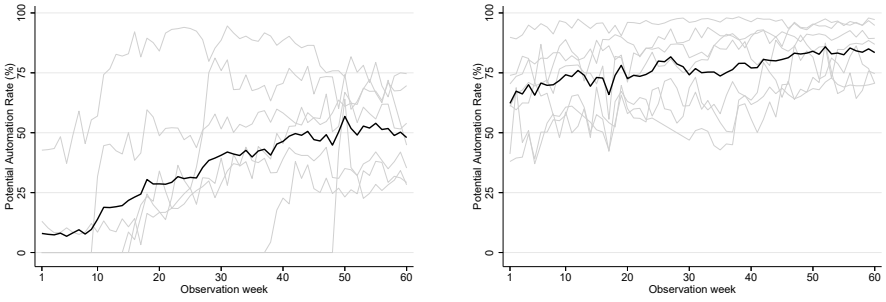
Notes: Each panel shows coefficient estimates when excluding one plant at a time for each model run including 95% confidence intervals, the Y-axis indicates which plant has been excluded. The red dashed line indicates the baseline estimate using all plants, the black dotted line indicates zero effect. All regressions include plant and week fixed effects with standard errors clustered at the plant level.



(a) Late Adopters (b) Early Adopters

Figure A3.3: Plant-Level Automation Rate Heterogeneity by Adopter Group

Note: Individual plant trajectories (thin gray lines) and group mean (thick black line) for automation rates over the 60-week observation period. Panel (a) shows late adopters (N=7 plants, deployed ≤8 months before sample start). Panel (b) shows early adopters (N=7 plants, deployed >8 months before sample start).



(a) Late Adopters (b) Early Adopters

Figure A3.4: Plant-Level Potential Automation Rate Heterogeneity by Adopter Group

*Note:*Individual plant trajectories (thin gray lines) and group mean (thick black line) for potential automation rates over the 60-week observation period. Panel (a) shows late adopters (N=7 plants, deployed ≤8 months before sample start). Panel (b) shows early adopters (N=7 plants, deployed >8 months before sample start).

Table A3.1: Sensitivity Analysis on Main Outcomes

	(1)	(2)	(3)	(4)	(5)
	Preferred Specification	Excl. Outliers	Excl. Missing Decisions	Excl. Missing Products	Excl. Missing Products & Decisions
Panel A – Δ Total Operator Decisions (Log)					
Δ AI Automation Rate	-1.167**	-1.558***	-0.876**	-1.415***	-1.139***
	(0.396)	(0.329)	(0.389)	(0.318)	(0.294)
N	810	728	758	740	702
Panel B – Δ Total QA Decisions (Log)					
Δ AI Automation Rate	0.425	0.085	0.751*	0.188	0.501*
	(0.382)	(0.396)	(0.351)	(0.320)	(0.275)
N	810	728	758	740	702
Panel C – Δ OK Rate					
			Total OK Rate		
Δ AI Automation Rate	0.076**	0.018	0.075**	0.079**	0.078**
	(0.028)	(0.013)	(0.032)	(0.030)	(0.035)
N	810	728	758	740	702
			Operator OK Rate		
Δ AI Automation Rate	0.009	-0.052**	0.007	0.009	0.005
	(0.019)	(0.021)	(0.022)	(0.019)	(0.022)
N	810	728	758	740	702

Notes: This table presents weighted two-way fixed effects regressions of weekly changes in labour demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the week-plant level with various sample restrictions applied. Column (1) reports the weighted estimates reported in Table 3.3; Column (2) removes the top and bottom 5% of changes in AI Automation Rate; Column (3) removes weeks with a high proportion of lines missing decisions; Column (4) removes weeks with a high proportion of missing products; Column (5) removes both. Standard errors are clustered at the plant level. All regressions include plant and week fixed effects. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

Table A3.2: Relation between changes in AI Automation Rate and Main Outcomes using Component Library Panel

	(1) Unweighted	(2) Weighted
Panel A – Δ Total Operator Decisions (Log)		
Δ AI Automation Rate	-2.066***	-2.071***
	(0.067)	(0.068)
Observations	24361	24361
Anticipation Effect		
Δ AI Automation Rate (Current)	-1.920***	-1.936***
	(0.075)	(0.075)
Δ AI Automation Rate (lead 1 month)	0.125***	0.142***
	(0.046)	(0.046)
Δ AI Automation Rate (lead 2 months)	0.206***	0.218***
	(0.043)	(0.043)
Observations	18750	18750
Panel B – Δ Total QA Decisions (Log)		
Δ AI Automation Rate	0.282***	0.277***
	(0.062)	(0.063)
Observations	24598	24598
Anticipation Effect		
Δ AI Automation Rate (Current)	0.373***	-0.190***
	(0.068)	(0.052)
Δ AI Automation Rate (lead 1 month)	0.115***	0.009
	(0.043)	(0.060)
Δ AI Automation Rate (lead 2 months)	0.190***	-0.038
	(0.042)	(0.045)
Observations	18909	18909
Panel C – Δ OK Rate		
Total OK Rate		
Δ AI Automation Rate	0.049***	0.048***
	(0.006)	(0.006)
Observations	24598	24598
Operator OK Rate		
Δ AI Automation Rate	-0.028***	-0.028***
	(0.007)	(0.006)
Observations	24361	24361

Notes: This table presents two-way fixed effects regressions of weekly changes in labour demand (Panel A), total QA throughput (Panel B), and OK Rates (Panel C) on changes in the AI Automation Rate at the month-component library level. All regressions include plant and week fixed effects. Column (1) reports unweighted estimates; Column (2) weights by total weekly decisions. Standard errors are clustered at the plant level. *** ($p < 0.01$); ** ($p < 0.05$); * ($p < 0.1$).

Appendix 3B: Identification Challenges

There are a number of factors that limit causal identification in our setup. At more granular levels than the plant month we encounter violations of the Stable Unit Treatment Value Assumption (SUTVA). Consider deployment at the component library level, where treatment is defined as whether AI has been approved for any product containing a component from that library. Because individual products contain components from multiple libraries, spillovers are inevitable. For example, if a PCB contains both a resistor (with AI deployed) and a diode (without AI deployed), AI adoption for the resistor affects outcomes for the diode as they are produced on the same production line. Operator workload may decrease due to reduced labour demands, but this could be offset by faster production times. They may also face increased cognitive demand due to a higher proportion of defective images, impacting their decision making on the diode. If we assigned treatment at the product level, we would be able to rule out this possibility. However, operators work across multiple lines. As we are unable to identify the lines on which operators work, we are unable to rule out the fact that decisions on untreated products may be affected by treated ones. Imagine a factory with two lines, each producing separate products. One product has the AI enabled, the other does not. An operator may be making decisions on products from both lines, thus the AI automating a proportion of decisions from one line could affect decisions made on products from another.

Aside from the SUTVA violations, we also encounter problems relating to the endogeneity of treatment timing. Plants do not deploy AI randomly across products or component libraries. As described in Section 3.2.5, deployment requires customer approval, trained models with sufficient labeled data, and plant-level confidence in the technology's performance. Critically, plants may strategically time AI deployment to coincide with increases in production volume associated with new orders or high-volume production periods. Figure A3.5 illustrates this problem using the estimator of the De Chaisemartin and D'Haultfœuille (2020) to trace the evolution of total QA decisions around AI deployment at the component library level. Total decisions increase sharply in the period leading up to and immediately after deployment, rising by

approximately 50% in the month of deployment. This spike demonstrates that deployment timing is correlated with production shocks rather than being exogenous. Because we cannot distinguish whether changes in operator decisions reflect genuine task displacement or compositional shifts in the product mix being inspected, we cannot construct a credible counterfactual for labour demand at the product or component library level, even with staggered difference-in-differences estimators designed to handle dynamic treatment effects.

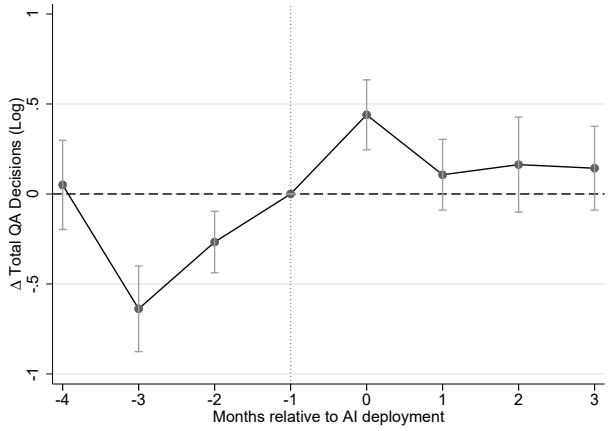


Figure A3.5: Endogeneity of AI Deployment Timing

Notes: This figure shows coefficients from the De Chaisemartin and D'Haultfœuille (2020) estimator at the component library level. The vertical line indicates AI deployment (t=1). Total QA decisions increase sharply following deployment, rising by approximately 50% at t=1. This spike demonstrates that deployment timing coincides with production increases, preventing causal identification at this level of aggregation.

4

How AI-Augmented Training Improves Worker Productivity

Abstract

We analyze the impact of AI-augmented training on worker productivity in a financial services company. The company introduced an AI tool that provides performance feedback on call center agents to guide their training. To estimate causal effects, we exploit the staggered roll out of the AI-tool. The AI-augmented training reduces call handling time by 10 percent. We find larger effects for short-tenured workers because they spend less time putting clients on hold. But the AI-augmented training also improves communication style with relatively stronger effects for long-tenured agents, and we find slightly positive effects on customer satisfaction.

This chapter is joint work with Didier Fouarge, Marie-Christine Fregin, Michael Stops, Simon Janssen, Mark Levels, Raymond Montizaan and Nicholas Rounding. My individual contribution to this paper involved contributing to the development of the research question and methodology, analysis of the data and interpretation of results, conducting qualitative interviews with firm representatives, and contributing to the writing of the manuscript. This study is part of the research project ai:conomics funded by the German Federal Ministry of Labour and Social Affairs (Bundesministerium für Arbeit und Soziales (BMAS)/Denkfabrik Digitale Arbeitsgesellschaft) by resolution of the German Bundestag. We thank the anonymous company for giving us access to the data and allowing to interview employees.

4.1 Introduction

Artificial intelligence (AI) has drawn significant attention for its potential to reshape economic activity. A growing body of research shows the productivity gains from AI, primarily in controlled laboratory settings, [1] but increasingly in real-world environments (Brynjolfsson et al., 2025b; Dillon et al., 2025; Dell'Acqua et al., 2025) Most of these studies focus on AI's direct productivity effects for workers who use general-purpose AI tools—such as ChatGPT and Microsoft Copilot—to support their work tasks. However, AI adoption may also generate indirect productivity effects that extend to workers who do not themselves use AI. For example, AI can improve managerial decision-making, feedback, and training processes, thereby raising the productivity of non-users through better guidance, training, and coordination within organizations (Luo et al., 2021; Riedl and Bogert, 2024).

Particularly, AI-augmented training can be valuable in sectors that handle large volumes of digital data, and face high worker turnover, such as the service industry. A recent survey[2] found that, among HR departments using AI, 43 percent apply it to learning and development. Moreover, companies like MetLife and Zoom are already using AI to improve training for their service agents. With their computational power, scalability, and cost-effectiveness, AI tools can process vast amounts of performance data and generate detailed, data-driven feedback tailored to individual workers' training needs (see, e.g., Council, 2019).

We study how introducing a labour-augmenting AI tool into the training system for customer contact agents impacts worker productivity. The intervention takes place at a major European financial services firm that manages over €500 million in assets and serves more than one million active customers. The company's setup is ideal for estimating causal effects. First, the company conducted a staggered introduction of the AI-augmented training system across geographically separated teams. Second, the company provides rich AI-generated data on worker productivity and call content for all agents in

[1] Choi and Schwarcz (2023); Dell'Acqua et al. (2023); Fogliato et al. (2022); Freeman et al. (2024); Noy and Zhang (2023)

[2] Society for Human Resource Management (SHRM) (2025)

both the treatment and control groups, covering the entire observation period—including the time before the use of the AI-generated data in training. Third, the company requires the agents to participate in the training, thereby allowing us to go beyond intention-to-treat effects and to estimate treatment-on-the-treated effects. Fourth, the company assigned each team a trainer who exclusively trained that team throughout the observation period. Fifth, agents handle only incoming calls, which are always randomly assigned to them.

More specifically, each call center agent is assigned to one of five teams that are located in geographically separated regions. The agents receive one-on-one training sessions of 30 to 60 minutes every two weeks from trainers who are each responsible for one of the five teams. Before using the AI-generated data throughout the training sessions, the trainers had simply selected three calls per agent for review and feedback. First in May and then in July 2023, the trainers from two teams received AI-generated performance data covering all of their agents' calls. While the overall training structure remained unchanged, this data enabled the trainers to provide more individualized and targeted feedback on the agents' productivity, such as handling time—but also quality aspects of the agent-client conversation such as communication style. We exploit this staggered introduction to estimate the causal effect of the AI-augmented training sessions on the agents' productivity, using Callaway and Sant'Anna (2021) event-study approach to account for dynamic and heterogeneous treatment effects.

Our data contains more than $180,000$ calls of approximately 150 agents throughout an entire year. The data includes quantitative information on call quality (e.g., total handle time, speaking time, hold time, and call frequency) and detailed information about the agents' communication style, particularly their usage of diminutives, hedge words, and filler words. In addition, we have detailed information about call topics and customer satisfaction ratings.

We find that the AI-augmented training reduced agents' Average Handle Time (AHT) by approximately 60 seconds, corresponding to a marginal effect of about 9 percent relative to their pre-treatment average. The largest effects—driven by a reduction in very long calls, during which agents put clients on hold to seek help from coworkers or information from the com-

pany's information systems—occurred at the upper end of the handle time distribution. However, we also observe substantial reductions in effective speaking time, with effects occurring more uniformly across the distribution.

As in previous studies (Brynjolfsson et al., 2025b), the productivity gains were significantly larger for short-tenured agents (17 percent) than for long-tenured ones (7 percent), suggesting that the AI-augmented training helped less-experienced workers move up the experience curve more rapidly. However, we also find meaningful effects for agents with longer tenure, although their improvements occurred along different margins of call quality. While long-tenured workers were more likely to reduce their effective speaking time by improving their communication style—using fewer filler words and hedge words—short-tenured workers showed greater improvements in handling previously difficult topics and reducing the duration of putting clients on hold.

We further find that productivity gains go hand in hand with improved customer satisfaction, although these effects are only marginally significant. Because the company's call center is an inbound center where agents cannot influence the frequency of incoming calls, we do not find a positive effect on the average number of handled calls per day. However, we observe large positive effects at the upper end of the distribution (at the 75th and 90th percentiles), suggesting that agents were able to handle more calls in periods with a high number of incoming inquiries.

Our results are robust to the use of other estimators and whether we use a larger sample—including more experienced never-treated agents—or a smaller sample of only not-yet-treated agents, who are likely more comparable to the treated group. Moreover, a specification that relies exclusively on treated and control agents located in geographically distant regions suggests that our results do not suffer from a violation of the stable unit value treatment assumption (SUTVA).

This paper contributes to the large body of research analyzing the productivity effects of technology adoption. While many earlier studies examined the impact of previous IT technologies on worker or company productivity (Bartel et al., 2007; Acemoglu et al., 2007; Bloom et al., 2014), more recent studies focused on the effects of AI on worker performance. This research

includes experiments with professionals performing tasks under controlled laboratory conditions (Peng et al., 2023; Noy and Zhang, 2023; Dell'Acqua et al., 2025; Cui et al., 2025; Agarwal et al., 2023; Choi and Schwarcz, 2023, e.g.,). Moreover, a smaller literature has evaluated the effects of AI adoption on worker productivity in real workplace settings, often exploiting staggered introductions of AI tools or field experiments within companies (e.g., Dillon et al., 2025; Brynjolfsson et al., 2025b; Kanazawa et al., 2025; Otis et al., 2024; Luo et al., 2021).

We contribute to this literature in several ways. First by providing novel evidence on an important but less-studied channel through which AI affects worker productivity: indirect exposure via AI-augmented training. While most studies examine the direct effects of generative AI tools that workers use in their day-to-day tasks, our study shows that non-generative AI tools can improve productivity in a similarly substantial way, even when operating only behind the scenes to enhance training content. If workers do not directly use AI tools in their daily routines, as in our case, they are less likely to suffer from either the cyborg effect—i.e., blindly following AI-generated instructions—or from information overload (Dell'Acqua et al., 2025; Luo et al., 2021).

Second, we contribute to this literature by providing more nuanced evidence on how AI-augmented training affects worker productivity across levels of tenure. For example, Brynjolfsson et al. (2025b) show that while AI particularly benefits younger and less able call center agents, it can slightly hamper performance among the most able workers who followed the AI tool's suggestions. Similarly, Peng et al. (2023) and Cui et al. (2025) find larger productivity gains for short-tenured and older programmers; Kanazawa et al. (2025) for lower-skilled taxi drivers; and Dell'Acqua et al. (2025) for lower-skilled consultants. In contrast, Otis et al. (2024) found that high-performing entrepreneurs benefited more from generative AI than their lower-performing peers.

A key contribution of our analysis is to show that short- and long-tenured agents benefit along different margins. Specifically, AI-augmented training helps less experienced agents by closing substantial knowledge gaps and reducing beginner mistakes. For long-tenured agents, our results suggest that AI-augmented training improves their productivity by enhancing the detail

and precision of their communication style. In this way, the AI-augmented training contributed to a reduction in productivity differences across workers.

Third, our study also relates to the literature on the productivity effects of management practices in general and the effects of worker training in particular. A large body of research has shown that targeted interventions—such as performance monitoring, personalized feedback, or structured training—can lead to substantial productivity gains. For example, Gosnell et al. (2020) show that performance monitoring and the provision of personalized target information significantly improved the productivity of airline pilots. Wheeler et al. (2022) shows that job readiness training for LinkedIn increases employment opportunities for job seekers, and Renée (2025) finds that career counseling and providing information improves the careers of students. Sauermann (2023) and Espinosa and Stanton (2022) document substantial productivity gains from on-the-job training, including spillover effects to coworkers and managers. We extend this literature by showing that training tailored to detailed, individual-specific performance measures can further improve productivity in a more scalable and targeted manner, highlighting the benefits of personalized, adaptive interventions over one-size-fits-all approaches.

The rest of the chapter is structured as follows. Section 4.2 describes the customer service department, and Section 4.3 describes the training program. Section 4.4 explains the identification and empirical strategy. Section 4.5 presents the data and provides summary statistics. Section 4.6 presents the results, and Section 4.7 concludes.

4.2 The Customer Service Department

We evaluate the productivity effects of an AI-augmented training program in the customer service department of a large financial asset manager in Western Europe. In terms of structure and scale, the company is comparable with global peers such as CalPERS (U.S.), Canada Pension Plan Investment Board (CPPIB), and Japan's Government Pension Investment Fund (GPIF). The company generates revenue primarily by charging fees for managing assets above 540 Billion on behalf of more than 3 Million participating customers.

The call center and client service model operated by this company is representative of broader practices among major financial companies for two key reasons: First, like other large financial intermediaries (e.g., pension funds and insurers), the company serves millions of beneficiaries, requiring a centralized and professionalized customer contact infrastructure. The call center handles high volumes of complex inquiries. Second, like its peers across Europe, this company uses service-level benchmarks, call tracking, and customer satisfaction metrics aligned with both internal performance standards and external supervisory expectations. The call center staff are well-trained public-facing professionals, covered by sectoral agreements, reflecting broader European employment practices in finance. The role of the call center is integrated into the broader administrative architecture.

During our observation period from February through October 2023, the department employed 147 agents, who handled a total of $187,839$ calls—an average of $24,000$ calls per agent per month, with an average call length of about 508 seconds. The agents' primary responsibility is to handle incoming customer calls involving questions, problems, or complaints related to the company's financial services. They also respond to emails and perform various administrative tasks. To ensure that customers receive and understand accurate information, the company requires that agents have a solid understanding of its services, efficient strategies for finding relevant information, and strong communication skills. As a result, the position requires a relatively high level of education—typically a vocational or a Bachelor's degree. Continuous changes to the rules governing the financial products offered by the company require agents to regularly update their knowledge to provide accurate and consistent information to clients. Workforce training therefore also plays a critical role. Thus the setting is ideal for the implementation of AI-augmented training. On the one hand, continuous training is essential for improving and maintaining agent productivity, making measures that reduce training costs or increase training effectiveness highly valuable to the company. On the other hand, the company records detailed data on all agent calls, providing a key condition for implementing AI-based solutions.

4.3 The Training Program

Upon joining the department, agents complete a four-week introductory training course that prepares them to handle calls and provides an overview of the most common call topics and related services. Once active in the contact center, all agents regularly receive on-the-job training through one-on-one sessions with a personal trainer. These sessions last between 30 and 60 minutes and occur weekly to monthly, depending on the agent's needs. On average, agents receive training every other week, meaning that they all regularly participate in training. Only the most experienced agents do not participate in the standard training program, instead receiving non-standard training on handling specific challenges they face. Although these agents never participate in the AI-augmented training and are therefore *never treated*, we still have access to AI-generated performance data for all agents for the entire observation period.

The trainers are former agents, each assigned to a dedicated team of call center agents, who provide training exclusively within this team and tailor sessions to each agent's individual needs. The main purpose of the training is to improve agents' overall performance by strengthening their communication skills, refining their problem-solving strategies, and deepening their knowledge of the company's products and processes.

4.3.1 Pre-AI training program

Prior to AI-augmented training, trainers did not have access to structured, computer-generated data on agent performance. Instead, they selected three calls per agent from the period (about a month) following the last training session. Using these calls, they (1) wrote a report highlighting effective or ineffective communication techniques, knowledge gaps, and other problems and (2) discussed these reports with their agents during the one-on-one training sessions.

Trainers selected calls either randomly or by call length. However, if they considered certain calls inappropriate or irrelevant for training, they selected others. This restrictive selection of training calls comes with significant short-

comings. The selected calls often failed to reflect agents' overall call quality, leading trainers to overlook important but less frequent weaknesses in agents' knowledge, performance, and communication style. In accompanying interviews we conducted with trainers and agents along with our quantitative analysis, both groups regularly reported that the selected calls were unrepresentative, resulting in poorly targeted training strategies. In some cases, the selections even caused disagreements between trainers and agents, further complicating the training process.

4.3.2 Post-AI training program

The company chose to improve training by implementing an AI tool that analyzes call center agents' past calls and comprehensively assesses their performance based on all calls. To examine all calls that each agent received, the AI tool relies on a supervised learning natural language processing (NLP) classification algorithm—a rule-constrained, non-generative NLP algorithm that processes call transcripts. The AI tool assigns predefined labels that reflect effective or ineffective communication techniques, the topic(s) of the call, and the duration of speaking, silence, and hold time. As the company defined all categories in advance, the AI tool did not identify patterns on its own (unlike, e.g., ChatGPT). Trainers access the AI-generated data through dashboards and reports in a proprietary software system. They then decide what to share with the agents, who never have direct access to the information.

The accompanying interviews revealed that both trainers and agents valued the ability of the AI tool to detect subtle quality issues that had previously gone unnoticed (Eijkenboom et al., 2025). They also appreciated its ability to confirm or challenge subjective assessments of agent performance, emphasizing the way it supported their feedback with detailed data. Moreover, when trainers showed agents the AI visualizations of their progress across a broader range of call quality dimensions, the agents welcomed this rich set of information. Thus both trainers and agents generally agreed that the AI tool led to improvements in agent productivity.

4.4 Identification and Estimation Strategy

4.4.1 The staggered implementation of the AI-augmented training regime

As previously mentioned, the introduction of the AI tool provides a unique setting for identifying the causal effect of the AI-augmented training on agents' productivity. Four factors make a causal assessment possible. First, the company introduced the AI-augmented training at the team level (five teams in total), with agents remaining in their assigned teams throughout the study period. As a result, they could not choose whether or when to enter the AI-augmented training program, thereby eliminating any self-selection bias and allowing us to go beyond the intention-to-treat effect by estimating the average treatment effect on the treated (ATT).

More specifically, the company incorporated AI-generated performance data into training through a staggered roll-out in May and July of 2023, with the assignment of teams to treatment and control decided earlier, in February 2023 (see Figure 4.1). Moreover, the company stated that it determined the order of the staggered introduction without any specific goal or purpose in mind. Teams one and two received the treatment in May; teams three and four, in July. The agents in team five constitute the control group, which received no AI-augmented training throughout the observation period. However, after our observation period ended, even team five received the treatment in November 2023.

Second, as calls are randomly distributed across agents, call quality is not mechanically related to the dynamics of the agents' productivity. Third, as each trainer is exclusively assigned to a specific team with the whole team consistently assigned to either the treatment or the control group throughout the observation period, spillover effects at the trainer level are unlikely. Fourth, as we have access to different types of control groups, and as teams of agents work in geographically separated regions, we can account for potential spillovers at the agent level. Specifically, agents in teams one and two, who received the May treatment, worked at the same location as the control group

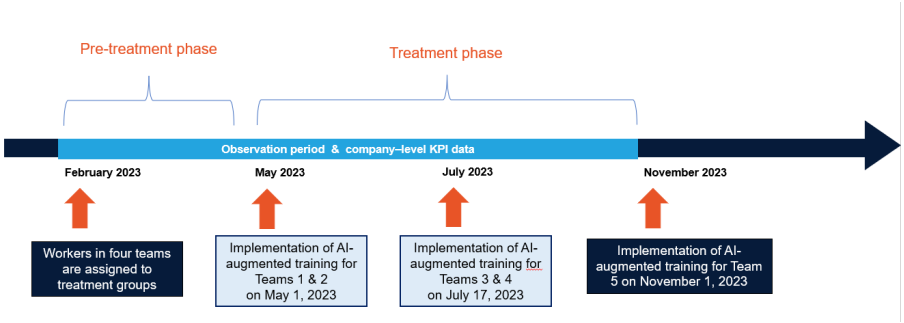


Figure 4.1: Rollout of AI-augmented training across teams

Notes: This figure illustrates the timeline of the field experiment. The observation period spans from week 5/2023 through week 44/2023, during which company-level KPI data was collected. The experiment was implemented in a staggered manner between week 18/2023 and week 44/2023. AI-augmented training was introduced for Teams 1 and 2 on May 1, 2023, followed by Teams 3 and 4 on July 17, 2023. Team 5 received AI-augmented training on November 1, 2023. Only after the technology was introduced in their respective teams did the treated agents begin receiving AI-augmented training.

Source: Authors' own depiction.

(team five). In contrast, the other treatment groups (teams three and four), who received the July treatment, were based at a different, distant location. Therefore, spillovers from the May to the July treatment are very unlikely, whereas spillover effects from the May treatment to the control group might be a concern.

To tackle this concern, we use various control groups available in our data. On the one hand, we rely on a group of very experienced agents who do not participate in the standard training program and are never treated. These agents are very proficient, with deep knowledge about virtually all aspects of their work. Given that they had substantially smaller potential for improvement than less experienced agents, we can plausibly assume that these experienced agents were too proficient to experience substantial spillover effects. On the other hand, we used the not-yet-treated workers from the July treatment as a control group for the May treatment. As these agents worked in a separate location, they were unlikely to be affected by spillovers. As we show in Appendix 4A,

our results remain stable regardless of which control group we use. Thus a violation of the stable unit treatment value assumption (SUTVA) is not likely to be a concern in our setting.

4.4.2 Estimation Strategy

We exploit the staggered introduction of AI-augmented training by estimating the following event study model.

$$y_{it} = \sum_{k=-3}^{3} \mathbf{1}(t = t_{is} + k)\delta_k + \mathbf{x}_{it} + \lambda_t + \theta_i + \epsilon_{it} \tag{4.1}$$

where y_{it} denotes the outcome, capturing different measures of productivity for the agent i at time t. The coefficients δ_k are the main effects of interest and reflect the relative difference in outcomes between treated and non-treated agents. The term t_{is} denotes the period of the implementation of the AI-augmented training for the agent i. The set $k = -3, -2, ..., 2, 3$ defines the time periods spanning from three months before the treatment to four months after. The last entire month without treatment is denoted as -1, and the AI tool was rolled out in period 0. Thus Equation (1) allows us to both analyze dynamic treatment effects and assess potential nonlinear pre-treatment trends in outcomes. $\mathbf{x}_{it}$ denotes a set of control variables that include a set of dummy variables measuring the agents' pre-treatment tenure. The term λ_t denotes a set of time fixed effects, θ_i captures individual agent fixed effects, and e_{it} is a normally distributed error term with zero mean. To account for the potential correlation of error terms within agents, standard errors are clustered at the agent level.

The key identification assumption of Equation 4.1 is that the outcomes of the treatment and the control groups would have evolved in a parallel way in the absence of the treatment, i.e., the introduction of the AI-augmented training program. However, even if this assumption held, classic naive OLS estimates of the presented event studies may fail to recover the ATT if the treatment effects are heterogeneous and dynamic. Instead, these classic estimators yield a

group and variance-weighted average of all possible combinations of treatment effects for which some weights might even be negative (Goodman-Bacon, 2021). Recent studies have provided a variety of approaches for overcoming this concern (Borusyak et al., 2024; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; De Chaisemartin and D'Haultfœuille, 2020).

Therefore, to carry out the majority of our estimates, we use the estimator proposed by Callaway and Sant'Anna (2021). This estimator chops the data into sub-samples of different valid 2x2 difference-in-differences (DiD) comparisons that only include valid combinations of treated, not yet treated, and never treated observations. The estimator then computes a weighted average of the DiD estimates, with weights based on the sample sizes of treated units in each treatment cohort. As a result, Callaway and Sant'Anna (2021) estimator recovers the ATT even when treatment effects are heterogeneous and dynamic. Another key advantage of this estimator is its ability to account for time-constant differences in pre-treatment characteristics through a re-weighting approach. In addition to Callaway and Sant'Anna (2021), we also provide robustness checks using regular two-way fixed effects models and other common estimators by Borusyak et al. (2024), Sun and Abraham (2021), and De Chaisemartin and D'Haultfœuille (2020).

4.5 Data and Summary Statistics

The company gave us access to a broad set of AI-generated performance data that trainers actively use in the AI-augmented training program.[3] A key advantage of our setting is that the firm used the AI tool to generate performance data on all calls—including those handled before the AI was implemented. We thus can track the same performance measures both before and after the performance data was used to train agents, and for all agents—regardless of whether they were ever treated.

The performance measures include quantitative indicators such as agents' call handle time, speaking time, silence time, and hold time. In addition,

[3]For privacy reasons, the data is stored on a protected server. Upon request, we will be able to make the data available for replication upon request.

the data include customer satisfaction, call topics, and nuanced measures of agents' communication styles. All measures are recorded at the call level. Because agents may take leave hours or work irregular schedules, we follow Brynjolfsson et al. (2025b) and aggregate the data at the monthly level to capture more complete work and performance histories.

Our main performance measure is the agent's Average Handle Time (AHT), which includes all parts of a call: (1) speaking time, reflecting communication efficiency; (2) silence time, signaling breakdowns in the conversation; and (3) hold time, when the agent searches for information or consults colleagues. Many comparable studies use AHT as a productivity measure (e.g., Brynjolfsson et al., 2025b; De Grip and Sauermann, 2012; Liu and Batt, 2007), and AHT is also the company's most important key performance indicator. While shorter AHT enables agents to handle more calls per day, in our context, this benefit materializes only on busy days with high call volumes—because we study an inbound call center. More importantly, a shorter handle time gives agents more time to conduct their other responsibilities, such as responding to emails or handling administrative duties.

The company's training explicitly aimed at reducing AHT below 600 seconds, both before and after the AI-generated data was introduced into the training process. Moreover, the company does not collect alternative performance measures, such as the resolution rate used in Brynjolfsson et al. (2025b), because virtually all incoming client queries are eventually resolved either by the agent alone or with support from others.

Our measures for communication style are dummy variables indicating filler words, diminutives, and hedge words. Filler words such as "just" or "um" add redundancy and increase speaking time. Diminutives such as "sec" for "second" or "info" for "information" can lead to misunderstandings, while hedge words such as "maybe" or "I think" signal uncertainty and may undermine customer confidence. These indicators capture more nuanced aspects of call quality than standard metrics (e.g., AHT) or hold time. Even highly capable agents who resolve calls efficiently and rarely place clients on hold may show subtle communication flaws—gaps that can be filled with AI-augmented training.

Moreover, we have detailed AI-generated information on call topics and agents' communication style. The AI tool identifies the main topic of each call (e.g., financial issues related to retirement, a deceased person, or unemployment), allowing us to pinpoint specific areas where individual agents may find challenging or have knowledge deficits. For example, some agents may have long call handle times when discussing death-related issues, while others frequently place clients on hold when handling received pension payments.

In addition to these outcome measures, the data includes information on teams and their operating locations, the type of contract, the training program (i.e., standard or non-standard in case of experienced workers), and an indicator for tenure. Unfortunately, due to the company's strict data protection requirements, we are only allowed to use a measure of tenure that distinguishes among agents with 0–4 months, 5–12 months, and more than 12 months of tenure. Moreover, we have no information on the agent's education, gender, or other personal data.

Table 4.1 summarizes the available individual and monthly observations in the sample. Over the observation period, the agents handled more than $180,000$ calls, and we observed a total of 147 agents: 84 are in the treated teams (teams 1 through 4), while 63 are in team 5, the control group. Overall, the data contains 957 person-month observations—440 in the treatment group and 517 in the control group. The last row shows that about 40 percent of the agents are experienced workers who were never treated because they were not part of the standard training program. All of these experienced workers were part of the regular control group in team 5.

Given their proficiency, these agents are on a flatter experience profile than short-tenured agents, making them an ideal control group for capturing unobserved common trends, such as seasonal shifts in call volume. However, one potential concern is that our estimates might be biased precisely because these agents follow different trends than less-tenured agents in the treatment group. To eliminate this concern, most of our specifications account for agents' pre-treatment tenure through reweighting, and we also report specifications that rely exclusively on not-yet-treated agents as controls.

Table 4.2 presents summary statistics for all agents in the baseline period, i.e.,

Table 4.1: Sample Characteristics

	(1) All	(2) Treatment Group	(3) Control Group
Number of calls	187,839	79,935	107,904
Number of agents	147	84	63
Share of agents with non-standard training	.41	-	.72
Number of monthly observations	957	440	517

Notes: This table reports summary statistics for all agents in the sample. Column 1 presents statistics for the full sample across all time periods. Column 2 presents summary statistics for treated agents who had AI-augmented training. Column 3 presents statistics for control agents who did not receive AI-augmented training at any point during the observation window.

one month before the treatment. The first three columns show results for the May treatment; the following columns for the July treatment. On average, agents in the treatment group have longer average handle time, speaking time, silence time, and hold time than those in the control group (Columns 1, 2, 4, and 5). They are also more likely to use filler words, diminutives, and hedge words. These differences are hardly surprising because the control group agents are more experienced than those in the treatment group. Although balancing pre-treatment characteristics is not required for identifying causal effects in a DiD framework, most of our specifications account for differences in agents' pre-treatment tenure through reweighting. Columns 3 and 6 show that applying the Callaway and Sant'Anna (2021) weights substantially improves balance across most characteristics.

Table 4.2: Agent and Call Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	May Treatment			**July Treatment**		
	Treatment	Control		Treatment	Control	
		Raw	Weighted		Raw	Weighted
Agent Tenure (Share of Agents):						
0–4 Months	0.455	0.279	0.469	0.455	0.082	0.500
5–11 Months	0.485	0.209	0.469	0.273	0.164	0.154
12+ Months	0.061	0.512	0.062	0.273	0.754	0.346
Average Call Characteristics (in seconds):						
Average Handle Time (AHT)	651.227	494.623	550.074	620.170	442.742	593.632
Speaking Time	386.723	311.014	320.166	352.495	302.949	380.276
Silence Time	69.421	53.628	58.664	64.393	47.134	60.932
Hold Time	195.082	129.981	171.244	203.282	92.659	152.423
Ineffective Communication Styles (Share of Calls):						
Filler Words	0.534	0.439	0.455	0.564	0.530	0.604
Diminutives	0.399	0.338	0.341	0.342	0.335	0.435
Hedge Words	0.386	0.376	0.419	0.484	0.358	0.430

Notes: This table reports sample statistics for the May and July treatment groups and their corresponding control groups. Columns 1 and 4 show means for treated agents; Columns 2 and 5, means for control agents. Columns 3 and 6 apply Callaway and Sant'Anna (2021) reweighting to adjust for observable differences.

4.6 Results

4.6.1 Main productivity effects on Average Handle Time (AHT)

Figure 4.2 presents estimates from a classic two-way fixed effects event study based on Equation (4.1), where the outcome variable is our main productivity measure, the AHT of calls. The horizontal axis displays the months relative to the implementation of the AI-augmented training, while the vertical axis shows the conditional difference in AHT before and after implementation. The figure reports results from two specifications of the regression model.

Gray triangles display the weekly estimates, with gray lines indicating their corresponding 95 percent confidence intervals. Black dots show the monthly estimates, with capped spikes representing their confidence intervals. Both models include agent and time fixed effects, and standard errors are clustered at the individual level.

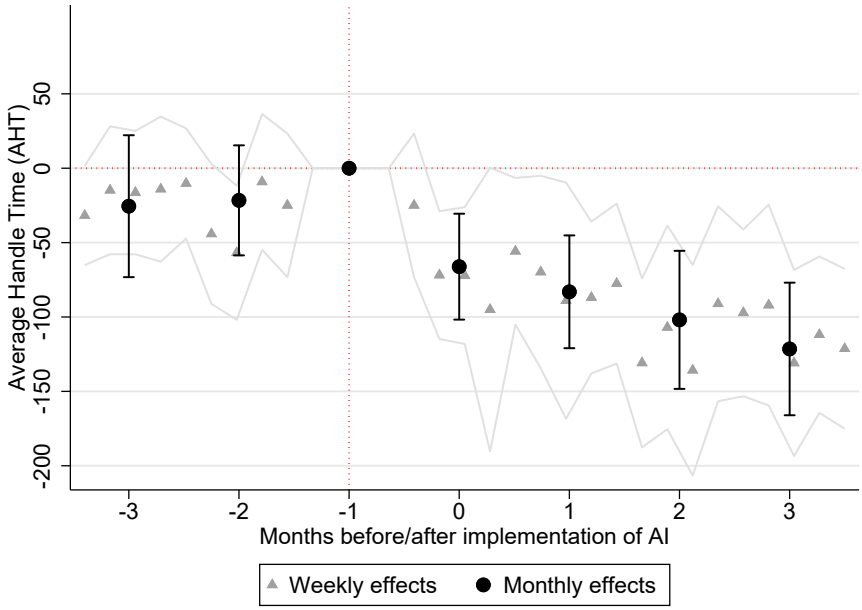


Figure 4.2: Event study estimates on average call handle time

Notes: The figure shows event study estimates of the AI-augmented training on workers' average call handle time (AHT). The x-axis represents the months relative to the implementation of the AI, while the y-axis shows the conditional gap in AHT before and after the implementation. The gray triangles represent weekly coefficient estimates, with gray lines showing their corresponding 95 percent confidence intervals. The black dots represent monthly coefficient estimates accompanied by capped spikes illustrating the estimates' 95 percent confidence intervals. All models include agent and month fixed effects. Standard errors are clustered at the agent level. The monthly effects are estimated based on 957 employee-month observations; the weekly effects, on 3,348 employee-week observations.

Our main identification assumption is that average outcomes for the treatment and control groups would have followed parallel trends in the absence of treat-

ment. Although we cannot test this assumption directly, Figure 4.2 provides support for its validity. Between treated and non-treated agents, we do not observe any pre-treatment AHT differences that are statistically different from zero at conventional significance levels. The weekly estimates in particular offer a more granular view of the pre-treatment period and—despite somewhat larger heterogeneity—show pre-treatment differences that are mostly close to zero.

In response to the treatment, a clear negative gap in AHT emerges between agents in the treatment and control groups—at both weekly and monthly levels. In the treatment month, the effect amounts to about -60 seconds, gradually increasing to about -120 seconds. The first column of Table 4.3 shows that the average effect of this specification is about -68 seconds, which corresponds to a marginal effect of -11 percent relative to the agents' pre-treatment AHT of 636 seconds.

As mentioned in Section 4.4.2, classic two-way fixed effects (TWFE) models may be biased because they rely on forbidden comparisons between already treated and not-yet-treated observations. To overcome this bias, Figure 4.3 presents estimation results based on the approach proposed by Callaway and Sant'Anna (2021) for each treatment separately. The estimation approach relies on both never treated and not-yet-treated agents as controls.

Both the May and July treatments yield qualitatively similar effects to the naive event study in Figure 4.2. However, the TWFE version of the Callaway and Sant'Anna (2021) estimator, reported in Column 2 of Table 4.3, shows a slightly larger negative average effect (approximately -92 seconds, or -14 percent).[4]

Unlike those in the treatment group, many agents in the control group are, on average, more experienced and, because of this, did not participate in the training program during our observation period (though they did participate in regular training earlier in their careers). Comparing these long-tenured

[4]OLS event study models with event-time dummies are often more similar to the estimator proposed by Callaway and Sant'Anna (2021) than the TWFE specification with a single post-treatment indicator. Event-time dummies allow treatment effects to vary over time and can partially mitigate bias from inappropriate comparisons across cohorts at different stages of treatment.

Table 4.3: Treatment effect on Average Handle Time (AHT)

	(1)	(2)	(3)	(4)	(5)	(6)
Treated x Post AI	-68.183***	-91.639***	-60.841***	-65.215***	-57.399***	-64.836***
	(17.364)	(18.618)	(20.342)	(20.854)	(16.925)	(18.569)
OLS TWFE	Yes	No	No	No	No	No
Callaway-Sant'Anna (CS)	No	Yes	Yes	Yes	Yes	Yes
Agent Tenure	No	No	Yes	Yes	Yes	Yes
Agents with non-standard training excluded	No	No	No	Yes	No	Yes
Entire Period	No	No	No	No	Yes	Yes
Pre-treatment average	636	636	636	636	589	589
Observations	957	956	956	579	764	402

Notes: This table presents the results of using the TWFE-OLS and robust DiD estimators introduced in Callaway and Sant'Anna (2021) on our main measure of productivity: Average Handle Time (AHT, in seconds). Columns 1 & 2 only include agent and month-fixed effects. Columns $3-6$ also include agent-tenure fixed effects. Columns 4-6 use alternative sample specifications, excluding both agents with non-standard training and workers who join or leave the department during our observation period. All standard errors are clustered at the call agent level. *** $p < 0.01$.

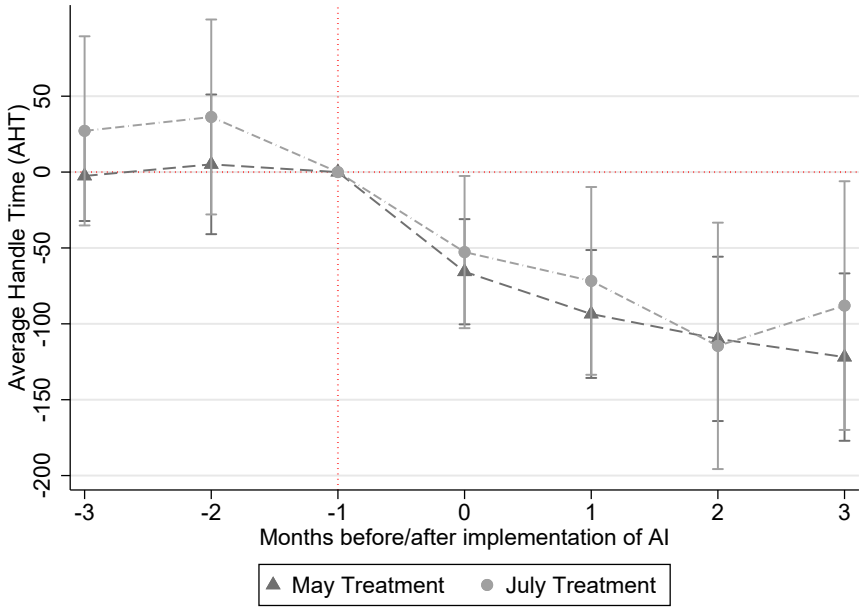


Figure 4.3: Event study estimates of Average Handle Time (AHT), May & July treatments

Notes: This figure presents an event study using the Callaway and Sant'Anna (2021) method to evaluate the effect of the implementation of AI-augmented training on Average Handle Time (AHT). The x-axis indicates months before and after AI implementation, and the y-axis shows changes in AHT (in seconds). Triangles represent coefficient estimates for the May treatment group; circles represent estimates for the July group. Vertical capped spikes indicate 95% confidence intervals. Each estimate includes agent and month fixed effects, and standard errors are clustered at the agent level.

agents (who are likely to be on a flat learning curve) to the less-tenured agents in the treatment group (who are likely to be on a steeper learning curve) might bias the coefficient estimate on AHT downward. Therefore, Column 3 of Table 4.3 controls for pre-treatment tenure and yields an estimated effect of about -60 seconds, or -9 percent. Overall, the specification of Column (3) is our preferred specification, yielding an effect size that closely matches that of Brynjolfsson et al. (2025b), who find that the direct introduction of

AI assistance in customer service reduces agents' chat time by 8.5 percent. Therefore, our results indicate that using AI insights to inform training—rather than intervening directly in real time—may be an equally effective approach to improving worker productivity.

Specifications 4 throughout 6 provide further robustness checks by excluding experienced never-treated agents who did not participate in the standard training program (Columns 4 and 6), and by excluding very short-tenured agents who were not present throughout the entire observation period (Columns 5 and 6).[5] All three specifications lead to quantitatively similar effects. Appendix 4A analyzes a potential violation of SUTVA by restricting the sample to the treatment group of the May treatment and using only the not-yet-treated agents of the July treatment as controls. This specification is unlikely to suffer from spillover effects because the affected teams are located in distant sites where interaction between agents hardly exists. Appendix 4A comfortingly shows that the results do not change. Appendix 4B shows that the most commonly used alternative approaches for correcting for bias in classic TWFE models yield virtually the same results.

4.6.2 Distributional effects at different productivity margins

A large literature has analyzed the distributional effects of modern technologies, particularly IT (Autor et al., 1998, 2003; Goos et al., 2014). That literature shows that while earlier technologies such as IT primarily replaced the routine tasks of less skilled workers, they complemented the non-routine tasks of more skilled workers. In contrast, the main advantage of AI lies in its ability to process large volumes of data to generate useful information on individual problems. In our case, the AI tool analyzed massive amounts of call data to deliver highly individualized performance feedback and to identify personal gaps in process and product knowledge. Consistent with this idea, interviews with both agents and trainers highlight two main channels through which the AI tool may have improved agents' productivity. First, the AI-generated feedback may have helped agents to refine their communication style. Second,

[5] Excluding these agents restricts the sample to only those agents with at least three months of tenure.

it may have supported trainers in more accurately identifying topic-specific knowledge gaps during training sessions.

In the following, we analyze the productivity effects of AI-augmented training across the entire outcome distribution and along three additional productivity margins—speaking time, hold time, and silence time. To analyze distributional effects, we transform the dependent variables using re-centered influence functions (RIF), as proposed by Firpo et al. (2009). In contrast to raw quantiles, RIFs are linear in expectation and can therefore be used as dependent variables in a regression framework analogous to the Callaway and Sant'Anna (2021) approach. The resulting coefficient estimates capture unconditional distributional effects, indicating whether the outcome distribution has shifted at different percentiles in response to the treatment. All regressions account for tenure, as well as individual agent and time fixed effects. Figure 4.4 presents the results.

The upper panel on the left side of the figure analyzes AHT—our main outcome measure. The y-axis shows the effect sizes of the coefficient estimates, as estimated using the Callaway and Sant'Anna (2021) version of Equation (4.1). The first coefficient estimate (black dot and dashed line) presents the average ATT, as already reported in Column 3 of Table 4.3.[6] The further coefficient estimates (gray markers) present the results for the 10^{th} throughout the 90^{th} percentiles.

According to the upper left panel of Figure 4.4 the AI-augmented training has a substantially larger impact at the upper tail of the AHT distribution than at the lower tail. While we find no effect at the lowest percentile, we observe large reductions of more than 160 seconds at the upper end of the distribution. Although the effects in the middle of the distribution are non-negligible, the AI-augmented training appears to have primarily reduced the occurrence of person-month observations with a very long AHT. These calls are most likely handled by agents who struggle throughout the conversation due to substantial knowledge gaps. In contrast, shorter calls are more likely handled by more skilled agents, who have less room for improvement.

[6]The same applies to the further panels of Figure 4.4. Appendix 4C presents results for the average effects for the other outcomes, each based on a specification that excludes all never-treated workers from the control group.

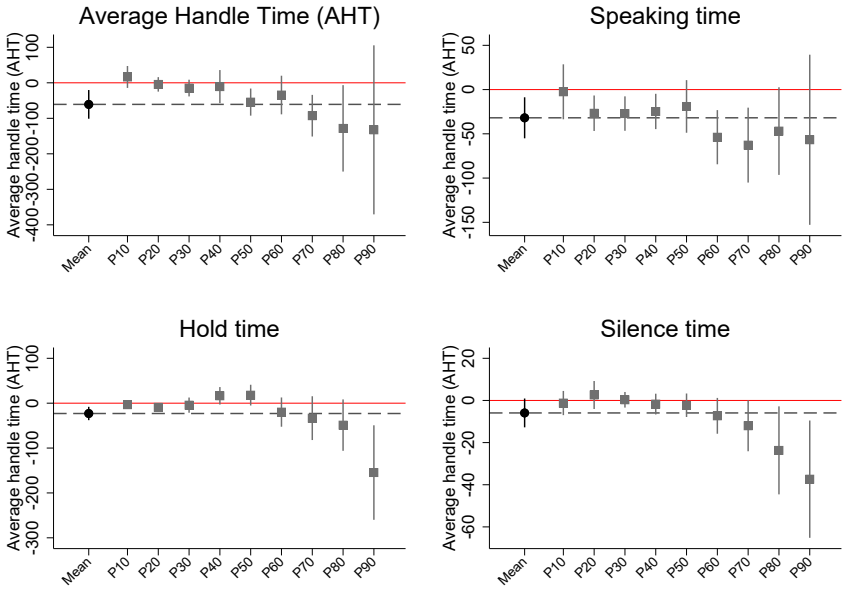


Figure 4.4: Effects on AHT and its elements at different percentiles of the distribution

Notes: This figure presents the results from using robust DiD estimators introduced in Callaway and Sant'Anna (2021) The black dots on the left side represent the average effects. The gray markers indicate the effects on various percentiles from Recentered Influence Function (RIF) regressions at different percentiles of the outcome distribution. The dependent variables are Average Handle Time (AHT) in seconds, speaking time, hold time, and silence time (all measured in seconds per call). All models include agent and month fixed effects and account for pre-treatment tenure categories through reweighting. Standard errors are clustered at the agent level.

The upper panel on the right side of Figure 4.4 analyzes agents' effective speaking time, which reflects how efficiently agents communicate with clients and is therefore likely related to their communication style. Yet longer speaking time may also indicate a lack of topic-specific knowledge, because agents might need more time to clarify issues in the dialogue with the clients. The average ATT on speaking time is around 32 seconds (black dot and dashed line)—half the size of the average effect on AHT. In contrast to the AHT results, the effects on speaking time are more uniform across the distribution. Although we do

not find any effect at the 10^{th} percentile, the remaining effects cluster around -50 seconds, with only a slight tendency for stronger effects at the upper end of the distribution. Thus the AI-augmented training program shifts almost the entire speaking time distribution to the left, i.e., a broader improvement in communication efficiency, not only among the lowest-performing agents.

The lower panel on the left analyzes agents' hold time, which captures the time they put clients on hold while requesting help from coworkers or searching for additional information. Thus hold time mostly reflects gaps in their product or procedural knowledge. The lower panel on the right examines silence time during calls, indicating a complete breakdown in the conversation between agent and client, and likely reflecting beginner mistakes. The average effect on hold time amounts to around -23 seconds, thereby explaining a large fraction of the overall effect on AHT. However, a substantial reduction in outliers with very long average hold times in a given month appears to be driving this effect.

The ATT on silence time, around -6 seconds, is not statistically significant at conventional levels, contributing very little to the overall effect on AHT. However, as with hold time, we find stronger and statistically significant effects in the upper tail of the distribution, suggesting that outlier conversations marked by complete breakdowns are driving this result.

Overall, the results suggest that the AI-augmented training improved agent productivity through two main channels. The first entails shortening excessively long calls—primarily by reducing hold and silence time, likely among agents with knowledge gaps. The second entails increasing communication efficiency, as reflected in reduced speaking time, likely affecting both less and more able workers. However, in contrast to earlier waves of IT technology, the AI-augmented training appears to benefit agents at the lower end of the experience distribution more than those at the upper end.

4.6.3 Heterogeneous Treatment Effects by Tenure

As previously mentioned, the impact of technological change on workers with different skill levels is of central interest. As short-tenured workers have

had less time to accumulate skills and human capital than their long-tenured peers, tenure is a useful indicator—available in our data—for analyzing the heterogeneous impact of AI along the skill distribution. Moreover, studies have shown that generative AI can help move short-tenured workers up the experience curve, while its effect on longer-tenured workers has been much weaker in comparison (see, e.g., Brynjolfsson et al., 2025b)

Table 4.4 presents the productivity results for agents in different tenure categories. The first column reports the results for agents with fewer than four months of tenure, while the second column shows the results for those with more than four months of tenure.[7] While the new training reduced the AHT of short-tenured agents by approximately 131 seconds (-17 percent relative to their average pre-treatment AHT), it reduced the AHT of long-tenured agents by only approximately 41 seconds (7 percent). Therefore, the AI-augmented training had a substantially larger effect on short-tenured agents than on long-tenured ones. We find a similar tenure gap in agents' speaking time (15 percent for short-tenured vs. 9 percent for long-tenured agents). The tenure gaps for hold time (19 percent vs. 5 percent) and silence time (17 percent vs. 6 percent) are even more pronounced, and for long-tenured agents, the effects are not statistically significant at conventional levels.

Thus, the largest relative difference between short- and long-tenured agents appears for hold time. This result is consistent with the idea that less-skilled workers, who are more likely to have knowledge gaps, need to ask trainers or coworkers for help, whereas more skilled (i.e., longer-tenured) workers may benefit relatively more from nuanced improvements. Such improvements might include communication style refinements, which could be reflected in reduced speaking time.

4.6.4 Knowledge gaps and communication style

This subsection analyzes how the AI-augmented training affects agents differently, depending on their experience level, by focusing on differences in

[7] Our limited number of observations prevents us from providing a more granular investigation of the tenure effects.

Table 4.4: Heterogeneous treatment effects on time spent during call

Experience group	(1) *0–4 Months Tenure*	(2) *5+ Months Tenure*
	Average Handle time (AHT)	
Treated x Post AI	-130.92**	-40.87**
	(40.18)	(20.24)
Pre-treatment average	777	555
Observations	199	757
	Speaking Time	
Treated × Post AI	-65.68***	-29.46**
	(22.54)	(12.98)
Pre-treatment average	430	335
Observations	199	757
	Hold Time	
Treated × Post AI	-51.641***	-7.99
	(18.644)	(7.34)
Pre-treatment average	266	161
Observations	199	757
	Silence Time	
Treated × Post AI	-13.596**	-3.41
	(5.917)	(2.36)
Pre-treatment average	81	59
Observations	199	757

Notes: This table presents the results of using robust DiD estimators introduced in Callaway and Sant'Anna (2021) and applied to disaggregated measures of AHT: speaking time, hold time, and silence time all measured in seconds. Columns compare agents with short tenure (0–4 months, Column 1) and long tenure (5+ months, Column 2). All regressions include agent and month fixed effects. Standard errors are clustered at the agent level. *** $p < 0.01$; ** $p < 0.05$.

knowledge gaps and communication style. Table 4.5 analyzes how the AI-augmented training influences the agents' communication style. To do so, we rely on three dummy variables measuring whether the agents have used filler words, diminutives, or hedge words in their communication (see Section 4.5 for details). All of these variables indicate nuances of poor communication styles.

The first row of Table 4.5 shows how the AI-augmented training affected agents' use of filler words. Although we find no statistically significant effects for short-tenured agents, we find large effects for long-tenured ones—a 5 percentage points (9 percent) reduction in the use of filler words. Similarly, in the second row, we find a 3 percentage points (10 percent) reduction in the use of diminutives among long-tenured agents, and, in the third row, a significant reduction in the use of hedge words; roughly 4 percentage points (10percent). All other coefficient estimates in the first three rows are statistically insignificant at conventional levels.

However, given that effect sizes and pre-treatment averages are very similar for short- and long-tenured agents, we cannot definitely conclude from these results that the communication style of long-tenured agents did not improve in response to the AI-augmented training. The absence of statistical significance among short-tenured agents might instead reflect greater effect heterogeneity or limited statistical power within this smaller sample.

Thus to further explore the effects of the AI-augmented training on agents' communication style, the last row of Table 4.5 analyzes a version of speaking time—residualized on the communication style dummies—as the dependent variable.[8] This outcome captures speaking time variation that is unrelated to our communication style indicators. After residualization, while the treatment effect on speaking time decreases by a third for short-tenured agents and is statistically significant at conventional levels, it drops by more than 50 percent for long-tenured agents, with the coefficient turning insignificant at conventional levels (see Table 4.4 for comparison).

[8] In Appendix 4D we show that poor communication styles significantly increase agents' speaking time, while being much less correlated with hold or silence time. Besides, we also show a negative effect on customer satisfaction.

Table 4.5: Heterogeneous treatment effects on communication styles

Experience group	(1) *0–4 Months Tenure*	(2) *5+ Months Tenure*
	Filler Words	
Treated × Post AI	-0.051	-0.044**
	(0.033)	(0.018)
Pre-treatment average	0.6	0.52
Observations	199	757
	Diminutives	
Treated × Post AI	-0.026	-0.030*
	(0.034)	(0.017)
Pre-treatment average	0.45	0.33
Observations	199	757
	Hedge Words	
Treated × Post AI	-0.049	-0.044***
	(0.040)	(0.015)
Pre-treatment average	0.45	0.42
Observations	199	757
	Speaking Time (resid)	
Treated × Post AI	-45.54**	-11.03
	(18.70)	(11.73)
Pre-treatment average	430	355
Observations	199	757

Notes: This table presents the results using robust difference-in-differences estimators introduced in Callaway and Sant'Anna (2021) applied to communication style metrics: the share of calls containing filler words, diminutives, or hedge words. Each outcome is estimated separately for agents with short tenure (0–4 months, Column 1) and long tenure (5+ months, Column 2). All regressions include agent and month fixed effects. Standard errors are clustered at the agent level. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

These results suggest that changes in observable communication style indicators largely drive the reduction in speaking time among long-tenured agents. In contrast, the smaller decline for short-tenured agents after residualization indicates that their improvements in speaking time are less closely tied to the specific communication style measures we capture—possibly reflecting broader gains such as generally more efficient calls or faster recall of product knowledge.

4.6.5 Call topics

In the following, we analyze whether the AI-augmented training specifically improves agents' performance—particularly that of short-tenured ones—on topics which they had previously found challenging. Since all calls are randomly assigned, agents must handle a wide range of customer inquiries, some of which are more technically complex or emotionally demanding than others. For example, agents may need to speak with a client's relatives about the client's death or discuss a client's recent inability to work. We therefore identify these topics discussed during the calls. Table 4.6, which presents the key characteristics of each topic, shows that inquiries about unemployment are, on average, associated with significantly longer durations across all call dimensions—from AHT to silence time. In contrast, inquiries about pension payments tend to take much less time across nearly all dimensions of the call.

Moreover, certain topics occur much more frequently, allowing agents to acquire more knowledge through on-the-job experience, whereas other topics arise so rarely that agents have little opportunity to learn about them. For example, inquiries related to *unemployment* appear in only 1 percent of calls, whereas inquiries about *retirement* and *pensions received* occur in 31 percent of them. This may be also reflected in the average handle times: an average call on *unemployment* reveals the longest handle time of 898 seconds whereas a call on *pensions received* needs only 487 seconds.

To identify topics that agents found challenging before the AI-augmented training, trainers had to rely on selected calls or agent feedback. After the AI augmentation, trainers could more efficiently find calls related to those topics

Table 4.6: Call characteristics by topic of calls

Call Topic	(1) Share of Total Calls (%)	(2) Average Handle Time (AHT)	(3) Speaking Time	(4) Hold Time	(5) Silence Time
Death	24	632	408	157	66
Disability	6	735	458	200	77
Cohabitation or marriage	3	726	474	176	76
Moving	5	533	347	125	61
Pension accrual	19	682	445	165	71
Receiving pension	31	487	313	119	55
Retirement	31	682	447	161	74
Separation or divorce	13	660	418	174	68
Unemployment	1	898	573	231	94

Notes: This table presents average call characteristics (in seconds) by topic of the call. "Share of Total Calls" reflects the proportion of all calls attributed to each topic.

by filtering dashboards and reports. Using this information, a trainer could then discuss with each agent why a particular topic might be challenging and develop a personalized development plan for mitigating the problem. Therefore, this AI-augmented information now allows trainers to more precisely target problematic topics, discuss them with the agent, and tailor support accordingly.

We now analyze how the AI-augmented training improved call performance on topics agents had previously found challenging and examine how these effects vary by agent tenure. To do so, we leverage the more granular call-level data by first residualizing pre-treatment call duration for each call, controlling for tenure, week, day of the week, time of day, and agent fixed effects. This approach removes unobserved factors related to both agents' average ability and detailed temporal dynamics. We then compute the average pre-treatment residual per topic for each agent and identify the top two topics with the highest average residuals as the agent's most challenging topics.

Panel A in Table 4.7 presents the estimates based on Equation (4.1) for AHT and other outcomes, using only calls related to each agent's most challenging topics (Panel A), compared to all other topics (Panel B). The first four rows of Panels A and B report the results for short-tenured agents; the second four rows, for long-tenured agents.

These results suggest that AI-augmented training is particularly effective at narrowing knowledge gaps, especially for short-tenured agents and in topics they found challenging earlier. The sharper reductions in AHT and improvements in call components for these agents indicate that the training helped close basic skill gaps. However, we do not observe that long-tenured agents benefited from AI-augmented training.

4.6.6 Calls per day and customer satisfaction

Although AHT is the company's most important productivity measure, other dimensions of agent performance exist. On one hand, a shorter call handle time allows agents to manage more calls per day, thereby boosting overall productivity. On the other hand, the AI-augmented training might have

Table 4.7: Heterogeneous treatment effects on different topics by tenure

	(1) Average Handle Time (AHT)	(2) Speaking Time	(3) Hold Time	(4) Silence Time
		Panel A: Challenging Topics		
		0–4 Months Tenure		
Treated x Post AI	-148.69***	-71.08***	-60.75**	-16.86***
	(46.57)	(24.63)	(27.77)	(4.95)
Pre-treatment average	943	529	314	99
Observations	131	131	131	131
		5+ Months Tenure		
Treated x Post AI	-87.33	-64.49	-17.87	-4.97
	(57.89)	(50.97)	(15.24)	(6.27)
Pre-treatment average	757	470	210	77
Observations	751	751	751	751
		Panel B: Normal Topics		
		0–4 Months Tenure		
Treated x Post AI	-63.28	-24.51	-31.70	-7.07
	(38.66)	(16.55)	(20.39)	(6.20)
Pre-treatment average	620	348	205	67
Observations	131	131	131	131
		5+ Months Tenure		
Treated x Post AI	-28.56	-16.80	-9.63	-2.14
	(22.08)	(11.05)	(11.05)	(2.94)
Pre-treatment average	528	313	159	56
Observations	752	752	752	752

Notes: This table presents the results using robust difference-in-differences estimators introduced in Callaway and Sant'Anna (2021) applied on different topics, i.e., topics agents found challenging vs normal. Each outcome, measured in seconds, is estimated separately for agents with short tenure (0–4 months) and long tenure (5+ months). All regressions include agent and month fixed effects. Standard errors are clustered at the agent level.*** $p < 0.01$; ** $p < 0.05$.

increased pressure on them to shorten interactions at the expense of service quality. For a more comprehensive assessment, this subsection examines two additional outcome measures: the number of calls handled per day and customer satisfaction scores.

The first row of Table 4.8 presents coefficient estimates of the effect on the average number of calls handled per day. The first column of this first row reports the effect on the mean, while the next five columns show the effects at different percentiles of the distribution. The last two columns show the effect on the average number of calls separated for short- and long-tenured agents.

Although the average effects, for all agents and separated for short- and long-tenured agents, are not statistically different from zero (Columns 1, 7 and 8), the table reveals strong positive effects at the upper end of the distribution (Columns 5 and 6).[9] These results suggest that the AI-augmented training shifted the upper tail of the distribution—indicating that a few agents were able to handle exceptionally high call volumes during certain months—rather than affecting typical or median performance. Given that the company operates as an inbound call center and that agents cannot influence the number of incoming calls, these gains likely reflect improvements in efficiency or call handling speed, rather than changes in call availability. These results suggest that the training enabled some agents to handle more calls during particularly busy days or months while potentially freeing time for other responsibilities only during periods of lower demand.

The second row of Table 4.8 presents the results on customer satisfaction. Larger values of customer satisfaction indicate larger satisfaction. Because clients often decline to rate call agents—and are more likely to do so only when they are very satisfied or dissatisfied—only about 7 percent of all calls receive customer satisfaction ratings. Such a low response rate is common in inbound call centers. Moreover, as most agents in our data receive at least one customer satisfaction rating per month, we lose only a small number of person-month observations in our sample.

[9] The low average effect may be attributable to the weakly significant, negative effect for the lower end of the distribution. Evidently, after the training, agents with an extremely low number of calls had an even lower number.

Table 4.8: Treatment effects on customer satisfaction and calls per day

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mean			*Percentiles*			*0-4 Months Tenure*	*5+ Months Tenure*
		10	25	50	75	90		
Calls per Day	0.69	-5.18*	-3.53	0.31	6.23**	5.29**	0.79	1.99
	(1.35)	(2.81)	(2.33)	(1.95)	(2.48)	(2.06)	(1.55)	(1.54)
Pre-treatment average	22.19						24.1	20.91
Observations	956	956	956	956	956	956	199	757
Customer Satisfaction	0.44*	0.27	0.27	0.41**	0.21	0.56**	0.46**	-0.30
	(0.24)	(0.54)	(0.31)	(0.21)	(0.19)	(0.25)	(0.21)	(0.21)
Pre-treatment average	8.23						8	8.39
Observations	738	738	738	738	738	738	146	500

Notes: This table presents results on Calls per Day and Customer Satisfaction using robust difference-in-differences estimators introduced in Callaway and Sant'Anna (2021). Columns labeled 10 to 90 report results from Recentered Influence Function (RIF) regressions at the 10th, 25th, 50th, 75th, and 90th percentiles, respectively. The final two columns show treatment effects by employee tenure: 0–4 months and 5+ months. All models include agent and month fixed effects and account for pre-treatment tenure categories through reweighting. Standard errors are clustered at the agent level. monitoring systems. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

On average, customer satisfaction increased by about 0.4 points (corresponding to about 5 percent). The effect is only marginally significant at the 10 percent level (Column 1), but it is large and statistically stronger in the upper tail of the distribution, suggesting that agents received more exceptionally positive ratings.

This effect is mainly driven by calls that short-tenured agents handled; for these agents, the customer satisfaction significantly increased after the AI-augmented training (Column 7), whereas the effect on customer satisfaction related to calls that long-tenured agents handled revealed no significant changes.

4.7 Conclusion

This paper provides causal evidence on the productivity effects of AI-augmented training in the real-world setting of a customer service provider. We evaluate the introduction of an AI tool that delivered individual-level performance feedback to support the training of customer contact agents. Exploiting the staggered roll-out of this tool, we show that AI-augmented training reduced call handling time by approximately 10 percent. While the particularly large gains were for short-tenured agents, who improved by closing knowledge gaps and reducing avoidable mistakes, more experienced agents also benefited from the training through improvements in communication style and efficiency. Importantly, we find that the gains did not come at the expense of service quality: indeed, customer satisfaction ratings slightly improved, and agents were able to handle more calls on particularly busy months or days.

Our results contribute to a growing literature on the labour market effects of AI adoption, emphasizing an important but understudied channel: the role of AI in augmenting training, as opposed to augmenting only day-to-day task execution. While recent studies focus primarily on direct interactions between workers and AI tools, our results show that non-generative AI systems operating in the background can also meaningfully enhance performance, particularly when used for supporting targeted and individualized training.

Our results also speak to broader questions about the distributional consequences of AI. Although short-tenured agents benefit most in absolute terms, we find that both less and more experienced workers improve, albeit along different dimensions, i.e., more efficient call handling for short-tenured workers especially, and improved communication style for long-tenured workers. In this way, AI-augmented training may reduce productivity gaps across workers, rather than widening them.

Beyond our specific setting, the mechanisms we identify are likely transferable to other data-rich environments where performance data can inform personalized training. These environments include manufacturing, where machine-level data and incident reports support operator learning; education, where learning analytics help teachers tailor instruction; healthcare, where clinician-patient interactions inform communication training; and logistics or retail, where process data can optimize workflows and service. As key drivers of impact—structured feedback, targeted skill development, and improved learning efficiency— are common across sectors, AI tools can also improve training systems. Gains might be particularly strong for low-skilled or short-tenured workers in high-turnover roles, where faster onboarding matters most.

Notably, our findings suggest that AI-augmented training can be particularly beneficial for young labour market entrants, helping them close skill gaps and accelerate learning during the critical, and early stages of their careers. Yet despite these potential benefits, how do prospective workers themselves perceive AI's impact on their future job prospects? And what policies do they demand in response to AI-driven labour market changes? The next chapter addresses these questions by examining the policy preferences of young prospective workers in response to perceived labour market implications of AI.

Appendix 4A: Potential violation of SUTVA

As mentioned in Section 4.4, the agents affected by the May intervention work in the same location as the members of the control group. As a result, spillover effects from the May treatment group to the control group may be a concern. To address this concern, Figure A4.1 presents results from a specification that compares the May treatment group to all not-yet-treated agents who are affected by the treatment in July and are located in a geographically distant region. Because the second treatment occurred in July, we can estimate effects only for the first (May) and second (June) months after treatment. Nonetheless, both the original and this geographically restricted comparison yield quantitatively and qualitatively similar results.

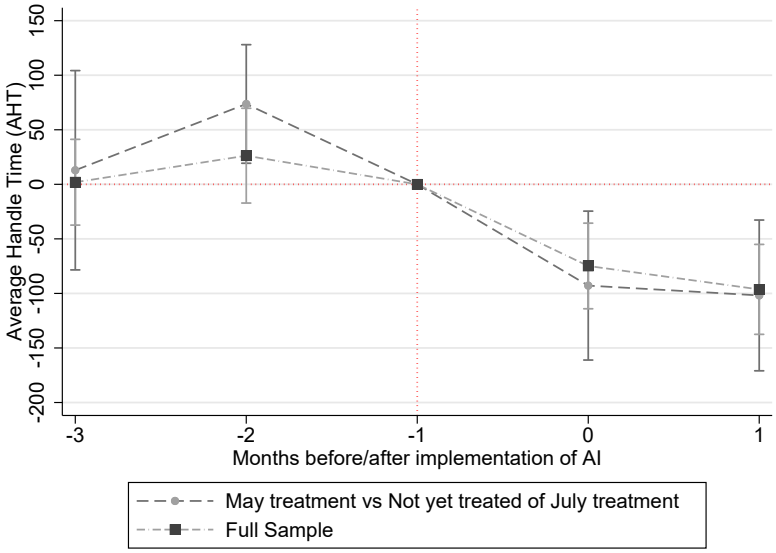


Figure A4.1: Event study estimates for spillover effects

Notes: This figure presents an event study using the Callaway and Sant'Anna (2021) method to evaluate the effect of the implementation of AI-augmented training on Average Handle Time (AHT). The x-axis indicates months before and after AI implementation, and the y-axis shows changes in AHT. They gray dots represent the coefficient estimates for the May treatment using only not-yet-treated agents from the July treatment as controls. The black squares represent coefficients for the May treatment using all agents as controls. Vertical capped spikes indicate 95% confidence intervals. Each estimate includes agent and month fixed effects, and standard errors are clustered at the agent level.

Appendix 4B: Alternative approaches to overcome the bias in classical TWFE estimation

Figure A4.2 presents estimation results of our preferred model Callaway and Sant'Anna (2021) in comparison with three other common approaches to account for potential estimation biases in classical TWFE models. Sun and Abraham (2021), De Chaisemartin and D'Haultfœuille (2020), and Borusyak et al. (2024). All results show quantitatively and qualitatively similar results.

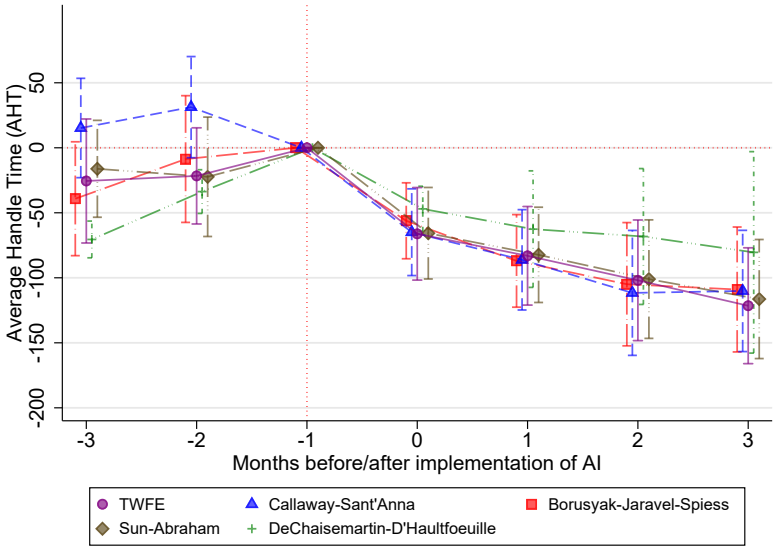


Figure A4.2: Event study estimate, Average Hangle Time (AHT), different model estimators

Notes: The figure compares different event study methods that correct biases from heterogeneous and dynamic treatment effects. The x-axis represents the months relative to the implementation of the AI, while the y-axis shows the conditional gap in AHT before and after the implementation. Purple represents the classical TWFE approach, blue represents the approach by Callaway and Sant'Anna (2021), green the one by Borusyak et al. (2024), red the one by De Chaisemartin and D'Haultfœuille (2020), and black the one by Sun and Abraham (2021). All capped spikes indicate 95 percent confidence intervals. All models include agent and month fixed effects. Standard errors are clustered at the agent level. The classical TWFE approach and Sun and Abraham (2021) rely on 957 observations while the remaining approaches exclude invalid comparisons, such that Callaway and Sant'Anna (2021) relies on 956, Borusyak et al. (2024) on 914, and De Chaisemartin and D'Haultfœuille (2020) on 664.

Appendix 4C: Treatment effects on time spent for different activities during the call

Table A4.1 shows the ATT for speaking, hold, and silence time. The upper panel shows results from a specification includes all agents from the control group. The lower panel shows results from a specification that excludes agents

with non-standard coaching from the control group.

Table A4.1: Treatment effects on time spent for different activities during the call

	(1) Speaking Time	(2) Hold Time	(3) Silence Time
Panel A: Whole Sample			
Treated x Post AI	-31.911**	-23.007***	-5.928
	(12.077)	(7.688)	(3.399)
Observations	956	956	956
Panel B: Never Treated Agents with Non-Standard Coaching Excluded			
Treated x Post AI	-35.317***	-24.190***	-5.708*
	(12.547)	(8.071)	(3.450)
Observations	579	579	579
Pre-treatment average	369	199	66

Notes: This table presents the results using robust difference-in-differences estimators introduced in Callaway and Sant'Anna (2021) on our disaggregated measures of AHT: speak time, hold time and silence times. All columns account for pre-treatment tenure through reweighing and include month and agent fixed effects. All standard errors are clustered at the agent level. ** $p < 0.05$, *** $p < 0.01$.

Appendix 4D: Use of communication styles and call length

Table A4.2 examines how our communication style metrics relate to AHT, speaking time, hold time, silence time, and customer satisfaction ratings using OLS. These metrics correlate strongly with all outcomes and explain the most variance in speaking time.

Table A4.2: Use of communication styles and call length

	(1) AHT	(2) Speaking Time	(3) Hold Time	(4) Silence Time	(5) Customer Satisfaction
Filler Words	202.26***	174.88***	25.56	1.82	0.51***
	(31.88)	(14.37)	(20.86)	(3.97)	(0.18)
Diminutives	340.23***	135.19***	154.21***	50.83***	-0.44**
	(32.81)	(14.79)	(21.47)	(4.09)	(0.19)
Hedge Words	360.41***	168.44***	169.70***	22.27***	-0.01
	(29.43)	(13.27)	(19.26)	(3.67)	(0.18)
Observations	1,159	1,159	1,159	1,159	801
Adjusted R^2	0.42	0.51	0.20	0.24	0.01

Notes: Each column reports coefficients from separate regressions of the listed outcome variable on communication-style indicators. ** $p < 0.05$; *** $p < 0.01$.

5

Generating Anxiety, but not Policy Demands? AI Competition and the Policy Preferences of Tomorrow's Workforce

Abstract

Developments in Artificial Intelligence (AI) have the potential to alter the demand for certain tasks and jobs, and thereby change the occupational structure of labour markets. How does the prospect of competition with technology affect policy preferences of young, educated prospective workers? In this study, we conduct a pre-registered online survey experiment where we randomly assign university students to manipulations about the potential of generative AI models and their implications for labour market prospects. Our objective is to assess the impact of (perceived) AI competition on preferences towards regulatory policies for these technologies. Our results indicate that the preferences of young people do not align with those of job incumbents. While the latter tend to demand more protection through welfare policies, prospective workers do demand for a regulation of AI but not through the welfare policies traditionally associated with labour protection. Policy demands regarding welfare appear to instead originate from pre-existing ideological positions. This occurs despite an increased anxiety about the job market resulting from the potential competition against artificial intelligence.

This chapter is joint work with Toon Van Overbeke, Michele Fenzl, Luana Russo and Darian Meacham. My individual contribution to this paper involved participating in joint discussions on the research questions, contributing to the development of the study design and survey instrument, administering the survey and collecting the data, performing the data analysis, and contributing to the writing of the manuscript. This study is based on a pre-registered experiment. The anonymous pre-registration with a pre-analysis plan (PAP) is available on the Open Science Framework (OSF) through this link.

5.1 Introduction

Over recent decades, automation and advanced robotics have been reshaping labour markets across affluent democracies (Autor and Dorn, 2013; Acemoglu and Restrepo, 2020a; Eurofound, 2018). The rapid advancement of Artificial Intelligence (AI), exemplified by models like ChatGPT and Gemini, further intensifies concerns about imminent labour market disruptions, sparking broad discussions about the benefits and risks of AI for the job market (Cedefop, 2025). Unlike previous waves of technological change, AI threatens not only routine manual work but also high-skilled knowledge jobs once considered insulated from automation.[1] While the future impact of this trend remains subject to debate (Yu and Qi, 2024), recent estimates suggest that up to 60% of workers are in high-exposure occupations in advanced economies (Cazzaniga et al., 2024; Acemoglu et al., 2022).

Given the high risks of AI and the widespread narratives around it, political scientists and political economists have started mapping the impact of automation on political behaviour and the policy preferences of labour market incumbents in rich democracies (Rueda and Thewissen, 2019; Im, 2020; Gallego and Kurer, 2022; Kurer and Häusermann, 2022; Gallego et al., 2022; Anelli et al., 2023; Magistro et al., 2024; Ahrens, 2024; Im, 2025). Their findings report that technologically-induced income and status anxiety are important drivers for far-right support (Gallego and Kurer, 2022; Anelli et al., 2023) as well as demands for expanded welfare state provisions among at-risk voters in the labour market (Rueda and Thewissen, 2019; Im, 2020; Kurer and Häusermann, 2022). Interestingly, Haslberger et al. (2025) found that, in an experimental setting, even if exposure to AI was not necessarily triggering risk perceptions in working age individuals, it generally generated a demand for stronger social policies.

What about those citizens who do not yet have a job? Recent emerging evidence suggests that AI adoption may disproportionately affect employment prospects for early-career workers (Brynjolfsson et al., 2025a; Klein Teeselink, 2025), although this interpretation remains contested (Iscenko and Curto

[1] For more information, see Chapter 1.

Millet, 2026). Nevertheless, despite the profound implications for younger generations–who are often even more exposed to both the potential of AI tools and discussions around them–existing studies offer limited insights into their perceptions and political responses to these structural changes. Given that policy preferences developed in the formative years of young adulthood may translate into long-term attitudes and behaviour[2], understanding how young students respond to the policy challenges posed by AI is crucial.

Would young citizens yet to enter the job market espouse the same views as older job incumbents? This seemingly simple question is everything but obvious. Previous literature in political science highlights differences in civic and political attitudes between generations (Dalton, 2020; Milkman, 2017; Twenge, 2010; Twenge et al., 2012). Similarly, scholars report age differences in vote choice (Geys et al., 2022; Wagner and Kritzinger, 2012, see, e.g.,). Considering these studies, the assumption that the preferences of (older) job-incumbents could simply align with those of younger citizens (who are yet to enter the job market) appears less trivial. This is especially true when considering preferences that result from changes in the job market caused by technological breakthroughs. Young citizens often exhibit greater confidence in new technologies and their ability to employ them in the workplace. It is, therefore, plausible that young citizens may develop preferences that are different from those of job incumbents about the policy instruments that should mitigate AI risks for the job market.

More specifically, there are contrasting conclusions that we can derive–theoretically–from previous studies. One set of studies indeed reports that young people, and especially the educated ones, show positive attitudes towards new technologies (Morris and Viswanath, 2000; de Koning and Gelderblom, 2006; Aubert et al., 2006; Beckmann, 2007; Behaghel and Greenan, 2007). Drawing from this literature, we may expect the young to demand *weaker* mitigation policies. This is especially true for university students, i.e., young people who have acquired a strong set of skills to navigate technological changes. Other studies, instead, suggest that young adults–and students in particular–significantly worry about the devaluation of their skills (Freeman,

[2]See, e.g., Dinas (2012) and the literature on habitual voting.

2025). Others find mixed evidence. For instance, CDT (2025) show that 41% of U.S. young adults believing AI will have both negative and positive effects on their future.

Building on the results of these studies, it becomes clear that we certainly cannot assume that the preferences of job incumbents would necessarily align with those of the (educated) young. To contribute to this question, we developed an online survey experiment with UK and US university students below the age of 30. We developed hypotheses drawn from previous studies about the response of job incumbents to AI threats and political economy models, who assume that citizens act to protect their material interests. However, our results do not show a systematic shift in policy preferences as a response to exposure to AI risks and to narratives about those risks. Although, as expected, individuals become more anxious about future job prospects, they, surprisingly, *do not* demand stronger risk-mitigating policies. In other words, the policy preferences of young students to address AI challenges do not correspond to those of job incumbents (see, e.g., Kurer, 2020; Gallego et al., 2022; Anelli et al., 2023; Ahrens, 2024). Our results indicate that demand for social policies are rather shaped by pre-existing political beliefs–a result not granted given the young age of our participants. Finally, our results show that while students do not demand *traditional* social policies that could mitigate job-market risks; they still prefer governments to slow down the development of AI. Taken together, these results suggest the voters of tomorrow might expect governments to develop strategies outside traditional paths to minimise the distruptive force of technological change on their lives.

5.2 Young Citizens, Their Preferences, and a Changing Labour Market

This study aims to understand how potentially disruptive AI innovations influence public policy preferences young citizens. Over the past decades automation technologies such as advanced robotics have drastically altered labour markets in rich democracies (Autor and Dorn, 2013; Autor, 2020), with

a polarizing effect. Recent advances in generative AI (such as ChatGPT or Gemini), meanwhile, promise to shake up these dynamics even further. Some studies have estimated that as much as 60% of all jobs will be affected by cutting-edge innovations such as ChatGPT (Cazzaniga et al., 2024).

These disruptive forms of technological change have generated a wide scholarship in political science and political economy. On the one hand, it is argued that citizens faced with technological risks shift to the left and make increased redistributive demands on the state (Rueda and Thewissen, 2019; Kurer and Häusermann, 2022; Im, 2020). In particular, risk-exposed voters are found to support spending on traditional welfare policies such as unemployment benefits (Kurer and Häusermann, 2022), but not on the implementation of more radical measures such as Universal Basic Income (Sacchi et al., 2020). On the other hand, technological risk has been predominantly associated with voter support for the far right (Kurer, 2020; Gallego et al., 2022; Anelli et al., 2023), leading to the hypothesis that technological change may also trigger a perception of social status risk. As Kurer (2020) argues, perceptions of relative decline among so-called survivors of automation–more than actual impoverishment among those who lose work–best explain support for the far right. A logic that Borwein et al. (2024) suspects is even amplified when the threat involves society as a whole and not an individual voter *per se*. This conclusion is also consistent with the findings of Ehret (2022) who shows that citizens are more likely to demand government intervention when they perceive a risk of unemployment resulting from AI.

Existing work has a strong empirical and theoretical emphasis on employed voters. This emphasis is understandable. Not only are employed citizens most immediately affected by disruptive innovations but many studies concentrated on the risk exposure of specific occupational categories (see e.g. Frey and Osborne, 2017a; Kurer and Gallego, 2019), making job incumbents the target population of interest. At the same time, this focus leaves a theoretical and empirical blind-spot for important groups, like students, who are not (yet) in the labour market. Indeed, it is not obvious that findings derived from employed voters can be straightforwardly generalized to young people who have not yet entered the labour market. Much of the existing literature links policy preferences to individuals' current occupational exposure to automa-

tion risk, realized labour market vulnerability, or status anxiety. By definition, students do not possess a well-defined labour market position to empart a either realized exposure or a crystallised social status. At most, they can form expectations about future risk. This distinction matters because the mechanisms identified in prior work—such as retrospective economic insecurity, income loss, or direct workplace disruption—may not operate in the same way for individuals without current employment. Moreover, young people differ systematically from the general working-age population in terms of income, asset ownership, family status, and institutional attachment to the labour market, all of which are known correlates of redistributional and labour market policy preferences. To the extent that existing findings are driven by these characteristics, their applicability to young cohorts is limited.

So, what about young citizens who are not yet employed? When we consider the possible response to AI of the young (and even more the educated young), several, yet contrasting, expectations can be developed. On the one hand, previous studies highlight that young people, and students in particular, are quicker to adopt emerging digital technologies such as AI (Morris and Viswanath, 2000; de Koning and Gelderblom, 2006). University students can be especially welcoming of innovation because, as a result of the skills they acquire through education, they are well-equipped to benefit from technological transitions (Aubert et al., 2006; Beckmann, 2007; Behaghel and Greenan, 2007). If university students think strategically about the future, they may develop even a preference for these technological disruptions–to gain a comparative advantage when compared to older job incumbents–and resist government intervention to slow down the process of innovation or its impact on the job market.

On the other hand, citizens rarely formulate political preferences based on complete information (Jones, 1999; Shin, 2018), a dynamic that becomes particularly relevant given the fast-moving nature of technological risks (Ahrens, 2024). In the context of their significant investment in time and resources, students may overlook the potential advantages of AI and become more anxious about the job market–and more uncertain about how their skills will lead them to a secure stable employment. This is particularly true in a context where technological advancements are always matched with narratives about

the benefits or risks of AI. This combination of perceived vulnerability and uncertainty can, in turn, foster a demand for greater government protection, aligning the policy preferences of young people with those of job holders (Sverke et al., 2002). Given that university students are a group even more exposed than the average to narratives about AI and are more likely, than other social groups, to be exposed to its use and potential, they might become even more careful about it than job incumbents. This conclusion is also consistent with the findings of Heinrich and Witko (2025), who observed that people with more knowledge–and especially knowledge about the potential effect of AI–tend to demand more government protection and stronger welfare policies.

In other words, it would be, in principle, possible to formulate hypotheses in either direction: Young citizens may see AI as an opportunity and demand a push towards technological change *matched with less* welfare protection; or they may focus on the risks of AI, and align their position with job incumbents–demanding *more* welfare protection. Since we consider that students, among young people, have heavily invested in education, we formulate the *a priori* expectation that they may overweight the risks and underweight the opportunities (see, e.g. Ahrens, 2024). In other words, we side with the literature that focuses on the protection of material interests (e.g. Gallego and Kurer, 2022), especially considering the high level of education of university students.

H1 *Exposure to AI risk increases demand for risk mitigating policies.*

In aligning our expectations with this political economy literature, we also consider that exposure to AI and its risks never occurs in a framing void. The positive or negative potential effect is now widely discussed in advanced economies. Among all social groups, university students are particularly exposed to discussions around AI. Previous research has highlighted that knowledge and the framing of AI are particularly relevant for the formation of policy attitudes (Ehret, 2022; Heinrich and Witko, 2025). We, therefore, distinguish between the potential effect of AI when the latter is presented as an opportunity (for instance as a tool that will enhance productivity) or as a risk (that may increase the chances of unemployment or worsen employment

conditions).

H2 *The effects for HI will be stronger when AI exposure is matched with narratives that underline the risks, and not the opportunities, of AI.*

In addition, we consider that policy demands can build on pre-existing ideological preferences. Governments can moderate the negative effects of AI through different strategies. These strategies can correspond to more left- or right-leaning approaches to the reduction of the negative effects of technological change on employment. While we expect that, overall, university students will demand more protection, we also hypothesise that *the type of policy that will be demanded* will map based on individuals' ideology even if these orientations are still consolidating (Neundorf and Smets, 2017).

H3 *Exposure to AI risk increases demands for risk mitigating policies that align with individuals' ideological orientation.*[3]

To summarise, while we recognise that there can be different expectations on how young, educated individuals may react to the challenge of AI for the labour market, we see particular merits in the literature highlighting the role of material interests protection. We believe that this literature should apply even more strongly to a group of people who are highly educated. We also expect narratives to matter, as well as preexisting ideological biases: framing AI more negatively should push people to demand for protection; and one's ideology should guide the choice of the specific recipe demanded to the government.

[3]We preregistered this hypothesis dividing between left-wing policies (spending on compensation for laid off workers, basic income, and/or taxing corporations) as HIa; and right-wing ones (active labor market programs, cutting immigration, paying companies to retain workers, and/or restricting international competition) as HIb in the pre-registration.

5.3 Experimental Design, Data and Methodology

5.3.1 Experimental Design

To assess the effect of exposure to AI risks and narratives about AI on the development of policy demands in young people, we fielded an online survey experiment in October 2025. The experiment involved 1328 UK and US university students below the age of 30 recruited on Prolific Academic.[4] Each participant was randomly assigned to one of the three conditions (with equal probability of assignment): (1) a control condition, (2) a treatment condition where they competed with AI in task *and* were informed about its potentially adverse effects on the labour market, or (3) a treatment condition in which they competed with AI in a task *and* received positive news on the complementary effects of AI on the job market.[5]

Students in the control group was asked to read an article about pollution in UK rivers (see 5B for more information).[6] Subsequently, participants were given up to 5 minutes to write a short summary of the article (max 200 words).[7] During the task, a virtual timer was visible to participants to help them keep track of their time. We later used the time students spent on the task to foster the idea of competition with the AI output. The objectives of both tasks (reading the article and providing a summary) was to expose the control group to a cognitive task similar to the treatment groups, without prompting considerations about AI.[8]

[4]The geographic focus on the UK and US means that the sample represents a students living in knowledge-intensive economies where AI is a salient issue. The comparatively limited welfare states of these two countries means we selected most-likely cases where citizens could express demands for policies to mitigate AI risks.

[5]We deliberately did not design a group receiving no information about AI as such as group could be theoretically interesting but not realistic. In real world scenarios, individuals are not just exposed to AI but also hear news and opinions about it.

[6]The newspaper article is from The Guardian and can also be found online here.

[7]We pre-tested the time participants need to complete the summary using a 150-person sample on Prolific. In this test, the median completion time for the summary was 4.5 minutes. Since our experimental sample of students is likely to be more familiar with this type of tasks, we concluded that we could reasonably expect completion within 5 minutes.

[8]They were then offered two cushioning questions aimed at making the task of reading the article more purposeful, otherwise participants may have thought of it as artificial in case of no follow-up questions. This artificiality may have induced demand effects.

Participants in the negative-information treatment started by reading the same article on pollution in UK rivers. Mirroring the control group, participants were asked to write a short summary of the article (max 200 words in up to 5 minutes). The participants were then confronted with a competing summary composed by ChatGPT. They were informed that the alternative text was written by an AI model and that the summary was achieved in 2 iterations and 7 seconds (using real information). This prompt was aimed at explicitly inducing a comparison between the participants' completion time and that (greatly shorter) of AI. The participants were then asked to compare the two summaries and rate their work *versus* to the one generated by AI, prompting–by design–a sense of competition.[9] To further reinforce the idea of competition between AI and labour, we then showed participants a screenshot from a real smartphone app displaying the use of AI in the role of 'lawyer', 'therapist', 'writer', 'language teacher', 'advertiser', and even 'doctor', among others (See Figure A5.3a for the presented image before the negative information on AI). The participants then read an excerpt of an IMF piece highlighting that AI will adversely affect around half of all jobs and likely increase inequality. The article stressed that high-skilled occupations are at risk and that policy can mitigate these risks.[10].

Participants in the positive-information treatment completed the same assignment, but was shown a screenshot of AI assistants that are less likely to be seen as competitors for future jobs (such as essay writer, e-mail writer, or math assistant, among others) (See Figure A5.3b for the presented image before the positive information on AI). Following this screen, they were asked to read excerpts from the same IMF article offered to the second treatment group, but they read selected paragraphs that highlighted the potential of AI to boost

[9]We recognize that the task of summarizing may be more worrying for students in some fields than in others that rely less on writing tasks. We, therefore, pre-registered a possible interaction between the participants' field of study and assignment to treatment. To delve even deeper, we also asked participants to rate how much they find summarizing information to be a pleasant or tedious task, and how relevant similar tasks may be for their prospective jobs. This is because those who perceive the task to be tedious might be more willing to accept AI assistance, especially if they believe that their future jobs will mainly demand other skills. See Table A5.4 for the heterogeneity analysis.

[10]The article can be retrieved Online and is included in the Annex. It is important to note that the article mentions the source of the information (IMF). We decided not to exclude this detail to increase the perceived reliability of the information.

economic growth and help workers (due to complementarity and increased productivity of labour).

The participants in all groups were then asked if AI will make future jobs easier, harder to find, more complex, and more creative. Each of these items was asked as a true or false survey item. Afterwards, participants had to self-evaluate their perception of the effect of AI on their job prospects, selecting among the possible options of excitement, anxiety, anger, or hope; and to rate (on a scale from 0 to 10) how interesting they found job tasks as summarizing information for their superiors and peers or data analysis.[11] Finally, we asked a recall question to further assess the level of attention spent on (and memory retention of) the article they read and the tasks they performed.

To assess how AI affects policy preferences (our main dependent variables of interest) we first asked participants whether they supported government intervention in monitoring the pace of technological changes, drawing on the item from Gallego et al. (2021). This question was followed by a more specific item on support for different policies from Magistro et al. (2024): (a) an expansion of social spending to support laid-off workers; (b) funding programs to re-skill workers for new jobs; (c) universal basic income; (d) subsidizing companies to retain displaced workers and guarantee jobs; (e) reducing the number of immigrants entering the country for work; (f) increasing trade barriers; and (g) taxing companies that replace workers with machines, robots or AI.[12]

Finally, we consider that the risks of technological change are not orthogonal, at the individual level, to personal circumstances. In particular, factors such as socio-economic status, educational track, previous experience with AI, the perceived relevance of the tested tasks for individuals' careers as well as concrete post-graduation plans could be important sources of heterogeneity

[11] We asked this item to evaluate whether the effect of the treatment might have been conditional on the possible taste of our participants for the completion of tasks that can now be instead performed by AI. The skills we included in the item also mentioned meeting deadlines and writing emails).

[12] See the Appendix 5C for the precise wording of these items. To increase clarity, each item included additional details. We also included additional questions for dependent variables of interest, and in particular trust in the political system and towards political and economic elites. In this paper, we focus on policy preferences.

which we explore in Appendix A5.2.

5.3.2 Data and Summary Statistics

Our sample consists of 1,328 university students who participated in the experiment.[13] Table 5.1 presents descriptive statistics for key demographic and political characteristics across three treatment conditions. The majority of participants (around 67% in control and negative treatment, 73% in positive treatment) are between 18 and 24 years old and most participants are pursuing bachelor's degrees (62–68% across groups), with a smaller share enrolled in master's or doctoral programs.

Regarding post-graduation plans, roughly 60% of students report intending to entering the labour market immediately after graduation, while approximately 30% plan to pursue further education. The sample is skewed toward higher-income backgrounds, with between 61% and 67% of participants self-identifying as coming from high-income households. In terms of political orientation, approximately 72–76% of participants reported having voted in the last national election, suggesting a politically engaged sample. Lastly, self-placement on the left–right political scale is relatively balanced across experimental conditions, with left-oriented respondents comprising between 51% and 54% of the sample.

To check the validity of the random assignment, Table 5.2 presents the results of balance checks comparing mean values of key demographic, behavioral, and survey engagement variables across the control group and the two treatment conditions. For each variable, we report an F-test of the null hypothesis that all three group means are equal.

The results demonstrate that randomization across all measured covariates were achieved successfully, as we observe no statistically significant differences across treatment groups in terms of age, gender, ethnicity, socioeconomic background, educational level, field of study, post-graduation plans, previous AI usage and political orientation.

[13] We excluded 34 participants who either did not give consent or reported not being currently enrolled as students, leaving us with a final sample of 1,328 respondents.

Furthermore, there are no significant differences in participant engagement, as seen by comparable survey completion times, completion time on the summary writing task, and number of clicks during the experiment. Overall, these findings increase our confidence that any differences in outcomes across treatment conditions can be attributed to the experimental manipulations rather than baseline differences between groups.

Lastly, to ensure data quality, we implemented several checks. First, we included an attention check question about the content of the river pollution article that all participants read at the beginning of the survey. As shown in Figure A5.1, the predicted probability of passing this attention check is high and comparable across all three experimental groups (approximately 0.70–0.75), indicating that participants paid attention to the reading task regardless of treatment assignment.[14] Second, we monitored participant engagement through behavioral metrics captured during the experiment using a Java-based monitoring tool that tracked participants' activity. Figure A5.2 displays two key engagement measures: the number of clicks participants made during the survey and the time they spent completing the summary writing task. Both metrics show consistency across treatment groups. This similarity in engagement patterns further confirms that any differences in outcomes cannot be attributed to differential attention or effort across experimental conditions.

5.3.3 Main Empirical Specification

Our empirical approach is designed to estimate the causal impact of AI-related narratives on young citizens' policy preferences and our identification relies on the random assignment of participants to one of three experimental conditions. We estimate intention-to-treat (ITT) effects, since not all participants may have fully processed or internalized the information provided during the experiment. Our baseline specification is as follows:

$$Y_i = \alpha + \beta_1(\textit{Positive Framing})_i + \beta_2(\textit{Negative Framing})_i + \varepsilon_i \qquad (5.1)$$

[14] We further conduct robustness checks by excluding participants who failed the attention check. The results, presented in Figure A5.7, remain robust.

where Y_i represents the outcome of interest for individual i. The main dependent variables capture respondents' assessment of their future job prospects, self-reported emotional responses, support for AI-related regulation, and support for government policies aimed at mitigating labour market risks. These outcomes are measured as binary indicators equal to one if the participant expresses support for a given outcome or policy option, and zero otherwise. *Positive Framing* and *Negative Framing* are treatment indicators equal to 1 if individual i was randomly assigned to receive positive or negative information about AI's labour market effects, respectively.

To test Hypothesis 3, that treatment effects are moderated by pre-existing ideological orientations, we re-estimate Equation 5.1 with interaction terms:

$$Y_i = \alpha + \beta_1 \textit{(Positive Framing)}_i + \beta_2 \textit{(Negative Framing)}_i + \beta_3 \textit{(Left-Right Scale)}_i + \beta_4 \textit{(Positive Framing} \times \textit{Left-Right)}_i + \beta_5 \textit{(Negative Framing} \times \textit{Left-Right)}_i + \varepsilon_i \tag{5.2}$$

where *Left-Right Scale* measures participants' self-placement on the political spectrum (higher values indicate more right-leaning positions). The interaction coefficients β_4 and β_5 test whether treatment effects vary systematically with ideological position. This specification allows us to examine whether left-leaning and right-leaning students demand different types of policies in response to AI exposure.

Table 5.1: Descriptive Statistics by Treatment Group

Variable	Control	Treatment 1 (Positive Framing)	Treatment 2 (Negative Framing)
Age			
18–24	0.67	0.73	0.67
25–30	0.33	0.27	0.33
Gender			
Male	0.38	0.36	0.40
Female	0.62	0.64	0.60
Education			
Bachelor	0.62	0.62	0.68
Master	0.26	0.23	0.21
PhD	0.05	0.07	0.03
Other	0.07	0.08	0.07
Future Plans			
Further Education	0.30	0.29	0.30
Finding a Job	0.61	0.59	0.59
Not Decided	0.08	0.11	0.10
Other	0.01	0.01	0.01
Income Level			
High	0.61	0.67	0.65
Low	0.39	0.33	0.35
Voting Behaviour			
Not Voted	0.28	0.24	0.26
Voted	0.72	0.76	0.74
Political Scale			
Left Oriented	0.54	0.51	0.54
Right Oriented	0.46	0.49	0.46
Observations	438	443	447

Notes: This table presents summary statistics for the 1,328 university students who participated in the experiment. Column 1 presents mean values for the control group, Column 2 for Treatment Group 1 (who received positive AI information), and Column 3 for Treatment Group 2 (who received negative AI information). All values represent the proportion of respondents in each category within the respective treatment group. Participants were randomly assigned to one of the three experimental conditions with equal probability. The sample consists of students aged 18–30 recruited via Prolific Academic from the United Kingdom and United States.

Table 5.2: Balance Check: Treatment vs. Control Groups

Variable	Control	Treatment 1 (Positive)	Treatment 2 (Negative)	P-Value	N(Control)	N(T1)	N(T2)
Demographics							
Age	23.0	22.7	22.9	0.447	432	430	438
Gender							
Male	0.40	0.40	0.40	0.534	430	428	436
Female	0.60	0.60	0.60	0.534	430	428	436
Ethnicity							
White	0.60	0.60	0.50	0.345	431	428	437
Black 2	0.20	0.10	0.10	0.279	431	428	437
Mixed 3	0.10	0.10	0.10	0.620	431	428	437
Asian	0.10	0.10	0.20	0.100	431	428	437
Education							
Degree Type							
Bachelor	0.60	0.60	0.70	0.072	438	443	447
Master	0.30	0.20	0.20	0.244	438	443	447
PhD	0.10	0.10	0.00	0.068	438	443	447
Other	0.10	0.10	0.10	0.769	438	443	447
Field of Study							
Arts and Humanities	0.10	0.10	0.10	0.536	438	443	447
Social Sciences	0.20	0.20	0.10	0.277	438	443	447
Natural Sciences	0.10	0.10	0.10	0.452	438	443	447
Engineering and Technology	0.20	0.20	0.20	0.251	438	443	447
Business and Economics	0.20	0.10	0.20	0.007	438	443	447
Health and Medicine	0.20	0.20	0.10	0.947	438	443	447
Law and Legal Studies	0.00	0.00	0.00	0.770	438	443	447
Education and Teaching	0.00	0.00	0.00	0.502	438	443	447
Post-Graduation Plans							
Further Education	0.30	0.30	0.30	0.946	438	443	447
Finding a Job	0.60	0.60	0.60	0.788	438	443	447
Not Decided	0.10	0.10	0.10	0.415	438	443	447
Other	0.00	0.00	0.00	0.622	438	443	447
Previous AI Experience							
Never	0.20	0.20	0.10	0.326	438	443	447
Only a couple of times	0.20	0.20	0.20	0.640	438	443	447
Sometimes but not too frequently	0.20	0.20	0.30	0.427	438	443	447
Frequently	0.30	0.20	0.20	0.877	438	443	447
Most of the times	0.10	0.10	0.10	0.283	438	443	447
Always	0.00	0.10	0.00	0.650	438	443	447
Socioeconomic Background							
Income Level							
High	0.60	0.70	0.60	0.140	438	443	447
Low	0.40	0.30	0.40	0.140	438	443	447
Political Characteristics							
Voting Behaviour							
Not Voted	0.30	0.20	0.30	0.486	433	438	445
Voted	0.70	0.80	0.70	0.486	433	438	445
Political Orientation							
Left	0.50	0.50	0.50	0.571	394	390	402
Right	0.50	0.50	0.50	0.571	394	390	402
Survey Engagement							
Survey Duration (in seconds)	1,430.8	1,410.0	1,394.9	0.983	438	443	447
Summary Completion Time (in seconds)	245.1	243.3	247.4	0.739	438	443	447
Number of Clicks	6.8	7.3	6.6	0.412	438	443	447

Notes: This table presents balance checks across experimental conditions. Columns 1–3 report mean values for the control group, positive AI framing treatment, and negative AI framing treatment, respectively. Column 4 reports p-values from one-way ANOVA F-tests testing the null hypothesis that all three group means are equal. Columns 5–7 report sample sizes for each group (note that sample sizes vary slightly across variables due to missing responses). The sample consists of students aged 18–30 recruited via Prolific Academic from the United Kingdom and United States.

5.4 Results

We begin with analyzing whether and how our treatments affected students' perceptions of their job prospects. Our manipulation checks indicate that the experimental design worked as intended. Figures 5.1 and 5.2, respectively, show the average agreement on four different perceptions of job prospects, as well as the mean of negative emotions that participants associate with AI and future careers (anger or anxiety). The differences across groups are large and statistically significant: Groups exposed to AI and narratives that highlight its risks reported more anger or anxiety about future prospects in the labor market. Instead, students who read a positive framing for AI advancements reported no significantly different emotions or career outlooks from the control group.

But how does this induced anxiety affect students' attitude towards the pace of technological change in the workplace? Figure 5.3 shows how the mean support in slowing the pace of AI. There are at least two takeaways here. First, while younger demographics are typically portrayed as enthusiastic supporters of new technologies, we observe surprisingly high levels of support for decelerating technological change. Both the overall mean support for slowing down change (40%) and the group means (36%, 42% and 44% for control, positive framing and negative framing groups, respectively) point in this direction. Second, participants in the negative framing conditions demanded a slowdown in the pace of AI beyond this already high baseline [15]. The estimate for the positive framing condition was equally positive, but not statistically significant. These results offer an interesting comparison with previous work by (Gallego et al., 2021) who equally found that priming on technological risks boosted support for slowing down the introduction of AI in the workplace, but reported much lower levels of general support for this policy among the general public.

[15] Participants exposed to the negative AI framing were about 8 percentage points more likely to support slowing down AI development compared to the control group.

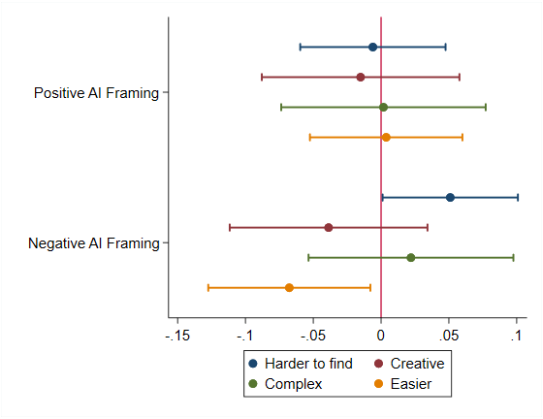


Figure 5.1: **Assessment of job prospects**

Notes: The figure shows estimated coefficients (95% confidence intervals shown) from the OLS models in which the dependent variable is a binary indicator for whether participants agreed with each statement about future job prospects. The control group represents the baseline and the coefficient estimates show whether being assigned to a treatment group affected participants' assessment of job market prospects.

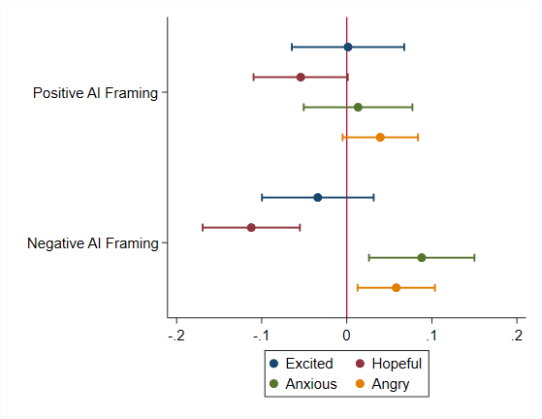


Figure 5.2: **Self-reported Emotions**

Notes: The figure shows estimated coefficients (95% confidence intervals shown) from the OLS models where the dependent variable is a binary indicator equal to 1 if participants selected a given emotion (e.g., *excited*, *hopeful*, *angry*, or *anxious*) when thinking about AI and the future job market. The control group represents the baseline.

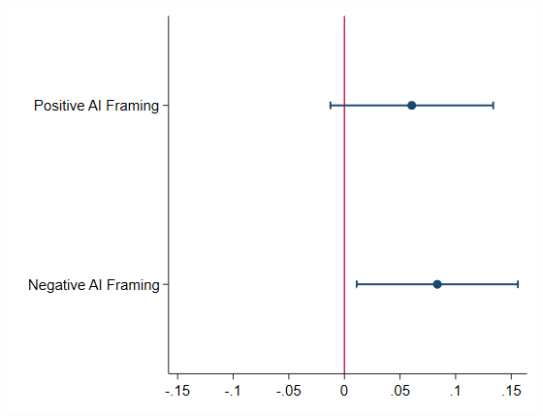


Figure 5.3: **Support for slowing down AI**

Notes: The figure shows estimated coefficients (95% confidence intervals shown) from the OLS model in which the dependent variable is a binary indicator for whether participants agreed with the statement that the government should act to slow down the development of AI. The control group represents the baseline and the coefficient, therefore, show whether being assigned to a treatment group affected participants' view of the role of their governments.

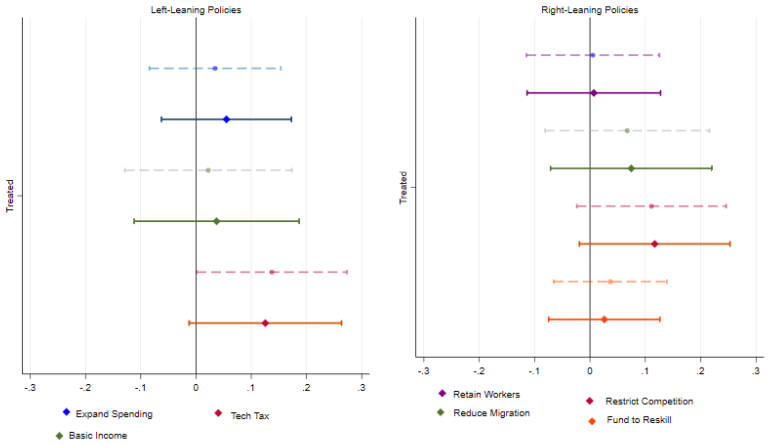


Figure 5.4: **Treatment Effect on Government Policies**

Notes: Treatment effects on left-leaning (on the left side) and right-leaning policy preferences (on the right side). Dashed bars with circles show coefficient estimates (95% confidence intervals shown) without controls, and solid bars with diamonds show estimates with controls. Control variables include: age, gender, nationality, income level, parental education and employment status, degree type, field of study, post-graduation plans, previous AI experience, and relevance of summary writing tasks for future career.

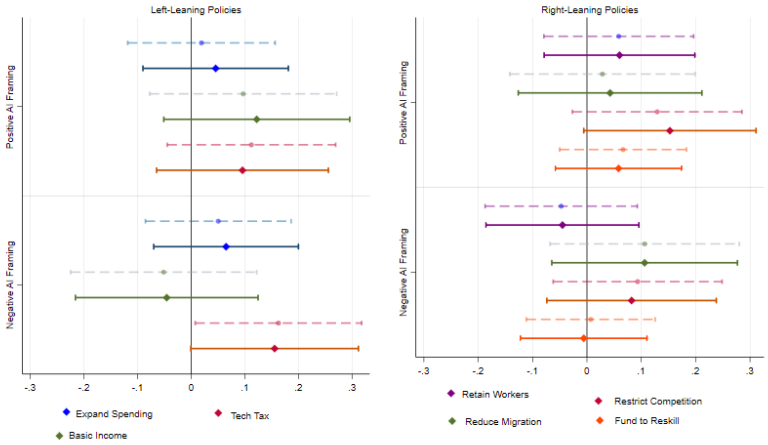


Figure 5.5: **Treatment Effect on Government Policies**

Notes: Treatment effects on left-leaning (on the left side) and right-leaning policy preferences (on the right side). Dashed bars with circles show coefficient estimates (95% confidence intervals shown) without controls, and solid bars with diamonds show estimates with controls. Control variables include: age, gender, nationality, income level, parental education and employment status, degree type, field of study, post-graduation plans, previous AI experience, and relevance of summary writing tasks for future career.

The picture is more nuanced when we consider specific policy options. Our first and second hypotheses, respectively, predicted that exposure to AI competition would increase demand for risk mitigating policies and that this effect would be greatest among individuals exposed to a negative framing about AI. To test our first hypothesis, we started by pooling both treatment conditions (positive and negative framing) and comparing them against the control group. As shown in Figure 5.4, we observe low to moderate baseline levels of support for all seven policy options, but the pooled treatment effect is not statistically significant for any policy. To test Hypothesis 2, Figure 5.5 reports disaggregated treatment effects by framing condition (positive versus negative). However, even when distinguishing between positive and negative framing, we find no significant effects of either treatment condition on any of the seven policies. In other words, merely exposing students to AI com-

petition, regardless of how information on AI is framed, does not translate into increased demand for risk-mitigating policies. These surprising results are robust even when we consider several sources of potential heterogeneity–students' educational track, socioeconomic background, previous experience with using AI tools, post-graduation plans, and the relevance of writing and summarizing tasks for their professional career.[16] Additionally, these findings hold even when distinguishing the UK and US samples: despite small differences, in neither country there's a systematic increase in demand for mitigating policies.[17] This is particularly interesting because we find a decisive difference from Haslberger et al. (2025). They focused on working age individuals and found that exposure to AI was not necessarily influencing risk perceptions but was nonetheless generating demands for social policies. Focusing on university students instead, we found the polar opposite: Exposure to AI triggers anxiety but the latter does not translate into a demand for social policies. While our design does not allow us to fully account for this contrasting result, we speculate that two explanations may underlie this difference. First of all, as previously mentioned, young individuals tend to be more optimistic about their ability to navigate technological transitions. Second, our design relies heavily on the activation of direct competition with AI (coupled with a positive/negative framings). It is, therefore, possible that our design activated more egocentric than sociotropic considerations–while Haslberger et al. (2025) underline that sociotropic reasoning seems to be the responsible factor activating policy demands in their study.

Instead, preferences over policies load onto preexisting political beliefs. That is, the liking or dislike of policies is driven more by their ideological orientations than by AI competition and narratives about it. Table 5.3 shows the results of the interaction effect between our treatment groups and participants' self-placement on the left-right scale for the seven risk-mitigating policies, with positive coefficients suggesting alignment with left-wing positions and nega-

[16] See Appendix A5.0.1 for additional details on the heterogeneity analysis.

[17] We however find that UK respondents may have a preference for active labour market policies (ALMPs), such as funding programs to re-skill workers, at least when AI competition is matched by a positive framing of AI. In the US, a positive framing of AI instead leads to increased support for a restriction in international competition. This last result may, however, have been influenced by the political discussions over tariffs in the time before and during our data collection. Figure A5.4 shows coefficient estimates by country.

tive coefficients with right-wing positions [18]. This table shows that left-wing individuals strongly prefer expanding social spending, taxing technologies, or introducing basic income. On the right, students support reducing migration and restricting international competition. Only funding re-skilling programs and paying to retain displaced workers did not align with our expectations. However, the overall support for mitigating AI risks appears to follow patterns that are broadly consistent with pre-existing political beliefs. These results suggest that ideological orientations provide an important anchor for young citizens' policy preferences, even in a novel domain such as AI.

Power analysis together with significant findings on other dimensions suggest this our results reflect a true null-effect. At 80% power and a two-sided $\alpha = 0.05$, our design is able to detect treatment effects of approximately 9 percentage points in policy support outcomes, corresponding to a standardized effect size of about Cohen's $d \approx 0.21$. This level of sensitivity is comparable to the magnitude of effects reported in related experimental research (e.g., Gallego et al., 2021), which finds impacts in the mid-single to low-teens percentage points on public support for AI and technology policy instruments. This allow us to confidently rule out moderate-to-large shifts with conventional confidence, even if very small effects cannot be excluded. It also aligns with the roughly 8 percentage point increase in support for slowing down AI observed under negative framing in our data. Importantly, we were able to pick up significant treatment effects on emotions and job prospects at similar or lower levels of sensitivity to the general battery of policy questions. Putting this together with the consistently null results across outcomes and subgroups suggests that the absence of significant effects reflects substantively meaningful ideological anchoring rather than a lack of statistical power.

To summarize, previous research consistently established a clear link between technological risks and political behaviour and/or policy preferences among the general public (Kurer, 2020; Gallego et al., 2022; Anelli et al., 2023; Ahrens, 2024). Our results indicate no such general effect when considering younger citizens, whose policy preferences are simply contingent on preexisting ideological beliefs.

[18] See Table A5.1 for the results including controls.

Table 5.3: Treatment Effects on Support for Government Policies

	Left-Leaning Policies			Right-Leaning Policies			
	(1) Social Spending	(2) Basic Income	(3) Tax Tech	(4) Retain Workers	(5) Reduce Migration	(6) Restrict Competition	(7) Fund to Reskill
Positive AI Framing	0.020	0.006	-0.015	0.125	0.160	0.051	0.124
	(0.099)	(0.126)	(0.114)	(0.103)	(0.119)	(0.111)	(0.088)
Negative AI Framing	0.043	-0.069	0.071	-0.012	0.153	0.046	0.144*
	(0.097)	(0.123)	(0.111)	(0.101)	(0.117)	(0.108)	(0.086)
Control × left-right	0.525***	0.634***	0.405***	0.343***	-1.008***	-0.911***	0.279***
	(0.101)	(0.129)	(0.116)	(0.105)	(0.122)	(0.113)	(0.090)
Positive AI Framing × left-right	0.550***	0.688***	0.498***	0.234**	-1.269***	-0.780***	0.185**
	(0.101)	(0.130)	(0.116)	(0.106)	(0.122)	(0.114)	(0.091)
Negative AI Framing × left-right	0.586***	0.561***	0.459***	0.282***	-1.145***	-0.768***	0.014
	(0.100)	(0.128)	(0.115)	(0.104)	(0.121)	(0.112)	(0.090)
Constant	18.420***	18.014***	18.611***	18.360***	18.552***	18.005***	18.962***
	(0.069)	(0.088)	(0.079)	(0.072)	(0.083)	(0.077)	(0.061)
Observations	1,183	1,185	1,184	1,185	1,185	1,185	1,185

Notes: This table reports OLS regression results examining whether treatment effects vary by participants' ideological orientation. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Left-right is a continuous scale measuring self-placement on the political spectrum, where higher values indicate more right-leaning positions. The interaction terms test whether treatment effects differ across the ideological spectrum. Columns 1–3 show left-leaning policies (expanding social spending, implementing basic income and taxing tech companies), while columns 4–7 show right-leaning policies (subsidizing companies to retain workers, reducing migration, restricting international competition, and funding reskilling programs). Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Why do students diverge from well-established patterns? We suggest four plausible explanations. First and foremost, our data indicates low levels of trust in the ability of political elites to navigate this challenge. Across both the UK and the US, participants not only put little trust in the government, but in fact trust business elites more when it comes to AI challenges. It stands to reason that citizens who lack confidence in their government, regardless of their age, will be less supportive of government's political initiatives. Secondly, we speculate that students could perceive the proposed policies as favouring labour market incumbents. If this is true, these policies could be seen as a barrier to an increasingly competitive labour market–a non-welcome addition to prospects also limited by AI. Thirdly, the fact that students are not yet fully integrated into the labour market may also matter. Entering employment is a prospect they anticipate rather than a lived reality, which makes the risks of AI more psychologically distant. Research in social psychology suggests that when threats are located in the future rather than in the present, people tend to construe them in abstract terms and are less likely to articulate concrete policy demands in response (Trope and Liberman, 2010; Zimbardo and Boyd, 2014). This temporal distance may help explain why students' views on AI remain largely filtered through broader ideological orientations rather than translated into immediate policy preferences. This retreat into established beliefs makes particular sense in the context of the high levels of uncertainty around the exact impact of AI. As noted earlier, previous survey evidence established that 41% of U.S. young adults believe that AI will likely have both negative and positive impacts on their work and lives (CDT, 2025). This suggests that large groups of young adults are (rightfully) uncertain about how this new technology will shape their careers. Finally, these results could echo previous conclusions by (Zhang, 2022) who argued that correcting beliefs about labour market prospects among citizens does necessarily lead to substantive policy shifts. In fact, there is a substantial body of literature which points out that such a "backfire" effect is suprisingly rare because changing factual beliefs do not straightforwadly influence how individuals use those facts to form political opinions (Wood and Porter, 2019; Guess and Coppock, 2020; Zhang, 2022).

5.5 Conclusion

This study examined how young citizens in advanced democracies respond to technological risks in the labour market. Prior research has documented that labour market incumbents demand stronger welfare policies to mitigate AI risks. Less is instead known about those yet to enter the workforce. We fielded an online survey experiment with university students in the UK and the US to explore whether exposure to AI risks and narratives underlying the advantages and perils of AI development shape perceptions of future employment and, consequently, the policy preferences of tomorrow's workforce.

Our findings reveal a nuanced picture. Exposure to competition with AI and negative information about it does trigger anxiety about job prospects, confirming that young citizens are not immune to concerns about technological disruptions. These students also report a much greater desire to slow the pace of technological change. However, contrary to our expectations based on previous studies on job incumbents and political economy assumptions, this does not translate into support for risk-mitigating policies. Instead, participants' policy preferences largely reflect preexisting ideological orientations: left-leaning students favour redistributive measures such as social spending or basic income, whereas right-leaning students prefer policies aligned with market-oriented or protectionist approaches. These patterns persist across treatments, suggesting that ideological predispositions, rather than situationally induced labour market anxiety, are the dominant drivers of policy preferences among the young.

We interpret these results as reflecting three key mechanisms. First, young citizens display low trust in governmental institutions' ability to manage AI-driven labour market risks, limiting their willingness to endorse specific policy interventions. Second, students may perceive proposed policies as primarily benefiting labour market incumbents rather than themselves, attenuating the link between personal anxiety and policy support. Third, owing to their limited labour market experience, students might approach risk-mitigating policies more theoretically, using their preexisting beliefs as a guide to evaluate potential interventions. Or, young individuals may demand policies other than traditional social recipes to address AI challenges to their future employment.

While our research design ensures internal validity across treatment groups, the generalizability of these findings warrants caution. The sample consists of university students recruited via Prolific Academic, who may systematically differ from the broader youth population in terms of digital literacy and political engagement. Accordingly, the findings should not be interpreted as population-level estimates of policy preferences, but rather as evidence on the underlying mechanisms through which AI-related information shapes anxiety and policy attitudes under conditions of uncertainty. Future research should therefore examine whether these patterns hold across different national and institutional contexts. At the same time, these insights underscore the importance of considering age and labour market status when evaluating political responses to technological change and suggest that interventions aimed at shaping public opinion on AI may need to account for generational differences.

Appendix 5A: Additional Figures & Tables

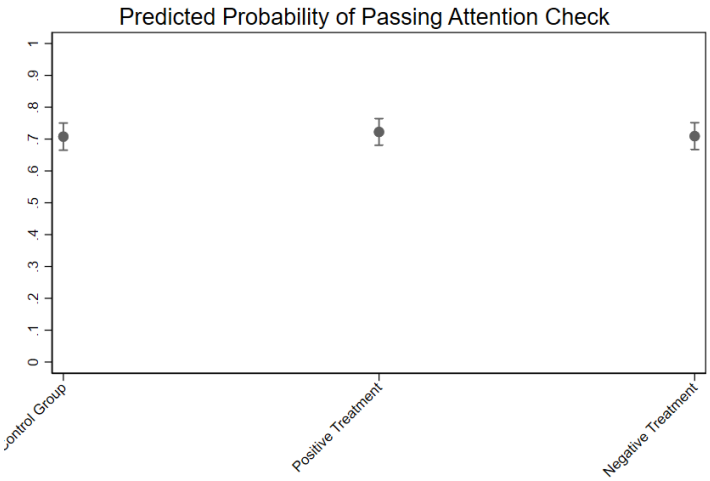


Figure A5.1: Attention Check - River article

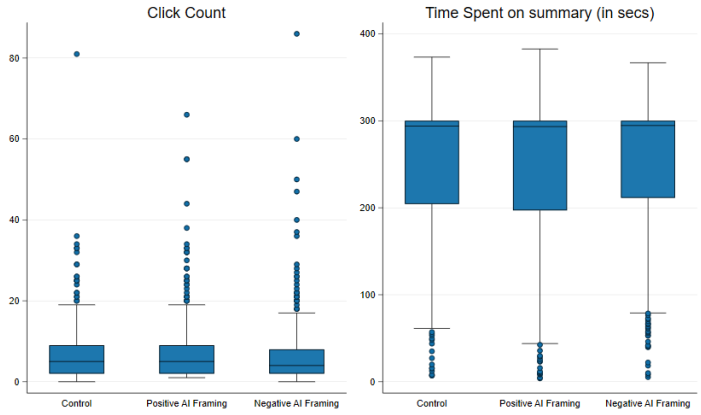
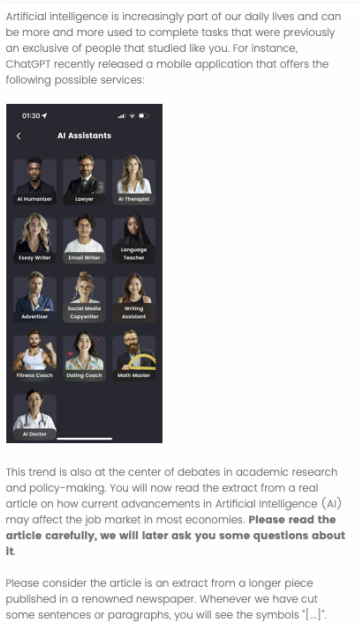
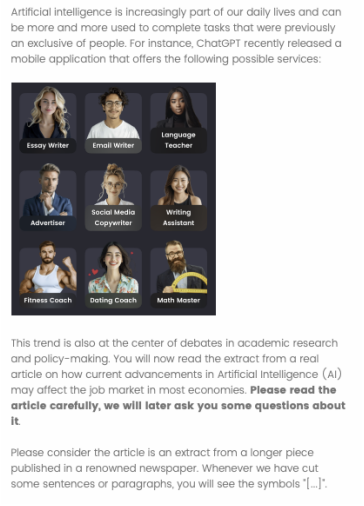


Figure A5.2: Engagement Measures by Treatment Group



(a) Image shown to Negative AI Framing Treatment Participants



(b) Image shown to Positive AI Framing Treatment Participants

Figure A5.3: Screenshots of AI Assistant Applications Shown to Treatment Groups

(a) UK Sample

(b) US Sample

Figure A5.4: Treatment Effects on Government Policies by Country

Notes: Treatment effects on left-leaning (on the left side of each panel) and right-leaning policy preferences (on the right side of each panel). Dashed bars with circles show coefficient estimates (95% confidence intervals shown) without controls, and solid bars with diamonds show estimates with controls. Control variables include: age, gender, nationality, income level, parental education and employment status, degree type, field of study, post-graduation plans, previous AI experience, and relevance of summary writing tasks for future career. Panel (a) shows results for UK students; Panel (b) shows results for US students. Sample consists of students aged 18–30 recruited via Prolific Academic from the United Kingdom and United States.

Table A5.1: Treatment Effects on Support for Government Policies (with Controls)

	Left-Leaning Policies			Right-Leaning Policies			
	(1) Social Spending	(2) Basic Income	(3) Tax Tech	(4) Retain Workers	(5) Reduce Migration	(6) Restrict Competition	(7) Fund to Reskill
Positive AI Framing	0.047	0.035	-0.062	0.092	0.197	0.073	0.132
	(0.100)	(0.127)	(0.116)	(0.105)	(0.120)	(0.113)	(0.089)
Negative AI Framing	0.056	-0.064	0.051	-0.049	0.168	0.040	0.130
	(0.097)	(0.123)	(0.113)	(0.102)	(0.117)	(0.109)	(0.086)
Control × left-right	0.497***	0.613***	0.360***	0.285***	-0.939***	-0.925***	0.309***
	(0.102)	(0.130)	(0.119)	(0.107)	(0.123)	(0.115)	(0.091)
Positive AI Framing × left-right	0.502***	0.647***	0.499***	0.209*	-1.238***	-0.805***	0.169*
	(0.103)	(0.131)	(0.120)	(0.108)	(0.125)	(0.117)	(0.092)
Negative AI Framing × left-right	0.553***	0.520***	0.405***	0.277***	-1.076***	-0.779***	0.045
	(0.102)	(0.129)	(0.118)	(0.107)	(0.123)	(0.115)	(0.091)
Age	0.023***	0.048***	0.018*	0.005	-0.001	0.018*	0.028***
	(0.008)	(0.011)	(0.010)	(0.009)	(0.010)	(0.010)	(0.008)
Gender: Female	-0.074	-0.015	0.062	0.114*	-0.373***	0.199***	-0.067
	(0.061)	(0.078)	(0.071)	(0.064)	(0.074)	(0.069)	(0.055)
Nationality	0.001	-0.010***	-0.007*	-0.002	0.009**	-0.001	-0.005*
	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
Income Level	0.098***	0.125***	0.067***	0.064***	-0.053**	-0.011	0.030
	(0.021)	(0.026)	(0.024)	(0.022)	(0.025)	(0.024)	(0.019)
Mother education level	0.076*	0.177***	0.063	0.002	-0.134***	-0.092**	-0.023
	(0.040)	(0.050)	(0.046)	(0.042)	(0.048)	(0.045)	(0.035)
Father education level	0.032	-0.031	-0.039	0.027	0.032	0.087**	0.009
	(0.036)	(0.045)	(0.042)	(0.038)	(0.043)	(0.040)	(0.032)
Mother employment status	0.008	0.040*	-0.022	0.011	0.001	-0.024	-0.007
	(0.017)	(0.022)	(0.020)	(0.018)	(0.021)	(0.020)	(0.015)
Father employment status	0.017	-0.027	0.023	0.012	0.013	0.023	-0.013
	(0.017)	(0.022)	(0.020)	(0.018)	(0.021)	(0.019)	(0.015)
Constant	17.034***	16.336***	18.112***	17.851***	18.833***	17.594***	18.522***
	(0.321)	(0.407)	(0.373)	(0.337)	(0.388)	(0.363)	(0.286)
Observations	1,148	1,149	1,148	1,148	1,149	1,149	1,149

Notes: This table reports OLS regression results including control variables examining whether treatment effects vary by participants' ideological orientation. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Left-right is a continuous scale measuring self-placement on the political spectrum, where higher values indicate more right-leaning positions. The interaction terms test whether treatment effects differ across the ideological spectrum. Columns 1–3 show left-leaning policies (expanding social spending, implementing basic income and taxing tech companies), while columns 4–7 show right-leaning policies (subsidizing companies to retain workers, reducing migration, restricting international competition, and funding reskilling programs). Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A5.0.1 Heterogeneity Analyses

Table A5.2: Treatment Effects on Support for Government Policies, by Income

	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Panel A: Low-Income Respondents							
Positive AI Framing	0.111	0.086	0.093	0.180	0.054	0.162	0.216**
	(0.128)	(0.165)	(0.152)	(0.126)	(0.139)	(0.146)	(0.110)
Negative AI Framing	-0.029	-0.197	-0.006	-0.025	-0.007	0.031	0.174
	(0.133)	(0.161)	(0.150)	(0.135)	(0.145)	(0.143)	(0.111)
Control × left-right	0.405***	0.588***	0.459***	0.308**	-1.063***	-0.885***	0.423***
	(0.143)	(0.178)	(0.152)	(0.138)	(0.159)	(0.157)	(0.130)
Positive AI Framing × left-right	0.387***	0.455**	0.369**	0.044	-1.348***	-0.891***	0.076
	(0.132)	(0.181)	(0.165)	(0.143)	(0.159)	(0.157)	(0.122)
Negative AI Framing × left-right	0.629***	0.739***	0.533***	0.326**	-1.054***	-0.759***	0.249**
	(0.128)	(0.165)	(0.156)	(0.143)	(0.169)	(0.148)	(0.115)
Constant	17.780***	16.908***	17.872***	18.285***	18.797***	17.290***	18.599***
	(0.397)	(0.488)	(0.441)	(0.388)	(0.449)	(0.440)	(0.316)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	666	666	665	665	666	666	666
Panel B: High-Income Respondents							
	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Positive AI Framing	-0.034	-0.011	-0.297	-0.018	0.430**	-0.135	0.013
	(0.189)	(0.230)	(0.200)	(0.189)	(0.205)	(0.191)	(0.161)
Negative AI Framing	0.199	0.161	0.170	-0.089	0.486**	0.047	0.033
	(0.167)	(0.211)	(0.182)	(0.172)	(0.201)	(0.178)	(0.140)
Control × left-right	0.680***	0.717***	0.258	0.299*	-0.730***	-0.975***	0.148
	(0.156)	(0.194)	(0.185)	(0.167)	(0.196)	(0.164)	(0.136)
Positive AI Framing × left-right	0.724***	0.927***	0.709***	0.449***	-1.158***	-0.667***	0.270*
	(0.170)	(0.203)	(0.189)	(0.173)	(0.192)	(0.188)	(0.141)
Negative AI Framing × left-right	0.473***	0.258	0.216	0.278*	-1.158***	-0.799***	-0.179
	(0.146)	(0.199)	(0.172)	(0.167)	(0.199)	(0.173)	(0.141)
Constant	17.347***	17.291***	19.500***	17.996***	17.999***	17.741***	18.866***
	(0.455)	(0.598)	(0.486)	(0.484)	(0.594)	(0.565)	(0.463)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	482	483	483	483	483	483	483

Notes: This table reports treatment effects on policy preferences by income level. Panel A shows results for respondents from low-income family backgrounds; Panel B shows results for high-income backgrounds. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A5.3: Treatment Effects on Support for Government Policies, by Degree Type

	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Panel A: Bachelor's Degree Students							
Positive AI Framing	0.055	-0.042	0.072	0.071	0.145	0.070	0.238**
	(0.130)	(0.161)	(0.152)	(0.134)	(0.156)	(0.147)	(0.116)
Negative AI Framing	0.071	-0.019	0.024	-0.137	0.260*	-0.022	0.206*
	(0.124)	(0.153)	(0.145)	(0.127)	(0.148)	(0.140)	(0.110)
Control × left-right	0.505***	0.770***	0.523***	0.234*	-0.814***	-0.819***	0.366***
	(0.131)	(0.162)	(0.153)	(0.135)	(0.157)	(0.148)	(0.117)
Positive AI Framing × left-right	0.582***	0.725***	0.369**	0.218	-1.138***	-0.882***	0.091
	(0.132)	(0.163)	(0.154)	(0.136)	(0.158)	(0.149)	(0.118)
Negative AI Framing × left-right	0.499***	0.462***	0.491***	0.315**	-1.150***	-0.693***	0.045
	(0.125)	(0.154)	(0.146)	(0.128)	(0.149)	(0.141)	(0.111)
Constant	16.760***	15.755***	17.743***	17.418***	18.711***	17.464***	18.017***
	(0.422)	(0.523)	(0.493)	(0.434)	(0.506)	(0.476)	(0.377)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	726	727	726	727	727	727	727
Panel B: Master's and PhD Students							
	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Positive AI Framing	0.091	0.173	-0.246	0.042	0.322	0.112	0.158
	(0.174)	(0.228)	(0.202)	(0.194)	(0.211)	(0.203)	(0.154)
Negative AI Framing	0.144	-0.206	0.236	0.146	-0.192	0.179	0.089
	(0.179)	(0.235)	(0.208)	(0.200)	(0.218)	(0.209)	(0.159)
Control × left-right	0.462**	0.351	0.159	0.386*	-1.244***	-1.151***	0.237
	(0.183)	(0.240)	(0.213)	(0.204)	(0.223)	(0.213)	(0.162)
Positive AI Framing × left-right	0.276	0.308	0.767***	0.126	-1.405***	-0.678***	0.152
	(0.190)	(0.248)	(0.220)	(0.211)	(0.231)	(0.221)	(0.168)
Negative AI Framing × left-right	0.483**	0.543**	0.043	0.023	-0.725***	-0.918***	-0.057
	(0.201)	(0.263)	(0.233)	(0.223)	(0.244)	(0.234)	(0.178)
Constant	17.740***	18.076***	19.012***	18.273***	19.398***	17.471***	19.204***
	(0.652)	(0.854)	(0.757)	(0.728)	(0.793)	(0.760)	(0.578)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	337	337	337	336	337	337	337

Notes: This table reports treatment effects on policy preferences by degree type. Panel A shows results for bachelor's degree students; Panel B shows results for master's and PhD students. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, income level, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A5.4: Treatment Effects on Support for Government Policies, by Field of Study

	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Panel A: Social Science Students							
Positive AI Framing	-0.010	-0.280	-0.265	-0.087	0.621***	0.020	0.029
	(0.174)	(0.212)	(0.185)	(0.174)	(0.192)	(0.189)	(0.144)
Negative AI Framing	0.139	-0.128	-0.223	0.012	0.479***	-0.109	0.138
	(0.161)	(0.197)	(0.172)	(0.161)	(0.178)	(0.175)	(0.133)
Control × left-right	0.539***	0.514**	0.113	0.221	-0.519***	-1.196***	0.087
	(0.175)	(0.214)	(0.187)	(0.175)	(0.193)	(0.190)	(0.145)
Positive AI Framing × left-right	0.631***	1.056***	0.340	0.477**	-1.505***	-0.955***	0.334**
	(0.195)	(0.238)	(0.208)	(0.194)	(0.215)	(0.212)	(0.161)
Negative AI Framing × left-right	0.444**	0.634***	0.304	0.139	-1.623***	-0.975***	0.149
	(0.184)	(0.225)	(0.197)	(0.184)	(0.203)	(0.200)	(0.152)
Constant	16.688***	15.521***	17.459***	17.455***	19.603***	17.239***	18.615***
	(0.536)	(0.654)	(0.572)	(0.535)	(0.591)	(0.582)	(0.443)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	401	401	401	400	401	401	401
Panel B: Other Fields							
	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Positive AI Framing	0.055	0.225	0.101	0.187	0.002	0.151	0.190*
	(0.123)	(0.159)	(0.150)	(0.133)	(0.153)	(0.142)	(0.114)
Negative AI Framing	0.004	-0.003	0.247*	-0.078	-0.031	0.140	0.124
	(0.123)	(0.157)	(0.148)	(0.132)	(0.151)	(0.141)	(0.113)
Control × left-right	0.440***	0.668***	0.509***	0.315**	-1.164***	-0.777***	0.441***
	(0.127)	(0.164)	(0.154)	(0.137)	(0.157)	(0.146)	(0.118)
Positive AI Framing × left-right	0.427***	0.446***	0.555***	0.089	-1.086***	-0.736***	0.091
	(0.122)	(0.157)	(0.148)	(0.132)	(0.151)	(0.140)	(0.113)
Negative AI Framing × left-right	0.590***	0.433***	0.417***	0.342**	-0.740***	-0.693***	0.003
	(0.124)	(0.158)	(0.149)	(0.133)	(0.152)	(0.142)	(0.114)
Constant	17.315***	16.663***	18.326***	18.046***	18.288***	17.624***	18.427***
	(0.409)	(0.526)	(0.495)	(0.441)	(0.506)	(0.471)	(0.379)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	747	748	747	748	748	748	748

Notes: This table reports treatment effects on policy preferences by field of study. Panel A shows results for social science students; Panel B shows results for students in other fields (including STEM, humanities, business, etc.). The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, income level, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A5.5: Treatment Effects on Support for Government Policies, by Post-Graduation Plans

	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Panel A: Planning Further Education/Training							
Positive AI Framing	0.052	-0.509**	-0.038	-0.098	0.210	0.010	0.218
	(0.190)	(0.230)	(0.217)	(0.184)	(0.240)	(0.218)	(0.172)
Negative AI Framing	0.159	-0.252	-0.133	-0.146	0.008	0.162	0.180
	(0.186)	(0.225)	(0.213)	(0.180)	(0.235)	(0.214)	(0.169)
Control × left-right	0.514***	0.324	0.169	0.370**	-0.584**	-0.790***	0.236
	(0.194)	(0.234)	(0.222)	(0.187)	(0.245)	(0.223)	(0.176)
Positive AI Framing × left-right	0.467**	1.169***	0.276	0.438**	-1.306***	-0.625***	0.043
	(0.203)	(0.245)	(0.232)	(0.196)	(0.256)	(0.233)	(0.184)
Negative AI Framing × left-right	0.628***	0.634***	0.513**	0.702***	-1.148***	-1.027***	-0.034
	(0.195)	(0.235)	(0.222)	(0.188)	(0.246)	(0.224)	(0.176)
Constant	16.292***	16.754***	18.865***	18.181***	16.622***	17.154***	18.112***
	(0.698)	(0.842)	(0.797)	(0.674)	(0.880)	(0.801)	(0.633)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	333	333	333	333	333	333	333
Panel B: Planning to Find a Job							
	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Positive AI Framing	0.114	0.280*	-0.040	0.233*	0.247*	0.147	0.022
	(0.130)	(0.165)	(0.149)	(0.140)	(0.148)	(0.146)	(0.114)
Negative AI Framing	0.127	0.061	0.139	0.024	0.288**	0.103	0.121
	(0.124)	(0.157)	(0.142)	(0.133)	(0.141)	(0.139)	(0.108)
Control × left-right	0.522***	0.743***	0.414***	0.230	-1.084***	-0.921***	0.334***
	(0.135)	(0.172)	(0.155)	(0.145)	(0.154)	(0.152)	(0.118)
Positive AI Framing × left-right	0.434***	0.442**	0.507***	0.078	-1.290***	-0.917***	0.255**
	(0.136)	(0.173)	(0.156)	(0.146)	(0.156)	(0.153)	(0.120)
Negative AI Framing × left-right	0.412***	0.489***	0.417***	0.051	-1.092***	-0.713***	0.039
	(0.134)	(0.170)	(0.153)	(0.143)	(0.152)	(0.150)	(0.117)
Constant	17.248***	16.372***	18.086***	17.520***	19.411***	17.636***	18.638***
	(0.401)	(0.510)	(0.459)	(0.431)	(0.457)	(0.451)	(0.351)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	686	687	686	686	687	687	687

Notes: This table reports treatment effects on policy preferences by post-graduation plans. Panel A shows results for students planning to pursue further education or training; Panel B shows results for students planning to enter the job market immediately after graduation. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, income level, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A5.6: Treatment Effects on Support for Government Policies, by Previous AI Experience

	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Panel A: Students Who Sometimes Use AI							
Positive AI Framing	0.153	0.171	-0.410	-0.086	0.435	-0.021	0.315
	(0.233)	(0.317)	(0.264)	(0.244)	(0.307)	(0.272)	(0.243)
Negative AI Framing	0.289	0.149	0.211	0.150	0.130	-0.186	0.291
	(0.207)	(0.281)	(0.234)	(0.216)	(0.272)	(0.241)	(0.215)
Control × left-right	0.652***	0.577**	0.700***	0.387*	-1.171***	-1.276***	0.496**
	(0.200)	(0.271)	(0.226)	(0.209)	(0.262)	(0.232)	(0.208)
Positive AI Framing × left-right	0.594**	0.560*	0.947***	0.462*	-1.264***	-0.748***	0.180
	(0.235)	(0.319)	(0.266)	(0.246)	(0.309)	(0.274)	(0.245)
Negative AI Framing × left-right	0.528***	0.468*	0.684***	0.168	-0.988***	-0.800***	0.070
	(0.203)	(0.276)	(0.230)	(0.212)	(0.267)	(0.237)	(0.212)
Constant	17.397***	17.245***	20.214***	17.718***	18.840***	17.211***	18.188***
	(0.700)	(0.950)	(0.793)	(0.732)	(0.921)	(0.816)	(0.729)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	229	229	229	229	229	229	229
Panel B: Students Who Frequently Use AI							
	Social Spending	Basic Income	Tech Tax	Retain Workers	Reduce Migration	Restrict Competition	Fund to Reskill
Positive AI Framing	0.085	-0.338	-0.025	0.178	0.431	0.307	0.053
	(0.287)	(0.388)	(0.327)	(0.337)	(0.326)	(0.335)	(0.296)
Negative AI Framing	-0.403	-0.430	-0.484	0.235	0.096	-0.138	-0.051
	(0.304)	(0.411)	(0.347)	(0.357)	(0.346)	(0.355)	(0.313)
Control × left-right	0.349	0.583	0.167	0.156	-0.158	-0.212	0.014
	(0.406)	(0.548)	(0.463)	(0.476)	(0.461)	(0.473)	(0.418)
Positive AI Framing × left-right	0.406	1.132***	0.140	0.082	-1.304***	-0.884**	0.415
	(0.317)	(0.428)	(0.362)	(0.372)	(0.360)	(0.370)	(0.326)
Negative AI Framing × left-right	0.790**	0.758*	0.801**	0.164	-1.154***	-0.569	0.616*
	(0.333)	(0.449)	(0.380)	(0.391)	(0.378)	(0.388)	(0.343)
Constant	17.221***	15.403***	17.411***	16.761***	17.873***	18.792***	18.140***
	(1.065)	(1.437)	(1.214)	(1.250)	(1.210)	(1.241)	(1.096)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	114	114	114	114	114	114	114

Notes: This table reports treatment effects on policy preferences by previous AI experience. Panel A shows results for students who sometimes use AI tools; Panel B shows results for students who frequently use AI. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, income level, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A5.0.2 Exclusion of participants failing attention check

Table A5.7: Treatment Effects on Support for Government Policies

	Left-Leaning Policies			Right-Leaning Policies			
	(1) Social Spending	(2) Basic Income	(3) Tax Tech	(4) Retain Workers	(5) Reduce Migration	(6) Restrict Competition	(7) Fund to Reskill
Positive AI Framing	0.131	-0.011	-0.163	0.088	0.329**	0.261**	0.191*
	(0.117)	(0.148)	(0.135)	(0.121)	(0.140)	(0.133)	(0.102)
Negative AI Framing	0.095	-0.072	-0.044	-0.058	0.208	0.082	0.182*
	(0.113)	(0.143)	(0.131)	(0.117)	(0.135)	(0.129)	(0.098)
Control × Left–Right	0.507***	0.591***	0.329**	0.336***	-0.885***	-0.786***	0.276***
	(0.122)	(0.154)	(0.142)	(0.126)	(0.146)	(0.139)	(0.106)
Positive AI Framing × Left–Right	0.461***	0.618***	0.608***	0.200	-1.342***	-1.023***	0.159
	(0.122)	(0.154)	(0.141)	(0.126)	(0.146)	(0.138)	(0.106)
Negative AI Framing × Left–Right	0.539***	0.563***	0.403***	0.258**	-1.097***	-0.701***	0.003
	(0.121)	(0.153)	(0.140)	(0.125)	(0.145)	(0.138)	(0.105)
Constant	16.890***	16.207***	18.313***	18.158***	18.804***	17.449***	18.451***
	(0.377)	(0.476)	(0.436)	(0.390)	(0.451)	(0.428)	(0.327)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	825	826	825	825	826	826	826

Notes: This table reports treatment effects on policy preferences excluding participants who failed the attention check on the river pollution article. The dependent variables are binary indicators of support for seven different policies to mitigate AI risks. Control variables include: age, gender, nationality, income level, and parental education and employment background. Robust standard errors are shown in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix 5B: AI-related Articles

Source: The Guardian

Full Article:

AI will affect 40% of jobs and probably worsen inequality, says IMF head. 'Crucial' that countries build social safety nets to mitigate impact on workers, says Kristalina Georgieva

Artificial intelligence will affect 40% of jobs around the world and it is "crucial" that countries build social safety nets to mitigate the impact on vulnerable workers, according to the head of the International Monetary Fund.

AI, the term for computer systems that can perform tasks usually associated

with human levels of intelligence, is poised to profoundly change the global economy with advanced economies at greater risk of disruption.

Analysis by the IMF, the international lender of last resort, says about 60% of jobs in advanced economies such as the US and UK are exposed to AI and half of these jobs may be negatively affected. But the technology will also help to enhance some humans' productivity as AI improves their performance, it said.

Horrified by Horizon? Then get ready to be totally appalled by AI. According to the IMF, the safest highly exposed jobs are those with a "high complementarity" to AI, meaning the technology will assist their work rather than displace it entirely. This includes roles with a high degree of responsibility and interacting with people – such as surgeons, lawyers and judges.

High-exposure jobs with "low complementarity" – meaning the potential for being displaced by AI – include telemarketing, or cold-calling people to offer goods or services. Low-exposure occupations include dish washers and performers, the IMF said.

AI jobs exposure is 40% in emerging market economies – defined by the IMF as states including China, Brazil and India – and 26% for low-income countries, with an overall total of just under 40%, according to the IMF.

Generative AI – the term for technology that can produce highly plausible text, images and even voice from simple hand-typed prompts – has risen up the political agenda since the emergence of tools such as the ChatGPT chatbot.

Kristalina Georgieva, the IMF managing director, said AI's ability to affect highly skilled jobs means that advanced economies face greater risks from the technology. She added that in extreme cases some jobs in major economies could disappear.

"Roughly half the exposed jobs may benefit from AI integration, enhancing productivity," wrote Georgieva, in a blogpost accompanying the IMF research. "For the other half, AI applications may execute key tasks currently performed by humans, which could lower labour demand, leading to lower wages and reduced hiring. In the most extreme cases, some of these jobs may disappear."

She added that in most scenarios AI would probably worsen overall inequality

across the global economy and could stoke social tensions without political intervention. AI is expected to feature prominently as a topic of discussion at the World Economic Forum in Davos this week, which top executives from the tech industry will attend.

"It is crucial for countries to establish comprehensive social safety nets and offer retraining programmes for vulnerable workers," Georgieva said. "In doing so, we can make the AI transition more inclusive, protecting livelihoods and curbing inequality."

The IMF analysis shows that higher-wage earners whose jobs have high complementarity with AI can expect an increase in their income, leading to an increase in inequality.

"This would amplify the increase in income and wealth inequality that results from enhanced capital returns that accrue to high earners," the IMF report said. "Countries' choices regarding the definition of AI property rights, as well as redistributive and other fiscal policies, will ultimately shape its impact on income and wealth distribution."

The report said workers in the UK, with its high proportion of graduates, might be better prepared to switch from jobs at risk of displacement to "high complementarity" jobs, although older workers may struggle to adapt and move to new jobs or retrain.

Satya Nadella, the chief executive of Microsoft, which is the biggest investor in OpenAI, the US company behind ChatGPT, said on Monday there would be jobs in the future but the question remained about what would be "the shape of these jobs". Speaking at an event hosted by the Chatham House thinktank in London, Nadella said AI could help with "mid-career transitions". He added: "I think this is the age where this is about expertise at your fingertips. So if anything, anyone can become an expert in anything because you have the AI assistant helping you."

Last year, the Organisation for Economic Co-operation and Development said the occupations at highest risk from AI-driven automation were highly skilled jobs and represented about 27% of employment across its 38 member countries, which include the UK, Japan, Germany, the US, Australia and

Canada. It said skilled professions such as law, medicine and finance were most at risk.

Negative Extract[19]

Artificial intelligence will affect 40% of jobs around the world and it is "crucial" that countries build social safety nets to mitigate the impact on vulnerable workers, according to the head of the International Monetary Fund.

AI, the term for computer systems that can perform tasks usually associated with human levels of intelligence, is poised to profoundly change the global economy with advanced economies at greater risk of disruption.

Analysis by the IMF, the international lender of last resort, says about 60% of jobs in advanced economies such as the US and UK are exposed to AI and half of these jobs may be negatively affected. [...]

But the technology will also help to enhance some humans' productivity as AI improves their performance, it said.

Horrified by Horizon? Then get ready to be totally appalled by AI. According to the IMF, the safest highly exposed jobs are those with a "high complementarity" to AI, meaning the technology will assist their work rather than displace it entirely. This includes roles with a high degree of responsibility and interacting with people – such as surgeons, lawyers and judges.

High-exposure jobs with "low complementarity" – meaning the potential for being displaced by AI – include telemarketing, or cold-calling people to offer goods or services. Low-exposure occupations include dish washers and performers, the IMF said.

AI jobs exposure is 40% in emerging market economies – defined by the IMF as states including China, Brazil and India – and 26% for low-income countries, with an overall total of just under 40%, according to the IMF.

[19]Participants will be informed that the "[...]" means a part of the text has been cut in order to make it shorter.

Generative AI – the term for technology that can produce highly plausible text, images and even voice from simple hand-typed prompts – has risen up the political agenda since the emergence of tools such as the ChatGPT chatbot.

Generative AI – the term for technology that can produce highly plausible text, images and even voice from simple hand-typed prompts – has risen up the political agenda since the emergence of tools such as the ChatGPT chatbot. Kristalina Georgieva, the IMF managing director, said AI's ability to affect highly skilled jobs means that advanced economies face greater risks from the technology. She added that in extreme cases some jobs in major economies could disappear. [...]

"Roughly half the exposed jobs may benefit from AI integration, enhancing productivity," wrote Georgieva, in a blogpost accompanying the IMF research.

"For the other half, AI applications may execute key tasks currently performed by humans, which could lower labour demand, leading to lower wages and reduced hiring. In the most extreme cases, some of these jobs may disappear."

She added that in most scenarios AI would probably worsen overall inequality across the global economy and could stoke social tensions without political intervention. AI is expected to feature prominently as a topic of discussion at the World Economic Forum in Davos this week, which top executives from the tech industry will attend.

"It is crucial for countries to establish comprehensive social safety nets and offer retraining programmes for vulnerable workers," Georgieva said. "In doing so, we can make the AI transition more inclusive, protecting livelihoods and curbing inequality." [...]

The IMF analysis shows that higher-wage earners whose jobs have high complementarity with AI can expect an increase in their income, leading to an increase in inequality.

"This would amplify the increase in income and wealth inequality that results from enhanced capital returns that accrue to high earners," the IMF report said. "Countries' choices regarding the definition of AI property rights, as well as redistributive and other fiscal policies, will ultimately shape its impact

on income and wealth distribution."

The report said workers in the UK, with its high proportion of graduates, might be better prepared to switch from jobs at risk of displacement to "high complementarity" jobs, although older workers may struggle to adapt and move to new jobs or retrain.

Satya Nadella, the chief executive of Microsoft, which is the biggest investor in OpenAI, the US company behind ChatGPT, said on Monday there would be jobs in the future but the question remained about what would be "the shape of these jobs". Speaking at an event hosted by the Chatham House thinktank in London, Nadella said AI could help with "mid-career transitions". He added: "I think this is the age where this is about expertise at your fingertips. So if anything, anyone can become an expert in anything because you have the AI assistant helping you."

Last year, the Organisation for Economic Co-operation and Development said the occupations at highest risk from AI-driven automation were highly skilled jobs and represented about 27% of employment across its 38 member countries, which include the UK, Japan, Germany, the US, Australia and Canada. It said skilled professions such as law, medicine and finance were most at risk.

Positive Extract

[...] AI, the term for computer systems that can perform tasks usually associated with human levels of intelligence, is poised to profoundly change the global economy with advanced economies at greater risk of disruption. Analysis by the IMF, the international lender of last resort, says about 60% of jobs in advanced economies such as the US and UK are exposed to AI [...]. According to the IMF, the safest highly exposed jobs are those with a "high complementarity" to AI, meaning the technology will assist their work rather than displace it entirely. This includes roles with a high degree of responsibility and interacting with people – such as surgeons, lawyers and judges. High-exposure jobs with "low complementarity" – meaning the potential for being displaced by AI – include telemarketing, or cold-calling people to offer goods or services. [...]

The IMF analysis shows that higher-wage earners whose jobs have high com-

plementarity with AI can expect an increase in their income, leading to an increase in inequality. "This would amplify the increase in income and wealth inequality that results from enhanced capital returns that accrue to high earners," the IMF report said. "Countries' choices regarding the definition of AI property rights, as well as redistributive and other fiscal policies, will ultimately shape its impact on income and wealth distribution." The report said workers in the UK, with its high proportion of graduates, might be better prepared to switch from jobs at risk of displacement to "high complementarity" jobs, although older workers may struggle to adapt and move to new jobs or retrain. Satya Nadella, the chief executive of Microsoft, which is the biggest investor in OpenAI, the US company behind ChatGPT, said on Monday there would be jobs in the future but the question remained about what would be "the shape of these jobs". Speaking at an event hosted by the Chatham House thinktank in London, Nadella said AI could help with "mid-career transitions". He added: "I think this is the age where this is about expertise at your fingertips. So if anything, anyone can become an expert in anything because you have the AI assistant helping you." [...]

Non-AI related article

Source: The Guardian

Is pollution in England's rivers really getting worse? There's more good news than you might think
Author- Michelle Jackson

Public outrage over river pollution has been heartening to see. Over the past few years, stories about sewage contamination in rivers have captured public attention, and prompted campaigns and protests, such as the forthcoming River Action UK march for clean water on 26 October in London. It is important to protect our rivers because they are biodiversity hotspots and essential for human health. However, as a freshwater ecologist, I know there is more nuance to the story than you may have been led to believe. From my perspective, there is some good news when it comes to our rivers. I would

even say that some rivers in England are in the best state they have been in for hundreds of years.

Many rivers in England are polluted, but we need to recognise that this is not an emerging issue but a much longer-standing one that has been largely ignored by the media and politicians for decades. Much of the recent furore over pollution has to do with increased awareness, rather than a sudden increase in pollution itself. It's only by understanding how these ecosystems have changed over time and reflecting on previous successes that we can make real progress.

It has been widely reported that English rivers have gone from 97% being classified as having "good chemical status" in 2016 to none in 2019. Across England, the Environment Agency monitors water quality by directly measuring some pollutants and characterising the life in the river. This data is then used to classify our rivers' water quality: "high", "good", "moderate", "poor" or "bad". However, rather than water quality actually declining this drastically within this timeframe, the way we measure it has changed – this perceived decline is an artefact of the data. From 2019, new chemicals were included in the testing, among them "ubiquitous, persistent, bioaccumulative, toxic substances", which accounts for this major change.

The media has also reported increases to the number of "sewage spills" – when raw sewage (rather than treated sewage) is pumped into our seas, lakes, and rivers. In 2016, data was only available from 862 storm overflows, with this increasing to 6,182 in 2018, 12,092 in 2020, and 13,080 in 2022. Of course, the thought of any raw sewage entering water bodies is disgusting and we should work to reduce this to zero by updating our Victorian wastewater infrastructure. The point, though, is that the problem may have not got worse in recent years – it is just that we are only now looking for it properly.

Other measures of water quality show that some things have actually improved. For instance, many invertebrates can only survive when pollution levels are low, so assessing what species are present provides a quick and cheap way of measuring water quality. Various metrics are calculated from data of invertebrate counts – and a higher score means the water is cleaner. Using this data (and associated metrics on algae, fish, hydrology and nutri-

ents), only 15% of rivers in England re currently classified as "good ecological status", which sounds pretty terrible, especially as this has declined from 24% in 2009.

However, there are a few factors to consider. First, a robust analysis of 223,300 sampling points of the invertebrate data by researchers at the UK Centre for Ecology & Hydrology, using records from 1989 to 2018, showed that invertebrate biodiversity has increased over time across England regardless of river size and land cover. Although this trend has slowed in some cases in the last 20 years, the headline here is that more wildlife is being found in our rivers now than there was three decades ago. For instance, the mayfly, stonefly, and caddisfly families (which are particularly sensitive to water quality) have increased by 300% during that time period.

Second, the way rivers are classified uses a "one out, all out" rule. This means that rivers only have to fail on one of 20 metrics to be downgraded from "good" to "moderate" ecological status. In other words, the metric with the worst score determines the overall status. So, overall water quality can be more positive than it initially seems. The increase in river invertebrate biodiversity across England is good news and it shows how far we have come in improving water quality in my lifetime.

At least in part, this can be linked to improvements in sewage treatment causing declines in nutrient pollution, with research showing a large decline in phosphate concentrations from 1997. This followed the European Union's 1991 urban waste water treatment directive, which tightened laws and forced water companies to improve the treatment of wastewater. However, nutrient levels are still higher today than they would naturally be due to sewage and agricultural practices, so there is still work to be done. Some of my group's own research has shown that additional reductions of phosphorus from improved sewage treatment would be the most effective way to further benefit biodiversity.

Water companies do need more scrutiny as there is evidence to suggest that they often still breach their licences. They are only supposed to release raw sewage when there is high rainfall. If this reaches such a critical level that the sewage would back up and risk coming out the wrong way, so to speak,

water companies are legally allowed to release the untreated waste. With the growing public concern, laws are starting to tighten on this, hence the increase in monitoring data. Only time will tell, as the data from the new network of event duration monitoring devices that are on all storm overflows (as of 2023) is analysed. While we need to limit the use of storm overflows, at the moment we cannot confidently claim that these are, in fact, being used more often now than they have been in the past. This is not to diminish the problem, only to say that it may not be a recent one.

There are still reasons to protest at events such as the River Action UK march. I will certainly be there. Despite the improvements to nutrient concentrations and biodiversity, there are emerging chemical pollutants and we have very little information on their impacts. This is where I think the concerned public should now turn their attention. These pollutants are typically household and medical products that go down the drain and are not removed at wastewater treatment plants, such as beauty products and pest treatments. Globally, freshwater ecosystems – which act as hotspots for biodiversity – are one of the most threatened ecosystems on the planet. They are also often overlooked in legislation – so don't stop campaigning, just make sure you represent the facts correctly so that we can make real actionable change.

AI-generated Article Summary

The article addresses the complexities surrounding pollution in England's rivers, suggesting that the situation is not as dire as public discourse often portrays. Freshwater ecologist Michelle Jackson argues that heightened awareness of pollution has contributed to the perception of worsening conditions, rather than a substantial increase in pollution levels. Recent changes in the methods for assessing water quality, such as the inclusion of persistent toxic substances in testing, have led to the reported decline in river health. Similarly, the rise in sewage spill reports is largely attributed to enhanced monitoring practices rather than a significant escalation in pollution events.

Despite these challenges, Jackson identifies notable improvements, particularly in river biodiversity. Sensitive species such as mayflies and stoneflies

have shown increased populations over the past decades, a positive trend linked to improvements in sewage treatment driven by European Union regulations. However, issues such as nutrient pollution from agriculture and emerging chemical contaminants remain pressing concerns. Jackson advocates for continued public engagement and protest to address these issues but stresses the importance of understanding the broader context of river health. A more nuanced perspective, she argues, will enable more informed and effective environmental action.

Appendix 5C: Survey Questionnaire

- Are you currently enrolled as a student?
- What is the level of the degree programme you are currently enrolled in?
- What is your main field of study/major (Please select the option that best matches your field of study)
- What are your plans after graduation?
- If you think about the future, which of the following jobs/sectors/roles do you think you might like doing (or in which job/sector are you employed in if you are an alumni)? Please note, this is not an exhaustive list. Please pick the best option among the listed ones.
- If you think about the future, which of the following jobs/sectors/roles do you think you might be doing (or in which job/sector are you employed in if you are an alumni)?
- In politics people sometimes talk of "left" and "right" to summarize their political views. On a scale from 0 to 10, where 0 means being at the most left and 10 at the most right, where would you place yourself?
- People have different ideas about what is the space for government in society. Do you agree or disagree with the idea that governments should reduce differences in income levels?

- Many people do not go to vote for many reasons. Some are not eligible (for instance, they are too young), some are busy, some do not believe their vote will matter, and others simply cannot make it for personal reasons. In the last national election of (Insert country), do you recall voting?
- When it comes to the employment (working) conditions, which of the following options best corresponds with the situation of your mother/caretaker?
- When it comes to the employment (working) conditions, which of the following options best corresponds with the situation of your father/caretaker?
- What is the highest level of education your mother or caretaker has completed?
- What is the highest level of education your father or caretaker has completed?
- Think of this ladder as representing where people stand in society. At the top of the ladder are the people who have most money, most education and the best jobs. At the bottom of the ladder are the people who have the least money, the lowest education and the worst jobs (or are unemployed). The higher one is in this ladder, the closer one is to the people who are at the top of society. The lower one is, the closer to the people at the bottom. Where would you place your family on this ladder, where 0 represents the lowest position and 10 the highest position?
- How important is it for you to live in a country that is governed democratically? On this scale where 0 means it is "not at all important" and 10 means "absolutely important", what position would you choose? If you are paying attention, please select the option two regardless of your actual attitude.
- Governments need to make policies in many different sectors. Which of the following policy sector you think is the most important for (insert country) right now?

- And what would be, for you, the second most important policy sector for (insert country) right now?
- Some people believe that the government should take measures to accelerate the pace at which new technologies are incorporated in workplaces. Other people think that the government should slow down the pace of adoption of new technologies. What do you think should the government of (insert country) do?
- How much do you agree or disagree with the following government policies? (Conditions: a. Expand social spending to support laid-off workers, and workers in similar positions b. Implement a basic income that gives every adult a set amount of money from the government on a regular basis. c. Pay to retain displaced workers and guarantee them jobs. d. Reduce the number of unskilled immigrants entering the country for work e. Restrict international competition by increasing trade barriers on goods and services to the (insert country). f. Fund programs to re-skill workers for new jobs. g. Directly tax companies that replace workers with machines, robots or AI)
- How much you trust different types of elites to navigating the challenges of AI?
- Thinking about the next national election of (insert country), how likely do you think it is that you will vote? Please use the following slider that goes from extremely unlikely to extremely likely.
- Most jobs consist of a range of different tasks. Thinking of your future career, how relevant you think each of the following task will be for your job position? Please use a scale where 0 means not relevant at all, and 10 means fundamental.
- Once again thinking about your future job (or one job option linked to your studies and that you think you'd be happy with), please consider the balance between writing tasks (summaries, briefings, but also E-mails) and technical ones (skills that are linked to your field and you have learned through your studies, including topic knowledge)... Which one you think will be more important for your future job?

- Thinking about the current advancements in AI capabilities, how do you feel about your job prospects? Please select the two adjectives that best describe how you feel. (Conditions: a. Excited b. Anxious c. Angry d. Hopeful)
- Do you agree or disagree with the following statements. AI will make future jobs: (Conditions: a. More creative b. Harder to find c. More complex d. Easier)
- You previously read the article on the quality and biodiversity of English rivers. Thinking back to it, can you recall what proportion of rivers in England is considered of "good ecological status"?
- You previously read the extract from an article on the potential impact of AI on labour markets across the world. Can you recall, according to the article, what proportion of jobs will likely be affected by AI in advanced economies?
- During your studies, how often would you say you have used generative AI tools such as ChatGPT or BARD?

6

Conclusion

Artificial Intelligence represents a fast-paced technological frontier that is significantly altering labour markets and the nature of work. Rather than resulting in economy-wide labour displacement, AI's economic impact primarily unfolds through changes in how tasks are performed, skills are deployed and work is organized within firms. This digital transformation poses new challenges for economies in retaining and developing skills, sustaining productivity growth and designing effective policies to support labour-market adjustments and workforce development.

This thesis explores the interplay between AI and the future of work by investigating how AI influences wage and employment dynamics, task allocation and productivity across different institutional and organizational contexts. By combining multiple data sources with diverse empirical approaches, it contributes to the literature on technological change by studying the mechanisms through which AI can affect labour markets across workers, firms and prospective labour-market entrants. A central contribution of this thesis is to demonstrate that AI's labour-market implications depend not on technological capability alone, but also critically on how AI systems are designed, adopted and integrated within workplaces.

A central discussion in the debate on AI and the future of work concerns

which workers are most affected by advances in new technologies. The first part of the thesis is therefore dedicated to understanding the potential implications of AI by examining the extent to which AI technologies can be applied to specific job tasks. **Chapter 2** contributes to this debate by using task-based technology exposure measures constructed by Webb (2020) to assess occupational susceptibility to AI and provides a descriptive comparison of observed patterns with exposure to two earlier technologies: software and robots. By linking these exposure scores to detailed employment biographies of workers in Germany, we investigate how AI exposure relates to employment and wage outcomes of workers in occupations with different task types and skill requirements. We observe that workers in highly skilled, analytical and non-routine occupations are more exposed to AI. This is in contrast with exposure to software and robotics, which are more strongly associated with routine and middle-skill occupations, suggesting that AI differs from earlier technological waves in terms of the types of workers that can potentially be impacted. Turning to labour-market outcomes, we find no evidence of changes in employment in occupations with higher AI exposure, indicating that AI susceptibility is not associated with job displacement between 2012-2019. Moreover, we demonstrate that while AI exposure is positively associated with wage growth, especially for workers in occupations characterized by a high routine task content, wage premium of high skills is unrelated to AI. Taken together, these findings suggest that AI has the potential to affect a broader set of tasks and a more diverse group of workers than previous technological waves, with potentially more widespread distributional consequences across the labour market.

With Chapter 2 investigating exposure to AI, a critical question remains: when firms actually adopt AI technologies, how does deployment unfold over time in practice and what happens to workers performing the automated tasks? **Chapter 3** explores these dynamics by studying an actual AI adoption within a high-tech manufacturing company. We analyze the deployment of an AI-based visual inspection system that automates parts of the quality-assurance process in an industrial production environment. The AI autonomously carries out a subset of decision-making tasks, taking over classifications that were previously performed by human operators. Drawing on detailed decision-level

data from both human operators and the AI system across 14 plants over a 15-month period, we examine how automation expands over time and how this expansion relates to labour demand within the affected inspection task. We show that AI adoption proceeds gradually rather than exponentially and we find substantial heterogeneity across plants in both the level and pace of deployment. Our analysis reveals a negative association between AI automation and operator activity: as automation intensifies, the volume of decisions performed by human operators declines, consistent with task-level substitution effects. However, we do not find any evidence that AI automation relates to changes in overall production activity. These findings suggest that technical capability alone does not guarantee rapid adoption: even when the AI system is technically capable of fully replacing human tasks, adoption unfolds incrementally in safety-critical environments where inspection accuracy and reliability are paramount. In such contexts, organizational governance, risk management and the need to build sustained trust in algorithmic decisions remain critical in shaping both the pace and the scope of automation.

Chapter 4 turns to the augmentative potential of AI and investigates how AI-enhanced training can improve worker performance in a call center in Europe. Exploiting the staggered introduction of an AI- augmented training tool, we provide evidence on the productivity effects of AI for call-center agents. The AI tool we study enhances the training process by analyzing recorded calls and generating precise, data-driven performance assessment of worker performance. This enables trainers to provide more targeted and individualized coaching, improving the quality and consistency of feedback that workers receive during their training. We find that AI significantly improves worker productivity: workers who receive AI-augmented training handle calls faster than those who receive standard training. These productivity gains are particularly pronounced for short-tenured workers, suggesting that AI helps less-experienced workers move up the experience curve more rapidly. For more experienced agents, we find AI also enables them to optimize their communication strategies and fine-tune conversation quality, further refining their expertise. Our results highlight that when AI is deployed to augment workers, especially in a context where it fosters

workplace training and learning, it can generate substantial productivity benefits and strengthen worker capabilities across different skill levels.

Lastly, **Chapter 5** shifts the focus from workers and firms to those yet to enter the labour market: university students who will enter the labour market at the next stage of their life. Given that these prospective workers will experience increasingly AI-intensive labour markets throughout their careers, this chapter studies how perceived competition with AI and information about its potential benefits and challenges shape their policy preferences. Using a pre-registered survey experiment with university students in the UK and the US, our analysis reveals a striking disconnect between emotional responses and policy demands. We find that experience with AI competition and negative narratives about its labour-market consequences significantly increases anxiety about future job prospects. However, this heightened concern does not translate into support for traditional risk-mitigating policies such as expanded social spending, universal basic income, or active labour market programs. Instead, students' policy preferences remain largely anchored in their pre-existing ideological orientations, with left-leaning students favoring redistributive measures and right-leaning students preferring market-oriented or protectionist responses. This contrasts sharply with research on job incumbents, who typically demand stronger welfare protections when confronted with technological threats. Our findings therefore suggest that young citizens may distrust government capacity to address challenges brought on by AI and perceive traditional policies as benefiting incumbents rather than new entrants.

In sum, findings of this thesis emphasize that the economic effects of AI are shaped less by its technological capabilities and more so by how various actors in the labour markets design, govern and respond to AI. While Chapter 2 shows that AI has the potential to reach to a broader and more diverse part of the workforce than earlier automation technologies, Chapters 3 and 4 document that its labour-market consequences crucially depend on the way AI systems are used. At the firm level, Chapter 3 shows that even technically mature AI systems are adopted cautiously in safety-critical environments, with organizational oversight and trust constraining the pace of automation and task replacement. At the worker level, Chapter 4 demonstrates that AI can enhance produc-

tivity when implemented to improve on-the-job learning rather than to replace tasks. Chapter 5 extends this perspective to future labour-market entrants, revealing that heightened concern about AI does not automatically translate into demand for protective policies. Taken together, rather than operating through a single dominant channel, findings from these chapters indicate that the transformative effects of AI can materialize through a series of context-specific incentives and decisions about how, where and for which aim AI technologies are designed and deployed.

Limitations and avenues for future research

While this thesis provides new evidence on how AI reshapes work and workers' labour market outcomes, it also highlights several important limitations and avenues for future research.

First, the current understanding of the long-term consequences of AI remains limited, particularly considering the early stages of development and deployment of latest AI systems and generative AI applications. While the trajectory of AI capabilities can be anticipated in the short to medium term, there is far less certainty about which tasks will remain beyond AI's reach in the long run. Although existing exposure measures, such as those used in Chapter 2 provide a useful way to highlight AI's potential to transform work across diverse occupational groups, they lack the nuance to capture the complexity and shifts in task changes and outcomes over time. Further research would benefit from combining longitudinal labour-market data on employment and wage trends with more refined, task-based frameworks that capture the dynamic nature of tasks, in order to assess AI's long-term implications for occupational restructuring and labour-market adjustments. In this regard, the expertise framework of Autor and Thompson (2025) provides a finer-grained lens, as it conceptualizes tasks as context-dependent and evolving—where the same task may require high expertise in one occupation but not in another and where automation can both eliminate existing tasks and create new ones. This dynamic can explain seemingly contradictory wage patterns: workers in AI-exposed occupations may initially experience earnings pressure as some components of their jobs are automated but subsequently see wage gains as they specialize in tasks that remain distinctly human or as new domains of expertise emerge (Autor and Thompson, 2025). Following this view, AI can redistribute work over

time, underscoring the importance of understanding how the value of human expertise shifts within tasks and occupations when assessing AI's potential long-run labour-market effects.

Second, advancing academic research on AI's implications on work requires robust causal evidence on how AI adoption affects workers, labour demand and productivity in workplaces. In practice, however, AI implementation rarely happens as a clean or exogenous intervention: it is embedded in complex organizational environments and often coincides with broader changes in workflows or production processes, making causal identification difficult (Fregin et al., 2025). As demonstrated in Chapter 3, where AI deployment timing is endogenous and suffers from spillovers, isolating AI's causal impact remains challenging even with detailed firm-level data. Moreover, AI adoption within firms can unfold gradually, raising questions about the applicability of quasi-experimental techniques that rely on clear treatment–control distinctions or sharp cutoffs. Nevertheless, as illustrated in Chapter 4, context-specific strategies, particularly how organizations structure their rollout, can still enable credible causal identification. Future research could therefore develop adaptable research designs, metrics and data collection strategies that examine AI adoption alongside organizational change, recognizing both the challenges and opportunities for causal inference in these complex settings. More research is needed to obtain detailed knowledge of how different stages of AI deployment unfolds within firms, for example when deployment decisions occur, how systems are adjusted over time and how AI is integrated in existing work processes. This often requires close collaboration with organizations, where internal data and institutional knowledge can help building applicable research designs and identification strategies to recover treatment effects. Such close collaboration, however, can be often demanding as researchers must navigate firm-specific constraints, organizational dynamics and data limitations that are not always fully under their control. In addition, given the complexity of AI adoption processes, future work would benefit from combining quantitative methods with qualitative approaches, such as interviews employed in Chapter 4, to capture dimensions of AI's impact that may not be observable through administrative or performance data alone.

Lastly, a further priority for future research lies in understanding the mechanisms through which AI can be steered toward augmenting, rather than replacing, human labour. This is particularly important for experimental research, which can go beyond estimating productivity effects in narrow tasks to investigate the conditions under which AI can complement workers across a diverse set of tasks and the design choices that shape these outcomes. As demonstrated in Chapter 4, when embedded in a training processes, AI can complement human skills and accelerate learning. In such settings, AI can help workers develop competencies that extend beyond AI's current capabilities, such as enhancing judgment, collaboration and creative problem-solving, thereby fostering higher-value work, human expertise and workforce adaptation (Autor, 2024). Advancing this research agenda requires a shift in emphasis from measuring AI's capabilities in isolation to evaluating its progress relative to human performance and learning. This entails collecting richer data on workers themselves; particularly on how workers learn alongside AI, which skills and competences they develop in AI-augmented settings and how teams reorganize as AI automates specific tasks or functions. Ultimately, understanding AI's role as a complement to human work depends on benchmarking AI metrics in ways that strengthen workers' comparative advantages and well-being, rather than focusing solely on advancing the performance of AI systems.

Bibliography

Abd Al Rahman, M. and Mousavi, A. (2020). A review and analysis of automatic optical inspection and quality monitoring methods in electronics industry. *Ieee Access*, 8:183192–183271.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Acciani, G., Brunetti, G., and Fornarelli, G. (2006). Application of neural networks in optical inspection and classification of solder joints in surface mount technology. *IEEE Transactions on industrial informatics*, 2(3):200–209.

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of economic literature*, 40(1):7–72.

Acemoglu, D. (2021). Harms of ai. Technical report, National Bureau of Economic Research.

Acemoglu, D., Aghion, P., Lelarge, C., Van Reenen, J., and Zilibotti, F. (2007). Technology, information, and the decentralization of the firm. *The Quarterly Journal of Economics*, 122(4):1759–1799.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2020a). Robots and jobs: Evidence from us labor markets. *Journal of political economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2020b). Unpacking skill bias: Automation and new tasks. In *AEA Papers and Proceedings*, volume 110, pages 356–361.

Agarwal, N., Moehring, A., Rajpurkar, P., and Salz, T. (2023). Combining human expertise with artificial intelligence: Experimental evidence from radiology. Technical Report 31422, National Bureau of Economic Research.

Ahrens, L. (2024). Labor market risks and welfare preferences: a bounded rationality approach. *Socio-Economic Review*, 22(1):327–348.

Albanesi, S., Dias da Silva, A., Jimeno, J. F., Lamo, A., and Wabitsch, A. (2025). New technologies and jobs in europe. *Economic Policy*, 40(121):71–139.

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour Economics*, 71:102002.

Anelli, M., Colantone, I., and Stanig, P. (2023). Individual vulnerability to industrial robot adoption increases support for the radical right. *PNAS*, 118(47):1–8.

Antoni, M., Ganzer, A., Vom Berge, P., et al. (2016). Sample of integrated labour market biographies (siab) 1975-2014. *FDZ-Datenreport*, 4:2016.

Arntz, M., Gregory, T., and Zierahn, U. (2017). Revisiting the risk of automation. *Economics Letters*, 159:157–160.

Asphjell, M. K., Letterie, W., Nilsen, Ø. A., and Pfann, G. A. (2014). Sequentiality versus simultaneity: Interrelated factor demand. *Review of Economics and Statistics*, 96(5):986–998.

Aubert, P., Caroli, E., and Roger, M. (2006). New technologies, organisation and age: firm-level evidence. *The Economic Journal*, 116(509):F73–F93.

Autor, D. (2020). The faltering escalator of urban opportunity. *Securing our Economic Future*, pages 108–136.

Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, 139(3):1399–1465.

Autor, D. and Thompson, N. (2025). Expertise. *Journal of the European Economic Association*.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the "other 99 percent". *Science*, 344(6186):843–851.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, 29(3):3–30.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review*, 103(5):1553–1597.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: have computers changed the labor market? *The Quarterly journal of economics*, 113(4):1169–1213.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Bachmann, R., Cim, M., and Green, C. (2019). Long-Run Patterns of Labour Market Polarization: Evidence from German Micro Data. *British Journal of Industrial Relations*, 57(2):350–376.

Bartel, A., Ichniowski, C., and Shaw, K. (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics*, 122(4):1721–1758.

Beckmann, M. (2007). Age-biased technological and organizational change: Firm-level evidence and management implications. *WWZ Discussion Paper*, (05/07).

Behaghel, L. and Greenan, N. (2007). Training and age-biased technical change. *LEA Working Paper*, (0705).

Bellmann, L., Lochner, B., Seth, S., and Wolter, S. (2020). Akm effects for german labour market data. Technical report, Institut für Arbeitsmarkt-und Berufsforschung (IAB), Nürnberg [Institute for ....

Bessen, J., Goos, M., Salomons, A., and Van den Berge, W. (2023). Automatic reaction-what happens to workers at firms that automate? *The Review of Economics and Statistics*, (Feb. 6, 2023).

Bloom, N., Garicano, L., Sadun, R., and Van Reenen, J. (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science*, 60:2859–2885.

Borland, J. and Coelli, M. (2017). Are robots taking our jobs? *Australian Economic Review*, 50(4):377–397.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

Borwein, S., Bonikoswski, B., Loewen, P., Lee-Whiting, B., and Magistro, B. (2024). Perceived technological threat and vote choice: evidence from 15 european democracies. *West European Politics*, 48(3):534–561.

Brynjolfsson, E., Chandar, B., and Chen, R. (2025a). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab, Stanford, CA.

Brynjolfsson, E., Li, D., and Raymond, L. (2025b). Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6370):1530–1534.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? In *AEA papers and proceedings*, volume 108, pages 43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230. Themed Issue: Treatment Effect 1.

Card, D. and DiNardo, J. E. (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of labor economics*, 20(4):733–783.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015.

Cazzaniga, M., Jaumotte, F., Lonji, L., Melina, G., Panton, A., Pizzinelli, C., Rockall, E., and Tavares, M. (2024). Gen-ai: Artificial intelligence and the future of work. *Staff Discussion Notes*, January.

CDT (2025). Teen and young adult perspectives on generative ai: Patterns of use, excitements, and concerns. *Center for Digital Thriving: Harvard Graduate School of Education*.

Cedefop (2025). Skills empower workers in the ai revolution first findings from cedefop's ai skills survey. publication office of the european union. *European Centre for the Development of Vocational Training Policy Brief*.

Choi, J. H. and Schwarcz, D. B. (2023). AI Assistance in Legal Analysis: An Empirical Study. *SSRN Electronic Journal*.

Council, J. (2019). Metlife says AI is improving its call centers. *The Wall Street Journal, WSJ PRO*.

Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2025). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. Technical report, SSRN.

Dalton, R. J. (2020). *The Good Citizen: How a Younger Generation Is Reshaping American Politics*. Washington, D.C.: CQ Press.

Daly, S. J., Wiewiora, A., and Hearn, G. (2025). Shifting attitudes and trust in ai: Influences on organizational ai adoption. *Technological Forecasting and Social Change*, 215:124108.

Dauth, W. and Eppelsheimer, J. (2020). Preparing the sample of integrated labour market biographies (SIAB) for scientific analysis: a guide. *Journal for Labour Market Research*.

De Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964—-96.

De Grip, A. and Sauermann, J. (2012). The Effects of Training on Own and Co-worker Productivity: Evidence from a Field Experiment. *Economic Journal*, 122(560):376–399.

de Koning, J. and Gelderblom, A. (2006). Ct and older workers: no unwrinkled relationship. *International Journal of Manpower*, 27(5):467–490.

del Rio-Chanona, R. M., Ernst, E., Merola, R., Samaan, D., and Teutloff, O. (2025). Ai and jobs. a review of theory, estimates, and evidence. *arXiv preprint arXiv:2509.15265*.

Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., and Lakhani, K. (2025). The cybernetic teammate: A field experiment on generative AI reshaping teamwork and

expertise. Working Paper 33641, National Bureau of Economic Research. NBER Working Paper No. 33641.

Dell'Acqua, F., McFowland, E. I., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Technical Report Working Paper No. 24-013, Harvard Business School Technology & Operations Management Unit. The Wharton School Research Paper, also available at SSRN.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Dengler, K., Matthes, B., and Paulus, W. (2014). Occupational tasks in the german labour market. *FDZ Methodenreport*, 12:1–36.

Dillon, E. W., Jaffe, S., Immorlica, N., and Stanton, C. T. (2025). Shifting work patterns with generative AI. Working Paper 33795, National Bureau of Economic Research. NBER Working Paper No. 33795.

Dinas, E. (2012). The formation of voting habits. *Journal of Elections, Public Opinion & Parties*, 22(4):431–456.

Domini, G., Grazzi, M., Moschella, D., and Treibich, T. (2021). Threats and opportunities in the digital era: Automation spikes and employment dynamics. *Research Policy*, 50(7):104137.

Domini, G., Grazzi, M., Moschella, D., and Treibich, T. (2022). For whom the bell tolls: the firm-level effects of automation on wage and gender inequality. *Research Policy*, 51(7):104533.

Ehret, S. (2022). Public preferences for governing ai technology: Comparative evidence. *Journal of European Public Policy*, 29(11):1779–1798.

Eijkenboom, D., Fregin, M.-C., de Grip, A., Horck, S., and Steens, S. (2025). Hybrid human-ai coaching: experiences of coaches and agents in a customer service environment. *Coaching: An International Journal of Theory, Research and Practice*, pages 1–19.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Espinosa, M. and Stanton, C. T. (2022). Training, communications patterns, and spillovers inside organizations. Working Paper 30224, National Bureau of Economic Research. NBER Working Paper No. 30224.

Eurofound (2018). Automation, digitisation and platforms: Implications for work and employment. *Publications Office of the European Union, Luxembourg*.

European Commission (2024). Special eurobarometer 554: Artificial intelligence and the future of work. April–May 2024.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

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

Felten, E. W., Raj, M., and Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *NYU Stern School of Business*.

Firpo, S., Fortin, N., and Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3):953–973.

Fogliato, R., Chappidi, S., Lungren, M., Fisher, P., Wilson, D., Fitzke, M., Parkinson, M., Horvitz, E., Inkpen, K., and Nushi, B. (2022). Who goes first? influences of human-AI workflow on decision making in clinical imaging. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, pages 1362–1374, New York, NY, USA. Association for Computing Machinery.

Freeman, B. S., Arriola, K., Cottell, D., Lawlor, E., Erdman, M., Sutherland, T., and Wells, B. (2024). Evaluation of task specific productivity improvements using a generative AI personal assistant tool. *arXiv*. arXiv:2409.14511 [cs.HC].

Freeman, J. (2025). Student generative ai survey 2025. *Higher Education Policy Institute*.

Fregin, M.-C., Rounding, N., van der Velden, R., Levels, M., Özgül, P., and Steens, S. (2025). Experiences with experimental research lessons from a co-creative research project on ai effects in companies.

Frey, C. B. and Osborne, M. (2013). The future of employment.

Frey, C. B. and Osborne, M. (2017a). Technological forecasting and social change the future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114(3):254–280.

Frey, C. B. and Osborne, M. A. (2017b). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280.

Frodermann, C., Ganzer, A., Schmucker, A., and Vom Berge, P. (2021). Sample of integrated labour market biographies regional file (siab-r) 1975-2019. Technical report.

Gallego, A., Kuo, A., Manzano, D., and Fernández-Albertos, J. (2021). Technological risk and policy preferences. *Comparative Political Studies*, 55(1):60–92.

Gallego, A. and Kurer, T. (2022). Automation, digitalization, and artificial intelligence in the workplace: implications for political behavior. *Annual Review of Political Science*, 25(1):463–484.

Gallego, A., Kurer, T., and Schöll, N. (2022). Neither left behind nor superstar: Ordinary winners of digitalization at the ballot box. *The Journal of Politics*, 84(1):418–436.

Gathmann, C., Grimm, F., and Winkler, E. (2024). Ai, task changes in jobs, and worker reallocation. Technical report, CESifo Working Paper.

Genz, S., Gregory, T., Lehmer, F., Janser, M., and Matthes, B. (2021). How do workers adjust when firms adopt new technologies?

Georgieff, A. and Hyee, R. (2022). Artificial intelligence and employment: New cross-country evidence. *Frontiers in artificial intelligence*, 5:832736.

Geys, B., Heggedal, T.-R., and Sørensen, R. J. (2022). Age and vote choice: Is there a conservative shift among older voters? *Electoral Studies*, 78:102485.

Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K., and Troszyński, M. (2025). *Generative AI and jobs: A refined global index of occupational exposure*. Number 140. ILO Working Paper.

Goldin, C. D. and Katz, L. F. (2008). *The race between education and technology*. Harvard University Press.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The review of economics and statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2009). Job Polarization in Europe. *The American Economic Review*, 99(2):58–63.

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

Goos, M., Rademakers, E., Salomons, A., and Vandeweyer, M. (2022). Job polarization: Its history, an intuitive framework and some empirical evidence. *Oxford University Press, The Oxford Handbook of Job Quality*.

Gosnell, G. K., List, J. A., and Metcalfe, R. D. (2020). The impact of management practices on employee productivity: A field experiment with airline captains. *Journal of Political Economy*, 128(4):1195–1233.

Grimm, F. and Gathmann, C. (2022). The diffusion of digital technologies and its consequences in the labor market.

Guess, A. and Coppock, A. (2020). Does counter-attitudinal information cause backlash? results from three large survey experiments. *British Journal of Political Science*, 50(4):1497–1515.

Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market.

Haslberger, M., Gingrich, J., and Bhatia, J. (2025). Rage against the machine? generative ai exposure, subjective risk, and policy preferences. *Journal of European Public Policy*, pages 1–28.

Heinrich, T. and Witko, C. (2025). Self-interest and preferences for the regulation of artificial intelligence. *Journal of Information Technology & Politics*, 22(3):306–321.

Heß, P., Janssen, S., and Leber, U. (2023). The effect of automation technology on workers' training participation. *Economics of Education Review*, 96:102438.

Humlum, A. and Vestergaard, E. (2025). Large language models, small labor market effects. Technical report, National Bureau of Economic Research.

Im, Z. J. (2020). Automation risk and support for welfare policies: how does the threat of unemployment affect demanding active labour market policy support? *Journal of International and Comparative Social Policy*, 37(1):76–91.

Im, Z. J. (2025). Technological change and support for redistributive politics. *Journal of European Public Policy*, pages 1–26.

Iscenko, Z. and Curto Millet, F. (2026). Looking for the ladder: Is AI Impacting Entry-Level Jobs? Working paper, Economic Innovation Group.

Jones, B. (1999). Bounded rationality. *Annual Review of Political Science*, 2:297–321.

Kanazawa, K., Kawaguchi, D., Shigeoka, H., and Watanabe, Y. (2025). AI, skill, and productivity: The case of taxi drivers. *Management Science*. Published online: June 9, 2025.

Kinder, M., Gimbell, M., Kendall, J., and Lee, M. (2025). New data show no ai jobs apocalypse—for now. *Brookings Institution*.

Klein Teeselink, B. (2025). Generative ai and labor market outcomes: Evidence from the united kingdom. *Available at SSRN*.

Kurer, T. (2020). The declining middle: Occupational change, social status, and the populist right. *Comparative Political Studies*, 53(10):1798–1835.

Kurer, T. and Gallego, A. (2019). Distributional consequences of technological change: Worker-level evidence. *Research & Politics*, 6(1):2053168018822142.

Kurer, T. and Häusermann, S. (2022). Innovation and intellectual property rights. In Marius, B., Achim, K., Kees, V. K., and Paul, M., editors, *Digitalization and the Welfare State*, chapter 8, pages 139–156. Oxford University Press, Oxford.

Lane, M. and Saint-Martin, A. (2021). The impact of artificial intelligence on the labour market: What do we know so far? *OECD Social, Employment, and Migration Working Papers*, (256):0_1–60.

Liu, X. and Batt, R. (2007). The economic pay-offs to informal training: Evidence from routine service work. *ILR Review*, 61(1):75–89.

Luo, X., Qin, M. S., Fang, Z., and Qu, Z. (2021). Artificial intelligence coaches for sales agents: Caveats and solutions. *Journal of Marketing*, 85(2):14–32.

Magistro, B., Loewen, P., Bonikowski, B., Borwein, S., and Lee-Whiting, B. (2024). Attitudes toward automation and the demand for policies addressing job loss: the effects of information about trade-offs. *Political Science Research and Methods*, 12(4):783–798.

Mann, K. and Püttmann, L. (2023). Benign effects of automation: New evidence from patent texts. *Review of Economics and Statistics*, 105(3):562–579.

Marken, S. and Nicola, T. (2023). Three in four americans believe ai will reduce jobs. Gallup. Accessed January 27, 2023.

Milkman, R. (2017). A new political generation: Millennials and the post-2008 wave of protest. *American sociological review*, 82(1):1–31.

Morris, M. and Viswanath, V. (2000). Age differences in technology adoption decisions: Implications for a changing work foce. *Personnel Psychology*, 53(2):375–403.

Nedelkoska, L. and Quintini, G. (2018). Automation, skills use and training. Technical Report 202, Organisation for Economic Co-operation and Development (OECD).

Neundorf, A. and Smets, K. (2017). Political socialization and the making of citizens. In *Oxford Handbook Topics in Politics*. Oxford University Press.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

OECD (2023). Oecd employment outlook 2023: Artificial intelligence and the labour market.

Otis, N., Clarke, R., Delecourt, S., Holtz, D., and Koning, R. (2024). The uneven impact of generative ai on entrepreneurial performance.

Ozgul, P., Fregin, M.-C., Stops, M., Janssen, S., and Levels, M. (2024). High-skilled human workers in non-routine jobs are susceptible to ai automation but wage benefits differ between occupations.

Paulus, W., Matthes, B., et al. (2013). The german classification of occupations 2010: structure, coding and conversion table.

Peede, L. and Stops, M. (2024). Artificial intelligence technologies, skills demand and employment: evidence from linked job ads data. Technical report, IAB-Discussion Paper.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.

Prasad, R. P. (2012). *Surface mount technology: principles and practice*. Springer Science & Business Media.

Prytkova, E., Petit, F., Li, D., Chaturvedi, S., and Ciarli, T. (2024). The employment impact of emerging digital technologies.

Rammer, C., Fernández, G. P., and Czarnitzki, D. (2022). Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Research Policy*, 51(7):104555.

Renée, L. (2025). The long-term effects of career guidance in high school and student financial aid: Evidence from a randomized experiment. *American Economic Journal: Applied Economics*, 17(2):165–183.

Riedl, C. and Bogert, E. (2024). Effects of ai feedback on learning, the skill gap, and intellectual diversity. Technical report.

Rounding, N., Arif, L. S., Berg, J., Cals, J., De Boer, D., De Bont, E., Dijksman, S., Findyartini, A., Fouarge, D., Fregin, M.-C., et al. (2025). Impact of llm assistance on physician decision-making: A multi-country randomized controlled trial. *medRxiv*, pages 2025–08.

Rueda, D. and Thewissen, S. (2019). Automation and the welfare state: Technological change as a determinant of redistribution preferences. *Comparative Political Studies*, 52(2):171–208.

Sacchi, S., Guarascio, D., Vannutelli, S., Careja, R., Emmenegger, P., and Giger, N. (2020). Risk of technological unemployment and support for redistributive policies. *The European Social Model under Pressure: Liber Amicorum in Honour of Klaus Armingeon*, pages 277–295.

Sauermann, J. A. N. (2023). Performance measures and worker productivity decision-making in policy and management. (April):1–12.

Shin, J. (2018). Relative deprivation, satisfying rationality, and support for redistribution. *Social Indicators Research*, 140:35–56.

Society for Human Resource Management (SHRM) (2025). From adoption to empowerment: Shaping the ai-driven workforce of tomorrow. Accessed: 2025-10-23.

Storm, E., Gonschor, M., and Schmidt, M. J. (2025). Ai in demand: How expertise shapes its (early) impact on workers.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Svanberg, M., Li, W., Fleming, M., Goehring, B., and Thompson, N. (2024). Beyond ai exposure: which tasks are cost-effective to automate with computer vision? *Available at SSRN 4700751*.

Sverke, M., Hellgren, J., and Näswall, K. (2002). No security: a meta-analysis and review of job insecurity and its consequences. *Journal of Occupational Health Psychology*, 7(3):242.

Trope, Y. and Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2):440.

Twenge, J. M. (2010). A review of the empirical evidence on generational differences in work attitudes. *Journal of Business and Psychology*, 25(2):201–210.

Twenge, J. M., Campbell, W. K., and Freeman, E. C. (2012). Generational differences in young adults' life goals, concern for others, and civic orientation, 1966–2009. *Journal of Personality and Social Psychology*, 102(5):1045.

Wagner, M. and Kritzinger, S. (2012). Ideological dimensions and vote choice: Age group differences in austria. *Electoral studies*, 31(2):285–296.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Manuscript / Working Paper.

Wheeler, L., Garlick, R., Johnson, E., Shaw, P., and Gargano, M. (2022). Linkedin(to) job opportunities: Experimental evidence from job readiness training. *American Economic Journal: Applied Economics*, 14(2):101–125.

Wood, T. and Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41(1):pp. 135–163.

Yu, J. and Qi, C. (2024). The impact of generative ai on employment and labor productivity. *Review of Business*, 44(1):53–67.

Zhang, B. B. (2022). Tno rage against the machines: Threat of automation does not change policy preferences. *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*.

Zimbardo, P. G. and Boyd, J. N. (2014). Putting time in perspective: A valid, reliable individual-differences metric. In *Time perspective theory; review, research and application: Essays in honor of Philip G. Zimbardo*, pages 17–55. Springer.

Impact

Advances in AI are rapidly transforming labour markets and economies worldwide. These developments hold the potential to enhance productivity and generate economic opportunities, yet they simultaneously raise concerns about job displacement, deskilling, widening income inequality, and disruptions to established patterns of work. Together, these advances demonstrate that the economic implications of AI extend beyond labour markets and carry important societal implications. Moreover, they raise fundamental questions for policymakers: how can necessary actions be taken to prepare for the future of work, and how can labour markets, training systems, and talent pathways adapt to keep pace with ongoing advances in AI?

This thesis studies how AI reshapes work and workers' labour market outcomes by assessing how AI influences employment and wage dynamics, task reallocation, and worker productivity. Across four chapters, it provides empirical evidence on how AI affects workers, firms, and prospective labour market entrants, highlighting the mechanisms through which technological change translates into changes in labour market outcomes and policy preferences. The findings show that the labour market effects of AI are not technologically predetermined but instead depend critically on how AI systems are designed, deployed, and integrated within organizational and institutional contexts. By documenting how context and implementation play a key role in shaping labour market outcomes brought about by AI, this thesis aims to provide valuable insights to inform current policy debates on fostering human-centered AI, workforce adaptability, and inclusive labour markets.

Chapter 2 examines which workers are most exposed to AI and how this exposure relates to employment and wage outcomes. Using a task-based measure on occupational AI susceptibility and administrative employment data from Germany, we document substantial heterogeneity in AI exposure across workers in occupational and skill groups. The findings show that AI exposure is concentrated among high-skilled workers in non-routine, analytical occupations, and that higher AI exposure is as-

sociated with heterogeneous wage growth patterns across occupations, with no evidence of displacement. The findings of this chapter suggest that policy responses to AI should not solely focus on predicting job losses and quantity of jobs impacted, but also on understanding how AI transforms the quality of the jobs and the nature of work itself. Assessing how work content and skill requirements evolve within occupations as well as changes in wages, working conditions and worker well-being is essential. This broader perspective is necessary to address AI's impact on inequality, as its distributional consequences across groups of workers depend on how task demands and compensation structures shift in response to AI adoption. To this end, continuous monitoring of labour market outcomes for at-risk and vulnerable worker groups is essential to identify where AI generates adjustment pressures and where it has the potential to create new opportunities. Tracking changes in task composition and skill demands within jobs allows policymakers to anticipate emerging skill bottlenecks and design timely training and reskilling strategies. The results further indicate that AI exposure is not concentrated in traditionally routine or low-skilled jobs. It rather has the potential to increasingly affect high-skilled, analytical occupations, emphasizing the importance of policies that strengthen complementarity between human labour and AI. Supporting skills that are less prone to automation such as interpersonal, creative, and adaptive competencies can enhance worker resilience. Targeted and occupation-specific training strategies, rather than one-size-fits-all programmes, are therefore central to supporting workforce adaptability in the face of AI-driven change.

Using firm-level evidence on AI adoption, **Chapter 3** shows that AI deployment unfolds gradually rather than through discrete automation shocks, even in a technically proven system. Studying the rollout of an AI-based visual inspection system across 14 plants in a multinational electronics manufacturer, the chapter reveals that firms expand AI automation incrementally, primarily by deploying AI to new products rather than intensifying its use on already-enabled components. Although AI adoption is associated with reduced labour demand for the specific inspection task, this substitution occurs only after deployment and is shaped by organizational constraints, trust-building processes, and customer-specific quality requirements. These findings underscore

the importance of developing policies and firm strategies that support workers facing job displacement and task adjustment. Early interventions that facilitate task reallocation within firms such as worker mobility pathways can lower adjustment costs. Adjustments to working hours or work schedules can also be useful to help workers adapt to reduced task demands while firms gradually restructure workflows. At the same time, retraining and reskilling policies remain essential, particularly in settings where automation replaces core job functions. Well-targeted transition support, including retraining programs aligned with local labour demand, can help workers move into new roles when displacement occurs. The chapter also highlights the continued role of human workers in monitoring and overseeing AI systems, reinforcing the importance of strengthening organizational and human oversight in shaping the pace and extent of AI deployment. This points to the need to develop regulatory frameworks and governance mechanisms that promote transparency and trustworthy AI use. Effective AI governance requires clear standards for performance validation and processes that involve workers' voice in deployment decisions to ensure that automation proceeds in ways that preserve safety, accountability, and workers' agency.

Chapter 4 investigates how AI can enhance worker productivity when deployed to augment rather than replace human labour. Studying the introduction of an AI-augmented training tool in a European call center, we find that AI significantly improves worker performance by enabling more targeted and individualized coaching. Workers who receive AI-augmented training handle calls faster and more efficiently than those receiving standard training. Productivity gains are particularly pronounced for less-experienced workers, suggesting that AI accelerates learning and helps newer employees move up the experience curve more rapidly. For more experienced agents, AI enables further optimization of communication strategies, allowing them to refine their existing expertise. These findings demonstrate that when AI is embedded in training processes rather than deployed solely to automate tasks, it can complement human skills and strengthen workforce capabilities across experience levels. Crucially, skills and continuous learning emerge as key drivers of AI's productivity impact, suggesting that investments in human capital development is central to realize AI's potential. This

points to the importance of fostering workplace learning and on-the-job training as central components of AI adoption strategies. However, achieving these productivity gains also requires ensuring broad and equitable access to AI-supported tools. Without adequate digital infrastructure, organizational capacity, and investment in training systems some workers might be excluded from AI-driven productivity gains, exacerbating existing inequalities. Our results show that less-experienced workers benefit most from AI-augmented training, suggesting that targeted deployment can reduce skill gaps and accelerate development, particularly for early career workers. Tailoring AI-augmented interventions to specific contexts and competencies can maximize their effectiveness. This requires policymakers and employers to prioritize investments in training infrastructure, ensure that AI-augmented learning tools are accessible and design programmes that align with workers' actual skill needs.

Chapter 5 switches the lens to future labour market entrants and examines how university students perceive AI-related labour market risks and how these perceptions shape their policy preferences. Using a pre-registered survey experiment with students in the UK and the US, we find that exposure to AI competition and negative narratives about AI's labour market consequences significantly increase anxiety about future job prospects. However, this heightened concern does not translate into support for traditional risk-mitigating policies such as expanded social spending, universal basic income, or active labour market programmes. Instead, students' policy preferences remain largely anchored in pre-existing ideological orientations, with left-leaning students favoring redistributive measures and right-leaning students preferring market -oriented or protectionist responses. These findings reveal a disconnect between emotional responses to AI and policy demands, suggesting that young citizens may distrust government capacity to address AI-related challenges or perceive traditional policies as benefiting incumbents rather than new entrants. Understanding these generational differences in risk perception and policy preferences is essential for building broad-based agreements to support inclusive labour markets. Policies should be designed and communicated in ways that address young workers' specific concerns about access, opportunity and mobility

in AI-driven labour markets. This may require rethinking traditional social safety nets and exploring alternative mechanisms, such as policies that facilitate early-career transitions that resonate more directly with the needs and priorities of future labour market entrants. One promising approach is further investing in structured learning pathways that combine education with practical experience which can help young workers build expertise and credibility. Policymakers should explore how to adapt and scale this apprenticeship style approaches across sectors, ensuring that young talent development keeps pace with AI adoption.

Overall, this thesis highlights the importance of deliberate design and policy interventions to foster pro-worker AI technologies that complement workers' skills, expertise and agency. This means harnessing AI's potential to activate groups facing barriers to work and lower barriers to entry, enable diverse forms of employment or enhance productivity, while protecting worker rights and ensuring gains are shared fairly. Furthermore, fostering human-centred AI will require grounding decisions in solid empirical evidence and maintaining ongoing dialogue among workers, firms, policymakers and researchers.

Summary

This thesis examines how Artificial Intelligence (AI) reshapes labour markets across workers, firms, and prospective labour market entrants. As AI capabilities advance rapidly, questions arise about which workers will be most affected, whether AI will complement or substitute human labour, and how societies should prepare for these transformations. This thesis investigates the mechanisms through which AI influences employment outcomes, wage dynamics, task allocation, and worker productivity. The overarching contribution is to demonstrate that AI's labour market implications are not technologically predetermined but depend critically on how AI systems are designed, deployed, and governed within organizational and institutional contexts.

The analysis begins with occupational exposure to AI. **Chapter 2** examines which workers are most susceptible to AI and how this exposure relates to employment and wage outcomes in Germany between 2012 and 2019. Using task-based technology exposure measures and administrative data, the chapter documents that AI exposure is concentrated among high-skilled workers in non-routine, analytical occupations. This contrasts sharply with earlier automation technologies (software and robotics), which primarily affected routine-intensive, middle-skill occupations. Despite high AI exposure, we find no evidence of employment displacement. Furthermore, we observe that AI exposure is positively associated with wage growth, particularly for workers in occupations with high routine task content. These findings challenge conventional predictions from skill-biased and routine-biased technological change frameworks and suggest that AI may affect a broader and more diverse set of workers than previous technological waves.

Chapter 3 shifts from potential exposure to actual adoption, providing firm-level evidence on AI deployment and its implications for labour demand. The chapter studies the roll-out of an AI-based visual inspection system across 14 plants in a multinational electronics manufacturer over a 15-month period, using detailed decision-level data from both human operators and the AI system.

The analysis reveals that AI adoption unfolds gradually rather than through discrete automation shocks, even when the technology is technically mature and rigorously validated. Plants expand automation primarily along the extensive margin—deploying AI on new products and component libraries—rather than increasing automation on already-enabled components. Using two-way fixed effects estimation, we find that a one percentage point increase in AI automation is associated with approximately a 1.2% decline in the growth rate of operator decisions, consistent with task-level labour displacement. However, this substitution occurs only after deployment and is mediated by organizational constraints, trust-building processes, and customer-specific quality requirements. These findings demonstrate that technical capability alone does not alone determine the scope of automation and that institutional and organizational governance play an important role in how AI is integrated into production environments.

Chapter 4 investigates whether AI can enhance worker productivity when deployed to augment human labour. Exploiting the staggered introduction of an AI-augmented training tool in a European call center, the chapter examines how AI-augmented coaching affects worker performance. The AI tool analyzes recorded calls and generates data-driven performance assessments, enabling trainers to provide more targeted and individualized coaching. Using difference-in-differences estimation, we find that workers who receive AI-augmented training handle calls significantly faster and more efficiently than those receiving standard training. Productivity gains are particularly pronounced for less-experienced workers, suggesting that AI accelerates learning and helps newer employees move up the experience curve more rapidly. For more experienced agents, AI enables further optimization of communication strategies and conversation quality. These findings demonstrate that AI's labour market effects depend critically on implementation choices: when embedded in training and learning processes rather than deployed solely to automate tasks, AI can complement human skills and strengthen workforce capabilities across experience levels.

Chapter 5 examines how future labour market entrants perceive AI-related risks and how these perceptions shape policy preferences. Using a pre-registered survey experiment with university students in the United Kingdom

and the United States, the chapter investigates whether exposure to AI competition and information about AI's labour market consequences influence anxiety about job prospects and demand for protective policies. The findings reveal that exposure to AI competition and negative narratives about AI significantly increases anxiety about future employment. However, this heightened concern does not translate into support for traditional risk-mitigating policies such as expanded social spending, universal basic income, or active labour market programmes. Instead, students' policy preferences remain largely anchored in pre-existing ideological orientations, with left-leaning students favouring redistributive measures and right-leaning students preferring market-oriented or protectionist responses. These patterns contrast with research on incumbent workers, who typically demand stronger welfare protections when confronted with technological threats, suggesting important generational differences in how labour market risks are perceived and what policy responses are deemed appropriate.

Taken together, the thesis demonstrates that AI's economic effects are shaped less by technological capabilities alone and more by how various actors design, govern, and respond to AI systems. While AI has the potential to reach a broader and more diverse segment of the workforce than earlier automation technologies, its labour market consequences depend fundamentally on implementation choices, organizational practices, and policy frameworks. The findings underscore the importance of deliberate design and policy interventions to foster pro-worker AI technologies that complement workers' skills and agency. Realizing AI's potential to enhance productivity, lower barriers to entry, and enable diverse forms of employment, while protecting worker rights and ensuring gains are shared fairly, requires grounding decisions in solid empirical evidence and maintaining ongoing dialogue among workers, firms, policymakers, and researchers.

Samenvatting

Dit proefschrift onderzoekt hoe kunstmatige intelligentie (AI) arbeidsmarkten hervormt: voor werknemers, voor bedrijven en voor toekomstige arbeidsmarktdeelnemers. Nu de capaciteiten van AI zich in hoog tempo ontwikkelen, rijzen er vragen over welke werknemers het meest geraakt zullen worden, of AI menselijke arbeid zal aanvullen of juist vervangen, en hoe samenlevingen zich op deze transformaties zouden moeten voorbereiden. Dit proefschrift onderzoekt de mechanismen waarlangs AI van invloed is op werkgelegenheid, loonontwikkeling, taakverdeling en arbeidsproductiviteit. De overkoepelende bijdrage is het aantonen dat de arbeidsmarkteffecten van AI niet technologisch voorbestemd zijn, maar in belangrijke mate afhangen van hoe AI-systemen worden ontworpen, ingezet en bestuurd binnen organisatorische en institutionele contexten.

De analyse begint bij de blootstelling van beroepen aan AI. **Hoofdstuk 2** onderzoekt welke werknemers het meest vatbaar zijn voor AI en hoe deze blootstelling samenhangt met werkgelegenheids- en loonuitkomsten in Duitsland tussen 2012 en 2019. Op basis van taakgebaseerde blootstellingsmaten en administratieve data documenteert het hoofdstuk dat AI-blootstelling geconcentreerd is bij hoogopgeleide werknemers in niet-routinematige, analytische beroepen. Dit staat in scherp contrast met eerdere automatiseringstechnologieën (software en robotica), die vooral routinematige beroepen in het middensegment van de arbeidsmarkt troffen. Ondanks de hoge blootstelling aan AI vinden we geen bewijs voor verdringing uit werkgelegenheid. Bovendien zien we dat AI-blootstelling positief samenhangt met loongroei, met name voor werknemers in beroepen met een hoge routinematige taakinhoud. Deze bevindingen dagen de gangbare voorspellingen van skill-biased en routine-biased technologische veranderingskaders uit en suggereren dat AI een bredere en meer diverse groep werknemers kan raken dan eerdere technologische golven.

Hoofdstuk 3 verschuift de aandacht van potentiële blootstelling naar daadw-

erkelijke implementatie en biedt bewijs op bedrijfsniveau over de uitrol van AI en de gevolgen daarvan voor de vraag naar arbeid. Het hoofdstuk bestudeert de uitrol van een AI-systeem voor visuele kwaliteitsinspectie in 14 fabrieken van een multinationale elektronicaproducent gedurende een periode van 15 maanden, op basis van gedetailleerde data op beslissingsniveau van zowel menselijke operators als het AI-systeem. De analyse laat zien dat AI-adoptie geleidelijk verloopt in plaats van via abrupte automatiseringsschokken, zelfs wanneer de technologie technisch volwassen en grondig gevalideerd is. Fabrieken breiden automatisering vooral uit langs de extensieve marge, door AI in te zetten op nieuwe producten en componentbibliotheken, in plaats van de automatiseringsgraad te verhogen op al vrijgegeven componenten. Met behulp van two-way fixed effects-schattingen vinden we dat een stijging van de AI-automatiseringsgraad met één procentpunt samenhangt met een daling van ongeveer 1,2% in de groei van het aantal operatorbeslissingen, consistent met verdringing van arbeid op taakniveau. Deze substitutie treedt echter pas op na implementatie en wordt gemedieerd door organisatorische belemmeringen, processen van vertrouwensopbouw en klantspecifieke kwaliteitseisen. Deze bevindingen tonen aan dat technische capaciteit op zich niet bepaalt hoe ver automatisering reikt, en dat institutionele en organisatorische governance een belangrijke rol speelt in de wijze waarop AI in productieomgevingen wordt geïntegreerd.

Hoofdstuk 4 onderzoekt of AI de productiviteit van werknemers kan verhogen wanneer het wordt ingezet om menselijke arbeid te versterken. Gebruikmakend van de gefaseerde invoering van een AI-ondersteund trainingsinstrument in een Europees callcenter onderzoekt het hoofdstuk hoe AI-ondersteunde coaching de prestaties van werknemers beïnvloedt. Het AI-instrument analyseert opgenomen gesprekken en genereert datagedreven prestatiebeoordelingen, waardoor trainers gerichtere en meer geïndividualiseerde coaching kunnen bieden. Met difference-in-differences-schattingen vinden we dat werknemers die AI-ondersteunde training ontvangen hun gesprekken aanzienlijk sneller en efficiënter afhandelen dan werknemers die de standaardtraining volgen. De productiviteitswinsten zijn met name uitgesproken voor minder ervaren werknemers, wat erop wijst dat AI het

leerproces versnelt en nieuwere medewerkers helpt om sneller de ervaringscurve te beklimmen. Voor meer ervaren agents maakt AI een verdere optimalisatie van communicatiestrategieën en gesprekskwaliteit mogelijk. Deze bevindingen laten zien dat de arbeidsmarkteffecten van AI sterk afhangen van implementatiekeuzes: wanneer AI wordt ingebed in trainings- en leerprocessen, in plaats van uitsluitend te worden ingezet om taken te automatiseren, kan het menselijke vaardigheden aanvullen en de capaciteiten van het personeelsbestand over alle ervaringsniveaus heen versterken.

Hoofdstuk 5 onderzoekt hoe toekomstige arbeidsmarktdeelnemers AI-gerelateerde risico's waarnemen en hoe deze percepties hun beleidsvoorkeuren vormgeven. Met een gepreregistreerd survey-experiment onder universiteitsstudenten in het Verenigd Koninkrijk en de Verenigde Staten onderzoekt het hoofdstuk of blootstelling aan concurrentie met AI en informatie over de arbeidsmarktgevolgen van AI de bezorgdheid over toekomstige baankansen en de vraag naar beschermend beleid beïnvloeden. De bevindingen laten zien dat blootstelling aan AI-concurrentie en negatieve berichtgeving over AI de bezorgdheid over toekomstige werkgelegenheid significant vergroten. Deze toegenomen bezorgdheid vertaalt zich echter niet in steun voor traditionele risicobeperkende beleidsmaatregelen zoals hogere sociale uitgaven, een universeel basisinkomen of actief arbeidsmarktbeleid. In plaats daarvan blijven de beleidsvoorkeuren van studenten grotendeels verankerd in reeds bestaande ideologische oriëntaties: linksgeoriënteerde studenten geven de voorkeur aan herverdelende maatregelen, terwijl rechtsgeoriënteerde studenten marktgerichte of protectionistische antwoorden prefereren. Deze patronen contrasteren met onderzoek naar zittende werknemers, die doorgaans om sterkere sociale bescherming vragen wanneer zij met technologische dreigingen worden geconfronteerd, wat wijst op belangrijke generatieverschillen in hoe arbeidsmarktrisico's worden ervaren en welke beleidsreacties passend worden geacht.

Alles bijeengenomen toont dit proefschrift aan dat de economische effecten van AI minder worden bepaald door technologische capaciteiten op zich, en meer door de wijze waarop verschillende actoren AI-systemen ontwerpen, besturen en erop reageren. Hoewel AI het potentieel heeft om een breder en

verscheidener deel van de beroepsbevolking te bereiken dan eerdere automatiseringstechnologieën, hangen de arbeidsmarktgevolgen fundamenteel af van implementatiekeuzes, organisatiepraktijken en beleidskaders. De bevindingen onderstrepen het belang van doordacht ontwerp en beleidsinterventies om pro-werknemers AI-technologieën te bevorderen die de vaardigheden en de autonomie van werknemers aanvullen. Het verwezenlijken van AI's potentieel om de productiviteit te verhogen, toetredingsdrempels te verlagen en uiteenlopende vormen van werk mogelijk te maken, en tegelijkertijd de rechten van werknemers te beschermen en een eerlijke verdeling van de opbrengsten te waarborgen, vereist dat beslissingen worden gegrond in solide empirisch bewijs en dat er een voortdurende dialoog wordt gevoerd tussen werknemers, bedrijven, beleidsmakers en onderzoekers.

Published work

- Fouarge, D., Fregin, M.-C., Janssen, S., Levels, M., Montizaan, R., Özgül, P., Rounding, N. & Stops, M. (2025). *How AI-Augmented Training Improves Worker Productivity*. IZA Discussion Paper Series No. 18224.

- Fregin, M.-C., Eijkenboom, D., Özgül, P., Pardesi, M. & Rounding, N. (2026). *Artificial Intelligence in Firms: A Synthesis of Micro-Level Evidence on Productivity, Work, and Skills*. ROA Policy Brief. Research Centre for Education and the Labour Market.

- Janssen, S., Stops, M., Steens, S., Özgül, P., Rounding, N., Dijksman, S., Montizaan, R., Levels, M., Fouarge, D., Eijkenboom, D., Graus, E. & Fregin, M.-C. (2025). *Künstliche Intelligenz in der Arbeitswelt: Wie KI-unterstütztes Training die Gespräche im Kundenservice verändert*. IAB-Forum.

- Fregin, M.-C., Rounding, N., van der Velden, R., Levels, M., Özgül, P. & Steens, S. (2025). *Experiences with Experimental Research: Lessons from a Co-creative Research Project on AI Effects in Companies*. ai:conomics Policy Brief. Research Centre for Education and the Labour Market.

- Fregin, M.-C., Koch, T., Malfertheiner, V., Özgül, P. & Stops, M. (2024). *On the Automation of Job Tasks: Occupational Exposure to Artificial Intelligence and Software*. ai:conomics Policy Brief. Research Centre for Education and the Labour Market. (Also available in German as IAB Kurzbericht 21/23.)

- Fleck, L., Rounding, N. & Özgül, P. (2022). *Artificial Intelligence in Hiring: Friend or Foe?* ai:conomics Policy Brief. Research Centre for Education and the Labour Market.

- Graus, E., Özgül, P. & Steens, S. (2021). *Artificial Intelligence: Shaping the Future of Work with Insights from Firm-Level Evidence*. ai:conomics Policy Brief. Research Centre for Education and the Labour Market.

About the author

Pelin Özgül was born on February 7, 1994 in Ankara (Turkey). She completed her BSc in Economics at Bilkent University, before obtaining two Master's degrees: the first in International Economics and Business from the University of Groningen, and the second in Advanced Studies in Economics (MASE) from KU Leuven.

In 2021, she joined the Research Centre for Education and the Labour Market (ROA) at Maastricht University as a PhD candidate. Her doctoral research was carried out within the *ai:conomics* project and examined how Artificial Intelligence reshapes work, workers, and their labour market outcomes, drawing on administrative, firm-level, and experimental data. During her PhD, she presented her work at international conferences and contributed to several academic papers and policy briefs on AI and the future of work. Since 2026, Pelin works as an Economist at the National Bank of Belgium.

ROA Dissertation Series

1. Lex Borghans (1993), *Educational Choice and Labour Market Information*, Maastricht, Research Centre for Education and the Labour Market.

2. Frank Cörvers (1999), *The Impact of Human Capital on International Competitiveness and Trade Performance of Manufacturing Sectors*, Maastricht, Research Centre for Education and the Labour Market.

3. Ben Kriechel (2003), *Heterogeneity Among Displaced Workers*, Maastricht, Research Centre for Education and the Labour Market.

4. Arnaud Dupuy (2004), *Assignment and Substitution in the Labour Market*, Maastricht, Research Centre for Education and the Labour Market.

5. Wendy Smits (2005), *The Quality of Apprenticeship Training, Conflicting Interests of Firms and Apprentices*, Maastricht, Research Centre for Education and the Labour Market.

6. Judith Semeijn (2005), *Academic Competences and Labour Market Entry: Studies Among Dutch Graduates*, Maastricht, Research Centre for Education and the Labour Market.

7. Jasper van Loo (2005), *Training, Labor Market Outcomes and Self-Management*, Maastricht, Research Centre for Education and the Labour Market.

8. Christoph Meng (2005), *Discipline-Specific or Academic? Acquisition, Role and Value of Higher Education Competencies*, Maastricht, Research Centre for Education and the Labour Market.

9. Andreas Ammermüller (2007), *Institutional Effects in the Production of Education: Evidence from European Schooling Systems*, Maastricht, Research Centre for Education and the Labour Market.

10. Bart Golsteyn (2007), *The Ability to Invest in Human Capital*, Maastricht, Research Centre for Education and the Labour Market.

11. Raymond Montizaan (2010), *Pension Rights, Human Capital Development and Well-being*, Maastricht, Research Centre for Education and the Labour Market.

12. Annemarie Nelen (2012), *Part-Time Employment and Human Capital Development*, Maastricht, Research Centre for Education and the Labour Market.

13. Jan Sauermann (2013), *Human Capital, Incentives, and Performance Outcomes*, Maastricht, Research Centre for Education and the Labour Market.

14. Harald Ulrich Pfeifer (2013), *Empirical Investigations of Costs and Benefits of Vocational Education and Training*, Maastricht, Research Centre for Education and the Labour Market.

15. Charlotte Büchner (2013), *Social Background, Educational Attainment and Labor Market Integration: An Exploration of Underlying Processes and Dynamics*, Maastricht, Research Centre for Education and the Labour Market.

16. Martin Humburg (2014), *Skills and the Employability of University Graduates*, Maastricht, Research Centre for Education and the Labour Market.

17. Jan Feld (2014), *Making the Invisible Visible, Essays on Overconfidence, Discrimination and Peer Effects*, Maastricht, Research Centre for Education and the Labour Market.

18. Olga Skriabikova (2014), *Preferences, Institutions, and Economic Outcomes: an Empirical Investigation*, Maastricht, Research Centre for Education and the Labour Market.

19. Gabriele Marconi (2015), *Higher Education in the National and Global Economy*, Maastricht, Research Centre for Education and the Labour Market.

20. Nicolas Salamanca Acosta (2015), *Economic Preferences and Financial Risk-Taking*, Maastricht, Research Centre for Education and the Labour Market.

21. Ahmed Elsayed Mohamed (2015), *Essays on Working Hours*, Maastricht, Research Centre for Education and the Labour Market.

22. Roxanne Amanda Korthals (2015), *Tracking Students in Secondary Education, Consequences for Student Performance and Inequality*, Maastricht, Research Centre for Education and the Labour Market.

23. Maria Zumbuehl (2015), *Economic Preferences and Attitudes: Origins, Behavioral Impact, Stability and Measurement*, Maastricht, Research Centre for Education and the Labour Market.

24. Anika Jansen (2016), *Firms' Incentives to Provide Apprenticeships – Studies on Expected Short- and Long-Term Benefits*, Maastricht, Research Centre for Education and the Labour Market.

25. Jos Maarten Arnold Frank Sanders (2016), *Sustaining the Employability of the Low Skilled Worker: Development, Mobility and Work Redesign*, Maastricht, Research Centre for Education and the Labour Market.

26. Marion Collewet (2017), *Working Hours: Preferences, Well-being and Productivity*, Maastricht, Research Centre for Education and the Labour Market.

27. Tom Stolp (2018), *Sorting in the Labor Market: The Role of Risk Preference and Stress*, Maastricht, Research Centre for Education and the Labour Market.

28. Frauke Meyer (2019), *Individual Motives for (Re-)distribution*, Maastricht, Research Centre for Education and the Labour Market.

29. Maria Ferreira Sequeda (2019), *Human Capital Development at School and Work*, Maastricht, Research Centre for Education and the Labour Market.

30. Marie-Christine Martha Fregin (2019), *Skill Matching and Outcomes: New Cross-Country Evidence*, Maastricht, Research Centre for Education and the Labour Market.

31. Sanne Johanna Leontien van Wetten (2020), *Human Capital and Employee Entrepreneurship: The Role of Skills, Personality Characteristics and the Work Context*, Maastricht, Research Centre for Education and the Labour Market.

32. Cécile Alice Jeanne Magnée (2020), *Playing the Hand You're Dealt, The Effects of Family Structure on Children's Personality and the Effects of Educational Policy on Educational Outcomes of Migrant Children*, Maastricht, Research Centre for Education and the Labour Market.

33. Merve Nezihe Özer (2020), *Essays on Drivers and Long-Term Impact of Migration*, Maastricht, Research Centre for Education and the Labour Market.

34. Inge Ingeborg Henrica Maria Hooijen (2021), *Place Attractiveness, A Study of the Determinants Playing a Role in Residential Settlement Behaviour*, Maastricht, Research Centre for Education and the Labour Market.

35. Alexandra Marie Catherine de Gendre (2021), *Behavioral Barriers to Success in Education*, Maastricht, Research Centre for Education and the Labour Market.

36. Kim van Broekhoven (2021), *From Creativity to Innovation: Understanding and Improving the Evaluation and Selection of Ideas in Educational Settings*, Maastricht, Research Centre for Education and the Labour Market.

37. François M.B.M. Molin (2022), *Using Digital Formative Assessments to Improve Learning in Physics Education*, Maastricht, Research Centre for Education and the Labour Market.

38. Bart Kasper de Koning (2022), *Empirical Studies on Information, Beliefs, and Choices in Education and Work*, Maastricht, Research Centre for Education and the Labour Market.

39. Alexander Dicks (2023), *NEET in the Netherlands*, Maastricht, Research Centre for Education and the Labour Market.

40. Lynn Lamberdina Johanna van Vugt (2023), *Different NEETs, Different Needs? Explaining Why Vulnerable Young People are More Likely to Become NEET*, Maastricht, Research Centre for Education and the Labour Market.

41. Maria-Eugenia Polipciuc (2023), *An Exploration of Trust, Betrayal, & Social Identity*, Maastricht, Research Centre for Education and the Labour Market.

42. Lies van de Kuilen (2024), *Governance of Complex Early Childhood Education and Care Systems: Promising Governance Strategies for Effective Implementation of Educational Equity Policy in Early Childhood*, Maastricht, Research Centre for Education and the Labour Market.

43. Pascal Heß (2024), *How Does Salient Information on Economic Prospects Shape Human Capital Investments?*, Maastricht, Research Centre for Education and the Labour Market.

44. Peter Hein Bles (2024), *Flexibility in the Dutch Education System: How Institutional Characteristics Modify Students' Educational Careers*, Maastricht, Research Centre for Education and the Labour Market.

45. Alina Shirshikova (2024), *Paths to Opportunity: Navigating Mobility-Related Barriers in Access to Education and Labor Market*, Maastricht, Research Centre for Education and the Labour Market.

46. Stijn Sebastiaan Horck (2024), *Failure as Catalyst for Learning in Healthcare: Investigations into Learning Responses in Different Contexts*, Maastricht, Research Centre for Education and the Labour Market.

47. Mélanie Nadège Winanda Monfrance (2024), *Exploring Non-Cognitive Skills and Inequality through the Magnifying Lens of a Pandemic*, Maastricht, Research Centre for Education and the Labour Market.

48. Luisa Sophie Caroline Minssen (2025), *Understanding Recruitment Choices: The Role of Applicant and Recruiter Characteristics*, Maastricht, Research Centre for Education and the Labour Market.

49. Babs Petronella Johanna Anne-Marie Jacobs (2025), *Key Skills for Lifelong Impact: Parental Resources, Family Dynamics, and Labour Market Outcomes*, Maastricht, Research Centre for Education and the Labour Market.

50. Elke Claes (2025), *Basic Skills*, Maastricht, Research Centre for Education and the Labour Market.

51. Tianyu Yang (2026), *Adaptation under Uncertainty: How Students Navigate Ed-*

ucational and Career Decisions during Major Disruptions, Maastricht, Research Centre for Education and the Labour Market.

52. José Victor Cremonesi Giarola (2026), *International Labor Mobility: Roots, Returns and Consequences*, Maastricht, Research Centre for Education and the Labour Market.

53. Carolin Linckh (2026), *Firms' Skill Demand: The Role of Minimum Wages, Tightness and Informal Training*, Maastricht, Research Centre for Education and the Labour Market.

54. Anna-Lena Wittich (2026), *Shaping Human Capital and Work Practices in a Changing Labor Market*, Maastricht, Research Centre for Education and the Labour Market.

55. Sandra Pérez Rodríguez (2026), *Occupational Choice: The Role of Skills, Preferences and Expected Labour Market Returns*, Maastricht, Research Centre for Education and the Labour Market.

56. Angela Messioui (2026), *Economic skills obsolescence*, Maastricht, Research Centre for Education and the Labour Market.

57. Pelin Özgül (2026), *AI at Work: How AI Reshapes Work and Workers' Labour Market Outcomes*, Maastricht, Research Centre for Education and the Labour Market.